FEM EIGENMODES AS SHAPE FEATURES

M.H-M. Syn and R.W. Prager

CUED/F-INFENG/TR 211

May 1995

Cambridge University Engineering Department
Trumpington Street
Cambridge CB2 1PZ
England

E-mail: `mhs@eng.cam.ac.uk, rwp@eng.cam.ac.uk`

Abstract

The Finite Element Method (FEM) solution of the *wave equation* which governs the behaviour of elastic structures leads to a generalised eigenproblem. The eigenvectors of this eigenproblem are known as *eigenmodes* or *mode shapes*, which we present as an ideal set of shape features for use in model-based 3D ultrasound imaging.

We derive from first principles a framework for the modelling of volumetric linear elastic structures, using the Principal of Virtual Work. This allows us to construct mass and stiffness matrices which describe the shape and physical properties of a shape model.

We go on to examine the properties of the FEM eigenmodes of an elastic shape model, and the suitability of such a model in describing shape changes in biological structures. We show that there is an intimate connection between this model, and a growth model based on diffusion processes.

Contents

Chapter 1

Introduction

1.1 Growth models

1.1.1 Diffusive growth

Work by Kimia, Tannenbaum & Zucker (1990), Pizer & Burbeck (1994), Amit, Grenander & Piccioni (1991) and others suggest diffusion systems as the most general approach to modeling shape. Zucker and Pizer follow shape change through scale and image space in analogy to theories of human visual perception, whereas Grenander formulates a more general theory of image transformation using *deformable templates*.

The diffusion metaphor for shape *variation* is particularly appropriate in models of biological growth, since diffusion is the principal physical mechanism driving the growth of soft-tissue (eg. spleen), branching (eg. vascular) and bony structures from their embryonic to fully-developed forms.

Fujita (1986) describes a computational model of central nervous system growth using 60 segments with separate growth parameters using a Gompertzian grwoth function, and simulates the growth of monkey, rat and mouse brains.

1.1.2 Elastic growth

In the absence of accepted computational models derived from considerations of diffusion kinetics however, linear elastic models of *shape variation between fully developed structures* are generally used.

Bookstein (1991) presents the use of landmarks in analysing biological shape variation, and uses the thin-plate spline as a shape interpolant in 2D. He proposes the bending energy of the plate as a measure of shape difference, and linearises the modes of shape deformation using the *principal warps* of the bending energy matrix.

Terzopoulos and many others have applied the constraint of minimal bending energy to

regularise the interpolation of image features using active contours or *snakes*. An extension of this approach have been applied to the surface segmentation of 3D medical datasets (Cohen & Cohen 1993) using an "active balloon".

Our work in 3D ultrasound imaging (Syn, Gosling, Prager, Berman & Crowley 1994, Syn & Prager 1995*b*, Syn & Prager 1995*c*) requires the use of strong prior constraints of expected shape and variation in shape, in order to facilitate the segmentation of 3D datasets in a particularly noisy modality. Another technical report in this series (Syn, Gosling, Berman & Prager 1995) describes the use of modes of elastic growth (cf. Bookstein's principal warps above) in constraining the frequency response of an elastic shape model, so that prior knowledge about shape variation is incorporated into the interpolating "spline".

1.2   Prospectus

In Chapter 2 we develop a framework for the modeling of volumetric linear elastic structures. The Principle of Virtual Work leads us to the finite element approximation of the governing equilibrium equation.

In Chapter 3 we present an efficient method of solving the matrix eigenproblem associated with diagonalisation of the governing equation.

In Chapter 4 we present eigenmodes computed for a number of synthetic shape models. These indicate that eigenmodes are stable for reasonable peturbations in shape, and therefore constitute an ideal set of shape features for use in comparison of shape models.

Please note that the results presented in this report are also available via anonymous-ftp. See Section 4.1.

Chapter 2

Finite Element Models

The following derivations follow closely those of Bathe (1982), using the following variables

$$
\begin{array}{lcl}
(x, y, z) & = & \text{global coordinate system} \\
(r, s, t) & = & \text{local coordinate system within element} \\
\mathbf{u} & = & \text{nodal displacements in global coordinates} \\
\mathbf{u}^e & = & \text{nodal displacements in local coordinates}
\end{array}
$$

2.1   Equilibrium

Consider the equilibrium of a body with the following forces acting upon it

$$
\begin{array}{lcl}
\mathbf{r}^{body} & = & \text{volume forces (per unit volume)} \\
\mathbf{r}^{surface} & = & \text{surface forces (per unit surface area)} \\
\mathbf{r}^{point} & = & \text{point forces acting on the body}
\end{array}
$$

The displacements of the body from its unloaded state, due to the action of these forces, are described by the vector $\mathbf{u}$ defined above. The plane and shear components of strain[1] corresponding to $\mathbf{u}$ are

$$
\boldsymbol{\epsilon} = \left[\begin{array}{cccccc} \epsilon_{xx} & \epsilon_{yy} & \epsilon_{zz} & \gamma_{xy} & \gamma_{yz} & \gamma_{zx} \end{array} \right]^T \tag{2.1}
$$

and the stress[2] components are

$$
\boldsymbol{\tau} = \left[\begin{array}{cccccc} \tau_{xx} & \tau_{yy} & \tau_{zz} & \tau_{xy} & \tau_{yz} & \tau_{zx} \end{array} \right]^T \tag{2.2}
$$

---

[1]Strain is ratio of displacement to original length in the direction of displacement.
[2]Stress is the force applied per unit area.

2.2 Principle of Virtual Work

In order to analyse the response of the body, we establish differential equations describing the equilibrium of the body subject to appropriate boundary and compatibility conditions. One such formulation employs the *Principle of Virtual Work*, which states that a body in equilibrium requires the total internal virtual work done to equal the total external virtual work done

$$\int_V \overline{\epsilon}^T \tau . dV = \int_V \overline{\mathbf{u}}^T \mathbf{r}^{body} . dV + \int_A \overline{\mathbf{u}}^T \mathbf{r}^{surface} . dA + \sum_i \overline{\mathbf{u}}^T \mathbf{r}_i^{point} \qquad (2.3)$$

where the overline signifies a virtual strain or displacement.

