


FIMH 2009 Overview



Nice June 3-5th 2009



Elsa Angelini

FIMH 2009

- Organization: N. Ayache, H. Delingette, M. Sermesant (INRIA, ASCLEPIOS).
- 5th edition.
- 150 attendees.
- Next edition: in New York – May 30th 2011 (organized by D. Metaxas & L. Axel).

Program overview

- Oral Session : (25 papers)
- Poster Session : (29 papers)
- Invited talks : 4

Keynote speakers



Prof. Reza Razavi

Professor of Paediatric Cardiovascular Science, Head of Division of Imaging Sciences, School of Medicine, King's College London, UK

"Translating biophysical models into the heart of the clinic"



Prof. Terry Peters

Professor, Medical Imaging, University of Western Ontario, and Scientist, Robarts Research Institute, London, Canada

"Virtual Environments to Guide Cardiac Interventions"

[--> Abstract](#)



Prof. Robert Howe

Director of Harvard Biorobotics Laboratory, USA

"Fixing the Beating Heart: Ultrasound Guidance for Robotic Intracardiac Surgery"

[--> Abstract](#)



Prof. Andrew McCulloch

Director of Cardiac Mechanics Research Group, UC San Diego, USA

"Multi-scale modeling of excitation-contraction coupling in the normal and failing heart"

[--> Abstract](#)

Themes


- Cardiac Imaging & Electrophysiology
- Cardiac Motion Estimation
- Cardiac Mechanics
- Cardiac Electrophysiology
- Cardiac Image Analysis
- Cardiac Anatomical and Functional Imaging

- Cardiac Research Platforms

A Statistical Approach for Detecting Tubular Structures in Myocardial Infarct Scars

Camille Vidal, Hiroshi Ashikaga and Elliot R. McVeigh, JHU

- The presence of an infarct scar in the heart generates abnormal electrical pathways that may trigger the occurrence of arrhythmic episodes.
- The infarct scar is perforated by tunnels of live tissue, which could generate abnormal activation pathways and therefore facilitate arrhythmia episodes.
- High res. MRI.
- Therefore, we propose to build low-level descriptors of the infarct scar to study this correlation more precisely and identify morphological markers of risk.



2 Tunnel Detector

2.1 Statistical Formulation

We model the image intensity by a real-valued random vector x , such that $x(s)$ represents the intensity value at voxel s . At each voxel of the image, we associate a discrete random variable $Y(s)$ that is equal to 0 if there are no tunnels in s and 1 if otherwise. $Y(s)$ is obtained using the Maximum A Posteriori (MAP) estimator at each pixel:

$$\forall s, \hat{Y}(s) = \arg \max_{j \in \{0,1\}} p(Y(s) = j|x), \quad (1)$$

Equivalently, we determine the presence or absence of a tunnel as follows:

$$\hat{Y}(s) = \begin{cases} 1 & \text{if } \mathcal{C}(s) > 1 \\ 0 & \text{otherwise} \end{cases}, \text{ where } \mathcal{C}(s) = \frac{p(Y(s) = 1|x)}{p(Y(s) = 0|x)},$$

Using Bayes' rule:


$$\mathcal{C}(s) = \frac{p(x|Y(s) = 1)p(Y(s) = 1)}{p(x|Y(s) = 0)p(Y(s) = 0)}. \quad (2)$$

We assume that there exists at most one tunnel per voxel, which means that the event $\{Y(s) = 1\}$ can be written as the union of independent events $\{T(s) = 1\} \cup \{T(s) = 2\} \cup \dots \cup \{T(s) = N\}$, where the discrete random variable T represents the type of tunnel to which voxel s belongs. We characterize a tunnel by its radius and its orientation and consider a finite set of N pairs of radii and orientations. To simplify the notations, we denote the absence of tunnel at voxel s by $T(s) = 0$, instead of $Y(s) = 0$. Equation (2) becomes:

$$\mathcal{C}(s) = \sum_{j=1}^N \frac{p(x|T(s) = j)p(T(s) = j)}{p(x|T(s) = 0)p(T(s) = 0)} = \sum_{j=1}^N \mathcal{C}(s, j) \frac{p(T(s) = j)}{p(T(s) = 0)}. \quad (3)$$

The tunnel type is estimated as follows:

$$\hat{T}(s) = \begin{cases} 0, & \text{if } \forall j, \mathcal{C}(s, j) \leq 1 \\ j, & \text{otherwise} \end{cases}, \text{ with } j = \arg \max_{j \in \{1, N\}} \mathcal{C}(s, j). \quad (4)$$



2.2 Tissue-Based Image Modeling


We assume that the voxel intensities are independent given the value of $T(s)$, hence the conditional probability can be written as a product (5). In (6), we introduce a segmentation label $z(s)$ at each voxel. $z(s) = 1$ if voxel s belongs to the infarcted tissue and 0 otherwise. The ratio becomes:

$$\mathcal{C}(s, j) = \frac{\prod_t p(x(t)|T(s) = j)}{\prod_t p(x(t)|T(s) = 0)}, \quad (5)$$

$$= \frac{\prod_t \sum_i p(x(t)|z(t) = i, T(s) = j) p(z(t) = i|T(s) = j)}{\prod_t \sum_i p(x(t)|z(t) = i, T(s) = 0) p(z(t) = i|T(s) = 0)}. \quad (6)$$

$$\mathcal{C}(s, j) = \frac{\prod_t \sum_i p(x(t)|z(t) = i) p(z(t) = i|T(s) = j)}{\prod_t \sum_i p(x(t)|z(t) = i) p(z(t) = i|T(s) = 0)}. \quad (7)$$

The distribution $p(x(t)|z(t) = i)$ represents the intensity distribution of tissue type i , while the geometry is encoded by the conditional distribution of the tissue type given the tunnel type, $p(z(t) = i|T(s) = j)$. The independence hypothesis allows us to model the intensity and the geometry separately.



2.3 Intensity Model

The computation of the ratio in (7) is quite intensive. It involves the computation at each voxel of a product of sums over all tissue types and all surrounding voxels. In order to simplify the computation of the ratio, we approximate the distribution $p(x(t)|z(t) = i)$ by its mode:

$$\forall t, i^*(t) = \arg \max_{i \in \{0,1\}} p(x(t)|z(t) = i),$$

$$\text{such that } p(x(t)|z(t) = i) = \delta(i^*(t) - i)p(x(t)|z(t) = i), \quad (8)$$

ments the image based on intensity. Each voxel is assigned a label: normal or infarcted tissue. In our application we model the intensity distribution as a mixture of Gaussian distributions. The segmentation is therefore obtained by the EM algorithm [17]. However, any intensity-based segmentation technique would be suitable.

