Active shape model based segmentation of abdominal aortic aneurysms in CTA images

Marleen de Bruijne, Bram van Ginneken, Wiro J. Niessen, J.B. Antoine Maintz, Max A. Viergever

Image Sciences Institute, University Medical Center Utrecht, The Netherlands

ABSTRACT

An automated method for the segmentation of thrombus in abdominal aortic aneurysms from CTA data is presented. The method is based on Active Shape Model (ASM) fitting in sequential slices, using the contour obtained in one slice as the initialisation in the adjacent slice. The optimal fit is defined by maximum correlation of grey value profiles around the contour in successive slices, in contrast to the original ASM scheme as proposed by Cootes and Taylor,[1] where the correlation with profiles from training data is maximised. An extension to the proposed approach prevents the inclusion of low-intensity tissue and allows the model to refine to nearby edges. The applied shape models contain either one or two image slices, the latter explicitly restricting the shape change from slice to slice.

To evaluate the proposed methods a leave-one-out experiment was performed, using six datasets containing 274 slices to segment. Both adapted ASM schemes yield significantly better results than the original scheme (p <0.0001). The extended slice correlation fit of a one-slice model showed best overall performance. Using one manually delineated image slice as a reference, on average a number of 29 slices could be automatically segmented with an accuracy within the bounds of manual inter-observer variability.

Keywords: active shape model, image segmentation, contour matching, CT, abdominal aortic aneurysm, blood vessels

1. INTRODUCTION

An abdominal aortic aneurysm (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, which usually results in death. Worldwide, approximately 100,000 surgical interventions for AAA repair are performed each year, of which 30 % are 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 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%.

To reduce analysis time and increase reproducibility, automated segmentation would be of great value. However, the images are difficult to segment. The thrombus boundary can be obscured by surrounding tissue of similar greyvalue, and many neighbouring structures induce strong edges in close proximity to the aneurysm wall. Quickly changing aneurysm radius in combination with partial volume effects makes the boundary hard to distinguish even in the absence of neighbouring structures. Thrombus texture and greyvalue can vary with the

Further author information: M.B.: E-mail: marleen@isi.uu.nl

Medical Imaging 2002: Image Processing, Milan Sonka, J. Michael Fitzpatrick,
Editors, Proceedings of SPIE Vol. 4684 (2002) © 2002 SPIE · 1605-7422/02/$15.00

presence of calcifications, graft metal, and intravenous contrast. In addition, the shape and size of aneurysms vary considerably between patients as well as in one patient over time. Figure 1 shows some examples of CTA slices, illustrating the diversity of aneurysm and background appearance.

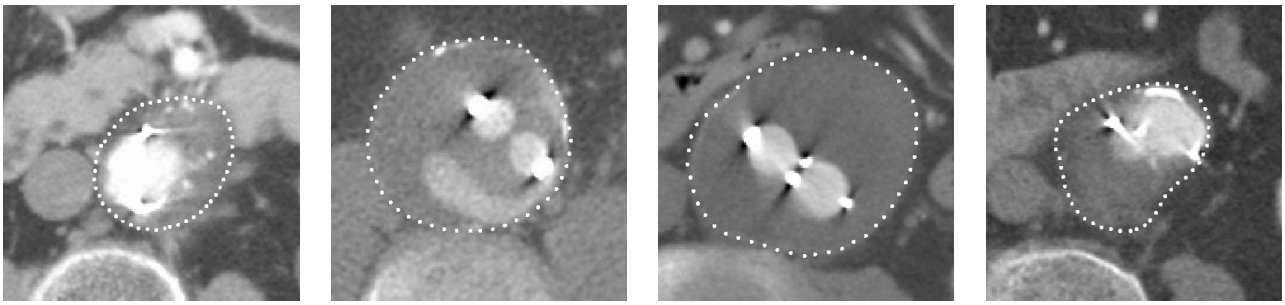

Figure 1: Four CTA slices of different patients, the aneurysm delineated manually.

Most publications on automated AAA segmentation have concentrated on lumen segmentation, either in pre-operative[6–9] or in post-operative CTA scans.[10] The more difficult problem of thrombus segmentation is less frequently addressed. Subasic et al.[11] 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 these images in combination with level set methods is that curvature constraints alone can not cope with the large regions lacking boundary evidence, and a more restrictive shape model is required.

