

Frequency-Based Nonrigid Motion Analysis: Application to Four Dimensional Medical Images

Chahab Nastar and Nicholas Ayache

Abstract—We present a method for nonrigid motion analysis in time sequences of volume images (4D data). In this method, nonrigid motion of the deforming object contour is dynamically approximated by a physically-based deformable surface. In order to reduce the number of parameters describing the deformation, we make use of a modal analysis which provides a spatial smoothing of the surface. The deformation spectrum, which outlines the main excited modes, can be efficiently used for deformation comparison. Fourier analysis on time signals of the main deformation spectrum components provides a temporal smoothing of the data. Thus a complex nonrigid deformation is described by only a few parameters: the main excited modes and the main Fourier harmonics. Therefore, 4D data can be analyzed in a very concise manner. The power and robustness of the approach is illustrated by various results on medical data. We believe that our method has important applications in automatic diagnosis of heart diseases and in motion compression.

Index Terms—Medical image analysis, nonrigid motion, deformable models, modal analysis, Fourier analysis, compression, dynamic data, four-dimensional images, cardiac imagery, automatic diagnosis.

1 INTRODUCTION

1.1 Motivation and Organization

IN this article, we propose a unified approach for nonrigid motion estimation from time sequences of three-dimensional images [5], i.e., 4D data, by taking into account both spatial and temporal frequencies of a deformable geometric model.

Our method has important applications in automatic diagnosis of heart diseases and in 4D data compression, as we shall see in the experimental section. Our method involves three steps:

- 1) Recover the deformation field between each pair of successive 3D images,
- 2) Express the modal coefficients (or amplitudes) of the deformation at each time t ,
- 3) Express Fourier coefficients of the time-varying modal amplitudes.

The information contained in the data can then be compressed, by discarding spatial modes and/or temporal Fourier harmonics.

The basic justification of the approach relies on the following observations: most smoothly deforming structures mainly have low-frequency excited modes; this justifies Step 2. Furthermore, for periodic motions like heart motion, modal amplitudes as a function of time are periodic and sine-like;¹ this justifies Step 3.

For Steps 1 and 2, it is possible to constrain the modes beforehand, by letting a deformable model evolve along low-order modes. The advantage is, of course, a reduced numerical complexity and a low-order smoothing of the deformation allowing its robust recovery. Note that Step 1 can be performed not only by our deformable model, but also by optical flow or any other technique providing the motion field [10].

The following example can serve as a motivation to our work. Let us consider (Fig. 1) the canine 4D heart data provided by the dynamic spatial reconstructor (DSR, a high speed X-ray CT scanner). We used this data as an input to our method. It consists of 18 volume (or 3D) images during a single cardiac cycle. Each volume image has a spatial resolution of $98 \times 100 \times 110$. Dye was injected into the left ventricle, which shows up as a light gray color. Note the difficulty of interpretation of such complex and huge data by physicians. The challenge was to analyze such a tremendous amount of information (19,404,000 bytes, with each voxel being coded on one byte) and supply physicians with a few quantitative parameters describing the motion. We shall see in the experimental section of this paper how we tried to meet these objectives.

The article is organized as follows:

In Section 2, we introduce our deformable model for nonrigid motion estimation, develop its governing equations and outline its properties.

In Section 3, we present the modal analysis of the model, allowing a closed-form recovery of the motion by a few parameters: the modal amplitudes. We show how these amplitudes may be used to characterize deformations. Section 3.2, which describes the *analytic modes*, may be required only for in depth reading.

• The authors are with INRIA, B.P. 105, 78153 Le Chesnay Cédex, France.
E-mail: {chahab.nastar; nicholas.ayache}@inria.fr.

Manuscript received Mar. 14, 1995; revised July 25, 1995.

Recommended for acceptance by A. Singh.

For information on obtaining reprints of this article, please send e-mail to: transpami@computer.org, and reference IEEECS Log Number P96085.

1. Note that our method is not restricted to periodic motions. For nonperiodic motions there is generally no need in performing the Fourier analysis (step 3).




Fig. 1. 4D data displaying a 3D dog ventricle in motion. Each 2D image is an axial slice in the $x y$ plane. The z coordinate is along the vertical axis, while the t coordinate is along the horizontal axis.

In Section 4, we propose a Fourier analysis for the time-varying main modal amplitudes; this is particularly well suited for cyclic motions (e.g., heart motion). An example involving 4D medical data is presented.

We conclude and propose future work in Section 5.

1.2 Related Work

A lot of research work is related to our work. First of all, we want to mention the pioneering work of Terzopoulos et al. on snakes and deformable physics-based models [24], [26], [42], [45]. This work was mainly dedicated to tracking and animation problems rather than analysis.

A second category of research work is related to the analysis of deformations. Our work was initially inspired by the pioneering work of Pentland et al. on the use of modal deformations to describe deformations [36], [37]. The main differences with our approach is that our modes have a simpler physical meaning (they can be interpreted as the harmonics of the free vibrations basis of the deformable object) while Pentland's modes are more dedicated to animation [38]. Also, we have derived a method to compute our modes beforehand, as a function of the topology of the deformable object only, allowing a much faster computation (Pentland's modes require the knowledge of the exact shape of the original object and, therefore, have to be computed online).

Our work can be seen as providing a coherent unifying framework between these first two categories of publications, allowing to both track and analyze a deformable 3D motion.

In addition to these references, we have to mention the work of Cootes et al. [15], [16], [17], done independently of, and concurrently with, ours which provides an interesting alternative to the solution of both the tracking and the analysis of the deformation. We also have to mention the work of Bardinet [7] and Metaxas [27], [35], among others, who use deformable parametric models like superquadrics to track and analyze deformable objects. Their work is limited to the class of objects which can be easily represented by such parametric models, which include the left cardiac ventricle. A number of methods to track 3D data were proposed in the past few years. Duncan et al. [2], [39], [40], and Benayoun et al. [9], [10], used differential constraints; Goldgof [23] and Chen [13] used a coarse-to-fine approach; Bookstein [12] and Szeliski and Lavalley [41] proposed energy based methods; Creswell [18] and Amini [1] suggested data-dedicated methods. All these approaches can be seen as complementary to ours, in the sense that the tracking result they provide could be completed by our proposed analysis. It would however be interesting to compare all these approaches on a specific and common database. This could be the topic of a future work.

2 A DEFORMABLE MODEL FOR MOTION ESTIMATION

In this section, we introduce our physically based deformable model which we use for tracking nonrigid motion of dynamic structures in time sequences of 2D or 3D medical images.

We consider both the surface and volume properties of the objects at hand. We restrict ourselves to elastic deformations, i.e., we assume that the object recovers its reference configuration as soon as all applied forces causing deformation are removed. In general, we seek a trade off between precise modeling and computational efficiency. Therefore, simplifying assumptions will be introduced in the modeling.

2.1 Mass-Spring Meshes

Modeling an elastic boundary can be achieved by a mesh of N virtual masses on the contour. Each mass is attached to its neighbors by perfect identical springs of stiffness K and natural length l_0 [23], [46].

