



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).

中心思想
'the central idea'

Chapter 5

Automated Contour Detection in X-Ray Left Ventricular Angiograms Using Multi-View Active Appearance Models and Dynamic Programming

This chapter was adapted from:

*Automated Contour Detection in X-Ray Left Ventricular Angiograms Using
Multi-View Active Appearance Models and Dynamic Programming*

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

IEEE Transactions on Medical Imaging, vol. 25, no. 9, pp. 1158-1171, 2006.

Abstract

This chapter describes a new approach to the automated segmentation of X-ray left ventricular (LV) angiograms, based on Active Appearance Models (AAMs) and Dynamic Programming. A coupling of shape and texture information between the end-diastolic (ED) and end-systolic (ES) frame was achieved by constructing a Multi-View AAM. Over-constraining of the model was compensated for by employing Dynamic Programming, integrating both intensity and motion features in the cost function. Two applications are compared; a semi-automatic method with manual model initialization, and a fully automatic algorithm. The first proved to be highly robust and accurate, demonstrating high clinical relevance. Based on experiments involving 70 patient data sets, the algorithm's success rate was 100% for ED and 99% for ES, with average unsigned border positioning errors of 0.68 mm for ED and 1.45 mm for ES. Calculated volumes were accurate and unbiased. The fully automatic algorithm, with intrinsically less user interaction was less robust, but showed a high potential, mostly due to a Controlled Gradient Descent in updating the model parameters. The success rate of the fully automatic method was 91% for ED and 83% for ES, with average unsigned border positioning errors of 0.79 mm for ED and 1.55 mm for ES.

5.1 Introduction

X-ray left ventricular (LV) angiography is a widely applied modality for the assessment of cardiac function. To visualize the heart with this modality, patients undergo a catheterization procedure in which the LV is filled with an X-ray opaque contrast dye. Acquisition can be made in bi-plane (combining the antero-posterior view and the lateral view or combining the 30° right anterior oblique view and the 60° left anterior oblique view) or in single-plane (generally 30° right anterior oblique view). Average acquisition time is around 8 to 10 seconds, covering 7 to 9 cardiac cycles. In the second or third cardiac cycle after the injection of the contrast fluid, irregular cardiac contractions generally have subsided and the distribution of the contrast dye within the LV is considered to be optimal. Hence, one of these cycles is selected for analysis of cardiac function by choosing the image frame in which the LV is fully filled (ED) and the first next image frame in which the ventricle is maximally contracted (ES). In both frames endocardial contours are drawn around the LV manually and used to determine surface areas of the projected LV, from which the ventricular volume in ED and ES can be estimated [1]. In addition, relevant clinical parameters such as regional wall motion and ejection fraction (EF) can be determined.

Currently, several packages are available that assist the cardiologists in manually drawing contours in LV angiograms. However, due to frequently occurring poor image quality, an expert examines an X-ray image sequence not only by considering the ED and ES image frame. To decide on the correct boundary locations neighboring frames around ED and ES are inspected. This way,

knowledge about contraction dynamics is used to improve the segmentation accuracy. This makes drawing contours by hand difficult, time-consuming and prone to inter- and intra-observer variability. The goal of the work presented here is to automate the contour detection process to reduce the cardiologist's workload and diminish inter- and intra-observer variability. Because of the aforementioned difficulties of interpreting X-ray angiograms, low level image processing tools have not shown to be sufficiently robust and integration of a priori knowledge is a necessity. Therefore, we aim to integrate the same priors that an expert uses in segmenting the left ventricle, i.e. shape, texture and contraction dynamics priors.

Several knowledge-based approaches for automated segmentation of the left ventricle have been reported, however none of these have been incorporated in daily clinical practice. Tehrani *et al.* [2], for example, combined fractions of possible LV edges, obtained by low level image processing, by means of blackboard architecture and statistical shape models, to come to a full LV delineation. Lilly *et al.* [3] used Dynamic Programming to fit a contour through a set of proposed candidate points, based on regional intensity information, edge information and multiple regional thresholding algorithms. Contour irregularities and contour drift due to insufficient image information were neutralized by template matching using a template library derived from manually traced contours. Figueiredo and Leitão [4] proposed an approach using maximum a priori probability in a Bayesian framework in combination with Markov random field contour modeling. However, this method was applied only on digital subtraction angiography (DSA), in which an image acquired before the injection of the contrast agent is subtracted from all following images, to remove the background of the images and only retaining the object of interest.

Recently McDonald and Sheehan [5] introduced a method using boosted decision trees for pixel classification based on feature images, containing geometry features and gray-level statistics of a sequence of images around the ED and ES frames. By using information from images around ED and ES, LV contraction dynamics is integrated in these feature images. This semi-automatic method requires 3 anatomical landmark points (the endpoints of the aortic valve and the apex) to be positioned manually.

Suzuki *et al.* [6] proposed a combination of edge detection by a modified multilayer neural network and a standard edge detection based on low-pass filtering and edge enhancement. The first method is able to localize less pronounced subjective edges and is trained on manually drawn LV contours, the latter detects rough edges in the images. After placing the two aortic valve points manually, the contours are traced automatically based on both edge detection methods. As in [4] this approach has been applied on DSA images only.

5.1.1 Contribution of this Work

Recent work [5-8] has shown the need for a priori knowledge of shape, image intensities and contraction induced motion in automated segmentation of LV angiograms. The majority of previously reported methods used only shape

knowledge, or was applied to DSA in which most of the intensity information is removed. Only the approach proposed by McDonald and Sheehan [5] uses knowledge about shape and image intensities and motion due to cardiac contraction.

The contribution of this work is fourfold:

- Multi-View AAMs are developed in which statistical information of different views of the same object is modeled simultaneously. The existing correlation in shape and texture between ED and ES is exploited. The more reliable LV information present in the ED images supports the segmentation of the frequently poorly defined LV in the ES images.
- To prevent the model from locking in on local minima, we propose a novel, Controlled Gradient Descent optimization, in which a limited number of model parameters is updated at a time. This greatly improves convergence robustness.
- Motion-based Dynamic Programming is applied to compensate for over-constraining by the model and thus to attain better local border delineation. The cost function is constructed from both image features and features from a subtraction image (ES minus ED). The latter incorporates contraction motion information in the algorithm.
- An elaborate evaluation of clinical efficiency of the algorithm is described based on 70 ED-ES image pairs. To our knowledge, this is the largest evaluation of an automated segmentation method for clinically realistic X-ray LV angiograms.

5.2 Background

Active Appearance Models, introduced by Cootes [9,10], are highly suitable to integrate knowledge in segmentation problems. AAMs are statistical models describing an object's shape and image texture. 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, the object contours are localized by minimizing the error between the model and the image, within the boundaries of statistically plausible deformations of the model.

