Fully Automated Endocardial Contour Detection in Time Sequences of Echocardiograms by Active Appearance Motion Models

JG Bosch[1], SC Mitchell[2], BPF Lelieveldt[1], F Nijland[3], O Kamp[3], M Sonka[2], JHC Reiber[1]

[1]Radiology, Leiden University Medical Center, Leiden, The Netherlands
[2]Electrical Engineering, University of Iowa, Iowa City, IA, USA
[3]Cardiology, Free University Medical Center, Amsterdam, The Netherlands

Abstract

*A novel fully automated border detection technique for phase-normalized echocardiographic image sequences is developed: Active Appearance-Motion Models (AAMM). AAMM finds shape and appearance eigenvariations of the heart over the full cardiac cycle from a set of examples, capturing typical motion patterns. AAMM segments sequences by adjusting eigenvariation coefficients to minimize model-to-target differences. This results in a time-continuous segmentation. The method was applied on 4-chamber sequences from 129 unselected patients, split randomly into training (TRN, n=65) and test set (TST, n=64). In all sequences, an independent expert manually drew endocardial contours (MAN). On TST, fully automated AAMM succeeded in 97% of cases (AUTO) and performed well (average contour distance 3.3 mm, area regression AUTO = 0.91*MAN + 1.7 cm^2, r = 0.87). Results outperformed single-frame AAM segmentation and human interobserver variabilities.*

1. Introduction

Robust automatic border detection in echocardiograms is highly desired for automated volume and wall motion analyses, particularly for stress echo, where the subjectivity of visual assessment of wall motion abnormalities is a limiting factor. Automatic endocardial border detection in echocardiographic image sequences remains a challenging problem. Various image artifacts like dropouts, noise, shadowing etc. conceal the true endocardial position. Endocardial border properties are very region dependent (especially in apical views) and the expert-defined border position often does not coincide with a strong image feature. Moreover, experts often employ the temporal coherence of the image data to solve ambiguities, e.g. by playing the image loop and checking temporal consistency of the borders. Automated detection techniques based on separate 2D images do not exploit this temporal coherence.

In this study, a new approach to fully automated detection of time-continuous contours in echocardiographic image sequences is developed and evaluated based on Active Appearance Models (AAMs), with improved robustness by employing temporal coherence in the data.

2. Methods

2.1. Active appearance models

2D Active Appearance Models as introduced by Cootes and Taylor [1] have demonstrated to be highly suitable for segmentation of noisy cardiovascular images such as cardiac MR images and echocardiograms, as we have shown in previous studies [2]. An AAM describes the shape and image appearance of an object over a set of examples as it is defined by experts. An AAM models the object shape and appearance as a multidimensional statistical distribution, consisting of the average shape and a shape-free average image patch, plus a set of eigenvariations that describe the most important shape and appearance variations over the set of examples. AAM processing is only described briefly here, see [1-3] for details.

After the object of interest has been outlined manually by an expert in all examples, each shape is expressed as a 2N-dimensional vector: $\mathbf{x} = (x_1, y_1, x_2, y_2, \cdots, x_N, y_N)^T$. All shape samples are aligned by Procrustes Analysis and an average shape $\bar{\mathbf{x}}$ is calculated. By applying a Principal Component Analysis (PCA) on the sample distribution, a matrix of eigenvectors $\mathbf{P}_s$ is found that describes all significant shape variations over the training set. Each aligned sample $\mathbf{x}$ within the distribution can now be approximated by the average shape with a linear combination of the eigenvectors superimposed: $\mathbf{x} \approx \bar{\mathbf{x}} + \mathbf{P}_s \mathbf{b}_s$ where $\mathbf{b}_s = \mathbf{P}_s^T(\mathbf{x} - \bar{\mathbf{x}})$ is a vector containing the coefficients for each of the eigenvectors.