1. Segment the image: $\forall s, i^*(s) \leftarrow \arg \max_i p(x(s)|z(s) = i)$,
2. For each tunnel type j , compute the matching function $\mathcal{C}(s, j)$, cf. (14),
3. At each voxel, determine the type of tunnel using decision rule (4),
4. Reconstruct the detected tunnels, based on their detected centers, orientations and radii.

by f_j the 3D filter of radius R_j , orientation α_j , and length h_j , such that:

$$\forall t, f_j(s, t) = \begin{cases} 0, & \text{if } t \in \text{Cyl}(s, R_j, h_j, \alpha_j), \\ 1, & \text{if } t \notin \text{Cyl}(s, R_j, h_j, \alpha_j), \end{cases} \quad (10)$$

$$\begin{aligned} \log p(z(t) = i^*(t)|T(s) = j) \\ = (1 - [i^*(t) - f_j(s, t)]^2) \log(1 - \epsilon_j(s, t)) + [i^*(t) - f_j(s, t)]^2 \log \epsilon_j(s, t), \end{aligned} \quad (11)$$

where ϵ_j controls the amount of tolerated mismatch between the image segmentation and the filter:

$$\forall t, \epsilon_j(s, t) = \begin{cases} 1/2, & \text{if } t \notin \text{Cyl}(s, \sqrt{2}R_j, h_j, \alpha_j), \\ \text{some constant} \ll 1/2, & \text{if } t \in \text{Cyl}(s, \sqrt{2}R_j, h_j, \alpha_j). \end{cases} \quad (12)$$




Fig. 2. Left: 1D profile of a tissue, Middle: 2D representation of the filter showing a cross-section of the cylinder, Right: 3D cylindrical filter

Remark: In order to filter an image, which is defined on a grid, we need to discretize the cylindrical filters (10). Since we are interested in structures of small radii (1 to 2 voxels), the discretization generates artifacts. To reduce this effect, we first build the filters at a finer resolution and downsample them, such that mixed voxels appear at the boundary of the filter, simulating the partial volume effect. In practice, the computation is unchanged, except that voxels at the tunnel boundary have a lower weight in the matching function.

enhancement MR images. The data is acquired as follow: in a male pig, the mid-left anterior descending coronary artery is occluded for 150 minutes using a balloon angioplasty catheter to create myocardium infarction. After 10 to 12 weeks, ex vivo high resolution MR Images are acquired to visualize the geometry of the infarct using Gd-DTPA (0.20 mmol/kg), to specifically mark the infarcted tissue. The heart is scanned in a 1.5-T MR scanner with a 3D gradient echo sequence (bandwidth, 130 Hz/pixel; flip angle, 20 degrees; echo time/repetition time, 4.02/9.7 ms, field of view, 100×100 mm; image matrix, 256×256; spatial resolution, 0.39×0.39×0.39 mm).

The detection proposed detection algorithm is used to detect tunnels of live myocardium in the infarct scar. We build a set of 13 filters with radius 1.5 voxels and different orientations. The cylinder has a length of 7 voxels. Figure 4 presents


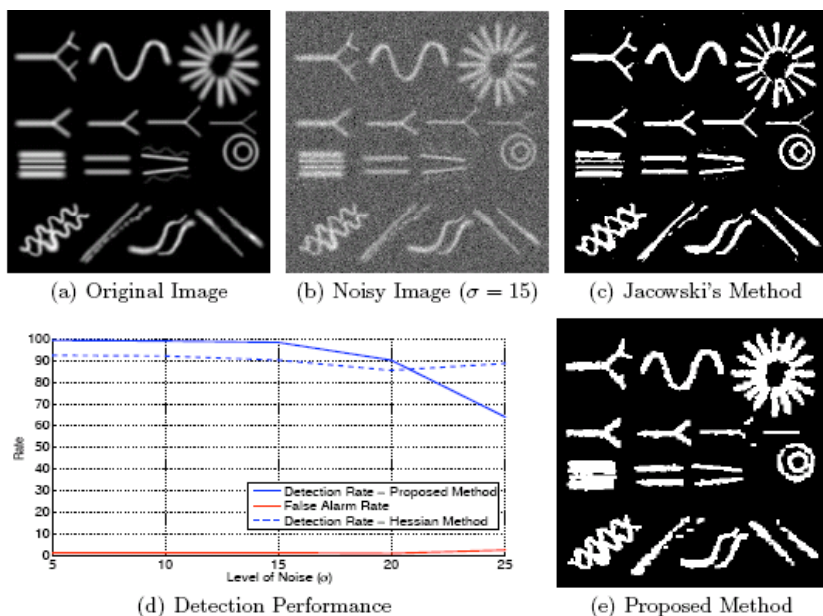


Fig. 5. 3D rendering of the infarct scar in red, superimposed on a short-axis view of the MR image. The detected tunnels are represented in blue/yellow/green depending on their detection score.

Quantitative Tool for the Assessment of Myocardial Perfusion during X-Ray Angiographic Procedures

Jean Liénard and Régis Vaillant , GE Healthcare

- During interventional procedures, cardiologists revascularize the ischemic heart by opening occluded arteries.
 1. Apply synchronized subtraction to enhance the tissue opacification by removing the background.
 2. compute time-density curves normalized with respect to the injection characteristics to quantify the microcirculation temporal behavior.
 3. Create parametric images that summarize in a single picture the dynamic information directly usable in the interventional setting.




Fig. 1. The raw image (*left*) shows the opacified coronary artery surrounded by chest and diaphragm (dark regions). As a result of synchronous subtraction (*right*), the background is removed and lets appear more clearly the perfused myocardium.




Fig. 4. The contrast wedge is a linearity insert (ref. 76-719 of Nuclear Associates); containing six iodinated discs, each with a different iodine thickness (0.5, 1, 2, 4, 10 and 20 mg/cm²). The phantom is acquired with two different angular positions, superimposed with different thickness PMMA block (2.5, 13.5, 24.5 cm) and images are subtracted (*left*) for assessing the linearity of the contrast signal (*right*).

of the patient. An example of TDC extracted from a subtracted coronary arteriogram is given in Fig. 3. Two main parameters differentiate the TDC: the peak opacification amplitude (POP) and the time to reach the peak value (TTP). Their values characterize the flow in the region-of-interest but they significantly rely on the way the contrast medium is injected into the arterial tree. Consequently, in order to limit the parameter variability due to difference in manual dye injection, an ostial ROI located at the catheter tip provides a TDC playing the role of reference in term of POP and TTP. We then normalize the amplitude of each TDC by division with the reference POP and we convert its time of maximum TTP in transit time (TT) by subtracting the reference TTP [10]. These two parameters, normalized amplitude and