Examples of successful application of AAMs in automatic segmentation in medical images, whether in 2D, 2D + time, 3D or 3D + time, are ample for a range of imaging modalities. An elaborate overview can be found in [11]. The general construction of an AAM and the conventional matching procedure are briefly introduced in this section. A detailed description can be found in [12].

5.2.1 AAM Training

An AAM is trained on a series of representative images, in which an expert manually segmented the object of interest. Contours are resampled in n corresponding points, and (for 2D) expressed as a vector of $2n$ elements:

$$\mathbf{x} = (x_1, y_1, x_2, y_2, x_3, y_3, \dots, x_n, y_n)^T \quad (5.1)$$

After Procrustes alignment [13] of the shape vectors to eliminate pose differences, a shape model is built by applying Principal Component Analysis (PCA) on the sample covariance matrix. Arranging the eigenvectors according to descending eigenvalues enables elimination of less significant eigenvectors.

The texture model is created by warping the intensities of the object of interest from the training images onto the mean shape. This way, an image patch is created, which is normalized for the shape of the individual samples; from this patch, pixel intensity vectors $\mathbf{g}$ are extracted. Typically, a thin strip around the object is included in the patch, to acquire information of pixel intensities outside the object's boundaries. Texture vectors are normalized to zero average and unit variance and PCA is performed on the sample covariance matrix, resulting in the statistical texture model.

To reduce model size, AAM models are generally constructed using only the eigenvectors corresponding to the largest n eigenvalues, capturing for example 98 % of the available variation. Using shape and texture models, the sample shapes $\mathbf{x}$ and textures $\mathbf{g}$ can be approximated from the respective models:

$$\mathbf{x} \approx \bar{\mathbf{x}} + P_s \mathbf{b}_s \quad \text{and} \quad \mathbf{g} \approx \bar{\mathbf{g}} + P_g \mathbf{b}_g \quad (5.2)$$

where $\bar{\mathbf{g}}$ and $\bar{\mathbf{x}}$ represent the average texture and shape vectors, P_g and P_s the texture and shape eigenvector matrices, and $\mathbf{b}_g$ and $\mathbf{b}_s$ the texture and shape parameters characterizing each training sample.

From the shape and texture models, an AAM is created by concatenating the shape and texture parameter vectors:

$$\mathbf{b} = \begin{pmatrix} W\mathbf{b}_s \\ \mathbf{b}_g \end{pmatrix} = \begin{pmatrix} WP_s^T (\mathbf{x} - \bar{\mathbf{x}}) \\ P_g^T (\mathbf{g} - \bar{\mathbf{g}}) \end{pmatrix} \quad (5.3)$$

W denotes a weight factor coupling the shape and texture coefficients.

After a final PCA over the set of appearance vectors $\mathbf{b}$ the resulting AAM can be written as:

$$\mathbf{b} = Q\mathbf{c} \quad (5.4)$$

in which Q is the matrix containing the combined shape/texture eigenvectors and c denotes the appearance parameters for the combined model.

5.2.2 Using AAMs for Segmentation

Matching the model to an unseen image involves minimizing the sum of squared pixel differences between the model gray-value patch and the normalized target image, within the boundaries of statistically plausible model limits. To drive the model matching iterations, the parameter update steps are computed from the residual images $\delta g_0 = g_s - g_m$, where g_s denotes the normalized target image, and g_m the model synthesized image. From the current estimate of the model parameters c_o and the parameter derivatives for the model and pose parameters, captured in the pre-computed gradient matrices R_c (for the model parameters) and R_p (for the pose parameters), Coates describes an iterative matching algorithm, consisting of the following steps [12,14]:

- 1) Normalize the target image patch to zero mean and unit variance
- 2) Calculate the residual between target image and model patch $\delta g_0 = g_s - g_m$
- 3) Calculate the error from the difference vector $E_0 = |\delta g_0|^2$
- 4) Using the pre-computed gradient matrices, determine the model parameter update $\delta c = R_c \delta g_0$ and pose update $\delta p = R_p \delta g_0$
- 5) Set $k = 1$
- 6) Determine new estimates for the model parameters $c_1 = c_0 - k \delta c$ and pose parameters $p_1 = p_0 - k \delta p$
- 7) Calculate a new model based on c_1 and p_1
- 8) Determine a new difference vector and calculate a new error E_1
- 9) If $E_1 < E_0$, select c_1 and p_1 as the new parameter vectors, else try $k = 1.5$, $k = 0.5$, $k = 0.25$ etc. and go to step 6

Repeat until convergence, either using a fixed number of iterations, or until no improvement is achieved.

As mentioned in Section 5.1, we propose a different matching strategy, which differs substantially from this approach. This novel approach will be discussed in Section 5.3.

5.3 Segmentation Method

The proposed method consists of three novel components: the Multi-View AAM, in which different views of the same object are modeled simultaneously, a parameter updating strategy that isolates the modes of variation yielding the largest criterion decrease, and a locally selective Dynamic Programming algorithm as post-processing to relax potential over-constraining model priors, increasing border localization accuracy.

5.3.1 Multi-View AAM

In many medical image segmentation problems one deals with multiple viewpoints, multiple cross sections or multiple time instances. Because all these views describe one single object, correlations in pixel intensities and contour shapes between views are present and can be exploited. This is particularly useful for X-ray LV angiography.

The Multi-View Active Appearance Model introduced here is specifically constructed to exploit the existing correlation between different views of the same object. The concept is derived from Cootes' work on coupled-view AAMs [15], where a frontal and a side view of a face are segmented simultaneously by building separate models for each view and a combined model for both views. During matching, segmentation is performed using the single view models, while shape constraints are applied from the combined model. Our method differs from this concept in that it models both organ shape and organ texture simultaneously for all views. Both during model training as during automatic segmentation, ED and ES scale, orientation and position are independent.

The key novelty in this approach is the direct modeling of existing correlations in image intensities, which is used to drive the segmentation of both images simultaneously. Contrary to Cootes' coupled-view AAM, the Multi-View AAM segments all views simultaneously using only one model. Because this model is constructed by concatenating, for every training sample, the shape and intensity vectors for all views, it directly exploits existing correlations, both in shape and in intensities, between the views.

The Multi-View model is constructed by aligning the training shapes for different views separately, and concatenating the aligned shape vectors x_i (see equation 5.1) for each of the N views. A Multi-View shape vector, x_{mv} , for N frames is defined as:

$$x_{mv} = (x_1^T, x_2^T, \dots, x_N^T)^T \quad (5.5)$$

By applying a Principal Component Analysis on the sample covariance matrix of the combined shapes, a shape model is computed for all frames simultaneously.