For Equation 2.3 to hold, the displacements must meet the following conditions

- they should be compatible and continuous between elements

- they must satisfy the displacement boundary conditions

- they must satisfy constitutive relationships (i.e. stresses can be evaluated from strains)

A finite element discretisation of the body allows us to use a system of interpolation functions, defined in the *displacement interpolation matrix* $\mathbf{H}$ to relate the displacement within each element to the displacements of the nodes bounding each element, $\hat{\mathbf{u}}$

$$\mathbf{u}^e(r, s, t) = \mathbf{H}^e(r, s, t)\hat{\mathbf{u}}^e(x, y, z) \qquad (2.4)$$

Strains can similarly be interpolated using the *strain interpolation matrix* $\mathbf{B}$

$$\epsilon^e(r, s, t) = \mathbf{B}^e(r, s, t)\hat{\mathbf{u}}^e(x, y, z) \qquad (2.5)$$

where $\mathbf{B}(r, s, t) =$

$$\begin{bmatrix} \frac{\partial}{\partial x} & 0 & 0 \\ 0 & \frac{\partial}{\partial y} & 0 \\ 0 & 0 & \frac{\partial}{\partial z} \\ \frac{\partial}{\partial x} & \frac{\partial}{\partial y} & 0 \\ 0 & \frac{\partial}{\partial y} & \frac{\partial}{\partial z} \\ \frac{\partial}{\partial x} & 0 & \frac{\partial}{\partial z} \end{bmatrix} \mathbf{H}(r, s, t) \qquad (2.6)$$

2.3　Mass and stiffness matrices

The relation between the strain $\boldsymbol{\epsilon}_j$ and initial stress $\boldsymbol{\tau}_j^0$ at node j is

$$\boldsymbol{\tau}_j = \mathbf{E}_j \boldsymbol{\epsilon}_j + \boldsymbol{\tau}_j^0 \qquad (2.7)$$

where $\mathbf{E}$ is the *elasticity matrix*[3] which relates local stresses to resultant local strains. Over a volumetric element[4]

$$\mathbf{E} = \frac{E(1-\nu)}{(1+\nu)(1-2\nu)} \begin{bmatrix} 1 & \frac{\nu}{1-\nu} & \frac{\nu}{1-\nu} & 0 & 0 & 0 \\ \frac{\nu}{1-\nu} & 1 & \frac{\nu}{1-\nu} & 0 & 0 & 0 \\ \frac{\nu}{1-\nu} & \frac{\nu}{1-\nu} & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & \frac{1-2\nu}{2(1-\nu)} & 0 & 0 \\ 0 & 0 & 0 & 0 & \frac{1-2\nu}{2(1-\nu)} & 0 \\ 0 & 0 & 0 & 0 & 0 & \frac{1-2\nu}{2(1-\nu)} \end{bmatrix} \qquad (2.8)$$

Summing over all m elements of the body, we can rewrite Equation 2.3

$$\sum_{j=1}^{m} \int_{V_j} \bar{\boldsymbol{\epsilon}}_j^T \boldsymbol{\tau}_j .dV_j \;\; = \;\; \sum_{j=1}^{m} \int_{V_j} (\mathbf{H}_j \bar{\mathbf{u}}_j)^T \mathbf{r}_j^{body} .dV_j + \sum_{j=1}^{m} \int_{A_j} (\mathbf{H}_j \bar{\mathbf{u}}_j)^T \mathbf{r}_j^{surface} .dA_j \quad (2.9)$$
$$+ \sum_i \bar{\mathbf{u}}_i^T \mathbf{r}_i^{nodes}$$

2.3.1　Stiffness matrix

By substituting the stress-displacement interpolation function from Equation 2.6 into Equation 2.7, and factoring out the initial stress state $\boldsymbol{\tau}^0$ later in Equation 2.18, we obtain using $\mathbf{u}$ instead of $\mathbf{u}_j^e$

$$\boldsymbol{\tau}_j = \mathbf{E}_j \mathbf{B}_j \mathbf{u} \qquad (2.10)$$

by noting that for a given element only the displacements at the nodes of the element affect the displacement and strain distributions within the element.

We can then substitute for $\bar{\boldsymbol{\epsilon}}_j$ and $\boldsymbol{\tau}_j$ in the left-hand side of Equation 2.9 to give

---

[3]The scalar E in one-dimension is also known as Young's Modulus

[4]ν is the Poisson ratio

$$\sum_{j=1}^{m} \int_{V_j} \overline{\boldsymbol{\epsilon}}_j^T \boldsymbol{\tau}_j . dV_j \quad = \quad \sum_{j=1}^{m} \int_{V_j} \left(\mathbf{B}_j \overline{\mathbf{u}}_j \right)^T \left(\mathbf{E}_j \mathbf{B}_j \mathbf{u} \right) . dV_j$$

$$= \quad \left(\sum_{j=1}^{m} \int_{V_j} \overline{\mathbf{u}}_j^T \mathbf{B}_j^T \mathbf{E}_j \mathbf{B}_j . dV_j \right) \mathbf{u} \tag{2.11}$$

From the right-hand side of equation 2.9

$$\sum_{j=1}^{m} \int_{V_j} \overline{\boldsymbol{\epsilon}}_j \boldsymbol{\tau}_j . dV_j \quad = \quad \sum_{j=1}^{m} \int_{V_j} \overline{\mathbf{u}}_j^T \mathbf{H}_j^T \mathbf{r}_j^{body} . dV_j \tag{2.12}$$

$$+ \sum_{j=1}^{m} \int_{A_j} \overline{\mathbf{u}}_j^T \mathbf{H}_j^T \mathbf{r}_j^{surface} . dA_j + \sum_i \overline{\mathbf{u}}_i^T \mathbf{r}_i^{nodes}$$

Comparing the right-hand sides of Equations 2.11 and 2.12 with the steady-state equilibrium Equation 2.3 under nodal loading $\mathbf{f}$

$$\mathbf{Ku} = \mathbf{f} \tag{2.13}$$

we have an expression for the *stiffness matrix*

$$\mathbf{K} = \sum_{j=1}^{m} \int_{V_j} \mathbf{B}_j^T \mathbf{E}_j \mathbf{B}_j . dV_j \tag{2.14}$$

and from Equations 2.12 and 2.13 the loads $\mathbf{f}$, which are composed of body, surface and nodal point components minus the vector of initial loads $\mathbf{f}^0$ (due to initial stresses $\boldsymbol{\tau}_j^0$)

$$\mathbf{f} = \mathbf{f}^B + \mathbf{f}^S + \mathbf{f}^N - \mathbf{f}^0 \tag{2.15}$$

where[5]

$$\mathbf{f}^B \quad = \sum_{j=1}^{m} \int_{V_j} \mathbf{H}_j^T \mathbf{r}_j^{body} . dV_j \tag{2.16}$$

$$\mathbf{f}^S \quad = \sum_{j=1}^{m} \int_{A_j} \mathbf{H}_j^T \mathbf{r}_j^{surface} . dA_j \tag{2.17}$$

$$\mathbf{f}^0 \quad = \sum_{j=1}^{m} \int_{V_j} \mathbf{B}_j^T \boldsymbol{\tau}_j^0 . dV_j \tag{2.18}$$

$$\mathbf{f}^N \quad = \mathbf{r}^{nodes} \tag{2.19}$$

---

[5]note that there are different $\mathbf{H}$ for surface and volume interpolations

