Fast segmentation, tracking, and analysis of deformable objects — **Source link** ⬀

C. Nastar, N. Ayache

Institutions: French Institute for Research in Computer Science and Automation

Related papers:

- Snakes : Active Contour Models

- Closed-form solutions for physically based shape modeling and recognition

- Dynamic 3D models with local and global deformations: deformable superquadrics

- Active shape models—their training and application

- Constraints on deformable models: recovering 3D shape and nongrid motion

Fast segmentation, tracking and analysis of deformable objects

Chahab Nastar, Nicholas Ayache

HAL Id: inria-00077023
https://hal.inria.fr/inria-00077023

Submitted on 29 May 2006

Rapports de Recherche

1 9 9 2

25ème
anniversaire

FAST SEGMENTATION, TRACKING, AND ANALYSIS OF DEFORMABLE OBJECTS

Chahab NASTAR
Nicholas AYACHE

RR.1783

Fast Segmentation, Tracking, and Analysis of Deformable Objects *

Chahab Nastar, Nicholas Ayache

November 10, 1992

INRIA

BP 105 78153 Le Chesnay Cédex

e-mail: Chahab.Nastar@inria.fr

Programme 4

Robotique, Image et Vision

Abstract

In this paper, we present a physically-based multidimensional deformable model, which can be used to track and to analyze non-rigid motion of dynamic structures in time sequences of 2D or 3D medical images.

The model considers an object undergoing an elastic deformation as a set of masses linked by springs, where the classical natural lengths of the springs is set equal to zero, and is replaced by a set of constant equilibrium forces, which characterize the shape of the elastic structure in the absence of external forces.

This model has the extremely nice property of yielding dynamic equations which are linear and decoupled for each coordinate, whatever the amplitude of the deformation.

Compared to the former work of Terzopoulos and his colleagues [11, 22, 20, 12] and Pentland and his colleagues [17, 16, 18, 10], our model provides a reduced algorithmic complexity, and a sound framework for modal analysis, which allows a compact representation of a general deformation by a reduced number of parameters.

The power of the approach to segment, track and analyze 2-D and 3-D images is demonstrated by a set of experimental results on various complex medical images.

Segmentation, suivi, et analyse du mouvement des objets déformables

Résumé : *Nous présentons un modèle physique d'objets déformables capable de segmenter, de suivre, et d'analyser les mouvements de structures non-rigides. Dans ce modèle, un objet subissant une déformation élastique est considéré comme un ensemble de masses ponctuelles reliées entre elles par des ressorts de longueur à vide nulle. Ce choix est compensé par un champ de forces d'équilibre, constant, qui caractérise la forme de l'objet en l'absence de forces extérieures. Une des propriétés intéressantes de ce modèle est que les équations différentielles qu'il fournit sont linéaires et découplées en chaque coordonnée, quelle que soit l'amplitude de la déformation. Par rapport aux approches précédentes de l'équipe de Terzopoulos [11, 22, 20, 12], et de l'équipe de Pentland [17, 16, 18, 10], notre modèle aboutit à une complexité algorithmique réduite, et fournit un cadre rigoureux pour l'analyse modale, qui permet une représentation compacte de la déformation d'un objet par un nombre réduit de paramètres. Enfin, nous présentons des résultats expérimentaux sur des images médicales 2D et 3D.*

2

1 Introduction

In this paper, we present a physically-based multidimensional deformable model, which can be used to track and to analyze non-rigid motion of dynamic structures in time sequences of 2D or 3D medical images.

The model considers an object undergoing an elastic deformation as a set of masses linked by springs, where the classical natural lengths of the springs is set equal to zero, and is replaced by a set of constant equilibrium forces, which characterize the shape of the elastic structure in the absence of external forces.

This model has the extremely nice property of yielding dynamic equations which are linear and decoupled for each coordinate, whatever the amplitude of the deformation.

Compared to the former work of Terzopoulos and his colleagues [11, 22, 20, 12] and Pentland and his colleagues [17, 16, 18, 10], our model provides a reduced algorithmic complexity, and a sound framework for modal analysis, which allows a compact representation of a general deformation by a reduced number of parameters.

We show how the model is built, and how to integrate the dynamic equations through time, with the possibility to take into account some tracking contraints, as the knowledge of sparse anatomical features. We then dicuss the use of two different types of modes to describe the deformation, and propose to use the one which best fits our physical model. Finally, we illustrate by a set of experiments on synthetic and real data, the validity and the power of our approach to track and analyze the motion of anatomical structures both in two and three dimensions.

Concerning the tracking and the analysis of deformable objects, our work differs from

but is related to the one of [13], who constrains locally the deformations to be conformal, and also to the one of [7, 1, 4], where tracking takes into account local differential properties of the surfaces. The principal warps analysis of [3] as well as the Fourier decomposition of [19] both show a similar spirit as our modal analysis. Finally the work of [5] and [8] provides an alternative way to segment volumetric images, but without explicitly trying to model the deformations.

2 The model

We consider both the surface and volumetric properties of the objects at hand. We restrict ourselves to elastic deformations, i.e. we assume the object recovers its reference configuration as soon as all applied forces causing deformation are removed. Modelling an elastic boundary M can be achieved by a mesh of n virtual masses on the contour, each mass being attached to its neighbors by springs of stiffness k and natural length l_0, as shown in figures 1 and 2. These springs model the elastic *surface* properties of the object.

We can improve the modelling by attaching extra springs between non-neighbor nodes in order to model some *volumetric* elastic properties inside the object (see appendix A). These springs constrain the general form of the object within its deformation. The boundary M modelled as above will also be called *structure*. Such a structure can be easily deformed to match the contour of an object of interest, thus performing a *segmentation* step. Now if we take a set of images displaying the deformation of the object, structure M can also achieve simultaneously both *segmentation* and *tracking* of the object's surface through time.
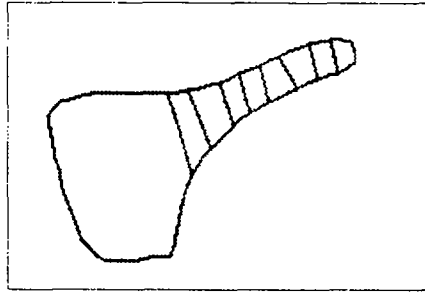
4

Figure 1: A 2D valve model with surface and volume springs
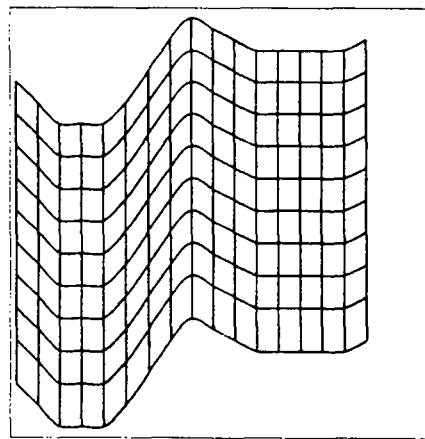
Figure 2: A 3D mass-spring mesh with surface springs
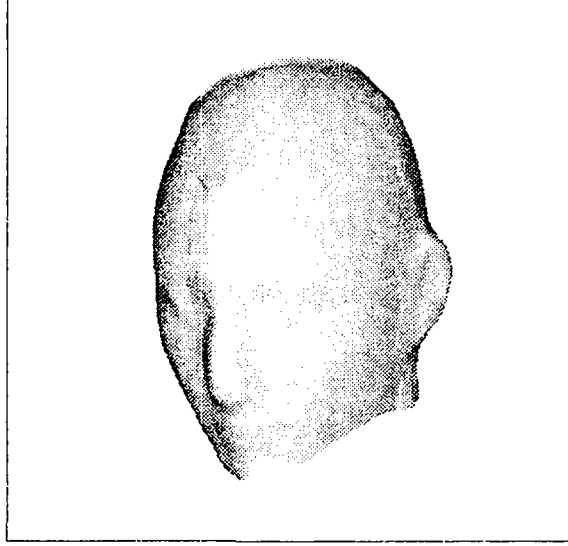
Figure 3: Segmentation of a human head from a 3D magnetic resonance image

3 The governing equation

The system under study is composed of n masses that are positionned at time t on the points $((M_1)_t, (M_2)_t, \ldots, (M_n)_t)$. Let :

$$M_t = [(M_1)_t, (M_2)_t, \ldots, (M_n)_t]^T$$

The evolution of the structure is governed by the fundamental equation of dynamics :

$$F_i = m_i a_i \qquad i = 1, \ldots, n$$

where m_i is the mass of point M_i and a_i its acceleration under total load F_i.