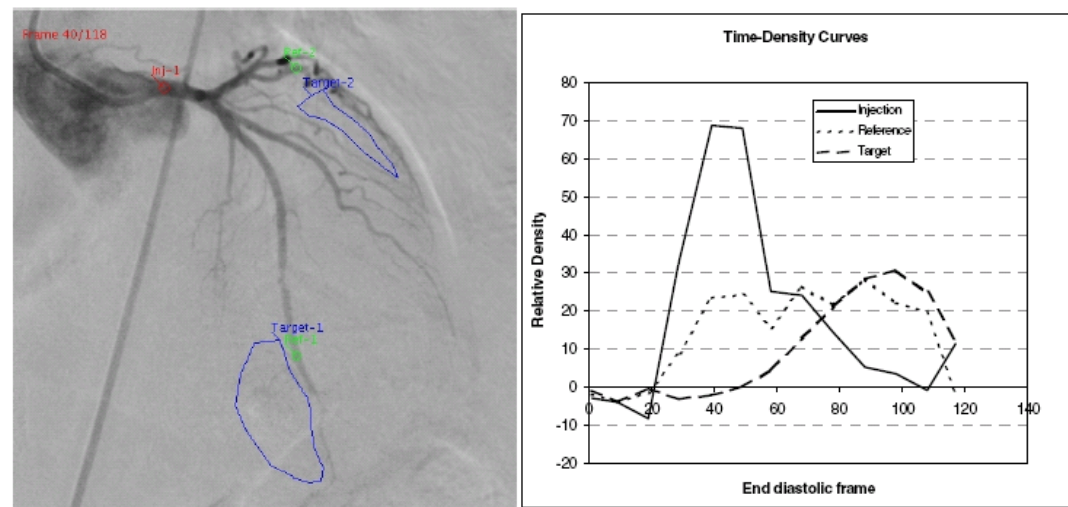


Fig. 3. Region-of-interest (ROI) are manually placed by the operator on the epicardial artery (ROI Inj-1 for Injection TDC) and on the myocardium (ROI Target-1 for Target TDC) (*left*). The corresponding densities are measured in consecutive end diastolic frames and plotted in Injection and Target Time-Density Curves (*right*).

the densitometric value in the successive diastolic frames and determines the POP and TTP parameters. In order to normalize the parameters like done with TDC, we designed an automatic detection algorithm to locate the arterial ostium where contrast injection occurs. This reference ROI is characterized by the earliest arrival of highly concentrated contrast medium in the half upper part of the field-of-view.




Fig. 4. Parametric image POP (*left*), TT (*center*), TMPG (*right*). The black square in the TMPG picture represents the automatic localization of the injection site. The densitometric color scales are relative. Each block in the TT color scale represents one cardiac cycle. Due to patient breathing at the end of the sequence, the diaphragm motion creates a semi-lunar artifact in the upper side of the images, without disturbing the reading of the apical myocardial area surrounding the LAD artery.

Local Cardiac Wall Motion Estimation from Retrospectively Gated CT Images

Jochen Peters, Olivier Ecabert, Holger Schmitt, Michael Grass and Jürgen Weese , Philips Research Aachen

- Automatic local tracking of the cardiac surfaces with sub-voxel accuracy in time series of CT data.
- 1. Multi-compartment mesh is adapted to a single reconstructed image.
- 2. Adapted mesh is then propagated through the whole time series.
- 3. Chamber volumes and local surface motion can be estimated at arbitrary time points.
- Experiments on 15 time series indicate that the system is very precise and robust with respect to the selection of the start phase of the propagation.
- The results can be used:
 - for dyssynchrony assessment
 - comparison with electromechanical models of the full heart.

2.1 General Segmentation Framework

To segment the cardiac surfaces we apply the shape-constrained deformable model adaptation framework of [9,4]. The endocardium of the four cardiac chambers, the left ventricular epicardium, and the truncated major vasculature (aorta, pulmonary artery and veins) are modelled by a triangulated mesh with $V = 7285$ vertices and $T = 14768$ triangles (see Fig. 1). Automatic segmentation starts by an approximate heart localization using the Generalized Hough Transform [10,4]. After positioning the mesh model, the surfaces are adapted to the corre-




Fig. 1. Cardiac mesh (seen from top right and bottom) and segmentation examples in axial CT slice for end-diastole and end-systole of the same patient. For the left ventricle, we model both the endo- and the epicardium [dark blue inside yellow]. All other chambers are modeled by endocardial walls. Vessels are arbitrarily truncated.

2.3 Improved Boundary Detection

Time series segmentation in the above framework with our *standard* cardiac model as published in [4] suffers from the limited accuracy of the boundary detection. Displacements between two successive images with phase increments of 2% range from 0.1–0.2mm up to 1–2mm (see also Sec. 3.2). To really track the cardiac surfaces for image volumes reconstructed at such small cardiac phase increments, sub-voxel accuracy is needed.

mesh surface. To localize the correct boundaries and to reject all (or most) misleading edges within the search range, we apply an edge detector and analyse additional image features surrounding all detected edges as detailed in [4,11]. In brief, the edge detector projects the image gradient ∇I onto the triangle normal (to suppress edges with false orientation). To encode the expected gray value transition (dark-to-bright or vice versa), the projected gradient is multiplied with ± 1 . To reject false boundaries, one or several image features such as gray values on both sides of the triangle are evaluated and compared to trained acceptance intervals. If any feature value violates its interval, the edge is rejected. Training, feature selection, and triangle-specific optimization are described in [11].

Our *standard* cardiac model had been developed in order to segment CT image volumes with sub-millimeter voxel extensions. On such images, we reach an average surface-to-surface error of 0.8mm [4]. This model uses a search profile with sampling points spaced at $\delta = 1\text{mm}$ and calculates the image gradient via tri-linear interpolation from the voxel data.

To increase the accuracy of the segmentation despite reduced image resolution, an *improved* model was trained. To detect boundaries with sub-voxel accuracy, a convolution with an approximated Gauss kernel ω is applied *instead* of the tri-linear interpolation scheme that is used by our standard model. We use a product of three \cos^2 functions extending over 3 voxels in each Cartesian direction. The image gradient is thus calculated by a convolution with the kernel gradient $\nabla\omega$:

$$\nabla[I * \omega](x) = [I * \nabla\omega](x) \quad (1)$$

This convolution is evaluated analytically at sub-voxel positions x : Each voxel's gray value is assigned to the complete voxel volume, i.e., we have a 3D gray value density (instead of gray values assigned to voxel centers). Convolution with the kernel gradient is performed by 3D integration. We do not apply a discretized convolution on the voxel grid as it is frequently done in image processing, e.g., for Sobel filters. Analytical 1D and 3D phantom studies with tissue transitions converted into a voxel image show that the gradient estimated via Eq. (1) reaches its maximum close to the true transition and is suitable to track moving edges with sub-voxel accuracy. Boundary detection in the *improved* model uses finer sampling with $\delta = 0.2\text{mm}$. To account for reduced image resolution in our test

2.4 Motion Estimation

images and the corresponding time points. To eliminate noise from the “raw trajectories” and to interpolate between segmented images, we establish local cubic fits over time. More specifically, within a time window of width $T\%$ of the cardiac cycle, a least squares fit of the form $x(t) = a + b \cdot t + c \cdot t^2 + d \cdot t^3$ is created through $x_{i,n}$. This fit and its time derivative yields a smoothed trajectory per triangle and the triangle velocity near the center of the time window. Including the cubic term $d \cdot t^3$ allows to approximate the trajectories for inflection points, e.g., at diastasis where motion (almost) stops and then continues along the same direction. These fits are established around densely placed evaluation times t and we use overlapping (or sliding) windows to define the contributing phases P_n and coordinates $x_{i,n}$ per fit. To “fade out” contributions of cardiac phases P_n that are close to the borders of the sliding window, the least-squares sum down-weights contributions from more distant P_n with a factor $\cos^2(\frac{P_n - P(t)}{T})$ with $P(t)$ being the phase corresponding to t .

