



Universiteit
Leiden
The Netherlands

Model-driven segmentation of X-ray left ventricular angiograms

Oost, C.R.

Citation

Oost, C. R. (2008, September 30). *Model-driven segmentation of X-ray left ventricular angiograms*. Retrieved from <https://hdl.handle.net/1887/13121>

Version: Corrected Publisher's Version

License: [Licence agreement concerning inclusion of doctoral thesis in the Institutional Repository of the University of Leiden](#)

Downloaded from: <https://hdl.handle.net/1887/13121>

Note: To cite this publication please use the final published version (if applicable).

心眼を開く
'gain insight'

Chapter 2

Left Ventricle Contour Detection in X-Ray Angiograms using Multi-View Active Appearance Models

This chapter was adapted from:

Left Ventricle Contour Detection in X-Ray Angiograms using Multi-View Active Appearance Models

E. Oost, B.P.F. Lelieveldt, G. Koning, M. Sonka, and J.H.C. Reiber

In: Proceedings of SPIE Medical Imaging, M. Sonka and J.M. Fitzpatrick Eds., vol. 5032, pp. 394-404, 2003.

Abstract

Automatic Left Ventricle (LV) border detection in X-ray angiograms for the quantitative assessment of cardiac function has proven to be a highly challenging task. The main difficulty is segmenting the End Systolic (ES) phase, in which much of the contrast dye has been squeezed out of the LV due to contraction, resulting in poor LV definition. 2D Active Appearance Models (AAMs) have shown utility for segmenting End Diastolic (ED) angiograms, but do not perform satisfactory in individual ES angiograms. In this work, we present a new Multi-View AAM in which we exploit the existing correlation in shape and texture between ED and ES phase to steer the segmentation of both frames simultaneously. Model scale, orientation and position remain independent, whereas appearance statistics are coupled. In addition, an AAM is presented in which the gray-value information of the inner part of the LV is not taken into account. This so-called boundary AAM is applied mainly to enhance local boundary localization performance. Both models are applied in a combined manner and are validated quantitatively. In 61 out of 70 experiments good convergence for both ED and ES segmentation was achieved, with average border positioning errors of 1.86 mm (ED) and 1.93 mm (ES).

2.1 Introduction

Left Ventricle (LV) angiography is a widely used modality for assessing left ventricular function. Because the heart itself is not visible in X-ray images, patients undergo a catheterization procedure, in which the left ventricle is filled with an X-ray opaque contrast dye. During this procedure 150 to 200 images are acquired in the right anterior oblique view, at a typical rate of 25 to 50 images per second. This generally covers 7 to 9 cardiac cycles. From this image sequence the second or third cardiac cycle is selected. Irregular muscle contractions due to the injection of the contrast agent are expected to have faded at this stage, while the contrast agent distribution within the left ventricle is expected to be optimal. The End Diastolic (ED) and End Systolic (ES) image frames from this cardiac cycle are the starting point for the quantification of left ventricular function. In both the ED and the ES image, a contour line is drawn around the left ventricle. Subsequently a surface area of the projected left ventricle is calculated for both phases, followed by an estimation of the ED and ES volumes, for which we used the volume estimation method proposed by Sandler and Dodge [1]. The ejection fraction is determined from both calculated volumes.

Currently, software packages are available that assist the cardiologists in drawing the contours around the left ventricle manually. But drawing these contours by hand is a difficult and time-consuming task. The poor image quality and varying distribution of the contrast agent complicate the drawing of correct contours. Also, manual contour drawing intrinsically introduces significant inter-observer and intra-observer variability.

Recognizing these difficulties, the need for a reliable and reproducible automatic method for segmenting the left ventricle becomes apparent. Automatic contour

detection in X-ray angiograms is a challenging task for which a reliable technique has not been developed yet to the point of robust clinical application. Several knowledge-based approaches have been proposed over the last two decades, starting with regular edge detection followed by a procedure to merge the separately found LV edges by means of a statistical shape model [2], or approaches using Dynamic Programming and template matching [3].