The principal model components represent shape variations, which are intrinsically coupled for all views.

For the intensity model, the same applies: an image patch is warped on the average shape for view i and sampled into an intensity vector g_i , the intensity vectors for each single frame are normalized to zero mean and unit variance, and concatenated:

$$g_{mv} = (g_1^T, g_2^T, \dots, g_N^T)^T \quad (5.6)$$

Analogous to the single frame AAM, a PCA is applied to the sample covariance matrices of the concatenated intensity sample vectors, and subsequently each training sample is expressed as a set of shape- and appearance coefficients. A combined model is computed from the combined shape/intensity sample vectors. Figure 5.1 demonstrates that in the combined model, the shape and appearance of both views are strongly interrelated.




Figure 5.1: First mode of variation for a left ventricle Multi-View AAM, constructed from 70 ED-ES X-ray LV angiograms. Upper row = ED, lower row = ES. From left to right the columns represent a standard deviation of minus two, minus one, zero, plus one and plus two sigma. The black dotted line represents the ventricle border. 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 between ED and ES.

Estimation of the gradient matrices for computing parameter updates during image matching is performed by applying perturbations on the model and pose parameters, and measuring their effect on the residual images [12,14]. Because of the correlations between views in the model, a disturbance in an individual model parameter yields residual images in all views simultaneously.

Although full 2D + time and 3D + time ASMs and AAMs have been reported [16-18], this work only reports a coupling of the ED and the ES frame in X-ray LV angiography, since they are the most clinically relevant frames. The ES image frame

in which much of the injected contrast agent has been ejected is especially difficult to interpret. ED images generally exhibit a better contrast. The existing correlations between both frames can be exploited by Multi-View AAMs to achieve better boundary delineation especially in the ES image. In practice, this will give the opportunity to exchange, for example, information about overlapping structures between the several views. When, for example in ED segmentation, such an overlapping structure is ‘recognized’, this information can be used in the simultaneous segmentation of the ES frame. This kind of image intensity information exchange between both views is crucial, in particular for difficult to interpret images like LV angiograms.

Stegmann and Pedersen [19] proposed a similar AAM for cardiac MRI, in which pose information was captured in the shape vector by concatenating the shapes of different views before Procrustes Alignment was carried out. This approach might also be useful for exploiting pose correlations between the ED and ES frames in LV angiography. It is however an extra factor that can over-constrain the model.

5.3.2 Controlled Gradient Descent

LV angiograms may exhibit many acquisition artifacts, such as overlapping anatomical structures (ribs, diaphragm) or strong shadows. In these cases, the regular AAM matching strategy described above, occasionally shows difficulties in converging to the true contour positions. A typical example of a diverging model is displayed in Figure 5.2. To cope with these situations during segmentation, images like this have to be included in the training data set, resulting in large allowed texture variations incorporated in the model.

Also the gradient matrix R_c which is used for updating the model parameters, is constructed during training by measuring the effect of disturbances on the isolated model parameters, while the pose parameters remain undisturbed. Therefore, fitting the model to the image using this matrix therefore will function optimally when the pose initialization is (nearly) correct. However, when the model is initialized far from the object of interest, or has a strongly deviating scale or orientation, the matrix R_c may lose its validity. This explains why a regular AAM initialized far from the correct position may converge to a local minimum, since the algorithm tries to reduce the difference between model and underlying image based on inaccurate gradient approximations.

Several methods have been described on constraining the energy function for improved convergence robustness; a review of such methods can be found in [20]. A common way to increase robustness is by limiting the number of simultaneous degrees of freedom during optimization; this parameter scheduling has been described mainly applied for minimizing physics-based energy functions, e.g. [21]; In the context of AAM fitting, we developed a Controlled Gradient Descent, updating a limited number of parameters at a time, where the free parameters are selected dynamically as follows. First all parameter updates, corresponding to the specific modes of variation, are sorted to descending magnitude and the largest




Figure 5.2: Example of spurious error criterion behavior. Black dotted lines represent correct LV borders, white dotted lines represent model results. From left to right the error criterion decreases.

single parameter update is executed. If this results in a lower error criterion, the proposed update is accepted, new model and pose parameters are calculated and a new vector of parameter updates is determined and ordered again in the next iteration. In case an update proposal does not lower the error criterion, the number of updated model parameters is incremented in the next attempt. If a new update proposal is successful, the next iteration will start again with only one model parameter.

5.3.3 Motion-Based Dynamic Programming

The power of the AAM algorithm is that it is able to still come to an acceptable global segmentation in an environment with partly invisible, or only vaguely perceptible features. AAM segmentation is based on minimizing the sum of squared pixel errors between a global model patch and the normalized underlying image and due to the global nature of this criterion it does not focus on local borders. Moreover, statistical models such as AAMs generally are over-constrained: the model freedom to deform is limited by the modes of variation derived from the training data set. Therefore, a shape that slightly deviates from characteristic shapes in the training set should also be considered as valid and a refinement of the contour is desirable. In previous work [8] an AAM contour refinement was done by applying a second AAM, in which only image intensities close to the contour were incorporated. This approach had a positive effect on the segmentation, but being model-based, it still intrinsically over-constrained the contours towards the training data. The same holds for a hybrid ASM/AAM segmentation approach, which has been reported for cardiac MR [22]. To allow for more shape flexibility a shape relaxation is required.

We developed a Dynamic Programming algorithm, in which the cost function is constructed from image and motion features, to mimic the experts' routine of including knowledge of contraction dynamics. A generalized Dynamic Programming algorithm that is used in X-ray angiography, for example for extracting coronary contours [23], searches for an optimal contour path through a cost matrix, based on first and second order derivatives, describing the edges of the object to be segmented. The cost function is generally defined by:




Figure 5.3: Additional information on true LV border position is available in the subtraction image (c): ES (b) minus ED (a). This specific example shows that the subtraction image results in a better definition of the mitral valve area in both the ED and ES frame. In addition it diminishes the influence of a diagonal shadow in the true image data.

$$C(i, j) = \alpha G(i, j) + (1 - \alpha) T(i, j) \quad (5.7)$$

in which $C(i, j)$ is the cost of element in row i and column j , $G(i, j)$ represents the gradient, $T(i, j)$ is the second order derivatives and α denotes the weighing factor between the first and second order derivatives.