In this work, 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 and Taylor,[1, 12, 13] 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.[14–21] Applications in 3D have either used a full 3D model[14] or a 2D model applied to image slices.[15, 20] Mitchell et al.[20] proposed the use of a hybrid active shape model/active appearance model, fitted to several cardiac MR image slices independently. Cootes et al.[15] reported successful slice-by-slice ASM segmentation of the ventricles in 3D MR datasets, using the fitted shape in one slice as the initialisation in the next slice.

In the case of AAA segmentation, the length of the objects to be segmented can vary considerably. To keep the number of landmarks in a full 3D model constant for all datasets, resampling along the aneurysm axis would be needed, either reducing the amount of information used for fitting — thus decreasing robustness — or inducing differences in greyvalue statistics. We therefore approach the problem of 3D thrombus segmentation in a slice-by-slice manner, similar to the method described by Cootes et al.[15]

Conventional ASMs use a model generated from greyvalue profiles in training images to fit the shape model to the image. AAAs, being situated in the abdomen, have such a diverse range of possible surrounding structures, varying in location, shape, intensity and texture, between patients as well as in one patient over time, that a linear model of greylevel structure about the contours may not be able to find the true contour in a new image. However, an important indicator for the image structure in one image slice is the structure in the adjacent slices. We propose a modified ASM scheme in which optimal landmark positions are defined by maximum greyvalue profile correlation with adjacent slices rather than by correlation with profiles from the training data. In addition, a penalty function that prevents the inclusion of low-intensity tissue, and a refinement step, which allows the model to attract to nearby edges, are suggested.

Section 2 briefly reviews the conventional ASM segmentation scheme. The adaptations we propose for the segmentation of AAA are discussed in section 3. The results of these methods as applied to six CTA scans, compared with original ASM segmentation and manually drawn contours, are presented in Section 4. A discussion and conclusions are given in Sections 5 and 6.

2. ACTIVE SHAPE MODELS

This section describes the ASM segmentation scheme as proposed by Cootes and Taylor. 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, which must be identified in each example. 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:

$$\mathbf{x} = (x_1, y_1, x_2, y_2, \ldots, x_n, y_n). \tag{1}$$

If we wish to allow translation, rotation, and scaling when fitting the model to a new image, the shape vectors should be aligned to remove variation which could be associated with these transformations, thus minimising non-linearities. However, the alignment can also be omitted. The model can then only generate shapes with a position, orientation and size as are observed in the training set.

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 S is

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

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

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

where $\mathbf{P}$ is the concatenation of eigenvectors of $\mathbf{S}$, i.e. $\mathbf{P} = (\phi_1|\phi_2|\ldots|\phi_{2n})$. The eigenvectors corresponding to the largest eigenvalues λ_i 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

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

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

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

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

In the original formulation of ASM, 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.

To enable coarse to fine fitting of the shape model, the profiles are sampled at 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 dissimilarity 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:

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

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

From coarse to fine resolution, the model is fitted to the image in an iterative procedure. The process initialises with the mean shape. For all landmarks, n_s 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 shape model and pose 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

This Section describes the application of ASM to AAA segmentation. The choices we made for the shape model are discussed in Section 3.1. The resulting model can be fitted to the image as was described in the previous Section. We propose two alternatives for the greylevel appearance model, which are discussed in Sections 3.2 and 3.3.

3.1. Shape modelling

We approach the problem of 3D thrombus segmentation in a slice-by-slice manner. The original CT-slices, which are perpendicular to the body axis and therefore always give approximately perpendicular cross-sectional views of the aorta, are used. One manually delineated image slice in the middle of the aneurysm is used as a reference. A plausible shape from the shape model is fitted in both adjacent slices, using the reference shape as initialisation. For the succeeding slices, the process is repeated, each time using the previous slice as the reference.

In the absence of well-defined anatomical landmark points, the landmarks are equidistantly placed along the object contours. The starting point of a contour is the point near the vertebra having the same y-coordinate as the centre of mass. All contours are aligned such that the centres of the adjacent contours coincide. In this way, translations between slices are contained in the model. No additional translation, rotation or scaling will be allowed during the fitting of the model.

Two approaches were tested: the first modelling all slices independently, the second combining two adjacent slices in one model. During fitting of the latter, the shape in the reference slice is kept constant and the deformation of the adjacent contour is constrained by the reference shape.

