Anisotropic filtering for model-based segmentation of 4D cylindrical echocardiographic images

J. Montagnat *, M. Sermesant, H. Delingette, G. Malandain, N. Ayache

INRIA—Epidaure, 2004 route des Lucioles, BP 93, 06902 Sophia Antipolis Cedex, France

Abstract

This paper presents a 4D (3D + time) echocardiographic image anisotropic filtering and a 3D model-based segmentation system. To improve the extraction of left ventricle boundaries, we rely on two preprocessing stages. First, we apply an anisotropic filter that reduces image noise. This 4D filter takes into account the spatial and temporal nature of echocardiographic images. Second, we adapt the usual gradient filter estimation to the cylindrical geometry of the 3D ultrasound images. The reconstruction of the endocardium takes place by deforming a deformable simplex mesh having an a priori knowledge of left ventricle shape and that is guided by a region-based data attraction force. The external force formulation improves the segmentation robustness against noise and outliers. We illustrate our method by showing experimental results on very challenging sparse and noisy ultrasound images of the heart and by computing quantitative measurements of the left ventricle volume.
© 2002 Elsevier Science B.V. All rights reserved.

Keywords: Ultrasound image; Cylindrical geometry; Anisotropic filtering; Deformable models; Echocardiography

1. Introduction

3D ultrasound imagery is a very promising image acquisition technology due to its low cost and non-invasive nature. Furthermore, the development of real-time 3D ultrasound probes creates numerous medical applications ranging from diagnosis to therapy planning and control. For instance, it allows the reconstruction of the left ventricle motion through the quantitative computation of the ventricle ejection fraction which is of major importance in detecting heart pathologies (Clarysse et al., 1997; Declerck et al., 1998). Unfortunately, the speckle noise in ultrasound images combined with a low image resolution of 3D ultrasound make it a very challenging task. In this paper, we propose a reconstruction algorithm of the left ventricle from echocardiographic images based on 3D deformable models. The input images are time series of 3D ultrasound (or 4D images) that are filtered to reduce noise before segmentation.

In some cases, 3D ultrasound systems are based on a moving 2D probe acquiring along non-aligned planes (Treece et al., 1998). In this paper,

---
* Corresponding author. Address: CREATIS, INSA 502, 20 bd A. Einstein, 69621 Villeurbanne Cedex, France.
E-mail addresses: johan@creatis.insa-lyon.fr (J. Montagnat), msermesa@sophia.inria.fr (M. Sermesant), hdeling@sophia.inria.fr (H. Delingette), greg@sophia.inria.fr (G. Malandain), na@sophia.inria.fr (N. Ayache).
URLs: http://www.creatis.insa-lyon.fr, http://www-sop.inria.fr/epidaure.
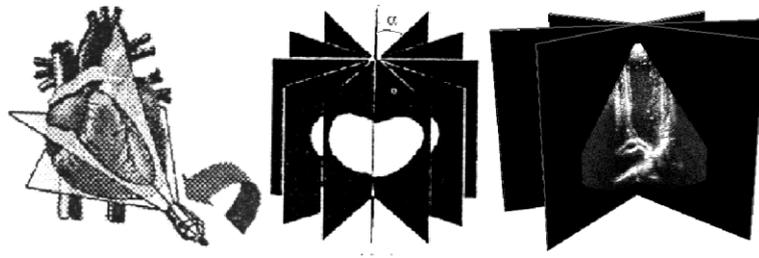
Fig. 1. The rotative probe and the resulting US images with cylindrical symmetry.

we consider rotative probes producing a set of 2D planes intersecting each other along a revolution axis. These images are acquired with a cylindrical geometry implying that the density of information decreases as a function of the axis distance (Fig. 1).

Previous work (Winterfeldt et al., 1997; Jacob et al., 1999) have considered 2D image segmentation and tracking approaches (Nastar and Ayache, 1996; McEachen and Duncan, 1997) without any spatial coherence between slices. Statistical information (mainly principal component analysis) has also been used to restrain the set of possible deformations of a left ventricle model (Jacob et al., 1999; Giachetti, 1998). Deformable surfaces have been recently used to reconstruct the heart motion from 3D echocardiographic images (Berger et al., 1999), but most of them rely on tagged MRI (Guttman et al., 1994; Reynard et al., 1995; Young et al., 1995; Park et al., 1996; Declerck et al., 1998).

Recent methods where proposed in the literature for the reconstruction of cardiac cycles in 4D ultrasound images. In (Papademetris et al., 2001), biomedical models are used for heart motion analysis. The segmentation stage is semi-automatic and simplified by using opened chest dog heart images. Surface-based methods have been proposed for left ventricle segmentation, either relying on region clustering (Sanchez-Ortiz et al., 2000) or prior shape and motion constraints (Montagnat and Delingette, 2000) to make the segmentation robust to ultrasound noise. To our knowledge, no other work deal with the native geometry of 3D ultrasound images acquired by a rotative probe.

Alternatively, several authors proposed to reduce speckle level by spatial compounding (Rohling et al., 1997) or to detect boundaries in US images by modeling speckle noise (Czerwinski, 1998). Seg-

mentation approaches based on voxels clustering after noise correction have been also proposed (Boukerroui et al., 1999; Stetten and Pizer, 2000).

In this paper, we propose a model-based segmentation procedure. Prior to left ventricle segmentation, the image is filtered to reduce speckle noise. A 4D anisotropic diffusion procedure benefits from the weak correlation of speckle through time. Image boundaries are then extracted using a differential operator specialized for the cylindrical geometry of the ultrasound images acquired using a rotative 2D probe. A deformable surface model including prior knowledge on the shape to recover is then used to segment each 3D volume composing the 4D sequence. The surface models compensate for the weak image resolution and data occlusions. Moreover, a region-based data term that benefits from the prior diffusion stage is used. Segmentation steps are illustrated in Fig. 2.

The main contributions of our approach are:

- to consider the full 3D nature of input images with a 3D model deformation,
- to introduce prior knowledge on the shape to recover,
- to take into account temporal consistency through 4D anisotropic diffusion,
- to adapt image filtering techniques to the ultrasound images cylindrical geometry,
- and to use a region-based approach for finding the boundaries of the left ventricle.

2. Image filtering

To get rid of image noise, many filtering techniques have been proposed based on the compu-
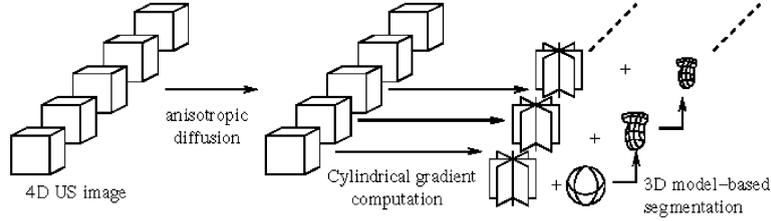
Fig. 2. Segmentation steps: anisotropic diffusion (Section 2), gradient computation in cylindrical geometry (Section 3) and model-based segmentation (Section 4).

tation of a mean intensity value on a local neighborhood at each image voxel. However, these techniques also tend to displace structures and blur their boundaries. This side effect is critical for segmentation algorithms relying on the detection of structures boundaries such as deformable model approaches. Anisotropic diffusion (Weickert, 1998; Gerig et al., 1992; Alvarez et al., 1992; Perona and Malik, 1990) has been developed to smooth an image, thus removing high frequency noise, while preserving the boundaries of structures of interest.