Now, what are the applied forces to the point M_i at time t? First, there is the elastic force due to its neighbors (assuming the 2D case where a mass has two neighbors) :

$$F_e(M_i, t) = -k[(M_{i-1}M_i)_t - l_0 \frac{(M_{i-1}M_i)_t}{\|(M_{i-1}M_i)_t\|}] - k[(M_{i+1}M_i)_t - l_0 \frac{(M_{i+1}M_i)_t}{\|(M_{i+1}M_i)_t\|}] \qquad (1)$$

6

where l_0 is the natural length of the springs.

Then, a damping force can be considered. This force is generally set proportional to the point velocity v_i :

$$F_d(M_i, t) = -c_i v_i$$

Moreover, an external load $F_{ext}(M_i, t)$ acts on each node M_i. Note that the structure must be given an initial position at time t_0. Since we want the structure to hold in this position, we have to apply on each node a force F_{eq} so that time t_0 is an equilibrium state of the structure. This force is similar to the force our fingers apply to an elastic in order to give a certain form to it. Thus, at initial time, we have :

$$F_{eq}(M_i) + F_e(M_i, t_0) = 0 \qquad (2)$$

that is, the sum of the forces acting on each is zero. We assume that at any future time this equilibrium force is constant.

Hence, the total load at time t is for each node M_i :

$$F(M_i, t) = F_e(M_i, t) + F_d(M_i, t) + F_{ext}(M_i, t) + F_{eq}(M_i)$$

The governing equation of the structure is :

$$F_e(M_i, t) + F_d(M_i, t) + F_{ext}(M_i, t) + F_{eq}(M_i) = m_i a_i$$

This equation, expressed for all n nodes, leads to a *nonlinear* system of *coupled* differential equations (for each node, the x, y and z displacements are coupled, and the displacement of a node depends on its neighbors' displacement, as it is clearly shown in equation 1). One possible approach is the solving of these nonlinear coupled equations by a rather costly

7

iterative procedure [20]. Another approach is the local approximation of the system as a linear one [18], an approximation which is only valid for small displacements, otherwise it must be recomputed at differents time steps.

We propose in this paper to set $l_0 = 0$, an assumption which does not restrict the generality of the problem because the constant equilibrium forces $F_{eq}(M_i)$ allow an arbitrary initial configuration. The advantage of this assumption is that we end up with a set of *linear* differential equations with node displacements *decoupled* in each coordinate, whatever the magnitude of the displacements.

Indeed, setting $l_0 = 0$ and denoting $U_i(t) = (M_i)_t - (M_i)_{t_0}$ we can express the elastic force in terms of nodal displacements :

$$F_e(M_i, t) = -k[(M_{i-1}M_i)_t + (M_{i+1}M_i)_t] = F_e(M_i, t_0) - k[2U_i(t) - U_{i+1}(t) - U_{i-1}(t)]$$

Remembering equation (2), the governing equation becomes :

$$-k[2U_i(t) - U_{i+1}(t) - U_{i-1}(t)] - c_i\dot{U}_i(t) + F_{ext}(M_i, t) = m_i\ddot{U}_i(t)$$

These equations can now be written in a matrix form :

$$M\ddot{U} + C\dot{U} + KU = F_t \tag{3}$$

where K is the stiffness matrix of the element (see appendix A), C its damping matrix, M its mass matrix, and F_t is the external force field at time t, and U is the nodal displacements vector ($U = M_t - M_{t_0}$). Note that in the above formulation the stiffness matrix is constant (i.e. the same at each time step), and no recomputation is necessary.

In 3D, the above matrix-form equation represents $3n$ coupled second-order ordinary differential equations. However, by setting $l_0 = 0$, this equation is separable in three n order

8

equations with respect to x, y and z. Therefore, from now on, except when specified, vectors and matrices will be of order n. For instance, U will refer indifferently to the x, y or z nodal displacements vector.

4   An elastic potential

4.1   General form of the potential

We make the common assumption that the external force field at time t, derives from a potential field V_t :

$$F_t = -\nabla V_t \tag{4}$$

Various potential fields can successfully solve the problem. In the "snake" formulation [11, 20], in order to have the snake attracted to contours with large image gradients, the potential on any point M_i is set to :

$$V(M_i) = -\|\nabla G_\sigma \star I(M_i)\|$$

where $G_\sigma \star$ denotes convolution with a gaussian smoothing filter of width σ, so that the edges of the image can attract the snake from a distance. Note that the farther the snake from the edges, the weaker the attraction force field.

In our formulation, we need a potential field that can *a posteriori* describe which nodal displacements are the most likely. One interesting choice for the potential is :

$$V(M_i) = \frac{1}{2} k_p \|P_i M_i\|^2 \tag{5}$$

where P_i is the closest boundary point to M_i. The algorithm for obtaining this potential on each point of an image is described in [6].

Another advantage of this potential is that it can be represented by a spring of natural length zero and of stiffness k_p joining M_i to P_i :

$$F_p(M_i) = -\nabla V(M_i) = k_p M_i P_i$$

Hence, the whole model remains elastic. Moreover, the farther the structure from the edges, the larger the attraction force. This procedure speeds up the segmentation of the image edges by the deformable structure.

Thus, we have defined n springs of natural length zero and of stiffness k_p joining structure nodal points M_i to their closest boundary point P_i. The external force field at time t is :

$$F_t = k_p D_t$$

where $D_t = [(M_1 P_1)_t, (M_2 P_2)_t, \ldots, (M_n P_n)_t]^T$ is a distance field which may not be unique (several structure points may have the same closest boundary point).

The governing equation is, for this choice of the potential field :

$$M\ddot{U} + C\dot{U} + KU = k_p D_t \tag{6}$$

4.2 Automatic selective potential

Suppose that the displacement of node i_0 is known (M_{i_0} matches M'_{i_0}) by an alternative method (for example by use of anatomical or artificial landmarks such as high curvature points, see [1, 4]). The displacement of landmark M_{i_0} has to be computed separately, in other terms, M_{i_0} must be submitted to another potential field that the non-landmark nodes, having it attracted by M'_{i_0}.

This can be achieved automatically with 2D and 3D feature extraction algorithms [14, 21]. We set the potential differently on every point $(M_i)_t$, depending on $(M_i)_t$ being a landmark

or not. If $(M_{i_0})_t$ is a landmark, we find its new deformed position, namely, the point M'_{i_0} that has the same feature (for instance, $(M_{i_0})_t$ and M'_{i_0} are both maximum curvature points on the surface). We submit M_{i_0} to the potential field :

$$V(M_{i_0})_t = \frac{1}{2}\alpha \|M'_{i_0}(M_{i_0})_t\|^2$$

For the non-landmark nodes $(M_i)_t$, the potential is computed as :

$$V(M_i)_t = \frac{1}{2}k_p\|(P_i)_t(M_i)_t\|^2$$

where $(P_i)_t$ is the closest boundary point to $(M_i)_t$ at time t, and $\alpha \gg k_p$. Intuitively, the physical interpretation of the landmark constraint for M_{i_0} can be seen as attaching a stiff spring between M_{i_0} and M'_{i_0}, and discarding the spring between M_{i_0} and P_{i_0} (figure 4).

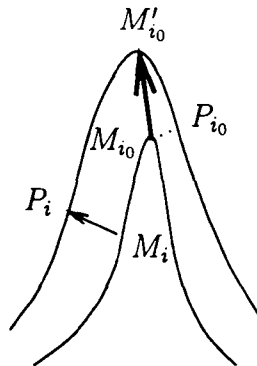Figure 4: Attraction force for landmark node M_{i_0} and for non-landmark node M_i

5 Direct integration of the governing equation

As the distance field may not be unique, there may be several acceptable solutions to the tracking problem. This implies that we have to give the structure a *reasonable* position at initial time $t_0 = 0$, that is, the initialization has to minimize the risk of conflicts.

The governing equation can then be directly integrated through time by various methods [2, 20]. One of the simplest and quickest ones is the explicit Euler method :

$$\begin{cases} \ddot{U}_t = M^{-1}(F_t - C\dot{U}_t - KU_t) \\ \dot{U}_{t+\Delta t} = \dot{U}_t + \Delta t \ddot{U}_t \\ U_{t+\Delta t} = \Delta t \dot{U}_{t+\Delta t} \end{cases}$$

where the damping, inertial and stiffness constants are chosen so that the system is over-damped.