2.3.2 Mass matrix

If a load is applied rapidly, we must include inertial effects using d'Alembert's principle [6]

$$\mathbf{M\ddot{u}} + \mathbf{Ku} = \mathbf{f}(t) \tag{2.20}$$

Assuming that the element accelerations are in the same directions as the element displacements, the contribution of the inertia force to the load vector is

$$\mathbf{f}^B = \sum_{j=1}^{m} \int_{V_j} \mathbf{H}_j^T \left[\mathbf{f}_j^B - \rho_j \mathbf{H}_j \ddot{\mathbf{u}} \right] .dV_j \tag{2.21}$$

$$= \sum_{j=1}^{m} \int_{V_j} \mathbf{H}_j^T \mathbf{f}_j^B .dV_j - \left(\sum_{j=1}^{m} \int_{V_j} \mathbf{H}_j^T \rho_j \mathbf{H}_j .dV_j \right) \ddot{\mathbf{u}} \tag{2.22}$$

where ρ is the material density, giving

$$\mathbf{M} = \sum_{j=1}^{m} \int_{V_j} \mathbf{H}_j^T \rho_j \mathbf{H}_j .dV_j \tag{2.23}$$

2.4 Equivalence of virtual work derivation to stationarity of potential

Assuming a linear elastic continuum, the total potential of the body is

$$\Pi = \frac{1}{2} \int_V \boldsymbol{\epsilon}^T \mathbf{E} \boldsymbol{\epsilon} .dV - \int_V \mathbf{u}^T \mathbf{r}^{body} .dV - \int_A \mathbf{u}^T \mathbf{r}^{surface} .dA - \sum_i \mathbf{u}^T \mathbf{r}_i^{point} \tag{2.24}$$

by using the stress-strain relation in Equation 2.7

$$\boldsymbol{\tau} = \mathbf{E}\boldsymbol{\epsilon} \tag{2.25}$$

For Π to be stationary, $\delta\Pi = 0$. Then since the elasticity matrix $\mathbf{E}$ is symmetric, we have

$$\int_V \delta(\boldsymbol{\epsilon}^T \mathbf{E} \boldsymbol{\epsilon}) .dV = \int_V \delta(\mathbf{u}^T \mathbf{r}^{body}) .dV + \int_A \delta(\mathbf{u}^T \mathbf{r}^{surface}) .dA + \sum_i \delta(\mathbf{u}^T \mathbf{r}_i^{point}) \tag{2.26}$$

---

[6]This treats the ma term in $f = ma$ as an inertial force, so that $f - f_{inertia} = 0$

Providing displacements satisfy boundary conditions, so will the corresponding strains in Equation 2.26. Stationarity of Π in Equation 2.24 is then equivalent to the principle of virtual displacements in Equation 2.3, with $\delta\boldsymbol{\epsilon} \equiv \overline{\boldsymbol{\epsilon}}$ and $\delta\mathbf{u} \equiv \overline{\mathbf{u}}$. In other words, the body tends towards a stable energetic state where its net potential is minimised.

2.5 Interpolation functions for a linear tetrahedral element

For a four-node tetrahedral element, we use *unit coordinates* for which the three spatial parameters (r, s, t) vary from 0 to 1. This gives nodal coordinates of

$$(r, s, t) = \{(0, 0, 0), (1, 0, 0), (0, 1, 0), (0, 0, 1)\} \tag{2.27}$$

In three dimensions, the simplest complete polynomial has three constants, so that this linear function can be derived from the four nodal constraints. We can then define a polynomial interpolation function within the tetrahedral element

$$p^e(r, s, t) = g_1^e + g_2^e.r + g_3^e.s + g_4^e.t = \begin{bmatrix} 1 & r & s & t \end{bmatrix} \begin{bmatrix} g_1^e \\ g_2^e \\ g_3^e \\ g_4^e \end{bmatrix} \tag{2.28}$$

where the superscript denotes *local parameters* in *local coordinates* within an element.

Substituting nodal constraints

$$\begin{bmatrix} p^e(0, 0, 0) \\ p^e(1, 0, 0) \\ p^e(0, 1, 0) \\ p^e(0, 0, 1) \end{bmatrix} = \begin{bmatrix} p_{node1}^e \\ p_{node2}^e \\ p_{node3}^e \\ p_{node4}^e \end{bmatrix} = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 1 & 1 & 0 & 0 \\ 1 & 0 & 1 & 0 \\ 1 & 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} g_1^e \\ g_2^e \\ g_3^e \\ g_4^e \end{bmatrix} \tag{2.29}$$

or

$$\mathbf{p}_{nodes}^e = \mathbf{G}\mathbf{g}^e \tag{2.30}$$

where $\mathbf{G}$ is the *geometry matrix*.

We then substitute back into Equation 2.28

$$p^e(r, s, t) = \begin{bmatrix} 1 & r & s & t \end{bmatrix} \mathbf{G}^{-1}\mathbf{p}_{nodes}^e = \mathbf{H}(r, s, t)\mathbf{u}_{nodes}^e \tag{2.31}$$

where $\mathbf{H}$ is the *interpolation matrix* containing the interpolation functions

$$\mathbf{H}(r,s,t) = \left[\begin{array}{cccc} H_1(r,s,t) & H_2(r,s,t) & H_3(r,s,t) & H_4(r,s,t) \end{array}\right] = \left[\begin{array}{cccc} 1 & -r & -s & -t \\ & r & & \\ & & s & \\ & & & t \end{array}\right]^T$$

$$(2.32)$$

2.5.1   Continuity

Along one edge of the tetrahedral element, we hold r and s constant. Thus the interpolation function varies linearly with t along this boundary. Since this linear function is uniquely determined by the nodal constraints at both ends of the edge, we have continuity along any interface between linear elements. This is not true if we use the global coordinates (x, y, z) to define the interpolation.

2.5.2   Isoparametric elements

Equation 2.3 is presented in global coordinates, and involve derivatives with respect to global coordinates. Isoparametric elements use the same (local coordinate) interpolation functions used to define any local parameter, to map to global coordinates.

For example, we can define the global coordinates of an element (x, y, z) in terms of its local coordinates (r, s, t)

$$\left[\begin{array}{c} x(r,s,t) \\ y(r,s,t) \\ z(r,s,t) \end{array}\right]^T = \mathbf{H}(r,s,t) \left[\begin{array}{ccc} \mathbf{x}_{nodes}^e & \mathbf{y}_{nodes}^e & \mathbf{z}_{nodes}^e \end{array}\right]$$

$$(2.33)$$

by interpolating between nodal values given as a function of global coordinates. This allows us to retain the interfacial continuity mentioned in the previous section.

2.5.3   Jacobian

Using the relationship just established between local and global coordinates, we can now find expressions for the Jacobian matrix $\mathbf{J}$, the determinant of which (the Jacobian) is used to effect a change of variables during integration in Section **??**.

