Active shape models exploiting slice-to-slice correlation in segmentation of 3D CTA AAA images

Marleen de Bruijne, Bram van Ginneken, Wiro J. Niessen, J.B. Antoine Maintz, Max A. Viergever
Image Sciences Institute
University Medical Center Utrecht, Rm E.01.334
Heidelberglaan 100, 3584 CX Utrecht, The Netherlands
E-mail: marleen@isi.uu.nl

Abstract

An automated method for the segmentation of thrombus in abdominal aortic aneurysms (AAA) from CTA data is presented. Three segmentation schemes, inspired by Active Shape Model (ASM) segmentation, were investigated. (1) The original ASM scheme as proposed by Cootes and Taylor [1], applied to sequential slices, using the contour obtained in one slice as the initial contour in the adjacent slice. (2) A similar approach, steered by profile greyvalue correlation with adjacent slices rather than by correlation with profiles from the training data and (3) as in (2), with additional attraction to nearby edges.

A leave-one-out experiment was performed, using five datasets containing 239 slices to segment. Both adapted ASM schemes yield considerably better results than the original scheme. Scheme 3 showed best overall performance. Using one manually delineated image slice as a reference, on average a number of 21 slices could be automatically segmented with an accuracy within the bounds of manual inter-observer variability.

1. Introduction

In this work we consider the clinically important task of thrombus segmentation in abdominal aortic aneurysm (AAA). An AAA is an enlargement of the infrarenal abdominal aorta, resulting from weakened arterial walls. Once present, AAAs continue to enlarge and, if left untreated, become increasingly susceptible to rupture, usually resulting in death. Worldwide, approximately 100,000 surgical interventions for AAA repair are performed each year, of which 30 % is endovascular. After endovascular aneurysm repair, in which a synthetic graft is placed inside the aorta, the process of aneurysm shrinkage, ongoing aneurysmal disease, and damage or fatigue of graft material may result in leakage, graft migration, and kinking or buckling of the graft, which can subsequently cause rupture or occlusion. Careful and frequent patient follow-up is therefore required [2]. Each patient is scanned every three to twelve months, depending on the state of the aneurysm.

It has been demonstrated that change of aneurysm volume is a good indicator for the risk of aneurysm rupture [3]. Currently, the gold standard for volume assessment is spiral Computer Tomography Angiography (CTA). The standard follow-up procedure includes a manual aneurysm delineation [4], generally referred to as 'thrombus segmentation', since after successful graft placement the excluded aneurysm sac surrounding the graft is completely filled with thrombus. The manual segmentation is a time-consuming process — it takes an experienced operator around 30 minutes — and suffers from inter- and intra-operator variations. Wever et al. [5] reported averaged inter-observer volume errors of 8.3% and intra-observer errors, for two observers, of 3.2% and 5.8%.

Therefore, automated segmentation would be of great value. However, the images are difficult to segment, as the thrombus edge can be obscured by surrounding tissue of similar greyvalue. The boundary can also be very weak owing to quickly changing aneurysm radius and partial volume effects, while many neighbouring structures induce strong edges in close proximity to the aneurysm wall. The texture and greyvalue of the thrombus can vary with the presence of calcifications, graft metal, and leakage of intravenous contrast. In addition, the shape and size of aneurysms can vary considerably between patients as well as in one patient over time.

Most publications on automated AAA segmentation have concentrated on lumen segmentation, either in pre-operative [6, 7, 8] or in post-operative CTA scans [9]. The more difficult problem of thrombus segmentation is less frequently addressed.

Subasic et al. [10] propose the use of 2D level sets for thrombus boundary segmentation, using the output

of a 3D level set lumen segmentation as initialisation. They define a stopping criterion, based on curve expansion speed, that should prevent the level set from growing into surrounding tissue. Our experience with this kind of images in combination with level set methods is that the regions lacking boundary information are too large to be accounted for by curvature constraints only, and a more restrictive shape model is required.