2.1. 4D anisotropic diffusion

The anisotropic diffusion theory in imaging raised from an analogy with fluids diffusion. It states that the image intensity I, seen as a fluid concentration, is evolving toward an equilibrium state according to the diffusion equation $\partial_t I = \mathrm{div}(\boldsymbol{D}\nabla I)$ where $\boldsymbol{D}$, a positive definite symmetric matrix (the diffusion tensor) and ∂_t denotes the derivation with respect to the time.

For ultrasound images filtering, we use a diffusion tensor that depends on the position of the voxel in the image (heterogeneous), that is computed from the evolving image (non-linear) and that depends on the direction of the gradient at this position (anisotropic). We use the diffusion function $g(x, a) = 1 - \mathrm{e}^{-(3.315/(x/a)^4)}$ (Weickert, 1998), with a the gradient threshold (the value above which the gradient is considered high enough for a voxel to be a boundary point). It acts as an edge-enhancing filter: where the image gradient is high, the diffusion is high in the direction parallel to the boundary and negative in the direction orthogonal to the boundary. Therefore, the intensity gap increases and this process leads to an

enhancement of boundaries. The anisotropic diffusion tensor is written as

$$\boldsymbol{D} = \mathrm{diag}\left(g\left(\frac{\partial I}{\partial x}, a_x \right), g\left(\frac{\partial I}{\partial y}, a_y \right), \right.$$
$$\left. g\left(\frac{\partial I}{\partial z}, a_z \right), g\left(\frac{\partial I}{\partial t}, a_t \right) \right)$$

With a tensor whose components are different in all directions, it is possible to vantage the diffusion in a direction (along boundaries) and, conversely, to prevent smoothing in another direction (orthogonal to a boundary).

The diffusion equation is discretized by the finite differences method. Let $\boldsymbol{A}$ denote the matrix of the spatial derivations, I^t the image (written as a vector) at time t and Δt the time step. If $\boldsymbol{D}_i^t$ is the value of $\boldsymbol{D}$ at a voxel i and at instant t, we have:

$$\boldsymbol{A}_{ij}(I^t) = \begin{cases} \dfrac{\boldsymbol{D}_i^t + \boldsymbol{D}_j^t}{2} & \text{if } j \in N(i), \\ -\sum_{n \in N(i)} \dfrac{\boldsymbol{D}_i^k + \boldsymbol{D}_n^k}{2} & \text{if } j = i, \\ 0 & \text{otherwise,} \end{cases}$$

where $N(i)$ is the neighborhood of i. Let $\boldsymbol{A}_d$ be the spatial derivation in direction d, we can write the diffusion equation in a discrete form as a *semi-implicit* scheme:

$$\frac{I^{t+\Delta t} - I^t}{\Delta t} = A(I^t)I^{t+\Delta t} \Rightarrow I^{t+\Delta t}$$
$$= \left(\boldsymbol{Id} - \Delta t \sum_{d=1}^{4} \boldsymbol{A}_d(I^t) \right)^{-1} I^t.$$

This scheme ensures stability with quite large time steps (around 1.0, compared to 0.001 for an explicit scheme) and a low enough computation cost, but as the dimension increases (3 or 4 dimensions)

the system solution becomes costly. The idea of numerical schemes based on *additive operator splitting* (AOS) is to separate dimensions and re-arrange image in order to only invert tridiagonal matrices ($N(i)$ is then only composed of the two neighboring voxels in the considered dimension):

$$I^{t+\Delta t} = \frac{1}{4} \sum_{d=1}^{4} (\boldsymbol{Id} - 4\Delta t \boldsymbol{A}_d(I^t))^{-1} I^t$$

AOS scheme has the same first-order Taylor expansion in Δt as the semi-implicit scheme (therefore the same convergence speed and stability) and it fits in the discrete multi-scale theory framework (Weickert, 1998).

As described in Section 3, there are different techniques to compute the image gradient direction. In the diffusion process, the gradient is computed at each iteration, so the computation must be fast. Moreover we want to have a symmetry between spatial and temporal data, as they are processed the same way and the AOS numerical scheme allows to parallelize the computation in all dimensions. For these reasons, the gradient is simply computed by central finite differences in the diffusion process.

One key parameter of the diffusion process is the gradient threshold a. As the diffusion process evolves, the homogeneity of image regions increases and the occurrence of high gradient voxels decreases. Thus the threshold must decrease with time. It must also be different for each dimension, as the information density and variation is different.

2.2. Ultrasound images filtering

Anisotropic diffusion is suited to filter 4D ultrasound images. It allows us to drastically reduce the speckle noise while enhancing the structure boundaries. It is a pre-segmentation tool that can noticeably improve the quality of model-based reconstruction.

Since ultrasound images speckle is not necessarily correlated in time, it appears as a high frequency noise in the time dimension. Although the speckle makes some parts of the boundaries disappear, the time diffusion allows to recover some

of them. Thus, it seems meaningful to perform anisotropic diffusion on 4D images (Elayadi, 1997; Sanchez-Ortiz et al., 1999). Brandt et al. (1999) similarly use time and space information for a better detection of ultrasound image boundaries.

2.3. Experiments

The 4D anisotropic diffusion is applied on the image sequences for three iterations of 1.0 time unit. The thresholds a_x, a_y, a_z and a_t are computed at each step as a fixed fraction (80%) of the cumulated histogram of the gradient component considered (as a percentage, it decreases when the number of edges is reduced, so the remaining boundaries are well preserved). Since boundaries correspond to high gradients, thresholds are chosen in the small upper part of the histogram. Threshold values are displayed in the table below:

	Initial image	First iteration	Second iteration
a_x	10	5	4
a_y	14	7	5
a_z	26	19	16
a_t	22	13	10

Thresholds are decreased by almost 50% after the first iteration for every dimension: as expected, the most important part of the smoothing is done during this iteration and thresholds are significantly different clearly showing the data anisotropy.