3 Experiments

3.1 Database

Our database covers time series of multi-slice computed tomography angiography (CTA) images acquired from 15 patients with 64-slice scanners in two clinics. Contrast medium is applied for coronary artery diagnosis. Bloodpool contrast in the chambers varies from patient to patient, in particular for the right chambers, and partly also within each chamber (see Fig. 2). For some time series acquired with dose modulation, SNR changes over the cardiac cycle. Heart rates vary from < 60 to 90 beats per minute. Images are reconstructed at phase increments of 2% of the RR interval using a gating window width of approximately 15%. Image resolution in the axial plane is 0.86mm (using a 256×256 voxel grid) and slice thickness is 2mm.




Fig. 2. Contrast variations for different patients (images around mid-diastole at 76%)




Fig. 3. (a) Cardiac chamber volumes [in mm^3] and (b) triangle velocity magnitude averaged over the chamber surfaces [in $\text{mm}/\%RR$] for a patient with low heart rate (53bpm). (LA, RA, LV, RV = left and right atrium and ventricle; results obtained from improved model of Sec. 2.3).




Fig. 4. As Fig. 3 for a patient with high heart rate (90bpm). Note that motion at mid-diastole (70–80%) have vanished.




Fig. 5. As Fig. 3 for a patient with irregular cardiac rhythm ($> 80\text{bpm}$)

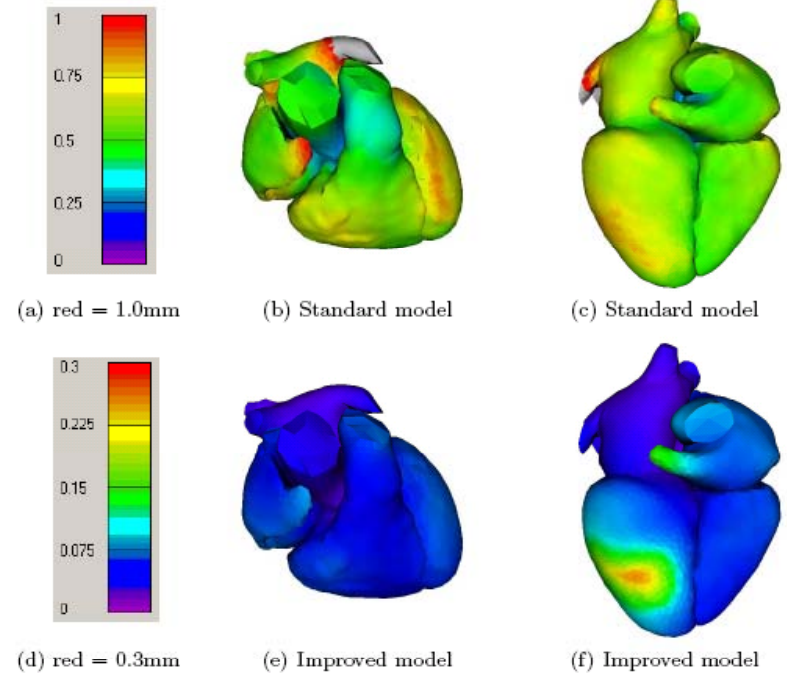



Fig. 7. Mean cyclic inconsistencies per triangle (averaged over 15 time series and 10 start phases per series). The color codes (a) for the standard model cover inconsistencies up to 1.0mm (larger values at the pulmonary veins are displayed in gray). For the improved model, color codes (d) represent a smaller interval. The region with increased inconsistencies at the bottom of the LV epicardium in (f) is due to one time series with confusing boundaries between LV epicardium, pericardium, and diaphragm (see text and last line in Tab. 1). For the improved model, mean inconsistencies at the LV endocardium are $< 0.1\text{mm}$ (not visible).

Physically-Constrained Diffeomorphic Demons for the Estimation of 3D Myocardium Strain from Cine-MRI

Tommaso Mansi, Jean-Marc Peyrat, Maxime Sermesant, Hervé Delingette, Julie Blanc, Younes Boudjemline and Nicholas Ayache , INRIA

- Image-based method to estimate myocardium strain from cine-MRI.
- To recover non-apparent cardiac motions, we improve the diffeomorphic demons NL registration algorithm, by adding two physical constraints.
- 1. First, myocardium near-incompressibility is ensured by constraining the deformations to be divergence free.
- 2. myocardium elasticity is modelled using smooth vector filters.
- The proposed physically-constrained demons are compared with the diffeomorphic demons and evaluated in a healthy subject against tagged MRI.



2 Methods

2.1 Non-linear Image Registration Using Diffeomorphic Demons (DD) Algorithm

Non-linear image registration aims at providing a transformation $\phi(\mathbf{x}) = \mathbf{x} + \mathbf{u}(\mathbf{x})$, where $\mathbf{u}(\mathbf{x})$ is the displacement vector field, that aligns a template image $T(\mathbf{x})$ with a reference image $R(\mathbf{x})$. The underlying principle lies in minimising an energy $\mathcal{E}(\phi)$ that comprises a similarity criterion $\mathcal{D}(T, R, \phi)$ and a regulariser $\mathcal{R}(\phi)$ ensuring smooth transformations and integrating prior knowledge.

Diffeomorphic demons (DD) algorithm [9] minimises $\mathcal{E}(\phi)$ iteratively through two decoupled steps. A correspondence field \mathbf{c} that matches the points between the two images, with a controlled uncertainty parametrised by σ_x^2 , is introduced in the energy: $\mathcal{E}(\mathbf{c}, \mathbf{u}) = 1/\sigma_i^2 \|R - T \circ \mathbf{c}\|^2 + 1/\sigma_x^2 \|\mathbf{u} - \mathbf{c}\|^2 + 1/\sigma_T^2 \|\nabla \mathbf{u}\|^2$. In this equation, σ_i^2 is a parameter related to the noise in the images and σ_T^2 compromises between the image data term and the regularisation.