In this paper, we investigate the applicability of statistical shape and appearance models to AAA segmentation.

We focus on Active Shape Models (ASM) as put forward by Cootes et al. [1], which combine statistical knowledge of object shape and shape variation with local appearance models near object contours. ASMs have been successfully applied to various segmentation tasks in medical imaging [11, 12, 13, 14, 15]. To our knowledge, no previous work has concentrated on ASM segmentation of vessel-like structures.

We present a method for AAA segmentation, based on automated 2D Active Shape Model (ASM) fitting in sequential slices. Cootes et al. [11] reported successful slice-by-slice segmentation of the ventricles in 3D MR datasets, using the original ASM scheme.

We propose an adaptation to this method, using greyvalue profile similarity between adjacent slices, rather than profile similarity with respect to training data, in the model fitting. In addition, a refinement step is proposed that allows the model to attract to nearby edges.

Section 2 briefly reviews the conventional ASM segmentation scheme. The adaptations we made for our application are discussed in section 3. The results of this method as applied to five CTA scans, compared with original ASM segmentation, are presented in Section 4. Conclusions and suggestions for future work are given in Section 5.

2. Active Shape Models

This section describes the ASM segmentation scheme as proposed by Cootes and Taylor [16]. Shape variations in a training set are described using a Point Distribution Model (PDM). The shape model is used to generate new shapes, similar to those found in the training set, that are fitted to the data using a model of local greyvalue structure.

2.1. Point Distribution Models

In PDMs, a statistical model of object shape and variation is derived from a set of s training examples. The shape of the training examples is described by a number of n landmark points, typically distributed evenly along the object contour. Variations in the coordinates of the landmark points describe the variation in shape and pose of the training examples.

A shape can be described by its shape vector containing all landmark coordinates, in 2D:

$$\mathbf{x} = (\mathrm{x}_1, \mathrm{y}_1, \mathrm{x}_2, \mathrm{y}_2, \ldots, \mathrm{x}_n, \mathrm{y}_n). \tag{1}$$

To reduce the dimensionality of the data, Principal Component Analysis (PCA) is applied, with which shapes are described by their deviation from the mean shape. The mean shape is given by

$$\overline{\mathbf{x}} = \frac{1}{s}\sum_{i=1}^{s}\mathbf{x}_i \tag{2}$$

and the $2n \times 2n$ covariance matrix is

$$\mathrm{S} = \frac{1}{s-1}\sum_{i=1}^{s}(\mathbf{x}_i - \overline{\mathbf{x}})(\mathbf{x}_i - \overline{\mathbf{x}})^T. \tag{3}$$

The eigenvectors λ_i of S provide the modes of shape variation present in the data. Any shape $\mathbf{x}$ in the dataset can be represented as

$$\mathbf{x} = \overline{\mathbf{x}} + \mathrm{P}\mathbf{b} \tag{4}$$

where P is the concatenation of eigenvectors of S, $\mathrm{P} = (\mathbf{p}_1|\mathbf{p}_2|\ldots|\mathbf{p}_{2n})$. The eigenvectors corresponding to the largest eigenvalues account for the largest shape variation. Usually, most of the variation can be explained by only a small number of modes. Each shape $\mathbf{x}$ in the set can thus be approximated by

$$\mathbf{x} \approx \overline{\mathbf{x}} + \Phi\mathbf{b} \tag{5}$$

where Φ consists of the eigenvectors corresponding to the t largest eigenvalues, $\Phi = (\phi_1|\phi_2|\ldots|\phi_t)$. The t-dimensional vector $\mathbf{b}$, containing the model parameters, indicates how much variation is exhibited with respect to each of the eigenvectors,

$$\mathbf{b} = \Phi^T(\mathbf{x} - \overline{\mathbf{x}}). \tag{6}$$

If we assume that the landmark points are normally distributed, λ_i is the variance of the i^{th} model parameter b_i across the training set, and 99% of all shapes that are in the training data can be described with b_i within the bounds of $\pm 3\sqrt{\lambda_i}$.