We can observe on Fig. 3 that the anisotropic filtering smooths the noise, but keeps the sharpness of boundaries, even enhancing it at some places. We can also observe that the 4D diffusion is better in filtering noise in homogeneous zones and tends to produce more continuous boundaries, which is, indeed, very interesting for model-based segmentation. Diffusion is an effective preprocessing but it remains computationally expensive. In a time series composed of five 3D images whose dimensions are $256 \times 256 \times 17$, the diffusion process takes 30 min while the deformation of the model in all image of the sequence takes less than 5 min using the same hardware.
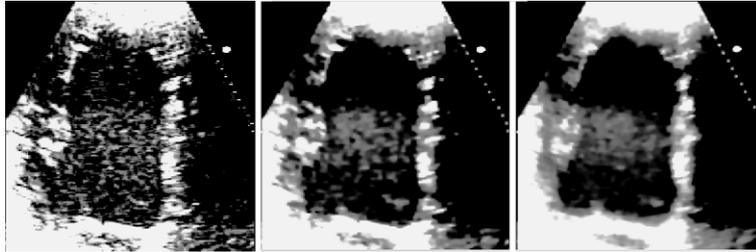
Fig. 3. 4D anisotropic diffusion in ultrasound images of the heart. Left: the original image. Center: the 3D diffused image. Right: the 4D diffused image (same diffusion parameters).

3. Gradient computation in cylindrical geometry

Model-based segmentation relies on boundaries detection to guide deformations. This usually involves gradient computation from the input images. In 3D images based on a regular lattice (Cartesian images), the computation of gradient vectors is mostly based on separable filters (e.g., Sobel operators) or recursive filtering (e.g., Canny–Deriche filters (Monga et al., 1991)). When considering images of cylindrical geometry, classical approaches can be used in each slice, computing 2D gradients and neglecting the tangential component. It is also possible to interpolate the image on a regular lattice before applying 3D operators. Unfortunately, this approach is hopeless due to the important angular resolution and the low signal-to-noise ratio (see left of Fig. 4 for a resampling example on a realistic echocardiographic image composed by nine planes). Instead, we compute a 3D gradient vector for each image voxel in the cylindrical geometry (see Herlin and Ayache (1992) for a similar approach in 2D).

Let M_0 be a point in Euclidean space with Cartesian coordinates (x_0, y_0, z_0) and cylindrical coordinates (r_0, θ_0, z_0) (see right of Fig. 4). Let I be the cylindrical image defined as $I : \Gamma \subset \mathbb{R}^3 \to \mathbb{R}$, where Γ is the discrete cylindrical grid over which the grey-level values are known. The convolution of a 3D signal S by a filter f at point M_0 is defined in Cartesian space as

$$(S \otimes f)(M_0) = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} f(x_0 - x, y_0 - y, z_0 - z) \times S(x, y, z) \, \mathrm{d}x \, \mathrm{d}y \, \mathrm{d}z. \tag{1}$$

To write this equation in cylindrical space we use the coordinate transformation $x = r\cos(\theta)$, $y = r\sin(\theta)$, and $z = z$. Let J be the Jacobian matrix corresponding to this transformation. A variable change in Eq. (1) leads to
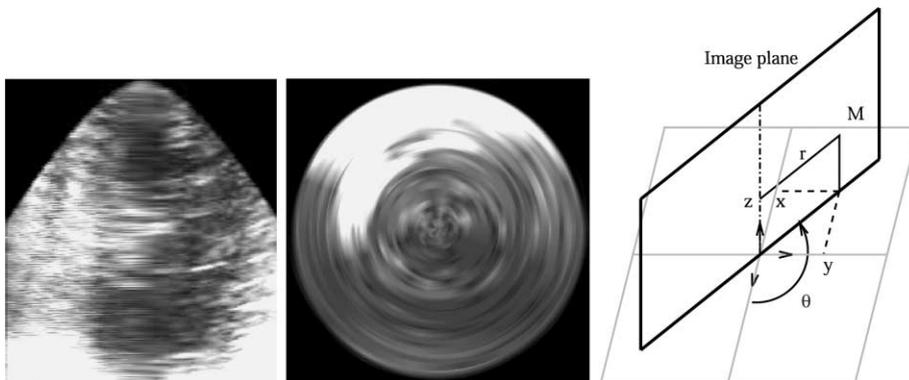
Fig. 4. Left: echocardiographic image interpolation on a Cartesian grid. Right: cylindrical geometry.
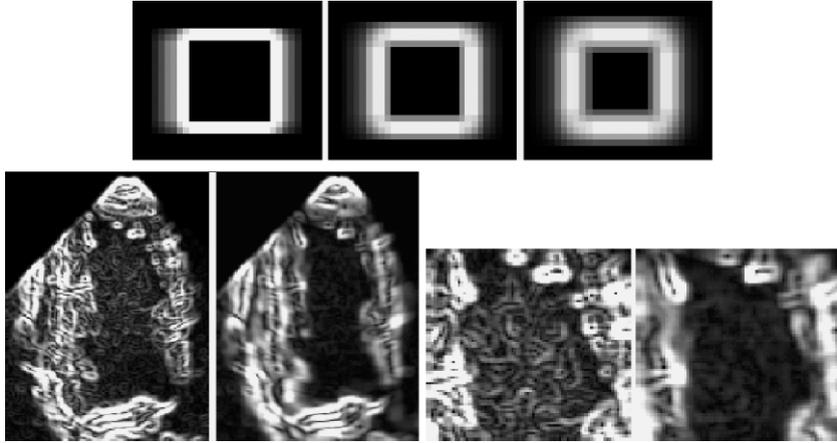
Fig. 5. Top row: cylindrical gradient filter response on a cube; Bottom row: comparison of Sobel operators (left) and the cylindrical gradient filter (right) on a slice. A magnification around the symmetry axis is given that clearly shows the filtering improvement.

$$(S \otimes f)(M_0) = \int_{-\infty}^{\infty} \int_{0}^{2\pi} \int_{0}^{\infty} f(x_0 - r\cos(\theta), y_0 - r$$
$$\times \sin(\theta), z_0 - z)S(r, \theta, z)|J| \, \mathrm{d}r \, \mathrm{d}\theta \, \mathrm{d}z. \quad (2)$$

Since image I is a discrete signal, we discretize Eq. (2). Let δr, $\delta\theta$, and δz be the dimensions of the discrete filter applied to I. The filtered value of I at point $M_0 \in \Gamma$, knowing that $|J| = r$, is:

$$(I \otimes f)(M_0) = \sum_{z=z_0-\delta z}^{z_0+\delta z} \sum_{\theta=\theta_0-\delta\theta}^{\theta_0+\delta\theta} \sum_{r=r_0-\delta r}^{r_0+\delta r} I(r, \theta, z)$$
$$\times f(x_0 - r\cos(\theta), y_0 - r\sin(\theta), z_0 - z)r.$$

We use Deriche filters (Deriche, 1987) to perform gradient computation in cylindrical geometry. Let D_d and S_d be the mono-dimensional Deriche derivative and smoothing filters respectively, in dimension $d \in \{x, y, z\}$: $D_d(u) = ue^{-\alpha|u|}$ and $S_d(u) = (\alpha|u| + 1)e^{-\alpha|u|}$. The directional gradient operators are defined as $G_x(u) = D_x(u)S_y(u)S_z(u)$, $G_y(u) = S_x(u)D_y(u)S_z(u)$, and $G_z(u) = S_x(u)S_y(u)D_z(u)$. The discrete masks are computed by sampling G_x, G_y, and G_z values then normalizing mask coefficients such that the sum of positive coefficients equals 1 and the sum of negative coefficients equals -1. This ensures that the derivative mask coefficients sum is null.