For a tetrahedral element with three independent coordinate variables (r, s, t)

$$\begin{bmatrix} \frac{\partial}{\partial r} \\ \frac{\partial}{\partial s} \\ \frac{\partial}{\partial t} \end{bmatrix} = \mathbf{J} \begin{bmatrix} \frac{\partial}{\partial x} \\ \frac{\partial}{\partial y} \\ \frac{\partial}{\partial z} \end{bmatrix} \tag{2.34}$$

$$\mathbf{J} = \begin{bmatrix} \frac{\partial x}{\partial r} & \frac{\partial y}{\partial r} & \frac{\partial z}{\partial r} \\ \frac{\partial x}{\partial s} & \frac{\partial y}{\partial s} & \frac{\partial z}{\partial s} \\ \frac{\partial x}{\partial t} & \frac{\partial y}{\partial t} & \frac{\partial z}{\partial t} \end{bmatrix} \tag{2.35}$$

The change of variables required to integrate globally over quantities defined in local coordinates can be achieved with

$$d(V) = dx.dy.dz = |\mathbf{J}|.dr.ds.dt \tag{2.36}$$

The Jacobian is usually computed numerically for elements other than the 4-node tetrahedron, which has a constant Jacobian $|\mathbf{J}| = 6 \times V$

Chapter 3

FEM eigenmodes

It is generally difficult to find the eigensolutions of Equation 2.3, even if an analytic description is available of its shape and material properties. One of the few examples is the closed-form eigensolution for surfaces which can be mapped onto a sphere, giving rise to *spherical harmonics* (Poli, Coppini & Valli 1994, Ballard & Brown 1982). Another is the solution of the wave equation on a vibrating square or circular membrane (Kreysig 1988, Ch. 11).

The purpose of the Finite Element Method (FEM) is to allow the solution of partial differential equations (PDEs) such as the one governing the dynamic equilibrium of a linear elastic structure with no damping, by discretisation over the continuous domain of the structural mass and stiffness operators to give

$$\mathbf{M\ddot{u}} + \mathbf{Ku} = \mathbf{f}(t) \tag{3.1}$$

where $\mathbf{M}$ and $\mathbf{K}$ are the mass and stiffness *matrices*, and $\mathbf{f}(t)$ is a time-dependant applied force.

The equation of undamped free (i.e. no external forces) vibration

$$\mathbf{M\ddot{u}} + \mathbf{Ku} = \mathbf{0} \tag{3.2}$$

has simple harmonic solutions, $\mathbf{u}_i(t) = \boldsymbol{\phi}_i e^{j\omega_i(t-t_0)}$, derived from the generalised eigenproblem

$$\mathbf{K\Phi} = \mathbf{\Omega}^2 \mathbf{M\Phi} \tag{3.3}$$

This can be solved using numerical techniques such as the Lanczos method to give $\mathbf{\Phi}$, the matrix of eigenvectors $\boldsymbol{\phi}_i$, and $\mathbf{\Omega}$, the diagonal matrix of eigenvalues ω_i.

3.1 FEM eigenmodes as shape features

FEM eigenmodes were first used as 2D shape features by Sclaroff & Pentland (1995), and this derived from their work in using elastic eigenmodes to constrain the recovery of shape (Pentland 1990, Pentland & Horowitz 1991, Pentland & Sclaroff 1991, Pentland & Williams 1989), using 2D boundary models and 3D surface models.

Brady & Scott (1988) also suggested eigenfunctions of other PDE formulations, over 2D solid shapes, as more generalised symmetric axes than the Symmetric-Axis Transform (SAT). SAT algorithms (e.g. grassfire) implicitly compute the eigenmodes of Laplace's equation with the appropriate boundary conditions. See Section 3.8 for more details.

In the medical imaging literature Nastar & Ayache (1993) use FEM eigenmodes in motion analysis. Staib & Duncan (1992) also use a more limited set of Fourier harmonics for motion tracking, and more recently Coppini, Poli & Valli (1995) have applied this to left-ventricular shape recovery in echocardiographic images.

Some early work in combining simple physical models of biological structures with statistical models appears in (Martin, Pentland & Kikinis 1994).

The application of eigenfunctions of a PDE to constrained "shape" recovery is well-established and commonly used in both the structural engineering (Bathe 1982, Ch. 9.3) and morphometric literature (Bookstein 1991, Ch. 7.5). The novelty in recent applications cited above, is the application of powerful modern workstations to the symmetric eigenproblem (Parlett 1980).

Our presentation in the following sections examines why elastic FEM eigenmodes have excellent scale-ordered and globally symmetric properties, when used as 3D shape features.

3.2 Mode shapes

These eigenvectors, ϕ_i, derived from structural FEM are usually termed *eigenmodes* or *mode shapes*. They describe mutually orthogonal modes of free vibration, each with natural frequency of vibration ω_i.

Figure 3.1 shows some non-rigid mode shapes of an ellipsoidal surface discretised into triangular elements, ranked in increasing order of natural frequency of vibration. It can be seen that the spatial frequency of the mode shapes increase with natural frequency.

These "harmonics" behave very much as the Fourier harmonics described for spheres and membranes at the start of Chapter 3. Indeed for a simple one dimensional example of a piece of linear elastic string with uniform material properties, the mode shapes are exactly the Fourier harmonics. This is illustrated in Figure 3.2. Note that while the *spatial* frequency of the harmonic mode shapes depends on the geometry of the string, the *natural* frequency of vibration depends only on the mass and stiffness properties of the string.

14

(a) 7th mode   (b) 9th mode   (c) 11th mode   (d) 13th mode   (e) 15th mode

Figure 3.1: FEM mode shapes of free vibration for an ellipsoidal surface in increasing order of natural frequency.


Figure 3.2: First and second mode shapes of a linear elastic string.

3.3 Rayleigh quotient characterisation

The minimax (Parlett 1980, Ch. 10.2) or Rayleigh quotient characterisation (Bathe 1982, Ch. 2.8) of the generalised eigenproblem leads to an equivalent definition for eigenmodes and natural frequencies

$$\omega_i = \max_{\mathbf{S},\dim(\mathbf{S})=n-i+1} \left\{ \min_{\mathbf{u} \in \mathbf{S}, \mathbf{u} \neq \mathbf{0}} \left\{ \frac{\mathbf{u}^T \mathbf{K} \mathbf{u}}{\mathbf{u}^T \mathbf{M} \mathbf{u}} \right\} \right\} \tag{3.4}$$

where the expression being locally minimised is the Rayleigh quotient $\rho(\mathbf{u})$, defined for symmetric mass and stiffness matrices, and the eigenvectors $\boldsymbol{\phi}_i$ are given by the vectors $\mathbf{u}$ corresponding to each ω_i.

The physical interpretation of $\rho(\mathbf{u})$ for structural FEM, is as the ratio of elastic potential to kinetic energy of the structure in free vibration. Successive and non-decreasing (Bathe 1982, Ch. 2.8) maxima in local minima of this energetic ratio therefore correspond to the eigenmodes of the structure in free vibration. All natural frequencies ω_i will be non-negative since $\mathbf{M} > \mathbf{0}$ and $\mathbf{K} \geq \mathbf{0}$.

3.4 Repeated eigenmodes

It follows from this that structures in three dimensions with either 2 or 3 planes of mirror symmetry, will have either 2 or 3 local minima of ρ with exactly the same energy, with respect to each of the planes of symmetry. These eigenmodes are known as *repeated eigenmodes* of multiplicity 2 or 3, and are of exactly the same shape about their respective planes of symmetry, as well as having exactly the same eigenvalues.

They are at the same time orthogonal to each other, and uniquely determined with respect to each other, although they are not individually uniquely determined with respect to other eigenmodes. Repeated eigenmodes therefore span an *eigenspace* uniquely defined with respect to other eigenmodes.