First, the correspondence field \mathbf{c} is updated using the diffeomorphic update rule $\mathbf{c} \leftarrow \mathbf{u} \circ \exp(\mathbf{v})$. $\exp(\mathbf{v})$ is the exponential operator defined in the Log-Euclidian framework and \mathbf{v} is the smooth stationary *velocity* field that results from the minimisation of the correspondence energy:

$$\mathcal{E}_{corr}(\mathbf{v}) = \frac{1}{\sigma_i^2} \|R - T \circ \mathbf{u} \circ \exp(\mathbf{v})\|^2 + \frac{1}{\sigma_x^2} \|\mathbf{u} - \mathbf{u} \circ \exp(\mathbf{v})\|^2 \quad (1)$$

$$\mathcal{E}_{corr}(\mathbf{v}) = \frac{1}{\sigma_i^2} \|R - T \circ \mathbf{u} \circ \exp(\mathbf{v})\|^2 + \frac{1}{\sigma_x^2} \|\mathbf{u} - \mathbf{u} \circ \exp(\mathbf{v})\|^2 \quad (1)$$

Derivating \mathcal{E}_{corr} yields $\mathbf{v} = -(R - T \circ (\text{Id} + \mathbf{u})) / (\|J\|^2 + \sigma_i^2/\sigma_x^2) J$, with $J = (\nabla R + \nabla(T \circ (\text{Id} + \mathbf{u}))) / 2$ and $\sigma_i = |R - T \circ (\text{Id} + \mathbf{c})|$ (see [9] for further details).

Second, the displacement \mathbf{u} that satisfies the regulariser $\mathcal{R}(\mathbf{u})$, knowing the new correspondence field \mathbf{c} , is computed. The second part of the demons energy $\mathcal{E}(\mathbf{c}, \mathbf{u})$, $\mathcal{E}_{reg}(\mathbf{c}, \mathbf{u}) = 1/\sigma_x^2 \|\mathbf{u} - \mathbf{c}\|^2 + 1/\sigma_T^2 \|\nabla \mathbf{u}\|^2$, is minimised with respect to \mathbf{u} , which amounts to convoluting $\mathbf{c} = \mathbf{u} \circ \exp(\mathbf{v})$ with a Gaussian kernel.

2.2 Adding Physical Constraints to the Diffeomorphic Demons

Velocity Correction: Myocardium incompressibility can be interpreted as conservation of myocardium mass throughout the cardiac cycle [8]. Therefore, the mass continuity equation $\partial\rho/\partial t + \text{div}(\rho\mathbf{v}) = 0$ applies, where ρ is the mass density, \mathbf{v} is the velocity and t is the time. Assuming ρ constant and uniform, the conservation of mass becomes conservation of volume and writes as $\text{div}(\mathbf{v}) = 0$.

$$\hat{\mathbf{v}} \in \{\mathbf{f} \in L^2(\Omega) \mid \text{div } \mathbf{f} = 0\} \quad \min_{\mathbf{f} \in \{\mathbf{f} \in L^2(\Omega) \mid \text{div } \mathbf{f} = 0\}} \max_{p \in H_0^1} \left(\frac{1}{2} \|\mathbf{v} - \hat{\mathbf{v}}\|^2 + \int_{\Omega} p \text{ div } \hat{\mathbf{v}} \right)$$

$$\hat{\mathbf{v}} = \mathbf{v} - \text{grad } p \quad (3)$$

under the constraint $\text{div } \hat{\mathbf{v}} = 0$. Equation (3) is the closed-form expression of the divergence-free projector $\Pi(\mathbf{v}) = \mathbf{v} - \text{grad } p$, where p is obtained by solving the Poisson equation $\Delta p = \text{div}(\mathbf{v})$ under 0-Dirichlet boundary conditions [15].

Position Correction:

For each point \mathbf{x} of the image space, the Jacobian determinant $J(\mathbf{x}) = \det(\text{Id} + \nabla \mathbf{u}(\mathbf{x}))$ relates to the local volume change due to the transformation $\phi(\mathbf{x}) = \mathbf{x} + \mathbf{u}(\mathbf{x})$. Hence, the volume is locally preserved if the non-linear constraint $J(\mathbf{x}) = 1$ is satisfied. However, since myocardium is *nearly* incompressible (5% of volume variation is usually observed [6,8]) and the displacements provided by the DD are smooth and derive from divergence-free velocities, the first-order approximation of $J(\mathbf{x}) = 1$, i.e. $\text{div}(\mathbf{u}) = 0$, can be used to correct the remaining volume drifts. As a result, the positions are corrected like the velocities. At each iteration of the demons algorithm, the displacement field \mathbf{u} is made divergence-free by using the above-mentioned projector Π (Equation (3)). It is worth stressing that

Elastic-Like Regularisation. Myocardium is commonly assumed to be visco-elastic. However, at the time-scale of the heart beat, one can reasonably approximate the visco-elasticity property by elasticity [8]. In our approach, elasticity is modelled as regularisation. We substitute the diffusion-like regularisation, which is unable to model elastic deformations, by an elastic regulariser. Linear elastic

filters, while ensuring smooth deformations. Minimising $\mathcal{E}_{reg}(\mathbf{c}, \mathbf{u})$ with this regulariser amounts to convoluting the deformation field \mathbf{u} by the 3D isotropic separable filter:

$$G_{\sigma, \kappa}(\mathbf{u}) = \frac{1}{(\sigma\sqrt{2\pi})^3(1+\kappa)} \left(\text{Id} + \frac{\kappa}{\sigma^2} \mathbf{u}\mathbf{u}^T \right) \exp\left(\frac{\mathbf{u}^T \mathbf{u}}{2\sigma^2}\right) \quad (4)$$

which can be performed very efficiently by using recursive Gaussian filter [17]. In Equation (4), σ is the standard deviation of the Gaussian function and κ is a parameter that acts as a Poisson ratio coefficient. The higher is κ , the stiffer is the deformation.

Physically-Constrained Demons (PCD) Algorithm

- Choose an initial spatial transformation $\phi(\mathbf{x}) = \mathbf{x} + \mathbf{u}(\mathbf{x})$
 - Iterate until $\mathcal{D}(R, T; \mathbf{u})$ stops decreasing
 1. Given the current deformation \mathbf{u} , compute the update velocity field \mathbf{v} by minimising $\mathcal{E}_{corr}(\mathbf{v})$ (Equation (1)).
 2. Fluid-like regularisation: $\mathbf{v} \leftarrow K_{fluid} \star \mathbf{v}$. K_{fluid} is typically a Gaussian kernel with standard deviation σ_{fluid} .
 3. (*) Correct the velocity to be divergence-free: $\mathbf{v} \leftarrow \Pi(\mathbf{v})$.
 4. Update the correspondence field: $\mathbf{c} \leftarrow \mathbf{u} \circ \exp(\mathbf{v})$
 5. (*) Elastic-like regularisation: $\mathbf{u} \leftarrow G_{\sigma, \kappa} \star \mathbf{c}$ (Equation (4))
 6. (*) Correct the deformation to be divergence-free: $\mathbf{u} \leftarrow \Pi(\mathbf{u})$.
-