3.2. Slice-to-slice correlation

We propose a modified ASM scheme in which optimal landmark positions are defined by maximum greyvalue profile correlation with adjacent slices rather than by correlation with profiles from the training data.

For each landmark, a profile of k pixels on either side of the landmark point in the normal direction is sampled from the image, at multiple levels of resolution, similar to the method described in 2.2. The optimal landmark position is 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 correlation, *Ncc*:

$$Ncc = \frac{\sum_{x=-k}^{k} I_{s-1}(x) \cdot I_s(x)}{\sqrt{\sum_{x=-k}^{k} I_{s-1}(x)^2 \cdot \sum_{x=-k}^{k} I_s(x)^2}} \tag{9}$$

where I_s and I_{s-1} are the pixel intensities in the examined profile and in the reference profile respectively. The *Ncc* is maximised in the fitting procedure.

3.3. Incorporating object intensity and gradient

As can be seen from Figures 1 and 2, the intensity distribution of the thrombus is relatively homogeneous. It can contain bright calcifications, graft metal or contrast enhanced blood, but does not, apart for some metal induced artefacts, contain tissue with greyvalues below a certain intensity threshold. This can be incorporated in the fit procedure by penalising model instances that include low intensity-tissue. We do not allow new landmark positions for which more than one of a number of n_i samples taken inwards has an intensity below a threshold T. If none of its evaluated new positions is allowed, the landmark is not moved. The threshold T could be obtained from the training data, but as thrombus intensity will not vary much between patients we keep it constant at 1150, leaving about 99% of the aneurysm labelled as 'possibly thrombus'.

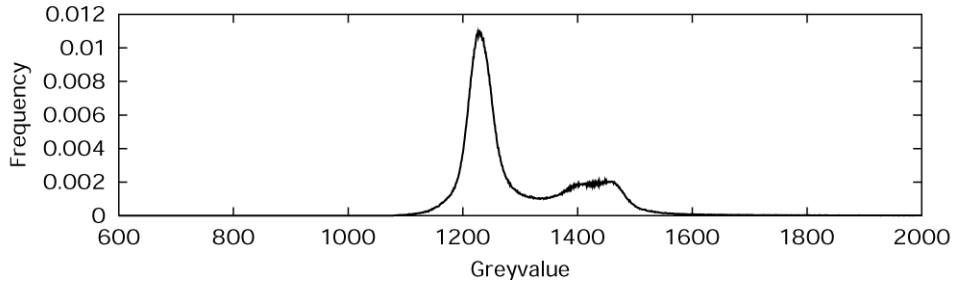

Figure 2. Intensity distribution in the aneurysm, obtained from six CTA scans. The greyvalues given are Hounsfield units, scaled between 0 and 4095.

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, although we know that the thrombus boundary 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 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 approximated using finite differences, and the function $f(x)$ 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 $0.485 \times 0.485 \times 2.0\,\text{mm}$. 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 six CTA images of six different patients, including one pre-operative and five post-operative scans. We performed a leave-one-out experiment: six shape models are built of five datasets, with which the sixth dataset, not included in the model, is segmented.

4.1. Shape model

Figure 3 shows the first six modes of the one-slice model trained with all six datasets. The amount of variance explained by the model is plotted as a function of the number of modes in Figure 4. To describe at least 99% of the total amount of variation present in the training set, 6 modes are needed in the one-slice model and 10 in the two-slice model.