These springs achieve a *polygonal approximation* and model the *surface properties* of the object.

Generalizing the model to 3D contours (surfaces) is straightforward: We can either model “quadrilateral” or “diagonal” meshes² (see Fig. 2).




Fig. 2. “Diagonal” mesh.

If necessary, we can improve the modeling by attaching extra springs between non-neighbor nodes in order to model *volume* properties inside the object. These springs constrain the general shape of the object within its deformation, making a template shape out of it [11], [48].

The boundary modeled as above will also be called a *structure*. Such a structure can be easily deformed to match the contour of an object of interest, thus performing a *segmentation* step. If we take a series of images displaying the deformation of the object through time, the structure can achieve simultaneously both *segmentation* and *tracking* of the object surface through time.

2.2 Governing Equations

The system under study is made of the N virtual masses located at time t at points $(M_1(t), M_2(t), \dots, M_N(t))$.

The fundamental equation of dynamics states that the vector addition of all applied forces on M_i is equal to its mass m_i multiplied by its acceleration. Let A be the origin of the reference system. Then:

$$\sum_j \mathbf{f}_j(M_i, t) = m_i \frac{d^2}{dt^2} \overrightarrow{AM_i} = m_i \ddot{\mathbf{M}}_i \quad (1)$$

Let us now describe the applied forces on each node M_i :

- Elastic force between M_i and its connected nodes.

$$\mathbf{f}_e(M_i, t) = -K \left\langle \sum_{j \in C_i} \mathbf{U}_{ij}(t) - l_0 \sum_{j \in C_i} \frac{\mathbf{U}_{ij}(t)}{\|\mathbf{U}_{ij}(t)\|} \right\rangle \quad (2)$$

2. The shear-resisting cross springs may be useful to avoid self-intersection problems during the deformation of the surface.

where K is the stiffness constant, $\mathbf{U}_{ij}(t) = \overrightarrow{M_j M_i}(t)$ is the vector separation of nodes M_j and M_i at time t , and C_i is the set of nodes connected to node M_i .

- Fluid damping force, proportional to node velocity:

$$\mathbf{f}_d(M_i, t) = -c_i \frac{d}{dt} \overrightarrow{AM_i} = -c_i \dot{\mathbf{M}}_i \quad (3)$$

where c_i is the damping constant.

- Image force $\mathbf{f}_{im}(M_i, t)$. This is the main external force that has the system attracted by image features. This force is defined in Section 2.3.
- Suppose the natural length l_0 of the springs is fixed. Since we wish to give the system an initial equilibrium configuration, we need to apply on each node a force balancing the action of the elastic force:

$$\mathbf{f}_{eq}(M_i) = -\mathbf{f}_e(M_i, t_0) \quad (4)$$

This force is similar to the force that our fingers apply to an elastic rubber to keep it in a specific shape. We assume that this force is constant over time.

Finally, (1) yields the governing equation:

$$\mathbf{f}_e(M_i, t) + \mathbf{f}_d(M_i, t) + \mathbf{f}_{im}(M_i, t) + \mathbf{f}_{eq}(M_i) = m_i \ddot{\mathbf{M}}_i \quad (5)$$

The governing 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 appears clearly in (2)).

One possible approach is the resolution of these complex equations by an iterative procedure [43]. In this paper, we propose to set $l_0 = 0$. This assumption does not restrict the arbitrary initial configuration of the structure because of the equilibrium force \mathbf{f}_{eq} . Indeed, this force keeps the structure inflated so that it does not shrink to a point. Thus, the natural state of the system is its initial configuration.

The advantage of this assumption is that our model can be considered within the framework of *linear elasticity*. As a consequence, we end up with a set of *linear* differential equations with node displacements *decoupled* in each coordinate, regardless of the magnitude of the displacements. Moreover, these linear equations are a prerequisite to further quantitative analysis of the motion (see Section 3).

On the other hand, our approximation is valid only if the spring orientations undergo small angular variations (typically less than 15 degrees), so that our assumption of constant equilibrium force \mathbf{f}_{eq} holds.³

Finally, in 3D, the deformation of the system is governed by the $3N$ -dimensional differential matrix equation:

$$\mathbf{M}\ddot{\mathbf{U}} + \mathbf{C}\dot{\mathbf{U}} + \mathbf{K}\mathbf{U} = \mathbf{F}(t) \quad (6)$$

where \mathbf{U} is a vector storing nodal displacements \mathbf{M} , \mathbf{C} , and \mathbf{K} are, respectively, the mass, damping, and stiffness matrices of the system, and \mathbf{F} is the image force which has the object attracted by image edges. Equation (6) is the finite element formulation of the deformation process. Note that the equilibrium forces do not explicitly appear in the governing equation.

3. Similar limitations can be found in the model described in [36], [37].

2.3 Image Force

In the original formulation of the "snake," the authors propose a convolution of the image with a smoothing filter that causes artificial blurring, so that the active contour can be attracted by the edges from a distance [24].

Unlike these methods, we introduce a force at each node M_i that points to the closest boundary point P_i in the image [14], [20], [22], [25]. Several Euclidean distance algorithms can help us extract this force in each voxel of the image [19], [47]. At node M , this force is set to:

$$\mathbf{f}_{im}(M, t) = \alpha \overrightarrow{M(t)P(t)} \quad (7)$$

where α is a constant scalar. This force can be seen as a virtual spring of natural length zero and of stiffness α joining M to P . Hence, both internal and external forces are elastic, making the modeling coherent.

The advantage is that we speed up the convergence of the model toward image edges; one can consider edge extraction and distance computation as a data-to-force transformation which is performed as a preprocessing.

2.4 Integration Scheme

In 3D, the $3N$ -order matrix equations decouple into three directional matrix equations of order N :

$$\begin{cases} \mathbf{M}\ddot{\mathbf{U}}_x + \mathbf{C}\dot{\mathbf{U}}_x + \mathbf{K}\mathbf{U}_x = \mathbf{F}_x(t) \\ \mathbf{M}\ddot{\mathbf{U}}_y + \mathbf{C}\dot{\mathbf{U}}_y + \mathbf{K}\mathbf{U}_y = \mathbf{F}_y(t) \\ \mathbf{M}\ddot{\mathbf{U}}_z + \mathbf{C}\dot{\mathbf{U}}_z + \mathbf{K}\mathbf{U}_z = \mathbf{F}_z(t) \end{cases} \quad (8)$$

where \mathbf{M} , \mathbf{C} , and \mathbf{K} are from now on N -order matrices, and \mathbf{U}_μ and \mathbf{F}_μ ($\mu = x, y, z$) are, respectively, the N -order displacement field and image force in the μ direction. From now on, we will omit the indexes, and all matrix equations will be of order N . It is assumed that three equations corresponding the three space directions have to be solved.