The image patch around each shape is warped to the average shape, yielding a shape-free intensity patch. Each sample patch is now expressed as a high-dimensional vector of intensities $\mathbf{g}$, and after an intensity normalization, a similar PCA will find the intensity eigenvectors. Shape and intensity coefficients for each sample are concatenated and a third PCA is applied to compute eigen-

0276-6547/01 $17.00 © 2001 IEEE

Computers in Cardiology 2001;28:93-96.

vectors describing simultaneous shape and intensity variations ('appearance'). This gives a compact model of the total appearance of the object: each sample can now be expressed as a linear combination of appearance eigenvectors. By varying the coefficients of the appearance eigenvectors, each of the examples can be resynthesized, as well as any statistically likely 'intermediate' appearance. Such a model can be used for analysis-by-synthesis: segmentation by comparing the synthesized object appearance to the image under investigation, and varying the model coefficients to minimize the difference. A gradient descent approach is used; the relation between a certain difference image and the disturbance in the coefficients is learned during a regression training stage.

2.2. Active appearance-motion models

However, sequential application of 2-D AAMs to a time sequence does not guarantee a time-continuous segmentation result and does not exploit temporal coherency in the data. We present a novel, 2D + time Active Appearance-Motion Model (AAMM) that models shape and appearance of the heart in combination with the dynamics of the cycle. In the AAMM, the appearance of the heart is modeled for the entire cardiac cycle by considering the sequence as a stack of 2D time slices. All single-beat sequences are phase-normalized into a fixed number of frames F (16) so that end-diastolic (ED) and end-systolic (ES) frames map to the same frame number (1, 9, and 16, respectively). When ED and ES frames in the original sequence are known, a relative cardiac phase can be calculated for each frame and the 16 frames for the desired points in the cardiac cycle are selected by nearest neighbor interpolation. In a training set, corresponding points on the endocardial contour are defined for each time slice based on expert drawn contours. The image appearance of the heart in each time frame is modeled as a vector of pixel intensity values in an image patch spanned by the manual contour. Intensity distributions are normalized non-linearly to deal with ultrasound-specific problems and allow Gaussian modeling of intensity [3]. The vectors for shape points for each time frame are concatenated in the order of their phase number (1..F).

$$\mathbf{x} = (\overbrace{x_{11}, y_{11}, \cdots, x_{1N}, y_{1N}}^{Phase\ 1}, \overbrace{x_{21}, y_{21}, \cdots, x_{2N}, y_{2N}}^{Phase\ 2}, \cdots, \underbrace{x_{F1}, y_{F1}, \cdots, x_{FN}, y_{FN}}_{Phase\ F})^T = (\mathbf{x}_1^T | \cdots | \mathbf{x}_F^T)^T$$

A similar approach is used for the intensity vectors.

Only 2-D coordinates of the points are considered, and the resulting shape and intensity vectors are F times as large as in a corresponding 2D AAM. Principal Component Analysis on the training samples renders the mean (Fig. 1) and eigenvariations of shape and intensity appearance of the complete cardiac cycle, and thus captures typical motion patterns associated to cardiac contraction (Fig. 2). Note that the shown eigenvariation models a clinically well-known pattern: a dilated apex (left) shows very little motion over the cycle.

The AAMM matching procedure resembles conventional 2D AAM matching. However, the RMS error criterion and the parameter regression matrices are calculated for the full image sequences, as opposed to a single 2D image frame in AAMs. Therefore the temporal coherence in the cardiac motion is preserved during the matching, ensuring a segmentation result which is largely consistent with the cardiac motion patterns observed in the training set. Typical processing time for a full-sequence match is ~ 6 sec. on an 800MHz PC.

Fig. 1. "Average heartbeat" in 4-chamber view over 65 patients, 5 of 16 phases shown. Left to right: ED (1), mid-systole (5), ES (9), mid-diastole (13), atrial filling (16).