In this work we have integrated features from a subtraction image (ES image minus ED image, see Figure 5.3) into the cost function. The cost function is constructed such that it searches a minimal cost path based on directional edges, while integrating both information of true image data and the subtraction image. The polarity of edges in the subtraction image is defined differently for ED and ES, making the cost function locally selective for each phase. In ED the area outside the contour should be dark and the area inside the contour should be light. For ES this edge polarity is opposite. The use of these directional edges is only possible, because the Multi-View AAM already produces a reliable global segmentation in each frame, close to the desired solution. This enables Dynamic Programming in a limited search space. The cost function used in our algorithm is similar to equation 5.7:

$$C(i, j) = \beta(\alpha_1 G_1(i, j) + (1 - \alpha_1) T_1(i, j)) + (1 - \beta)(\alpha_2 G_2(i, j) + (1 - \alpha_2) T_2(i, j)) \quad (5.8)$$

in which $C(i, j)$ is the cost of element in row i and column j , β denotes the weighing factor between the cost matrix constructed from the true image data and the cost matrix constructed from the subtraction image, $G_1(i, j)$ and $G_2(i, j)$ are the gradients of both images, $T_1(i, j)$ and $T_2(i, j)$ are the second order derivatives and α_1 and α_2 are weighing factors between the first and second order derivatives for the true image data and the subtraction image data respectively.

5.4 Clinical Evaluation

The purpose of this study is to determine the clinical utility of our approach and more specifically, whether Multi-View AAM segmentation results are of a comparable quality as manual segmentation results produced by clinicians. We evaluate the automatically generated contours by relating them to contours of three experts. Furthermore we determine the state of automation that can be achieved, by comparing a fully automatic method with a semi-automatic approach in which for both ED and ES the endpoints of the aortic valve and the apex are predefined by a user. In addition we compare the proposed techniques with conventional Active Appearance Models and with conventional Dynamic Programming.

5.4.1 Data Material

To examine the effectiveness of our methodology we have tested it on 70 single-plane ED-ES pairs. All angiograms were acquired in the 30° right anterior oblique view, using standard contrast agent (Iopamiro® 350). Average acquisition time was about 8 to 10 seconds, covering 7 to 9 cardiac cycles. All data stemmed from adult patients, suffering from one, two or diffuse coronary diseases. There was no pathology-based selection of the data set. Two exclusion criteria were applied: ventricles being not fully imaged and ventricles showing only extra-systolic contraction. As Figure 5.4 illustrates, many different acquisition artifacts were encountered. The situations displayed in Figures 5.4d and 5.4e were especially frequently occurring (20 out of 70 cases and 15 out of 70 cases respectively).

5.4.2 AAM Training

For all 70 paired ED-ES images a clinician has drawn manual contours which have been used to train the Active Appearance Model. Point correspondence was achieved by resampling every contour to 60 equidistant points, based on three specific landmark points: the upper aortic valve point, the lower aortic valve point and the apex. This resulted in a training sample description of 120 points, combining the ED and the ES shapes. 14 leave-five-out models were trained, on 65 out of the 70 ED-ES image pairs, leaving out 5 pairs for testing purposes. Using leave-five-out instead of leave-one-out was opted mainly for efficiency reasons, whereas a model trained on 65 data sets does not differ very much from a model trained on 69 data sets.

All models were constructed retaining 100% shape variance and 95% intensity variance. This corresponded with 64 shape modes and 27 intensity modes and resulted in 64 appearance modes of variation.

To speed up both AAM training and AAM matching, all training and matching experiments were executed after subsampling the images by a factor 4.




Figure 5.4: Typical examples of LV angiograms (upper row = ED; bottom row = ES): good contrast (a), poor contrast (b), uneven contrast agent distribution (c), partial overlap with the diaphragm (d) and substantial shadows caused by the shutter (e), expressed in this example as a diagonal dark band. The black dotted lines represent the LV boundaries as determined manually by an expert. Particular difficulty can be expected in the ES frame of image (d). In a normal situation the texture knowledge within the model consists of a relative dark ventricle in an embedding with higher pixel intensities. Locally in this image it is opposite, while the true LV contour apparently coincides more or less with the strong edge of the diaphragm.

5.4.3 Semi-Automatic Segmentation

In the semi-automatic segmentation, the model was initialized using three landmark points: upper aortic valve point, lower aortic valve point and apex. Initial model scale, orientation and position were estimated from these three points. Initial scale, orientation and position were furthermore used to constrain the pose parameters, which were allowed to differ 10% from the initial pose.

Model parameters were ordered and initially updated separately while the number of model parameters was incremented when a proposed model update was rejected. To ignore the influence of trivial model parameters, a maximum of 50% of the most pronounced modes of variation was employed. Possible local minima were evaded by following Cootes' forced update approach [12,14] twice after the model could not improve any further. The best values for model and pose parameters were stored and, if not improved in a subsequent attempt, were considered to be the final result.

5.4.4 Fully Automatic Segmentation

The fully automatic segmentation consisted of two stages. The first stage globally positioned the model, after initialization in the image center. The second stage was similar to the semi-automatic segmentation. However, without user interaction no knowledge about approximate scale, orientation and position was available and therefore pose constraints were not applied.

When large changes in pose parameters occur, it is likely that previous updates of

model parameters have been inaccurate, since the overlap between the model and the ventricle probably was limited. Therefore, it is better to discard these model parameter updates. In the global positioning of the model, model parameters were reset when the scaling or orientation deformation exceeded 10% or when the shift in x or y direction exceeded 5% of image width or height respectively.

5.4.5 Comparison with Conventional Methods

To determine the effect of all proposed contributions, the following comparisons were executed:

- A comparison between Single-View and Multi-View AAMs.
- A comparison between the Controlled Gradient Descent and a conventional AAM.
- A comparison between the hybrid AAM / Dynamic Programming algorithm and standalone AAM / Dynamic Programming.
- A comparison between conventional Dynamic Programming and Dynamic Programming in which subtraction image information is included.

5.4.6 Dynamic Programming Parameters

Table 5.1 summarizes the relevant Dynamic Programming parameters, used for all experiments. The parameters were selected after a pilot study on part of the data set. The allowed search area for ED is 16 mm on both sides of the AAM contour, for ES this area is set to 6 mm inside and 12 mm outside the AAM contour. Due to lack of visible contrast agent in ES and the resulting tendency of the AAM to produce an LV ES segmentation with a slightly too small surface area, the Dynamic Programming search area for ES is restricted more. The weighting factor β during Dynamic Programming is 0.5 for ED and 0.7 for ES. Due to a better definition of the ES contour in the subtraction image, information in the subtraction image is

	ED	ES
α	0.5	0.5
β	0.5	0.7
area outside [mm]	16	12
area inside [mm]	16	6

Table 5.1: Dynamic Programming parameters.

graded higher than the information from the true image data. For the Dynamic Programming in the ED frame and in the ES frame, edge convolution filters with opposite signs are used. This way the correct edge polarity is created.