The total variance in the dataset is given by $\sum\lambda_i$. The number t of modes in the model is chosen such that the model captures a certain proportion f_v of the total variance observed:

$$\sum_{i=1}^{t}\lambda_i \geq f_v\sum_{i=1}^{2n}\lambda_i. \tag{7}$$

2.2. Modelling local greyvalue structure

The appearance of the training images near object contours is modelled using greyvalue profiles around the landmark points. A profile $\mathbf{g}_i$ of k pixels is sampled on either side of the landmark, perpendicular to the contour. The effect of global intensity changes is reduced by sampling the first derivative and normalising the profile.

The profiles are sampled for multiple resolutions. The finest resolution uses the original image and a step size of one pixel, the next resolution is the image observed at scale σ of one pixel (obtained by convolution with a Gaussian kernel of width σ) and a step size of two pixels. Subsequent resolutions are obtained by doubling both the image scale and the step size. The normalised samples for one landmark and for all training shapes are assumed to be distributed as a multivariate Gaussian, and the mean $\overline{\mathbf{g}}$ and covariance $\mathbf{S}_g$ are computed.

The measure of similarity of a new profile $\mathbf{g}_s$ to the profiles in the distribution is given by the Mahalanobis distance $f(\mathbf{g}_s)$ from the sample to the model mean:

$$f(\mathbf{g}_s) = (\mathbf{g}_s - \overline{\mathbf{g}})^T \mathbf{S}_g^{-1} (\mathbf{g}_s - \overline{\mathbf{g}}). \quad (8)$$

Minimising $f(\mathbf{g}_s)$ is equivalent to maximising the probability that $\mathbf{g}_s$ originates from the profile distribution.

2.3. Active Shape Model fitting

The fitting of the shape model to the image initialises with the mean shape. From coarse to fine resolution, the model is fitted to the image in an iterative procedure. For all landmarks, a number n_s of possible new positions along the line perpendicular to the contour, are evaluated. The optimal position is the location where the Mahalanobis distance from the local profile to the model is minimal. The model parameters are adjusted to fit the optimal landmark positions. This process is repeated a fixed number of times N or until changes in the model parameters become negligible, whereupon it is repeated at the next level of resolution.

3. ASM applied to 3D AAA segmentation

3.1. Multislice 2D versus 3D

The idea of ASM segmentation can be readily extended to 3D. However, a few difficulties arise:

- The ordering of landmarks is not straightforward in 3D.
- To capture all possible shape variations of a 3D object, many datasets are needed.
- Existing manual segmentations of 3D objects are in practice obtained through delineation in 2D slices. Building a model from a stack of slices results in irregular slice-to-slice shape variations that should not be modelled.

Moreover, in our application, the length of the objects to be segmented can vary considerably. To keep the number of landmarks constant for all datasets, resampling along the aneurysm axis would be needed, inducing differences in greyvalue statistics.

On basis of these considerations, we propose to attack the problem of 3D thrombus segmentation in a slice-by-slice manner. The slices that we use are the original CT-slices, which are perpendicular to the body axis and therefore give approximately perpendicular cross-sectional views of the aorta. We start with one slice in the middle of the aneurysm, segmented manually, as a reference contour. On basis of correlation with local image structure around this reference contour, a plausible shape from the shape model is fitted in both adjacent slices. For the succeeding slices, the process is repeated, each time using the previous slice as the reference.

3.2. Landmarking and alignment

In the absence of well-defined anatomical landmark points, we choose the centre point of the central slice of the segmented object as a reference point. All training shapes are translated such that their reference points coincide. This accounts for the largest inter-subject difference. Any translation in other slices with respect to the reference point will be modelled.