6 Results on 2D objects

6.1 Segmentation

Figure 5 outlines the importance of the volume springs. A classical curvilinear snake cannot segment the simulated valve given this initial position : it collapses against the bottom boundary of the second valve contour (figure 5.a and 5.b). Volume springs prevent the elastic structure from collapsing (figure 5.c and 5.d). A hundred of nodal points were used for the modelling.
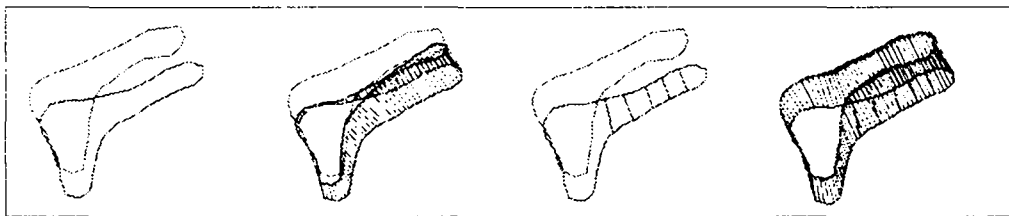

Figure 5: Segmentation of the mitral valve without (a,b) and with (c,d) volume springs

12

6.2 Tracking

Figure 6 shows the modification of the nodal displacements when the displacement of one of the nodes is known. This result is an alternative to the results obtained with another energy-minimizing procedure [4]. Sixty nodal points were used for the modelling.
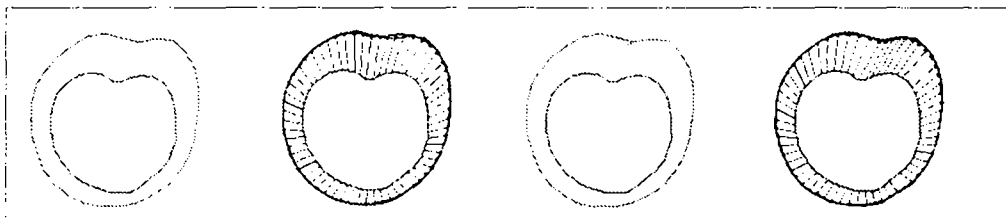
Figure 6: Unconstrained and constrained tracking of a 2D contour

7 Modal analysis

This well-known approach in the field of mechanical engineering [2], was brought to the field of computer vision by Pentland's team [18, 10]. We outline the general principle and present our formulation, which differs from the one of Pentland in being better adapted to the underlying physical model we are using.

7.1 Change of basis

Instead of solving directly the equilibrium equation (3), one can transform it by a change of basis :

$$U = P\tilde{U}$$

where P is the square n order nonsingular transformation matrix to be determined, and $\tilde{U}$ is referred to as the *generalized displacements* vector (see [2]).

13

The question is finding a suitable transformation matrix P that would reduce the bandwidth of the system matrices. One effective way of achieving this is using the displacement solutions of the free vibration equilibrium equations :

$$M\ddot{U} + KU = 0 \qquad (7)$$

where M is the mass matrix of the structure. The solution to (7) can be postulated to be of the form :

$$U = \phi \sin \omega(t - t_0)$$

where ϕ is a vector of order n, t the time variable, t_0 a time constant, and ω the frequency of vibration of the vector ϕ. We can now substitute this expression of U in (7) to determine ϕ and ω, which leads to the generalized eigenproblem :

$$K\phi = \omega^2 M\phi \qquad (8)$$

This equation yields n eigensolutions $(\omega_1^2, \phi_1) \ldots (\omega_n^2, \phi_n)$. It can be shown (see appendix B) that the eigenvectors ϕ_i (also called shape vectors) are M-orthonormal and K-orthogonal :

$$\phi_i^T M \phi_j = \begin{cases} 1; & i = j \\ 0; & i \neq j \end{cases}$$

$$\phi_i^T K \phi_j = \begin{cases} \omega_i^2; & i = j \\ 0; & i \neq j \end{cases}$$

$$0 \leq \omega_1^2 \leq \omega_2^2 \leq \ldots \leq \omega_n^2$$

The vectors $\left(\vec{\phi}_i\right)_{i=1,\ldots,n}$ can then be referred to as the structure's *eigenbasis*.

The structure's eigenbasis is determined as soon as the mass matrix and the stiffness matrix of the structure are defined.

14

The former equations can be rewritten in a matrix form :

$$\tilde{M} = \Phi^T M \Phi = I_n \tag{9}$$

$$\tilde{K} = \Phi^T K \Phi = \Omega^2 \tag{10}$$

where Φ is the matrix whose columns are the eigenvectors ϕ_i ($\Phi = [\phi_1, \phi_2, \ldots, \phi_n]$) and Ω^2 is a diagonal matrix which stores the eigenvalues ω_i^2 on its diagonal ($\Omega^2 = diag(\omega_1^2, \omega_2^2, \ldots, \omega_n^2)$). The matrix Φ is then chosen to be the transformation matrix P :

$$U = \Phi \tilde{U} = \sum_{i=1}^{n} \tilde{u}_i \phi_i \tag{11}$$

Equation (11) is referred to as the modal superposition equation. $\tilde{u}_i$ is the ith-mode *amplitude* within displacement U.
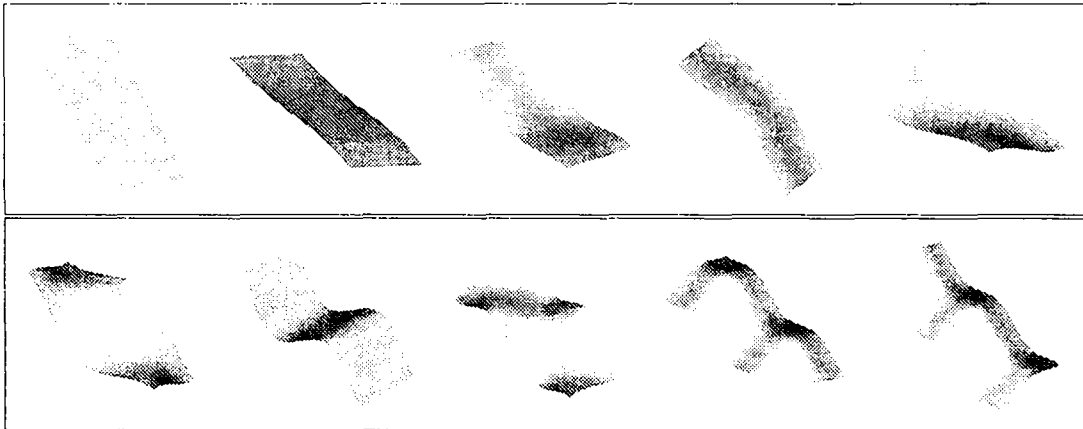

Figure 7: Energy-increasing eigenmodes of a flat element

15

7.2 Decoupling the equations of motion

Premultiplying the governing equation (3) by Φ^T and expressing the displacement vector in terms of the generalized displacement vector leads to :

$$\ddot{\tilde{U}} + \tilde{C}\dot{\tilde{U}} + \Omega^2 \tilde{U} = \Phi^T F \tag{12}$$

Hence, if the matrix $\tilde{C} = \Phi^T C \Phi$ is diagonal, then the above matrix-form equations decouple into n scalar equations :

$$\ddot{\tilde{u}}_{t,i} + \tilde{c}_i \dot{\tilde{u}}_{t,i} + \omega_i^2 \tilde{u}_{t,i} = \tilde{F}_{t,i} \qquad i = 1, \ldots, n. \tag{13}$$

Solving these equations at time t leads to $(\tilde{u}_{t,i})_{i=1,\ldots,n}$, and the displacement U_t of the structure's nodes is obtained by the modal superposition equation.

7.3 Modal approximation

In practice, we wish to approximate the nodal points displacements U by picking up p significant modes, where $p \ll n$:

$$U_t \approx \sum_{i=1}^{p} \tilde{u}_{t,\sigma(i)} \phi_{\sigma(i)} \tag{14}$$

where σ is a suitably chosen permutation.

The $\left(\vec{\phi}_{\sigma(i)}\right)_{i=1,\ldots,p}$ are the *reduced eigenbasis* of the structure.

This is the major advantage of modal analysis : it provides an approximate but quite accurate solution by selecting a few number of modes. The participation of each mode to the motion is ordered. In those terms, we can compare modal analysis to principal components analysis. Whereas in principal components analysis the object, which is in a high order space, is displayed after being projected into a lower order subspace maximizing

16

its spreading, modal analysis projects the nodal displacements vector U (which is in a n order space) into a subspace of much lower order p, with minimum loss of accuracy.

It is now apparent that modal analysis provides a compact description of the motion allowing straightforward interpretation and comparison of motions in terms of modal amplitudes. For a 3D image, the motion parameters are reduced from $3n$ (x,y, and z displacements of the n nodes) to $3p$ (x, y and z modal amplitudes). Typically, if n is about 1000, and we assume that 20 modes provide a good approximation of the motion, we reduce the number of parameters from 3000 to 60. This represents a compaction of 50 for the motion description. In 2D, although less spectacular, compaction rates of an order of magnitude are quite common (see figure 8).

