

Characterization of Neuropathological Shape Deformations

by

John William Martin II

Submitted to the Department of Nuclear Engineering
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

May 1995

© Massachusetts Institute of Technology 1995. All rights reserved.

Author

.....
Department of Nuclear Engineering
May 17, 1995

Certified by

.....
Alex P. Pentland
Associate Professor
Thesis Supervisor

Certified by

Jacquelyn C. Yanch
Associate Professor
Thesis Reader

Accepted by

.....
Allan F. Henry
Chairman, Departmental Committee on Graduate Students

ARCHIVES

MASSACHUSETTS INSTITUTE
OF TECHNOLOGY

JUN 07 1995



Characterization of Neuropathological Shape Deformations

by

John William Martin II

Submitted to the Department of Nuclear Engineering
on May 17, 1995, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy

Abstract

This thesis 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. Using this 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.

The 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 thesis, the focus is 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, it is possible to compensate for some of the nonpathological modes of shape variation. This allows the experimental characterization of modes of variation that are indicative of disease processes.

The technique is applied to magnetic resonance images of the brains of individuals with schizophrenia, Alzheimer's disease, and normal-pressure hydrocephalus, as well as to healthy volunteers. The results are used in three ways. First, the modes of variation are displayed in order to indicate the types of deformations caused by different pathologies. Second, projections onto the modes are fed into standard pattern recognition algorithms in order to perform shape-based diagnosis. Finally, the results are used in a database program, indicating that the brain structure shape representation developed in this thesis is an efficient metric for database search.

Thesis Supervisor: Alex P. Pentland
Title: Associate Professor

Thesis Reader: Jacquelyn C. Yanch
Title: Associate Professor

Acknowledgements

First and foremost, I wish to thank my thesis advisor, Sandy Pentland. His combination of technical expertise and friendly manner has made working on this thesis both interesting and enjoyable. In addition, Sandy has managed to surround himself with a remarkable group of people. Interacting with the faculty, students, and staff of the Perceptual Computing Section of the MIT Media Lab has been a rewarding and intellectually stimulating experience. In particular I would like to thank Irfan Essa for many interesting discussions, and for helping me to get up to speed on physical modeling. Both Laureen Chapman and Judy Bornstein have been a pleasure to work with. I am especially grateful to Stan Sclaroff, who helped me understand many of the more mathematical aspects of this thesis, and also gave me programming help on several different occasions. Without his input, this document would have been of much lower quality.

I have also been fortunate to be associated with the Surgical Planning Laboratory of Brigham and Women's Hospital. Ron Kikinis provided both medical advice as well as friendship, and shielded me from funding concerns. Under the guidance of Ron, and with the help of Ferenc Jolesz, the SPL was able to flourish into a world-class medical image processing lab. Hats off to both Adam and Mark for keeping the SPL running while the software and hardware seemed to change on a daily basis. The POETS of the SPL made the lab a fun place to work, and provided much-needed after hours distractions. I could complain about the Red Sox with Mark, and could complain about everything else in the world with Cindy. And getting to know I-Han was a truly wonderful experience.

Thanks also to the other members of my thesis committee, Jacquelyn Yanch and Eric Grimson. Jacquelyn took the time to understand my work and provide helpful input, even though the topic of this thesis is outside her own research areas. Eric gave useful advice throughout the development of this work. Many of his insightful comments found their way into the final version of this document.

I am also grateful for the friendship of current and past students of the Radiological Sciences Program: Sekhar, Mickey, Larry, Chris, Sonny, Steve, and Alex. Many thanks also go out to my good friends Brad, Harry, and Amy. They helped remind me that there are other things in life besides work.

Finally, I wish to thank my family for providing support both during my long haul at MIT, as well as during the preceding 26 years. It is hard to believe that I will no longer be hearing the ubiquitous "So, when are you graduating?"

Contents

1	Introduction	8
1.1	General Approach	9
1.2	An Example	11
1.3	Thesis Overview	13
2	Background	15
2.1	Registration	15
2.2	Segmentation	17
2.3	Shape Description	17
3	Mathematical Model	20
3.1	Physical Modeling	21
3.2	Experimental Modes	25
3.3	The Connection	26
3.4	Compound Modes	27
3.5	Summary	29
4	Application of Model to Neuropathology	30
4.1	Head Shape Normalization	30
4.2	Characterization of Disease States	37
4.3	Classifier Design	39
5	Experiments with Schizophrenia	42
5.1	Medical Background	42
5.2	The Putamen Data Set	43
5.3	Application of Our Method	43
5.4	Experimental Results	43
6	Experiments with Ventricular Disorders	48
6.1	Medical Background	48
6.2	The Ventricular Data Set	49
6.3	Application of Our Method	49
6.4	Experimental Results	49
7	Diagnostic Tool	52
7.1	Content-Based Searches	52
7.2	Photobook	54
7.3	Demonstration	54

7.4 Experiments	55
8 Discussion	65
8.1 Effect of Head Shape Normalization	65
8.2 Model Accuracy	67
8.3 Statistical Analysis	69
8.4 Number of Principal Components	71
8.5 Preset Parameters	73
8.6 Other Issues and Limitations	75
8.7 Implementation	78
8.8 Suggestions for Future Research	80
9 Summary	83

List of Figures

1-1	Schematic representation of head and ventricles	10
1-2	3-D reconstruction of left putamen	12
1-3	3-D reconstruction of cerebral ventricles	12
3-1	Physical modes of a square	25
3-2	Decomposition of deformations	28
4-1	Physical deformation modes of the ICC	32
4-2	Normalizing ventricular shape by cranium shape - side view	37
4-3	Normalizing ventricular shape by cranium shape - top view	38
5-1	Putamen data set - side view	45
5-2	Putamen data set - front view	46
5-3	Putamen deformation modes	47
6-1	Ventricular data set	50
6-2	Ventricular deformation modes	51
7-1	Photobook functionality	57
7-2	Ventricular database in Photobook	58
7-3	Photobook ranking after selecting NPH1	59
7-4	Photobook ranking after selecting AD2 and female	60
7-5	Putamen database in Photobook	61
7-6	Photobook ranking after selecting Normal10	62
7-7	Photobook ranking after selecting Schiz4	63
7-8	Photobook ranking after selecting Schiz3	64
8-1	Projections onto putamen modes	66
8-2	Projections onto ventricle modes	67
8-3	Mean square distance between ventricles	68
8-4	Projections onto ventricle modes 5 and 6	74
8-5	Unsuccessful warping of cranial contents - side view	76
8-6	Unsuccessful warping of cranial contents - top view	76

List of Tables

Putamen volume classification rates	13
Putamen Gaussian linear classification rates	47
Ventricle Gaussian linear classification rates	51
Putamen nearest neighbor classification rates	56
Putamen Gaussian linear classification rates	56
Ventricle nearest neighbor classification rates	56
Ventricle Gaussian linear classification rates	56
Robustness of # PCs - Gaussian linear classifier, ventricle database	71
Robustness of # PCs - Gaussian linear classifier, putamen database	71
Robustness of # PCs - nearest neighbor classifier, ventricle database	72
Robustness of # PCs - nearest neighbor classifier, putamen database	72

Chapter 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. The aim of this thesis is to quantitatively describe these pathological shape deformations.

Previous studies of neuropathological morphology suffer from two drawbacks. First, these studies have used just linear [6, 44, 34, 83, 29], planar [64], and/or volumetric [29, 69, 15, 58, 70, 14, 3, 72, 13, 55, 71, 40] 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 [34, 46] and/or smoothness [34]. None of these previous shape descriptions is both general and quantitative.

The second drawback of previous work involves the method of normalizing for non-pathological inter-patient 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. While this is the correct first-order approximation, a more thorough normalization would take into account the complete brain shape.

This thesis overcomes 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. The 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, can be used in pattern recognition algorithms to classify diseases based on shape changes, and can be used as similarity metrics in order to search through a medical image database.

1.1 General Approach



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

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



To make this precise, we represent the shape of an average, healthy brain structure as a set of 3-D point positions \mathbf{X}_A . This list will contain one entry for every spatial location included in the model. This could be every voxel in the volume, just the surface voxels, or even just a small set of landmarks. The particular choice depends on the implementation, and also on exactly what is expected to be measured. 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:

$$\mathbf{u}_{p,i}(\mathbf{x}_i) = \mathbf{u}_{p,i}^H(\mathbf{x}_i) + \mathbf{u}_{p,i}^D(\mathbf{x}_i) \quad (1.1)$$



where $\mathbf{u}_{p,i}^H(\mathbf{x}_i)$ is the deformation due to global effects that are correlated with overall head



(a)



(b)



(c)

Figure 1-1: 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.

shape, and $\mathbf{u}_{p,i}^D(\mathbf{x}_i)$ is the deformation caused by disease and individual local variation.¹ For the entire point set we have

$$\mathbf{U}_p = \mathbf{U}_p^H + \mathbf{U}_p^D \quad (1.2)$$

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

$$\mathbf{U}_p = \begin{bmatrix} \mathbf{u}_{p,1}(\mathbf{x}_1) \\ \mathbf{u}_{p,2}(\mathbf{x}_2) \\ \cdot \\ \cdot \\ \cdot \\ \mathbf{u}_{p,V}(\mathbf{x}_V) \end{bmatrix} \quad (1.3)$$

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

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

We apply *modal analysis* to the physical modeling, and *principal component analysis* to the experimental observations. Both are eigenanalysis techniques, and represent shape in terms of deformation modes [57, 67, 26]. These modes represent unique, natural coordinates in which to express the shape and deformation of brain structures.

1.2 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). Figure 1-2 shows

¹Equation 1.1 is valid only if $\mathbf{u}_{p,i}^H(\mathbf{x}_i)$ and $\mathbf{u}_{p,i}^D(\mathbf{x}_i)$ are measured in the same coordinate system.




Figure 1-2: Reconstructions of the left putamen created from MR images. (a) Normal, healthy adult. (b) Patient with schizophrenia.




Figure 1-3: 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.

the left putamen of a healthy volunteer and of a patient with schizophrenia, while Figure 1-3 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 [40], and that both AD and NPH cause the ventricles to enlarge [53]. While these studies treated just volumetric changes, we seek to explore other pathological deformations in addition to just volume.

As a reference point for the methodology to be developed in the following sections, we consider here possible ways of classifying patients into the two classes shown in Figure 1-2. Given a data set consisting of samples from these two categories, the most straightforward classification procedure would be to use just one feature, putamen volume. However, since a person with a larger head will tend to have larger putamen, even if healthy he or she may be misclassified as schizophrenic. Therefore the second procedure to try would be to normalize each person's putamen volume by his or her overall intra-cranial cavity (ICC) volume.

Using each of these two features, putamen volume and putamen volume normalized

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

TABLE I
Putamen Volume Classification Rates

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

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

1.3 Thesis Overview

Chapter 2 reviews previous work in this field. Because work specific to neuropathology has used only very simple shape measures, and has already been briefly discussed in this chapter, Chapter 2 focuses on the application of shape models to medical imaging in general.

Chapter 3 presents the mathematical model that forms the basis of this thesis. The model consists of two types of deformation modes. The first type of mode is physical, coming from finite element modeling and modal analysis. The second type of mode is statistical, being derived from experimental observations using principal component analysis.

Chapter 4 applies the mathematical model developed in Chapter 3 to the application area of this thesis, neuropathology. The physical modes of the model are used to normalize a patient's cranial contents for his or her particular head shape, while the statistical modes of the model are used to characterize the shape deformations caused by neuropathology.

The next two chapters describe the experimental results obtained by applying the model to two specific medical data sets. Chapter 5 applies the model to the putamen of schizophrenics, while Chapter 6 applies it to the cerebral ventricles of patients with Alzheimer's disease and patients with normal-pressure hydrocephalus. In each case, the recovered modes are first displayed in order to show the sorts of deformations undergone by the afflicted brain structures, and then input into a pattern recognition algorithm in order to perform shape-based diagnosis.

Chapter 7 uses the results of the previous two chapters as input into a database search tool. Basically, the recovered experimental modes, which serve as a shape description for a particular brain structure, are used to form a similarity metric upon which the database tool can search. A few simple experiments are run in order to indicate the functionality of this tool when applied to the medical domain.

Many issues arise during the above chapters; Chapter 8 organizes and discusses these issues, and brings up other open questions. Also, current limitations as well as possible future research directions are considered.

Finally, Chapter 9 briefly summarizes the work done for this thesis.

Chapter 2

Background

This thesis develops a shape representation method consisting of both physical and statistical deformation modes. Once developed, the method is used to characterize neuropathological shape deformations.

There are thus two areas of previous work that relate to this thesis. This first involves the application area of this thesis, neuropathology. As discussed in Chapter 1, however, previous studies of the shape effects of various neuropathologies have used either very simple measures such as volume, or when the measures have been more complex, they have been qualitatively provided by a user. To our knowledge, there are no previous neuropathological shape studies that employ a quantitative, general shape methodology.

In this chapter we therefore focus on the second area, previous work that employs methods similar to those developed in this thesis. 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 chapter we review relevant literature from all three application areas, and draw comparisons to our work on shape description.

2.1 Registration

Bajcsy [4] used an elastic model, combining it with cross correlation measures in order to align raw gray scale brain scans with a simplified brain atlas. Christensen *et al.* [19]

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.

Volumetric deformations have also been computed without resorting to a physical model. Similar to Bajcsy, Collins *et al.* [23, 25] 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. Szeliski and Lavallée [78] used a least-squares approach to compute the minimum distance between two sets of 3-D surface points. Although surfaces are matched, the deformation is modeled as a volumetric spline. Their multiresolution implementation, based on *octree splines*, is computationally very efficient. Face data sets from a range scanner were registered to one another, as were vertebra data sets from a computed tomography (CT) scanner. Feldmar and Ayache [32] also matched 3-D surface data in a non-physical fashion, using *locally affine deformations*. The distance function they minimize incorporates 3-D distance, surface normals, and principal curvatures. Faces, brains, and hearts were registered.

Although all of the above techniques are formulated differently, it has been shown that physically-based and probabilistic approaches are in fact very similar to one another [77, 10]. Thus the important differences between these techniques lie more in the implementation details, and thus in speed and convergence properties, than in the particular formulation employed.

While all these methods have shown promising results, they all suffer from one major drawback. The priors used in their deformation models, whether in the form of elastic constants, prior probabilities, or allowable geometric transformations (e.g. affine), are typically chosen in an ad hoc fashion, often for numerical convenience. The result is either that the physical and geometric models deform in a physically unrealistic way, or that the prior probabilities in the Bayesian models are poor approximations of the true priors.

2.2 Segmentation

Both physical and probabilistic methods have also been used for medical image segmentation. The primary way in which physical models have been employed has been through the use of deformable models such as snakes [80] and their variants. Cohen [22] augmented the original snake formulation with a balloon force to help it avoid local minima. Staib and Duncan [73] 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 from a training set of manually traced contours, and then applied Bayes rule in order to find the best set of parameter values. Székely *et al.* [76] extended this approach to 3-D using the spherical harmonic technique developed by Brechbühler [12], with statistical deviations in the shape parameters derived from 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 [73] and Székely *et al.* [76] 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 [57]. 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.

2.3 Shape Description

2.3.1 Previous work

Because of the above limitations of mesh-like approaches, researchers have developed other physically-based shape representations. Pentland and Sclaroff [57], on which this thesis 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

[8] required corresponding landmarks, more recent efforts [9] 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 [82] 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.* [27] used principal components to experimentally describe the modes of variation inherent in a training set of 2-D heart images. Hill *et al.* [37] extended this technique to 3-D and analyzed the cerebral ventricles for purposes of segmentation and representation.

Along with ourselves [52], other researchers have also begun investigating the relationship between physical and experimental modeling. Cootes [26] 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 [85] have also considered both physical and statistical shape models, in the context of representing and recognizing objects from their 2-D silhouettes.

2.3.2 Comparison to our method

Like modal analysis, all of the above shape descriptors satisfy the requirements of being unique and compact. Physical deformation modes, however, 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 [66]. While Bookstein's thin-plate modes are also physical, they come from a 2-D model and are derived from a finite difference formulation. A finite element model is more general, and has better convergence properties. Also, 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 thesis is the connection between FEM modes and principal component analysis, presented in Section 3.3. 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 [26].

The final original feature of this thesis 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 already mentioned, previous work that has examined pathological brain morphology has used just linear, planar, and/or volumetric shape measurements.

Chapter 3

Mathematical Model

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, in this chapter 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 set. 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 experimental observation.

3.1 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 behavior. The solution to these equations gives the elastic deformation that the object undergoes due to applied forces.

3.1.1 The finite element method

The most common numerical approach for solving elastic deformation problems is the *finite element method* (FEM) [7]. 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:

$$\mathbf{x}(\mathbf{r}) = \mathbf{H}(\mathbf{r})\hat{\mathbf{X}} \quad (3.1)$$

$$\mathbf{u}(\mathbf{r}) = \mathbf{H}(\mathbf{r})\hat{\mathbf{U}}. \quad (3.2)$$

Here $\hat{\mathbf{X}}$ and $\hat{\mathbf{U}}$ denote the nodal position and displacement vectors of the shape model, respectively, and are defined in the (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 thesis, a vector with a “hat” ($\hat{}$) denotes a set of FEM nodal positions or displacements, while a vector without a “hat”

denotes positions or displacements at a set of non-nodal points.)