Symmetries of multiplicity 3 (eg. sphere, unit cube) are highly unlikely, so we consider only two-dimensional eigenspaces which define a unique plane (rather than vector) at each of the structure's nodes. We can use the normal to this plane as the unique vector given by a pair of repeated eigenmodes, so that both repeated or simple eigenmodes can in future be treated in the same way when used as shape features.

The block Lanczos algorithm which we use is designed to detect such repeated eigenmodes when solving the generalised eigenproblem (see Section 4.2).

3.5 You cannot hear the shape of a drum

Kac (1966) asked the question "Can one hear the shape of a drum", and proved that one

can determine the area and circumference of a membrane from its eigenvalue spectrum. For a polygonal membrane, he also proved the connectivity of the membrane can also be determined. Gordon, Webb & Wolpert (1992) have shown that two different domains can have the same eigenvalue spectrum however, hence one cannot determine the shape of the drum from its eigenvalue spectrum.

Even when we have information about eigenvectors as well as eigenvalues, this is still insufficient to uniquely determine the structure's shape. Information in the mass and stiffness matrices is reduced in dimensionality when the generalised eigenproblem is solved, so this transformation is necessarily irreversible.

Reduction in dimensionality is a desirable property for shape features, since we are trying to distill shape information into a handful of meaningful and representative quantities.

Natural frequencies have properties similar to moment invariants (Hu 1962) when used as shape features. They are global features and hence not fault tolerant, and their representation of actual geometric features is very weak. They are invariant to linear transformation and scaling however, and shape features derived from moments can be useful when used for rough discrimination between shapes. Syn & Prager (1995a) presents an application for moment features.

3.6   The diffusion equation in elastic structures

There is an intimate connection between the *wave equation* governing the behaviour (or displacement, u) of elastic structures

$$\frac{\partial^2 u}{\partial t^2} = c^2 \nabla^2 u \tag{3.5}$$

and the *diffusion equation* governing the behaviour of diffusing structures

$$\frac{\partial u}{\partial t} = c^2 \nabla^2 u \tag{3.6}$$

For the wave equation, c^2 is equal to the stiffness operator divided by the mass operator. For the diffusion equation applied to heat flow for example, the stiffness operator in c^2 is replaced with the ratio of thermal conductivity to specific heat (Kreysig 1988, Ch. 11)[1].

For a separation of variables solution

$$u = f(x, y, z)g(t) \tag{3.7}$$

with the spatial and temporal derivatives noted with primes and dots respectively

---

[1]We follow Kreysig's notation in using c^2 to show that the operator is always positive semi-definite

$$\frac{\partial^2 u}{\partial t^2} = f\ddot{g} \tag{3.8}$$

$$\nabla^2 u = f''g \tag{3.9}$$

and substituted into the wave equation in Equation 3.5 we have

$$\frac{\ddot{g}}{c^2 g} = \frac{f''}{f} = k(\text{constant}) \tag{3.10}$$

The spatial component of the above equality is the same for both wave and diffusion equations

$$\nabla^2 u + ku = 0 \tag{3.11}$$

which leads to the spatial eigenproblem first derived in Finite Element matrix form in Equation 3.3. See (Syn & Prager 1995a) for details on how the FEM formulation maps to the continuous formulation.

In other words, the eigenmodes of Equation 3.3 can also be interpreted as the principal modes of growth of a structure, the shape and properties of which are described only in c^2, which is changing shape by process of diffusion (e.g. accretion of mass in bony structures, or subdivision of cells in soft tissue structures).

3.7   Scale ordering of eigenmodes

The analytic solution for mode shapes of a spherical membrane is not only harmonic in time, being of form $e^{j\omega t}$, but also harmonic in spatial shape. Spherical harmonics (Ballard & Brown 1982, pp. 270–274) parameterised in (θ, ϕ), are modulated by $e^{jn\theta}$ and $e^{jn\phi}$.

This is also true of other structures such as those in Figure 3.1, and this observation accords with the Rayleigh quotient characterisation of mode shapes. The smaller the spatial *wavelength* of each mode shape, the higher the elastic potential energy, since ω_i^2 is effectively the modal elasticity constant (Bathe 1982, Ch. 4.2.6, p. 173). The kinetic energy tends to decrease with spatial wavelength.

This natural scale ordering is another desirable property of eigenmode shape features, since higher (natural) frequency modes are unimportant in determining larger-scale response of the freely vibrating structure. We can hence discard eigenmodes in principled order, particularly since the Nyquist cutoff (spatial) frequency of the mode shapes is determined by the sampling resolution of the model's finite element mesh.

Modal Elasticity is square of the natural frequency

Elastic energy is higher and kinetic energy is lower for mode shape of smaller spatial scale

Figure 3.3: Scale ordering of modes of string.

3.8 Medial axes

Brady & Scott (1988) present the Symmetric-Axis Transform (SAT), commonly also known as the Medial-Axis Transform, as a highly parallelisable means of shape description. The use of medial axes in medical imaging has been extensively explored[2], and the most recent research involves tracking these axes through diffusion scale space eg. Pizer's method of "cores" (Pizer & Burbeck 1994) . This extension attempts to build a description which draws support from neighbourhoods of varying extent[3] in order to localise features on the shape's boundaries.

The SAT is designed to emphasise symmetry, but the most important criticism of this representation, which scale-space extensions partially alleviate, is that it cannot encompass both mirror and rotational symmetries. Brady in fact speculates as to the greater flexibility of the normal modes of a membrane shape model in this respect, and suggests that these modes (which can be derived in more than one way for the same shape depending upon the PDE) are in fact more generalised descriptors of shape and symmetry than medial axes.

Figure 3.4 from (Brady & Scott 1988) compares the eigenmodes and medial axes of a square plate

The medial axes can be computed using a "grassfire" algorithm. A simple analogy is that of placing a uniform distribution of electric charge along the boundary, so that the ridges of maximum (or minimum) potential describe the medial axes. This is a solution of Laplace's equation, $\nabla^2 u = 0$, with the given boundary conditions. The medial axes are also the ridges of maximum concentration in a homogeneous mass within the shape boundaries

---

[2]It has been found to be particularly appropriate as a means of describing branching (eg. vascular or arterial) structures!().

[3]The standard medial axis description typically has global support

<div align="center">(a) Eigenmodes   (b)
Medial Axes</div>

Figure 3.4: Comparing eigenmodes and medial axes of a square plate

undergoing diffusion (Pizer, Oliver & Bloomberg 1987).

We have seen in Section 3.6 that there is an intimate connection between the the wave equation and the heat equation. This Section also shows that the eigenmodes of the wave equation are superior shape features to the medial axes which are derived from a particular formulation of the diffusion equation.

3.9   Symmetry of eigenmodes

The first six eigenmodes of the freely vibrating structure are those of rigid-body motion. These rigid-body eigenmodes account for any translational or rotational motion of the body about its centre of mass. There are no supports or physical boundary conditions, so after an initial impulse acts upon the body, there are no external forces acting upon the freely vibrating body. Therefore none of the other (nonrigid) mode shapes can have a net translation of the centre of mass, or a net rotational moment.