Until now, the most promising results were obtained in segmenting the ED image, while results in ES segmentation remained poor. Because much of the contrast agent is pumped out of the LV during contraction, the ES image quality and LV definition in general is rather poor, hampering the automatic segmentation of the ventricle. In this chapter, we overcome this by using additional knowledge: for an individual patient, the left ventricular shape and appearance in end diastolic images and end systolic images are highly correlated, although the position and orientation of the ventricle can change significantly due to cardiac motion and contraction. The emphasis of this research lies in the improvement of the ES segmentation results, by using this correlation between the ED phase and the ES phase, resulting in an overall improved automatic assessment of cardiac function.

The methodology is an extension on the Active Appearance Models (AAMs), introduced by Cootes *et al.* [4,5], which in recent years have proven to be highly successful in automatic object segmentation in medical images. AAMs, derived from the earlier introduced Active Shape Models (ASMs), are statistical models describing the shape and the appearance of an object. For both shape and gray-values, an average and a series of eigenvectors is computed, from which the “modes of variation” of the model are determined. When matching the model to an unseen image, it searches the LV contours by minimizing the error between the model and the image, within the boundaries of statistically plausible deformations, as defined by the model.

Direct application of AAMs to ED LV angiograms already shows great potential for automatic segmentation and thus towards automatic quantitative evaluation of LV function. However, segmenting ES images using general Active Appearance Models still does not result in satisfactory results, because of the poor ventricle definition in ES. To exploit the correlation between the LV shape in ED (usually with a better image contrast) and ES, we have developed the Multi-View AAM, which models multiple views of an object. By training the model on the available shape and texture information of both ED and ES simultaneously, the existing correlation between the two phases is preserved and used to steer the segmentation. The idea presented here resembles Cootes’ Coupled View AAMs for tracking faces and estimating head pose [6], but differs greatly in both the training and matching procedure.

In addition, Active Appearance Models have proven to be less accurate than Active Shape Models [7], when it comes to precise border localization. Because global appearance is modeled, accurate border characteristics are less pronounced and image intensity in the middle of the ventricle may negatively influence boundary localization performance. Therefore, we have developed a matching refinement step, where an AAM is applied in which the gray-value information of the inner part of the LV is not taken into account. This “boundary AAM” is also applied in a Multi-View manner: appearance optimization is coupled, while pose is refined for

each frame independently. The principal difference between the boundary AAM and the classic Active Shape Model is the simultaneous optimization of shape and local appearance using a single criterion function, whereas in ASMs, shape and local appearance are optimized independently.

To summarize the novel aspects of this work:

- We have developed the Multi-View Active Appearance Model, allowing to exploit existing correlations between different views of the same object.
- We have introduced a refinement matching procedure, consisting of a boundary model that only takes gray-value information near the suspected LV border into account.

2.2 Active Appearance Models

An Active Appearance Model is a statistical description of appearance of an image in terms of object shape and texture. Applying AAMs consists of two parts: training and matching. Both the construction of the AAM and the matching procedure are briefly introduced in this section. An elaborate explanation can be found in [5].

2.2.1 Active Appearance Model Training

An AAM is trained on a series of representative images, all containing the specific object to segment, in our case X-ray LV angiograms. For each image in the training




Figure 2.1: The first three modes of variation for a left ventricle AAM, end diastolic phase. From top to bottom, the rows represent mode 1, 2 and 3, respectively. From left to right the columns represent a standard deviation of minus two sigma, minus one sigma, zero, plus one sigma and plus two sigma. In this example the first mode of variation describes the shape variation in elongation while the texture emphasis changes from the lower part of the ventricles embedding to the upper part. The second mode of variation mainly describes the shape variation in the bottom part of the ventricle and the brightness in the valve area. The third mode of variation is a combination of elongation and skewing variations.

set, an expert segments the object of interest manually, the drawn contour is sampled in n points and defined (for 2D) as a vector of $2n$ elements, identifying image coordinates. This vector is defined as

$$x = (x_1, y_1, x_2, y_2, x_3, y_3, \dots, x_n, y_n)^T \quad (2.1)$$

Before statistical models of shape and texture are computed, all sample point vectors are aligned using Procrustes Analysis. By applying Principal Component Analysis (PCA) on the sample covariance matrix, a statistical shape model can be built. The eigenvectors and eigenvalues are calculated and the eigenvectors are ordered following descending eigenvalues. Arranging the eigenvectors in this manner enables elimination of less significant eigenvectors, resulting in a statistical model based on the most dominant eigenvectors. The statistical shape model can be formulated as

$$x \approx \bar{x} + P_s b_s \quad (2.2)$$

where any shape vector x in the training set can be approximated by a linear combination of the mean shape $\bar{x}$, and the eigenvectors in P_s , which are weighted by the shape coefficients in parameter vector b_s .