The shape matrix \mathbf{H} is given by [7]

$$\mathbf{H}(\mathbf{r}) = \begin{bmatrix} h_1(r) & 0 & 0 & h_2(r) & 0 & 0 & \cdots & h_n(r) & 0 & 0 \\ 0 & h_1(s) & 0 & 0 & h_2(s) & 0 & \cdots & 0 & h_n(s) & 0 \\ 0 & 0 & h_1(t) & 0 & 0 & h_2(t) & \cdots & 0 & 0 & h_n(t) \end{bmatrix}. \quad (3.3)$$

This set of polynomial shape functions is used to relate the displacement at any particular point to the displacements at the FEM nodes. The fundamental properties of the h_i are:

- $h_i = 1.0$ at node i and equals 0 at every other node.
- $\sum_{i=1}^n h_i = 1.0$ at any point in the object.

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:

$$\epsilon(\mathbf{x}) = \mathbf{B}(\mathbf{x})\hat{\mathbf{U}} \quad (3.4)$$

where the strain displacement matrix \mathbf{B} is computed by taking the appropriate derivatives of the interpolation matrix \mathbf{H} . (See [30] for more details.) Because \mathbf{B} is a function of \mathbf{x} , and \mathbf{H} is a function of \mathbf{r} , the chain rule must be invoked in order to perform the differentiation. This requires the use of the Jacobian matrix \mathbf{J} :

$$\mathbf{J} = \begin{bmatrix} \frac{\partial x}{\partial r} & \frac{\partial y}{\partial r} & \frac{\partial z}{\partial r} \\ \frac{\partial x}{\partial s} & \frac{\partial y}{\partial s} & \frac{\partial z}{\partial s} \\ \frac{\partial x}{\partial t} & \frac{\partial y}{\partial t} & \frac{\partial z}{\partial t} \end{bmatrix}. \quad (3.5)$$

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

$$\mathbf{K}\hat{\mathbf{U}} = \mathbf{R} \quad (3.6)$$

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:

$$\mathbf{K} = \int_V \mathbf{B}^T \mathbf{C} \mathbf{B} dV, \quad (3.7)$$

where \mathbf{C} is the *material matrix*, which expresses the material's particular stress-strain law. See Bathe [7] for more details on setting up FEM integrals and equations.

3.1.2 Modal analysis

It is often more convenient to represent the nodal displacement vector $\hat{\mathbf{U}}$ in the *modal* coordinate system, in which displacements are represented as linear combinations of an object's free vibration modes. These modes provide a unique, natural, and compact coordinate system in which to represent shape and shape change, are computationally efficient to calculate, and have convenient robustness properties with respect to sampling irregularities and measurement noise [57].

To compute the free vibration modes, Equation 3.6 is diagonalized via an orthogonal transform Φ :

$$\hat{\mathbf{U}} = \Phi \tilde{\mathbf{U}} \quad (3.8)$$

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 3.8 into Equation 3.6 and premultiplying by Φ^T yields

$$\tilde{\mathbf{K}} \tilde{\mathbf{U}} = \tilde{\mathbf{R}} \quad (3.9)$$

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

$$\mathbf{K} \phi_i = \lambda_i \phi_i \quad (3.10)$$

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} ,

$$\Phi = [\phi_1, \phi_2, \phi_3, \dots, \phi_{3N}], \quad (3.11)$$

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

$$\tilde{\mathbf{K}} = \Phi^T \mathbf{K} \Phi = \Lambda = \begin{bmatrix} \lambda_1 & & & & & \\ & \lambda_2 & & & & \\ & & \cdot & & & \\ & & & \cdot & & \\ & & & & \cdot & \\ & & & & & \lambda_{3N} \end{bmatrix}. \quad (3.12)$$

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

$$\mathbf{U} = \Phi \mathbf{a} = \sum_{i=1}^{3N} \phi_i a_i \quad (3.13)$$

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

The modal decomposition is illustrated by the two-dimensional example shown in Figure 3-1. 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 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.

See [57] 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`.

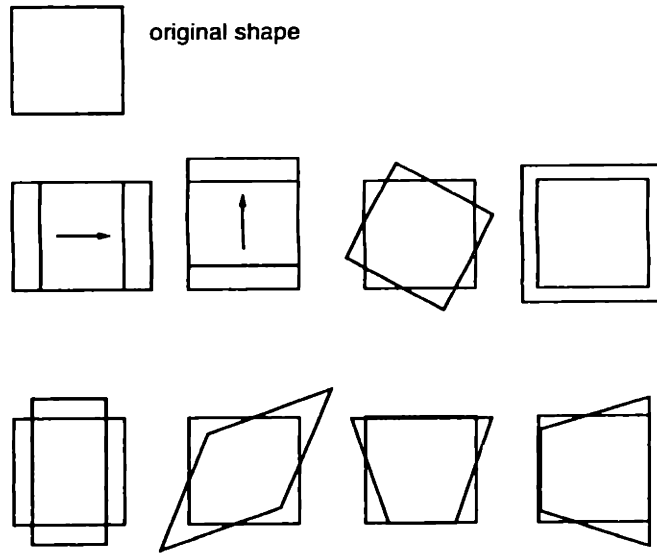


Figure 3-1: Physical modes of a square.

3.2 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:

$$\mathbf{S} = \frac{1}{P-1} \sum_{p=1}^P (\mathbf{U}_p - \bar{\mathbf{U}})(\mathbf{U}_p - \bar{\mathbf{U}})^T, \quad (3.14)$$

where \mathbf{U}_p is the nodal displacement vector of sample P , and $\bar{\mathbf{U}}$ is the mean of these displacement vectors.

Diagonalizing \mathbf{S} by performing a principal component analysis gives the experimental modes of variation ψ_i , 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 modes¹:

$$\mathbf{U} = \Psi \mathbf{b} = \sum_{i=1}^{3N} \psi_i b_i. \quad (3.15)$$

In this case it is the modes whose associated eigenvalues are *large* that are typically

¹This 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 (i.e. variation that resides in the nullspace of the training set) will not be represented.

most useful. This is because the directions of highest variability in the data are given by the eigenvectors of the largest eigenvalues. Also, these eigenvectors are the most reliably estimated [33].

Software that calculates and analyzes experimental modes is available from whitechapel.media.mit.edu in the file `/pub/face-recognition.tar.Z`.

3.3 The Connection

The connection between physical and experimental modes comes from the close relationship between mechanical and probabilistic prior models [60, 77, 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} .

To derive this relationship, we begin by interpreting Equation 3.6 and its solution

$$\hat{\mathbf{U}} = \mathbf{K}^{-1}\mathbf{R} \quad (3.16)$$

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

$$\Sigma_{\hat{\mathbf{U}}} = \mathbf{K}^{-1}\Sigma_{\mathbf{R}}(\mathbf{K}^{-1})^T \quad (3.17)$$

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

$$\Sigma_{\hat{\mathbf{U}}} = \mathbf{K}^{-1}\Sigma_{\mathbf{R}}\mathbf{K}^{-1}. \quad (3.18)$$

Under the assumption² that the elements of \mathbf{R} are independent and have common variance σ^2 ,

$$\Sigma_{\hat{\mathbf{U}}} = \sigma^2\mathbf{K}^{-2}. \quad (3.19)$$

We can form the sample covariance matrix $\mathbf{S} \approx \Sigma_{\hat{\mathbf{U}}}$ from a set of observations of $\hat{\mathbf{U}}$, then use Equation 3.19 to obtain the estimate

$$\mathbf{S} \approx \sigma^2\mathbf{K}^{-2}. \quad (3.20)$$

²The validity of this assumption is discussed in Section 8.8.1.

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 3.1.2 that diagonalized \mathbf{K} , we have

$$\begin{aligned}
 \Phi^T \mathbf{S} \Phi &\approx \sigma^2 \Phi^T \mathbf{K}^{-1} \mathbf{K}^{-1} \Phi \\
 &= \sigma^2 (\Phi^T \mathbf{K}^{-1} \Phi) (\Phi^T \mathbf{K}^{-1} \Phi) \\
 &= \sigma^2 \Lambda^{-1} \Lambda^{-1} \\
 &= \sigma^2 \begin{bmatrix} \frac{1}{\lambda_1^2} & & & & \\ & \frac{1}{\lambda_2^2} & & & \\ & & \ddots & & \\ & & & \ddots & \\ & & & & \frac{1}{\lambda_{3N}^2} \end{bmatrix} \tag{3.21}
 \end{aligned}$$

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.

3.4 Compound Modes

Consider the square shown at the left of Figure 3-2. Let us assume that the deformation we observe in the square is the result of two independent physical processes \mathbf{P}_1 and \mathbf{P}_2 , so that its total deformation is just a combination of two independent deformations. If we




Figure 3-2: Total deformed shape = original shape + deformation due to physical process P_1 + deformation due to physical process P_2 .

have good physical models for both processes, then we could use modal FEM 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 second situation is the one applicable to our problem.

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 U^{P_1} that it causes. Also assume that we do not have detailed knowledge of the physics of the second process, so we model its deformation U^{P_2} by collecting samples and using experimental modes. The total deformation U 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 3.13 and 3.15 we write the total deformation U as

$$U = U^{P_1} + U^{P_2} \quad (3.22)$$

$$= \Phi \mathbf{a} + \Psi \mathbf{b} \quad (3.23)$$

$$= \sum_{i=1}^{3N} \phi_i a_i + \sum_{i=1}^{3N} \psi_i b_i \quad (3.24)$$

$$= \sum_{i=1}^{6N} \gamma_i c_i \quad (3.25)$$

$$= \Gamma \mathbf{c} \quad (3.26)$$

where

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

and

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

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

3.5 Summary

In this chapter, two types of deformation modes were derived and illustrated. Physical modes come from finite element modeling and modal analysis, while experimental modes are obtained from sample data and principal component analysis. Both types of modes serve as a convenient basis set in which to represent shape deformation.

After deriving the connection between them, an example was given that illustrated how physical and experimental modes can be applied in a serial manner in order to represent shape deformation in terms of the compound modes of a system. This approach decomposes total observed deformation into modes of shape variation that are the result of different physical processes.

In the following chapter, this serial approach will be applied to neuropathology. The total deformation observed in a brain structure will be decomposed into physically-based nonpathological modes of shape variation, and experimentally-based pathological modes of shape variation. The nonpathological modes will be used to normalize brain structures for an individual's head shape, while the pathological modes will be used to characterize diseases.

Chapter 4

Application of Model to Neuropathology

In this chapter, our mathematical model is applied to structures within the brain. We show how the two types of modes can be applied in a serial manner in order to improve disease classification.

4.1 Head Shape Normalization

To characterize the global deformations due to head shape, we model the ICC as a linear elastic material, and then set up equations describing its behavior. One reason for using a physically-based model is that we can formulate an approximate physical model for the ICC as a whole. The accuracy of this physical model is discussed in Section 8.2. 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.

Our procedure for head shape normalization consists of three steps. First, a finite element model of the average ICC is created. This is accomplished by constructing an average ICC from our patient database, and then warping a finite element model to approximate its shape. Next, the same finite element model is warped to approximate the shape of a particular patient's ICC. Finally, using these two sets of warping coefficients, the cranial

contents of the patient are normalized for his or her particular head shape. These three steps are described in more detail in the following three sections.

4.1.1 FEM model of average ICC

In order to create an FEM model of the average ICC, we first construct a voxel-based average from patient data sets. This is done by automatically segmenting the ICC from each of the data sets [21], performing a rigid body alignment between all the extracted ICCs [31], and then averaging spatial occupancy over the aligned ICCs.

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

$$\mathbf{X}_A^{ICC} = [x_1, y_1, z_1, \dots, x_M, y_M, z_M]^T. \quad (4.1)$$

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:

$$\mathbf{f}_m = \mathbf{d}(\mathbf{x}_m, \text{FEM surface}) \quad (4.2)$$

where $\mathbf{d}(\cdot)$ is the displacement vector between the given point and the nearest point on the FEM surface. In general, a force \mathbf{f}_m will not act at an FEM node; however we can use the FEM interpolation functions \mathbf{H} to distribute the force to the surrounding nodes in a physically meaningful way. The load vector of Equation 3.6 is therefore constructed as

$$\mathbf{R}_A = \mathbf{H}^T \mathbf{F}_A. \quad (4.3)$$

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

Equation 3.6 can then be solved for the nodal displacements that give the FEM model a shape approximating that of the average ICC:

$$\hat{\mathbf{U}}_A^{ICC} = \mathbf{K}^{-1} \mathbf{R}_A. \quad (4.4)$$

These nodal displacements can be added to the original nodal positions to obtain the nodal




Figure 4-1: The average ICC (a) and several of the ICC's physical deformation modes (b-f).

positions of the average ICC:

$$\hat{\mathbf{X}}_A^{ICC} = \hat{\mathbf{X}} + \hat{\mathbf{U}}_A^{ICC}, \quad (4.5)$$

where $\hat{\mathbf{X}}$ is the original, undeformed nodal position vector of Equation 3.1.

We can then transform into the modal coordinate system via Equation 3.8:

$$\tilde{\mathbf{U}}_A^{ICC} = \mathbf{\Phi}^T \hat{\mathbf{U}}_A^{ICC} \quad (4.6)$$

where the modal coefficient vector $\tilde{\mathbf{U}}_A^{ICC}$ specifies how much of each deformation mode is contained in the shape of the average ICC. If only the modal coefficients $\tilde{\mathbf{U}}_A^{ICC}$ and not the nodal displacements $\hat{\mathbf{U}}_A^{ICC}$ are required, then we can skip solving for $\hat{\mathbf{U}}_A^{ICC}$, and instead solve directly for $\tilde{\mathbf{U}}_A^{ICC}$, using Equation 3.9. As mentioned earlier, because the system of equations is decoupled in the modal coordinate system, using Equation 3.9 directly is much faster.

Figure 4-1a shows the average ICC, while Figures 4-1(b-f) show several of its physical deformation modes. These modes represent various ways in which the ICC can deform; the ICC shape of a particular patient can be approximated as a linear combination of these and other modes that are not shown. Figure 4-1b, for example, is the basis function that

represents a front-to-back squishing along with a top-to-bottom stretching. Any patient's ICC that tended toward this shape would have a large projection onto this particular basis function.

In order to obtain additional computational advantages, these modes were computed using the “idealized modes” technique described in [57]. 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 [68] will recover modes that are closer to the actual ICC modes.

4.1.2 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:

$$\hat{\mathbf{U}}_p^{ICC} = \mathbf{K}^{-1}\mathbf{R}_p, \quad (4.7)$$

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

$$\hat{\mathbf{X}}_p^{ICC} = \hat{\mathbf{X}} + \hat{\mathbf{U}}_p^{ICC}. \quad (4.8)$$

Alternatively, we can recover the modal amplitudes $\hat{\mathbf{U}}_p^{ICC}$ via Equation 3.9.

Note that although each $\hat{\mathbf{U}}_p^{ICC}$ (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.

4.1.3 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 [84]. Because this will produce non-integer coordinates in the patient space, interpolation is necessary in order to calculate an intensity value for the voxel. Repeating this procedure for every voxel position in the average ICC space completely fills up that coordinate system with values from the ICC of patient p .

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

From average ICC space to parameter space

Given a voxel position \mathbf{x} in the average ICC coordinate system, the first step is to transform it into the parametric coordinate system. This can be done with Equation 3.1, using the nodal positions of the average ICC:

$$\mathbf{x}(\mathbf{r}) = \mathbf{H}(\mathbf{r})\hat{\mathbf{X}}_A^{ICC}. \quad (4.9)$$

Note, however, that Equation 4.9 must be inverted - given \mathbf{x} we need to find \mathbf{r} .

This is accomplished as follows. According to Equation 4.9, once the nodal positions $\hat{\mathbf{X}}_A^{ICC}$ are known, then \mathbf{x} is just a function of \mathbf{r} :

$$\mathbf{x} = \mathbf{g}(\mathbf{r}) \quad (4.10)$$

where \mathbf{g} is the system of three polynomials given in Equation 4.9. To solve for \mathbf{r} , \mathbf{g} must

be inverted:

$$\mathbf{r} = \mathbf{g}^{-1}(\mathbf{x}). \quad (4.11)$$

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

$$\mathbf{J}^k (\mathbf{x}^{k+1} - \mathbf{x}^k) + \mathbf{g}^k = \mathbf{0} \quad (4.12)$$

where \mathbf{J} is the Jacobian matrix defined in Equation 3.5, and k represents iteration [75, 30].

From parameter space to patient space

Once the parametric coordinates \mathbf{r} are known for \mathbf{x} , they must be converted into the coordinate system of patient p . Once again this is accomplished using Equation 3.1, but this time we use it directly, with the known \mathbf{r} and the ICC nodal positions of patient p :

$$\mathbf{x}'(\mathbf{r}) = \mathbf{H}(\mathbf{r})\hat{\mathbf{X}}_p^{ICC}. \quad (4.13)$$

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.