Top row of Fig. 5 illustrates the filter response on a slice of a synthetic cube image for a 3^3 (left), a 5^3 (center), and a 7^3 (right) filter. Bottom of Fig. 5 compares the output of the 2D Sobel filters (left) and a 3D cylindrical operator (right) in one image plane. A magnification of the center part is also shown. The cylindrical operator improves significantly the edge detection and reduces the speckle effect by filtering along the axial direction, especially close to the rotation axis where dense data is available.

4. Model representation

Many deformable surface representations have been proposed for model-based segmentation of medical images (McInerney and Terzopoulos, 1996). The left ventricle was first modeled using parameterized representation such as surface harmonics (Schudy and Ballard, 1979), superquadrics (Bardinet et al., 1996), or generalized primitives (Park et al., 1996). Discrete surfaces have also been used for their capabilities to model complex shapes (McInerney and Terzopoulos, 1995; Montagnat and Delingette, 2000). Among existing representations, we use the discrete simplex meshes (Delingette, 1999) for their simple geometry and their ability to define shape constraints in a computationally efficient manner.

4.1. Law of motion

A two-simplex mesh is a discrete non-parametric representation of a surface in $\mathbb{R}^3$ defined by a set of vertices $\{\boldsymbol{p}_i\}_i$ and their connectivity. It is a regular three-connected mesh. Its topology is dual to a triangulation, and it can represent surfaces of any topology. The geometry of simplex meshes allows us to define discrete quantities such as mean curvature at each vertex (see Delingette (1999) for details). All vertices follow a Newtonian law of motion. Three kinds of forces are computed at each vertex $\boldsymbol{p}_i$. The internal force f_i^{int} enforces some regularization behavior. The external force f_i^{ext} pulls the model toward the ventricle boundaries. The global displacement force f_i^{global} is computed from a global transformation with few degrees of freedom designed to constrain all vertices displacements. The law of motion is discretized using finite differences and an explicit iterative scheme:

$$\boldsymbol{p}_i^{t+1} = \boldsymbol{p}_i^t + (1-\gamma)(\boldsymbol{p}_i^t - \boldsymbol{p}_i^{t-1}) + \lambda(\alpha f_i^{\text{int}} + \beta f_i^{\text{ext}}) + (1-\lambda)f_i^{\text{global}}, \tag{3}$$

where $\boldsymbol{p}_i^t$ denotes the position of vertex $\boldsymbol{p}_i$, at time t, γ is a damping coefficient, α and β are the internal and external force weights, and λ is a locality weight. The discrete time step is hidden in the α, β and γ coefficients. We use a coarse-to-fine approach, limiting model deformation capabilities by setting λ to zero at the beginning of the deformation process, then gradually increasing it to allow more local deformations. The α coefficient value is set to 1 thus having β variations weight the respective influence of the internal and the external forces. The β coefficient varies from 1 to 0.1 as λ increases. The damping value γ is fixed to 0.35 based on some empirical results showing this value usually provides a fast convergence speed. An automatic algorithm for governing λ evolution was proposed in (Montagnat and Delingette, 1998).

At the beginning of the deformation process. $\lambda = 0$ and, due to the computation of f_i^{global} (Montagnat and Delingette, 1998), Eq. (3) exactly corresponds to the iterative closest point (ICP) algorithm (Besl and McKay, 1992; Zhang, 1994). Thus a rigid registration of the model is first

achieved. This drastic reduction of the model number of degrees of freedom constrains the deformation capability of the model and make the model far less sensitive to its initial position. When rigid registration is achieved, affine registration is applied before allowing local model deformations ($\lambda > 1$).

4.2. Surface regularization

Due to its discrete nature, the regularization of a simplex mesh is not based on the evaluation of surface partial derivatives but on the relative position of a vertex with respect to its neighbors. When a reference shape similar to the structure to segment is known, we rely on a shape regularizing constraint. The internal force is computed from the geometric parameters of simplex meshes.

Let us consider a mesh vertex $\boldsymbol{p}_i$ and its three neighbors $\boldsymbol{p}_{\text{ngh}_j(i)}$, $j \in [1,3]$, defining a tangent plane $\mathscr{P}_i$ with normal $\boldsymbol{n}_i$. The position of $\boldsymbol{p}_i$ is related to the position of its neighbors by the equation:

$$\boldsymbol{p}_i = \varepsilon_i^1 \boldsymbol{p}_{\text{ngh}_1(i)} + \varepsilon_i^2 \boldsymbol{p}_{\text{ngh}_2(i)} + \left(1 - \varepsilon_i^1 - \varepsilon_i^2\right) \boldsymbol{p}_{\text{ngh}_3(i)} + H(\boldsymbol{p}_{\text{ngh}_j(i)}, \varepsilon_i^1, \varepsilon_i^2, \varphi_i) \boldsymbol{n}_i$$

where ε_i^1 and ε_i^2 are two barycentric coordinates called *metric parameters*, φ_i is a curvature related angle called *simplex angle*, and H is an elevation function above $\mathscr{P}_i$ (see Delingette (1999) for details). It can be shown that a simplex mesh shape is defined up to a similarity transformation by the set of its metric parameters and simplex angles $\{\varepsilon_i^1, \varepsilon_i^2, \varphi_i\}_i$.

Each vertex $\boldsymbol{p}_i$ is attracted toward a point $\tilde{\boldsymbol{p}}_i$ on a smoother mesh. Let $\tilde{\varepsilon}_i^j$, $\tilde{\varphi}_i$ and $\tilde{\boldsymbol{p}}_i^\perp$ denote the metric parameters, the simplex angle and the projection of $\tilde{\boldsymbol{p}}_i$ on $\mathscr{P}_i$ respectively. The internal force can be decomposed as the sum of a tangential and normal component:

$$\boldsymbol{f}_{\text{int}}(\boldsymbol{p}_i) = (\tilde{\boldsymbol{p}}_i^\perp - \boldsymbol{p}_i^\perp) + (H(\boldsymbol{p}_{\text{ngh}_j(i)}, \varepsilon_i^1, \varepsilon_i^2, \varphi_i) - H(\boldsymbol{p}_{\text{ngh}_j(i)}, \tilde{\varepsilon}_i^1, \tilde{\varepsilon}_i^2, \tilde{\varphi}_i)) \boldsymbol{n}_i.$$

The tangential component of the internal force controls the vertex spacing over the surface. To ensure uniformly spread vertices, metric parameters

are all set equal: $\tilde{\varepsilon}_i^1 = \tilde{\varepsilon}_i^2 = 1/3$. The normal component constrains the mean curvature of the surface through the simplex angle. The definition of $\tilde{\varphi}_i$ depends on the level of geometric regularity that should be enforced. Let $\{\varphi_i^0\}_i$ be the set of simplex angles defining the reference shape of an anatomical structure. Setting $\tilde{\varphi}_i = \varphi_i^0, \forall i$ constrains the surface to converge toward the reference shape in the absence of external forces. This internal force definition introduce shape prior knowledge in the deformation process. Reference simplex angles φ_i^0 are measured on the reference mesh used for segmentation. This implies deforming a first mesh in a reference image. The reference mesh can be guided manually if needed as this initialization stage needs only to be done once.