This does not necessitate any symmetry in the (nonrigid) degrees of freedom that the structure has in motion, but for those which are eigenmodes and therefore local minima in the Rayleigh quotient, symmetric deformations allow for a minimal penalty in elastic energy while obeying the constraints set out above

Additionally most biological structures are simply connected (ie. "bloblike") and map straightforwardly to an ellipsoid. The eigenmodes of the ellipsoid retain symmetry when mapped.

Each mode shape therefore tends to emphasise symmetries in the structure, at the characteristic spatial scale of that mode. Eigenmode features require global support, and cannot therefore easily cope with occluded part structures. This is irrelevant in model-based imaging, since the shape model is always constructed in full.

3.10 Summary

We have presented FEM eigenmodes as ideal shape features for models used in 3D ultra-sound imaging. They have a natural ordering in spatial scale (cf. Medial Axes), and draw on global support to emphasise structural symmetries at each characteristic scale.

These shape features can be derived for any structure, using any of the available elements in structural FEM. There are no complications element mesh generation because we are not interested in a stress-strain analysis, and there is therefore no need to refine meshes at points of high stress. Uniform meshing is in fact recommended to minimise the number of elements needed for the desired accuracy in modelling (Hitchens 1992, Ch 4.7)[4].

It is standard practice in numerical FEM to allow a minimum of eight elements to span the smallest feature of interest. We are therefore justified in using only linear elements, which have the computational advantage of having a constant Jacobian matrix, since a piecewise linear approximation to, say, a sine wave discretised to eight elements per half-wavelength is more than sufficient for our needs. We need then only use mode shapes of low spatial frequency out of this rich feature set, since the high frequency ones are inaccurate beyond the Nyquist sampling frequency.

It is also an ideal representation for determining the equilibrium configuration of the structure given a set of applied forces, and has been used by Pentland & Horowitz (1991) for overconstrained shape fitting, where he uses only as many degrees of freedom as there are applied point forces.

The rigid-body eigenmodes, which are orthogonal to the other non-rigid mode shapes, can therefore be used to separate the translation and rotation response of the body from nonrigid shape deformation. We have applied these eigenmodes in the efficient computation of the response of a linear elastic shape model when used as an interpolating spline in segmenting 3D ultrasound images (Syn et al. 1995).

---

[4]Mode shapes cause large displacement amplitudes all over the body, so uniform meshing is most sensible.

Chapter 4

Results

4.1 Anonymous FTP

The datasets presented here are available in animated form for the following platforms

```
Linux
SunOS/Solaris
IRIX
HP-UX
```

from our anonymous-ftp server

`svr-ftp.eng.cam.ac.uk:/pub/data/geomdemo.*.tar.gz`

4.2 Computing eigenmodes

We have implemented a block Lanczos solver to compute eigenmodes and natural frequencies. The Lanczos method was first proposed as a method for tridiagonlisation of matrices, but this later developed into a method of extracting extreme eigenpairs in the symmetric eigenvalue problem, because tridiagonal systems can be solved efficiently (Golub & van Loan 1989, Ch. 9) (Parlett 1980, Ch. 13).

The *block* Lanczos algorithm allows us to detect repeated eigenpairs (see Section 3.4) by generalising the Lanczos algorithm to work with a subspace instead of a vector.

4.2.1 Sparse indexing

The sparsity structure of the mass and stiffness matrices of an ellipsoidal volume are shown in Figure 4.2.1(a)

(a)

Sparsity of mass and stiffness matrices

(b)

Sparsity of Cholesky factor of mass matrix

Figure 4.1: Sparsity plots for an ellipsoidal volume (nz is the number of non-zero elements)

Dense storage of this structure is very expensive, both in memory usage as well as in computation. Sparse indexing technology (Saad 1992, Chapter 2) allows us to efficiently store sparse matrices, so that more matrix elements can be kept in core memory. When properly designed, a sparse indexing system also allows for efficient computation of matrix-vector products by exploiting the fact that many of the elemental multiplications are guaranteed to be zero, since the matrices are sparse.

Bathe (1982, App. 2.3) recommends renumbering nodes in order to reduce matrix bandwidth, which directly affects the computational cost of solution methods. Modern sparse matrix-vector product routines incorporate this consideration.

4.2.2 Factorisation of mass matrix

Given a standard eigensolver, the naive approach to solving the generalised eigenproblem

$$\mathbf{K\Phi} = \mathbf{M\Lambda\Phi} \qquad (4.1)$$

is to invert the mass matrix to form the standard eigenproblem, which destroys the sparsity structure of the mass and stiffness matrices.

We can perform a Cholesky factorisation on a positive-definite mass matrix

$$\mathbf{M} = \mathbf{L}\mathbf{L}^{\mathbf{T}} \qquad (4.2)$$

and it can be seen in Figure ??(b) that the Cholesky factor preserves sparsity.

Given this factorisation, we proceed to solve Equation 4.1 in standard form[1]

$$\mathbf{L}^{-1}\mathbf{K}\mathbf{L}^{-\mathbf{T}}\mathbf{\Phi} = \mathbf{\Lambda}\mathbf{\Phi} \qquad (4.3)$$

The Lanczos solver allows the user to specify the matrix implicitly by providing a matrix-vector product routine

$$
\begin{aligned}
\mathbf{v} \; &= \; \mathbf{H}\mathbf{u} \\
&= \; \mathbf{L}^{-1}\mathbf{K}\mathbf{L}^{-\mathbf{T}}\mathbf{u} \qquad (4.4)
\end{aligned}
$$

An efficient method of computing Equation 4.4 is to use sparse triangular-solve and matrix-vector product routines in the following procedure

$$
\begin{aligned}
\mathbf{L}^{\mathbf{T}}\mathbf{w} &= \quad \mathbf{u} & \text{(giving } \mathbf{w}) \\
\mathbf{y} &= \quad \mathbf{K}\mathbf{w} & \text{(giving } \mathbf{y}) \\
\mathbf{L}\mathbf{v} &= \quad \mathbf{y} & \text{(giving } \mathbf{v}) \qquad (4.5)
\end{aligned}
$$

4.2.3 Spectral shift

See Grimes, Lewis & Simon (1994).

4.3 Mass lumping

Eigenmodes for volumetric models are much more stable with respect to small shape changes than for a surface model. This is simply because there are more constraints internally due to volumetric elasticity.

Eigenmodes can be made even more invariant by using a *lumped mass* approximation to the *consistent mass* representation derived in Section 2.3.2. This approximation gives a diagonal mass matrix instead of a fully populated one, by lumping the element's distributed mass at nodal positions.

There are a number of lumping schemes, the simplest of which Hughes (1987) calls *special lumping*. This discards off-diagonal elements, and scales the remaining ones to give the same total elemental mass. This scheme has the advantage of ensuring that the mass matrix is positive definite.

---

[1]Using $\mathbf{L}^{-1}\mathbf{K}\mathbf{L}^{-\mathbf{T}}$ preserves symmetry.