About the algorithm parameters. PCD algorithm is controlled by four parameters: the maximum step length [9], σ_{fluid} , σ and κ . In all our experiments, we have used standard values according to previous experiments [9,17]: the maximum step length is 1, σ_{fluid} is 0.5, σ is 2.0 and κ is 0.5.

composed transformation $\phi_{I_n \rightarrow I_{n-1}} \circ \phi_{I_{n-1} \rightarrow I_0}$. Finally, Lagrangian finite strain tensor $\mathbf{E} = 1/2 (\nabla \mathbf{u} + \nabla \mathbf{u}^T + \nabla \mathbf{u}^T \nabla \mathbf{u})$ is computed from the recovered deformation $\phi(\mathbf{x}) = \mathbf{x} + \mathbf{u}(\mathbf{x})$. Radial, circumferential and longitudinal strains are estimated in the heart prolate coordinate system [18].

both ventricles using 1.5T MR scanner (Achieva, Philips Medical System) (seven 10mm-thick slices; 1.37x1.37mm in-plane resolution, 30 frames). No longitudinal cine-MRI are available. Visual inspection of the images revealed no spatial mis-registration. To improve registration performances, the resolution of the cine-MR images is made isotropic at the in-plane resolution.

First, the incompressibility constraint is evaluated by computing the Jacobian determinant of the recovered myocardium deformation at each frame of the cardiac sequence. Without any constraint, mean Jacobian equals 1.09 ± 0.07 (mean \pm SD), maximum deviation from unity: 0.22. Adding the constraints yields $J(\mathbf{u}_{myo}) = 0.99 \pm 0.02$ (mean \pm SD) with a maximum deviation from unity of 0.05, which relates to the observed 5% of myocardium volume variation [8].

the global strains with minimum and maximum zone values. DD mainly recovers the visible radial motion of the heart, resulting in over-estimated radial strains, implausible positive longitudinal strains and weak circumferential strains. On the other hand, the physical constraints enable recovering non-apparent motions by reorienting the estimated myocardium deformations to ensure incompressibility. Radial strain is closer to clinical observations and more homogeneous across the myocardium [18], thus suggesting that the physical constraints also perform global regularisation. Longitudinal strain is clinically plausible yet slightly

under-estimated with respect to reported values, mainly because of the coarse through-plane resolution. Circumferential strain is improved but remains close to DD estimations. A probable reason is the low temporal resolution, circumferential motion occurring in few time frames.

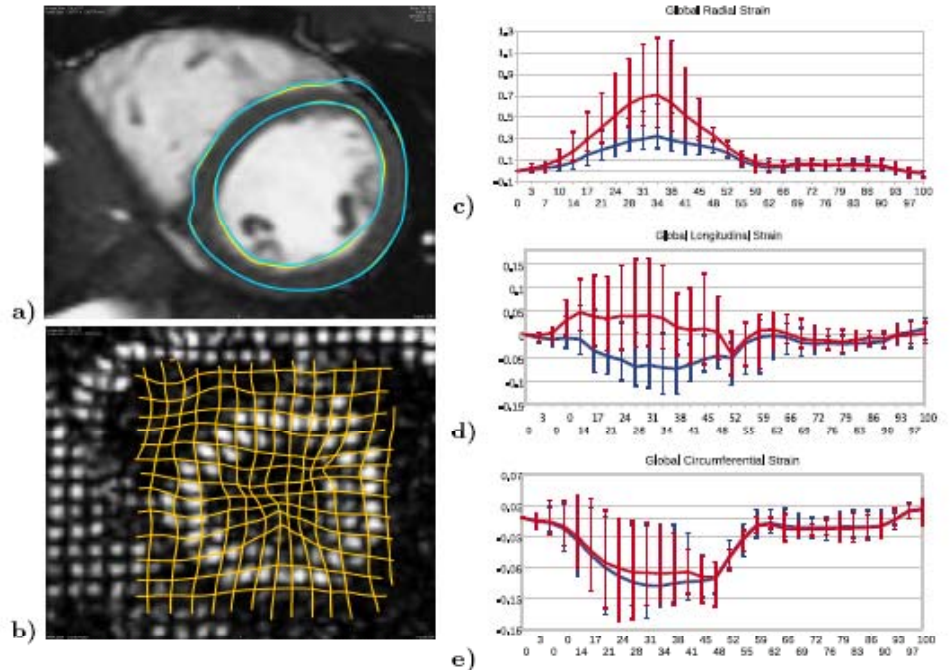


Fig. 1. a) End-systole myocardium mask propagated back to the end-diastole image using DD (in yellow) and PCD (in blue). Both algorithms provide similar registration performances. b) Tagged-MRI at end-systole and tag grid deformed by the cine-MRI transformation estimated using PCD. c-e) Global strain with maximum and minimum zone values. Time is normalised in %.

euHeartDB: A Web-Enabled Database for Geometrical Models of the Heart

Daniele Gianni, Steve McKeever, and Nic Smith

Computing Laboratory

Oxford University

Parks Road, OX1 3QD, Oxford, UK

{daniele.gianni, steve.mckeever, nic.smith}@comlab.ox.ac.uk

These non functional characteristics include:

- Decentralised control of database content (R8)
- Web accessibility (R9)
- Secure control access (R10)
- Semantic-based search of models (R11)
- Graphical methods (R12)
- Interoperable and standard XML data format for model exporting (R13)
- Support for popular model data formats (R14).

Acknowledgements

This work has been supported the European Commission (FP7-ICT-2007-224495: euHeart). The authors would like to thank the Auckland Bioengineering Institute, the ASCLEPIOS team, at INRIA Sophia-Antipolis, the CILAB, at Universitat Pompeu Fabra, and the X-ray Imaging Systems department of Philips Research Aachen for providing the first geometrical models uploaded in euHeartDB.

