

Segmentation of myocardial perfusion MR sequences with multi-band Active Appearance Models driven by spatial and temporal features

N. Baka^{a,b,d}, J. Milles^a, E.A. Hendriks^{a,b}, A. Suinesiaputra^a, M. Jerosh Herold^c,
J.H.C. Reiber^a, B.P.F. Lelieveldt^{a,b}

^aDivision of Image Processing, Department of Radiology,
Leiden University Medical Center, Leiden, the Netherlands

^bMediamatics Department, Delft University of Technology, Delft, the Netherlands

^cDepartment of Radiology, Brigham & Women's Hospital, Boston, USA

^dBiomedical Imaging Group Rotterdam, Erasmus Medical Center, Rotterdam,
the Netherlands

ABSTRACT

This work investigates knowledge driven segmentation of cardiac MR perfusion sequences. We build upon previous work on multi-band AAMs to integrate into the segmentation both spatial priors about myocardial shape as well as temporal priors about characteristic perfusion patterns. Different temporal and spatial features are developed without a strict need for temporal correspondence across the image sequences. We also investigate which combination of spatial and temporal features yields the best segmentation performance. Our evaluation criteria were boundary errors wrt manual segmentations, area overlap, and convergence envelope. From a quantitative evaluation on 19 perfusion studies, we conclude that a combination of the maximum intensity projection feature and gradient orientation map yields the best segmentation performance, with an average point-to-curve error of 0.9-1 pixel wrt manual contours. We also conclude that addition of different temporal features does not necessarily increase performance.

Keywords: multi-band AAM, cardiac perfusion sequences, temporal features, gradient orientation map, maximum intensity projection

1. INTRODUCTION

Ischemic heart disease, and more specifically myocardial infarction are highly prevalent diseases in western, developed countries. The disease is characterized by a decreased blood supply within the heart muscle. The main cause of ischemic heart disease is the obstruction of a coronary artery (coronary atherosclerosis). First-pass myocardial perfusion magnetic resonance imaging (MRI) is commonly used for early ischemia detection in coronary heart disease. During first pass myocardial perfusion MRI, a sequence of images is acquired from short axis cross sections of the heart. The sequence typically consists of approximately 50 images, each taken in the same cardiac phase in several cardiac cycles. Acquisition is triggered by the ECG signal. In this manner, the circulation of a bolus of extracellular paramagnetic contrast agent can be followed. Typically, the contrast agent enters the right ventricle (RV) first, subsequently it reaches the Left Ventricle (LV), arrives at the myocardium, and afterward the first pass ends. The intensities appearing in the scans are related to the blood flow, as the bolus (DTPA) does not enter the cells and remains in the extracellular space. The size of the first pass non-perfused area in perfusion images was shown to correlate with the size of the infarct area.¹ Figure 2 shows an example of a cardiac MR perfusion sequence.

Usual practice for hypoperfusion assessment is to mark regions of interest (ROI) in all frames. Time intensity curves are computed for each pixel, and subsequently averaged for the marked regions for noise suppression. Such a signal intensity curve is shown in Figure 1 for two regions: the Left Ventricle (LV); and the myocardium. In first pass imaging the first pass of the bolus through the heart is captured. The upslope parameter of the myocardial tissue is a proven indicator of the coronary artery disease.² A manual segmentation of the myocardial

ROI and assessment of the obstruction severity is a time-demanding process, therefore an automated procedure is desired.

Automated quantitative analysis of myocardial perfusion is hampered by strong motion artifacts caused by breathing of the patient during the scanning, and minor inaccuracies in triggering. Therefore, post-processing should consist of two steps: 1) breathing motion correction, and 2) segmentation. In this work, we focus on the second element: segmentation of already registered data, where motion has been eliminated using a previously developed motion correction algorithm³ or manual registration. However, even for motion-corrected image data, segmentation is challenging due to the rapidly changing contrast variations caused by the circulating contrast agent, by noise, fuzzy edges and image artifacts. Also, the desired segmentation does not always coincide with the strongest gradients, especially in the area around the papillary muscles. Integration of prior knowledge is therefore essential to robustly segment the myocardium in these image sequences.