Creating the texture model consists of the following steps. First, a convex hull is constructed from every sample point vector, from which an image patch is sampled. For every training image this image patch is warped onto the mean shape, creating a shape free patch, from which the texture vectors g are extracted. Usually the area of this image patch is dilated to a certain extent, in order to incorporate gray-value information of the direct object surroundings in the model. The elements of a texture vector g represent the pixel intensities in the image patch. All texture vectors are normalized to zero average and unit variance and a PCA is performed on the sample covariance matrix, resulting in the statistical texture model. Analogous to the shape model, each texture sample g is approximated by

$$g \approx \bar{g} + P_g b_g \quad (2.3)$$

with mean texture vector $\bar{g}$, texture eigenvector matrix P_g and the set of texture parameters b_g .

From the shape and texture models, an Active Appearance Model is created by concatenating the shape parameter vector and the texture parameter vector, derived from equations (2.2) and (2.3):

$$b = \begin{pmatrix} W b_s \\ b_g \end{pmatrix} = \begin{pmatrix} W P_s^T (x - \bar{x}) \\ P_g^T (g - \bar{g}) \end{pmatrix} \quad (2.4)$$

W denotes a weight factor coupling the shape and texture coefficients. After a final Principal Component Analysis over the set of appearance vectors b the resulting Active Appearance Model can be written as

$$b = Qc \quad (2.5)$$

in which Q is the matrix containing the eigenvectors and c denotes the appearance parameters. The modes of variation of the AAM display the characteristic variations in shape and gray-value of the model. Figure 2.1 shows the first three modes of variation of an AAM trained on end diastolic LV angiograms.

The last part of training the AAM is to estimate the parameter update steps required to drive the model matching iterations. These are computed from the residual images $\delta g_0 = g_s - g_m$, where g_s denotes the target image, and g_m the model synthesized image. By applying parameter perturbations on the model, pose and texture parameters for model samples with known parameters, gradient matrices R_c , R_p and R_t can be estimated for the model, pose and texture respectively. In our approach, we followed Cootes' direct gradient approach, as recommended in [8].

2.2.2 Active Appearance Model Matching

When matching the model to an unseen image, the model searches the LV contours by minimizing the relative mean square error between the model and the image, within the boundaries of statistically plausible deformations of the model. Based on the model image patch, the image to segment, the current estimate of the model parameters c_0 and the parameter derivatives for the model, texture and pose parameters (matrices R_c , R_t & R_p respectively), Cootes describes an iterative matching algorithm, consisting of the following steps [5]:

- 1) Calculate the difference-vector between the image and model patch
 $\delta g_0 = g_s - g_m$
- 2) Calculate the root mean square (RMS) error of the difference-vector
 $E_0 = |\delta g_0|^2$
- 3) Determine the predicted model parameter updates $\delta c = R_c \delta g_0$, pose update $\delta p = R_p \delta g_0$ and texture update $\delta t = R_t \delta g_0$
- 4) Set $k = 1$ and determine a new estimate for the model parameters
 $c_1 = c_0 - k \delta c$, pose parameters $p_1 = p_0 - k \delta p$ and texture parameters
 $t_1 = t_0 - k \delta t$
- 5) Calculate a new model based on c_1 , p_1 & t_1




Figure 2.2: First mode of variation for a left ventricle Multi-View AAM for X-ray LV angiography. Upper row denotes ED, lower row denotes ES. The correlation in shape between ED and ES is clearly visible. Also the texture variation, describing mainly the local contrast between the LV and its embedding around the mitral valve, shows clear similarities for ED and ES.

- 6) Determine a new difference-vector and calculate its RMS error E_t
- 7) If $E_t < E_o$, select c_t , p_t & t_t as the new parameter vectors, else try $k = 1.5$, $k = 0.5$, $k = 0.25$ etc. and go to step 4

The stop criterion of this algorithm is determined by a fixed number of passes through step 7.