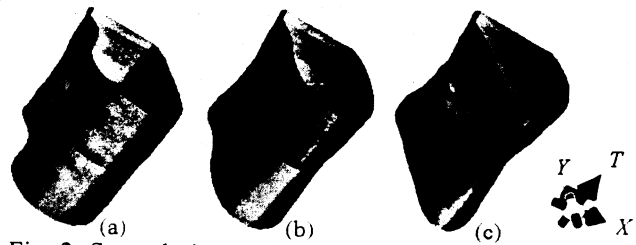

Fig. 2. Second eigenvariation of AAMM shape. The objects represent motion patterns, stacks of contours for the full cardiac cycle (time axis bottom left to top right; shape constriction in the middle corresponds to ES, open ends to ED): a) average shape motion pattern minus three standard deviations; b) average motion pattern; c) average plus three standard deviations.

3. Experiments and material

3.1. Patient material

Transthoracic echo (TTE) 4-chamber sequences were acquired from 129 unselected infarct patients participating in a clinical trial. No patients were excluded for reasons of image quality. Images were digitized from videotape at a resolution of 768*576 pixels (calibration factors 0.28 to 0.47 mm/pixel). ED and ES frames were marked by an expert observer. All single-beat sequences were phase-normalized to F=16 frames as described. The expert observer manually outlined endocardial contours in all frames of all sequences (2064 contours). The set was split randomly into a training set (65 patients) and a testing set (64 patients).

3.2. Intra- and interobserver variability

Inter- and intraobserver variability of manual drawing was determined on 20 randomly selected patients. Contours were drawn independently by a second observer, and again (6 months interval) by the first observer. Contour differences were determined as described below. One patient with an extremely dilated ventricle formed an outlier and was excluded.

3.3. AAMM evaluation

The AAMM created from the training set was applied to segment the 64 test sequences. All models were initialized to the same fixed initial pose (average pose obtained from the training set). For comparison, matching was also performed on the training set.

To compare the automatically detected and manual contours, five indices were calculated per patient (Table 1). Endocardial border positioning error (point distance) was defined as the average of all 592 unsigned distances between corresponding contour points over the full cycle. Success rate was determined as the fraction of patients with point distance < 8 mm.

Endocardial percent area errors were determined separately for each phase of the cardiac cycle. Error in Area Ejection Fraction ($AEF=(1-\frac{Area_{ES}}{Area_{ED}})$) was the difference between computer and manual AEF. All mean errors were tested against zero and the mean of the test set. Statistical significance was determined with a one-sided t-test, where $p<0.01$ was considered significant. Moreover, linear area regressions were performed.

3.4. Comparison to single-frame AAMs

We hypothesized that an AAMM could better exploit the temporal continuity of the image data than AAMs. To test this, we compared the results of the 16-phase AAMM with the combined results of 16 single-phase AAMs, each trained and tested on images of only one phase, with all other parameters identical. Average and standard deviation of endocardial border positioning errors for the 16-phase AAMM were compared to the pooled average and standard deviation of the set of single-phase AAMs. The AAMM average error was tested for being significantly lower using a single-sided Z-test. Percent area errors were compared between each single-phase AAM and the corresponding phase of the 16-phase AAMM.

4. Results

In 62 of all 64 tested patients (97%), the AAMM fully automatically found borders that agreed well with the manual standard (Fig. 3, Table 1) with average point distance of 3.35 mm (9.3 pixels). In two cases the matching failed (point distance >8mm); these were excluded from further analysis (if the failed matches were included, point distance was 3.54 ± 1.62mm). Fig. 4a demonstrates a good correlation of the LV endocardial areas (r=0.87). Endocardial percent area error averaged over all phases was -2.9%, showing a slight but significant negative bias of the AAMM. Fig. 4b shows that endocardial percent area error does not change substantially along the cardiac cycle. Mean AEF error was small: 0.66%.