Figure 1. Signal-Intensity curve of myocardial and left ventricle regions over 50 frames. The square markers show the frame positions shown on the right in Figure 2.




Figure 2. Frames from a mid-ventricular first pass perfusion sequence using a Gadolinium bolus. From left to right: the baseline state without any bolus (frame 2); contrast bolus arrived in the right ventricle (frame 10); bolus appears in the left ventricle (frame 15); bolus is in the left ventricle (frame 18); perfusion of the myocardium (frame 22); start of the second pass of the contrast agent (frame 40)

In this work we build upon previous work on multi-band Active Appearance Models to:

- develop a segmentation method that enables the integration of both spatial as well as temporal priors in the segmentation of the myocardium.
- evaluate which combination of spatial and temporal image features leads to the best segmentation performance.

The paper is organized as follows. Section 2 gives a short overview of Active Appearance Models (AAMs) in general, a brief literature review, and the proposed multi-band AAM. This is followed by a basic introduction of the sequence characteristics, and the proposed feature images. Section 3 describes the experimental validation protocol, the evaluation data, the error criteria, and the experiments. Discussion of the results follows in Section 4. Finally Section 5 summarizes the presented work.

2. METHODS

2.1 AAM Algorithm

AAMs consist of a statistical shape model and an intensity model trained from a set of manually annotated example images. The model can be constrained to generate only plausible object appearances. Segmentation is established via matching the model texture to the target image. As shape and texture variation are coupled in the model, this will yield the desired segmentation of the original unseen image.

Shape is modeled as a point distribution model (PDM):

$$x_s = \bar{x}_s + \Phi_s b_s$$

with \bar{x}_s is the mean shape, Φ_s represents the eigenvectors of the covariance matrix, and b_s is the vector of shape parameters. The shape x_s can thus be represented as the mean shape and a linear combination of the principle axes. For a more detailed description we refer to.⁴

Intensities are handled similarly. Each image is sampled on corresponding positions, and is represented by an intensity vector. Point correspondence is obtained by deriving a pixel mask from the mean shape, and subsequently using a thin plate spline registration of the shapes to this mean. After intensity normalization a model is built by PCA analysis.

$$h = \bar{h} + \Phi_h b_h$$

where h denotes the intensity vector, \bar{h} is the mean texture, Φ_h are the principle components derived from the training intensities, and b_h is the texture parameter vector.

The two models are combined by concatenating the model parameters for the same training samples. As the units of shape (PDM) and intensities (texture model) differ, the shape parameters are weighted (W_s) so that both models contain the same total variance. We define the total variance as the sum of the variances of the principal components. A final PCA is performed on these vectors to determine the principal directions of the combined shape-appearance vectors.

$$b = \begin{pmatrix} W_s b_s \\ b_h \end{pmatrix} = \Phi_{cc} \quad (1)$$

During matching a gradient descent algorithm is used to find the set of parameters that minimize the intensity difference between model and target image. The metric to be minimized is the root-mean-square error between the model texture vector and the target texture vector.

2.2 Related literature

Segmentation of cardiac motion sequences was investigated by Bosch et al.⁵ and van der Geest et al.⁶ Their method, called Active Appearance Motion Model (AAMM) combines the shapes and intensities of different time-points into one vector, and further uses the AAM algorithm. It requires temporal correspondence, which is established by manually marking the end systolic and end diastolic frames, and interpolating in between to derive a fixed number of frames. This approach was shown to perform better than multiple AAM-s applied to the same data set. A drawback of this method is the high model dimensionality, which necessitates a sufficiently large number of training images. The authors also detected a tendency toward finding a 'too normal' fit.

Stegmann et al.⁷ presented an AAM based approach for myocardial perfusion sequence segmentation which is called the cluster segmentation aware-AAM. Over the sequence the shape remains nearly stationary, while there is high variation in the texture, therefore he argues to leave out the final PCA step of the AAM, and handle distinct shape and texture parameters. The assumption of PCA modeling of texture, namely that the set of texture vectors can be modeled as a simple multivariate Gaussian, is in this case not valid. The proposed solution is to compute k ($k = 5$) separate PCA models of the available training samples, using prior unsupervised clustering of the training frames. A new image is matched to all models, and the best match is selected. The cluster aware AAM is therefore quite computationally expensive, and it does not include any information about the bolus pass into the segmentation.