The landmarks are placed equidistant along the object contours, starting at the point with the largest y-coordinate, which is usually the point nearest to the vertebra. Any rotation or scaling with respect to the reference point is contained within the model.

3.3. Maximising slice-to-slice correlation

For each landmark, an image patch of k_n pixels on either side of the landmark point in the normal direction and k_t pixels in the tangential direction is sampled from the image, at multiple levels of resolution, similar to the method described in 2.2. The optimal landmark position will now be the position where the sampled patch is most similar to the corresponding patch sampled from the reference slice for that landmark. The similarity measure used is the normalised cross corre-

lation, *Ncc*:

$$Ncc = \frac{\sum_{x=0}^{2k_t+1} \sum_{y=0}^{2k_n+1} I_{s-1}(x,y) \cdot I_s(x,y)}{\sqrt{\sum_{x=0}^{2k_t+1} \sum_{y=0}^{2k_n+1} I_{s-1}(x,y)^2 \cdot \sum_{x=0}^{2k_t+1} \sum_{y=0}^{2k_n+1} I_s(x,y)^2}} \tag{9}$$

where I and I_0 are the pixel intensities in the examined patch and in the reference patch respectively. The *Ncc* is maximised in the fitting procedure.

3.4. Incorporating gradient information

The nature of the images provides that, moving along the axis, almost all landmark points will encounter several slices where image edge evidence is absent, whereas it may be present again in the following slice. As we maximise local similarity with respect to the reference slice, we risk the landmarks being pushed away from the 'new' edge, while in fact we know that the thrombus wall does coincide with an edge in most cases. We try to overcome this problem by employing local gradient information in the fitting procedure. Since the thrombus wall can be either a transition to lighter (bone, contrast enhanced vessels) or to darker tissue, we will use the absolute gradient value.

First, the model is fitted as described in the previous section, using local intensity correlation with the reference slice, at all levels of resolution. Subsequently the model is allowed to refine to nearby strong edges at the finest resolution. A number n_{rs} of possible new landmark positions along the line normal to the boundary is evaluated. The derivative along the normal is computed using finite differences, and the function to be maximised in the fitting procedure is given by:

$$f(x) = |I(x+1) - I(x-1)| \cdot Ncc. \tag{10}$$

4. Experiments and results

Our database consists of CTA scans with accompanying binary volume segmentations that were obtained through manual delineation by an expert. The scan resolution is $0.485 \times 0.485 \times 2.0\,\mathrm{mm}$. The images consist of circa 125 slices of 512×512 voxels, of which 25 to 60 slices contain aneurysmal tissue. In this work, we have analysed five CTA images of different patients, including both pre- and post-operative scans. We performed a leave-one-out experiment: five shape models are built of four datasets, with which the fifth dataset, not included in the model, is segmented. Figure 1 visualises the first six modes of the shape model trained with all five datasets.

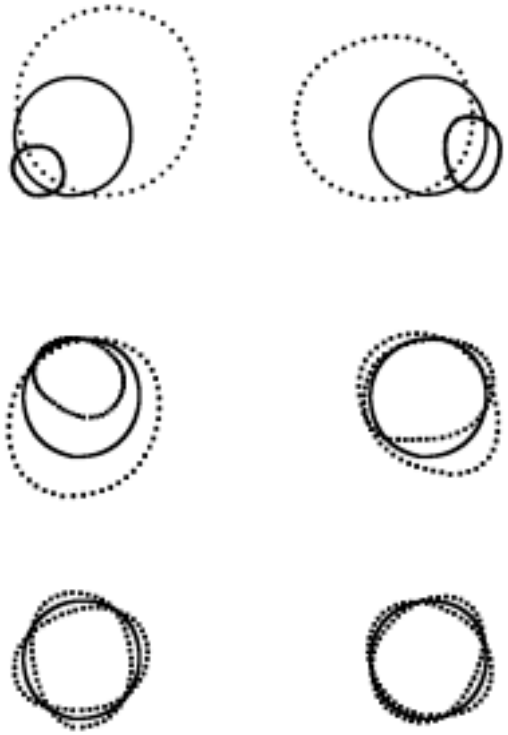

Figure 1: The effect of varying the first six modes of shape variation individually in a model built from all five datasets, using a total of 239 slices. The mean shape (solid line) and both allowed extremes $\pm 3\sqrt{\lambda_i}$ (dotted lines) are shown.