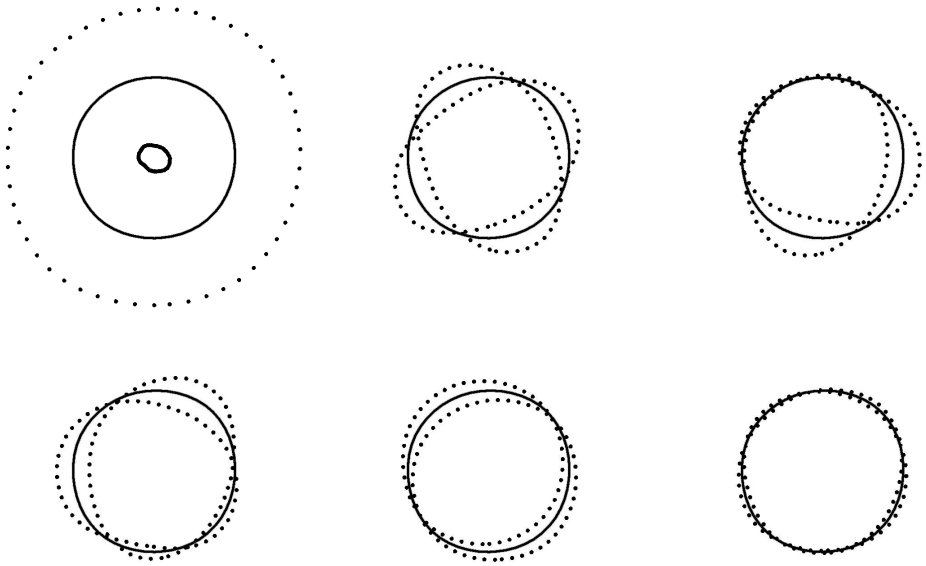

Figure 3. The effect of varying the first six modes of shape variation individually in a one-slice model built from all six datasets, using a total of 274 slices. The mean shape (solid line) and both allowed extremes $\pm 3\sqrt{\lambda_i}$ (dotted lines) are shown.

The validity of the shape model is tested by fitting it directly to the contours that were manually drawn, which gives an upper bound for the accuracy that can be obtained when the model is fitted to new image data. The error in landmark position is defined as the minimum distance from the landmark to the fitted contour. Figure 4 shows the mean error as a function of the number of modes retained. The fit accuracy is better for the one-slice model, but for both models the errors are small. For a model containing 10 modes the mean error was 0.81 voxels (0.39 mm) for the one-slice model and 0.97 (0.24 mm) voxels for the two-slice model. The averaged maximum errors per slice were 2.3 and 2.7 voxels and the overall maximum errors were 11 and 12 voxels.

For a model containing 25 modes, the error as a function of the number of datasets used for training is shown in Figure 5. These are the results for all possible combinations of training sets out of the six datasets, for fit errors on new data as well as on datasets that were contained in the model. With only two training sets, comprising about 100 slices, reasonable results can be obtained. The mean landmark positioning error in that case was 0.44 voxels for the one-slice model and 0.62 for the two-slice model. However, if more datasets were

added, the fit error for new shapes would eventually equal the error for shapes already present in the training set. This indicates that our results could still improve if more training sets were used.

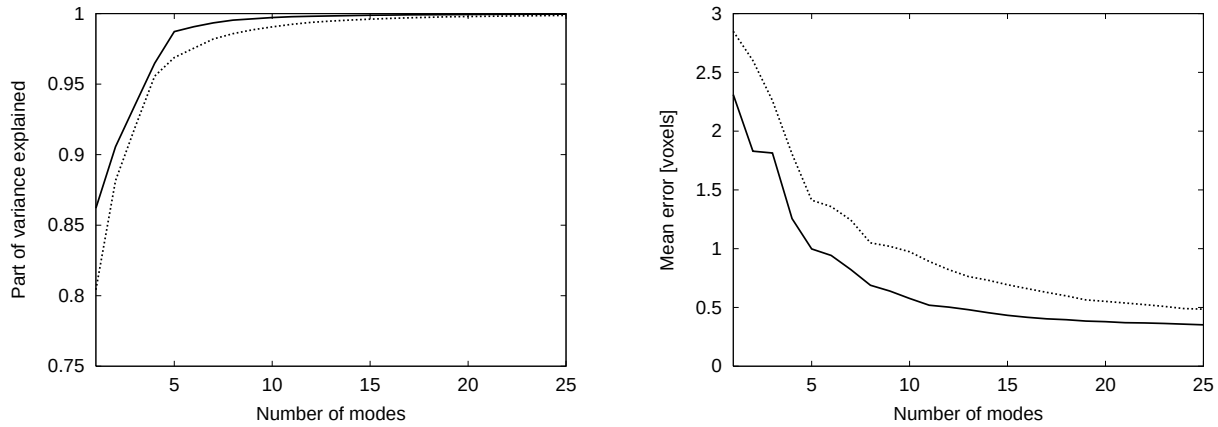

(a) (b)

Figure 4. (a) The amount of variance explained by the model as a function of the number of modes, for the one-slice model (solid line) and the two-slice model (dotted line).
(b) Mean error as a function of the number of modes contained in the model, for the one-slice model (solid line) and the two-slice model (dotted line).