Stegmann and Larsen⁸ extended the AAM framework to model objects in images with multiple intensity bands, such as RGB color images. The different color bands of an image are concatenated into one texture

vector, subsequently AAM is performed using this input. Further processing, such as the normalization of the texture vectors, and the update scheme is identical to the original AAM. The authors state, that regarding the dissimilar statistics of the sub-bands of the different representations, the joint normalization is not optimal.

The work of Scott et al.⁹ relies on the same idea. They use different features, like edgeness, cornerness, gradient structure, instead of color bands of an image. After trying different combinations of features they conclude that the intensity AAM could be improved by the combined local feature AAM in all tested cases. They do however not investigate the performance outcome of combined intensity and local features AAM. As all of the features used in the experiments have similar units, they also do not consider using different normalizations for the different features.

2.3 Proposed algorithm

In this work, we explore the potential of multi-band AAMs to integrate spatial and temporal information in the segmentation of cardiac MR perfusion sequences. To realize this, we make two adaptations to the multi-band AAM framework:

1. Instead of inserting RGB bands in the multi-band AAM, we compute several combinations of temporal and spatial feature images from the MR perfusion sequences, which serve as intensity bands in the multi-band AAM. We selected temporal features which express information about the temporal perfusion profile: Maximum Intensity Projection (MIP), average intensity over time, maximum upslope and time to peak intensity. As spatial features, we adopted the gradient orientation maps of the average intensity feature image from.⁹ Figure 3 displays all these feature images for one case.
2. In multi-band AAMs, the intensity model is constructed by expressing the image intensities in a training sample as a concatenation of the intensity vectors for each intensity band. However, the dimensionality of the vector space may become prohibitively large with an increasing number of feature images: each intensity vector may consist of thousands of pixels. To avoid this dimension expansion in the multi-band AAM for different feature combinations, we first perform a dimensionality reduction of the feature images, and subsequently apply the multi-band AAM to the reduced feature vectors. This dimensionality reduction is performed by constructing a separate appearance model for each feature image after computing the reduced eigenvector matrix from the the sample covariance matrix. Each sample is then expressed as a linear combination of these eigenvectors. Furthermore, such model construction allows different normalizations for the separate features.

Let s be a shape vector and $f(i)$ is a feature vector ($i = 1 \dots N$ with N is the number of features). Then, the shape and intensity of a training sample can be combined for all input features by concatenating the different vectors b , while applying a weighting vector for each feature to accommodate the differences between the shape and intensity features:

$$b = \begin{pmatrix} W_s b_s \\ W_{f(1)} b_{f(1)} \\ \vdots \\ W_{f(N)} b_{f(N)} \end{pmatrix} = \Phi_c c = \begin{pmatrix} \Phi_{cs} \\ \Phi_{ct} \end{pmatrix} c \quad (2)$$

where W_s and $W_{f(i)}$ are the weighting matrices for the shape and the i th feature vector, respectively. The vectors b represent the projection of the a particular feature on the reduced eigenvector set, i.e. the shape s , feature images $f(i)$. Applying a PCA to the covariance matrix of combined b -vectors yields the eigenvectors of the combined model.

Weighting matrices W_s and $W_{f(i)}$ were defined conform the weighting scheme reported in Cootes and Taylor,⁴ as a diagonal matrix with elements equal to the ratio between the total intensity variation for a particular feature image (trace of the feature model eigenvalue matrix) and the total variation of a feature used as reference. Motivations behind such weighting schemes are:

1. To eliminate differences caused by diverse units, hence enabling combination of different characteristics of the object.
2. To accommodate for different noise characteristics in the various features: a down-weighting of features containing a high amount of noise may be desired.