Modal warping

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

In the modal coordinate system, instead of Equation 3.1, the interpolation and warping is done by combining Equations 3.2 and 3.8:

$$\mathbf{u}(\mathbf{r}) = \mathbf{H}(\mathbf{r})\Phi\tilde{\mathbf{U}}. \quad (4.14)$$

To further increase computational efficiency, the polynomial deformations of Equation 4.14 are approximated by a 3×3 *modal deformation matrix* $\mathbf{D}(\mathbf{r}, \tilde{\mathbf{U}})$ [57], which is used to map from parametric to element coordinates:

$$\mathbf{x}(\mathbf{r}) = \mathbf{D}(\mathbf{r}, \tilde{\mathbf{U}})\mathbf{r}. \quad (4.15)$$

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

$$\mathbf{x}(\mathbf{r}) = \mathbf{D}(\mathbf{r}, \tilde{\mathbf{U}}_A^{ICC})\mathbf{r} \quad (4.16)$$

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

$$\mathbf{x}'(\mathbf{r}) = \mathbf{D}(\mathbf{r}, \tilde{\mathbf{U}}_p^{ICC})\mathbf{r} \quad (4.17)$$

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

Equation 4.16, Equation 4.17, 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.

4.1.4 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 Figures 4-2 and 4-3. Figures 4-2a and 4-2b show the average ICC and ventricles, while Figures 4-2c and 4-2d show the ICC and ventricles of one of the healthy subjects. This particular subject's ICC is larger than average, particularly in the front-to-back direction (left-to-right in the figure). This ICC shape difference is propagated down to the ventricles, where we see similar shape differences between the two ventricular systems. Calculating the ICC physical deformation modes that warp this subject's ICC to the shape of the average ICC, and then applying that warping to the subject's ventricles, produces the warped ventricles in Figure 4-2e. As can be seen, the warping has increased the shape similarity between patient and average ventricles.




Figure 4-2: 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).

4.2 Characterization of Disease States

4.2.1 Shape representation

As described in Equations 1.1-1.3, our shape representation for a brain structure of a particular patient p is a list of displacements \mathbf{U}_p away from a set of average point positions \mathbf{X}_A . These average positions are found by first constructing a volumetric average of the structure from a group of patient data sets, and then extracting the surface voxels.

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




Figure 4-3: Same structures as in Figure 4-2, but viewed from above. The front-to-back direction is now bottom-to-top. See the caption of Figure 4-2 for details.

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 U_p^D of Equation 1.2.

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

4.2.2 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 maxi-

mum variability inherent in a data set. When applied to 2-D outline or 3-D surface data, the principal components are called the *eigenshapes* of the structure under study. Unlike physical modes, eigenshapes are derived solely from a data set, without the aid of an underlying physical model.

The eigenshapes of our data sets are found using Equation 3.14. From our P patients, each with a $3V \times 1$ pathological displacement vector \mathbf{U}_p^D , the sample covariance matrix is formed as

$$\mathbf{S} = \frac{1}{P-1} \sum_{p=1}^P (\mathbf{U}_p^D - \bar{\mathbf{U}}^D)(\mathbf{U}_p^D - \bar{\mathbf{U}}^D)^T, \quad (4.18)$$

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:

$$\mathbf{U}_p^D = \Psi \mathbf{b}_p, \quad (4.19)$$

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 therefore 24 principal components. Because it has been our experience that approximately the top $\frac{1}{4}$ principal components of a data set are reliably estimated, the projections onto the first six principal components of each data set will serve as the input to the classifier.

We now turn to designing an actual classifier. If the feature vectors of each class are characterized by a Gaussian density function, then a set of discriminant functions that minimizes the probability of error is given by^{1,2}

$$g_i(\mathbf{b}_p) = -(\mathbf{b}_p - \mathbf{m}_i)^T \Sigma_i^{-1} (\mathbf{b}_p - \mathbf{m}_i) - \ln |\Sigma_i|. \quad (4.20)$$

¹Equation 4.20 assumes that the prior probabilities of all classes are equal.

²If the underlying distributions are not Gaussian, then these discriminant functions do not necessarily minimize the probability of error. They can, however, be interpreted as the discriminant functions that are best matched to the first and second moment statistics of the random vectors \mathbf{b}_p [81].

In our case \mathbf{b}_p is a vector of the projections of patient p onto the first six principal components (see Equation 4.19), \mathbf{m}_i is the 6×1 mean vector computed by averaging the projections over disease class i , and Σ_i is the 6×6 covariance matrix of 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 . Because Equations 4.20 are quadratic functions of \mathbf{b}_p , this form of classifier is referred to as a *Gaussian quadratic classifier* [81].

The covariance matrix of each class, Σ_i , must be estimated from the samples contained within that class. In our two data sets, however, we have only between 7 and 13 members per disease state. This is not enough information to reliably estimate a 6×6 covariance matrix. The quadratic classifier will consequently be overparametrized, with the result that the training set will not be a good predictor for new cases.

Instead we will use a common covariance matrix, which can be estimated using all 25 data samples. When a common Σ is substituted for Σ_i in each of the Equations 4.20, the final term is the same for all the discriminant functions and can thus be dropped. By expanding what remains, more common terms appear. When these are dropped, the set of discriminant functions simplifies to

$$g_i(\mathbf{b}_p) = 2\mathbf{m}_i^T \Sigma^{-1} \mathbf{b}_p - \mathbf{m}_i^T \Sigma^{-1} \mathbf{m}_i \quad (4.21)$$

which are linear functions of \mathbf{b}_p . The resulting classifier is thus called a *Gaussian linear classifier* [81]. This is what will be applied to our data.

In order to separate the training stage from the classification stage, we use the *leave-one-out* [33] 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 are:

- For each of the 25 patients:
 1. Use the \mathbf{b}_p vectors of the remaining 24 patients to estimate the common covariance matrix Σ and the individual class means \mathbf{m}_i .

2. Using the estimates Σ and \mathbf{m}_i from Step 1, and the \mathbf{b}_p of the current patient to be classified, compute the linear discriminant functions g_i of Equations 4.21.
 3. Choose as the winner the disease state i with the largest g_i . Classify the current patient as having this disease.
- After doing the above for all 25 patients, count up the number of patients that were correctly classified. Divide this number by 25 to obtain the classification rate.

Chapter 5

Experiments with Schizophrenia

5.1 Medical Background

Schizophrenia is a disabling mental disorder that affects approximately 1% of the population. Its symptoms include thought disorders such as disintegrated thinking and hallucinations, as well as certain mood and behavior changes. The etiology of schizophrenia is poorly understood, although it is probably a combination of genetic, biological, and social factors. Prognosis varies, with treatment consisting of antipsychotic medication and psychotherapy [1].

The presence of psychosis in an individual can generally be established on the basis of clinical symptoms alone. However, different types of psychoses can present very similar symptoms, making the differential diagnosis between psychoses difficult. As an example, making the diagnosis of schizophrenia instead of manic-depressive psychosis can require additional factors such as a family history of schizophrenia and a poor work history.

By investigating the structural changes caused by different psychoses, our hope is that shape information derived from a patient's medical scans can provide clinicians with additional discriminating factors. Although the present data set contains just schizophrenics and normals, the same concepts illustrated in this section can be applied to future data sets consisting of different types of psychoses.

In addition to its potential for improving diagnosis, examining shape changes can also lead to an improved understanding of both the etiology and the effects of schizophrenia. As mentioned earlier, previous research on neuropathological structural changes has focused primarily on volume. The pathological deformation modes that we calculate in this thesis

extend this analysis to shape changes in general, and provide a unique, natural coordinate system in which to express and view these changes.

5.2 The Putamen Data Set

Thirteen schizophrenic patients and 12 healthy control subjects, matched for gender, age, and handedness, underwent an MR brain scan. As part of an ongoing volumetric study [69, 40], the basal ganglia were manually segmented from these scans. Figures 5-1 and 5-2 show the resulting data set of 25 left putamen, viewed from the side and front, respectively.

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.

5.3 Application of Our Method

First, using the techniques of Section 4.1, the cranial contents of each patient were warped in order to normalize the database for head shape. Next, the pathological displacement vector U_p^D was calculated for each patient's left putamen, as described in Section 4.2.1. A principal component analysis of the 25 displacement vectors was then performed.


5.4 Experimental Results

Figure 5-3 shows the first two putamen principal components. The first mode (Figure 5-3(a-e)) is a contrast between the size of the top of the putamen and the size of the bottom. This can be seen by noting that Figure 5-3a, with a negative modal amplitude, shows the top of the putamen expanded relative to the undeformed putamen (zero modal amplitude) shown in Figure 5-3c. Figure 5-3e, with a positive modal amplitude, shows a (slight) expansion near the bottom of the putamen.


The second mode (Figure 5-3(f-j)) contrasts the size of the front to the size of the back of the putamen (left and right in the figure). These deformations are easier to see than those of the first mode. The deformation associated with a negative modal amplitude, illustrated

in Figure 5-3f, is an expansion of the front of the putamen. Likewise Figure 5-3j, with a positive modal amplitude, shows an expansion at the back of the putamen.

Next we input the top six principal components (modes) into our Gaussian linear classifier. Table II shows the results, along with the classification rate using just putamen volume. A 12% improvement in the classification rate occurs when using the putamen principal components instead of just volume. With our small sample size, this is not a statistically significant result. Instead the data only suggests that the method is working. Topics related to statistical significance are discussed in Section 8.3.




(a)




(b)

Figure 5-1: The left putamen data set, with the patients facing to the right. (a) Healthy control subjects. (b) Patients with schizophrenia.



(a)



(b)

Figure 5-2: The left putamen data set, with the patients facing forwards. (a) Healthy control subjects. (b) Patients with schizophrenia.




Figure 5-3: (a)-(e) First principal component of the putamen data set. The amplitude of the mode is increasing from (a) to (e). (f)-(j) Second principal component of the putamen data set. The amplitude of the mode is increasing from (f) to (j).

TABLE II
Gaussian Linear Classification Rates
Putamen Database

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

Chapter 6

Experiments with Ventricular Disorders

6.1 Medical Background

Alzheimer's disease (AD) is a degenerative disorder of the nervous system, whose initial major symptom is the gradual development of forgetfulness. As the disease progresses, the dementia becomes more and more severe, finally ending with the patient's death. Normal-pressure hydrocephalus (NPH), on the other hand, is a disorder of the cerebrospinal fluid. While the exact causes can be quite varied, they all involve problems in CSF flow and possibly absorption. Like AD, one of the major clinical symptoms of NPH is dementia [59].

Both disorders cause ventricular enlargement. In NPH, it comes from the disruption of normal CSF flow characteristics. In AD, ventricular enlargement is a secondary effect. The brain atrophies due to neuronal destruction, and because the cranium remains fixed, the amount of CSF in both the subarachnoid space and the ventricles increases.

Unlike AD, which is untreatable, NPH can often be treated surgically by implanting a shunt in order to redirect CSF flow. Because of the invasiveness of this procedure, differential diagnosis between AD and NPH is critical. Since both cause dementia and ventricular enlargement, measuring ventricular volume is not always helpful for differential diagnosis. Therefore quantitative measures of characteristic shape changes may assist in differential diagnosis.

As with schizophrenia, there are other possible uses of our method in addition to serving

as a diagnostic aid. By analyzing the shape deformations caused by AD and NPH, insights into the causes and effects of these poorly-understood disorders can hopefully be made.

6.2 The Ventricular Data Set

Nine patients with Alzheimer's disease, 7 patients with normal-pressure hydrocephalus, and 9 healthy control subjects, all matched for age, underwent an MR brain scan as part of a previous volumetric study [53] in which the ventricles were segmented using a semi-automatic procedure. Figure 6-1 shows the resulting data set of 25 ventricular systems. As can be seen, there is a greater difference between the three ventricular classes than there was between the two putamen classes shown in Figures 5-1 and 5-2.


6.3 Application of Our Method

As with the schizophrenic data set, the first step was a rigid body alignment of the 25 ICCs [31]. Next the ICCs were used to perform modal fitting [57], and then the recovered modal coefficients were used to warp each patient's cranial contents to a standard ICC shape. The principal components of both the unwarped and warped ventricle surface voxel displacements were then computed. Because we again have 25 patients, the top six eigenvectors will be the most reliably estimated, and we will use just these in our analysis.


6.4 Experimental Results

Figure 6-2 shows the first two principal components. The first mode (Figure 6-2(a-e)) is just a measure of the overall size of the ventricles, while the second (Figure 6-2(f-j)) is the degree of extension of the posterior horns of the lateral ventricles.


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



(a)



(b)



(c)

Figure 6-1: The ventricular data set. (a) Healthy control subjects. (b) Patients with Alzheimer's disease. (c) Patients with normal-pressure hydrocephalus.




Figure 6-2: (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).

TABLE III
Gaussian Linear Classification Rates
Ventricle Database

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

Chapter 7

Diagnostic Tool

In Chapters 5 and 6, patients were normalized for head shape, and then different brain structures underwent a statistical shape decomposition. The resulting deformation modes were first displayed to indicate the sorts of deformations caused by neuropathologies, and then input into a linear classifier in order to perform shape-based diagnosis.

Instead of using shape features, automatic diagnosis can use other features derived from images, or non-image data such as clinical symptoms. Whatever the information used, however, removing the human doctor from the decision-making process should not be the goal. Rather, the results of any computer-based analysis should serve as an additional source of information for the clinician. One such form of computer-assisted diagnosis would be a database tool that allows for efficient searches based on medical image *content*.

7.1 Content-Based Searches

Medical images are used by clinicians for diagnosis and for planning surgery and radiation therapy, by researchers in attempting to understand normal and pathological morphology, and by educators in teaching both anatomy and pathology. These images come from an ever-increasing variety of sources, in ever-increasing amounts. Currently, in major hospitals, on the order of one terabyte of medical imagery is produced every year [86].

While the organization and retrieval of medical images is in itself an important and actively researched problem [42, 62, 51, 50, 61, 43, 11, 38, 39, 5, 74, 79, 16, 41, 45, 49, 48], it has so far focused primarily on text-based searches (for example, by patient name), and on the construction of the appropriate hospital computer networks. Commonly called picture

archiving and communication systems (PACS), the main goal of these projects is to provide clinicians with immediate access to images, thereby improving diagnosis and patient care.

There are many applications, however, in which it is more appropriate to search a medical image database using the *content* of the images themselves, rather than just the accompanying text annotation. A radiologist, for example, may want to see previous patient studies in which the liver “looks like” the liver of a current patient. By then consulting the final diagnoses of those previous patients, the radiologist may be able to improve the diagnosis of the current patient. Without knowing the patient names in advance, there is currently no way for the radiologist to acquire this sort of information.

A recent workshop [86] co-sponsored by the National Cancer Institute, the Agency for Health Care Policy and Research, and the National Library of Medicine focused on the challenges presented by medical image databases. Among their specific suggestions were:

- In addition to textual indexing, there should be a means for content-based indexing.
- A descriptive language for the geometric content of medical images should be developed.
- Special tools that allow for flexible shape matching will be needed. Image databases are needed that can retrieve and sort images by shape and geometry in addition to conventional text indexing.
- The search and analysis processes should be fast enough for supervision by a human investigator.
- The user must be able to define which relations and which similarity measures are appropriate for the problem under consideration.
- By exploring a database, the user will collect images that are “interesting”. These will suggest new formulations, so that the user will need to alter the search in a customized way and on the fly.

These issues all involve searching through a medical database using image content, and being able to do so in an efficient, interactive manner. However, general methods of searching through image databases based on image content do not yet exist. The main problem is that it is unclear what image features ought to be used as similarity metrics.

In the case of a text database, the database elements themselves - strings of characters - can be directly compared in the search process. Since most often the bits of one image are not directly comparable to the bits of another image (even when the images contain similar objects), direct comparison is not a viable option for image databases. Rather, some sort of high-level features must be derived from the images before performing database search [65].

7.2 Photobook

We address several of the above issues by applying the Photobook image database tool [56] to the two medical data sets used in this thesis. The key contributions of Photobook are that it searches a database based on image *content*, and that it does so in real-time.

Photobook is a general image database tool that has been applied to faces, texture, and 2-D shape. As will be seen, the underlying functionality and similarity metrics used by Photobook when searching databases using image content extend naturally to the problems encountered in the medical domain. Although the experiments in the following sections were run on our two relatively small databases of 25 patients each, Photobook has been shown to maintain its real-time speed on a database of over 7500 face images [56].

7.3 Demonstration

To demonstrate Photobook's utility, we ran it using our two medical image databases. The first database consists of the 25 patients with cerebral ventricular disorders. Figure 7-1 shows this database as initially displayed in Photobook. To each side of the displayed images are user interface buttons and menus, illustrating some of the functionality of Photobook. Below the images is the Search filter field, which allows the user to implement text-based searches instead of or in addition to content-based searches.

Figure 7-2 is a magnified version of Figure 7-1, showing just the images and the Search filter field. Note that the patients with normal-pressure hydrocephalus (NPH) tend to have the largest ventricles, the patients with Alzheimer's disease (AD) have mid-sized ventricles, and the normal control subjects have the smallest ventricles.

For the first example, we selected one of the NPH patients, and had Photobook rank all the other patients based on a shape similarity metric. For this metric we used Euclidean

distance in the feature space defined by the first six principal components of normalized ventricular shape (the first six components of \mathbf{b}_p of Equation 4.19). Figure 7-3 shows this ranking. As can be seen, Photobook chose four of the other NPH patients as the most similar to the patient selected.