2.2.3 Medical Applications of Active Appearance Models

Several examples exist regarding the application of Active Appearance Models in medical image segmentation. Cootes has demonstrated the application of 2D AAMs on finding structures in brain MR images [9], and knee cartilage in MR images [5]. In 2D cardiac MR images, Mitchell *et al.* applied AAMs to segment the left and right ventricle [10]. Thodberg [11] applied a 2D AAM to reconstruct bones in hand radiographs. Bosch *et al.* successfully used 2D + time AAMs in order to segment endocardial borders in echocardiography [12]. Beichel *et al.* describe a semi 3D AAM extension applied to segmentation of the diaphragm dome in 3D CT data [13]. Mitchell *et al.* describe a full 3D AAM extension, and apply it to 3D cardiac MR data and 2D + time echocardiograms [14]. Research on utilizing AAMs in other modalities is ongoing, but for X-ray LV angiography, no trials have been reported yet.

Active Appearance Models have shown to outperform other segmentation approaches in MRI and echocardiography and are also believed to outperform other methods in LV angiography because of the following advantages:

- Because in many cases the image quality of left ventricular angiograms is poor, it is required to use a priori knowledge about intensity characteristics. Active Appearance models make good use of all available gray-value

information. By supplying it with a sufficiently varied set of LV images, the AAM is adapted to the image appearances that can be expected. This also means that, when supplied with sufficient examples in the training set, AAMs are capable of recognizing pathological shape variations and texture variations caused by acquisition artifacts.

- The left ventricular border does not necessarily coincide with the strongest edges in the image. Because an Active Appearance Model is a statistical description of shape *and* appearance, it is able to copy the drawing characteristics of the cardiologist(s) that drew the contours in the training images. This implies that the model does not necessarily follow the strongest edges, but is adaptive to observer preferences.

2.3 New AAM Extensions

2.3.1 Multi-View Active Appearance Models

The most challenging aspect of automatic border detection in LV angiography is segmenting the left ventricle in the end systolic phase. Due to contraction of the papillary muscles, contrast fluid is squeezed out of the LV in the ES phase, seriously hampering the ES LV visualization. Since manual segmentation of ES images is already highly subjective, automatic segmentation of individual ES images appears to be very difficult. In manual segmentation, the cardiologist reviews the images in cine-mode, instinctively coupling the shape characteristics and texture information of the ventricle and its embedding in ED phase and in ES phase. This observation is our main motivation to pursue simultaneous segmentation of both phases. The Multi-View Active Appearance Model presented here is designed to exploit this existing correlation between ED and ES LV appearance and therefore potentially produces better segmentation results in ES images.

In the Multi-View Active Appearance model, the LV shape is modeled by aligning the training shapes for ED and ES separately, and concatenating the aligned shape vectors for each view. In this application, a shape vector is defined as:

$$\mathbf{x} = (x_{1ED}, y_{1ED}, x_{2ED}, y_{2ED}, \dots, x_{nED}, y_{nED}, x_{1ES}, y_{1ES}, x_{2ES}, y_{2ES}, \dots, x_{nES}, y_{nES})^T \quad (2.6)$$

For the intensity vectors, the same applies: all intensity vectors are separately normalized to zero mean and unit variance, and subsequently concatenated:

$$\mathbf{g} = (i_{1ED}, i_{2ED}, i_{3ED}, \dots, i_{nED}, i_{1ES}, i_{2ES}, i_{3ES}, \dots, i_{nES})^T \quad (2.7)$$

Analogous to the single frame AAM, a PCA is applied to the concatenated sample covariance matrices for shape and appearance, and subsequently a combined

model is computed. In the combined model, the shape and appearance of both views are strongly interrelated, as is illustrated in Figure 2.2.

Estimation of the gradient matrices for computing parameter updates during image matching is performed by applying perturbations on the model, pose and texture parameters. Because of the correlations between views in the model, a parameter disturbance in the model parameters yields difference images in both views simultaneously. The pose parameters however, are perturbed for each view separately, hence the model is trained to accommodate for trivial differences in object pose in each view, whereas the shape and intensity gradients are trained for all views.