5.4.7 Evaluation Indices

Experiments were carried out using the data set of 70 ED-ES image pairs. None of the test images was included in the model that was used for AAM segmentation. Calculated ED volume, calculated ES volume and calculated EF were compared with the manually defined independent standard. Volumes were calculated by using the area-length equation introduced by Sandler and Dodge [1]:

$$V = \frac{8A^2}{3\pi L_A} \quad (5.9)$$

in which A denotes the projected surface area and L_A is the distance from upper aortic valve point to apex. Linear regression was used to determine relationships between manually traced and computer determined values. A two-tailed paired samples t-test was applied to volume measurements from manual and automatic contours to investigate systematic errors. A p-value smaller than 0.05 was considered significant.

In addition, the point-to-curve errors of the contours and the percentage of the automatic contour requiring manual correction by an expert were determined. When a contour part of at least 3 successive points had a point-to-curve error larger than 2 mm, it was considered necessary to redraw this contour part. This parameter for contour editing was not selected to discriminate between proper segmentation and failure, but to distinguish between drawing preferences of different experts in practice. A quantitative evaluation method from recently published work [6] has been adopted for comparison:

$$E_C = \frac{\sum_{x,y \in R_E} \{a_P(x,y) \otimes a_D(x,y)\}}{\sum_{x,y \in R_E} a_D(x,y)} \quad (5.10)$$

$$E_A = \frac{\left| \sum_{x,y \in R_E} a_D(x,y) - \sum_{x,y \in R_E} a_P(x,y) \right|}{\sum_{x,y \in R_E} a_D(x,y)} \quad (5.11)$$

with

$$a_P(x,y) = \begin{cases} 1, & (x,y) \in R_P \\ 0, & \text{otherwise} \end{cases} \quad (5.12)$$

$$a_D(x,y) = \begin{cases} 1, & (x,y) \in R_D \\ 0, & \text{otherwise} \end{cases} \quad (5.13)$$

in which R_P is the region within the automatically drawn contour, R_D is the region within the manually drawn contour, R_E is the region of evaluation and $\otimes$ denotes the logical exclusive OR operator. Equation 5.10 represents a contour error defined as the summation of pixels that are exclusively located within the boundaries of either one of the two compared contours, divided by the total sum of pixels of the evaluation area (in this case the surface area of the manual contour). Equation 5.11 denotes the difference in surface area between automatic and manual contour, in respect to the surface area of the manual contour.

It is well known that the difficulty in interpreting LV angiograms results in a large inter- (and intra-) observer variability. When comparing, for example, the contours of the three experts contributing to this project, differences in estimated ES volumes up to a factor 2.6 were observed. The highest observed average unsigned point-to-curve differences amounted to 5 mm. Therefore it is difficult to define a gold standard or to define the notion of success of an automatic LV segmentation algorithm. To this end, we first investigated inter-observer variability for the ejection fraction, ED volume calculation, ES volume calculation, ED point-to-curve values and ES point-to-curve values. The cut-off parameter thresholds for outliers were defined as the average of the three maximum differences in the parameters between the three experts. These obvious segmentation failures were noted and excluded in further contour evaluation indices.

The performance of our algorithms was tested by comparing obtained results with the manual contours, that were used to train the 14 AAMs (expert #1 contours). Because the leave-five-out setup was used, none of the tested image pairs was included in the specific model that was used for segmentation. To assess the clinical relevance of the method, calculated contours were compared with manually drawn contours of three experts. In addition, to establish the accuracy of the manually defined standard, different experts were compared to each other.

5.5 Results

5.5.1 Semi-Automatic Segmentation

The semi automatic algorithm yielded borders that agreed closely to the manual expert contours. The success rate of the algorithm is 100% for ED and 99% (1 outlier) for ES. After removal of the image pair with this partial failure, both ED and ES contour errors, calculated areas and calculated volumes were, to our knowledge, better than any previously reported method [2-6]. Figure 5.5 displays representative examples of obtained contours, proving that accurate segmentation is also feasible in images with acquisition artifacts.

Border positioning errors were generally small. Average unsigned point-to-curve errors were 0.68 ± 0.37 mm for ED and 1.45 ± 0.76 mm for ES. All quantitative results (model vs. expert #1) are summarized in Figure 5.6, together with a




Figure 5.5: Successful matches for ED and ES generated with the semi-automatic algorithm. Black dotted lines denote the manual contour, white dotted lines represent the semi-automatic contours. Semi-automatic contours correspond closely with manual contours, also in images with acquisition artifacts such as low contrast (a & b), overlapping diaphragm (c & d), strong shadows (e & f). Contours are particularly good in images without acquisition artifacts (g & h).




Figure 5.6: Point-to-curve distances (a & f), contour errors (see equation 5.10) (b & g), area errors (see equation 5.11) (c & h), volume errors (d & i) and the percentage of contour that needs to be redrawn (e & j) for ED (a-e) and ES (f-j). All bars in the graphs denote average values, error ranges span 1 standard deviation in both directions. Six comparisons are displayed: model vs. expert #1, model vs. expert #2, model vs. expert #3, expert #1 vs. expert #2, expert #1 vs. expert #3 and expert #2 vs. expert #3.

comparison of semi-automatically generated contours with contours drawn by expert #2 and expert #3 and a mutual comparison of all three experts. For the mutual comparison of experts, only 43 samples of the original 70 paired ED-ES data were available.

Excellent correlation between volumes based on manual and semi-automatic contours was achieved, as shown in Figure 5.7. In a paired samples t-test differences between manually and semi-automatically calculated ED volume, ES volume and ejection fraction were found statistically insignificant ($p=0.13$, $p=0.76$ and $p=0.15$ respectively). Table 5.2 provides an overview of errors in ED volume, ES volume and ejection fraction. The semi-automatic algorithm compared to expert #1 has the overall best results. In particular, the differences in calculated ES volume and EF are remarkably good, better than any of the inter-expert differences.




Figure 5.7: Volume regression plots for ED (a), ES (b) and ejection fraction (c) for the semi-automatic algorithm.

	samples	ED error [%]	ES error [%]	EF error [%]
model vs exp 1	69	-1.56	-0.88	-1.20
model vs exp 2	42	-0.86	12.79	-6.20
model vs exp 3	42	0.30	9.54	-4.26
exp 1 vs exp 2	43	0.70	12.11	-4.95
exp 1 vs exp 3	43	1.85	9.45	-2.98
exp 2 vs exp 3	43	1.15	-3.03	1.96

Table 5.2: Comparison of semi-automatic contours with 3 experts and comparing the experts mutually: relative ED volume and ES volume errors and the absolute ejection fraction error.




Figure 5.8: Volume regression plots for ED (a), ES (b) and ejection fraction (c) for the automatic algorithm.