We have investigated three schemes for evaluating the optimal landmark positions:

1. Statistical model of local image structure (original ASM, Section 2.2)
2. Greyvalue correlation with adjacent slices (Section 3.3)
3. Greyvalue correlation with adjacent slices, combined with local gradient information (Section 3.4)

The parameters used in all experiments are listed in Table 1.

The similarity, or relative volume of overlap, of two volumes A and B is defined by

$$2\,\frac{A \cap B}{A + B}. \tag{11}$$

Figure 2 shows the similarity with respect to the manual segmentations that were obtained for the three schemes, for all scans and slices. The horizontal

Table 1: Parameters for the three segmentation schemes described; the values used in this report are given between parentheses.

General	
s	Number of training image slices, differing per segmented volume (~ 200)
Shape model	
n	Number of landmark points (50)
f_v	Part of shape variance to be explained by the model (0.999), controlling the number of modes t (9–11)
m	Bounds on eigenvalues λ_i (3)
Appearance model	
k	Number of points in profile to sample on either side of the landmark point, scheme 1 (5)
k_n	Number of points in patch to sample on either side of the landmark point in the normal direction, scheme 2 and 3 (5)
k_t	Number of points in patch to sample on either side of the landmark point in the tangential direction, scheme 2 and 3 (1)
Fitting algorithm	
n_s	Number of new landmark positions to examine per iteration on either side of the current position (3)
n_{rs}	Number of new landmark positions to examine per refinement iteration on either side of the current position in scheme 3 (2)
L	Number of resolution levels (3)
N	Number of iterations per resolution level (5)

line indicates the level of inter-observer reproducibility (91.7%)[5]. Note, that the inter-observer errors as quoted are not completely comparable to the relative volume of overlap in our experiments. The study in [5] describes the relative volume error per volume instead of per slice, without taking the overlap into account.

In one case (dataset IV) the original ASM scheme yields results that are comparable to those of the other two schemes. In all four other datasets the results obtained by the local correlation schemes are considerably better.

Incorporating the local gradient term in the fit-function improved the results substantially for some datasets (sets I and V, an example is given in Figure 3), the results were similar in sets III and IV, but less good in set II. Closer investigation of set II showed that this is coincidental, as is shown in Figure 4. This particular dataset shows strong intensity variations inside the object, near the edge. Both profile correlation based schemes are sensitive to this kind of distortions,

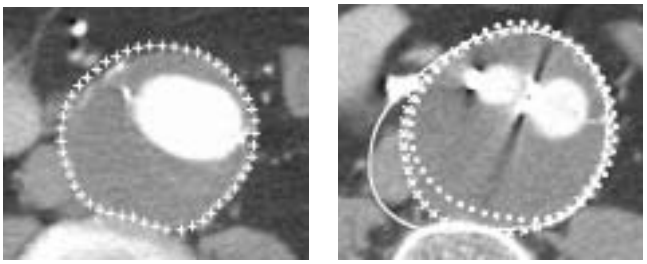

Figure 3: An example where the local gradient term in the fit function improves the results. The left image is the reference slice of set V, the image on the right is taken six slices lower. Scheme 1 (solid line) has drawn to the edge of the neighbouring vena cava, which is more similar to the mean thrombus edge for this landmark than the true edge that is hardly distinguishable. Scheme 2 (circles) is pushed away from the nearby vertebra, which shows up brighter than in the previous slices. Lastly, scheme 3 (crosses) draws the model to the nearest edge, which is the proper one.