It is also useful to be able to combine text-based searches with content-based searches. If the patient under study is female, for instance, the clinician may want to search only for females in the database. Photobook allows for these types of combinations. The text to search on is entered at the bottom, in the Search filter field. Several of these can be applied at once, in various boolean combinations. These text searches can then be combined with a content-based search. Figure 7-4 shows the result when the user selects one of the AD patients, and requests that only females in the database be considered when searching.

The next set of demonstrations uses the putamen database. Figure 7-5 shows this database as initially displayed by Photobook. As mentioned earlier, the putamen of schizophrenia patients tend to be slightly larger than those of healthy controls.

First we selected one of the smaller healthy putamen, and had Photobook rank all the other patients. As before, the shape similarity metric used was Euclidean distance as measured in the feature space defined by the first six principal components of normalized putamen shape. As Figure 7-6 shows, the top three and the first six out of seven putamen chosen were those of other healthy subjects.

For the next example, one of the larger putamen from a patient with schizophrenia was selected. Again, Photobook correctly found putamen belonging to the same class as the patient selected. This result is shown in Figure 7-7.

Finally, we selected another schizophrenia patient, but one with a relatively small putamen. This caused Photobook to incorrectly choose healthy subjects as being the most similar to the patient we selected. Photobook's ranking in this case is shown in Figure 7-8.

7.4 Experiments

By selecting each patient in turn from a database, and examining the resulting Photobook rankings, more quantitative experiments can be run. The most straightforward experiment is to simply record the class membership (diagnosis) of the patient that Photobook picks as being the most similar to the one selected. By repeating this procedure for each patient

in the database, and tallying the results, a classification rate is obtained.

Because Photobook uses Euclidean distance as its similarity metric, the above procedure is equivalent to *nearest neighbor classification*. This type of classification, unlike the Gaussian approach used earlier, is nonparametric since it assumes nothing about the mathematical form of the underlying class probability distributions.

Table IV shows the results of applying this procedure to the putamen database. For comparison, the classification rate obtained using just putamen volume is also shown. A 12% improvement occurs when using principal components instead of volume.

For further comparison, Table II from Chapter 5 is repeated here. This table shows the classification accuracy obtained when a Gaussian linear classifier is applied to the putamen database. As can be seen, the nearest neighbor algorithm produces lower classification rates for both volume and principal components.

TABLE IV
Nearest Neighbor Classification Rates
Putamen Database

Feature(s)	%Correct
Putamen volume	52
Putamen principal components	64

TABLE II
Gaussian Linear Classification Rates
Putamen Database

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

Table V shows the classification rates obtained when the nearest neighbor algorithm is applied to the ventricle database. The results of using a Gaussian linear classifier are also repeated here, in Table III. In this case, using the nearest neighbor algorithm has increased the classification rate for the principal components, and decreased it for ventricular volume.

TABLE V
Nearest Neighbor Classification Rates
Ventricle Database

Feature(s)	%Correct
Ventricle volume	68
Ventricle principal components	92

TABLE III
Gaussian Linear Classification Rates
Ventricle Database

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

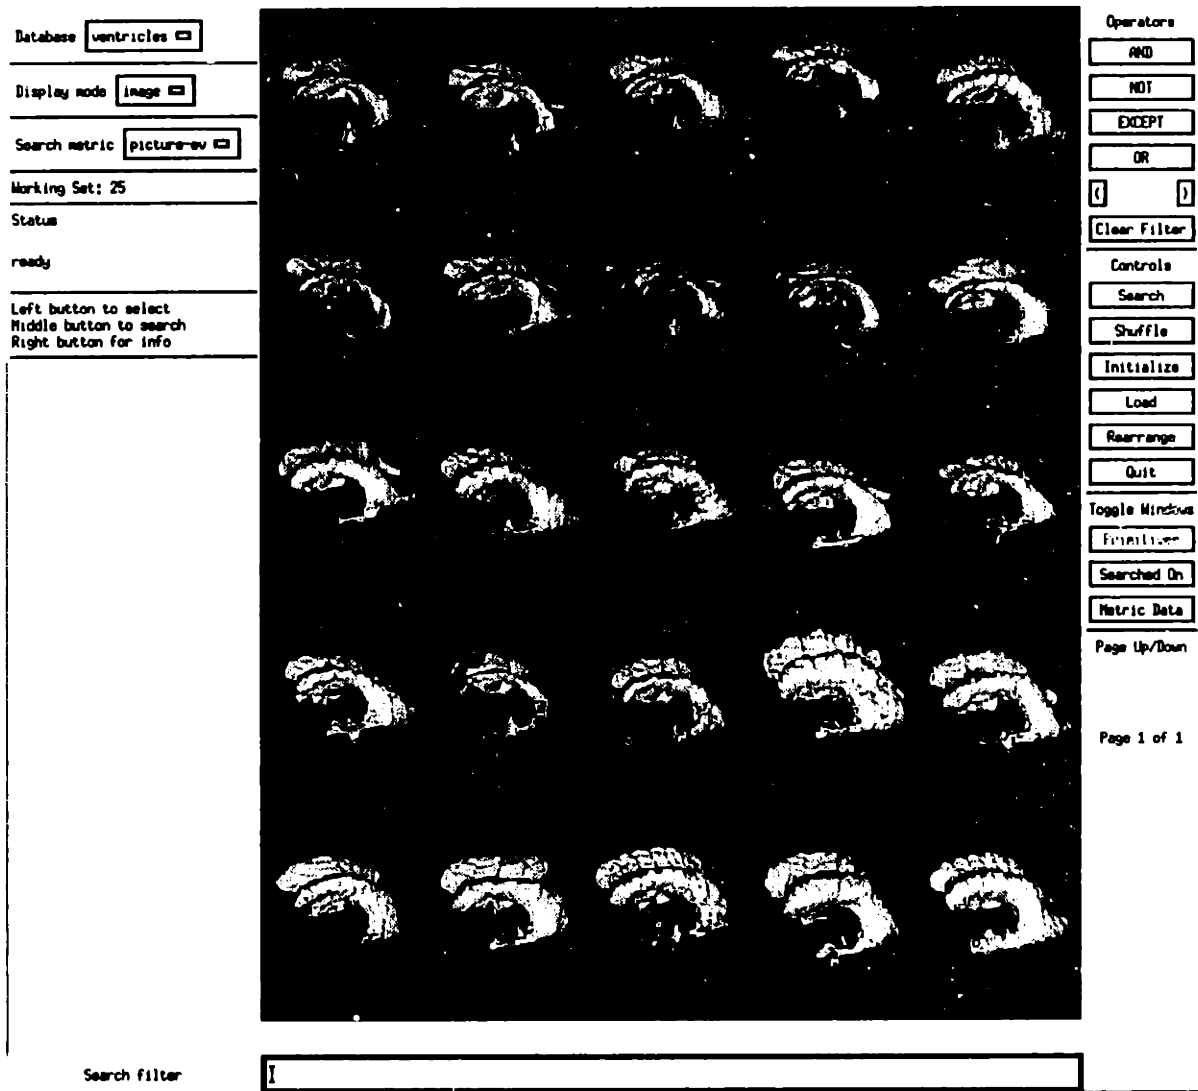


Figure 7-1: Initial configuration of the ventricular database in Photobook.

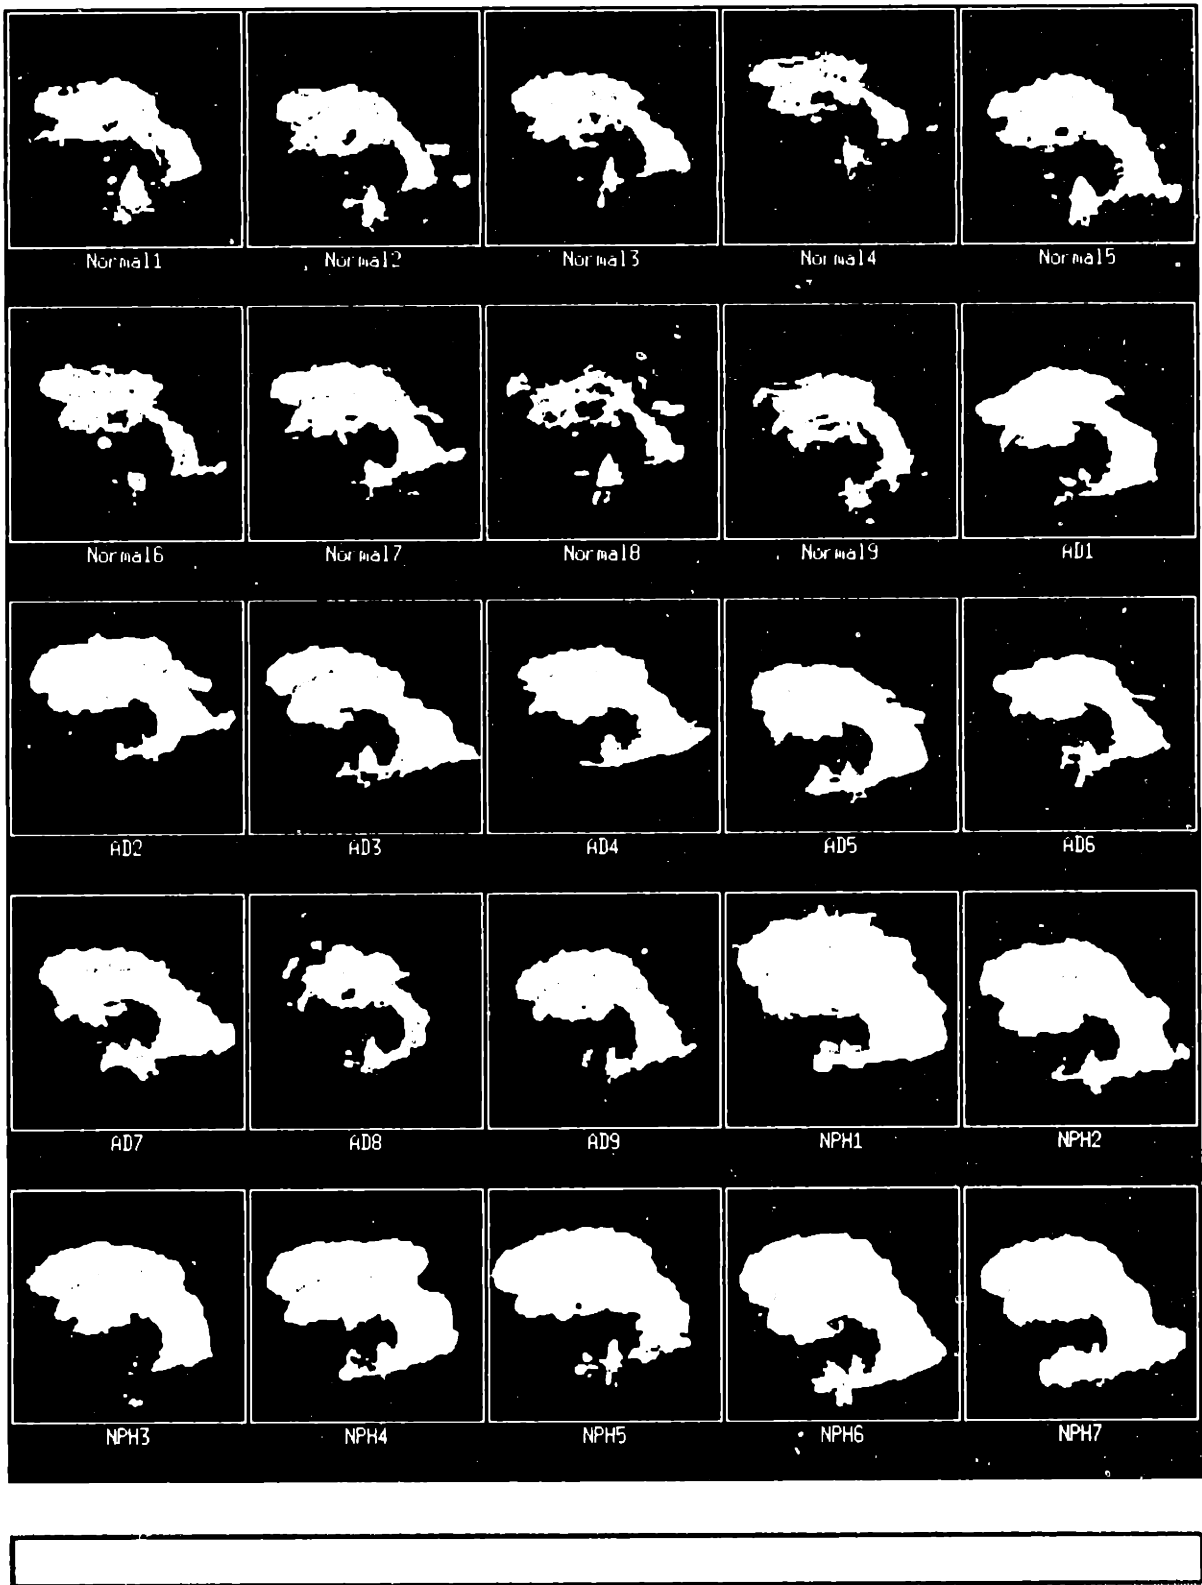
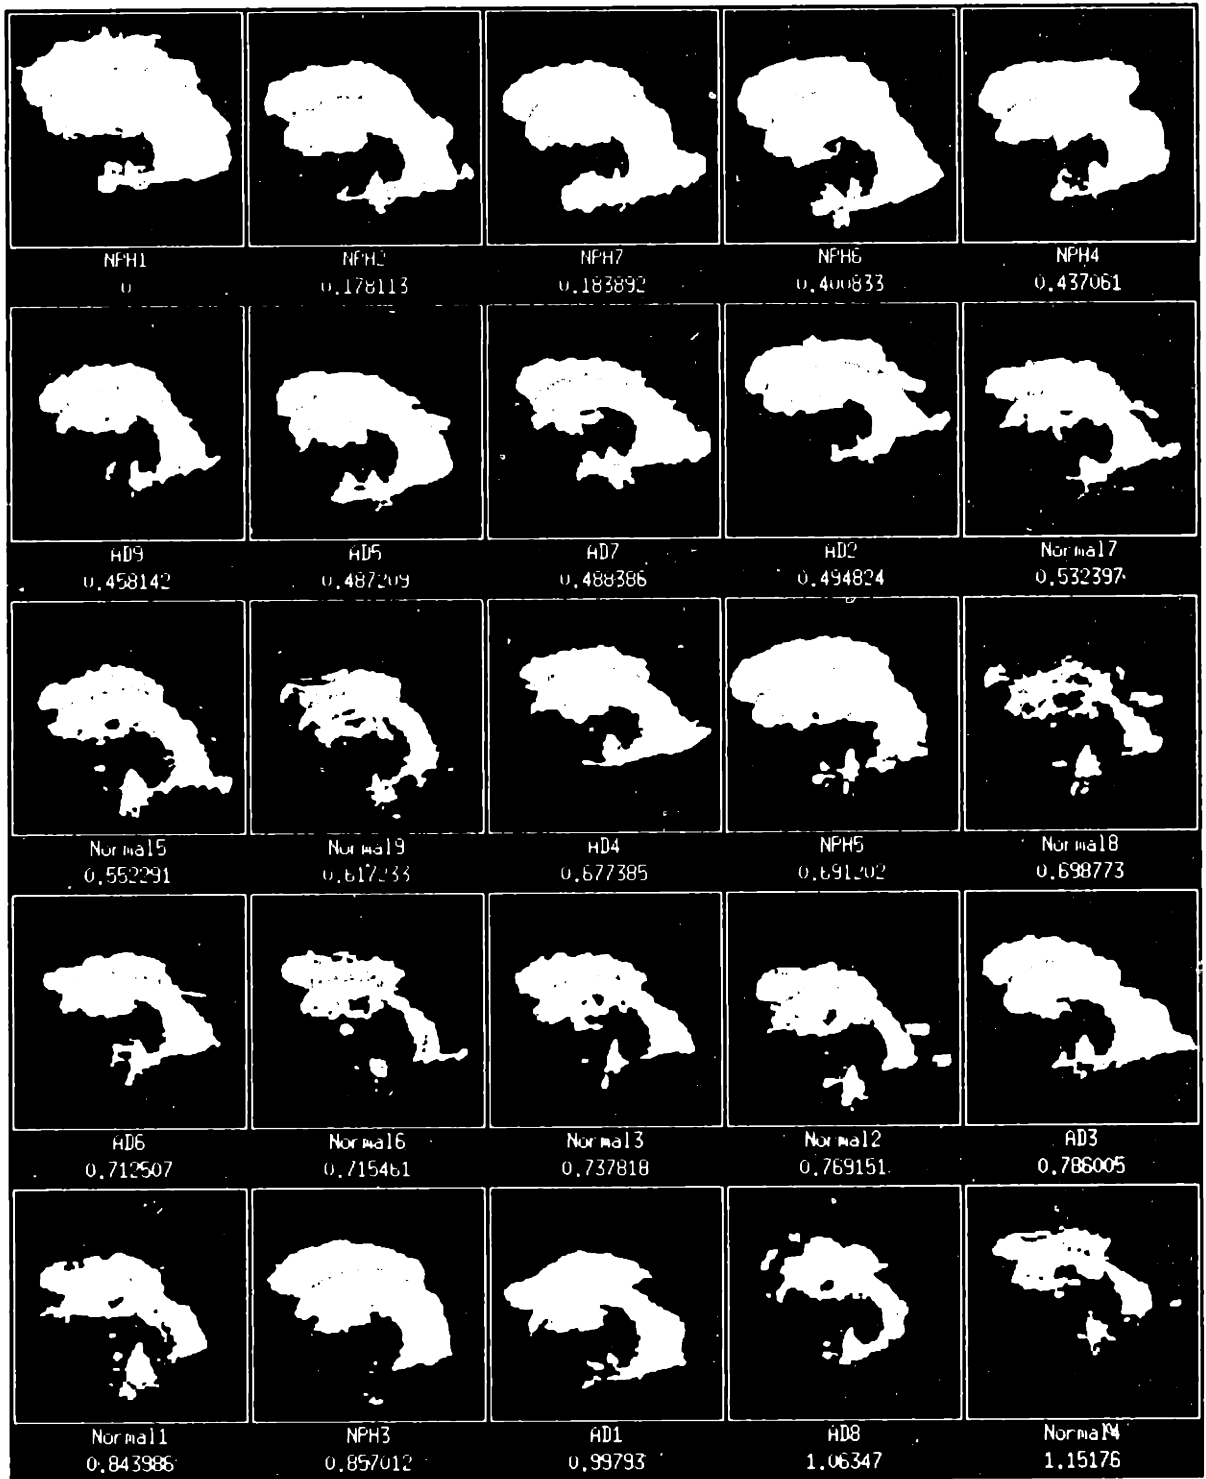


Figure 7.2: Magnified version of Figure 7.1, showing just the database images and the Search filter field. The label under each image indicates the disease the patient has been diagnosed with. Normal indicates a control subject, AD stands for Alzheimer's disease, and NPH stands for normal pressure hydrocephalus.