Special lumping is a very good approximation for a uniformly tessellated structure with uniformly distributed mass and stiffness for which only low-frequency response is required (Hitchens 1992, Ch. 1). In addition, the FEM tends to over-estimate natural frequencies, whereas the lumped mass approximation tends to compensate for this (see Section 4.3.1 below). The associated mode shapes tend to have cumulative error, which does not matter since we are not interested in dynamic response.

Qualitatively, special lumping "conditions" the landscape of the Rayleigh Quotient so that local minima representing eigenvectors are better separated and made "deeper", so that peturbed shapes still tend to give the same mode shapes.

The inertial effects of distributed mass are effectively muted by mass lumping, so that the structure's frequency response is determined mainly by the stiffness matrix. This is why lumped-mass mode shapes are good shape features: the stiffness matrix is effectively a shape "connectivity matrix" for a uniformly elastic volume.

4.3.1 Invariance to mesh density

The FEM converges monotonically to the true mode shapes as the mesh gets finer (Bathe 1982, Chapter 4). We have observed that this can lead to "flipping" in the eigenvalue sequence, so that mode shapes can no longer be reliably matched for two tessellations of the same shape. Mass lumping tends to reduce this effect, thus making mode shapes more invariant to changes in mesh density.

4.3.2 Eigenmode computation

A diagonal mass matrix is trivial to invert, hence Equations 4.5 can be replaced with a single matrix vector product using $\mathbf{M}^{-1}\mathbf{K}$. This reduces computation time substantially, since code profiling suggests that up to 65% of CPU cycles are spent in the two sparse triangular-solve routines in Equations 4.5

4.4 Mode shapes

The following six Chapters 5–10 show the mode shapes of the following volumetric models

(a) Uniform meshing

(b) Non-uniform meshing

(c)
Non-uniform meshing and bending

(d) Sparse meshing

(e)
High spatial-frequency noise

26

(f)
Low spatial-frequency noise

Figure 4.2: Ellipsoid shape models with peturbations

Chapter 5

Ellipsoid with uniform tessellation

Figures 5.1 to 5.9 show the mode shapes for an ellipsoidal volume tessellated uniformly with tetrahedral elements.

Figure 5.1: Snapshots of 1st non-rigid mode shape

Figure 5.2: Snapshots of 2nd non-rigid mode shape. This is a repeated version of the 1st non-rigid mode shape, since the ellipsoidal volume has 2 planes of mirror symmetry.

Figure 5.3: Snapshots of 3rd non-rigid mode shape

Figure 5.4: Snapshots of 4th non-rigid mode shape

Figure 5.5: Snapshots of 5th non-rigid mode shape

Figure 5.6: Snapshots of 6th non-rigid mode shape

Figure 5.7: Snapshots of 7th non-rigid mode shape

Figure 5.8: Snapshots of 8th non-rigid mode shape

Figure 5.9: Snapshots of 9th non-rigid mode shape

Chapter 6

Ellipsoid with non-uniform tesselation

Figures 6.1 to 6.9 show the mode shapes for an ellipsoidal volume sampled more densely near the poles, both on the surface and internally.

(e.g. recovered from MR, or other previously built models, or from a tesselation technique, or manually sampled?)

Figure 6.1: Snapshots of 1st non-rigid mode shape

Figure 6.2: Snapshots of 2nd non-rigid mode shape

Figure 6.3: Snapshots of 3rd non-rigid mode shape

Figure 6.4: Snapshots of 4th non-rigid mode shape

Figure 6.5: Snapshots of 5th non-rigid mode shape


Figure 6.6: Snapshots of 6th non-rigid mode shape



Figure 6.7: Snapshots of 7th non-rigid mode shape

Figure 6.8: Snapshots of 8th non-rigid mode shape

Figure 6.9: Snapshots of 9th non-rigid mode shape

Chapter 7

Bent ellipsoid with non-uniform tesselation

Figures 7.1 to 7.9 show the mode shapes for a bent version of the ellipsoid in Section **??**
    (shows stability of modes of large spatial-scale shape change) eg from ultrasound to MR ground truth. note that 2d ultrasound probes never reconstruct scanned images perfectly. an equatorial cross-section of a pingpong ball is reconstructed as ellipsoidal.

Figure 7.1: Snapshots of 1st non-rigid mode shape

Figure 7.2: Snapshots of 2nd non-rigid mode shape

Figure 7.3: Snapshots of 3rd non-rigid mode shape

31

Figure 7.4: Snapshots of 4th non-rigid mode shape

Figure 7.5: Snapshots of 5th non-rigid mode shape

Figure 7.6: Snapshots of 6th non-rigid mode shape

Figure 7.7: Snapshots of 7th non-rigid mode shape

Figure 7.8: Snapshots of 8th non-rigid mode shape

Figure 7.9: Snapshots of 9th non-rigid mode shape

Chapter 8

Ellipsoid with sparse tessellation

Figures 8.1 to 8.9 show the mode shapes for sparsely sampled ellipsoid. This shows that mode shapes are stable for models with different sampling densities (e.g. MR and ultrasound).

Figure 8.1: Snapshots of 1st non-rigid mode shape

Figure 8.2: Snapshots of 2nd non-rigid mode shape

Figure 8.3: Snapshots of 3rd non-rigid mode shape

Figure 8.4: Snapshots of 4th non-rigid mode shape

Figure 8.5: Snapshots of 5th non-rigid mode shape

Figure 8.6: Snapshots of 6th non-rigid mode shape

Figure 8.7: Snapshots of 7th non-rigid mode shape

Figure 8.8: Snapshots of 8th non-rigid mode shape

Figure 8.9: Snapshots of 9th non-rigid mode shape

Chapter 9

Ellipsoid with noise of high spatial frequency

Figures 9.1 to 9.9 show the mode shapes for noisily sampled ellipsoid. The noise added to each sample point is correlated Gaussian with a standard deviation of 1 element width. This simulates the effect of errors in sampling, in landmark specification (see (Syn et al. 1995)), in proprioception etc.

Figure 9.1: Snapshots of 1st non-rigid mode shape

Figure 9.2: Snapshots of 2nd non-rigid mode shape

Figure 9.3: Snapshots of 3rd non-rigid mode shape

Figure 9.4: Snapshots of 4th non-rigid mode shape

Figure 9.5: Snapshots of 5th non-rigid mode shape

Figure 9.6: Snapshots of 6th non-rigid mode shape


Figure 9.7: Snapshots of 7th non-rigid mode shape

Figure 9.8: Snapshots of 8th non-rigid mode shape

Figure 9.9: Snapshots of 9th non-rigid mode shape

Chapter 10

Ellipsoid with noise of low spatial frequency

Figures 10.1 to 10.9 show the mode shapes for an ellipsoid with noise of large spatial correlation added to each node. This simulates the effect of imaging distortions in 2D ultrasound acquisition[1], or interpatient variability in structures being scanned.