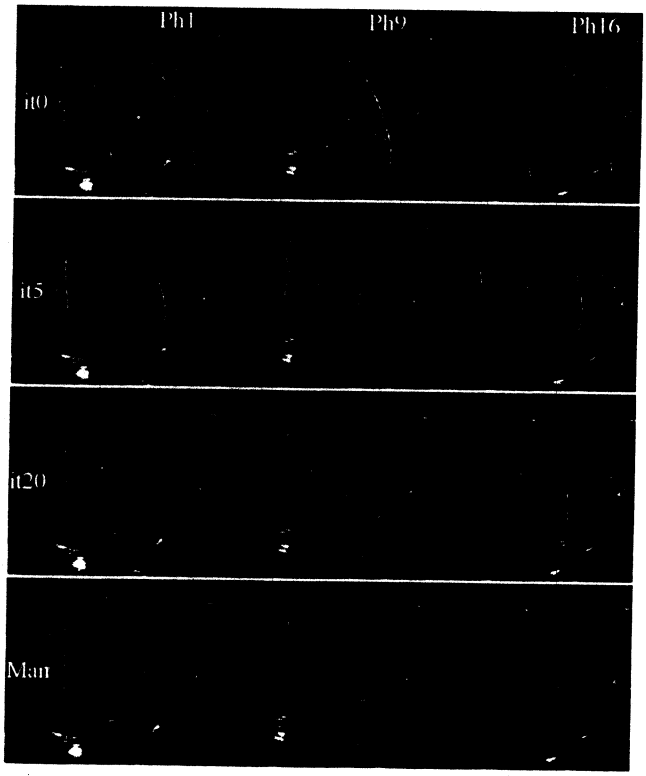

Fig. 3. Example of fully automated AAMM segmentation of a phase-normalized sequence from the test set. Left to right: partial images from phase 1, 9 and 16. Row 1: initial AAMM model positioned on phase images 1, 9, and 16. Row 2: AAMM match after 5 iterations. Row 3: final match after 20 iterations. Row 4: Observer-identified endocardial contours shown for comparison.

Comparison with inter- and intraobserver variabilities (Table 1) clearly shows that the fully automated AAMM detection generally outperforms a second independent observer and is therefore clinically acceptable. However, it does not match the intraobserver variability level. Intraobserver variability can be seen as a measure for inaccuracies in the training data of the AAMM and therefore as a plausible lower bound of the AAMM accuracy.

For comparison, matching was performed on the training set as well. Here, 4 of 65 patients showed a failed match and were left out; for all others the match was nearly perfect (Table 1). All failed matches could be attributed to the model's initialization; manual re-initialization closer to the true pose gave a near-perfect result.

Table 1. AAMM results on test set and comparison to training set and manual inter/intraobserver results.

Results (mean ± SD)	Test set (n = 62 of 64)	Training set (n = 61 of 65)	Manual Intraobserver (n = 19)	Manual Interobserver (n = 19)
Success score (%)	96.9%	93.8%	-	-
Point distance (mm)	3.35 ± 1.22	0.12 ± 0.36 †	2.32 ± 0.75 †	3.82 ± 1.44 †
Fractional area error (%)	-2.89 ± 10.2	0.07 ± 1.08 *†	0.92 ± 6.19 †	-4.39 ± 10.3
Ejection Fraction err (%)	0.66 ± 5.5 *	-0.05 ± 0.41 *	-1.71 ± 2.84	0.88 ± 3.15 *
Area regression (cm^2) (y=computer, x=manual)	y=0.91x+1.73 r = 0.87	y=1.00x+0.001 r = 0.999	y=0.87x+3.57 r = 0.96	y=0.95x+0.15 r = 0.88

* : mean not significantly ≠ 0; † : mean significantly larger/smaller than mean of test set