Figure 5.9: Successful matches for ED and ES generated with the automatic algorithm. Black dotted lines denote the manual contour, white dotted lines represent the automatic contours. Acceptable results can be achieved for good quality images (a & b) as well for images with artifacts such as low contrast (c & d), overlapping diaphragm (e & f) and poor contrast distribution (g & h). Also when the LV scaling is extremely large (c & d) or when the LV is relatively far from the image center (g & h), where the model is initialized, acceptable segmentation is feasible. Note that in image (h) user information on the apex location is indispensable for correct segmentation in the apex area.

5.5.2 Fully Automatic Segmentation

The success rate of the fully automatic algorithm is 91% for ED and 83% for ES. 6 complete failures were observed, in which both ED and ES segmentation diverged, and 6 partial failures in which only ES segmentation failed. After the removal of these outliers, unsigned point-to-curve errors were 0.79 ± 0.43 for ED and 1.55 ± 0.66 for ES, comparable to the semi-automatic results.

Linear regression plots are presented in Figure 5.8, showing acceptable results. Only ED volume comparison between manual and automatic contours was statistically significant ($p=0.03$). Errors in ES volume and ejection fraction were found statistically insignificant ($p=0.33$ and $p=0.72$ respectively). Examples of successful fully automatic segmentations are shown in Figure 5.9.

5.5.3 Comparison with Conventional Methods

The three proposed technical contributions (Multi-View AAM, Controlled Gradient Descent and the hybrid AAM / Dynamic Programming) have been compared with conventional AAMs and with conventional Dynamic Programming. Table 5.3 gives an overview of the experiments. All results in this table are based on the entire data set, without removal of obvious segmentation failures. The values denote averages $\pm$ standard deviation. ED_Ec and ES_Ec denote contour errors for ED and ES, as defined by equation 5.10. ED_Ev and ES_Ev denote relative volume errors for ED and ES.

The first two lines in Table 5.3 show the difference between Multi-View AAMs and Single-View AAMs for the semi-automatic algorithm, when no Controlled Gradient Descent and no post-processing by means of Dynamic Programming is applied. When initialization is done manually (first 2 lines), the results for Multi-View and

	ED_EC [%]	ED_EV [%]	ES_EC [%]	ES_EV [%]
Multi-View (semi)	13.2 $\pm$ 6.2	15.7 $\pm$ 13.5	22.6 $\pm$ 14.8	25.8 $\pm$ 37.3
Single-View (semi)	12.8 $\pm$ 6.7	15.3 $\pm$ 13.3	22.2 $\pm$ 13.9	24.7 $\pm$ 29.4
Multi-View (fully)	20.4 $\pm$ 15.8	18.7 $\pm$ 20.4	34.7 $\pm$ 24.2	33.1 $\pm$ 40.3
Single-View (fully)	26.6 $\pm$ 21.6	26.7 $\pm$ 32.8	48.5 $\pm$ 37.2	51.4 $\pm$ 63.9
CGD	13.4 $\pm$ 8.1	14.3 $\pm$ 12.7	25.6 $\pm$ 16.4	25.3 $\pm$ 23.8
no CGD	20.4 $\pm$ 15.8	18.7 $\pm$ 20.4	34.7 $\pm$ 24.2	33.1 $\pm$ 40.3
AAM + DP	4.1 $\pm$ 1.9	3.6 $\pm$ 3.5	14.1 $\pm$ 12.1	17.4 $\pm$ 34.7
only AAM	13.7 $\pm$ 6.0	15.7 $\pm$ 13.4	23.3 $\pm$ 14.2	26.0 $\pm$ 32.8
only DP	4.8 $\pm$ 3.1	4.5 $\pm$ 5.5	19.2 $\pm$ 17.1	28.7 $\pm$ 52.0
AAM + 1 img DP	5.8 $\pm$ 3.9	6.2 $\pm$ 7.6	15.5 $\pm$ 15.1	21.6 $\pm$ 42.2

Table 5.3: Comparison with regular AAMs and Dynamic Programming.




Figure 5.10: The benefit of integrating both true image information and subtraction image information in Dynamic Programming. Black dotted lines denote the manual contour, white dotted lines represent the computed contours. Results after AAM segmentation (left) show a proper segmentation of the anterior wall, while the posterior wall is delineated poorly. With regular Dynamic Programming as post-processing tool (middle) using only the true image information, the segmentation of the posterior wall is corrected. However, due to the strong shadow in the image, the contour near the anterior wall is attracted to the wrong edge. Including also subtraction image information (right) results in a proper segmentation of both anterior and posterior wall (The subtraction image information for this example can be found in Figure 5.3c).

Single-View are comparable, although the standard deviation for ES_Ev is significantly smaller for the Single-View Model. The added value of the Multi-View model becomes obvious when comparing results for fully automatic segmentation (lines 3 and 4). For all 4 parameters the Multi-View Model clearly outperforms the Single-View models.

To determine the effect of the Controlled Gradient Descent, experiments were repeated while using a regular Multi-View AAM instead of the proposed algorithm. When applying a regular Multi-View AAM in semi-automatic segmentation, performance and accuracy remained similar. However, a significant difference in performance occurred when applying a regular Multi-View AAM in fully automatic segmentation (lines 5 and 6).

The lines 7-9 in Table 5.3 show the need for a hybrid algorithm, combining AAMs and Dynamic Programming. When providing the model with the correct positions of the upper valve point, lower valve point and apex point, for both ED and ES, the hybrid algorithm obviously outperforms both methods, when applied separately. Furthermore, when the subtraction image information is not incorporated in Dynamic Programming (last line of the table), results deteriorate. Figure 5.10 shows a typical example.

5.6 Discussion

The method proposed in this chapter introduces three novel elements: a Multi-View AAM is developed to exploit the coherence between different image frames, a

locally selective Dynamic Programming is introduced to relax over-constraining inherent to statistical shape models and a Controlled Gradient Descent strategy is proposed to improve convergence and lock-in range. The Multi-View AAM robustly detects the initial models for the Dynamic Programming, enabling the integration of locally selective motion features in the cost function. The proposed method proved to be a robust and accurate tool for automatic segmentation of the left ventricle in angiographic images. It uses knowledge about shape, texture and motion from a large variety of training images and inherently mimics the drawing behavior of the clinical expert. Even in images that are difficult to interpret the algorithm produces reliable results.




Figure 5.11: The only, partial, segmentation failure when using the semi-automatic algorithm. Black dotted lines denote the manual contour, white dotted lines represent the semi-automatic contours.