In the matching procedure, the pose transformation for each view is also performed separately, whereas the model coefficients intrinsically influence multiple frames at a time. Hence, the allowed shape and intensity deformations are coupled for all frames, whereas pose parameter vectors for each view are optimized independently. This is a significant difference as compared to the coupled view AAM by Cootes *et al.*, where separately trained 2D models are matched to each separate view, and subsequently the appearance model constraints are imposed from a combined appearance model [6].

2.3.2 Boundary Active Appearance Models

Because global appearance is modeled in AAMs, accurate border characteristics are less pronounced, causing possible deterioration of boundary localization performance. Therefore we have developed a boundary Active Appearance Model, which, besides shape information, takes only gray-value information in the vicinity of the detected border into account. By eliminating the gray-value information of the inner part of the ventricle, we aim to improve the border detection accuracy. Figure 2.3 shows the first 3 modes of variation for a boundary AAM.




Figure 2.3: First three modes of variation for a left ventricle boundary AAM, end diastolic phase. From top to bottom, the rows represent mode 1, 2 and 3, respectively. From left to right the columns represent a standard deviation of minus two, minus one, zero, plus one and plus two sigma.

The boundary AAM approach presented here resembles the original ASM formulation, however it differs in the fact that more intensity information in the vicinity of the boundary is utilized: this way, the relation between the manually drawn contour and the underlying image intensity patch is preserved. Also, the boundary AAM matching intrinsically relies on the simultaneous optimization of shape and appearance using a single criterion function, as in ASMs this is performed separately. Applying Active Shape Models for border positioning refinement is not likely to improve results, because of the poor LV border definition, especially in the ES image.

2.4 Experiments and Results

2.4.1 Experimental Setup

To test the clinical efficacy of the Multi-View and boundary AAM, 70 ED-ES pairs of representative LV angiograms from infarct patients were collected. Apart from high quality images with good LV definition in both ED and ES, images were selected, in which frequently appearing acquisition artifacts were present (poor LV contrast, inhomogeneous distribution of the contrast agent, presence of an overlap of the diaphragm with the LV). Figure 2.4 displays a few examples of images that were used in our models. An expert manually defined contours in both frames, and point correspondence was defined based on three prominent landmarks: both aorta valve points and the apex. Every contour was equidistantly resampled to 60 points. 14 leave-five-out models were trained on 65 out of 70 ED-ES image pairs, leaving out 5 sets for testing purposes. To speed up the training and matching process and to reduce model dimensionality, all images were subsampled with a factor 4.




Figure 2.4: LV example images for ED (upper row) and ES (lower row). From left to right: well defined LV, poor contrast, inhomogeneous distribution of the contrast agent (most apparent in ED) and presence of an overlap of the diaphragm with the LV.

In this validation, we aimed to minimize user interaction, therefore all model matching experiments were initialized from a fixed position in the image (the image center), assuming that there is sufficient overlap between the model and the true location of the LV. To minimize the risk of convergence to local minima, sequential application of several AAMs was performed, where the amount of preserved AAM variation was increased from 80% to 99%. For the first model matching, a coupling between the pose and scale parameters in ED and ES was enforced. In subsequent matches, this constraint was released, and pose and scale were allowed to vary for each frame separately.

In addition, the added value of the boundary AAM was investigated by executing an additional model run with the boundary AAM after the regular Multi-View model and comparing their performances.

2.4.2 Evaluation Method

A matching experiment was considered a success when the distance of more than two thirds of the total number of contour points was within 5 pixels of the manual contour. Because there was no calibration factor available for every image, we used a calibration factor of 0.25 mm per pixel, which is representative for most LV angio acquisition systems. In the four-times-subsampled 512 x 512 images that we used, an error of 5 pixels corresponds with 5 mm. Failed matches were repeated with manually set initial position, and cases still yielding matching failure were reported and excluded from further quantitative analysis. These matching failures were divided into three classes: failure in only ED, failure in only ES and failure in both frames. On the successful matches, a statistical analysis was performed, in which the following metrics were used to compare the automatically generated contours with the manual reference contours:

- the border positioning errors (point-to-curve) for the ED and ES contours separately.
- the error in the surface area enclosed by the contours for the ED and ES frames separately.
- the error in estimated ED and ES volumes, as obtained with the Sandler and Dodge volume estimation method [1].
- the error in calculated ejection fraction.