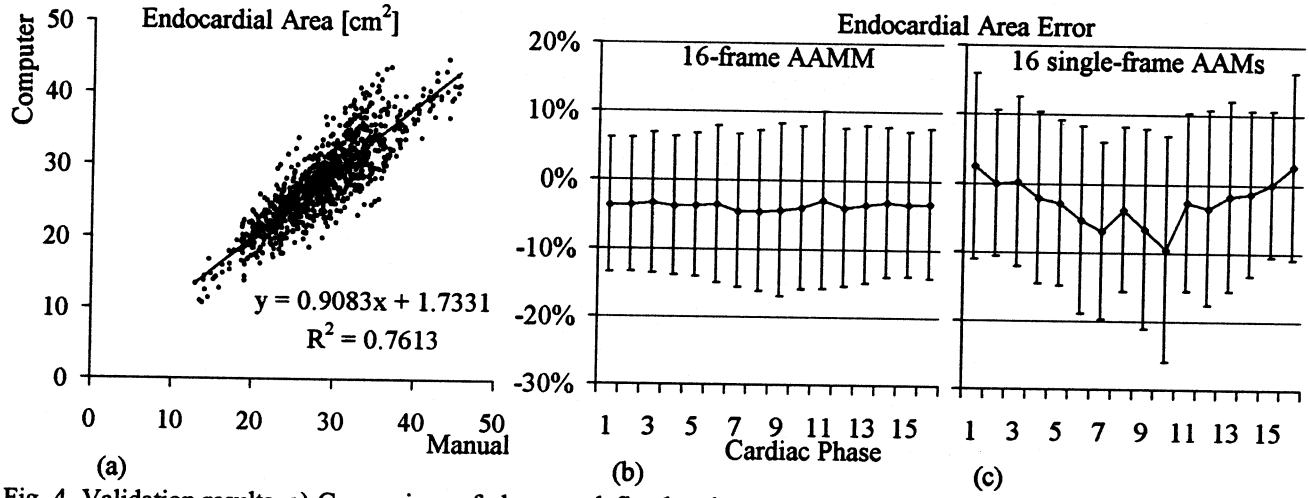

Fig. 4. Validation results. a) Comparison of observer-defined and computer-determined LV endocardial areas in the 992 test set images from 62 out of 64 patients (all 16 cardiac phases). b) Endocardial percent area error as a function of cardiac phase (mean error ± standard deviation). c) Endocardial percent area error for 16 single-phase AAMs.

The pooled average border positioning error for the set of 16 single-frame AAMs was 4.27 ± 2.52 mm (n=1024, 16*64 frames in test set) while for the 16-frame AAMM it was 3.54 ± 1.62 mm. This 17% improvement is highly significant ($p<<0.001$). The percent endocardial area errors in Fig. 4b and 4c show that the AAMM (Fig. 4b) generates much more consistent results over the phases, with generally 15% lower standard deviation.

5. Discussion

An AAM for time sequences (Active Appearance-Motion Model) has been developed that models cardiac motion and image appearance and performs fully automated fast contour detection on phase normalized echocardiographic time sequences in a practically applicable manner. It generates time-continuous segmentation results consistent with cardiac dynamics, and superior to a single-frame approach. It has demonstrated robustness and accuracy in a large clinical study on 4-chamber TTE.

Acknowledgements

This research was supported in part by the Technology Foundation (NWO-Technologiestichting STW), The Netherlands, under grant nr. LGN4349.

References

[1] Cootes TF, Edwards GJ, Taylor CJ. Active Appearance Models. In: Burkhardt H, Neumann B, eds. Proc. Eur Conf Computer Vision 1998. Berlin: Springer, 1998:484-98.

[2] Mitchell SC, Lelieveldt BPF, van der Geest RJ, Bosch HG, Reiber JHC, Sonka M. Multistage hybrid Active Appearance Model matching: segmentation of left and right ventricles in cardiac MR images. IEEE Trans Med Imaging 2001;20:415-23.

[3] Bosch HG, Mitchell SC, Lelieveldt BPF, Nijland F, Kamp O, Sonka M, Reiber JHC. Active Appearance Motion Models for endocardial contour detection in time sequences of echocardiograms. In: Sonka M, Hanson K. Proc. SPIE Medical Imaging 2001 (Vol. 4322). San Diego: SPIE, 2001:257-68.

Address for correspondence.

Johan G. Bosch
Division of Image Processing, Department of Radiology, room 1C2-S, Leiden University Medical Center, P.O. Box 9600, 2300 RC Leiden, The Netherlands, phone +31-71-5262270
E-mail j.g.bosch@lumc.nl