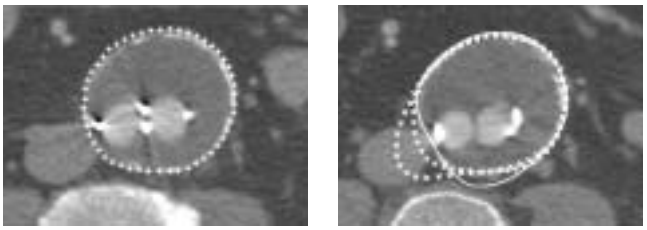

Figure 4: An example, taken from set II, where both scheme 1 (solid line) and scheme 2 (circles) outperform scheme 3 (crosses). The left image is the reference slice, the image on the right is taken several slices lower. The bright spots within the aneurysm are radiopaque markers on the graft, which have been absent in one slice. Both profile correlation based methods are pushed away from the marker, but the one with incorporating gradient information has more iterations and therefore has more time to expand from the initialisation, while the edge that it should attract to is at too far a distance. Scheme 1 uses normalised gradient profiles and is therefore less sensitive to this kind of intensity variations.

but scheme 3 uses extra iterations in the refinement step that allow the model to expand further. Overall, the edge-seeking scheme performs best.

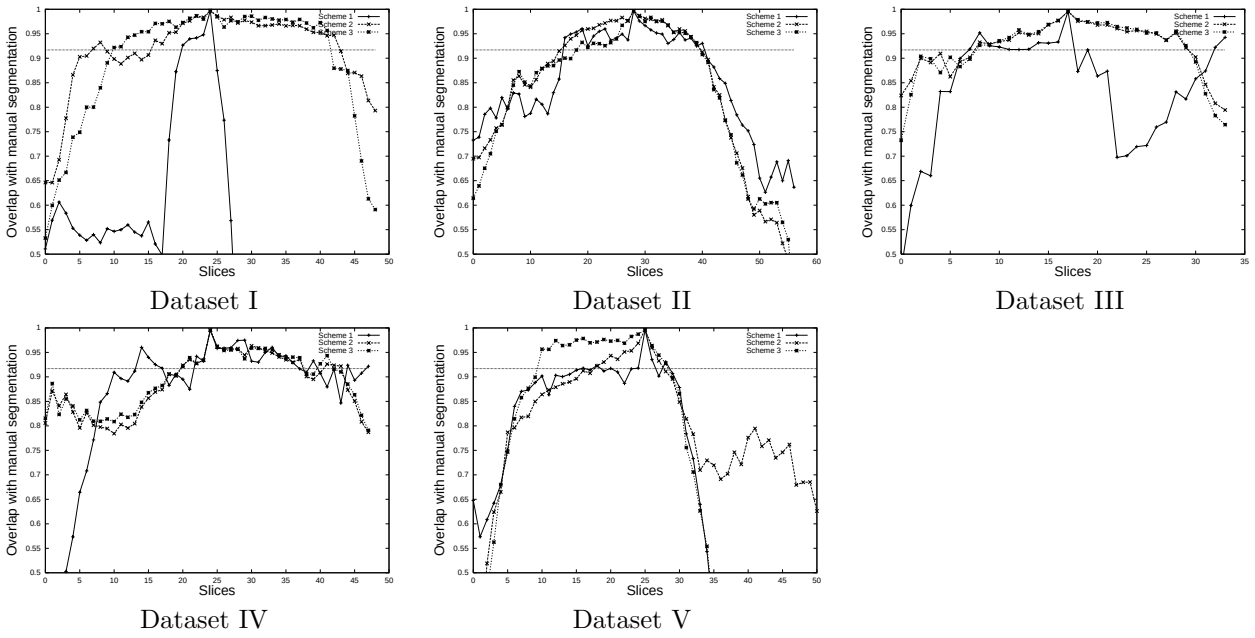

Figure 2: Results for datasets I through V. The relative volume of overlap with respect to the manual segmentation is plotted for each slice, for the segmentations obtained by all three schemes. The slice in the middle of the plot, with overlap 1, is the reference slice from which sequential slices are segmented in both directions. The horizontal line indicates the level of inter-observer reproducibility of the manual segmentations.