Equation (2.8) describes the volume estimation according to Sandler and Dodge, based on the projected surface Area A and the distance L_A from the upper aortic valve point to the apex:

$$V = \frac{8A^2}{3\pi L_A} \quad (2.8)$$

2.4.3 Results

When matching the LV with a fixed initial position, 57 out of 70 cases (81%) showed convergence. In 4 cases a failure in the calculated ED contour occurred, in 4 other cases the computed ES contour was unsatisfactory and in 5 other cases there was failure in determining both the ED and the ES contour. These 13 failures in LV delineation were matched again, but this time the initialization of the model was done manually. This way 9 failures remained (resulting in a success rate of 87%), 2 in ED, 4 in ES and 3 in ED and ES. Border positioning errors for the 61 successful segmentations are summarized in Table 2.1. For the ED matching results, 72% of the calculated LV contour points was positioned within 2 mm from the expert contour, another 22% was positioned between 2 and 5 mm from the expert contour and the remaining 6% were more than 5 mm away from the expert contour. For the ES matching results, the distribution was 67% (error < 2 mm), 27% (2 mm < error < 5 mm) and 6% (error > 5 mm).




Figure 2.5: Comparison of manual and model results for ED Area ($r = 0.93$; $y = 0.87x + 3.57$ [cm²]), ES Area ($r = 0.93$; $y = 0.94x + 0.31$ [cm²]), ED Volume ($r = 0.93$; $y = 0.83x + 17.74$ [cm³]) and ES Volume ($r = 0.94$; $y = 0.88x + 2.62$ [cm³]).

	Average positioning error [mm]	Average upper valve error [mm]	Average lower valve error [mm]	Average apex error [mm]
ED	1.86 ± 2.45	4.47 ± 3.23	4.60 ± 4.06	4.54 ± 4.63
ES	1.93 ± 2.11	4.46 ± 2.90	3.68 ± 2.67	6.50 ± 5.11

Table 2.1: LV border positioning error and errors in the location of upper valve, lower valve and apex. Errors are expressed as unsigned average $\pm$ standard deviation.

	Average positioning error [mm]	Average upper valve error [mm]	Average lower valve error [mm]	Average apex error [mm]
ED	1.87 ± 2.50	4.47 ± 3.25	4.70 ± 4.22	4.39 ± 4.34
ES	1.97 ± 2.17	4.56 ± 3.24	3.64 ± 2.62	6.44 ± 5.26

Table 2.2: LV border positioning error and errors in the location of upper valve, lower valve and apex, after application of the boundary AAM. Results denote unsigned average $\pm$ standard deviation.

Good correlation has been achieved in comparing the areas and volumes calculated from expert contours and the areas and volumes calculated from the model-determined contours (Figure 2.5). Comparing the calculated ejection fraction for expert contours and model contours, the average signed error was -0.26% with a standard deviation of 14.89% . Still large errors occur, varying from -82% to $+33\%$. Figure 2.6 shows examples of a good matching result and a rejected matching result.

The results for the boundary AAM, also based on the 61 cases of successful segmentation, are summarized in Table 2.2. The average signed error between calculated ejection fractions for expert contours and model contours is 0.24% with a standard deviation of 13.22% . Minimum and maximum errors still can be very large ranging from -71% to $+32\%$. The performance of the boundary AAM is comparable to the results mentioned above: the average boundary positioning errors did not differ statistically significantly as compared to the Multi-View AAM alone. Figure 2.7 displays the potential refinement achievement of the boundary AAM.

2.5 Discussion

The reported method shows promising results in automatic segmentation of left ventricles in X-ray angiograms. The method uses coupled statistical information on

shape and image intensities in the ED and ES images. The matching procedure is based on expert-drawn LV contours and copies human drawing behavior. Although the method has not yet evolved to clinical applicability and especially has difficulties with the segmentation of the LV in poor-quality images, it is the first method that performs well in a large variety of LV shapes in X-ray angiograms for both ED and ES, even when using a fixed initialization in the image center.