5.6.1 Automatic versus Semi-Automatic Segmentation

The semi-automatic algorithm shows a high success rate of 100% for ED and 99% for ES. The only failure occurred when the ES image showed an extremely slim and elongated shape (see Figure 5.11). Because of the elongated shape, the model is initialized with a relatively large scale. This, in combination with the object's shape extremity and the applied pose constraints, resulted in a failure. In general, pose constraints are a useful guarantee to prevent the model from diverging, but this example shows that in extreme cases it restricted the model too much. The less restricted fully automatic algorithm actually performs better in this example.

Correlation between manually determined LV volumes and semi-automatically calculated LV volumes was excellent. In particular, the ES results were very good compared to previously reported segmentation methods for this modality [2,3] which can be attributed to the large amount of knowledge, incorporated in the model. The correlation values shown in Figure 5.7 are, to our knowledge, the best values reported until now. However, correlation values ($R^2 = \{0.99; 0.95; 0.84\}$) do not match inter-observer correlations ($R^2 = \{0.99; 0.98; 0.93\}$). Due to a lack of image information, ES volumes are generally slightly underestimated.

Quantitative evaluation results of the semi-automatic algorithm, displayed in Figure 5.6 and Table 5.2, proved to be within boundaries of inter-observer variability. The average difference and standard deviation in comparing the semi-automatic method with expert #1 contours (the expert who produced the training contours) were comparable to values obtained when comparing different experts. The ability to mimic expert drawing behavior is most particularly pronounced in Table 5.2. Differences between the semi-automatic algorithm and expert #1 are generally smaller than differences between the three experts. Furthermore, comparing line 2 with line 4 of Table 5.2 and comparing line 3 with line 5, indicates that the model has comparable behavior as expert #1. The proposed method produces results that are within limits of inter-observer variability and therefore is considered to be clinically relevant.

Also, our method outperforms other recently published methods. Suzuki's neural edge detector [6], trained on 12 ED and 12 ES images, achieved average contour errors E_C of 6.2% and 17.1% for ED and ES respectively and average area errors E_A of 4.2% and 11.6% for ED and ES respectively. The semi-automatic approach presented in this chapter needs a similar amount of user interaction and produces (after removal of the single ES segmentation failure) E_C values of 4.1% and 12.8% and E_A values of 1.9% and 6.4%, comparing favorably to Suzuki's results on all indices. Moreover, more than five times as much data was used in our evaluation.

Both the success rate and the quantitative results for the fully automatic algorithm were not as good as the semi-automatic approach. The major difficulty in fully automatic segmentation is the location of the three landmark points; upper aortic valve point, lower aortic valve point and apex. Errors for these landmarks are 3.8 mm, 4.1 mm and 3.0 mm respectively for ED and 4.4 mm, 3.8 mm and 6.2 mm for ES. Errors in these landmarks strongly influence the volume estimates using the area-length method (Equation 5.9).

In terms of accuracy, the fully automatic algorithm provided very good segmentation results, as Figure 5.9 points out. However, the rate of matching failures is significantly higher than for the semi-automatic algorithm. After removal of these failures, quantitative results are comparable to the results obtained with the semi-automatic algorithm. Although fully automatic segmentation results are promising, the semi-automatic algorithm is more robust, while the amount of user interaction remains limited.

5.6.2 Behavior of the Proposed Algorithm

Multi-View versus Single-View

Given optimal initialization, the Multi-View algorithm displays comparable results as Single-View models, as Table 5.3 points out. For ES segmentation, the Single-View model even appears to be marginally better. The benefit of the Multi-View AAM becomes evident when the initialization is not perfect. Table 5.3 clearly shows that the average contour error and the average volume error are significantly higher for the Single-View models, than for the Multi-View model. More pronounced is the difference in the standard deviations of the errors. This all corresponds with the higher amount of obvious segmentation failures, observed for the Single-View

models. One can conclude that a model that is dedicated to a specific view is only slightly better than a model dealing with multiple views, *only* when the initialization is perfect. When the initialization is less good, Multi-View AAMs are much more robust than Single-View AAMs. Segmentation of the ES frame is improved in particular, due to the coupling of shape and intensity information from both the ED phase and the ES phase.

In LV angiography strong correlation in the ventricle's pose between the ED image and the ES image is apparent. However, strict coupling of the ED and ES pose parameters seemed to excessively tie the pose parameters of the two views, hampering a correct convergence.

Controlled Gradient Descent

Using the Controlled Gradient Descent further constrains model convergence: The model adapts only one or a few model parameters, while pose parameters remain unaffected. The constraint, introduced by the Controlled Gradient Descent, results in a gradual updating of the model parameters, preventing the model from locking in on the direct underlying image features.

As expected, the Controlled Gradient Descent clearly outperforms a regular AAM, as Table 5.3 points out. It is more robust than a regular AAM and it is a proper instrument to cope with texture ambiguities and structures that are in some cases present in the image, and in some cases not.

Other solutions to this problem may be possible, such as the one proposed in [19], in which the regions in short-axis MR images where papillary muscles could possibly be visible were excluded from the model. Since in X-ray LV angiography the entire posterior wall can be obstructed by the diaphragm and the entire anterior wall can be obstructed by either a rib or a shutter-induced shadow, such a pruning technique is not suitable for automatic segmentation of X-ray angiograms. A parameter update strategy, comparable to the Controlled Gradient Descent, is presented in [21], for fitting blended 3D deformable models. First global parameters of one model are adapted, followed by parameters that blend different 3D models, and subsequently parameters that change the model's topology. In the final stages of fitting the model all parameters are released simultaneously. In this updating scheme the parameters are updated in a fixed order. The Controlled Gradient Descent, however, first decides in every iteration which parameter(s) should be regarded most significant, based on the features found in the underlying image. These parameters are updated first, resulting in a tailor-made parameter updating strategy.

Hybrid AAM and Dynamic Programming

The proposed AAM is a robust tool for global segmentation. However, because of its global nature, it lacks sufficient accuracy for correct border delineation. The result contour computed by the AAM on the other hand, creates a solid base for subsequent local segmentation by means of Dynamic Programming. Since the AAM segmentation result is already close to the desired solution (see for example the left image in Figure 5.10), the Dynamic Programming can be constrained to a small

specific area. Therefore, the proposed hybrid algorithm outperforms both an AAM and Dynamic Programming, when applied separately, as Table 5.3 indicates. Including the subtraction image into the Dynamic Programming cost function is a way of capturing cardiac motion dynamics. The hyper-intensity area in this image strongly correlates with the area in which the ventricle boundary progresses during contraction and release. This information is essential to make Dynamic Programming a proper post-processing tool for contour refinement, as shown in Figure 5.10.

5.6.3 Clinical Applicability

The presented results prove that the semi-automatic algorithm is a highly robust and accurate method for segmentation of the left ventricle in angiograms. Both end-diastolic and end-systolic segmentation results were within boundaries of inter-observer variability. Accurate results were obtained, even when acquisition artifacts were present, such as poor contrast, overlapping diaphragm or strong shadows in the image.