I

Figure 7.3: Ranking of the patients using a shape similarity measure, after the user has selected NPH1. The similarity decreases left to right and top to bottom. Photobook has picked four of the other NPH patients as the most similar to the one selected.



















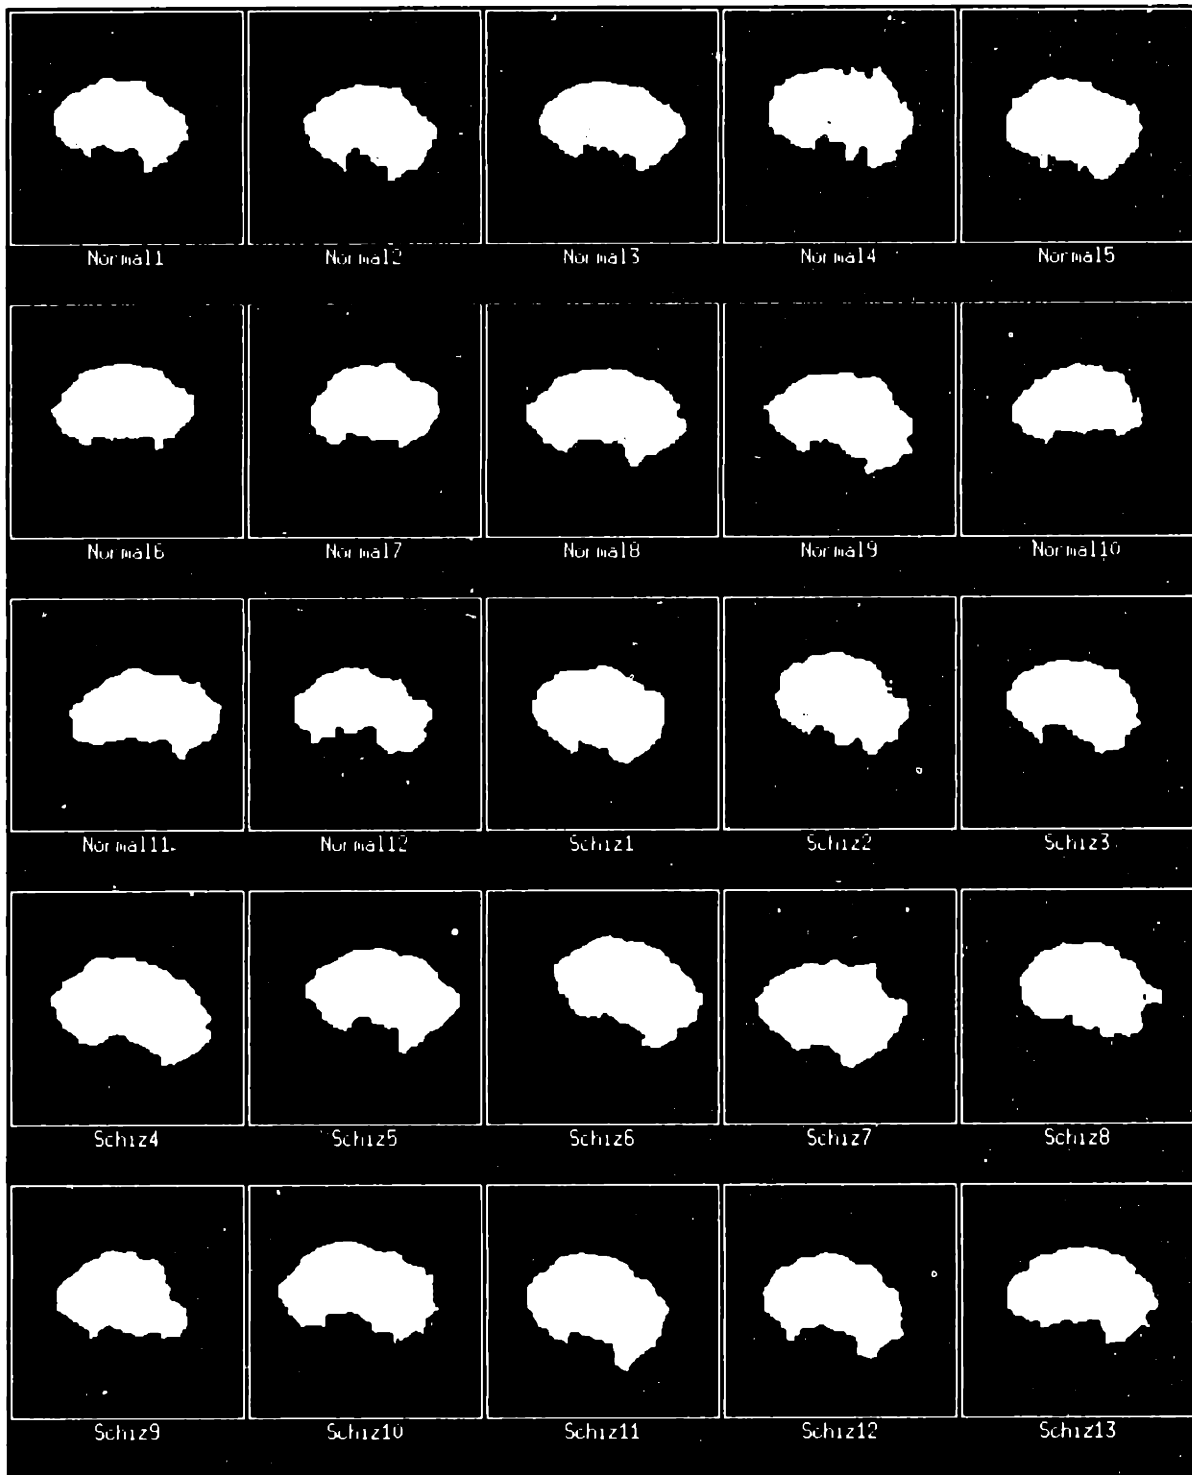
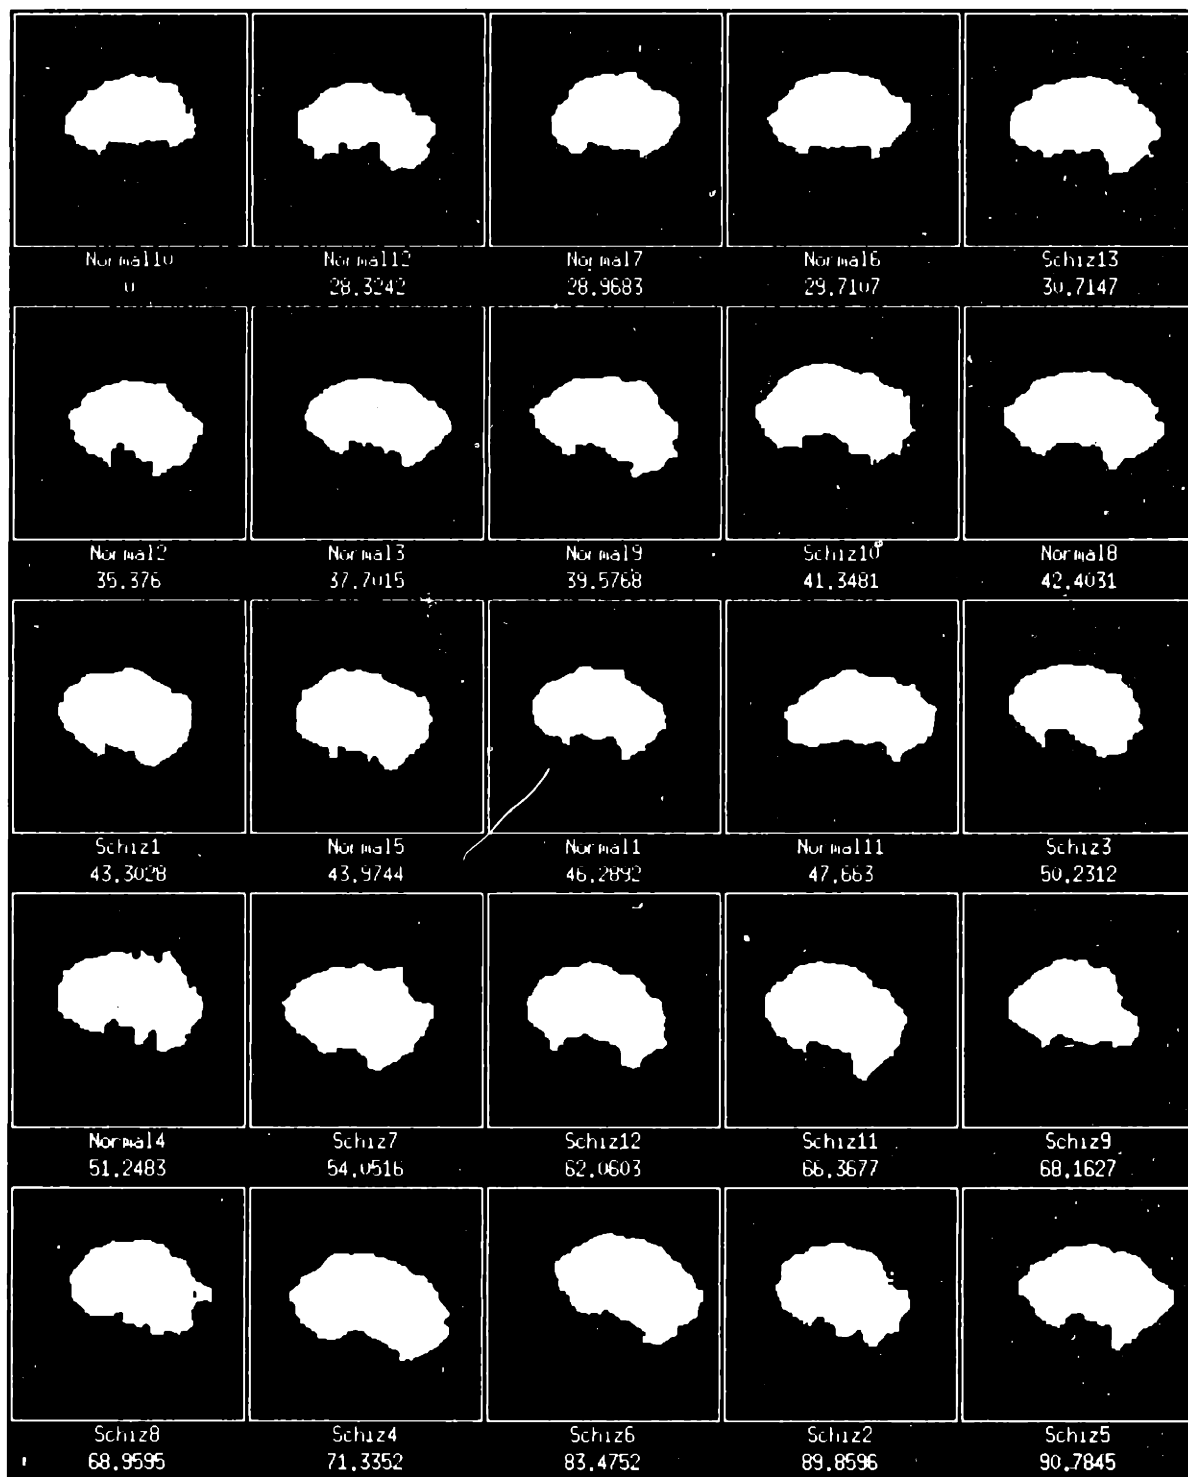
				
AD2 0	AD6 0.377725	Normal17 0.382879	AD9 0.397567	AD5 0.46351
				
NPH5 0.466009	NPH7 0.484665	AD7 0.494493	NPH1 0.494824	Normal15 0.575993
				
Normal18 0.646959	AD4 0.669208	Normal13 0.679122	NPH3 0.776139	Normal11 0.790438
				
AD3 0.906771	Normal14 0.950334	AD8 1.10258		
"female"				

Figure 7-4: Ranking of just the female patients, after the user has selected AD2. By entering the text string "female" in the Search filter field, the user has told Photobook to consider only the female patients in the database.



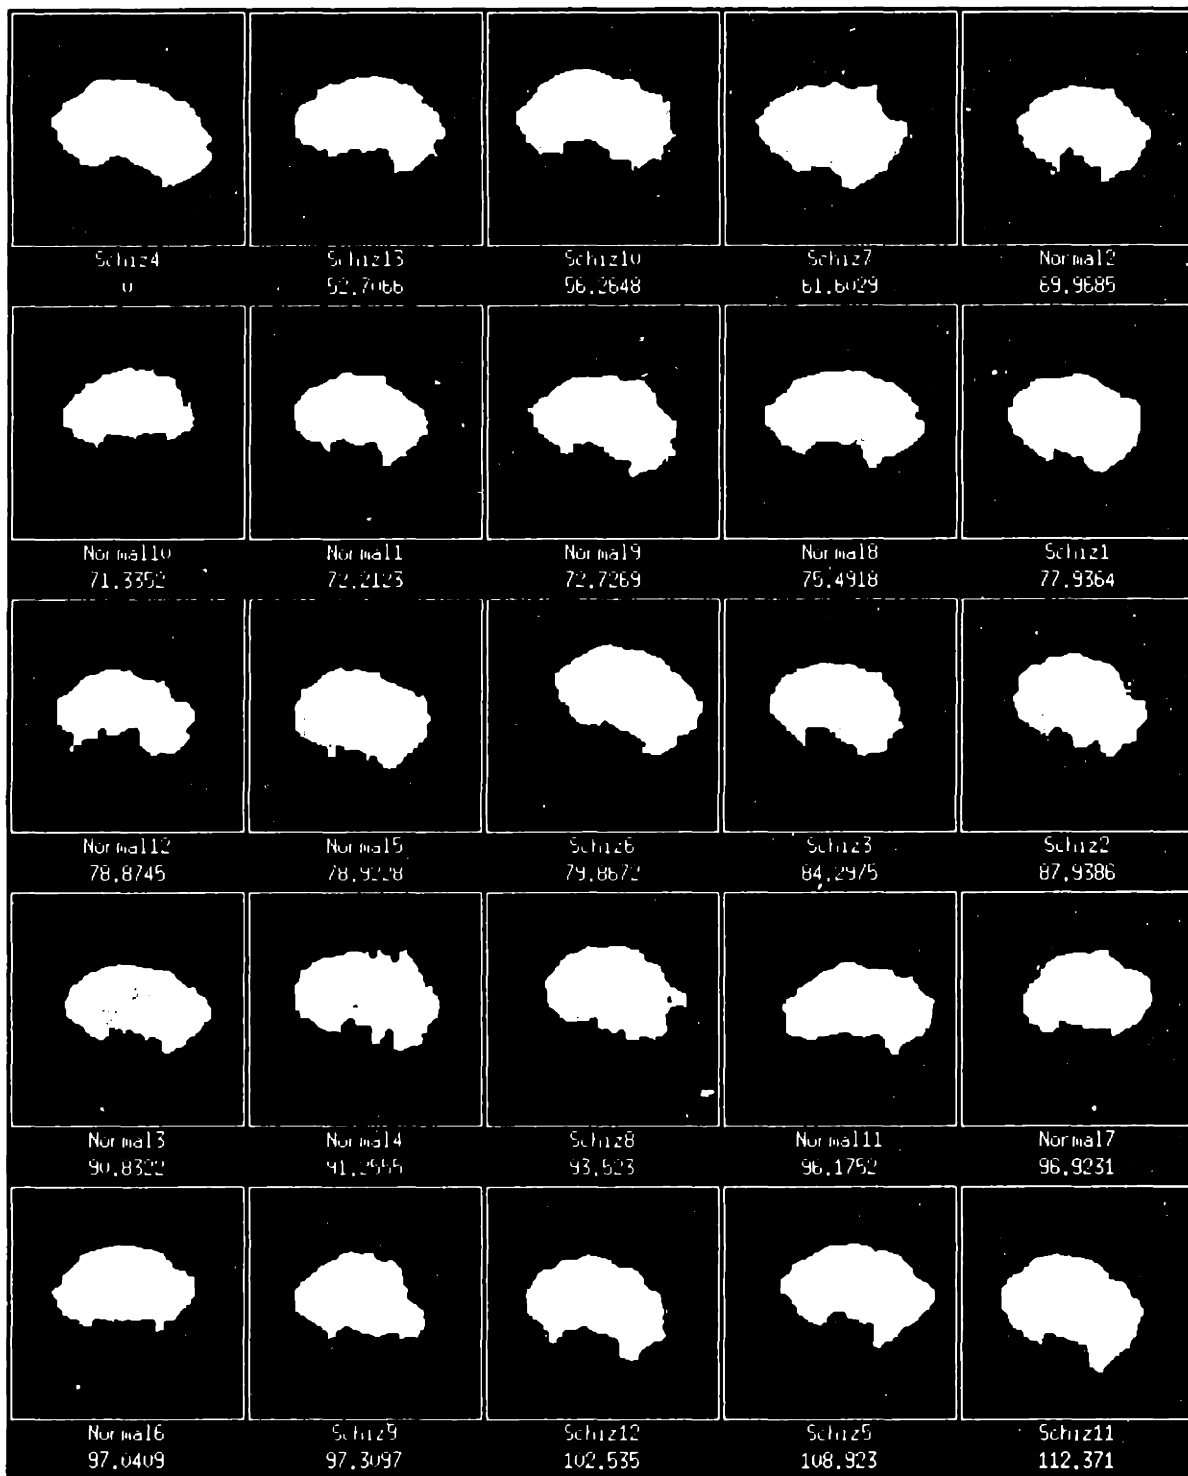
I

Figure 7.5: Initial configuration of the putamen database in Photobook. Normal indicates a control subject, and Schiz stands for schizophrenia.