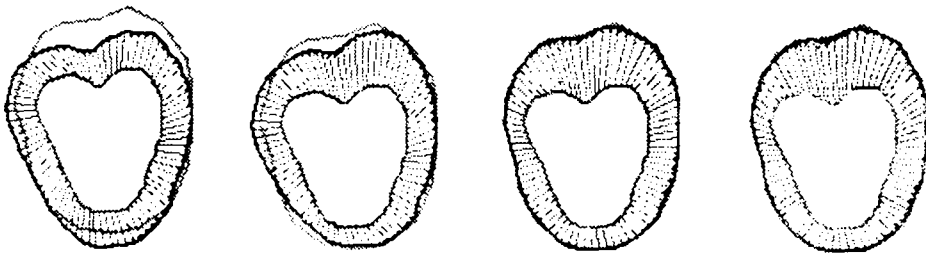Figure 8: Superposition of high amplitude modes. a. 2 modes (compaction : 80, recovery : 50.6 %) b. 4 modes (compaction : 40, recovery : 68.7 %) c. 20 modes (compaction : 8, recovery : 98.2 %) d. All 160 modes (compaction : 1, recovery : 100.0 %)

7.4 Selecting the most significant modes

In this section we concentrate on the choice of the p modes that can most accurately describe the object's motion.

The most accurate choice is the choice of the p modes of *highest amplitude*, having omitted $\tilde{u}_{t,i}$ by equation (13). Instead of choosing the number p of significant modes, we can rather choose a constant μ $(0 < \mu \leq 1)$ which represents the confidence we need for the approximation of the displacement :

$$\frac{\sum_{i=1}^{p} \tilde{u}_{t,i} \sigma_i}{\sum_{i=1}^{n} |\tilde{u}_{t,i}|} \geq \mu \tag{15}$$

Instead of sorting the modal amplitudes, and remembering equation (13), we can choose the p *lowest frequency modes* to approximate the displacement ($\sigma = Id$, see appendix C). Other advantages of this method are : first, whatever the motion, the reduced eigenbasis is constant for a given structure ; then, only $3p$ equations instead of $3n$ of the form of equation 13 are to be solved.

7.5 Qualitative modes

In this section vectors and matrices are of order $3n$.

One can choose another suitable matrix P to perform the change of basis :

$$U = P\hat{U}$$

A modal superposition can be performed, if we set $P = [\psi_1, \ldots, \psi_{3n}]$ where it is assumed that each mode ψ_i is orthogonal to all others, but is not necessarily an eigenmode of the element :

$$U = \sum_{i=1}^{3n} \hat{u}_i \psi_i$$

The modes ψ_i can be set to be elementary displacements referred to as : translation, rotation, scaling, shearing... (see appendix D). They are "qualitative modes".

18

In general, q ($q < n$) qualitative modes are considered. Denoting $\Psi = [\psi_1, \ldots, \psi_q]$ the change of basis leads to the equation

$$\Psi^T M \Psi \ddot{\hat{U}} + \Psi^T C \Psi \dot{\hat{U}} + \Psi^T K \Psi \hat{U} = \Psi^T F \qquad (16)$$

But the change of basis is interesting when matrices $\Psi^T K \Psi$, $\Psi^T M \Psi$ and $\Psi^T C \Psi$ are diagonal, so that the q above equations decouple. Pentland's team proposes to set directly these matrices to some diagonal matrices [18, 10], since the stiffness matrix representing the nodes connections has not been calculated.

As we have calculated the stiffness matrix, we prefer to make use of the calculated node displacements vector and project it in the q-dimensional space defined by the q qualitative modes :

$$proj(U) = \sum_{i=1}^{q} \hat{u}_i \psi_i$$

The amplitudes $\hat{u}_i$ can then be easily determined, since the orthonormality of qualitative modes (with respect to scalar product) can generally be ensured [15] :

$$\hat{u}_i = proj(U).\psi_i$$

As the qualitative modes often represent global geometrical transforms, they can be assumed to be of low natural frequency, and the approximation of U by $proj(U)$ is satisfactory, mainly if the displacement norm $\|U\|$ is small.

7.6 Eigenmodes versus qualitative modes

Qualitative modes provide a good physical interpretation of non-rigid motion. They are derived from our common vocabulary, and can be quite easily computed via a polynomial

19

approximation [18]. They seem to be well dedicated to computer graphics, where certain types of deformations are to be synthetized.

On the other hand, the eigenmodes form a complete basis where the motion can be expressed accurately or approximately. They lack a physical meaning, but they can be used *without any motion assumption* ("small" displacement, linearized rotation...). Moreover, our model provides decoupled eigenmodes with respect to x, y and z, reducing the dimensions of the handled matrices from $3n \times 3n$ to $n \times n$. The eigenmodes seem to be the best modal basis for computer vision.

In figures 9 and 10, the motion of the valve with volume springs is analyzed in modal space. The eigenmodes and their superposition, as well as the qualitative modes and their superposition are displayed. Note the error in the recovery of the motion using the qualitative modes (figure 10) versus the accuracy of the superposition eigenmodes using the same number of modes (figure 9, compare to figure 5). This comes from the fact that our eigenmodes better correspond to the physical model we have developped for the motion. On the other hand, qualitative modes are expressed in a reference frame intrinsic to the object, a nice property that our own eigenmodes can share provided that we choose the intrinsic reference frame of the model (using the center of inertia and the axes of inertia).

7.7 Applications

One of the applications of modal approximation is *motion correction*. In order to maintain an object's geometry within its displacement, we can approximate the nodal displacements vector by superimposing a few number of low-frequency modes, then substitute the approximated displacement to the original one. The displacement field is smoothed and the
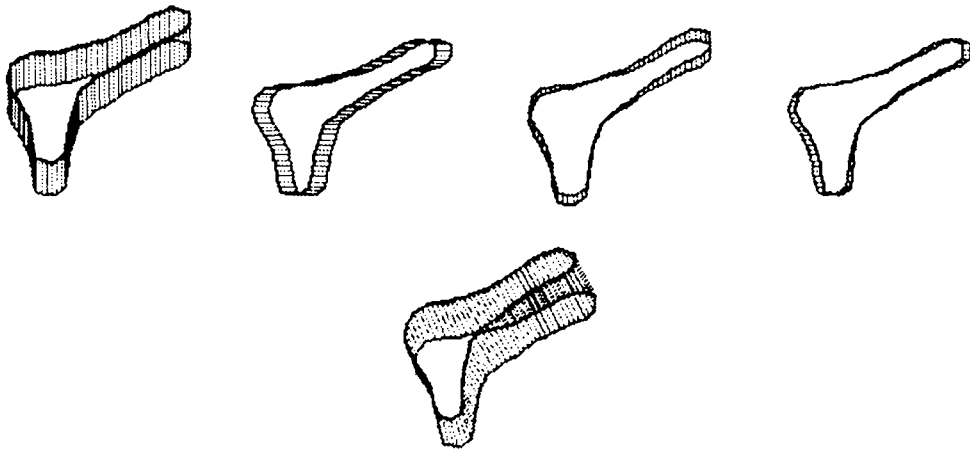
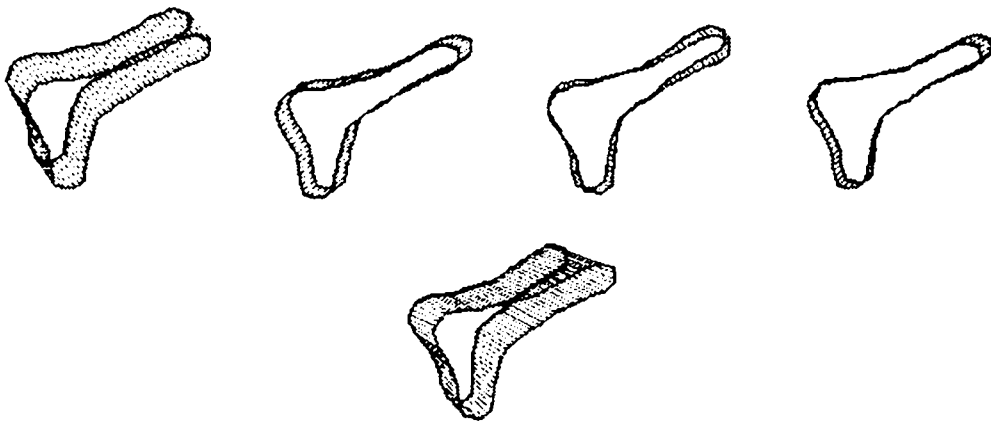Figure 9: Four high amplitude eigenmodes of the valve and their superposition

Figure 10: Four qualitative modes of the valve and their superposition

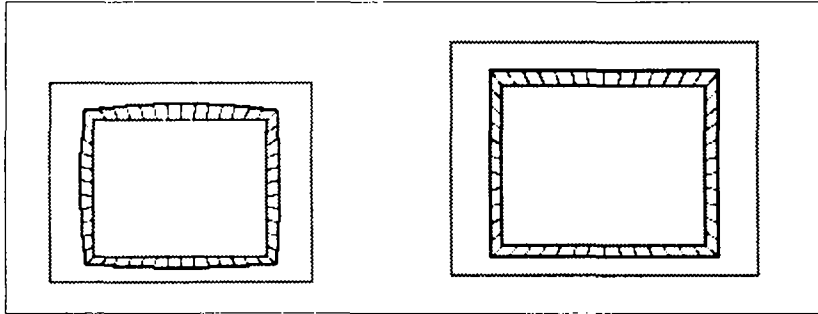deformed object has global geometric properties similar to the ones of the undeformed object (figure 11).