Shape and feature images can be reconstructed directly from the final parameters by the following equations:

$$\begin{aligned} x &= \bar{x} + \Phi_s W_s^{-1} \Phi_{cs} c \\ h_{f(i)} &= \bar{h}_{f(i)} + \Phi_{f(i)} W_{f(i)}^{-1} \Phi_{cf(i)} c \quad \text{for all } i = 1..N, \end{aligned}$$

where $\Phi_{cf(i)}$ and Φ_{cs} are sub-matrices of Φ_c , as described in Equation 2.

The image matching scheme is equivalent to the basic AAM, with the extension that the error vector is computed from the concatenation of the different feature image error vectors, using the same weighting matrices W used during model training. The minimized measure is the L_2 norm of the joint error vector. With such arrangement we minimize all errors at once. The optimization is stopped when changes between iterations drop below a minimum threshold.

Using this structure, we are interested in the influence of different features on the segmentation performance of myocardial first pass perfusion sequences. For this several features are proposed.

2.4 Features

We examined four temporal features characterizing the signal intensity curve without a strict need of temporal correspondence, and one spatial feature. All features are calculated pixel-wise from the sequences, as we assume to have priorly registered image series. Image examples of all features are presented in Figure 3, being

- **Maximum intensity projection (MIP)**

The maximum intensity projection (MIP) relates to the maximum blood flow through the cardiac tissue. This feature makes the myocardium separable for both right, and left ventricle. Low perfused myocardial areas are represented by a lower MIP value than normally perfused.

- **Average intensity over time**

A noise-robust feature to segment LV from myocardium is based on determining the average intensity over time. Though the average of the signal naturally varies with the selected time window (unlike the MIP), it is robust to small differences in this window selection. We selected the time window to approximately match the first pass.

- **Maximum up-slope**

The maximum upslope characterizes the rate of diffusion of the blood through the tissue. We derived this property by using a smoothed gradient filtering in the temporal dimension. Myocardium has lower slope values than RV or LV. Hypo-perfused myocardial areas are signified by a lower value than normal myocardium because of the occlusions hampering the blood flow.

- **Time to peak**

The time to Peak feature is based on the anatomical fact, that blood flows out from the LV through the coronary arteries into the myocardium. An obstruction of the coronary artery delays the bolus arrival, thus higher time to peak values indicate hypo-perfused myocardium.

We compute time to peak wrt the first frame of each sequence. This is possible because during model construction all feature images are normalized to zero mean, which compensates for offsets in the temporal dimension.

- **Orientation map**

We adopted the orientation map representation suggested by Cootes and Taylor in.^{9,10} This feature represents the orientation of the structures in the average image. Orientation is represented as $(\sin \phi, \cos \phi) =$

$(g_x/g; g_y/g)$ where $g = \sqrt{g_x^2 + g_y^2}$ is the length of the gradient. Noise deteriorates the gradient direction on planar areas, therefore a nonlinear weighting function is used:

$$F(g) = \frac{|g|}{|g| + g_0} \quad (3)$$

where g_0 is the mean of g . This sigmoid weighting suppresses homogeneous regions, and presumes edge regions. The orientation map results in two images, weighted gradient image in x , and y direction respectively. Examples are shown in Figure 3(e) and 3(f).




Figure 3. Examples of feature images: 3(a) a MIP image in the intensity range; 3(b) the average intensity over time; 3(c) maximum up-slope; 3(d) time to peak; 3(e) and 3(f) spatial gradient features in the x and the y direction respectively.




Figure 4. A MIP image with the indication of the grid of initial matching positions. Squares with a bright value did not converge, dark squares mark starting positions from which the model converged.

3. EXPERIMENTAL SETUP

We set up a validation experiment to answer two questions:

1. Which combination of features yields the best segmentation results in the multi-band AAM?
2. Is the achieved optimal performance good enough for clinical application?

To this end, 19 cardiac MR sequences were collected. The images were T1-weighted, with fast gradient echo pulse acquisition (TR = 2 ms, TE = 1.2 ms, flip angle 18 degrees). Participants were asked to hold their breath for 12 to 18 seconds. A Gadolinium- DTPA (Gd-DTPA) bolus of 0.04 mmol per kg of body weight was injected, starting at the third or fourth heartbeat, followed by a saline flush of 10 ml. The in-plane resolution of the scans was $1.37 \times 1.37 \text{mm}^2$, resulting in an image size of 256×204 pixels. In total 35 heart beats were processed for each participant, representing the first pass of the contrast agent. The myocardium of the mid-ventricular slices was segmented by medical experts, where the endocardial and epicardial contours were defined on one reference frame of the sequence, and subsequently adjusted to the different frames. As a pre-processing step, breathing motion was corrected by aligning the centers of gravity of the endocardial contours in the sequence.