We integrate the governing equations with an explicit Euler scheme:

$$\begin{cases} \ddot{\mathbf{U}}(t) = \mathbf{M}^{-1}(\mathbf{F}(t) - \mathbf{C}\dot{\mathbf{U}}(t) - \mathbf{K}\mathbf{U}(t)) \\ \dot{\mathbf{U}}(t + \Delta t) = \dot{\mathbf{U}}(t) + \Delta t \ddot{\mathbf{U}}(t) \\ \mathbf{U}(t + \Delta t) = \mathbf{U}(t) + \Delta t \dot{\mathbf{U}}(t + \Delta t) \end{cases} \quad (9)$$

where Δt is the time step of the simulation. The initial values of displacement and velocity are generally set to zero. Note that the mass matrix is diagonal and therefore its inversion is trivial and only has to be processed once. The damping constant is chosen so that the system is slightly overdamped [44]. Convergence is achieved once the image force balances the internal elastic forces.

In any numerical integration scheme, numerical stability requires the time step Δt to be inferior to a critical value Δt_{cr} which is defined by the mass and stiffness properties of the system:

$$\Delta t \leq \Delta t_{cr} = \frac{T_N}{\gamma} \quad (10)$$

where $T_N = \frac{2\pi}{\omega_N}$ is the smallest period of the finite element mesh, and γ is a constant depending on the scheme. The Euler method is conditionally stable, i.e., needs small time steps ($\gamma \geq 2$). In order to make a more accurate computation, we decrease the time step and set $\gamma \geq 10$. Note that a

larger time step can be used when high-frequency components of the deformable model are discarded (see Section 3).

2.5 Results on Medical Data

We have tested our method on a set of ultrasound images of the left ventricle of a human heart. The tracking of the mitral valve is indeed a problem of major interest in medical imaging since heart-attacks can generally be predicted from abnormal motion of the valve. Each image has a resolution of 256×256 pixels. First, a polar edge extraction is performed on the images [6]. Then, for each image, the distance field is computed on every pixel. Thus we can segment the initial valve. We then add volume springs for making our model a template; this will make the tracking more robust. Finally we track the valve through time and display the estimated displacement field on the valve surface (Fig. 3). The program runs in *real-time* on 2D data.




Fig. 3. Initial segmentation and tracking of the mitral valve.




Fig. 4. Convergence of an initial sphere toward the human head.

The method has similarly been tested on a 3D Magnetic Resonance image of the human head, with a resolution of $158 \times 158 \times 158$ voxels. We make use of the 3D generalization of Canny's edge detector [28], [29], then compute the 3D distance maps and finally run the program. The convergence of an initial sphere toward the human head is shown in Fig. 4. It runs at *interactive rates* on a DEC-ALPHA workstation; convergence is achieved in eight seconds. Note that this segmentation example does not really make use of the power of our dynamic approach; it is provided here to demonstrate model properties.

3 MODAL ANALYSIS AND APPLICATIONS

In this section we focus on a quantitative analysis of the model deformation, using a frequency-based technique called modal analysis.

Modal analysis is a standard engineering technique allowing more effective computations and a closed-form solution of the deformation process [8]. It was first introduced in computer vision by Pentland's team [36], [37]. Let us now explain this technique.

3.1 General Approach

Instead of solving directly the equilibrium equation (6), one can transform it by a change of basis:

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

where \mathbf{P} is the square nonsingular transformation matrix of order N to be determined, and $\tilde{\mathbf{U}}$ is referred to as the *generalized displacements* vector. One effective way of choosing \mathbf{P} is setting it to Φ , a matrix whose entries are the eigenvectors of the generalized eigenproblem:

$$\mathbf{K}\phi = \omega^2 \mathbf{M}\phi \quad (12)$$

$$\mathbf{U}(t) = \Phi \tilde{\mathbf{U}} = \sum_{i=1}^N \tilde{u}_i(t) \phi_i \quad (13)$$

Equation (13) is referred to as the modal superposition equation. ϕ_i is the i th mode, \tilde{u}_i its *amplitude*, and ω_i its *frequency*. Fig. 5 displays a sample of six frequency-increasing modes of a cylinder, all of them having the same amplitude.




Fig. 5. A sample of frequency-increasing modes of a cylinder (constant amplitude).

Note that the modes of a generalized eigenproblem involving real symmetric matrices can be chosen to be orthonormal vectors [8]. The new modal basis simultaneously diagonalizes \mathbf{M} and \mathbf{K} , and provided that matrix $\tilde{\mathbf{C}} = \Phi^T \mathbf{C} \Phi$ is diagonal as well,⁴ the governing matrix-form equations decouple into N scalar equations:

$$\ddot{\tilde{u}}_i(t) + \tilde{c}_i \dot{\tilde{u}}_i(t) + \omega_i^2 \tilde{u}_i(t) = \tilde{f}_i(t) \quad (14)$$

The amplitudes $(\tilde{u}_i(t))_{i=1, \dots, N'}$ are obtained by solving these equations at time t , and the displacement of the structure nodes is obtained by the modal superposition equation.

In practice, we wish to approximate nodal displacements $\mathbf{U}(t)$ by $\hat{\mathbf{U}}(t)$, the truncated addition of the p low-frequency modes, where $p \ll N$.

$$\hat{\mathbf{U}}(t) \approx \sum_{i=1}^p \tilde{u}_i(t) \phi_i \quad (15)$$

Vectors $(\phi)_{i=1, \dots, p}$ form the *reduced modal basis* of the system. This is the major advantage of modal analysis: It allows a closed-form solution by selecting a few number of low-frequency modes [8], [31], [32]. Therefore a compact description of the motion is provided by *spatial smoothing*. Fig. 6 illustrates modal superposition on the real example of the mitral valve.

4. This condition, called the Rayleigh condition, is satisfied as soon as the damping matrix \mathbf{C} is a linear combination of the mass and stiffness matrices. Rayleigh damping is generally assumed for any standard engineering problem [8].




Fig. 6. Nine low-order amplitude-moderated eigenmodes describing the motion of the valve; their superposition (bottom-right) is a good approximation of valve motion.

Note that superimposing p low-frequency equations of the type (14) allows the time step to be larger; we choose:

$$\Delta t \leq \Delta t_{cr} = \frac{T_p}{\gamma} \quad (16)$$

where $T_p = \frac{2\pi}{\omega_p} \gg T_N$. As a consequence, the convergence of the model is numerically much faster with modal analysis than with direct integration of the governing equations.

Due to orthonormality of the modes, the approximation error is:

$$(\mathbf{U} - \hat{\mathbf{U}})^2 = \sum_{i=p+1}^N \tilde{u}_i^2 \quad (17)$$

This means that the approximation error is a rapidly decreasing function of the truncation frequency p .




Fig. 7. Approximation error distribution for the low-order deformation of a sphere into the human head. Darker regions indicate more important errors, while lighter regions represent best-recovered ones. a. Computation using 60 modes. b. 243 modes. c. 468 modes. d. 2790 modes. (Exact computation needs 37,683 modes).

The distribution of the approximation error along the surface for increasing values of p is shown in Fig. 7. Darker colors outline regions with bigger error. The model being

initially a sphere, one can observe that the error is more important in the regions that need a locally important deformation of the sphere (nose, eyes, ears), while more spherical regions (top of the head) are better approximated, as expected. Note however that the level of accuracy can be easily controlled, at the expense of additional modes.