Figure 11: Motion correction of a rectangle using low frequency mode shapes

The reduced eigenbasis is then a good low order basis enabling straightforward *compar-isons* of different types of motions. Moreover, the representation of U in modal space is stable when the number n of structure points varies. In other terms, the modal representation is not too sensitive to the sampling of the boundary (which is always arbitrary), which is a very important property (figure 12).

Note that the mode shapes depend upon the original model ; in other words, had we chosen another way of connecting the nodal points (another K matrix), the eigenmodes would have been different.

7.8 Computational cost

We consider the general case where matrix K is only symmetric. The choice of $l_0 = 0$ which enables decoupled equations in x, y and z reduces the matrix vector multiplication complexity from $O(9n^2)$ for $3n$-order matrices to $O(3n^2)$ for 3 n-order multiplications. This means that, at each step, we reduce CPU time by a factor 3 in computing the elastic force.
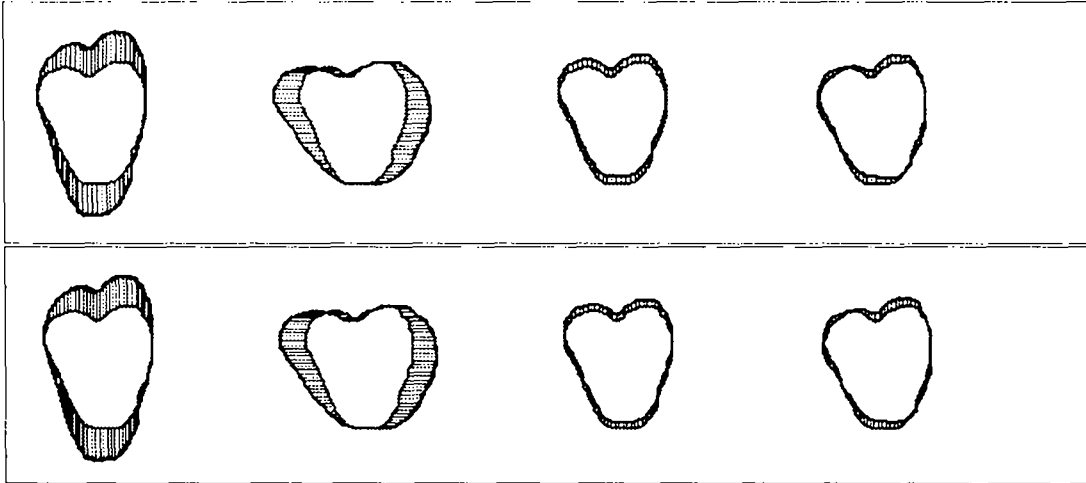
22

Figure 12: Four high amplitude eigenmodes with 80 and 95 nodal points

Then, computing eigenvalues and eigenvectors of an n order matrix is in $O(n^3)$ versus $O(27n^3)$ for a $3n$-order matrix.

On the whole, computational cost is reduced by a factor of $3 \times 27 = 81$ compared to a calculation without decoupling.

Finally, we make use of the Euler method for the resolution of our differential equations, which needs only one calculation at each step. Although the stability of this method is limited, it is less costly that a more stable method like Runge-Kutta, which can be used if necessary.

8   Experimental results

8.1   3D segmentation of human head

A mass spring mesh is used to segment a 3D magnetic resonance image of a head. The resolution of the 3D image is $158 \times 158 \times 158$.

Figures 13, 14 and 15 show the segmentation of the human head by a deformable cylindric

mesh of $159 \times 70 = 11130$ nodes. The initial mesh is given by the user as a 2D curve, that will be repeated on the 70 plans to form a cylindric mesh. The solution is displayed after 50 iterations, each iteration being performed in a few seconds CPU on a workstation, although no optimization of the code is yet done.
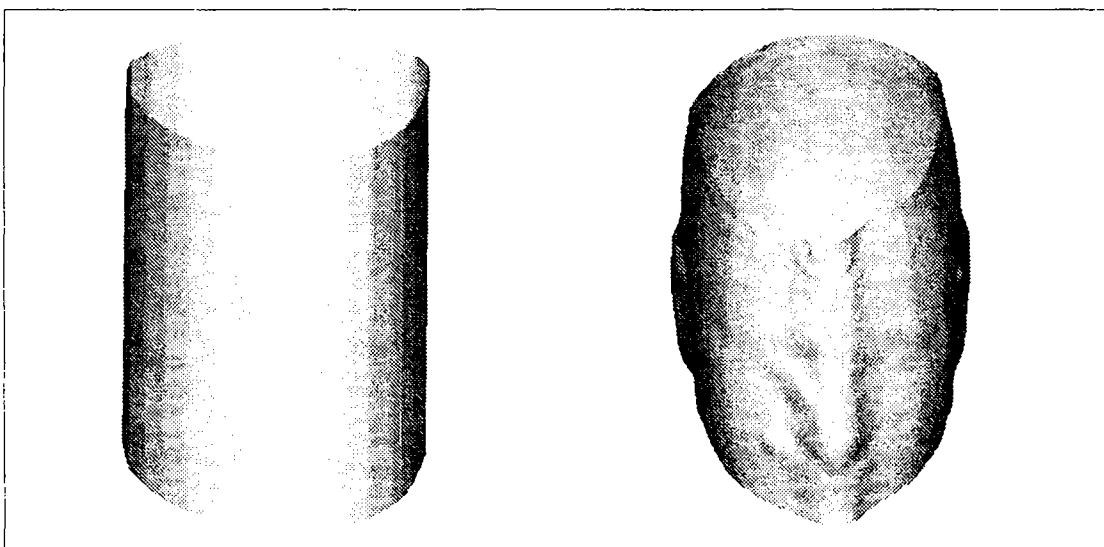
Figure 13: Initial mesh and solution after 5 iterations

8.2 2D tracking of human left ventricle

We have tested our method on a set of ultrasound images of the human heart's left ventricle. The tracking of the mitral valve is indeed a problem of major interest in medical imaging. First, a polar edge extraction is performed on the images (see [9]). Then, for each image, the distance field is computed on every pixel, using the algorithm described in [6].

The segmentation of the valve is shown in figure 16. It can be performed in one step. Then we can track the valve through time (figure 17 and 18).