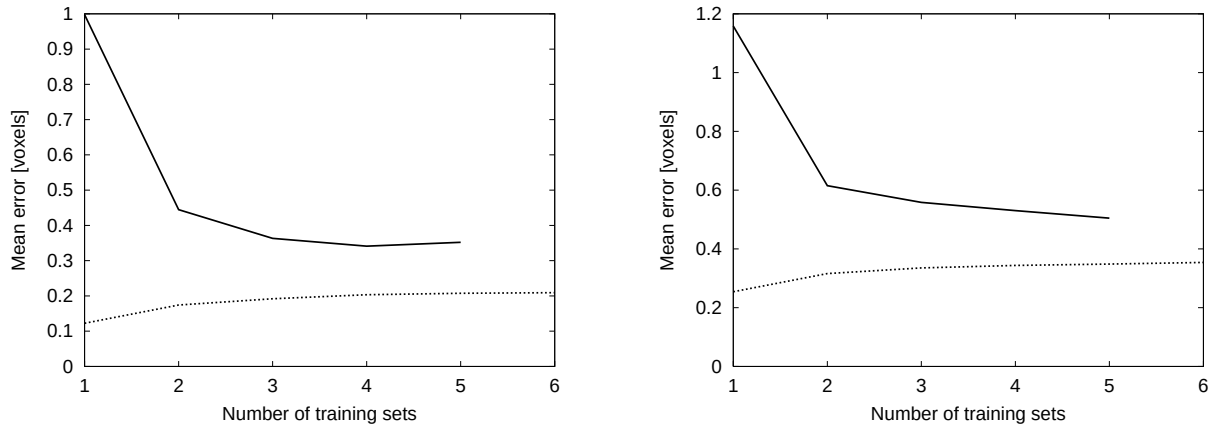

Figure 5. Mean error as a function of the number of datasets, for the one-slice model (left) and the two-slice model (right). The solid line shows the result for the leave-one-out experiment; the dotted line is the result if the training set contains the dataset to segment.

4.2. Fitting the model

We have investigated three schemes for evaluating the optimal landmark positions, as was discussed in Section 3:

1. Statistical model of local image structure (original ASM, Section 2.2)
2. Greyvalue correlation with adjacent slices (Section 3.2)
3. Greyvalue correlation with adjacent slices, combined with local intensity and gradient information (Section 3.3)

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

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

General	
s	Number of training image slices, differing per segmented volume ($\sim$ 230)
Shape model	
n	Number of landmark points (50)
f_v	Part of shape variance to be explained by the model (0.995), controlling the number of modes t (9–14)
m	Bounds on eigenvalues λ_i (3)
Appearance model	
k	Number of points in profile to sample on either side of the landmark point (3)
n_i	Number of points to sample inwards to detect different-intensity tissue, scheme 3 (6)
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)

The segmentations obtained are compared to manual segmentations. The relative volume of overlap of two volumes A and B is defined by

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

Figure 6 shows the relative overlap with the manual segmentations that were obtained for the three schemes, for all scans and slices. The horizontal 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 referred study describes the relative volume error per volume instead of per slice, without taking the overlap into account.

In datasets II and IV the original ASM scheme yields results that are comparable to those of the first scheme using correlation between slices. In all four other datasets the results obtained by the slice correlation scheme are considerably better. The number of slices segmented correctly, i.e. within the bounds of manual inter-observer reproducibility, was 86 (32%) for the conventional ASM scheme and 105 (39%) for the slice correlation scheme. Due to the strong influence of the reference slice on the slice to segment, errors are propagated through the dataset. Once the optimal contour is lost, it is not likely to be recovered in the remaining slices. The incorporation of thrombus intensity and gradient increases the probability that the error is corrected in the next slice and thus improves performance of the slice correlation scheme substantially; now 176 slices (66%) are segmented correctly. A paired t-test on the overlap values of the first 20 slices, 10 on each side of the initialisation slice, showed that both improvements are significant ($p < 0.0001$). An example of the segmentations obtained is given in Figure 7.

Figure 8 shows the results of the two-slice model together with those of the one-slice model, both fitted with the optimal scheme which combines slice correlation with intensity and gradient information. The graphs reveal that similar results are obtained in many slices, but if the methods diverge the one-slice model usually performs better ($p < 0.0001$ in a paired t-test). An example of the segmentations obtained is given in Figure 9.