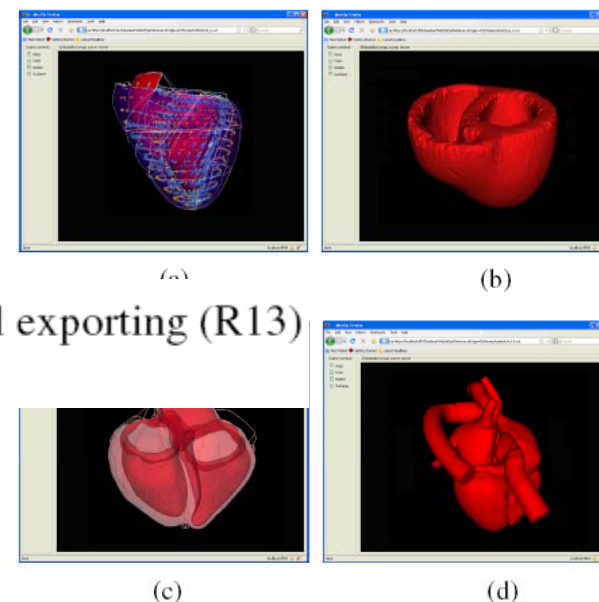


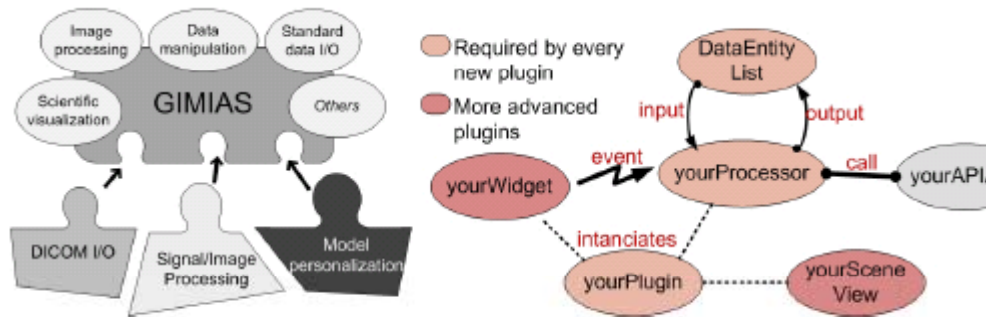
Fig. 2. euHeartDB Visualization of the Heart Models in [3] – (a), [7] – (b), [6] – (c), [27] – (d)

GIMIAS: An Open Source Framework for Efficient Development of Research Tools and Clinical Prototypes

Ignacio Larrabide^{1,2,*}, Pedro Omedas^{2,1}, Yves Martelli^{2,1}, Xavier Planes^{2,1},
Maarten Nieber^{2,1}, Juan A. Moya^{2,1}, Constantine Butakoff^{2,1},
Rafael Sebastián^{2,1}, Oscar Camara^{2,1}, Mathieu De Craene^{1,2},
Bart H. Bijnens^{3,2,1}, and Alejandro F. Frangi^{2,1,3}

Abstract. GIMIAS is a workflow-oriented environment for addressing advanced biomedical image computing and build personalized computational models, which is extensible through the development of application-specific plug-ins. In addition, GIMIAS provides an open source framework for efficient development of research and clinical software prototypes integrating contributions from the Virtual Physiological Human community while allowing business-friendly technology transfer and commercial product development. This framework has been fully developed in ANSIC++ on top of well known open source libraries like VTK, ITK and wxWidgets among others. Based on GIMIAS, in this paper is presented a workflow for medical image analysis and simulation of the heart.

Name	Distribution license	Brief list of features
BrainVISA	Open source software. CeCILL License.	- scientific visualization - signal and image processing
CMGUI	Open source software. Mozilla Public License.	- scientific visualization - scripting capabilities - model personalization - signal and image processing
MedINRIA	Free for non-commercial use.	- scientific visualization - model personalization - signal and image processing
MeVisLab	Basic version free for non-commercial use. SDK version requires a commercial license.	- scientific visualization - scripting capabilities - plug-in architecture - image processing
OpenMAF	Open source software. BSD-style License.	- generic application framework - scientific visualization - model personalization - image processing
ParaView	Open source software. BSD-style License.	- scientific visualization - scripting capabilities - plug-in architecture
Slicer 3D	Open source software. BSD-style License.	- scientific visualization - plug-in architecture - image processing



- **DataEntity:** This class stores and encapsulates a single data object (medical image, transformation, surface mesh, model simulation result, etc.).
- **DataHolder:** This class allows objects to register as observers of data instances, to be notified when the data is changed.
- **WorkingData:** This class holds a reference to the DataHolders associated to the data required and is currently used by the plug-in (selected in output, etc.).
- **DataEntityList:** This class contains a list populated with the data elements currently available for processing in the application. The input data that will be processed by the plug-in and also their outputs are stored in this list.
- **FrontEndPlugin:** This is the base class for all the plug-ins, every new plug-in must extend this class. This base class provides access to the rendering classes, but delegates the creation of the rendering scene to the external class' implementation.
- **Processor:** This class executes one algorithm and given one or more input data and generates one or more output data (usually coming from DataEntityList). A plug-in may contain more than one Processor class. No data is stored inside the Processor.

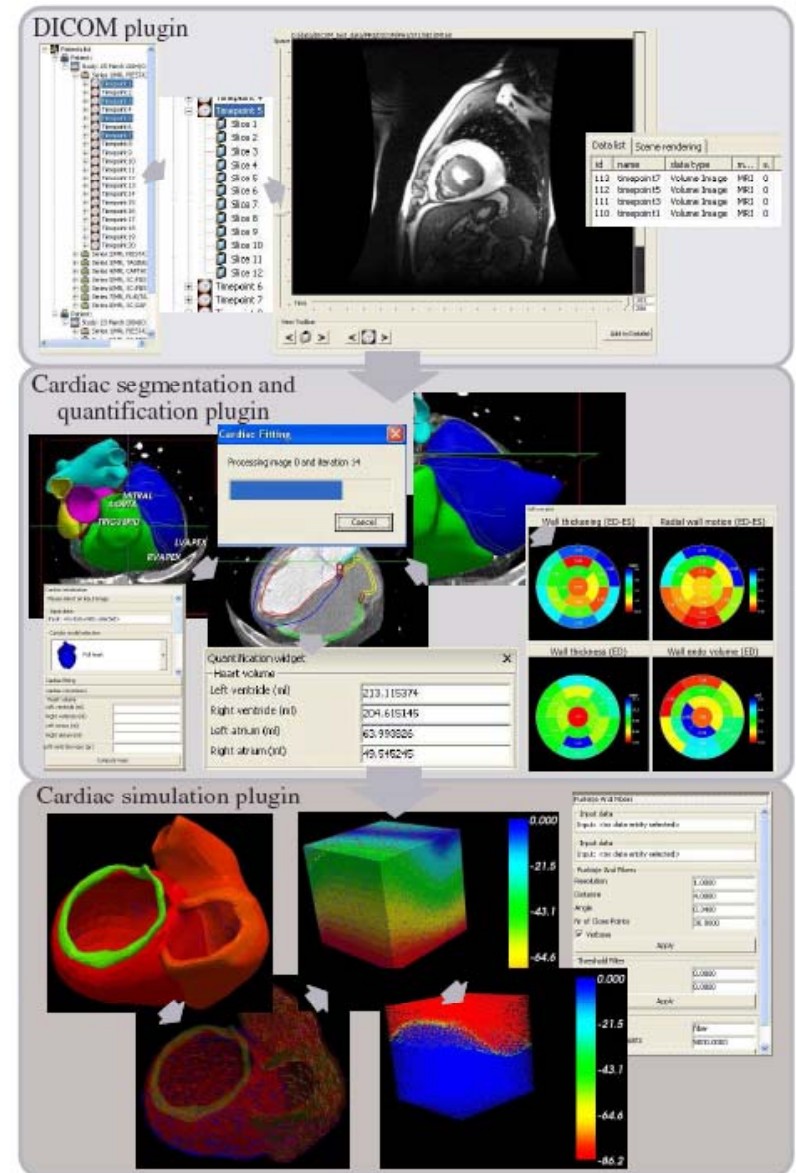


Fig. 2. Cardiac modeling workflow as implemented in GIMIAS. From top to bottom are presented the DICOM, the Cardiac segmentation and quantification, and the Cardiac simulation plug-ins. After desired images are selected from the time sequence, an atlas is used to segment the images and adjust the segmentation result where required. Then, from the resulting heart geometry, a model for simulation of the heart electrophysiology is created. Finally, the propagation front in a selected region of the cardiac muscle is presented.