Figure 14: Solution after 20 and 50 iterations

Figure 15: Segmentation of the human head (11130 nodes)
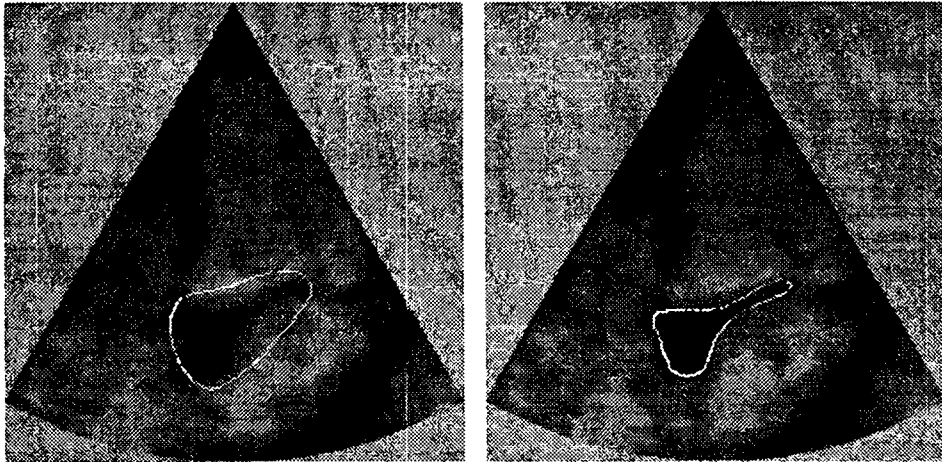
Figure 16: Initial segmentation of the valve's contour
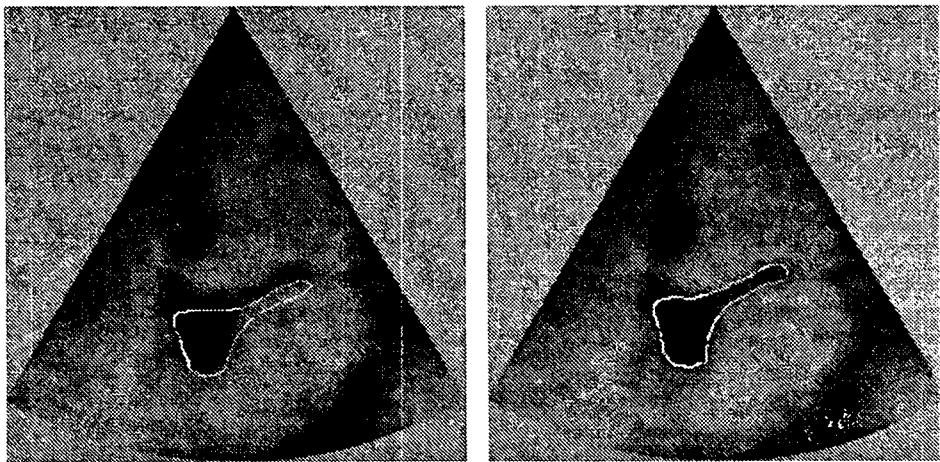
Figure 17: Tracking of the mitral valve
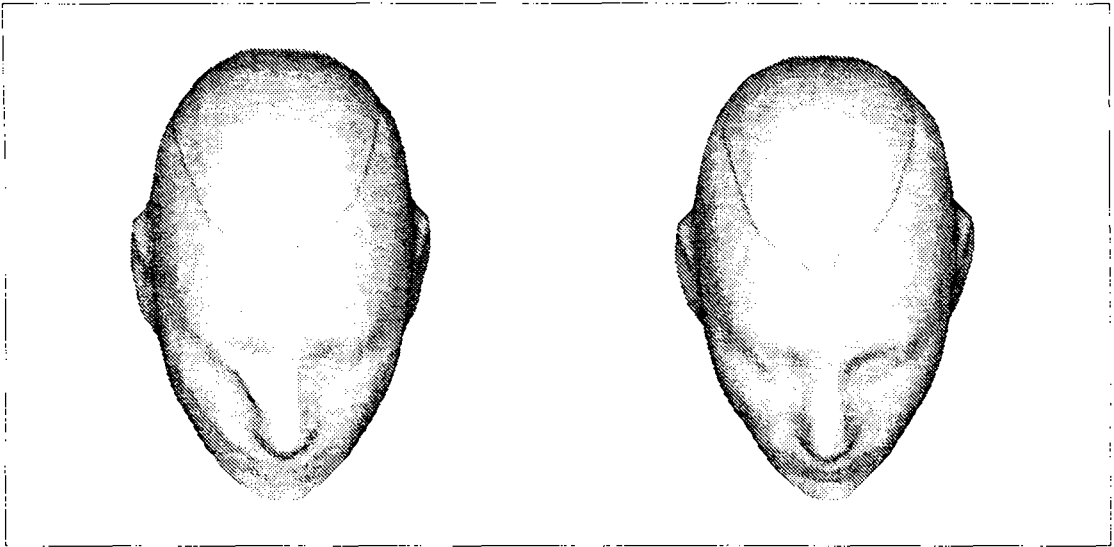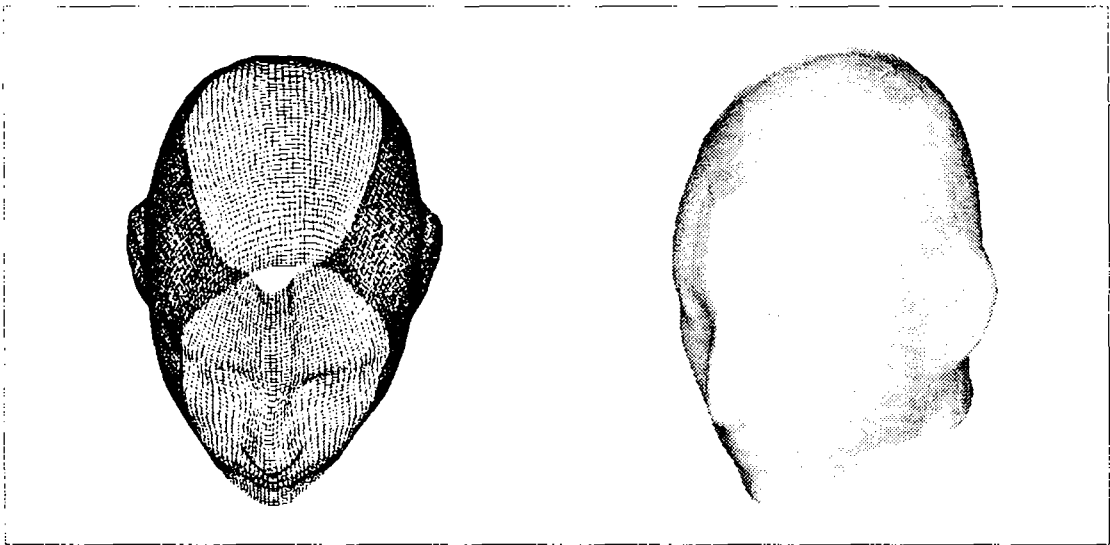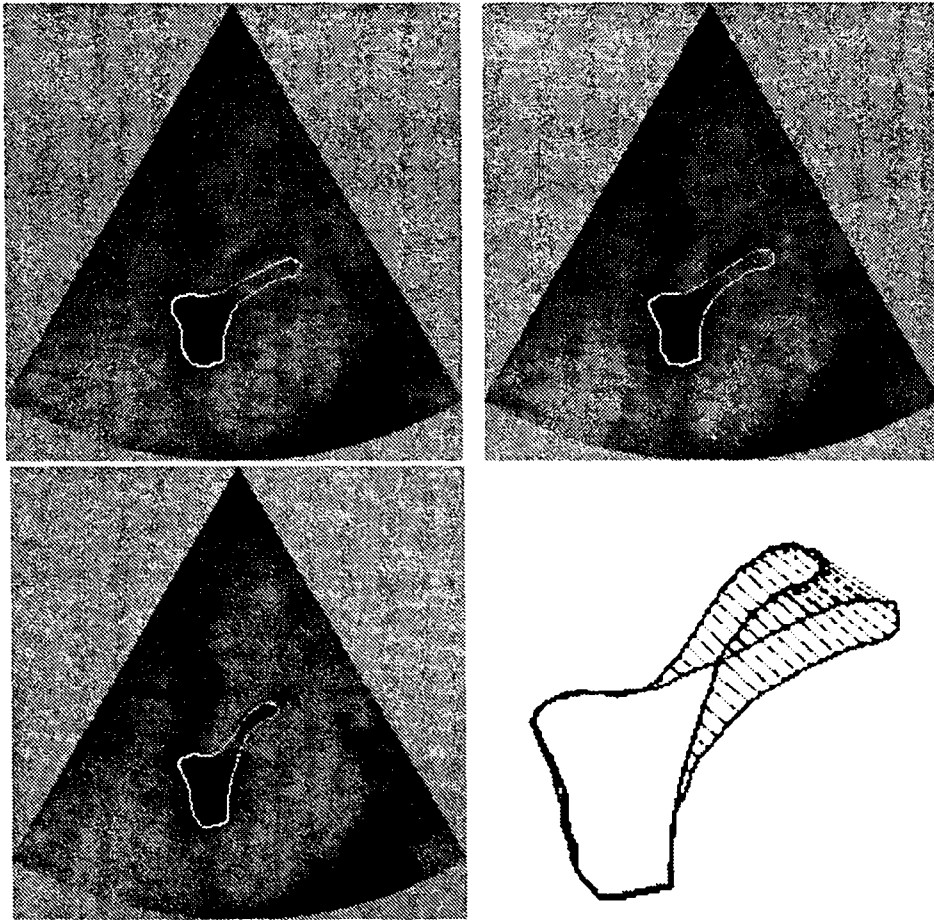
Figure 18: Tracking of the mitral valve

9   Conclusion

We have presented an elastic model enabling fast segmentation and tracking of 2D and 3D images. The high flexibility of the method makes it easy to use for modelling various types of elastic objects. The stiffness matrix remains constant at all iterations, and the computations are decoupled in x,y and z. The external potential field is selective, depending on the node being a landmark or not. Thanks to the introduction of equilibrium forces F_{eq}, it is possible to set the natural lengths of the springs to zero ; this yields a set of *linear* differential equations *decoupled* in x, y and z, even for large deformations. Therefore the algorithmic complexity is reduced by almost two orders of magnitude. Also, we showed how to take into account additional tracking constraints between irregular features. Finally, we adopted modal analysis of the motion to better fit the underlying physical model, and showed its power for a compact description and/or smoothing of a complex deformation.