5. Conclusions and discussion

A method has been implemented to automate the segmentation of thrombus in abdominal aorta aneurysms in CTA datasets. The method is inspired by the ASM segmentation scheme. Two modifications with respect to the conventional ASM approach were evaluated:

1. Correlation with greyvalue profiles of adjacent slices, rather than greyvalue profiles obtained from the training set, was used to determine the optimal landmark positions.

2. A refinement step to locally adjust the position of the landmarks to points with maximum gradient was added to the scheme using greyvalue correlation of adjacent slices.

In a leave-one-out experiment with five datasets, comprising a total number of 239 aneurysm slices, it was shown that the introduction of both modifications outperformed the conventional ASM in four out of five cases. On average 21 slices could be segmented with an accuracy at least as good as the inter-observer reproducibility of manual delineation, while the user had to draw only one contour. These results suggest that the modified ASM scheme exploiting correlation between neighbouring slices could be valuable in speeding up the manual segmentation process.

Still, several improvements are possible.

- The used similarity measure in scheme 2 and 3, *Ncc*, may not be the best measure for maximising the probability that both profiles are located at the same anatomical edge in sequential slices. Two patches with a small intensity difference for all pixels may yield a lower *Ncc*, thus a better fit, than two patches with a large intensity difference for a few pixels (for example a new background structure showing up) and equal intensities for the remaining pixels. The use of other similarity measures, for example histogram-based measures, might improve the results. What is a good similarity measure could be derived from the training data.

- It was shown that the original ASM scheme, using local profile models, does not perform well on these images. This can be explained since the measure of fit in the original scheme, the Mahalanobis distance to the model of profiles, assumes a normal profile distribution. This assumption does not hold in the presence of distinct background structures of varying shape and brightness. The use of

non-linear models could be an improvement.

- Scheme 1 on the one hand and schemes 2 and 3 on the other hand are extreme cases, where in scheme 1 no information of neighbouring slices is used, while in schemes 2 and 3 no model information is used. As a consequence, segmentation errors are propagated through the dataset in schemes 1 and 2. Once the optimal contour is lost, it is not likely to be recovered in the remaining slices. A hybrid scheme combining both model information and information of neighbouring slices could improve overall performance. The variation of profiles between sequential slices should also be taken into account.

- In the presented scheme, the contours in all slices are taken as completely independent, apart from their initialisation, while many shapes in the model could be excluded on basis of the shape in the previous slice. The possible shape variation between slices could be modelled.

- Apart from correlation with reference or average profiles, or optimising edge information, there are many other ways to decide whether a point is located on an anatomical edge. For instance, the intensity and textural appearance inside the thrombus is only implicitly employed in the schemes described, while this seems to be an important clue for a human observer in the thrombus segmentation task. An ASM scheme steered by optimal features that were deduced from the training set, as proposed in [17], may be more suitable.

References

[1] T. F. Cootes, C. J. Taylor, D. H. Cooper, and J. Graham, "Active shape models – their training and application," *Computer Vision and Image Understanding*, vol. 61, no. 1, pp. 38–59, 1995.

[2] I. A. M. J. Broeders, J. D. Blankensteijn, A. Gvakharia, J. May, P. R. F. Bell, J. Swedenborg, J. Collin, and B. C. Eikelboom, "The efficacy of transfemoral endovascular aneurysm management: A study on size changes of the abdominal aorta during mid-term follow-up," *European Journal of Vascular and Endovascular Surgery*, vol. 14, no. 2, pp. 84–90, 1997.