Figure 10.1: Snapshots of 1st non-rigid mode shape

Figure 10.2: Snapshots of 2nd non-rigid mode shape

Figure 10.3: Snapshots of 3rd non-rigid mode shape

---

[1]2D ultrasound probes never reconstruct scanned images with perfect linearity.

Figure 10.4: Snapshots of 4th non-rigid mode shape

Figure 10.5: Snapshots of 5th non-rigid mode shape

Figure 10.6: Snapshots of 6th non-rigid mode shape

Figure 10.7: Snapshots of 7th non-rigid mode shape

Figure 10.8: Snapshots of 8th non-rigid mode shape

Figure 10.9: Snapshots of 9th non-rigid mode shape

Bibliography

Amit, Y., Grenander, U. & Piccioni, M. (1991), 'Structural image restoration through deformable templates', *Journal of the American Statistical Association* **86**(414), 376–387.

Ballard, D. & Brown, C. (1982), *Computer Vision*, Prentice-Hall.

Bathe, K.-J. (1982), *Finite Element Procedures in Engineering Analysis*, Prentice-Hall.

Bookstein, F. (1991), *Morphometric tools for landmark data : geometry and biology*, 1st edn, Cambridge University Press.

Brady, M. & Scott, G. (1988), Parallel algorithms for shape representation, *in* I. Page, ed., 'Parallel Architectures and Computer Vision', Clarendon Press, Oxford, pp. 97–118.

Cohen, L. & Cohen, I. (1993), 'Finite-element methods for active contour models and balloons for 2D and 3D images', *IEEE Trans. Pattern Analysis and Machine Intelligence* **15**(11), 1131–1147.

Coppini, G., Poli, R. & Valli, G. (1995), 'Recovery of the 3-d shape of the left ventricle from echocardiographic images', *IEEE Transactions on Medical Imaging* **14**(2), 301–317.

Fujita, S. (1986), Synthetico-analytic approach to mechanism of shaping of the brain through simulation with 3D computer graphics, *in* S. Ishizaka, ed., 'Science on Form: Proceedings of the First International Symposium for Science on Form', KTK Scientific Publishers, Tokyo, pp. 339–352.

Golub, G. & van Loan, C. (1989), *Matrix computations*, 2nd edn, The Johns Hopkins University Press.

Gordon, C., Webb, D. & Wolpert, S. (1992), 'You cannot hear the shape of a drum', *Bulletin of the American Mathematical Society*.

Grimes, R., Lewis, J. & Simon, H. (1994), 'A shifted block Lanczos algorithm for solving sparse symmetric generalized eigenproblems', *SIAM Journal Matrix Anal. Appl.*

Hitchens, D., ed. (1992), *A Finite Element Dynamics Primer*, NAFEMS, Birniehill, East Kilbride, Glasgow.

Hu, M. (1962), 'Visual pattern recognition by moment invariants', *IEEE Transactions on Information Theory* **8**, 179–187.

Hughes, T. (1987), *The Finite Element Method*, Prentice-Hall, Englewood Cliffs, New Jersey.

Kac, M. (1966), 'Can one hear the shape of a drum?', *The American Mathematical Monthly* **73**(4), 1–23.

Kimia, B., Tannenbaum, A. & Zucker, S. (1990), Toward a computational theory of shape: an overview, *in* '1st ECCV 90', Vol. 427 of *Lecture Notes in Computer Science*, Springer Verlag, pp. 23–27.

Kreysig, E. (1988), *Advanced Engineering Mathematics*, 6th edn, John Wiley and Sons.

Martin, J., Pentland, A. & Kikinis, R. (1994), Shape analysis of brain structures using physical and experimental modes, Technical Report 276, MIT Media Laboratory, Perceptual Computing.

Nastar, C. & Ayache, N. (1993), Non-rigid motion analysis in medical images : a physically based approach, *in* 'Proceedings of the 13th International Conference Information Processing in Medical Imaging', Vol. 687 of *Lecture Notes in Computer Science*, Springer-Verlag.

Parlett, B. (1980), *The Symmetric Eigenvalue Problem*, Prentice-Hall, Englewood Cliffs, New Jersey.

Pentland (1990), 'Automatic extraction of deformable part models', *International Journal of Computer Vision* **4**, 107–126.

Pentland, A. & Horowitz, B. (1991), 'Recovery of nonrigid motion and structure', *IEEE Trans. Pattern Analysis and Machine Intelligence* **13**(7), 730–742.

Pentland, A. & Sclaroff, S. (1991), 'Closed-form solutions for physically based shape modeling and recognition', *IEEE Trans. Pattern Analysis and Machine Intelligence* **13**(7), 715–729.

Pentland, A. & Williams, J. (1989), 'Good vibrations: Modal dynamics for graphics and animation', *Computer Graphics* **23**(4), 215–222.

Pizer, S. & Burbeck, C. (1994), Cores as the basis for object vision in medical images, *in* 'Image Perception (1994)', Vol. 2166 of *SPIE*, pp. 191–198.

Pizer, S., Oliver, W. & Bloomberg, S. (1987), 'Hierarchical shape description via the multi-resolution symmetric axis transform', *IEEE Trans. Pattern Analysis and Machine Intelligence* **PAMI-9**, 505–511.

Poli, R., Coppini, G. & Valli, G. (1994), 'Recovery of 3D closed surfaces from sparse data', *CVGIP : Image Understanding* **60**(1), 1–25.

Saad, Y. (1992), *Numerical methods for large eigenvalue problems*, Manchester University Press, Oxford Road, Manchester.

Sclaroff, S. & Pentland, A. (1995), 'Modal matching for correspondence and recognition', *IEEE Trans. Pattern Analysis and Machine Intelligence* **17**(6), 545–561.

Staib, L. & Duncan, J. (1992), 'Boundary finding with parametrically deformable models', *IEEE Trans. Pattern Analysis and Machine Intelligence* **14**(11), 1061–1075.

Syn, M. & Prager, R. (1995*a*), Bayesian registration of models using fem eigenmodes, Technical Report CUED/F-INFENG/TR213, Cambridge University Engineering Department, Trumpington Street, CB2 1PZ.

Syn, M. & Prager, R. (1995*b*), Mesh models for three-dimensional ultrasound imaging, Technical Report CUED/F-INFENG/TR210, Cambridge University Engineering Department, Trumpington Street, CB2 1PZ.

Syn, M. & Prager, R. (1995*c*), A model based approach to three-dimensional ultrasound imaging, *in* 'Proc. XIVth Information Processing in Medical Imaging 95', Brest, France.

Syn, M., Gosling, J., Berman, L. & Prager, R. (1995), Software for model-based segmentation in three-dimensional freehand ultrasound imaging, Technical Report CUED/F-INFENG/TR214, Cambridge University Engineering Department, Trumpington Street, CB2 1PZ.

Syn, M., Gosling, J., Prager, R., Berman, L. & Crowley, J. (1994), Tracking the interframe deformation of structures in 3d ultrasound imaging, *in* R. Robb, ed., 'Proc. Visualization in Biomedical Computing 94', Vol. 2359 of *SPIE*, SPIE, Rochester, Minnesota, pp. 170–178.