I

Figure 7.6: Ranking of the patients after the user has selected Normal10. Photobook has picked three of the other healthy subjects as the most similar to the one selected.



I

Figure 7.7: Ranking of the patients after the user has selected Schiz1. Photobook has picked three of the other schizophrenia patients as the most similar to the one selected.

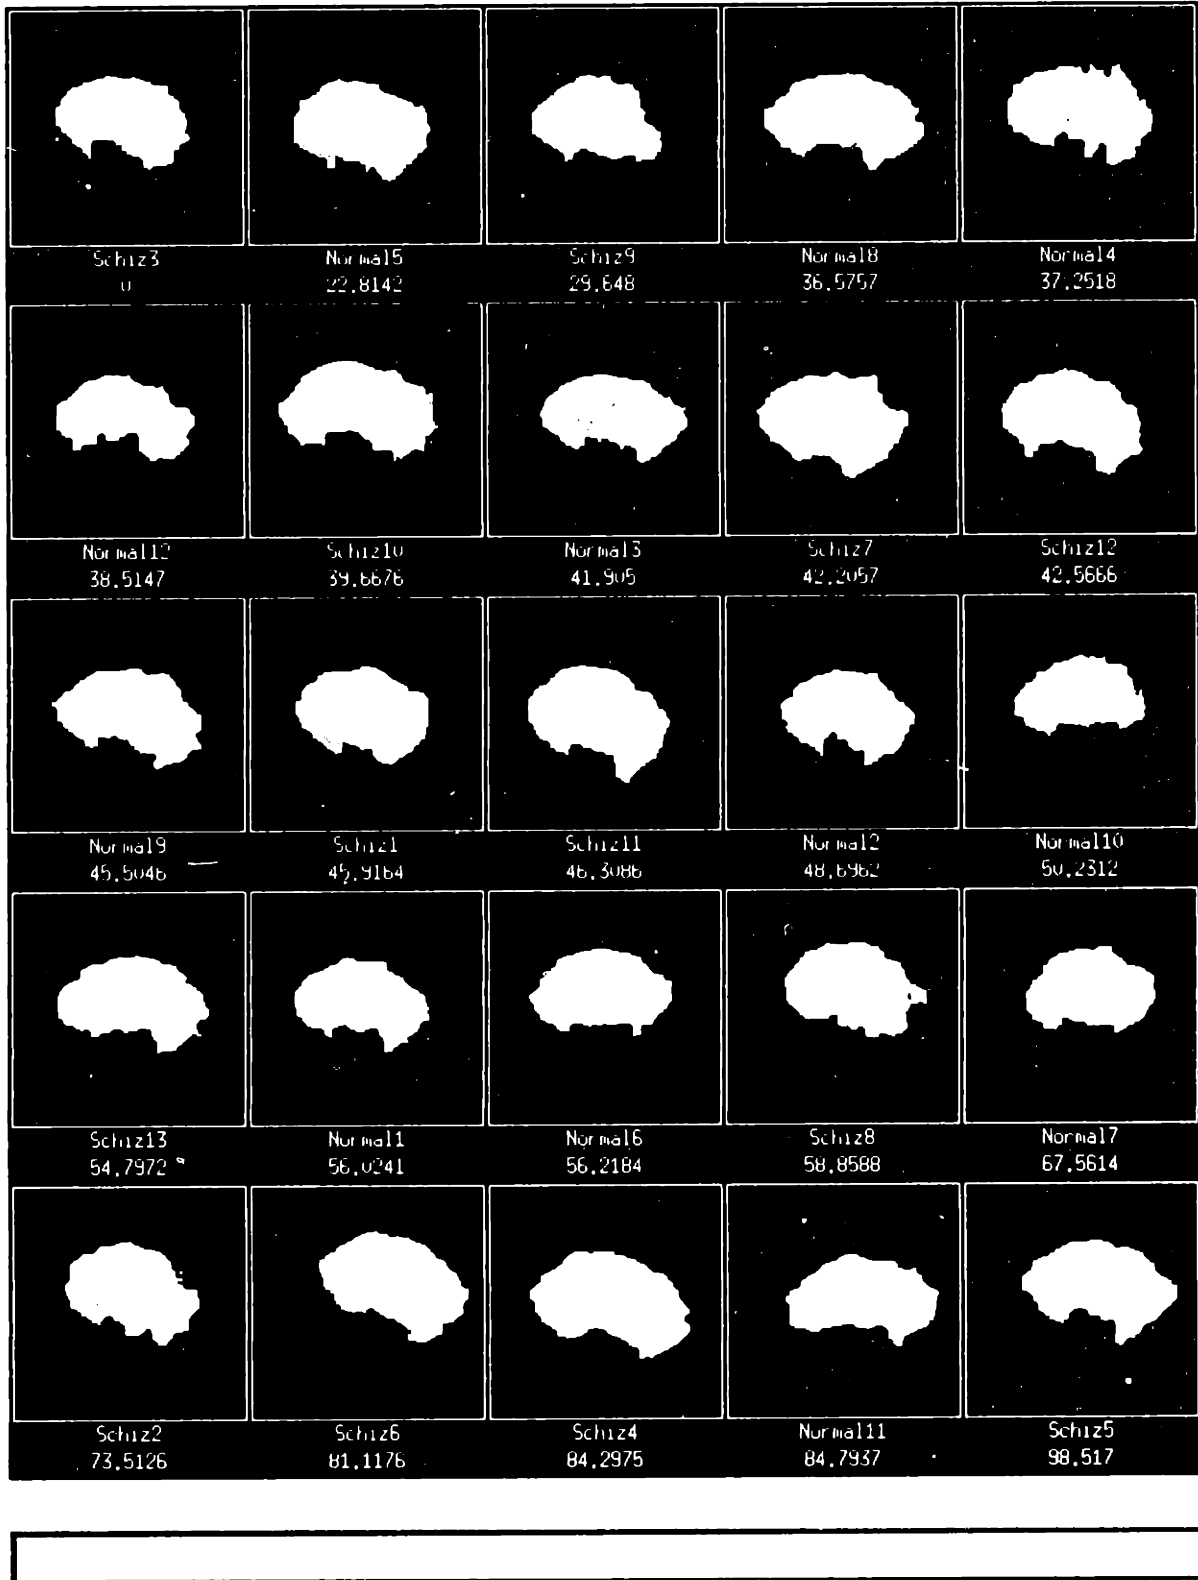


Figure 7.8: Ranking of the patients after the user has selected Schiz3. Photobook has incorrectly identified primarily healthy subjects as being the most similar to the schizophrenia patient selected.

Chapter 8

Discussion

8.1 Effect of Head Shape Normalization

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

Since the classification is performed in a six-dimensional space, it is difficult to visualize. We therefore plotted just the top two principal components for all 25 patients in the putamen data set, both without and with the head shape normalization. Figure 8-1a shows the projections onto the top two modes of the unwarped data, while Figure 8-1b shows the projections onto the top two modes computed after warping. Note that b_1 and b_2 of Figure 8-1b are the first two components of the projection vector \mathbf{b} of Equation 4.19, and that the eigenvectors onto which they are projected were shown earlier in Figure 5-3. Similarly, a_1 and a_2 of Figure 8-1a are projections onto the two highest-variance eigenvectors of the sample covariance matrix of the original, unwarped putamen.

Consistent with the decrease in classification rate, Figure 8-1a shows less class separability than Figure 8-1b. This is because the head shape normalization has removed some of the nonpathological shape variation between patients. This nonpathological deformation is represented by the projections onto the first two ICC deformation modes, shown in Figure 8-1c. 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.

Ideally, Figure 8-1b would show each disease state in a very compact cluster, with




Figure 8-1: Projections onto modes computed from the schizophrenia database. Each schizophrenic patient is denoted by a +, and each healthy volunteer is denoted by a o. (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.

complete separability between the clusters. There are several reasons why these conditions are not achieved. First, some nonpathological shape variation certainly remains. This shape variation might, for instance, be correlated with individual variation that is more local than what can be measured by overall head shape. Second, the amount of shape deformation caused by a disease is not necessarily going to be enough to result in a completely separable class. Finally, there will be some amount of variation in the shape effects of a particular disease. This will cause additional scatter within each of the disease classes.

We repeated the above procedure for the ventricle database. This time the classification rate did not drop, staying at 88%. Figure 8-2 shows the projections onto the top two modes for the unwarped ventricles, the warped ventricles, and the ICC. No improvement is seen in class separability between Figures 8-2a and 8-2b.

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

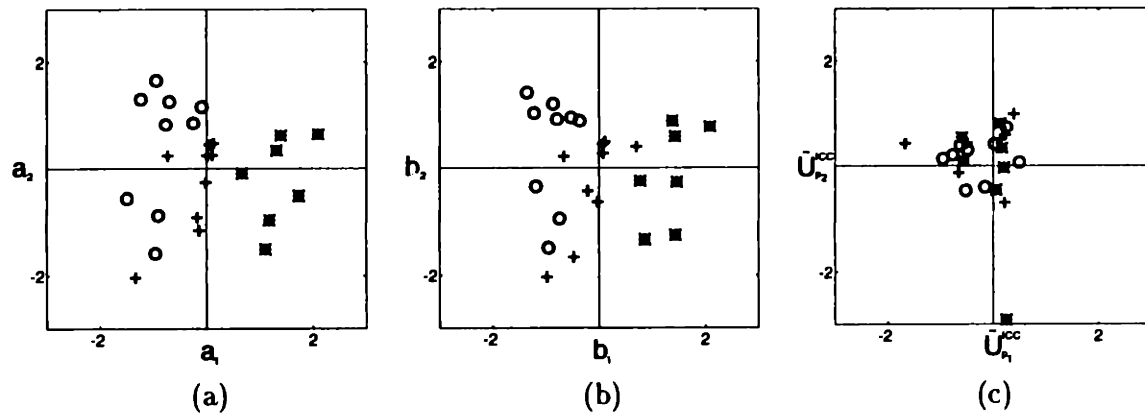


Figure 8-2: Projections onto modes computed from the ventricle database. Each Alzheimer's patient is denoted by a +, each normal-pressure hydrocephalus patient by an *, and each healthy volunteer by a o. (a) Ventricle principal components, computed from the original data. (b) Ventricle principal components, computed after first normalizing for head shape. (c) ICC physical deformation modes.

same amount [53]. Since normalizing an entire database to one standard head shape does not take into account gender-based shape variation, this gender-based variation may be interfering with the analysis of pathological shape differences.

To further investigate this problem, we constructed a mean square measure of ventricular similarity, by computing the distance between each voxel on the average ventricular surface and the nearest point on a patient's ventricular surface. The sum of these for a particular patient is a measure of the similarity of the patient to the average. If the ICC warping is truly removing some of the variation between ventricles, then this measure should decrease. Figure 8-3 shows the results for our 25 ventricular patients. Only a moderate decrease is seen with warping.

8.2 Model Accuracy

Another 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 basal ganglia, 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 [17, 20, 35, 63]. 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

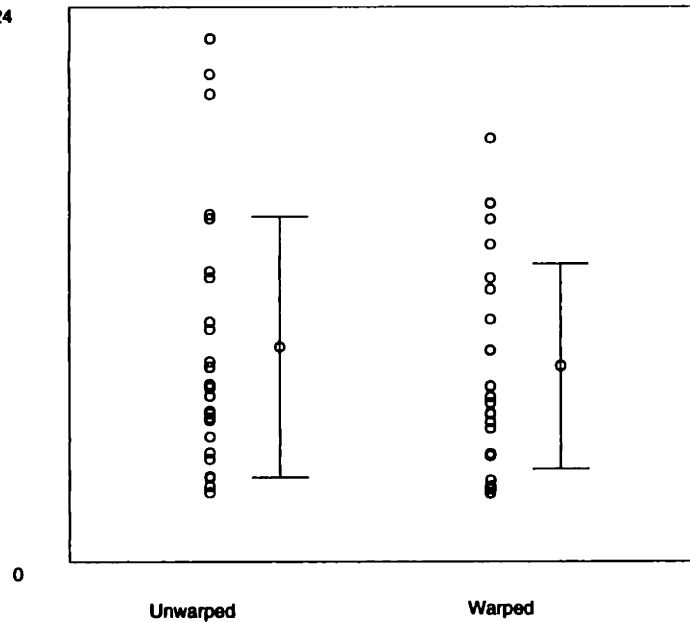


Figure 8-3: Mean square distance (in mm^2) between ventricles before and after head shape normalization.

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 [20, 63]. Because most of these models were created 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 [4, 18, 24], 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.

8.3 Statistical Analysis

In this thesis, we have reported experimental results (e.g. classification rates) without providing any sort of statistical analysis. Our results could therefore only *suggest* that the method works.

With larger data sets, statistical analysis becomes feasible. Obtaining accurate estimation errors of quantities derived from high-dimensional data is, however, very difficult. In this section we review some of the statistical measures that are available. Much of the development that follows is taken from [33].

8.3.1 Eigenvector estimation

The principal components of our data sets were first used to visually examine modes of deformation, and then used as inputs into a linear classifier. Because they are the eigenvectors of the sample covariance matrix, these sample principal components are random variables, with their own sampling distributions.

Unfortunately, the results that are available from multivariate statistical theory regarding the distribution of sample eigenvectors are complicated, asymptotic, and assume a multivariate normal distribution for the underlying population. For example, the variation of the sample principal component $\hat{\psi}_i$, measured in terms of the expectation of the mean-square error between it and the population principal component ψ_i , is given approximately by

$$E \left\{ \|\hat{\psi}_i - \psi_i\|^2 \right\} \approx \frac{1}{N} \sum_{\substack{j=1 \\ j \neq i}}^n \frac{\lambda_i \lambda_j}{(\lambda_i - \lambda_j)^2} \quad (8.1)$$

where λ_i is the eigenvalue corresponding to the eigenvector ψ_i , n is the dimensionality of the observation vector, and N is the number of samples [33]. Note that this expression requires knowledge of the population eigenvalues, which themselves can only be estimated.

The applicability of estimates such as Equation 8.1 is severely limited by the assumptions needed to derive the estimates [2]. For this reason, various ad hoc rules concerning which eigenvectors can be reliably estimated from a set of data have been proposed. In this thesis, we kept those eigenvectors corresponding to the largest $\frac{1}{4}$ eigenvalues.

8.3.2 Tests of normality

Equation 8.1, as well as some of the results to be presented in the next section, assume that the underlying distribution is multivariate normal. Methods of testing whether or not high-dimensional data comes from a normal distribution are not generally available, however. The best that can be done is to perform normality tests on the univariate marginal distributions, and also perhaps on a few linear combinations of the data. If the univariate tests are passed, then the hope is that the joint distribution of the original data is multivariate normal [47].

8.3.3 Classification accuracy

For each patient p , the projections \mathbf{b}_p onto the principal components were calculated. These were then fed into linear discriminant functions of the form

$$g_i(\mathbf{b}_p) = 2\mathbf{m}_i^T \Sigma^{-1} \mathbf{b}_p - \mathbf{m}_i^T \Sigma^{-1} \mathbf{m}_i \quad (8.2)$$

where \mathbf{m}_i is the mean of class i and Σ is the common covariance matrix. Because we did not know the underlying distributions, sample estimates were used for \mathbf{m}_i and Σ . The classifier is thus a function of these sample estimates, and its output is subject to sampling fluctuations. Expressions can be derived for the resulting bias and variance of the sample classification rate. However, even under simplifying assumptions, these estimates are very complicated [33].