4.3. Statistical shape variations

By using simplex meshes, we are using zero-order statistics on shape (i.e. a mean model) with reference curvature information to introduce prior shape knowledge in the segmentation process. Several authors reported positive results in using PCA to compute first-order statistics (i.e. shape variations) (Cootes et al., 1995). However, we found difficult to apply this result in our case. Training on reference surfaces in the 3D case is not straightforward due to the difficulty to match corresponding points on a set of surfaces (except for a very few curvature extrema points). Moreover, relevant statistics can only be extracted from large enough image databases. This is even more critical in 3D were the variation space dimensionality is much higher. In the case of 4D imagery, statistical variation modes would only be relevant for the segmentation of one, say the first, image. Each following image of the sequence would require a new statistical model.

5. External forces for 3D US image segmentation

Our deformable simplex mesh relies on external forces for segmenting the left ventricle in 3D US images. These forces should attract each vertex toward the closest boundary voxel. In echocardiographic images several difficulties arise for finding object boundaries. The *speckle*, inherent to the image formation of US imagery, and the image *noise* mainly originating from the digitization process, are responsible for causing false contour detections (false positives). The *variable information density*, due to the cylindrical geometry of the image, and the *missing information*, due to the low energy of the ultrasound beam reflected at the organ interfaces and the rib occlusions, make the ventricle contours hard to detect (false negatives). We rely on the model regularizing behavior and the robust force expressions proposed below to deal with false positives and false negatives.

5.1. Scan-line algorithm

Closest boundary voxels are computed for each vertex by scanning the image along the vertex normal direction until a boundary voxel is reached. Indeed, it has been shown in (Kimia et al., 1992) that the class of surface deformations is not restricted by only considering displacement along each vertex normal. Thus, we restrict the search for boundary points along the normal direction $\boldsymbol{n}_i$ of each vertex $\boldsymbol{p}_i$.

We proceed by scanning all cylindrical voxels intersected by the normal line $(\boldsymbol{p}_i, \boldsymbol{n}_i)$ in the image volume within a fixed range centered on each vertex (see left of Fig. 6). The image intensity $I(\boldsymbol{p})$ at any normal line point $\boldsymbol{p}$ is tri-linearly interpolated from the eight closest image voxel intensities I_0–I_7 (see center of Fig. 6). This algorithm outputs a list of voxels for which we store their interpolated intensity value and their Euclidean distance from $\boldsymbol{p}_i$. The scan-line range is fixed as a percentage of the image size. We used a 10 voxel length scan line in our experiments.

5.2. Gradient norm extremum

We first use the extracted scan line in order to find the gradient norm extrema along the normal direction. In practice, the normal direction at a gradient vertex quickly converges toward the direction of the gradient vector in presence of a strong boundary. In order to improve this external force computation, we propose to add two addi-
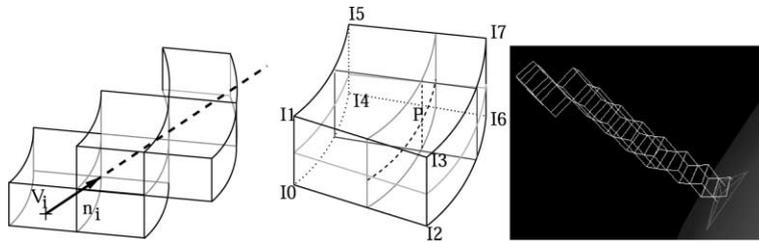
Fig. 6. Normal line and the scanned voxels: Tri-linear interpolation on a cylindrical voxel; 3D display of the scanned voxels.

tional constraints for a voxel to be considered as the closest boundary voxel. Since the cavity of left ventricle is surrounded by structures that appear brighter, we use a *gradient orientation constraint* to keep only boundaries whose gradient vector direction is roughly the same as the normal direction. Also, we constrain the boundary voxel to belong to a given *range of intensity values* in order to eliminate spurious boundary information.

5.3. Region-based algorithm

The previous algorithm may still fail due to the presence of strong edges in the ventricle boundaries vicinity. Errors occur if such a neighboring edge has a higher gradient than the ventricle edge and if it meets the two additional constraints listed above. In order to improve the discrimination of the boundary search algorithm, we propose to use both notions of region and boundary. Region-based image segmentation has been proposed in (Ronfard, 1994) and has proven to be more powerful than gradient-based approaches. We first extract the intensity profile of a vertex in the filtered image. Prior anisotropic diffusion ensures rather homogeneous regions in the image. The intensity profile is scanned from the inside to the outside to find a set of consecutive voxels within a given intensity range and of length greater than a minimum threshold. Once the beginning of a region has been found, we search for the first voxel value that does not belong to the region intensity range thus corresponding to the end of the region. Given that such a voxel has been found we then look for a voxel whose gradient norm is above a given threshold in the vicinity of the region end. This algorithm requires to set appropriate values for the region intensity range and the gradient

norm threshold parameters, based on the segmented image properties.

6. Experimental results

6.1. Gradient computation

This experiment compares three cardiac left ventricle reconstructions from a 3D echocardiographic image using two different expressions of the external force. The first reconstruction relies on the gradient of the gradient norm as the external force, thus following the approach originally introduced by Kass et al. (1988). The second and third reconstructions relies on gradient extremum and region-based external forces as introduced in Section 5. The left ventricle model is initialized as an ellipsoid roughly centered inside the ventricle. Fig. 7 shows the reconstructed surfaces (top row) and the intersection of each surface with one of the image planes (bottom row). The template image quality is high (4° of arc resolution) and gradients are compute using a 7^3 cylindrical filter. The gradient of the gradient norm (left) is not reliable enough due to the large amount of speckle which introduces too many false positives. The scan-line algorithm (middle) provides better results for both the gradient extremum and the inner dark regions search. The region approach (right) better smooths the boundary and is more accurate as can be clearly seen in the mitral valve area (lower part of image).