3.2 Analytic Modal Analysis

Even as a precalculation, solving the generalized eigenproblem is very costly as soon as we consider 3D boundaries (surfaces). For instance, if we consider a mesh of 100×100 nodes, a generalized eigenproblem where the size of the matrices is $10,000 \times 10,000$ has to be solved. It is clear that the analytic expression of the modes would noticeably reduce the computations [30]. This leads us to consider the theory of solid state physics, where similar problems are encountered at a microscopic level (ionic vibrations of a crystal lattice). If we parameterize our deformable curves by arc length, and similarly our deformable surfaces by natural coordinates, we get periodic boundary conditions which depend on the surface topology. This allows the analogy between our deformable model and a crystal lattice.

3.2.1 Free Vibrations of a Chain

The classical theory of vibration of a crystal lattice is based on the *harmonic approximation*, a theory which assumes that the first non-vanishing correction to the equilibrium potential energy is quadratic [3]:

$$V^{harm} = \frac{1}{2} \sum_{\substack{\vec{R}, \vec{R}' \\ \mu, \nu}} u_{\mu}(\vec{R}) D_{\mu\nu}(\vec{R} - \vec{R}') u_{\nu}(\vec{R}') \quad (18)$$

where $u_{\mu}(\vec{R})$ is the displacement in the μ direction of the ion whose mean position is \vec{R} , and \mathbf{D} is the Hessian matrix of the interaction energy.

Consider now a set of ions distributed along a *closed chain* at points separated by a distance a , so that the lattice vectors are $R = na$ for $n \in \{1, \dots, N\}$. If only neighboring ions interact, we may take the harmonic potential energy to have the form:

$$V^{harm} = \frac{1}{2} K \sum_{n=1}^N [u(na) - u((n+1)a)]^2 \quad (19)$$

where $K = v''(a)$ is the stiffness constant of the system, $v(x)$ being the interaction energy of two ions at distance x along the chain. The free vibrations of the lattice are governed by:

$$M \ddot{u}(na) = - \frac{\partial V^{harm}}{\partial u(na)} = K(u((n+1)a) + u((n-1)a) - 2u(na)) \quad (20)$$

These are precisely the equations that would be satisfied if each ion were connected to its neighbors by perfect massless springs of stiffness K (and equilibrium length a , although the equations are in fact independent of the equilibrium length of the spring).

We seek solutions to (20) of the form $Ae^{i(kna - \omega t)}$. This yields the dispersion equation which gives the relationship between spatial (k) and temporal (ω) frequencies:

$$\omega^2(p) = \frac{4K}{M} \sin^2 \left(\frac{k(p)a}{2} \right) \quad (21)$$

The periodicity of the closed chain is expressed by $u[(n + N)a] = u(na)$. We now obtain the set of independent solutions:

$$k(p)a = \frac{2p\pi}{N} \quad p \in \mathcal{B}(N) \tag{22}$$

where $\mathcal{B}(N)$ is the first Brillouin zone.

$\mathcal{B}(N)$ is equal to $\{-\frac{N}{2} + 1, \dots, \frac{N}{2}\}$ for N even, and $\{-\frac{N-1}{2}, \dots, \frac{N-1}{2}\}$ for N odd.

The general solution of the free vibrations of the closed chain is the linear combination of the former solutions:

$$u(na, t) = \sum_{p \in \mathcal{B}(N)} A(p)e^{i(k(p)na - \omega(p)t)} \tag{23}$$

$$= \sum_{p \in \mathcal{B}(N)} A(p)e^{-i\omega(p)t} e^{ik(p)na} \tag{24}$$

The case of the open chain is very similar. We sum up the expression of the free vibrations for both types of chains in Table 1.

TABLE 1
DISPLACEMENT OF THE NODES OF A FREE VIBRATING CHAIN

	$k(p)a$	$p \in$	$u(na, t)$
closed	$\frac{2p\pi}{N}$	$\mathcal{B}(N)$	$\sum_p A(p)e^{-i\omega(p)t} e^{ik(p)na}$
open	$\frac{p\pi}{N}$	$\{0, \dots, N-1\}$	$\sum_p A(p)e^{-i\omega(p)t} e^{i\frac{k(p)a}{2}} \cos\left(k(p)na - \frac{k(p)a}{2}\right)$

3.2.2 Nonlinear Waves in Discrete Media

When k is small compared with π/a (i.e., when the wavelength is large compared to the interparticle spacing), ω is linear in k (from (21)):

$$\omega = a\sqrt{\frac{K}{M}}|k| \tag{25}$$

This is the type of behavior we are accustomed to in the case of light waves and ordinary sound waves. If ω is linear in k , then the group velocity is the same as the phase velocity (equal to $c = a\sqrt{\frac{K}{M}}$), and both are independent of frequency. Note that if we approximate finite differences by derivatives:

$$\begin{aligned} u((n+1)a) - u(na) &\approx a\partial u / \partial x(na) \\ u((n-1)a) - u(na) &\approx a\partial u / \partial x((n-1)a) \\ u((n+1)a) - u(na) + u((n-1)a) - u(na) &\approx a^2\partial^2 u / \partial x^2 \end{aligned} \tag{26}$$

in (20), we end up with a wave equation of velocity c :

$$\frac{\partial^2 u}{\partial x^2} = \frac{1}{c^2} \frac{\partial^2 u}{\partial t^2} \quad \text{with } c = a\sqrt{\frac{K}{M}} \tag{27}$$

One of the characteristic features of waves in discrete media, however, is that the nonlinearity ceases to hold at wavelengths short enough to be comparable with the interparticle spacing. In the present case ω falls below ck as k increases, and the group velocity drops to zero when $|k|$ reaches π/a .

3.2.3 Analytic Modes for Curves

In the more general case of damped and forced vibrations (which is the case of our governing equations), time dependency is not harmonic, and has to be computed separately. For a closed chain, (23) becomes:

$$u(na, t) = \sum_{p \in \mathcal{B}(N)} \tilde{u}_p(t) e^{ik(p)na} \tag{28}$$

Comparing (28) with the modal superposition (13) yields the analytic expression of the modes for a closed chain. Note that since time and space dependency are separate in (28), modal analysis is the decomposition of the displacement in a basis of *standing waves* which are the vibration modes of the system.

TABLE 2
ANALYTIC EIGENVALUES AND EIGENVECTORS FOR CLOSED AND OPEN CURVES

	$\omega^2 \times M/4K$	$\phi(p)$
closed	$\sin^2\left(\frac{p\pi}{N}\right)$	$\left[\dots, \cos\frac{2p\pi n}{N}, \dots\right]^T$
open	$\sin^2\left(\frac{p\pi}{2N}\right)$	$\left[\dots, \cos\frac{p\pi(2n-1)}{2N}, \dots\right]^T$

Table 2 sums up the frequencies (eigenvalues) and modes (eigenvectors) for closed and open curves.