5. DISCUSSION

It was shown that, using an ASM based approach, a large number of contours can be deduced from one initialisation contour drawn manually. However, performance deteriorates with distance from the reference contour. Up to 42 slices could be segmented using one reference contour, but this number could be less in the

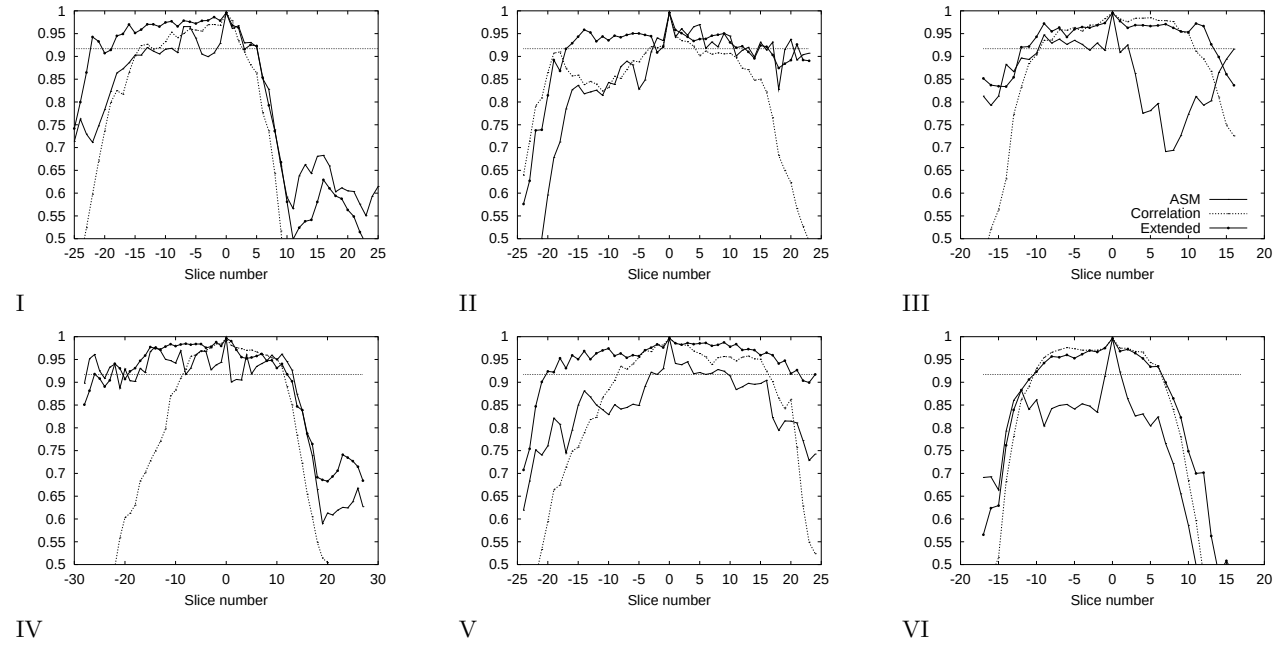

Figure 6. Results for each of the six datasets, obtained by the one-slice model. The relative volume of overlap with respect to the manual segmentation is plotted for each slice, for the segmentations obtained using the three fitting schemes. The slice in the middle of the plot, with overlap 1, is the reference slice from which sequential slices are segmented in both directions. The horizontal line indicates the approximate level of inter-observer reproducibility of the manual segmentations (91.7%).