The 2D shapes were represented by 60 landmark points. Simple arc length point correspondence was used because of the low complexity of the shape. For texture modeling, a mask containing an outer band, and the left ventricle was used.

Evaluation was performed in a leave-one-out manner, where the left out sample was used to test the performance of a model that was trained on the remaining data samples. The performance of the multi-band AAM for segmentation of the perfusion images was evaluated for the following feature combinations: MIP, MIP + time to peak intensity, MIP + maximum upslope, MIP + average intensity, MIP + maximum upslope + average intensity, MIP + orientation map. The proposed weighting was applied to the MIP + time to peak combination because of the large noise present in the time to peak images. These were thus down-weighted so as to describe the same variance as the MIP feature. Weighting was also applied at the MIP + Orientation map combination, as the two different normalizations (linear for MIP, and non-linear for Orientation map) did result in a total variance difference of orders of magnitude.

Segmentation performance was evaluated on the following criteria:

- point-to-point error (endo- and epicardial contour) to evaluate matching performance
- point-to-curve error (endo- and epicardial contour) to evaluate the clinical contour quality
- area error (for endo- and epicardial contour separately) to evaluate the accuracy of area estimates
- statistical significance of the abovementioned errors was assessed using a paired samples t-test
- robustness to model initialization was tested to measure convergence envelopes around the initial position. This was tested by using a 7x7 grid of initial positions centered around the manually set initialization midpoint, with a grid spacing of 5 pixels. Thus, in total $19 \times 7 \times 7 = 931$ matching runs were performed for each tested feature combination, and the total number of matching failures was recorded. Matching failure was declared when the endocardial mean point to curve error exceeded 4 pixels. The threshold was determined empirically, by visual inspection of converging and non-converging cases.

All evaluation criteria except convergence tests were calculated for converging cases only. Each test was initialized by the mean model parameters, zero rotation, and average scaling.

4. RESULTS AND CONCLUSIONS

Table 1 summarizes the evaluation of the segmentation performance of the multi-band AAM with different feature combinations. From Table 1, it can be concluded that

- the MIP+orientation map feature combination yields the best performance in terms of convergence rate (lowest number of failures), lowest point-to-curve errors and lowest area errors; this feature combination clearly outperforms the other tested feature combinations. Improvements compared with the standalone MIP feature are mainly visible in the epicardial contours.
- the inclusion of multiple temporal features does not necessarily improve the segmentation results. This may be due to a certain redundancy between the MIP feature and the other temporal features, but also due to the noisy nature of features such as time-to-peak, which may actually deteriorate segmentation performance.
- the multi-band AAM driven by MIP+orientation map features yields a mean point-to-curve error of less than one pixel, which is in the same order of magnitude as inter-observer variation in manual cardiac MR contour drawing;¹¹ therefore, the achieved results can be regarded as adequate for clinical application.

A study with a down-weighted MIP and up-weighted orientation map combination showed that the stand-alone orientation map feature gives similar boundary positioning errors, though with a considerably smaller convergence radius (convergence failure in 401 cases).

Some examples of segmentation results are shown in Figure 5. The red contours in the left column show the starting positions of the segmentations displayed in the right columns respectively. The manual segmentation is drawn with yellow, while the results of model matching are represented by green color.