3.2.4 Analytic Modes for Surfaces

The generalization to surface meshes is done by mixing all possible pairs of boundary conditions; it yields three different surface topologies:

- open and open \rightarrow torus topology,
- closed and closed \rightarrow plane topology,
- closed and open \rightarrow cylinder topology

The analytic expressions for these topologies are summed up in Table 3, where the set of variation of the mode parameters p and p' is $\mathcal{B}(N)$ or $\{0, \dots, N-1\}$, depending on the boundary conditions.

Note that we implicitly develop the surface expressions for quadrilateral meshes, and that the modes have to be normalized to unity.

Finally, analytic modal analysis has both theoretical and practical implications. Theoretically, it shows that modal analysis is a specific form of Fourier decomposition of the deformation in a basis of standing waves (28). From a more practical point of view, analytic expressions of the modes are an efficient tool for *real time* eigenvector extraction as soon as the surface topology and the mass and stiffness properties of the model are defined.

3.2.5 Analytic Modes for Volumes

Suppose we wish to model a deformable volume. This may be typically the case for nonrigid motion recovery in a whole 3D image [10]. Therefore we model a volume mesh of size $N \times N' \times N''$. The volume has "open" boundary conditions in all three directions.

Thus, it is easy to generalize the results for the deformable plane, defining this time the modes with three pa-

TABLE 3
ANALYTIC EIGENVALUES AND EIGENVECTORS
FOR DIFFERENT SURFACE TOPOLOGIES

	$\omega^2 \times M/4K$	$\phi(p)$
plane	$\sin^2 \frac{p\pi}{2N} + \sin^2 \frac{p'\pi}{2N'}$	$\left[\dots, \cos \frac{p\pi(2n-1)}{2N} \cos \frac{p'\pi(2n'-1)}{2N'}, \dots \right]^T$
torus	$\sin^2 \frac{p\pi}{N} + \sin^2 \frac{p'\pi}{N'}$	$\left[\dots, \cos \left(\frac{2p\pi n}{N} + \frac{2p'\pi n'}{N'} \right), \dots \right]^T$
cylinder	$\sin^2 \frac{p\pi}{2N} + \sin^2 \frac{p'\pi}{N'}$	$\left[\dots, \cos \frac{p\pi(2n-1)}{2N} \cos \frac{2p'\pi n'}{N'}, \dots \right]^T$

parameters p , p' , and p'' varying, respectively, in $\{0 \dots N-1\}$, $\{0 \dots N'-1\}$, and $\{0 \dots N''-1\}$.

The eigenvalues are:

$$\omega^2(p, p', p'') = \frac{4K}{M} \left(\sin^2 \frac{p\pi}{2N} + \sin^2 \frac{p'\pi}{2N'} + \sin^2 \frac{p''\pi}{2N''} \right) \quad (29)$$

while the eigenvectors have the following expression:

$$\phi(p, p', p'') = \left[\dots, \cos \frac{p\pi(2n-1)}{2N} \cos \frac{p'\pi(2n'-1)}{2N'} \cos \frac{p''\pi(2n''-1)}{2N''}, \dots \right]^T \quad (30)$$

with $n \in \{1, \dots, N\}$, $n' \in \{1, \dots, N'\}$, and $n'' \in \{1, \dots, N''\}$.

3.3 The Deformation Spectrum and Its Applications

DEFINITION. *The deformation spectrum of a motion is the graph representing the value of the modal amplitudes as a function of mode rank: $\tilde{u}_i(t) = f(i)$ [33]. The deformation spectrum is initially drawn for a deformation occurring between two image frames: It describes which modes are excited, and how, in order to deform one object into another. It also gives an indication of the strain energy, as we have:*

$$E_{\text{strain}} = \frac{1}{2} \mathbf{U}^T \mathbf{K} \mathbf{U} = \frac{1}{2} \sum_{i=1}^N \omega_i^2 \tilde{u}_i^2 \quad (31)$$

Note that rigid motion has zero strain energy. Let us now define the term *similar deformations*. Two deformations are similar when the corresponding displacement fields \mathbf{U}_1 and \mathbf{U}_2 are similar within a rigid transform, i.e., if we can find a rotation matrix \mathbf{R} and a translation vector \mathbf{T} such that:

$$\| \mathbf{U}_1 - (\mathbf{R} \mathbf{U}_2 + \mathbf{T}) \| < \epsilon \quad (32)$$

for a small value of ϵ .

TABLE 3

Therefore, it is natural to state that, provided that the dimensionality p of the reduced modal basis is suitably chosen, *two similar deformations have similar deformation spectra*. However, we cannot compare the deformations of two objects placed in arbitrary configurations in the same global reference frame. Thus the modal computations have to be developed in the object reference frame, defined by its *center and axes of inertia*.

In the following, we get rid of the rigid modes in order to study exclusively the deformations. Thus we will refer to the

deformation spectrum as the set of total amplitudes of the first p deformable and low-frequency modes. Once the spectra are computed, we can define a distance measure between the spectra. We choose the Euclidean distance d , such that lower amplitudes are given less importance than higher ones:

$$d(D_1, D_2) = \frac{1}{p} \sqrt{\sum_{i=1}^p (\tilde{u}_i(D_1) - \tilde{u}_i(D_2))^2} \quad (33)$$

where D_1 and D_2 are the labels of the deformations. Distance d gives a *relative* value of how different the two deformations are.

3.3.1 Similarity

Fig. 8 shows two similar deformations and their spectrum; in this example, \mathbf{R} is a 90-degree rotation matrix and \mathbf{T} is arbitrary.

These are the deformations of the initial deformable model for segmenting the mitral valve contour in Fig. 3.




Fig. 8. Similar deformations D_1 and D_2 and their spectrum.

3.3.2 Robustness

In order to test the robustness of the spectra in the presence of noise, we add a Gaussian noise of standard deviation σ to the initial 3D diastole, then deform the diastole into the systole, draw the spectrum and compare it to the spectrum of the true deformation.

The results are very promising (Fig. 9): For $\sigma = 0.1$ the difference between the original spectrum and the corrupted one is almost invisible ($d = 0.01$); for $\sigma = 1$ very slight differences appear ($d = 0.10$); for $\sigma = 5$ the main excited modes are still the same ($d = 0.44$).

Note the very chaotic visual representations of the corrupted shapes, while the modal spectra still succeed in extracting the significant deformation information.

3.3.3 Classification

In order to test the matching of a specific deformation to a group of predefined deformations, we can consider several admissible deformations X_i of a reference shape like the valve (Fig. 10). Under the assumption of Gaussian distribution, we can then classify a test deformation Y as belonging or not to the set $\mathcal{D} = \{ \dots, X_i, \dots \}$ by using the Mahalanobis distance [4], [21]:

$$d_M^2(Y, \mathcal{D}) = (Y - \bar{X})^T W^{-1} (Y - \bar{X}) \tag{34}$$

where W is the covariance matrix of the admissible deformations \mathcal{D} , and \bar{X} the mean admissible deformation.