A predominant feature of our algorithm is that it is trained on a large data set of clinical images and manually drawn expert contours. Therefore it inherently mimics the drawing behavior of the clinical expert. Because the model is trained on patient data the majority of pathologies can be recognized and matched during segmentation.

The method is fast (1-2 seconds per case) and needs minimal user input. After setting 6 seed points, the model produces the ED and ES contours. Subsequent manual editing of the contours should be possible in clinical practice, but little need for editing is expected.

Estimations based on experiments are that on average less than 5% of editing is needed for an automatically generated ED contour and approximately 20% editing is needed for an automatically generated ES contour. To put these numbers in perspective, based on the same conditions, expert #1 on average would redraw 2% of an ED contour drawn by expert #2, 1% of an ED contour drawn by expert #3, 14% of an ES contour drawn by expert #2 and 12% of an ES contour drawn by expert #3.

In conclusion, the semi-automatic algorithm can be a helpful and reliable tool in automated segmentation of LV angiograms in daily clinical practice. As Table 5.2 points out, a further refinement of the algorithm can be achieved by training the model on expert data from several different experts. This will result in a more solid base for creating a quality standard in segmenting LV angiograms.

5.7 Conclusions

A new algorithm for semi-automatic segmentation of the left ventricle in X-ray LV angiograms is presented. The method is a combination of a Multi-View Active Appearance Model and a locally selective Dynamic Programming approach and

exploits knowledge about LV shape, image texture and contraction dynamics. The algorithm is capable of mimicking the drawing behavior of a clinical expert and therefore provides excellent results. ES segmentation results have especially improved significantly with respect to previously reported methods, due to the coupling of statistical data for both frames. Local AAM border positioning is refined by a Dynamic Programming step in which both image information and knowledge of contraction dynamics was integrated in the cost function. Furthermore, the robustness in fully automatic segmentation has improved significantly by introducing a Controlled Gradient Descent approach in updating the model parameters, evading local minima.

The algorithm has been tested on 70 paired ED-ES images from patient studies, showing high robustness and accuracy, when compared to expert contours. Results were within boundaries of inter-observer variability and derived volume calculations were accurate and unbiased.

Although less robust than the semi-automatic algorithm, the proposed fully automatic method has shown high potential. The proposed Controlled Gradient Descent has especially shown its benefit in fully automatic segmentation.

The semi-automatic method has proven to have high relevance for daily clinical practice, in segmenting the left ventricle angiograms for the quantitative assessment of cardiac function.

References

- [1] H. Sandler and H.T. Dodge, "The use of single plane angiocardigrams for the calculation of left ventricular volume in man," *American Heart Journal*, vol. 75, no. 3, pp. 325-334, 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] M. A. T. de Figueiredo and J. M. N. Leitão, "Bayesian estimation of ventricular contours in angiographic images," *IEEE Transactions on Medical Imaging*, vol. 11, no. 3, pp. 416-429, 1992.
- [5] J. A. McDonald and F. H. Sheehan, "Ventriculogram segmentation using boosted decision trees," *Proceedings of SPIE Medical Imaging*, vol. 5370, pp. 1804-1814, 2004.
- [6] K. Suzuki, I. Horiba, N. Sugie, and M. Nanki, "Extraction of left ventricular contours from left ventriculograms by means of a neural edge detector," *IEEE Transactions on Medical Imaging*, vol. 23, no. 3, pp. 330-339, 2004.
- [7] C. R. Oost, B. P. F. Lelieveldt, M. Üzümcü, H. Lamb, J. H. C. Reiber, and M. Sonka, "Multi-view active appearance models: application to x-ray LV angiography and cardiac MRI," *Proceedings of Information Processing in Medical Imaging*, C. J. Taylor and J. A. Noble Eds., vol. 2732, pp. 234-245, 2004.
- [8] E. Oost, B. P. F. Lelieveldt, G. Koning, M. Sonka, and J. H. C. Reiber, "Left ventricle contour detection in x-ray angiograms using multi-view active appearance models," *Proceedings of SPIE Medical Imaging*, vol. 5032, pp. 394-404, 2003.
- [9] G. J. Edwards, C. J. Taylor, and T. F. Cootes, "Interpreting Face Images using Active Appearance Models," *Proceedings of the 3rd International Conference on Automatic Face and Gesture Recognition*, pp. 300-305, 1998.
- [10] 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.

- [11] M. B. Stegmann, "Generative interpretation of medical images," *Kgs. Lyngby: PhD thesis Informatics and Mathematical Modelling*, Technical University of Denmark. 232 p., 2004.
- [12] 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
- [13] J. C. Gower, "Generalized procrustes analysis," *Psychometrika*, vol. 40, no. 1, pp. 33-51, 1975.
- [14] 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.
- [15] 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.
- [16] 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.
- [17] R. J. van der Geest, B. P. F. Lelieveldt, E. Angelíé, M. Danilouchkine, C. Swingen, M. Sonka, and J. H. C. Reiber, "Evaluation of a new method for automated detection of left ventricular boundaries in time series of magnetic resonance images using active appearance motion model," *Journal of Cardiovascular Magnetic Resonance*, vol. 6, no. 3, pp. 609-617, 2004.
- [18] G. Hamarneh and T. Gustavsson, "Deformable spatio-temporal shape models: extending active shape models to 2D+time," *Image and Vision Computing*, vol. 22, no. 6, pp. 461-470, 2004.
- [19] M. B. Stegmann and D. Pedersen, "Bi-temporal 3D Active Appearance Models with Applications to Unsupervised Ejection Fraction Estimation," *Proceedings SPIE Medical Imaging*, vol. 5747 pp. 336-350, 2005.
- [20] J. Montagnat, H. Delingette, N. Ayache, "A review of deformable surfaces: topology, geometry and deformation," *Image and Vision Computing*, vol. 19, no. 14, pp. 1023-1040, 2001.
- [21] D. DeCarlo and D. Metaxas, "Blended deformable models," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 18, no. 4, pp. 443-448, 1996.

- [22] 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.
- [23] J. H. C. Reiber, L. R. Schiemanck, P. M. J. van der Zwet, B. Goedhart, G. Koning, M. Lammertsma, M. Danse, J. J. Gerbrands, M. J. Schalijs, and A. V. G. Bruschke, "State of the Art in Quantitative Coronary Arteriography as of 1996," In: J. H. C. Reiber and E. E. van der Wall editors. *Cardiovascular Imaging*, Dordrecht: Kluwer Academic Publishers, pp. 39-56, 1996.