6.2. Echocardiographic images

The temporal evolution of the cardiac left ventricle is extracted from a series of eight 3D echocardiographic images covering the heart systole
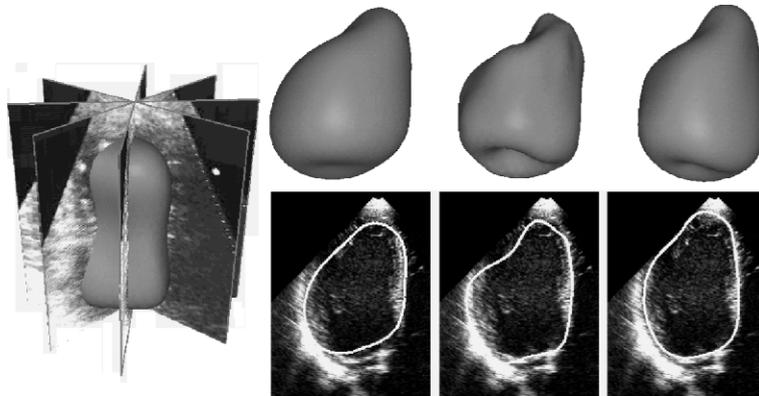
Fig. 7. Gradient of the gradient norm (left), gradient extremum (center), and region-based (right) reconstruction.

(contraction). The original images quality is poor due to a low resolution (20° of arc between planes) speckle, video signal digitization artifacts, and occlusions due to the reflection of the ultrasound beam by ribs. A region-based approach is used for computing the external image attraction term. In order to study the benefit of anisotropic diffusion, the same 4D image is segmented both with and without prior diffusion.

The model built in the previous experiment is used as a template for the segmentation of the first volume in time. At each following time step, the model is initialized from its position at the previous time instant. The same deformation parameters have been used for original and diffused images segmentation. Two deformation stages occur for each time instant: a global rescaling of the model using a global affine constraint ($\beta = 1$, $\lambda = 10\%$, scan-line range $r = 10$ voxels) followed by a local deformation stage that compensates for non-rigid deformations ($\beta = 0.2$, $\lambda = 70\%$, $r = 5$ voxels).

Fig. 8 displays the surface model extracted at all 8 time instants in the native (top row) and diffused image (bottom row). Even though they appear to be similar, the graph shows the model volume variations through time for the original (dashed line) and diffused (solid line) images. For detailed view of the differences, Fig. 9 shows the intersection of the models extracted from the original data onto slices of the original image, while Fig. 10 shows the models extracted from the diffused data. The horizontal axis represents time (all 8 time instants of the sequence). The vertical axis is the rotation angle of the displayed slice (from top to bottom: 0°, 20°, 40°, 60°, 80°, 100°, 120°, 140°, 160° of arc).

A last segmentation example on a complete heart cycle sequence is shown in Fig. 11. The sequence has a high resolution: 4° of arc angle
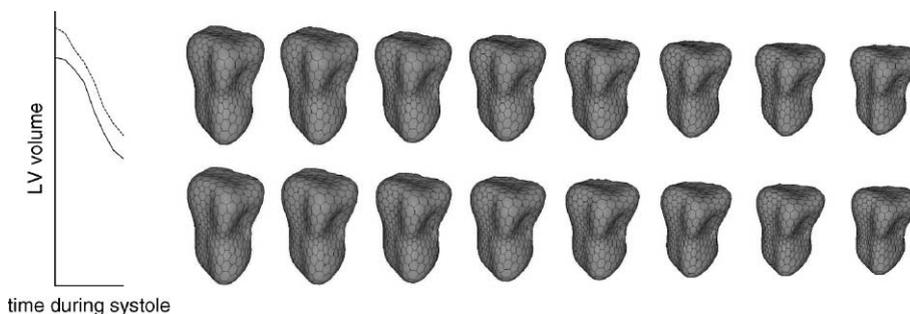
Fig. 8. Left: volume variation extracted from the original (- - -) and diffused (—) images. Right: surface models extracted from the original (top row) and diffused (bottom row) images.
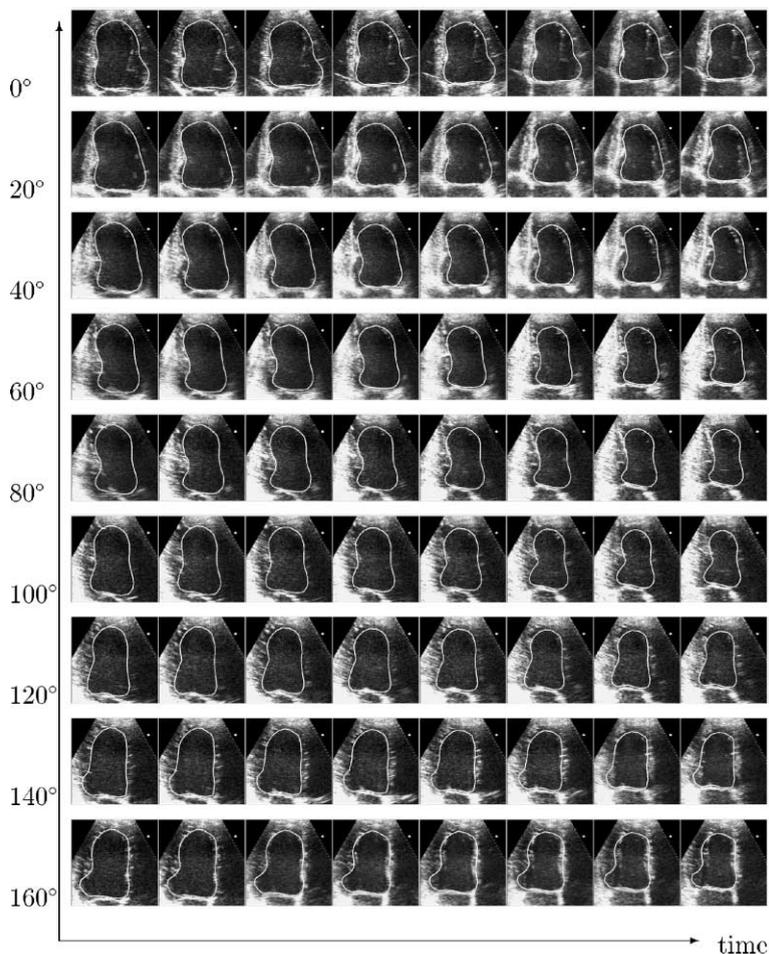
Fig. 9. Intersection of the deformed models with the original image.

between slices, 22 time instants. There is a visible breathing artifact. Fig. 11 shows the volume variation curve and the reconstructed model superimposed on one slice of the sequence.

7. Discussion

The anisotropic diffusion procedure improves both the gradient computation and the extraction of homogeneous regions along the model normal directions. As a consequence, the model better fits the visible boundaries of the left ventricle in echocardiographic images. For the two segmentations shown above, the same parameters of motion have been used.