We divided the 70 image pairs that were used in this study into four categories: good contrast (28 cases), poor contrast (30 cases), overlap between LV and diaphragm (4 cases) and both poor contrast and overlap (8 cases). When using a fixed initialization, in 57 out of 70 experiments (81%), successful segmentation was achieved, without requiring user interaction. 13 matching failures occurred from which 1 image pair was found to have good contrast, 8 belonged to the poor contrast subset, 1 image pair had an overlap between LV and diaphragm and 3 image pairs both had an overlap and poor contrast. After repeating the 13 failures with manual initialization, 9 cases of failure remained (success rate = 87%), in which 6 image pairs with poor contrast and 3 image pairs with both poor contrast and an overlap between LV and diaphragm.

Figures 2.6a, 2.6b, 2.6e and 2.6f show a representative example of a successful segmentation. From a fixed starting position in the image center, with a scaling that corresponds to the average LV size in the training set, the model finds the correct LV location and deforms within statistically plausible limits to match the underlying LV contours as close as possible. For this specific example, the average border positioning errors are 1.20 ± 1.07 mm and 1.00 ± 0.97 mm for ED and ES




Figure 2.6: Examples of a good matching result (b and f) and a rejected matching result (d and h) for both ED (upper row) and ES (lower row). a and e show fixed initializations in the center of the image, c and g show manual initialization.

respectively. A critical remark that can be made about these successful segmentation results concerns the errors in valve positions in the ED frame and the error in the apex position in the ES frame, which can have a large effect on the volume estimation that is based on the distance between upper valve and apex [1]. Figures 2.6c, 2.6d, 2.6g and 2.6h show a representative example of a failed segmentation. In this case the ED contour was unsatisfactorily detected; the entire region between upper valve and apex should have been drawn wider. This is clearly the result of poor image contrast in this region. Hence, most failures can be assigned to poor contrast conditions and overlap between LV and diaphragm. In the latter case the LV contour will, either locally or entirely, drift away from the true location of the LV.

The success rate of our method suggests that it is capable to cope with a variety of LV shapes. Image pairs with large infarcted areas around the apex however were a minor category in the training set and therefore resulted in a failure in segmentation. When considering the 61 cases of successful segmentation, the average point-to-curve border positioning error was 1.86 ± 2.45 mm for ED images and 1.93 ± 2.11 mm for ES images, which for automatic LV angiographic segmentation is a good performance. Especially the results for ES images are excellent compared to previously explored X ray LV angio segmentation methods [2,3]

The regression lines in Figure 2.5 are based on the 61 occurrences of successful convergence and show a good correlation between surface area and volume computation based on manually drawn contours and based on model-generated contours. Correlation coefficients are 0.93 for ED area, ES Area and ED Volume and 0.94 for ES Volumes, which demonstrates that automatic ES segmentation can accomplish the same level of performance as automatic ED segmentation performance. The error in calculated ejection fraction is -0.26% with a standard deviation of 14.89 %. For both calculated volumes and for the ejection fraction, the systematic errors are close to zero. The accompanying high standard deviations on the other hand show that large errors may still occur, and an improvement is still required to meet clinically demanded error margins.

The three landmark points that are used for calculation of LV volume (upper valve point, lower valve point and apex) have a considerably higher average positioning error than the average border positioning error over the entire contour. Because these three points are of high influence on the calculation of LV volume, future research must aim on decreasing border positioning errors of these three points.

The main difficulty in interpreting the obtained results is that there is no ‘gold standard’, describing the characteristics of a well-drawn LV contour. In clinical practice different medical doctors will have different signatures in drawing contours and therefore an automatic technique is only clinically reliable when the difference in performance between the automatic technique and a medical expert is comparable to the difference in performance between several medical experts. In order to achieve this for Active Appearance Models one may use multiple contours from several different medical experts in model training.




Figure 2.7: Boundary AAM performance examples for ED (a = AAM result, b = boundary AAM result) and ES (c = AAM result, d = boundary AAM result).

The performance criterion of 67% of the calculated contour points to be within 5 millimeters of the actual LV border has not yet been compared with figures for inter-observer variability. This standard is likely to be insufficient in clinical practice, although inter-observer variability can also amount to high values, especially in the difficult images, where observers tend to disagree. Therefore, in the current stage of development of an automatic method for LV segmentation in X-ray angiograms it is legitimate to use this classification to define convergence to an acceptable contour.