A somewhat simpler, though not as useful, approach is to construct measures of class separability. For two classes, a common measure of separability is the Bhattacharyya distance, given by

$$\mu = \frac{1}{8}(\mathbf{m}_2 - \mathbf{m}_1)^T \left[\frac{\Sigma_1 + \Sigma_2}{2} \right]^{-1} (\mathbf{m}_2 - \mathbf{m}_1) + \frac{1}{2} \ln \frac{|\frac{\Sigma_1 + \Sigma_2}{2}|}{\sqrt{|\Sigma_1||\Sigma_2|}}. \quad (8.3)$$

Although the Bhattacharyya distance is derived assuming the underlying distributions are normal, it can also be used as a reasonable estimate for the separability of two non-normal distributions.¹

Like the discriminant functions of Equation 8.2, in practice the Bhattacharyya distance

¹If there are more than two classes, then other measures of class separability can be used. Most of these measures involve comparing between-class to within-class scatter matrices.

is computed using sample estimates of the unknown population parameters \mathbf{m} and Σ . Thus the Bhattacharyya distance becomes a random variable, with its own sampling distribution. Fukunaga [33] has derived an expression for the variance of this distribution. Defining μ_1 and μ_2 as the first and second terms of Equation 8.3, respectively, we have²

$$\text{Var}(\hat{\mu}_1) \approx \frac{1}{4N} \left[\sum_{i=1}^n \frac{m_i^2}{1 + \lambda_i} + \sum_{i=1}^n \sum_{j=1}^n \frac{m_i^2 m_j^2 (1 + \lambda_i \lambda_j)}{2(1 + \lambda_i)^2 (1 + \lambda_j)^2} \right] \quad (8.4)$$

and

$$\text{Var}(\hat{\mu}_2) \approx \frac{1}{2N} \sum_{i=1}^n \left[\left(\frac{1}{1 + \lambda_i} - \frac{1}{2} \right)^2 + \left(\frac{1}{1 + \lambda_i} - \frac{1}{2\lambda_i} \right)^2 \lambda_i^2 \right]. \quad (8.5)$$

8.3.4 Conclusion

In the ideal situation, the underlying probability distribution for each class would be known. The Bayes classifier could then be used, with the Bayes error giving the theoretically optimal classification rate. For most applications, however, the probability distributions are not known. Obtaining accurate error estimates for quantities such as the sample classification rate and the Bhattacharyya distance then becomes complicated, as illustrated in this section. In practice, the best that can be hoped for is to have available large sample sizes, enabling accurate empirical estimates of these quantities to be made.

8.4 Number of Principal Components

As mentioned earlier, estimating the sampling distribution of principal components (PCs) is difficult, and a reasonable rule of thumb is to keep those PCs corresponding to the largest $\frac{1}{4}$ eigenvalues. For each of our two 25-patient data sets, then, the top six PCs were kept.

To test the robustness of this parameter, we varied the number of PCs used in classification. Table VI shows the ventricle classification rate obtained using the Gaussian linear classifier, with the number of PCs used ranging from four to eight. As can be seen, the classification rate is not that stable. Table VII shows the effect of varying the number of PCs when the Gaussian linear classifier classifier is applied to the putamen database. With five or more PCs, the classification performance is very stable.

²To prevent these expressions from becoming more cumbersome, the covariance matrices were simultaneously diagonalized to \mathbf{I} and Λ , and a coordinate shift was applied to move the mean vector of the first class to $\mathbf{0}$ [33]. This can be done without loss of generality. Thus m_i refers to component i of the mean vector

TABLE VI

Gaussian Linear Classification Rates
 Ventricle Database
 Number of Principal Components Varied

# PCs	% Correct
4	64
5	60
6	88
7	80
8	84

TABLE VII

Gaussian Linear Classification Rates
 Putamen Database
 Number of Principal Components Varied

# PCs	% Correct
4	60
5	76
6	72
7	72
8	72

Tables VIII and IX show the effects of varying the number of PCs input into the nearest neighbor classifier. Just as for the Gaussian linear classifier, the number of PCs used is more robust for the putamen database than it is for the ventricle database.

As can be seen in both Tables VI and VIII, part of the instability in ventricle classification rate is due to the large improvement in performance that occurs when PC 6 is included. This suggests that PC 6 contains shape variation that is useful for class discrimination. Figure 8-4, which shows the patient projections onto PCs 5 and 6, verifies this. As the figure shows, PC 6 helps to distinguish Alzheimer's patients from both normal-pressure hydrocephalus patients as well as from normal subjects. Since classification errors are most often made either between Alzheimer's patients and normal-pressure hydrocephalus patients, or between Alzheimer's patients and normal subjects, it is not surprising that inclusion of PC 6 makes a large difference in classification performance.

of the second class, and λ_i refers to diagonal component i of the covariance matrix of the second class.

TABLE VIII

Nearest Neighbor Classification Rates
 Ventricle Database
 Number of Principal Components Varied

# PCs	% Correct
4	72
5	68
6	92
7	80
8	88

TABLE IX

Nearest Neighbor Classification Rates
 Putamen Database
 Number of Principal Components Varied

# PCs	% Correct
4	64
5	68
6	64
7	72
8	68

8.5 Preset Parameters

In this section, we briefly discuss each of the preset parameters used in the different steps of the implementation.

8.5.1 Rigid body alignment

The first step in our system is a rigid body alignment of two ICCs. To accomplish this, we used the technique developed by Ettinger *et al.* [31]. In their method, each ICC is represented by a list of its surface points. One of the 3-D point sets is then iteratively transformed in order to match the other, using a combination of Gaussian smoothing, stochastic sampling, and the minimization of an objective function based on a distance measure between the two point sets.

The preset parameters used in their technique are the widths of the smoothing Gaussians, the cutoff distance at which a point is considered to be an outlier, and the threshold




Figure 8-4: Projections onto ventricle principal components 5 and 6. Each Alzheimer's patient is denoted by a +, each normal-pressure hydrocephalus patient by an *, and each healthy volunteer by a o.

used when deciding to halt the iteration and declare the two point sets matched. The values of these parameters have been chosen based on both theoretical and empirical considerations. Because this work is still under development, however, the precise effects of changing these parameters are presently unclear.

8.5.2 Modal fitting

Once the two ICCs have undergone a rigid body alignment, the nonrigid components of deformation are calculated using the modal fitting method of Pentland and Sclaroff [57]. The preset parameters of their technique are the object's material properties. In their implementation, the material properties are chosen to correspond to those of a homogeneous, isotropic, extremely elastic material. Modeling the ICC with these values is certainly an over-simplification, and more accurate values ought to lead to better results. However, as discussed in Section 8.2, accurate values for the material parameters of actual ICCs are not available.

8.5.3 Principal component analysis

The final step in our system involves computing the principal components of the warped structures. The only parameter value that needed to be set was the number of principal components to use. The robustness of this parameter was investigated in Section 8.4.

8.6 Other Issues and Limitations

In addition to the topics discussed so far in this chapter, there are several secondary issues and limitations that are also worth exploring.

8.6.1 Correlations between ICC and interior structure shape

Because the shapes of the ICC and interior structures are not perfectly correlated, even a totally accurate physical model of the ICC will not remove all nonpathological shape variation when used to warp interior structures. While some correlation exists, the exact amount and types of it are not known. A study using linear measurements of the ventricles and cranium, for example, found that while the size of the lateral ventricles was correlated to cranium size, the size of the third ventricle was not [36].

Even if a correlation exists in general, it will not necessarily hold for all individuals. In the example of Section 4.1.4, for instance, shape correlation was present and we were able to successfully remove some of it. There will be cases, however, when there is little or no correlation present; warping the cranial contents according to ICC shape will therefore not help in normalizing the shape of interior structures.

Figures 8-5 and 8-6 provide an example where head shape normalization does not increase ventricular shape similarity. Even though the subject's ventricles (Figure 8-5d) are significantly larger than average (Figure 8-5b), their ICC (Figure 8-5c) is not that different from average (Figure 8-5a). Using the ICC shape difference to warp their cranial contents, therefore, does not produce ventricles (Figure 8-5e) that are any closer to the mean.

8.6.2 Correlations between ICC shape and disease state

Another important correlation is that between the shape of the ICC and the presence of pathology. An assumption throughout this thesis is that normalizing cranial contents by ICC shape will remove only nonpathological shape variation. If the disease under study has affected the shape of the ICC, however, then some pathological shape variation will also be taken out. Early childhood hydrocephalus, for example, often causes enlargement of the head, and thus the increase in ventricular size is accompanied by an increase in ICC size. In this case, normalizing the cranial contents using ICC shape will remove some pathological shape deformation.




Figure 8-5: An example where normalizing ventricular shape by cranium shape does not remove any nonpathological shape variation. (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 8-6: 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 Figure 8-5 for details.

The assumption of independence between pathology and head shape is probably valid for both Alzheimer’s disease and normal-pressure hydrocephalus, since their onset occurs in adulthood. Schizophrenia, on the other hand, may be present during early development, thereby increasing the chances of schizophrenia-induced effects in overall brain and ICC shape. However, both the improved feature space (Figure 8-1b versus 8-1a) and classification rates (72 versus 64%) obtained by first normalizing for ICC shape, as well as the class-independent distribution of schizophrenic and normal projections onto the ICC modes (Figure 8-1c), suggest that primarily nonpathological deformation was removed.

8.6.3 Disease duration and etiology

Two points concerning pathology should be noted. First, in medical studies researchers must be certain that their subjects actually have the diseases under investigation. Therefore typically only chronic cases are included in these studies. This is helpful when trying to understand structural changes caused by a disease, since deformation modes are best estimated from a data set of chronic patients. However, since a database of chronic cases will have better inter-class separability than a truly random sample, classification rates based on chronic cases (as in this thesis) will be greater than those generally attainable.

Second, many diseases have a wide etiologic diversity. Schizophrenia, for example, is thought to consist of a group of closely related, but presently unknown, “subdisorders”, all presenting similar clinical symptoms [1]. The same is true of normal-pressure hydrocephalus, although its group of subdisorders is probably more diverse than that of schizophrenia [59]. If the induced structural changes vary with the particular subdisorder present, then the modes calculated from a data set of NPH patients, for example, will be combinations of the modes of the individual NPH subdisorders. If the data set’s projections onto these modes show subclustering, then it may be possible to infer the existence of NPH subdisorders.

8.6.4 Correspondence

Computing the various deformation modes requires knowing the correspondence between points on surfaces. In this thesis we used a closest-match technique to approximate this correspondence. For calculating the correspondence between two ICCs or two putamen, this procedure is adequate. For structures as complicated as 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 in the estimated 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 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 [54].

8.6.5 Segmentation

Our approach assumes that the data has been pre-segmented into the structures of interest. The ventricular data set used in this thesis was semi-automatically segmented [53], while the putamen data set was segmented manually [69, 40].

The quality of the segmentation will obviously affect our results. Because of the low contrast and high noise levels typical of medical imagery, segmentation is a difficult and largely unsolved problem. While the above studies did test the reliability of their segmentation procedures, not all of the segmentation results were equally dependable. In the ventricular study, for example, the segmentation of the third ventricle was less accurate than that of the lateral and fourth ventricles.

8.7 Implementation

In this section, we discuss the speed at which the various components of our method run. All reported running times are on a Sun SPARC 2 with a standard configuration.

8.7.1 Rigid body alignment

The rigid body alignment stage [31] is the most time-consuming step in our current implementation, requiring on the order of 15 minutes per patient. Its relative slowness is caused primarily by its stochastic optimization procedure.

8.7.2 Modal fitting

The modal fitting [57] takes only a few minutes per patient. Note, though, that even with a database of only 25 patients, this step takes on the order of one hour to complete.

8.7.3 Principal component analysis

As described in Sections 1.1 and 4.2.2, a principal component analysis is performed on a set of 25 displacement vectors, with each vector being of size $3V \times 1$. Because we use all the surface points in our shape representation, $3V$ is on the order of several thousand. Thus it would appear that we must compute the eigenvectors of an extremely large matrix.

Fortunately, because the number of patients is much less than the dimensionality of the original feature space, the amount of required computation can be greatly reduced [82]. With only 25 patients, for instance, there will be only 24 non-zero eigenvalues, and thus 24 meaningful eigenvectors. It then becomes possible to transform the original problem into one of finding the eigenvalues and eigenvectors of a 25×25 matrix. This requires only a few seconds of CPU time.

8.7.4 Classification and database search

By the time we are ready to perform classification and database search, the dimensionality of the original space has been reduced to a feature space dimension of just six. Thus both classification and database search [56] take just a second or two.

8.7.5 Conclusion

The only parts of our implementation that are slow enough to warrant further optimization are the rigid body alignment and the modal fitting. Fortunately, these are the only two steps in which each patient is processed independently of all other patients. Thus parallelization of these steps would be relatively easy, since only an executive module that assigned each patient to a different processor would be needed. The core code of each algorithm would not have to change.

If these two steps were parallelized, then the entire procedure, from start to finish, would take on the order of minutes instead of hours. This would make it possible to do more experimentation with the preset parameters discussed in Section 8.5.

8.8 Suggestions for Future Research

8.8.1 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, as discussed in Section 8.2. However, it is possible to bypass this dependence on physical assumptions by using the connection between the eigenshapes calculated using principal component analysis and the physical deformation modes computed via modal analysis (see Section 3.3).

Using this connection, 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 the nodes of the model, Equations 3.16-3.20 can then be used to estimate the stiffness matrix.

Here we see the real power of the FEM and Equations 3.16-3.20. 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 3.16-3.20 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 thesis relies on experimental observation of material properties, and then builds the stiffness matrix in the usual FEM fashion. The required assumption is that the ICC behaves as a homogeneous, linear elastic material.

The method proposed in this section creates the stiffness matrix directly from experimental observations of ICCs, via Equation 3.20. Although assumptions about the ICC's material properties are not necessary in this case, certain assumptions regarding the load vector \mathbf{R} were needed in order to derive the connection between the covariance and stiffness

matrices.

Which of the two methods 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 3.20. Whether the first or second method is used, or even if the two sources of information are combined (using, for instance, the methods described in [26]), 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.

8.8.2 Normal shape correlations

In addition to its use in statistically calculating an ICC stiffness matrix, a large database of normal subjects could also be used to study correlations between ICC and interior structure shape. Once the amount and type of these correlations are known, more appropriate warping functions could be constructed.

A more ambitious task would be to search for structures other than the ICC with which to perform shape normalization. This thesis used the shape of the ICC for normalization because it is easy to segment, and because it is unlikely that ICC shape is affected by the diseases that we studied. There is no reason to assume, however, that the ICC is the best structure to use when normalizing every interior structure. It may be that other, more nearby structures are better suited for this purpose.

This idea can be investigated using a large database of normal subjects in which many of the interior structures have been segmented. For each particular structure whose shape is to be studied, various normalizations can be attempted, using different nearby structures or combinations of structures. To make this task computationally feasible, it may be necessary to use knowledge from developmental neuroanatomy in order to select those nearby structures which are most likely to produce effective shape normalizations.

8.8.3 Pathological shape correlations

In addition to studying normal shape correlations, pathological shape correlations should also be investigated. Doing so will require a large database of segmented scans for each disease of interest.

There are several aspects of pathology that could be studied. First, as mentioned earlier, little is known about possible correlations between ICC shape and disease state. By combining the pathological database with a normal control database, correlations between pathology and ICC shape can be found.

Second, studies similar to the work in this thesis could be performed. With the larger database, however, precise estimates can be made of quantities such as the improvement in classification rates obtainable when using principal component analysis instead of just volume.

Lastly, it would be possible to investigate the use of structures other than the ICC for performing shape normalization. Recall, however, that one of the benefits of using the ICC is that its shape is unaffected by most diseases. Since this is less likely to be true for nearby interior structures, using these structures for nonpathological shape normalization may also remove pathological shape variation. It would therefore be necessary to assess the amount of correlation between the disease under study and any structures that are candidates for use in the shape normalization procedure.

8.8.4 Segmentation versus shape description

Robust segmentation of medical imagery requires the use of prior anatomical knowledge. In addition to models of the average shape and location of various structures, “allowable” deformations in these models are also needed. Unfortunately, in order to obtain these shape descriptions, accurate segmentation of training data is required.

One possible way around this chicken-and-egg problem is to perform some sort of iteration between segmentation and shape description. Manually segmented training data can be used to calculate initial shape descriptions. These shape descriptions can then be used to segment other data [73, 76, 28]. The new segmentation results can then be used to update the shape descriptions, and so on. Both theoretical work that investigates the convergence properties of this approach, as well as experimental work that validates it, are needed.

Chapter 9

Summary

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

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

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

The limitations of our method involve the accuracy of physical models of brain stiffness, the ability to determine the correct correspondence between points on structures, and the degree to which nonpathological and possibly pathological morphology are correlated with

overall head shape. Overcoming these limitations will require better implementation, further investigation of the brain's material properties, and shape correlation studies involving large numbers of patients.

In summary, there are three 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. Finally, by incorporating our shape descriptions into an image database tool, we have indicated that it is possible to accurately and efficiently search through medical image databases using similarity metrics derived from the content of the images themselves.