The region-based force defines the model boundary as a bright region whose length is at least 5 voxels. However, due to the intensity smoothing caused by the diffusion process, the intensity threshold defining ''bright regions'' is different in the original and the diffused data. The definition of the ventricle boundaries is sensitive to this threshold and this results in a displacement of the apparent boundaries that explains the difference in volume estimation between the original and diffused data (see left of Fig. 8). In the absence of ground truth segmentation, it is very difficult to set the region threshold so that the reconstructed surface corresponds to the physical boundaries of the heart ventricle and that the computed volume corresponds to the actual volume.
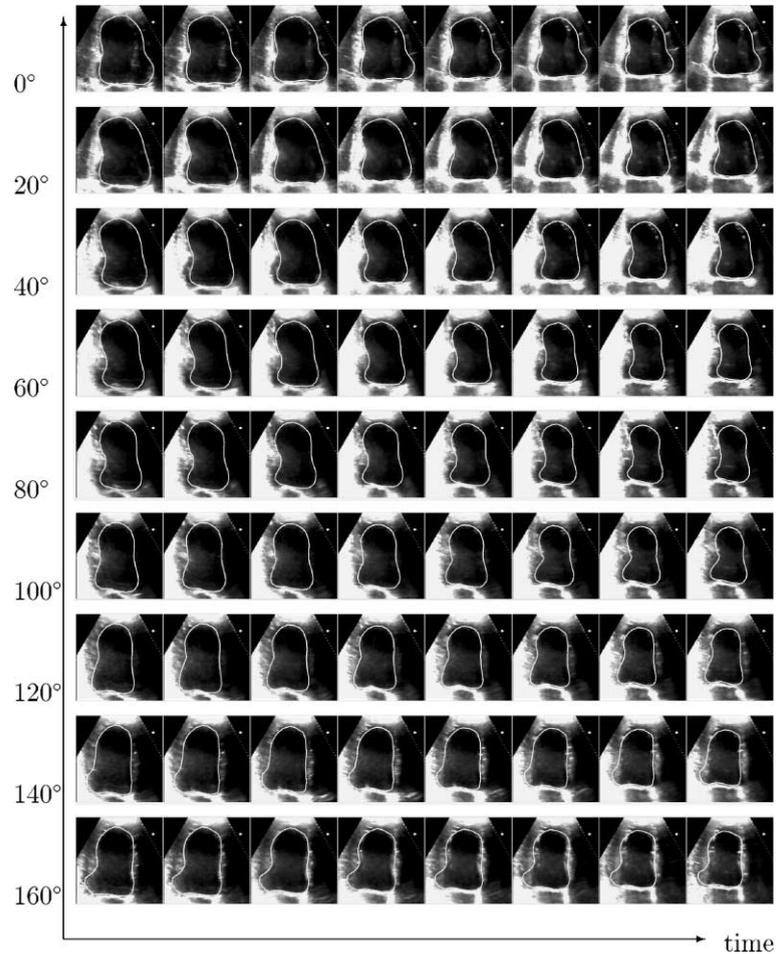
Fig. 10. Intersection of the deformed models with the diffused data.
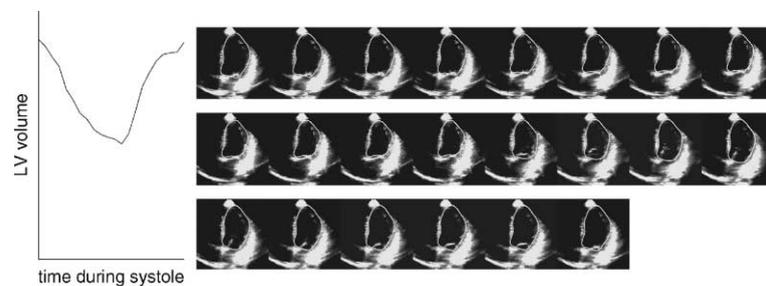
Fig. 11. Left: model volume variation over a complete heart cycle. Right: intersection of a the surface with one of the sequence slices through time.

However, the computation of the heart left ventricle ejection fraction if far less sensitive to the intensity threshold since it is a ratio parameter.

The ejection fraction ef measures the fraction of blood ejected by the ventricle at each heart cycle: $ef = (v_{max} - v_{min})/v_{max}$ where v_{max} and v_{min} are the

maximal and minimal ventricle volume respectively. In the first experiment shown above, we measured ejection fraction values of 41% in the original data and 44% in the diffused data. This values compare to the 45% ejection fraction estimated from hand-made segmentation of the image by a cardiologist. These results are very satisfying since the inter-operator variability in the computation of the ejection fraction on that kind of images is in the order of a few percents.

The model shape constraints regularize the surface deformation. This is visible in the second experiment (Fig. 11) where shape constraints attenuate the effect of the breathing artifact on the surface reconstruction.

8. Conclusion

In this paper we have demonstrated that segmentation of cylindrical echocardiographic images can be performed using deformable surfaces. The model provides enough intrinsic (shape) and extrinsic (grey-level range) prior knowledge on the data, to constrain the deformations properly even in the presence of very sparse and noisy data. We are now focusing on introducing more prior knowledge in the deformation process using temporal continuity in the deformation process (Montagnat and Delingette, 2000).

Prior filtering of the image reduces the speckle level and improves the region detection. The 4D anisotropic diffusion preprocessing improves spatial and temporal consistency of the time series of images. The image geometry is also taken into account for a better detection of structure boundaries. The model-based segmentation enables an accurate computation of the left ventricle ejection fraction when compared to hand-made segmentation by an expert.

Acknowledgements

References

Alvarez, L., Lions, P.-L., Morel, J.-M., 1992. Image selective smoothing and edge detection by nonlinear diffusion. SIAM Numer. Anal. 29 (3), 845–866.

Bardinet, E., Cohen, L., Ayache, N., 1996. Tracking and motion analysis of the left ventricle with deformable superquadrics. Med. Image Analy. 1 (2), 129–149.

Berger, M.-O., Winterfeldt, G., Lethor, J.-P., 1999. Contour tracking in echocardiographic sequences without learning stage: application to the 3D reconstruction of the 3D beating left ventricle. In: Med. Image Computing and Computer-Assisted Intervention (MICCAI'99). Lecture Notes on Computer Science, Vol. 1679. Springer, Cambridge, UK, pp. 508–515.

Besl, P., McKay, N., 1992. A method for registration of 3D shapes. IEEE Trans. Pattern Anal. Machine Intell. 14 (2), 239–256.

Boukerroui, D., Basset, O., Baskurt, A., Noble, A., 1999. Segmentation of echocardiographic data. Multiresolution 2D and 3D algorithm based on grey level statistics. In: Med. Image Computing and Computer-Assisted Intervention (MICCAI'99). Lecture Notes on Computer Science, Vol. 1679. Springer, Cambridge, UK, pp. 516–523.

Brandt, E., Wigström, L., Wranne, B., 1999. Segmentation of echocardiograhic image sequences using spatio-temporal information. In: Med. Image Computing and Computer-Assisted Intervention (MICCAI'99). Lecture Notes on Computer Science, Vol. 1679. Springer, Cambridge, UK, pp. 410–419.

Clarysse, P., Friboulet, D., Magnin, I., 1997. Tracking geometrical descriptors on 3-D deformable surfaces: application to the left-ventricular surface of the heart. IEEE Trans. Med. Imaging 16 (4), 392–404.