The boundary AAM that was introduced for refinement of the segmentation results shows results comparable with the regular Multi-View AAM. However, overlooking all 61 cases, only 41 ED segmentations (67%) and 34 ES (56%) segmentations showed improvement after the boundary AAM. Figure 2.7 shows the potential capacity of the boundary AAM. In Figure 2.7a and 2.7b the segmentation of the lower part of the LV in ED phase has clearly improved and the location of the upper valve is more accurate. Figure 2.7c and 2.7d show improvements in the ES segmentation after application of the boundary AAM, especially expressed in a better segmentation of the apex region.

References

- [1] H. Sandler and H. T. Dodge, "The Use of Single Plane Angiocardiograms for the Calculation of Left Ventricular Volume in Man", *American Heart Journal*, vol. 75, no.3, pp. 325-34, 1968.
- [2] S. Tehrani, T. E. Weymouth, and G. B. J. Mancini, "Model generation and partial matching of left ventricular boundaries," *Proceedings of SPIE Medical Imaging*, vol. 1445, pp. 434-445, 1991.
- [3] P. Lilly, J. Jenkins, and P. Bourdillon, "Automatic contour definition on left ventriculograms by image evidence and a multiple template-based model," *IEEE Transactions on Medical Imaging*, vol. 8, no. 2, pp. 173-185, 1989.

- [4] T. F. Cootes, G. J. Edwards, and C. J. Taylor, "Active appearance models," *Proceedings of the European Conference on Computer Vision*, H. Burkhardt and B. Neumann, Eds., vol. 2, Berlin: Springer Verlag, 1998, pp. 484-498.
- [5] T. F. Cootes and C. J. Taylor, "Statistical models of appearance for computer vision," Online available:
http://personalpages.manchester.ac.uk/staff/timothy.f.cootes/Models/app_models.pdf
- [6] T.F. Cootes, G.V. Wheeler, K.N. Walker, C.J. Taylor, "View-based active appearance models," *Image and Vision Computing*, vol. 20, no. 9-10, pp. 657-664, 2002.
- [7] T.F. Cootes, G. J. Edwards and C. J. Taylor, "Comparing Active Shape Models with Active Appearance Models," *Proceedings of British Machine Vision Conference*, T.Pridmore and D.Elliman, Eds., vol. 1, pp. 173-182, 1999.
- [8] T.F.Cootes, P.Kittipanya-ngam, "Comparing Variations on the Active Appearance Model Algorithm," *Proceedings of British Machine Vision Conference*, vol.2, pp. 837-846, 2002.
- [9] T. F. Cootes, C. Beeston, G. J. Edwards, and C. J. Taylor, "A Unified Framework for Atlas Matching Using Active Appearance Models," *Proceedings of Information Processing in Medical Imaging*, Kuba, Samal and Todd-Pokropek, Eds, vol. 1613, pp. 322-333, 1999.
- [10] S. C. Mitchell, B. P. F. Lelieveldt, R. J. Van der Geest, H. G. Bosch, J. H. C. Reiber, and M. Sonka, "Multistage hybrid active appearance model matching: segmentation of left and right ventricles in cardiac MR images," *IEEE Transactions on Medical Imaging*, vol. 20, no. 5, pp. 415-423, 2001.
- [11] H.H. Thodberg, "Hands on Experience with Active Appearance Models," *Proceedings of SPIE Medical Imaging*, vol. 4684, pp. 495-506, 2002.
- [12] J. G. Bosch, S. C. Mitchell, B. P. F. Lelieveldt, F. Nijland, O. Kamp, M. Sonka, and J. H. C. Reiber, "Automatic segmentation of echocardiographic sequences by active appearance motion models," *IEEE Transactions on Medical Imaging*, vol. 21, no. 11, pp. 1374-1383, 2002.
- [13] R. Beichel, G. Gotschuli, E. Sorantin, F. Leberl, M. Sonka, "Diaphragm Dome Surface Segmentation in CT Data Sets: A 3D Active Appearance Model Approach," *Proceedings of SPIE Medical Imaging*, vol. 4684, pp. 475-484, 2002.
- [14] S.C. Mitchell, J.G. Bosch, B.P.F. Lelieveldt, R.J. van der Geest, J.H.C. Reiber, M. Sonka, "3D Active Appearance Models: Segmentation of Cardiac MR and Ultrasound images," *IEEE Transactions on Medical Imaging*, vol. 21, no. 9, pp. 1167-1178, 2002.