[3] J. J. Wever, J. D. Blankensteijn, W. P. Mali, and B. C. Eikelboom, "Maximum aneurysm diameter follow-up is inadequate after endovascular abdominal aortic aneurysm repair," *European Journal of Vascular and Endovascular Surgery*, vol. 20, no. 2, pp. 177–182, 2000.

[4] R. Balm, R. Kaatee, J. D. Blankensteijn, W. P. T. M. Mali, and B. C. Eikelboom, "CT-angiography of abdominal aortic aneurysms after transfemoral endovascular aneurysm management," *European Journal of Vascular and Endovascular Surgery*, vol. 12, no. 2, pp. 182–188, 1996.

[5] J. J. Wever, J. D. Blankensteijn, J. C. van Rijn, I. A. J. M. Broeders, B. C. Eikelboom, and W. P. T. M. Mali, "Inter- and intra-observer variability of CTA measurements obtained after endovascular repair of abdominal aortic aneurysms," *American Journal of Roentgenology*, vol. 175, no. 5, pp. 1297–1282, 2000.

[6] M. Fiebich, M. M. Tomiak, R. M. Engelmann, and J. M. K. R. Hoffman, "Computer assisted diagnosis in CT angiography of abdominal aortic aneurysms," in *Proc. SPIE*, vol. 3034, 1997.

[7] V. Juhan, B. Nazarian, K. Malkani, R. Bulot, J. M. Bartoli, and J. Sequeira, "Geometrical modelling of abdominal aortic aneurysms," in *Proc. CVRMed and MRCAS*, no. 1205 in Lecture Notes in Computer Science, pp. 243–252, Springer Verlag, Berlin, 1997.

[8] O. Wink, W. J. Niessen, and M. A. Viergever, "Fast delineation and visualization of vessels in 3-D angiographic images," *IEEE Transactions on Medical Imaging*, vol. 19, no. 4, pp. 337–346, 2000.

[9] M. de Bruijne, W. J. Niessen, J. B. A. Maintz, and M. A. Viergever, "Semi-automatic aortic endograft localisation for post-operative evaluation of endovascular aneurysm treatment," in *Proc. SPIE Medical Imaging: Image Processing*, 2001.

[10] M. Subasic, S. Loncaric, and E. Sorantin, "3D image analysis of abdominal aortic aneurysm," in *Proc. SPIE Medical Imaging: Image Processing 2001*, 2001.

[11] T. F. Cootes, A. Hill, C. J. Taylor, and J. Haslam, "The use of active shape models for locating structures in medical images," *Image and Vision Computing*, vol. 12, no. 6, pp. 355–366, 1994.

[12] S. Solloway, C. Taylor, C. Hutchinson, and J. Waterton, "The use of active shape models for making thickness measurements from MR images," in *Proceedings of the 4th European Conference on Computer Vision*, pp. 400–412, 1996.

[13] N. Duta and M. Sonka, "Segmentation and interpretation of MR brain images: An improved active shape model," *IEEE Transactions on Medical Imaging*, vol. 17, no. 6, pp. 1049–1067, 1998.

[14] G. Behiels, D. Vandermeulen, F. Maes, P. Suetens, and P. Dewaele, "Active shape model-based segmentation of digital X-ray images," in *Medical Image Computing and Computer-Assisted Intervention*, pp. 128–137, 1999.

[15] G. Hamarneh and T. Gustavsson, "Combining snakes and active shape models for segmenting the human left ventricle in echocardiographic images," *IEEE Computers in Cardiology*, vol. 27, pp. 115–118, 2000.

[16] T. F. Cootes and C. J. Taylor, "Statistical models of appearance for computer vision," tech. rep., University of Manchester, 2000.

[17] B. van Ginneken, *Computer Aided Diagnosis in Chest Radiography*. PhD thesis, University of Utrecht, March 2001.