Systems	w	fail	endocardial						epicardial						area error		
			point-point			point-curve			point-point			point-curve			mean	std	s
			mean	std	s	mean	std	s	mean	std	s	mean	std	s			
MIP	-	290	3.37	2.86		1.02	0.50		3.93	3.65		1.14	0.81		257.44	116.81	
MIP-T2P	-	436	3.13	2.05	0	<i>1.06</i>	<i>0.34</i>	1	3.54	2.67	0	1.03	0.43	0	253.32	76.65	0
MIP-T2P	+	372	2.95	1.66	0	0.99	0.37	0	3.40	2.37	0	1.04	0.50	0	246.41	81.87	0
MIP-Average	-	297	3.09	2.09	0	1.01	0.44	0	3.57	2.96	0	0.99	0.72	0	241.39	130.69	0
MIP-Slope	-	327	2.87	1.82	0	0.99	0.46	0	3.35	2.53	0	1.12	0.61	0	257.08	114.13	0
MIP-Average-Slope	-	308	<i>3.36</i>	<i>1.86</i>	1	<i>1.25</i>	<i>0.61</i>	1	3.84	2.69	0	<i>1.23</i>	<i>0.69</i>	1	<i>296.54</i>	<i>159.15</i>	1
MIP-Orientation map	+	261	3.10	1.96	0	0.99	0.47	0	3.60	2.84	0	0.88	<i>0.44</i>	1	224.85	<i>92.25</i>	1

Table 1. Evaluation results of the different feature combinations. The column w denotes if weighting was applied when building the joint model. The columns s state whether the changes are statistically significant using the paired t-test at a 5% confidence level. The statistically significant changes are marked in italics; the best performances of each metric are marked by bold numbers. The column $fail$ indicates the number of trials in which the system did not converge. Boundary point errors are expressed in pixels, areas in pixel².




Figure 5. Segmentation examples of the MIP and the MIP-orientation map combinations. Both systems were started with the same initial pose parameter setup, shown with the red curve in the first column. Model parameters were initialized to the mean values. The yellow curves are the manual segmentations regarded as correct, the green curves are the AAM segmentation results. The first row shows the 5th leave one out case, the second the 6th, the third the 7th. The last row shows a bad match: leave one out case 16. This indicates, that such narrow contour shapes were probably not included in the training set.

Figure 6 shows the epicardial paired mean point to curve errors of MIP and MIP + Orientation map combined features for each leave-one-out case. At leave one out case 7 we see a decrease in the performance. Depending

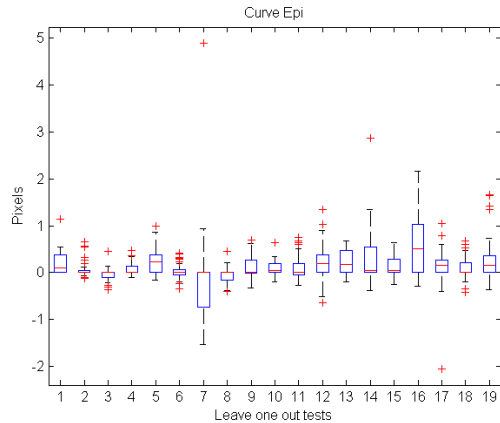


Figure 6. Box plot of the paired mean point to curve epicardial error differences between MIP and MIP + Orientation map combination. The value zero means no change compared to the MIP feature, positive values indicate performance increase, negative values indicate performance decrease.




Figure 7. Distribution of the point to curve errors of the MIP + Orientation map combination plotted on the mean shape. The radii of the circles drawn on the landmark points correlate with their mean point to curve error. The mean errors range from 0.7 to 1.2 pixels.

on the initialization, the model update tends to stop in a local minimum, and thus may converge to a false edge.

Figure 7 shows the distributions of the errors for all landmark points on the shape. The radii of the circles drawn on the landmark points correlate with their mean point to curve error. The mean errors range from 0.7 to 1.2 pixels, and show an approximately even distribution along the contour. The point to point errors are on average about 3 times as high as the point to curve errors in all tested systems. This is the consequence of a poor angular fit. The angular fit could be improved by adding the right ventricle to the model. Alternatively, whiskers may be added to the appearance model. Whiskers are line segments, which extend the intensity mask perpendicular to the contour. They were shown to increase robustness and lock-in range.⁷

AAMs rely on global optimization, meaning that a refinement at one landmark point will lead to changes in all landmark points. Therefore segmentation could also be improved by local refinement steps applied to the AAM contour. To perform the local refinements Oost¹² uses dynamic programming, whereas Stegmann¹³ uses simulated annealing as post-processing.