By comparing with a χ^2 table of q degrees of freedom, where q is the rank of the covariance matrix,⁶ we can determine a confidence measure for acceptance ($d_M^2(Y, \mathcal{D}) < \epsilon$) or rejection of the test deformation as being part of the predefined admissible deformations. Fig. 11 shows the classification of four deformations. The confidence that we have in our classification (derived from the χ^2 table) is also indicated. Note that visually, we would classify the four deformations the way the system has done it (that is, reject the first three deformations, and accept the fourth). Note finally that though this set of deformations has been artificially generated, the application of the method to a clinically-significant case is straightforward.

4 TIME EVOLUTION OF THE MAIN MODAL AMPLITUDES

A single deformation spectrum gives a static information about the spatial frequencies of the motion, whereas the time parameter t is not really taken into account.

However, temporal evolution is really what we are interested in. A dynamic deformation process can be much better described and interpreted if we have a time sequence of images. Therefore, we can draw T consecutive deformation spectra for $T + 1$ frames of images showing the temporal evolution of the deformation process. For a chosen mode i , we are interested in the time signals: $\tilde{u}_i(t)$ for different values of i .

5. By deformation we mean the vector of total modal amplitudes, in other terms the deformation spectrum.

6. If $q < p$ then W^{-1} is the pseudo-inverse of W .




Fig. 9. Corrupted initial diastole deformed to the systole (not shown here), and its spectrum, for different values of noise σ ($\sigma = 0.1, \sigma = 1, \sigma = 5$).




Fig. 10. Ten admissible deformations of the valve.




Fig. 11. Classification of four deformations, with the confidence values.

Let us track the moving ventricle (cf. Fig. 1) through all image frames. We wish to find a minimum number of parameters describing this sequence. Can we compress the 4D data information for its further analysis, storage, or transmission?

Let us denote by (P_0, \dots, P_T) the positions of this surface during the $T + 1$ frames of the cardiac cycle. Therefore, $3N \times (T + 1)$ parameters describe the 4D data (i.e., 345,600 parameters with $N = 6,400$ and $T = 17$).

Let us now track the ventricle in the reduced modal space. We obtain T truncated spectra that store the low-frequency modal amplitudes through the cycle. 4D data is now stored in $3N + T \times p$ parameters (i.e., 20,985 parameters with $p = 105$). The numerical value of p is chosen so that each truncated spectrum has 90% of the energy of the corresponding non-truncated spectrum.

We wish to discard as many modes as possible. Therefore we define a criterion for keeping only the most excited modes among the p low-frequency modes. We compute the energy of an amplitude through the cycle:

$$L^2(i) = \int_0^T \tilde{u}_i(t)^2 dt \quad (35)$$

The larger L^2 , the more important the contribution of the corresponding amplitude. In our experiment with this data, the first nine values of L^2 were much larger than the other ones. $q = 9$ was then an obvious threshold for us to choose.

Let us draw the temporal evolution of the q selected modal amplitudes. They are represented by groups of three in Fig. 12.

These curves have, globally, a single period sine shape, which is an expectable result for low-frequency modes during a cardiac cycle. Their shape is quite like the ventricle volume curve as a function of time during a cardiac cycle. The shape of these curves encourage us to perform a fast Fourier Transform: Most probably only a few Fourier harmonics will describe the time evolution of the curves.




Fig. 12. Top-left, top-right and bottom-left: Each figure displays the time evolution of three main modal amplitudes through a cardiac cycle. In total, the nine main amplitudes are displayed. Bottom right: time evolution of a particular modal amplitude.

In order to illustrate this point, let us choose one among these $q = 9$ spectral components (or modal amplitudes). Fig. 12 (bottom right) shows the time evolution of this particular amplitude. Fig. 13 displays the real (top left) and imaginary (top right) parts of the corresponding Fourier Transform; we observe that indeed only low-frequency Fourier harmonics are excited. Therefore we keep the harmonics of rank 0, 1, and 2 of the Fourier spectra (thus, with symmetry considerations, a total of $H = 5$ harmonics, see Fig. 13 bottom left and bottom right); we then reconstruct the time signal in Fig. 14.




Fig. 13. Fourier (top) and Truncated Fourier (bottom) spectra of the modal amplitude displayed in Fig. 12, bottom-right.




Fig. 14. Reconstruction of the time evolution of the modal amplitude in Fig. 12 bottom-right from truncated Fourier spectra.




Fig. 15. Space-time evolution of the left ventricle model. In each row: Left: Finite element formulation. Center: Modal recovery. Right: Modal plus Fourier recovery.

As for motion compression, instead of keeping a whole 3D image at every frame (that makes $t_x t_y t_z = 1,078,000$ parameters per frame) we now keep only $H \times q / (T + 1) = 2.5$ parameters per frame. This compression of 4D data is impressive: It means that, provided that we keep the first shape, we are able to synthesize the motion with 2.5 parameters per frame. This 4D data is then described by $3N + H \times q = 19,245$ parameters instead of $(T + 1) \times t_x t_y t_z = 19,404,000$, which makes a compression of more than 10^3 , for a cardiac cycle whose duration is approximately one second. If the analysis was extended over a much longer period (e.g., 10 minutes), the compression rate would continuously increase toward the asymptotic value of $1,078,000/2.5$, i.e., approximately 4.10^5 per frame. If we compare the total compression with the transmission of all the 3D images over the 10 minutes, we would get an asymptotic compression factor superior to 10^9 .

On the top of the compression ability, the synthetic information that our spatiotemporal processing provides allows comparison of nonrigid motion by comparing very few parameters. The method is indeed useful not only for compression, but also for analysis of dynamic motion for diagnosis purposes.

The evolution of the left ventricle model over a cardiac cycle is displayed in Fig. 15. In each row, we observe: Left, the evolution of the model in the real space (6); center, the approximation with $q = 9$ modes (spatial smoothing); right, spatiotemporal processing with nine modes and five Fourier harmonics.

Finally, Fig. 16 displays the mean Euclidean error between the real-space mesh (complete computation) and respectively spatial smoothing (modal recovery) and spatiotemporal smoothing (modal and Fourier recovery). Note that this error (in voxels) is extremely small. The position of the two peaks indicate that the error is maximum in the middle of diastolic and systolic shapes: This is because the change in the overall shape is maximum in those frames. An interesting topic would be the study of the error range, and its relationship with the number of modes and Fourier harmonics.




Fig. 16. Evolution of the approximation error (in voxels) during a cardiac cycle, for modal recovery (solid line) and modal plus Fourier recovery (dotted line).

5 CONCLUSION

In this article, we presented a frequency-based analysis of nonrigid motion by coupling modal analysis (outlining the main spatial frequencies of the motion), and Fourier analysis (outlining the main temporal frequencies, mainly for cyclic motions).

Nonrigid motion is first estimated by a physically based deformable model.

The modal decomposition of the deformation, made real-time and precise by using the analytic expressions of the modes, leads to the definition of the deformation spectrum, which is a compact description of the deformation allowing its straightforward comparison and classification (identifying pathologic and normal deformations).

We then introduced a temporal analysis of nonrigid motion from 4D data by recovering the temporal evolution of the main modal amplitudes by a Fourier analysis.