Bibliography

- [1] R.D. Adams and M. Victor. *Principles of Neurology*. McGraw-Hill, Inc., 3rd edition, 1985.
- [2] T.W. Anderson. *An Introduction to Multivariate Statistical Analysis*. John Wiley & Sons, 2nd edition, 1984.
- [3] E.H. Aylward, J.D. Henderer, J.C. McArthur, P.D. Brettschneider, G.J. Harris, P.E. Barta, and G.D. Pearlson. Reduced basal ganglia volumes in HIV-1-associated dementia: Results from quantitative neuroimaging. *Neurology*, 43:2099–2104, October 1993.
- [4] R. Bajcsy and S. Kovacic. Multiresolution elastic matching. *Comput. Vision Graphics Image Process.*, 46:1–21, 1989.
- [5] A.R. Bakker. HIS, RIS, and PACS. *Computerized Medical Imaging and Graphics*, 15(3):157–160, 1991.
- [6] D. Bartelt, C.E. Jordan, E. Strecker, and A.E. James. Comparison of ventricular enlargement and radiopharmaceutical retention: A cisternographic-pneumoencephalographic comparison. *Radiology*, 116:111–115, July 1975.
- [7] K. Bathe. *Finite Element Procedures in Engineering Analysis*. Prentice-Hall, Inc., 1982.
- [8] F.L. Bookstein. Principal warps: Thin-plate splines and the decomposition of deformations. *IEEE Trans. Pattern Anal. Machine Intell.*, 11(6):567–585, June 1989.
- [9] F.L. Bookstein and W.D.K. Green. The biometrics of landmarks and edgels: A new geometry of prior knowledge for medical image understanding. In *Proc. AAAI Sympo-*

sium on Applications of Computer Vision in Medical Image Processing, pages 134–137, March 1994.

- [10] T.E. Boulton, S.D. Fenster, and T. O'Donnell. Physics in a fantasy world vs robust statistical estimation. In *Proc. NSF Workshop on 3D Object Recognition*, New York, NY, November 1994.
- [11] C.M Breant, R.K. Taira, and H.K. Huang. Integration of a voice processor machine in a PACS. *Computerized Medical Imaging and Graphics*, 17(1):13–19, 1993.
- [12] C. Brechbühler, G. Gerig, and O. Kübler. Parametrization of closed surfaces for 3-D shape description. *Computer Vision, Graphics and Image Processing: Image Understanding*, to appear, 1995.
- [13] G.D. Cascino, C.R. Jack, Jr., F.W. Sharbrough, P.J. Kelly, and W.R. Marsh. MRI assessments of hippocampal pathology in extratemporal lesional epilepsy. *Neurology*, 43:2380–2382, November 1993.
- [14] F. Cendes, F. Andermann, F. Dubeau, P. Gloor, A. Evans, M. Jones-Gotman, A. Olivier, E. Andermann, Y. Robitaille, I. Lopes-Cendes, T. Peters, and D. Melanson. Early childhood prolonged febrile convulsions, atrophy and sclerosis of mesial structures, and temporal lobe epilepsy: An MRI volumetric study. *Neurology*, 43:1083–1087, June 1993.
- [15] F. Cendes, F. Andermann, P. Gloor, A. Evans, M. Jones-Gotman, C. Watson, D. Melanson, A. Olivier, T. Peters, I. Lopes-Cendes, and G. Leroux. MRI volumetric measurement of amygdala and hippocampus in temporal lobe epilepsy. *Neurology*, 43:719–725, April 1993.
- [16] K.K. Chan and R.K. Taira. Systems integration for PACS. *Computerized Medical Imaging and Graphics*, 15(3):177–181, 1991.
- [17] M. Chan, C. Ward, D. Schneider, and S. Adams. Relative importance of skull deformation. In *Proc. Biomechanics Symposium*, 1981.
- [18] 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.

- [19] C.E. Christensen, R.D. Rabbitt, and M.I. Miller. 3D brain mapping using a deformable neuroanatomy. *Physics in Medicine and Biology*, 39:609–618, 1994.
- [20] C. Chu, M. Lin, and M. Lee. Finite element analysis of cerebral contusion. *Journal of Biomechanics*, 27(2):187–194, 1994.
- [21] H.E. Cline, W.E. Lorensen, R. Kikinis, and F.A. Jolesz. 3-D segmentation of MR images of the head using probability and connectivity. *Journal of Computer Assisted Tomography*, 14(6):1037–1045, 1990.
- [22] L. Cohen. On active contour models and balloons. *Computer Vision, Graphics and Image Processing: Image Understanding*, 53(2):211–218, 1991.
- [23] D.L. Collins, P. Neelin, T.M. Peters, and A.C. Evans. Automatic 3D intersubject registration of MR volumetric data in standardized Talairach space. *Journal of Computer Assisted Tomography*, 18(2):192–205, 1994.
- [24] D.L. Collins, T.M. Peters, W. Dai, and A.C. Evans. Model based segmentation of individual brain structures from MRI data. I. *SPIE: Visualization in Biomedical Computing*, pages 10–23, 1992.
- [25] D.L. Collins, T.M. Peters, and A.C. Evans. An automated 3D non-linear image deformation procedure for determination of gross morphometric variability in human brain. In *SPIE: Visualization in Biomedical Computing*, pages 180–190, 1994.
- [26] T.F. Cootes. Combining point distribution models with shape models based on finite element analysis. In *Proc. British Machine Vision Conference*, 1994.
- [27] 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.
- [28] T.F. Cootes, A. Hill, C.J. Taylor, and J. Haslam. Use of active shape models for locating structures in medical images. *Image and Vision Computing*, 12(6):355–365, July/August 1994.
- [29] M.J. de Leon, A.E. George, B. Reisberg, S.H. Ferris, A. Kluger, L.A. Stylopoulos, J.D. Miller, M.E. La Regina, C. Chen, and J. Cohen. Alzheimer’s disease: Longitudinal

- CT studies of ventricular change. *American Journal of Neuroradiology*, 10:371–376, March/April 1989.
- [30] I.A. Essa, S. Sclaroff, and A.P. Pentland. *Physically-based Modeling for Graphics and Vision*. Chapter in *Directions in Geometric Computing 1992*, Information Geometers, 1992. R. Martin, Ed.
- [31] G. Ettinger, E. Grimson, and T. Lozano-Perez. Automatic registration for multiple sclerosis change detection. In *CVPR Workshop on Biomedical Image Analysis*, pages 297–306, June 1994.
- [32] J. Feldmar and N. Ayache. Locally affine registration of free-form surfaces. In *Proc. CVPR*, pages 496–501, 1994.
- [33] K. Fukunaga. *Introduction to Statistical Pattern Recognition*. Academic Press, 2nd edition, 1990.
- [34] T. El Gammal, M.B. Allen, Jr., B.S. Brooks, and E.K. Mark. MR evaluation of hydrocephalus. *American Journal of Neuroradiology*, 8:591–597, July/August 1987.
- [35] 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.
- [36] C. Gyldensted. Measurements of the normal ventricular system and hemispheric sulci of 100 adults with computed tomography. *Neuroradiology*, 14:183–192, 1977.
- [37] 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.
- [38] B.K.T. Ho. Automatic acquisition interfaces for computed radiography, CT, MR, US, and laser scanner. *Computerized Medical Imaging and Graphics*, 15(3):135–145, 1991.
- [39] B.K.T. Ho, O. Ratib, and S.C. Horii. PACS workstation design. *Computerized Medical Imaging and Graphics*, 15(3):147–155, 1991.
- [40] H. Hokama, M.E. Shenton, P.G. Nestor, R. Kikinis, J.J. Levitt, C.G. Wible, B.F. O'Donnell, D. Metcalf, F.A. Jolesz, and R.W. McCarley. Caudate, putamen, and globus pallidus volume in schizophrenia: A quantitative MRI study. Submitted to *Psychiatric Research: Imaging*, 1994.

- [41] S.C. Horii, S.K. Mun, B. Levine, B. Lo, B.S. Garra, R.K. Zeman, M. Freedman, C. Left-ridge, D. Schellinger, J. Keyes, L.P. Elliott, R. Fielding, and H. Benson. PACS clinical experience at Georgetown University. *Computerized Medical Imaging and Graphics*, 15(3):183-190, 1991.
- [42] H.K. Huang, R.L. Arenson, S.-L. Lou, A.W.K. Wong, K.P. Andriole, T.M. Bazzill, and D. Avrin. Multimedia in the radiology environment: Current concept. *Computerized Medical Imaging and Graphics*, 18(1):1-10, 1994.
- [43] H.K. Huang, R.K. Taira, S.-L. Lou, A.W.K. Wong, C. Breant, B.K.T. Ho, K.-S. Chuang, B.K. Stewart, K. Andriole, R. Tecotzky, T. Bazzill, S.L. Eldredge, J. Tagawa, Z. Barbaric, M.I. Boechat, T. Hall, J. Bentson, and H. Kangarloo. Implementation of a large-scale picture archiving and communication system. *Computerized Medical Imaging and Graphics*, 17(1):1-11, 1993.
- [44] C.P. Hughes and M. Gado. Computed tomography and aging of the brain. *Radiology*. 139:391-396, May 1981.
- [45] G. Irie and K. Miyasaka. Clinical experience - 16 months of HU-PACS. *Computerized Medical Imaging and Graphics*, 15(3):191-195, 1991.
- [46] C.R. Jack, Jr., B. Mokri, E.R. Laws, Jr., O.W. Houser, H.L. Baker, Jr., and R.C. Petersen. MR findings in normal-pressure hydrocephalus: Significance and comparison with other forms of dementia. *Journal of Computer Assisted Tomography*, 11(6):923-931, November/December 1987.
- [47] R.A. Johnson and D.W. Wichern. *Applied Multivariate Statistical Analysis*. Prentice Hall, 2nd edition, 1988.
- [48] H. Kangarloo. PACS - clinical experience at UCLA. *Computerized Medical Imaging and Graphics*, 15(3):201-203, 1991.
- [49] H.L. Kundel, S.B. Seshadri, and R.L. Arenson. Clinical experience with PACS at the University of Pennsylvania. *Computerized Medical Imaging and Graphics*, 15(3):197-200, 1991.
- [50] H.U. Lemke, G. Faulkner, and M. Krauss. Development towards multimedia medical workstations. *Computerized Medical Imaging and Graphics*, 18(2):67-71, 1994.

- [51] Y. Ligier, O. Ratib, M. Logean, C. Girard, R. Perrier, and J.R. Scherrer. Object-oriented design of medical imaging software. *Computerized Medical Imaging and Graphics*, 18(2):125–135, 1994.
- [52] J. Martin, A.P. Pentland, and R. Kikinis. Shape analysis of brain structures using physical and experimental modes. In *Proc. CVPR*, pages 752–755, 1994.
- [53] M. Matsumae, A.V. Lorenzo, R. Kikinis, F.A. Jolesz, and P.McL. Black. Intra and extraventricular CSF volumes in normal and ventriculomegalic patients assessed by MR computerized image processing segmentation. In *American Association of Neurological Surgeons*, San Francisco, 1992.
- [54] 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.
- [55] A.M. Murro, Y.D. Park, D.W. King, B.B. Gallagher, J.R. Smith, F. Yaghmai, V. Toro, R.E. Figueroa, D.W. Loring, and W. Littleton. Seizure localization in temporal lobe epilepsy: A comparison of scalp-sphenoidal EEG and volumetric MRI. *Neurology*, 43:2531–2533, December 1993.
- [56] A.P. Pentland, R.W. Picard, and S. Sclaroff. Photobook: Tools for content-based manipulation of image databases. In *SPIE Conf. Storage and Retrieval of Image and Video Databases II, No. 2185*, February 1994.
- [57] 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.
- [58] B. Peterson, M.A. Riddle, D.J. Cohen, L.D. Katz, J.C. Smith, M.T. Hardin, and J.F. Leckman. Reduced basal ganglia volumes in Tourette’s syndrome using three-dimensional reconstruction techniques from magnetic resonance images. *Neurology*, 43:941–949, May 1993.
- [59] M. Matsumae P.L. Gleason, P.McL. Black. The neurobiology of normal pressure hydrocephalus. *Neurosurgery Clinics of North America*, 4(4):667–675, October 1993.

- [60] R.W. Preisendorfer. *Principal Component Analysis in Meteorology and Oceanography*. Elsevier Science Publishers B.V., 1988.
- [61] O. Ratib. From multimodality digital imaging to multimedia patient record. *Computerized Medical Imaging and Graphics*, 18(2):59–65, 1994.
- [62] O. Ratib, Y. Ligier, and J.R. Scherrer. Digital image management and communication in medicine. *Computerized Medical Imaging and Graphics*, 18(2):73–84, 1994.
- [63] 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.
- [64] T. Sandor, M. Albert, J. Stafford, and S. Harpley. Use of computerized CT analysis to discriminate between Alzheimer patients and normal control subjects. *American Journal of Neuroradiology*, 9:1181–1187, November/December 1988.
- [65] S. Sclaroff. *Modal Matching: A Method for Describing, Comparing, and Manipulating Digital Signals*. PhD thesis, Massachusetts Institute of Technology, 1995.
- [66] S. Sclaroff. Personal communication. 1995.
- [67] S. Sclaroff and A.P. Pentland. Modal matching for correspondence and recognition. *IEEE Trans. Pattern Anal. Machine Intell.* To appear.
- [68] S. Sclaroff and A.P. Pentland. On modal modeling for medical images: Underconstrained shape description and data compression. In *CVPR Workshop on Biomedical Image Analysis*, June 1994.
- [69] M.E. Shenton, R. Kikinis, F.A. Jolesz, S.D. Pollak, M. LeMay, C.G. Wible, H. Hokama, J. Martin, D. Metcalf, M. Coleman, and R.W. McCarley. Abnormalities of the left temporal lobe and thought disorder in schizophrenia: A quantitative magnetic resonance imaging study. *The New England Journal of Medicine*, 327(9):604–612, August 1992.
- [70] H.S. Singer, A.L. Reiss, J.E. Brown, E.H. Aylward, B. Shih, E. Chee, E.L. Harris, M.J. Reader, G.A. Chase, R.N. Bryan, and M.B. Denckla. Volumetric MRI changes in basal ganglia of children with Tourette’s syndrome. *Neurology*, 43:950–956, May 1993.

- [71] H.S. Soininen, K. Partanen, A. Pitkänen, P. Vainio, T. Hänninen, M. Hallikainen, K. Koivisto, and P.J. Riekkinen, Sr. Volumetric MRI analysis of the amygdala and the hippocampus in subjects with age-associated memory impairment: Correlation to visual and verbal memory. *Neurology*, 44:1660–1668, September 1994.
- [72] S.S. Spencer, G. McCarthy, and D.D. Spencer. Diagnosis of medial temporal lobe seizure onset: Relative specificity and sensitivity of quantitative MRI. *Neurology*, 43:2117–2124, October 1993.
- [73] L.H. Staib and J.S. Duncan. Boundary finding with parametrically deformable models. *IEEE Trans. Pattern Anal. Machine Intell.*, 14(11):1061–1075, November 1992.
- [74] B.K. Stewart, J.C. Honeyman, and S.J. Dwyer, III. Picture archiving and communication system (PACS) networking: Three implementation strategies. *Computerized Medical Imaging and Graphics*, 15(3):161–169, 1991.
- [75] G. Strang. *Introduction to Applied Mathematics*. Wellesley-Cambridge Press, 1986.
- [76] G. Székely, A. Kelemen, C. Brechbühler, and G. Gerig. Segmentation of 3D objects from MRI volume data using constrained elastic deformations of flexible fourier surface models. Submitted to *CVRMed*, 1995.
- [77] R. Szeliski. *Bayesian Modeling of Uncertainty in Low-level Vision*. Kluwer Academic Publishers, 1989.
- [78] R. Szeliski and S. Lavallée. Matching 3-D anatomical surfaces with non-rigid deformations using octree-splines. In *CVPR Workshop on Biomedical Image Analysis*, pages 144–153, June 1994.
- [79] R.K. Taira, B.K. Stewart, and U. Sinha. PACS database architecture and design. *Computerized Medical Imaging and Graphics*, 15(3):171–176, 1991.
- [80] D. Terzopoulos, A. Witkin, and M. Kass. Symmetry-seeking models and 3D object reconstruction. *Int. J. Comput. Vision*, 1:211–221, 1987.
- [81] C.W. Therrien. *Decision, Estimation, and Classification: An Introduction to Pattern Recognition and Related Topics*. John Wiley & Sons, 1989.

- [82] M. Turk and A.P. Pentland. Eigenfaces for recognition. *Journal of Cognitive Neuroscience*, 3(1):71–86, 1991.
- [83] C. Wikkelsö, H. Andersson, C. Blomstrand, M. Matousek, and P. Svendsen. Computed tomography of the brain in the diagnosis of and prognosis in normal pressure hydrocephalus. *Neuroradiology*, 31:160–165, 1989.
- [84] G. Wolberg. *Digital Image Warping*. IEEE Computer Society Press, 1990.
- [85] S.C. Zhu and A.L. Yuille. Forms: A flexible object recognition and modelling system. Technical Report 94-1, Harvard Robotics Laboratory, 1993.
- [86] S. Zink and C.C. Jaffe. Medical imaging databases. *Investigative Radiology*, 28(4):366–372, April 1993.