Acknowledgements

Appendices

A Formulation of the stiffness matrix

For a two-dimensional structure with identical springs of stiffness k joining a mass to its neighbors, the monodimensional n order stiffness matrix is respectively for an open and for a closed structure :

$$K_{open} = \begin{bmatrix} k & -k & & & & \\ -k & 2k & -k & & & \\ & & \ddots & \ddots & \ddots & \\ & & & -k & 2k & -k \\ & & & & -k & k \end{bmatrix} \tag{17}$$

$$K_{closed} = \begin{bmatrix} 2k & -k & & & & -k \\ -k & 2k & -k & & & \\ & & \ddots & \ddots & \ddots & \\ & & & -k & 2k & -k \\ -k & & & & -k & 2k \end{bmatrix} \tag{18}$$

In the 3D case, the stiffness matrix comprehends submatrices that have the form described above.

We can model the structure in a more complex way, for instance by attaching extra springs of stiffness k_{extra} between non-neighbor nodes i and j. The stiffness matrix is then modified as follows :

$$(K_{ii})_{extra} = K_{ii} + k_{extra}$$

$$(K_{jj})_{extra} = K_{jj} + k_{extra}$$

$$(K_{ij})_{extra} = K_{ij} - k_{extra}$$

$$(K_{ji})_{extra} = K_{ji} - k_{extra}$$

Note that the stiffness matrices are non-definite positive ($KT=0$, with $T = [1, 1, \ldots, 1]^T$).

B   Orthogonality of mode shapes

An important property of mode shapes is orthogonality, which is discussed here.

For the ith and jth natural frequencies ω_i and ω_j and the ith and jth mode shapes ϕ_i and ϕ_j, we have :

$$K\phi_i = \omega_i^2 M\phi_i \tag{19}$$

$$K\phi_j = \omega_j^2 M\phi_j \tag{20}$$

Premultiplying equation (19) by ϕ_j^T and equation (20) by ϕ_i^T leads to :

$$\phi_j^T K\phi_i = \omega_i^2 \phi_j^T M\phi_i \tag{21}$$

$$\phi_i^T K\phi_j = \omega_j^2 \phi_i^T M\phi_j \tag{22}$$

Taking the transpose of equation (22), and remembering that mass and stiffness matrices are symmetric, we obtain :

$$\phi_j^T K\phi_i = \omega_j^2 \phi_j^T M\phi_i \tag{23}$$

Comparing now equation (21) and equation (23) yields :

$$(\omega_i^2 - \omega_j^2)\phi_j^T M\phi_i = 0 \tag{24}$$

Since $\omega_i^2 \neq \omega_j^2$, we have :

$$\phi_j^T M\phi_i = 0 \qquad \forall i \neq j$$

M being a positive definite matrix, we can norm the mode shapes to obtain :

$$\phi_i^T M \phi_i = 1 \qquad \forall i$$

The mode shapes are then M-orthonormal.

Using M-orthonormality in equation (21) leads to the K-orthogonality of mode shapes :

$$\phi_j^T K \phi_i = 0; \qquad \forall i \neq j$$

$$\phi_i^T K \phi_i = \omega_i^2; \qquad \forall i$$

C   Selecting the low natural frequencies

Using an easy example, we discuss here the choice of mode shapes that have the lowest natural frequencies for the approximation of the displacement vector. For the detailed equations, see [2].

Consider a clamped beam. Figure 19 shows the aspect of a low order and a high order mode.
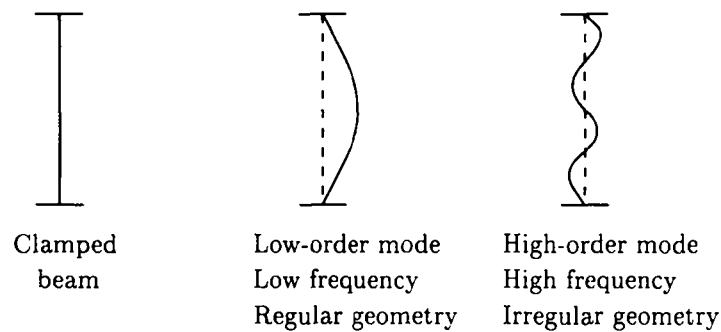

Clamped   Low-order mode   High-order mode
beam   Low frequency   High frequency
               Regular geometry   Irregular geometry

Figure 19: a clamped beam

C.1 Frequential concordance

Consider a concentrated load in the middle of the beam. If the load varies slowly, the natural high frequencies are not amplified. This is illustrated in figure 20.
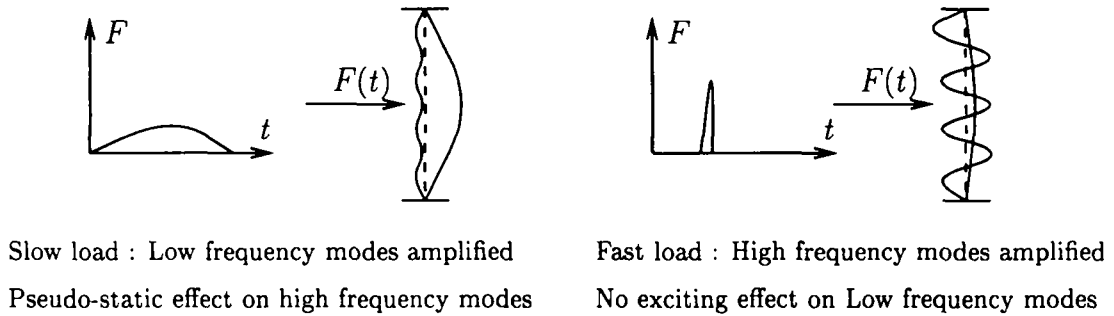
Slow load : Low frequency modes amplified
Pseudo-static effect on high frequency modes

Fast load : High frequency modes amplified
No exciting effect on Low frequency modes

Figure 20: Frequential concordance

C.2 Geometric participation

If the load is equally distributed along the beam, the natural high frequencies, whose geometry is irregular, are not excited. The amount of work for force F (that is $\Phi^T F$ for displacement Φ) is negligible for these frequencies (figure 21).

Equally-distributed
load

Low-frequency mode
Important work for F

High-frequency mode
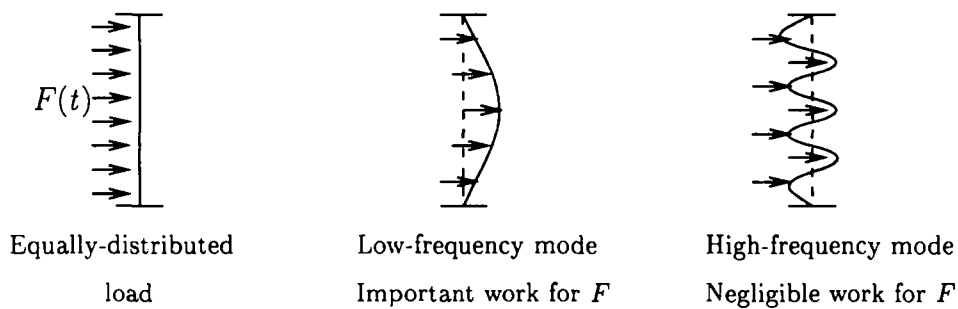Negligible work for F

Figure 21: Geometric participation

D Qualitative modes

As described before, the qualitative modes ψ_i of an object are generally not the solution of the eigenproblem $K\phi = \omega^2 M\phi$.

Their main advantage is that they can be described by common terms, such as translation, rotation (rigid-body modes), scaling, shearing (first order modes), tapering, pinching, bending (second-order modes)...