A further validation in terms of clinical perfusion parameters is ongoing, as well as a study into the effects of non-rigid shape deformations that may occur due to large breathing motions.

5. SUMMARY

This work introduced a modular architecture for joining multiple features of the object of interest in the AAM framework. Several MR perfusion features were tested in a multi-band AAM configuration using a set of 19 perfusion sequences. The key question we investigated was whether combination of temporal and spatial features would improve the matching over the basic intensity AAM. The proposed architecture was shown to perform well in joint texture representation and segmentation. Experiments showed that the combined orientation map and intensity model improved the segmentation performance on first pass myocardial perfusion sequences. The resulting mean point to curve error is comparable with the intra-observer variability. The combination of two or more temporal features did not significantly improve segmentation accuracy. This might be because of the amount of noise present in the features, but also because of the high correlations between different temporal features. Further improvements could arise with a different update scheme, normalization, and integration of the RV.

Acknowledgements

The authors gratefully acknowledge the financial support of the Netherlands Organisation for Scientific Research (NWO), grant number 612.065.618.

REFERENCES

1. T.-H. Lim and S. I. Choi, "Mri of myocardial infarction," *Journal of Magnetic Resonance Imaging* **10**, pp. 686–693, 1999.
2. M. Jerosch-Herold, R. T. Seethamraju, C. M. Swingen, N. M. Wilke, and A. E. Stillman, "Analysis of myocardial perfusion mri," *Journal of Magnetic Resonance Imaging* **19**, pp. 758–770, 2004.
3. J. Milles, R. van der Geest, M. Jerosch-Herold, J. Reiber, and B. Lelieveldt, "Fully automated registration of first-pass myocardial perfusion mri using independent component analysis," in *Lecture Notes in Computer Science, IPMI*, **4584**, pp. 544–555, 2007.
4. T. Cootes and C.J.Taylor, "Statistical models of appearance for computer vision," technical report, Imaging Science and Biomedical Engineering, University of Manchester, Manchester M13 9PT, U.K., 2004.
5. J. G. Bosch, S. C. Mitchell, B. P. 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* **21**, pp. 1374–1384, November 2002.
6. R. J. van der Geest, B. P. F. Lelieveldt, E. Angelié, D. Mikhail, 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 an active appearance motion model," *Journal of Cardiovascular Magnetic Resonance* **6**, pp. 609–617, July 2004.
7. M. B. Stegmann, *Generative Interpretation of Medical Images*. PhD thesis, Technical University of Denmark, Denmark, 2004.
8. M. Stegmann and R. Larsen, "Multi-band modelling of appearance," in *First International Workshop on Generative Model-Based Vision GMBV*, **6**, pp. 101–106, 2002.
9. I. Scott, T. Cootes, and C. Taylor, "Improving appearance model matching using local image structure," in *Proc. Information Processing in Medical Imaging*, pp. 258–269, 2003.
10. T. Cootes and C. Taylor, "On representing edge structure for model matching," in *IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, 2001.
11. M. Üzümcü, R. van der Geest, M. Sonka, H. J. Lamb, J. H. C. Reiber, and B. P. F. Lelieveldt, "Multiview active appearance models for simultaneous segmentation of cardiac 2- and 4-chamber long-axis magnetic resonance images.," *Investigative Radiology* **40(4)**, pp. 195–203, 2005.
12. E. Oost, G. Koning, M. Sonka, P. V. Oemrawsingh, J. H. C. Reiber, and B. P. Lelieveldt, "Automated contour detection in x-ray left ventricular angiograms using multiview active appearance models and dynamic programming," *IEEE transactions on medical imaging* **25**, pp. 1158–1171, September 2006.
13. M. B. Stegmann, R. Fisker, and B. K. Ersbøll, "Extending and applying active appearance models for automated, high precision segmentation in different image modalities," in *Proc. 12th Scandinavian Conference on Image Analysis - SCIA 2001, Bergen, Norway*, I. Austvoll, ed., pp. 90–97, NOBIM, (Stavanger, Norway), jun 2001.