The method provides a very nice description of 4D data by very few parameters (a few modal amplitudes and a few Fourier harmonics). It has important applications in medical analysis of nonrigid motion (mainly for automatic diagnosis purposes), and also provides a tremendous compression of multidimensional nonrigid motion, mainly for storage or transmission purposes.

Our future work will focus on the clinical use and validation of this framework, and its transfer toward concrete medical applications, in particular in the cardiology field.

ACKNOWLEDGMENTS

We wish to thank Dr. Richard Robb at Biomedical Imaging Resource, Mayo Foundation/Clinic, for providing the DSR data. This work was supported in part by a grant from Digital Equipment Corporation.

REFERENCES

- [1] A.A. Amini, R. Curwen, R.T. Constable, and J.C. Gore, "MR Physics-Based Snake Tracking and Dense Deformations from Tagged Cardiac Images," *AAAI 1994 Spring Symp. Series. Application of Computer Vision in Medical Image Processing*, Stanford Univ., March 1994.
- [2] A.A. Amini and J.S. Duncan, "Bending and Stretching Models for LV Wall Motion Analysis Form Curves and Surfaces," *Image and Vision Computing*, vol. 10, no. 6, pp. 418-430, 1992.
- [3] N.W. Ashcroft and M.D. Mermin, *Solid State Physics*. Saunders College Publishing Int'l Ed., 1976.
- [4] N. Ayache, *Artificial Vision for Mobile Robots—Stereo-Vision and Multisensory Perception*. MIT Press, 1991.
- [5] N. Ayache, "Medical Computer Vision, Virtual Reality and Robotics—Promising Research," *Image and Vision Computing*, vol. 13, no. 4, pp. 295-313, May 1995.
- [6] N. Ayache, I. Cohen, and I. Herlin, "Medical Image Tracking," *Active Vision*, chapter 17. MIT Press, Dec. 1992.
- [7] E. Bardinet, L.D. Cohen, and N. Ayache, "Superquadrics and Free-Form Deformations: A Global Model to Fit and Track 3D Medical Data," *First Conf. Computer Vision, Virtual Reality and Robotics in Medicine (CVRMed '95)*, N. Ayache, ed., vol. 905, Nice, France, Apr. 1995. Springer Verlag.
- [8] K.J. Bathe, *Finite Element Procedures in Engineering Analysis*. Prentice-Hall, 1982.
- [9] S. Benayoun, N. Ayache, and I. Cohen, "Adaptive Meshes and Nonrigid Motion Computation," *Int'l Conf. Pattern Recognition*, Jerusalem, Israel, March 1994.
- [10] S. Benayoun, C. Nastar, and N. Ayache, "Dense Non-Rigid Motion Estimation in Sequences of 3D Images Using Differential Constraints," *First Conf. Computer Vision, Virtual Reality and Robotics in Medicine (CVRMed '95)*, N. Ayache, ed., vol. 905, Nice, France, Apr. 1995. Springer Verlag.
- [11] A. Blake, R. Curwen, and A. Zisserman, "Affine-Invariant Contour Tracking with Automatic Control of Spatiotemporal Scale," *IEEE Proc. Fourth Int'l Conf. Computer Vision*, pp. 66-75, Berlin, Germany, May 1993.
- [12] L. Bookstein, "Principal Warps: Thin-Plate Splines and the Decomposition of Deformations," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 11, no. 6, pp. 567-585, June 1989.
- [13] C.W. Chen, T.S. Huang, and M. Arrott, "Modeling, Analysis, and Visualization of LV Shape and Motion by Hierarchical Decomposition," *IEEE Trans Pattern Analysis and Machine Intelligence*, vol. 16, no. 7, Apr. 1994.
- [14] I. Cohen, L.D. Cohen, and N. Ayache, "Using Deformable Surfaces to Segment 3D Images and Infer Differential Structures," *Computer Vision, Graphics, and Image Processing*, vol. 26, no. 2, pp. 242-263, 1992.
- [15] T.F. Cootes, C.J. Taylor, D.H. Cooper, and J. Graham, "Active Shape Models—Their Training and Application," *Computer Vision and Image Understanding*, vol. 61, no. 1, pp. 38-59, Jan. 1995.
- [16] T.F. Cootes, A. Hill, C.J. Taylor, and J. Haslam, "The Use of Active Shape Models for Locating Structures in Medical Images," *Proc. 13th Int'l Conf. Information Processing in Medical Imaging*, Flagstaff, Ariz., June 1993.
- [17] T.F. Cootes and C.J. Taylor, "Combining Point Distribution Models with Shape Models Based on Finite Element Analysis," *Image and Vision Computing*, vol. 13, no. 5, 1995.
- [18] L.L. Creswell, S.G. Wyers, J.S. Pirollo, W.H. Perman, M.W. Vannier, and M.K. Pasque, "Mathematical Modeling of the Heart Using Magnetic Resonance Imaging," *IEEE Trans. Medical Imaging*, vol. 11, no. 4, pp. 581-589, 1992.
- [19] P.E. Danielsson, "Euclidean Distance Mapping," *Computer Vision, Graphics, and Image Processing*, vol. 14, pp. 227-248, 1980.
- [20] H. Delingette, M. Hebert, and K. Ikeuchi, "Shape Representation and Image Segmentation Using Deformable Surfaces," *IEEE Proc. Computer Vision and Pattern Recognition*, pp. 467-472, Lahaina, Maui, Hawaii, June 1991.
- [21] R.O. Duda and P.E. Hart, *Pattern Classification and Scene Analysis*. John Wiley and Sons, 1973.
- [22] A. Guézic, "Large Deformable Splines: Crest Lines and Matching," *Proc. Fourth Int'l Conf. Computer Vision*, Berlin, May 1993.
- [23] W.-C. Huang and D.B. Goldgof, "Adaptive-Size Physically Based Models for Nonrigid Motion Analysis," *Proc. Computer Vision and Pattern Recognition. IEEE Computer Society Conf.*, pp. 833-835, Champaign, Ill., June 1992.
- [24] M. Kass, A. Witkin, and D. Terzopoulos, "Snakes: Active Contour Models," *Int'l J. Computer Vision*, vol. 1, pp. 321-331, 1987.
- [25] F. Leitner, I. Marque, S. Lavallée, and P. Cinquin, "Dynamic Segmentation: Finding the Edge with Snake-Splines," *Proc. Int'l Conf. Curves and Surfaces*, pp. 1-4, Chamonix, France, June 1990. Academic Press.
- [26] T. McInerney and D. Terzopoulos, "A Dynamic Finite Element Surface Model for Segmentation and Tracking in Multidimensional Images with Application to Cardiac 4D Image Analysis," *J. Computerized Medical Imaging and Graphics*, vol. 19, no. 1, pp. 69-83, 1995.
- [27] D. Metaxas and D. Terzopoulos, "Shape and Non-Rigid Motion Estimation Through Physics-Based Synthesis," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 15, no. 6, pp. 580-591, 1993.
- [28] O. Monga and R. Deriche, "3D Edge Detection Using Recursive Filtering," *Proc. Computer Vision and Pattern Recognition*, San Diego, June 1989.
- [29] O. Monga, R. Deriche, G. Malandain, and J-P. Cocquerez, "Recursive Filtering and Edge Closing: Two Primary Tools for 3D Edge Detection," *Proc. First European Conf. Computer Vision (ECCV)*, Antibes, France, Apr. 1990.
- [30] C. Nastar, "Vibration Modes for Nonrigid Motion Analysis in 3D Images," *Proc. Third European Conf. Computer Vision (ECCV '94)*, Stockholm, May 1994.
- [31] C. Nastar and N. Ayache, "Fast Segmentation, Tracking, and Analysis of Deformable Objects," *Proc. Fourth Int'l Conf. Computer Vision (ICCV '93)*, Berlin, May 1993.
- [32] C. Nastar and N. Ayache, "Non-Rigid Motion Analysis in Medical Images: A Physically Based Approach," *Proc. 13th Int'l Conf. Information Processing in Medical Imaging (IPMI '93)*, Flagstaff, Ariz., June 1993.