Random Forest Classification for Automatic Delineation of Myocardium in Real-Time 3D Echocardiography

Victor Lempitsky, Michael Verhoek, J. Alison Noble and Andrew Blake,
Microsoft Research & Oxford Univ.

- Segmentation of real-time 3D US data.
- Machine learning approach treating this problem as a two-class 3D patch classification task.
- Solving such task using *random forests*, which are the discriminative classifiers developed recently
- *Computational time*: matter of seconds (on a CPU) or even in real-time (on a GPU) for the entire 3D volume.

Trees construction.

Random forests = collection of randomized binary trees.

Construction proceeds in a top-down fashion, starting from the root.

At each node:

- *a training collection of labeled voxels received from its parent is considered (the root node receives a randomly drawn subset of the size N_{train} of our full training set).*
- *A binary test is then chosen to minimize the uncertainty of the class labels in the subsets,*
- The test splits the training collection into (the exact choice method is discussed below).
- Once the binary test is chosen, the training collection of voxels is split in the two subsets according to this binary test.
- The two children nodes are then created and the two subsets are passed to them as their training collections, so that the recursion proceeds.

- The recursive tree construction stops either when a certain maximal tree depth D_{\max} is reached or when all voxels in the training collection have the same label.

Binary tests. Let us now discuss the way the binary tests are chosen. As the voxel description has two components, namely, the position in the volume and the neighborhood appearance, we consider two types of tests. The *appearance tests* are defined as:

$$t_{\text{appearance}}^{a_1, a_2, a_3, \tau_a}(\mathbf{x}, \mathbf{I}) = \begin{cases} 0, & \text{if } \sum_{p_1=x_1-a_1}^{x_1+a_1} \sum_{p_2=x_2-a_2}^{x_2+a_2} \sum_{p_3=x_3-a_3}^{x_3+a_3} \mathbf{I}(p_1, p_2, p_3) > \tau_a \\ 1, & \text{otherwise.} \end{cases} \quad (1)$$

An appearance test thus compares a sum of the intensities over an axis-aligned box of the size $(2a_1+1) \times (2a_2+1) \times (2a_3+1)$ centered at the voxels with some threshold τ_a , where the parameter values a_1, a_2, a_3, τ_a are chosen in a randomized fashion as discussed below. Note that such box sums can be evaluated extremely efficiently using just 7 additions and subtractions irrespective of the size of the box using the *integral image* technique [25].

The second type of the tests is the *position test* defined as:

$$t_{\text{position}}^{i, \tau_p}(\mathbf{x}) = \begin{cases} 0, & \text{if } x_i > \tau_p \\ 1, & \text{otherwise.} \end{cases} \quad (2)$$

This test, thus, simply compares the i th coordinate of the voxel with a threshold τ_p , where i, τ_p are chosen randomly as discussed below.

Given a collection of labeled voxels $\mathcal{C} = \{\mathcal{V}^i, c^i\}$, we define the *uncertainty* of its labels as:

$$U(\mathcal{C}) = |\mathcal{C}| \cdot \text{Entropy}(\{c^i\}) = |\mathcal{C}| (\bar{c} \cdot \log \bar{c} + (1 - \bar{c}) \cdot \log(1 - \bar{c})) \quad (3)$$

where \bar{c} denotes the average class label over collection. Thus, the uncertainty is maximal for the collections where all voxels are from the same class and minimal (for a given size) for the collections where class labels are split into equal halves.

When choosing a split for a training collection \mathcal{C} , we generate N_{tests} random tests. We choose two thirds of our tests to be the appearance tests and one third to be the position tests. When generating an appearance test, we sample a_1, a_2, a_3 independently and uniformly from the interval $[0; R]$ from zero to the neighborhood radius R . We then pick the threshold τ_a by averaging the respective box sums from (1) for the two randomly chosen samples from the training collection. When generating a position test, we pick one of the three dimensions with equal probability and we pick the threshold τ_p equal to the respective coordinate of a randomly chosen sample from the training collection.

For each of the generated tests, we consider the two subsets, it splits the training collection into. We then pick a test, which has the smallest sum of the uncertainties (3) for these subsets. Such a choice ensures that the uncertainty in labels decreases quickly towards the leaves of the tree.

There are two properties worth noting. Firstly, despite the extreme simplicity of the position tests the superposition of such tests at different depths of the tree allows to learn a non-trivial shape prior. Secondly, as the position and appearance tests are interleaved within a tree, we are able to learn the inter-dependencies between the shape and appearance.

Annotating a voxel using a random forest. Given a random forest, i.e. several trees that were independently trained as discussed above, we can classify a new voxel \mathcal{V} located at the position \mathbf{x} in a previously unseen echocardiogram \mathbf{I} . To do that we simply pass the voxel $\mathcal{V} = (\mathbf{x}, \mathbf{I})$ through every tree. Given a set of leaves the voxel ends up in different trees, we simply average the class probabilities recorded in these leaves.

and used them for validation, while the first eleven echocardiograms were used for training. As a result of the validation step, the values of parameters were as follows: $N_{\text{train}} = 100000$, $D_{\text{max}} = 16$, $N_{\text{tests}} = 30$, $R = 32$.

After the parameter tuning, we have evaluated our method on the first 6

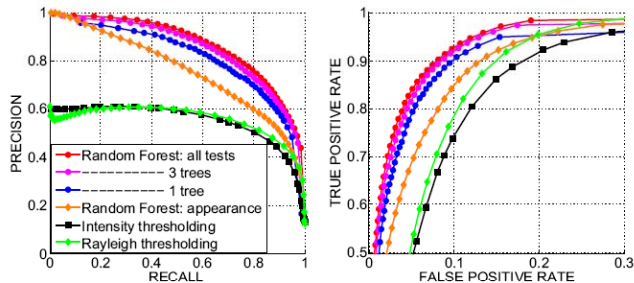



Fig. 2. Quantitative results. Recall-precision (left) and ROC (right) curves demonstrate the significant advantage of our method over the thresholding of raw intensity and of the shifted Rayleigh parameter. As an aside, removing the position tests (Random Forest-A curve) deteriorates the performance of our method much more than reducing the number of trees.



sample short-axis slices expert annotations random forest results