To intoduce the qualitative modes, consider the general non-rigid motion of an object. It can be represented by :

$$\begin{cases} M(x, y, z) \rightarrow M'(x', y', z') \\ M' = rd\ M + t \end{cases}$$

Where r is a rotation matrix of angles $\theta_x, \theta_y, \theta_z$, t is a translation vector ($t = t_x + t_y + t_z$), and d is a symmetric deformation matrix.

Suppose now that *the displacement is "small"*. Then rotation can be linearized ($|\theta| \ll \pi$) :

$$r = I + \rho = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} + \begin{bmatrix} 0 & -\theta_z & -\theta_y \\ \theta_z & 0 & -\theta_x \\ \theta_y & \theta_x & 0 \end{bmatrix}$$

And deformation can be decomposed into :

$$d = I + \epsilon$$

where ϵ is the symmetric strain matrix.

A first order approximation of the motion leads to :

$$\begin{cases} M(x, y, z) \rightarrow M'(x', y', z') \\ M' - M = (\rho + \epsilon)\ M + t \end{cases}$$

More precisely, in $3D$, we have :

$$\begin{cases} x' - x = t_x - \theta_z y - \theta_y z + \epsilon_{xx} x + \epsilon_{xy} y + \epsilon_{xz} z \\[2mm] y' - y = t_y + \theta_z x - \theta_x z + \epsilon_{xy} x + \epsilon_{yy} y + \epsilon_{yz} z \\[2mm] z' - z = t_z + \theta_y x + \theta_x y + \epsilon_{xz} x + \epsilon_{yz} y + \epsilon_{zz} z \end{cases}$$

Suppose the deformation is *homogeneous*, that is, matrix ϵ components does not depend on point positions. The displacement of the points on the object is then controlled by 12 parameters : $t_x, t_y, t_z, \theta_x, \theta_y, \theta_z, \epsilon_{xx}, \epsilon_{yy}, \epsilon_{zz}, \epsilon_{xy}, \epsilon_{xz}, \epsilon_{yz}$.

Now if the object points before deformation are positionned at points $(x_i, y_i, z_i), i = 1, \ldots, n$, we can derive the twelve low-order modes :

$$\psi_1^T = [1, 0, 0, 1, 0, 0, \ldots, 1, 0, 0] \qquad x \text{ translation}$$

$$\psi_2^T = [0, 1, 0, 0, 1, 0, \ldots, 0, 1, 0] \qquad y \text{ translation}$$

$$\psi_3^T = [0, 0, 1, 0, 0, 1, \ldots, 0, 0, 1] \qquad z \text{ translation}$$

$$\psi_4^T = [-y_1, x_1, 0, -y_2, x_2, 0, \ldots, -y_n, x_n, 0] \qquad z \text{ rotation}$$

$$\psi_5^T = [-z_1, 0, x_1, -z_2, 0, x_2, \ldots, -z_n, 0, x_n] \qquad y \text{ rotation}$$

$$\psi_6^T = [0, -z_1, y_1, 0, -z_2, y_2, \ldots, 0, -z_n, y_n] \qquad x \text{ rotation}$$

$$\psi_7^T = [x_1, 0, 0, x_2, 0, 0, \ldots, x_n, 0, 0] \qquad x \text{ scaling}$$

$$\psi_8^T = [0, y_1, 0, 0, y_2, 0, \ldots, 0, y_n, 0] \qquad y \text{ scaling}$$

$$\psi_9^T = [0, 0, z_1, 0, 0, z_2, \ldots, 0, 0, z_n] \qquad z \text{ scaling}$$

$$\psi_{10}^T = [y_1, x_1, 0, y_2, x_2, 0, \ldots, y_n, x_n, 0] \qquad z \text{ shearing}$$

$$\psi_{11}^T = [z_1, 0, x_1, z_2, 0, x_2, \ldots, z_n, 0, x_n] \qquad y \text{ shearing}$$

$$\psi_{12}^T = [0, z_1, y_1, 0, z_2, y_2, \ldots, 0, z_n, y_n] \qquad x \text{ shearing}$$

Thus, in 3D, assuming that the displacement is small and the strain matrix is homogeneous, only twelve modes are necessary for parameterizing the displacement. Note that these as-

sumptions are restrictive, and more polynomial modes have generally to be considered [16],

and their orthonormality must be ensured.

References

[1] A. Amini, R. Owen, L. Staib, P. Anandan, and J. Duncan. *Non-rigid motion models for tracking the left ventricular wall.* Lecture notes in computer science: Information processing in medical images. Springer-Verlag, 1991.

[2] Klaus-Jurgen Bathe. *Finite Element Procedures in Engineering Analysis.* Prentice-Hall, 1982.

[3] Fred L. Bookstein. Principal warps: Thin-plate splines and the decomposition of deformations. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, PAMI-11(6):567–585, June 1989.

[4] Isaac Cohen, Nicholas Ayache, and Patrick Sulger. Tracking points on deformables objects. In *Proceedings of the Second European Conference on Computer Vision (ECCV) 1992*, Santa Margherita Ligure, Italy, May 1992.

[5] Isaac Cohen, Laurent D. Cohen, and Nicholas Ayache. Introducing deformable surfaces to segment 3−D images and infer differential structures. In *Proc. Computer Vision and Pattern Recognition*, Hawai,USA, June 1991.

[6] P. E. Danielsson. Euclidean distance mapping. *Computer Vision, Graphics, and Image Processing*, 14:227–248, 1980.

[7] J.S. Duncan, R.L. Owen, L.H. Staib, and P. Anandan. Measurement of non-rigid motion using contour shape descriptors. In *Proc. Computer Vision and Pattern Recognition*, pages 318–324. IEEE Computer Society Conference, June 1991. Lahaina, Maui, Hawaii.

[8] André Gueziec. Large deformable splines : Crest lines and matching. Technical Report 1782, INRIA, October 1993.

[9] I.L. Herlin and N. Ayache. Features extraction and analysis methods for sequences of ultrasound images. In *Proceedings of the Second European Conference on Computer Vision (ECCV) 1992*, Santa Margherita Ligure, Italy, May 1992.

[10] Bradley Horowitz and Alex Pentland. Recovery of non-rigid motion and structure. In *Proc. Computer Vision and Pattern Recognition*, pages 325–330. IEEE Computer Society Conference, June 1991. Lahaina, Maui, Hawaii.

[11] Michael Kass, Andrew Witkin, and Demetri Terzopoulos. Snakes: Active contour models. In *Proceedings of the First International Conference on Computer Vision*, pages 259–268, London, June 1987.

[12] Dimitri Metaxas and Demetri Terzopoulos. Constrained deformable superquadrics and nonrigid motion tracking. In *Proc. Computer Vision and Pattern Recognition*, pages 337–343. IEEE Computer Society Conference, June 1991. Lahaina, Maui, Hawaii.

[13] Sanjoy K. Mishra, Dmitry B. Goldgof, and Thomas S. Huang. Motion analysis and epicardial deformation estimation from angiography data. In *Proc. Computer Vision and Pattern Recognition*, pages 331–336. IEEE Computer Society Conference, June 1991. Lahaina, Maui, Hawaii.

[14] Olivier Monga, Serge Benayoun, and Olivier D. Faugeras. From partial derivatives of 3d density images to ridge lines. In *Proc. Computer Vision and Pattern Recognition*, pages 354–359. IEEE Computer Society Conference, June 1992. Champaign, Illinois.

[15] G.N Pande, G. Beer, and J.R. Williams. *Numerical Methods in Rock Mechanics.* John Wiley and Sons, 1990.

[16] A. Pentland and Williams J. Good vibrations : Modal dynamics for graphics and animation. In *Computer Graphics,* 1989.

[17] A. Pentland and Williams J. The perception of non-rigid motion : Inference of material properties and force. In *Proceedings of the International Joint Conference on Artificial Intelligence,* August 1989.

[18] Alex Pentland and Stan Sclaroff. Closed-form solutions for physically based shape modelling and recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence,* PAMI-13(7):715–729, July 1991.

[19] Lawrence H. Staib and James S. Duncan. Deformable fourier models for surface finding in 3d images. In *Proceedings of Visualization in Biomedical Computing,* Chapell Hill,USA, October 1992.

[20] Demetri Terzopoulos and Keith Waters. Physically-based facial modelling, analysis, and animation. *The Journal of Visualization and Computer Animation,* 1:73–80, 1990.

[21] Jean-Philippe Thirion and Alexis Gourdon. The 3d marching lines algorithm and its application to crest lines extraction. Technical Report 1672, INRIA, May 1992.

[22] Demitri Terzopoulos Andrew Witkin and Michael Kass. Symmetry-seeking models for 3-D object reconstruction. In *Proceedings of the first International Conference on Computer Vision,* pages 269–276, June 1987.