- [33] C. Nastar and N. Ayache, "Classification of Nonrigid Motion in 3D Images Using Physics-Based Vibration Analysis," *Proc. IEEE Workshop Biomedical Image Analysis*, Seattle, June 1994.
- [34] C. Nastar and N. Ayache, "Time Representation of Deformations: Combining Vibration Modes and Fourier Analysis," *Lecture Notes in Computer Science 994: Object Representation in Computer Vision*. Springer-Verlag, 1995.
- [35] J. Park, D. Metaxas, and L. Axel, "Volumetric Deformable Models with Parametric Functions: A New Approach to the 3D Motion Analysis of the LV from MRI-SPAMM," *Proc. Fifth Int'l Conf. Computer Vision (ICCV '95)*, Boston, June 1995.
- [36] A. Pentland and B. Horowitz, "Recovery of Non-Rigid Motion and Structure," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 13, no. 7, pp. 730-742, July 1991.
- [37] A. Pentland and S. Sclaroff, "Closed-Form Solutions for Physically-Based Shape Modeling and Recognition," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 13, no. 7, pp. 715-729, July 1991.
- [38] A. Pentland and J. Williams, "Good Vibrations: Modal Dynamics for Graphics and Animation," *Computer Graphics*, 1989.
- [39] P. Shi, G. Robinson, A. Chakraborty, L. Staib, R. Constable, A. Sinuasa, and J. Duncan, "A Unified Framework to Assess Myocardial Function from 4D Images," *Proc. First Conf Computer Vision, Virtual Reality, and Robotics in Medicine (CVRMed '95)*, N. Ayache, ed., vol. 905, Nice, France, Apr. 1995. Springer Verlag.
- [40] P. Shi, A. Amini, G. Robinson, A. Sinuasa, C.T. Constable, and J. Duncan, "Shape-Based 4D Left Ventricular Myocardial Function Analysis," *Proc IEEE Workshop on Biomedical Image Analysis*, Seattle, Wash., June 1994.
- [41] R. Szeliski and S. Lavallée, "Matching 3D Anatomical Surfaces with Non-Rigid Volumetric Deformations," *Proc. AAAI 1994 Spring Symp. Series. Application of Computer Vision in Medical Image Processing*, Stanford Univ., March 1994.
- [42] D. Terzopoulos and K. Fleischer, "Deformable Models," *The Visual Computer*, vol. 4, pp. 306-331, 1988.
- [43] D. Terzopoulos and K. Waters, "Physically Based Facial Modeling, Analysis, and Animation," *J. Visualization and Computer Animation*, vol 1, pp. 73-80, 1990.
- [44] D. Terzopoulos and K. Waters, "Analysis and Synthesis of Facial Image Sequences Using Physical and Anatomical Models," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 15, no. 6, pp. 569-579, 1993.
- [45] D. Terzopoulos, A. Witkin, and M. Kass, "Constraints on Deformable Models: Recovering 3D Shape and Nonrigid Motion," *AI J.*, vol. 36, pp. 91-123, 1988.
- [46] M. Vasilescu and D. Terzopoulos, "Adaptive Meshes and Shells," *Proc. Computer Vision and Pattern Recognition, Proc. IEEE Computer Society Conf.*, June 1992, pp. 829-832, Champaign, Ill.
- [47] Q.Z. Ye, "The Signed Euclidean Distance Transform and Its Applications," *Int'l Conf. Pattern Recognition*, pp. 495-499, 1988.
- [48] A.L. Yuille, D.S. Cohen, and P.W. Hallinan, "Feature Extraction from Faces Using Deformable Templates," *Proc. Computer Vision and Pattern Recognition*, San Diego, June 1989.



Chahab Nastar is a senior research scientist at Institut National de Recherche en Informatique et Automatique (INRIA), Rocquencourt, France.

He received the *Diplôme d'Ingénieur* from *l'Ecole Nationale des Ponts et Chaussées* in Paris in 1991 and the PhD degree in mathematics and computer science in 1994. His research, conducted at INRIA, focused on medical image analysis.

In 1994-95, Dr. Nastar joined the MIT Media Laboratory as a postdoctoral fellow, developing new techniques for image matching and recognition. Dr. Nastar returned to INRIA in 1996, where his current research interests include computer vision, eigentechniques, object recognition, face recognition, content-based image retrieval, and augmented reality.



Nicholas Ayache is a research director at INRIA, Sophia-Antipolis, France, where he has lead the EPIDAURE research group on medical image analysis and robotics since 1989. He is also teaching graduate courses on computer vision at the University of Paris XI, *Ecole Centrale*, *Ecole des Ponts*, *Ensta*, and *Ecole des Mines*. He is also consulting for a number of private companies, including *Matra-Cap-Systemes* and *Focus-Medical*.

Dr. Ayache received his PhD in 1983, and his *Thèse d'Etat* in 1988, both in computer science from the University of Paris XI, on topics related to model-based object recognition, passive stereovision, and multisensor fusion. His current research interests are medical image processing and analysis, shape and motion representation, rigid and nonrigid registration, tracking and analysis of deformable objects, and image guided and simulated therapy.

Dr. Ayache is currently coeditor-in-chief of *the Journal of Medical Image Analysis* (Oxford University Press), a member of the editorial board of the *International Journal of Computer Vision* (Kluwer), associate editor of *Medical Imaging* (IEEE), and area editor of the *Journal of Computer Vision and Image Understanding* (Academic Press). He also served on the editorial board of the *IEEE Transactions on Robotics and Automation* (1990-93).

Dr. Ayache is the author of the books, *Artificial Vision for Mobile Robots* (MIT Press) and *Vision stéréoscopique et perception multisensorielle* (Inter-Editions). He chaired the first International Conference on Computer Vision, Virtual Reality and Robotics in Medicine (CVRMed), held in Nice, France, in April 1995 and serves on the editorial board of a number of conferences on medical imaging, computer vision, visualization, and robotics, including ECCV, ICCV, VBC, MRCAS, and MMVR.