Cootes, T., Taylor, C., Cooper, D., Graham, J., 1995. Active shape models, their training and application. Comput. Vision Image Understanding 61 (l), 38–59.

Czerwinski, R.N., 1998. Line and boundary detection in speckle images. IEEE Trans. Image Process. 7 (12), 1700–1714.

Declerck, J., Feldmar, J., Ayache, N., 1998. Definition of a 4D continuous planispheric transformation for the tracking and the analysis of LV motion. Med. Image Anal. 2 (2), 197–213.

Delingette, H., 1999. General object reconstruction based on simplex meshes. Intl. J. Comput. Vision 32 (2), 111–146.

Deriche, R., 1987. Using Carmy's criteria to derive a recursively implemented optimal edge detector. Int. J. Comput. Vision 1 (2), 167–187.

Elayadi, M., 1997. Équations aux dérivées partielles et réseaux de neurones pour le traitement d'images. Ph.D. Thesis, Université Joseph Fourier, Grenoble, France.

Gerig, G., Kübler, O., Kikinis, R., Jolesz, F., 1992. Nonlinear anisotropic filtering of MRI data. IEEE Trans. Med. Imaging 11 (2), 221–232.

Giachetti, A., 1998. On-line analysis of echocardiographic image sequences. Med. Image Anal. 2 (3), 261–284.

Guttman, M., Prince, J., McVeigh, E., 1994. Tag and contour detection in tagged MR images of the left ventricle. IEEE Trans Med. Imaging 13 (1), 74–88.

Herlin, I., Ayache, N., 1992. Feature extraction and analysis methods for sequences of ultrasound images. In: Second European Conference on Computer Vision, pp. 43–57.

Jacob, G., Noble, A., Mulet-Parada, M., Blake, A., 1999. Evaluating a robust contour tracker on echocardiographic sequences. Med. Ira. Anal. 3 (1), 63–75.

Kass, M., Witkin, A., Terzopoulos, D., 1988. Snakes: active contour models. Int. J. Comput. Vision 1, 321–331.

Kimia, B., Tannenbaum, A., Zucker, S., 1992. On the evolution of curves via a function of curvature i. The classical case. J. Math. Anal. Appl. 163, 438–458.

McEachen, J., Duncan, J., 1997. Shaped-base tracking of left ventricular wall motion. IEEE Trans. Med. Imaging 16 (3), 270–283.

McInerney, T., Terzopoulos, D., 1995. A finite element model for 3D shape reconstruction and nonrigid motion tracking. In International Conference on Computer Vision, Virtual Reality and Robotics in Medicine (CVRMed'95). Lecture Notes on Computer Science, Vol. 905, Springer-Verlag, pp. 518–523.

McInerney, T., Terzopoulos, D., 1996. Deformable models in medical image analysis: a survey. Med. Image Anal. 1 (2), 91–108.

Monga, O., Deriche, R., Malandain, G., Cocquerez, J.-P., 1991. Recursive filtering and edge tracking: two primary tools for 3d edge detection. Image Vision Comput. 9 (4), 203–214.

Montagnat, J., Delingette, H., 1998. Globally constrained deformable models for 3D object reconstruction. Signal Process. 71 (2), 173–186.

Montagnat, J., Delingette, H., 2000. Space and time shape constrained deformable surfaces for 4D medical image segmentation. In: Med. Image Computing and Computer-Assisted Intervention (MICCAI'00). Lecture Notes on Computer Science, Vol. 1935. Springer, Pittsburgh, USA, pp. 196–205.

Nastar, C., Ayache, N., 1996. Frequency-based nonrigid motion analysis: application to four dimensional medical images. IEEE Trans. Pattern Anal. Machine Intell. 18 (11), 1067–1079.

Papademetris, X., Sinuas, A., Dione, D., Duncan, J., 2001. Estimation of 3D left ventricular deformation from echocardiography. Med. Image Anal. 5 (1), 17–28.

Park, J., Metaxas, D., Axel, L., 1996. Analysis of left ventricular motion based on volumetric deformable models and MRI-SPAMM. Med. Image Anal. 1 (1), 53–71.

Perona, P., Malik, J., 1990. Scale-space and edge detection using anisotropic diffusion. IEEE Trans. Pattern Anal. Machine Intell. 12, 629–639.

Reynard, D., Blake, A., Azzawi, A., Styles, P., Radda, G., 1995. Computer tracking of tagged [1]HMR images for motion analysis. In: International Conference on Computer Vision, Virtual Reality and Robotics in Medicine (CVR-Med'95). Lecture Notes on Computer Science, Vol. 905. Springer-Verlag, pp. 272–276.

Rohling, R., Gee, A., Berman, L., 1997. Three-dimensional spatial compounding of ultrasound images. Med. Image Anal. 1 (3), 177–193.

Ronfard, R., 1994. Region-based strategies for active contour models. Int. J. Comput. Vision 13 (2), 229–251.

Sanchez-Ortiz, G., Declerck, J., Mulet-Parada, M., Noble, J., 2000. Automating 3D echocardiographic image analysis. In: Med. Image Computing and Computer-Assisted Intervention (MICCAI'00). Lecture Notes on Computer Science, Vol. 1935. Springer, Pittsburgh, USA, pp. 687–696.

Sanchez-Ortiz, G.I., Rueckert, D., Borger, P., 1999. Knowledge-based tensor anisotropic diffusion of cardiac magnetic resonance images. Med. Image Anal. 3 (1), 77–101.

Schudy, R., Ballard, D., 1979. A computer model for extracting moving heart surfaces from four-dimensional cardiac ultrasound images. In: International Conference on Computer Vision (ICCV'79), pp. 366–376.

Stetten, G., Pizer, S., 2000. Medial-guided fuzzy segmentation. In: Med Image Computing and Computer-Assisted Intervention (MICCAI'00). Lecture Notes on Computer Science, Vol. 1935. Springer, Pittsburgh, USA, pp. 226–235.

Treece, G., Prager, R., Gee, A., Berman, L., 1998. Fast surface and volume estimation from non-parallel cross-sections, for freehand 3-D ultrasound. Technical Report 326, Engineering Department, Cambridge University.

Weickert, J., 1998. Anisotropic diffusion in image processing. Teubner Verlag.

Winterfeldt, G., Berger, M., Lethor, J., Handschuhmacher, M., 1997. Expert model based 3D reconstruction of the left ventricle using transtho-rasic echographic images. In: Computers in Cardiology, pp. 89–99.

Young, A., Kraitchman, D., Dougherty, L., Axel, L., 1995. Tracking and finite element analysis of stripe deformation in magnetic resonance tagging. IEEE Trans. Med. Imaging 14 (3), 413–421.

Zhang, Z., 1994. Iterative point matching for registration of free-form curves and surfaces. Intl. J. Comput. Vision 13 (2), 119–152.