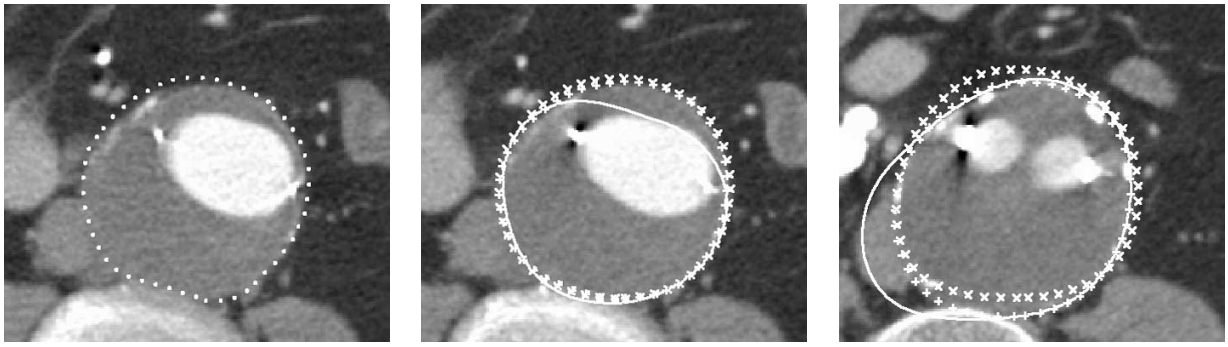

Figure 7. A typical example of segmentations obtained. The left image shows the reference slice of dataset I, with the manual initialisation. The middle image is of one slice below the reference, the right is taken ten slices lower. The original ASM scheme (solid line) has drawn to a different boundary, the slice correlation scheme (crosses) gave a good segmentation in the first slice but has slided away after several slices, and the extended slice correlation scheme (pluses) yields the right segmentation.

aneurysm top and bottom, where quickly decreasing aneurysm radius in combination with anisotropic voxels decrease contrast. The method as presented would, in an interactive setting, increase segmentation speed and reproducibility, but is not yet accurate enough for fully automated segmentation.

We found that modelling two successive slices, thus restricting the possible shape changes between slices, yields less good results than modelling all slices independently. Although fit errors are less likely to occur, if they do occur they are less likely to be corrected in an adjacent slice. However, Figure 5 reveals that the minimum

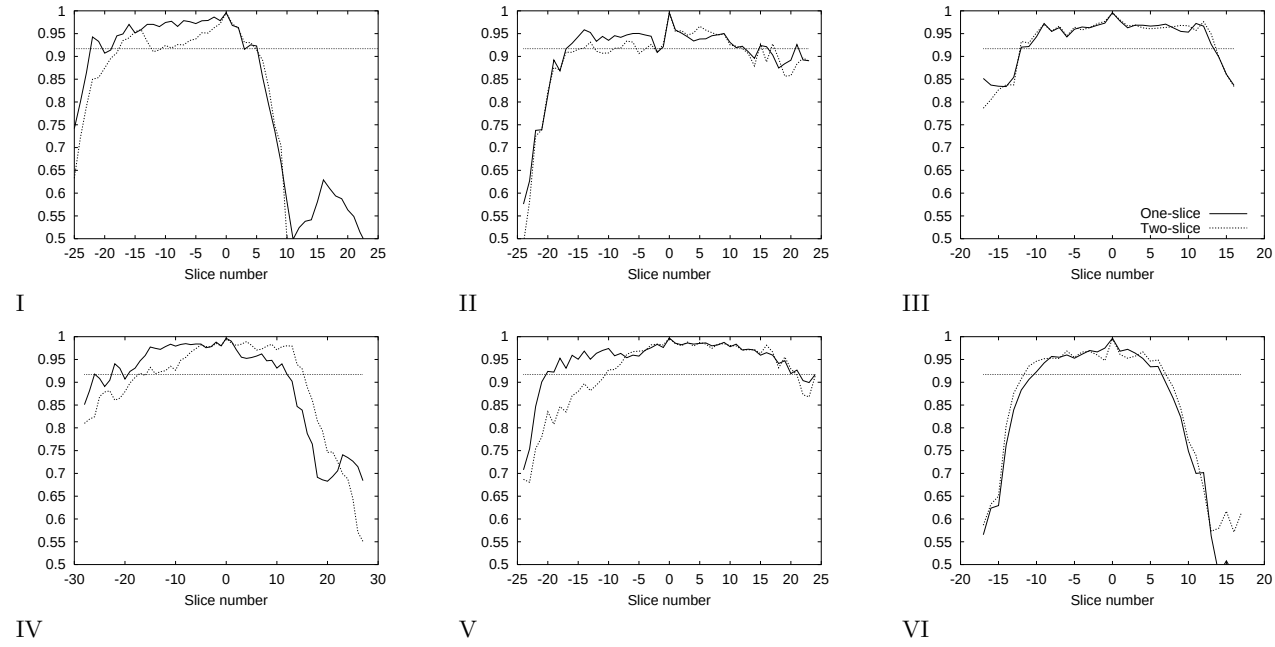

Figure 8. The one-slice and two-slice models compared. The relative volume of overlap with respect to the manual segmentation is plotted for each slice, for the segmentations obtained using both models and the extended slice correlation scheme.

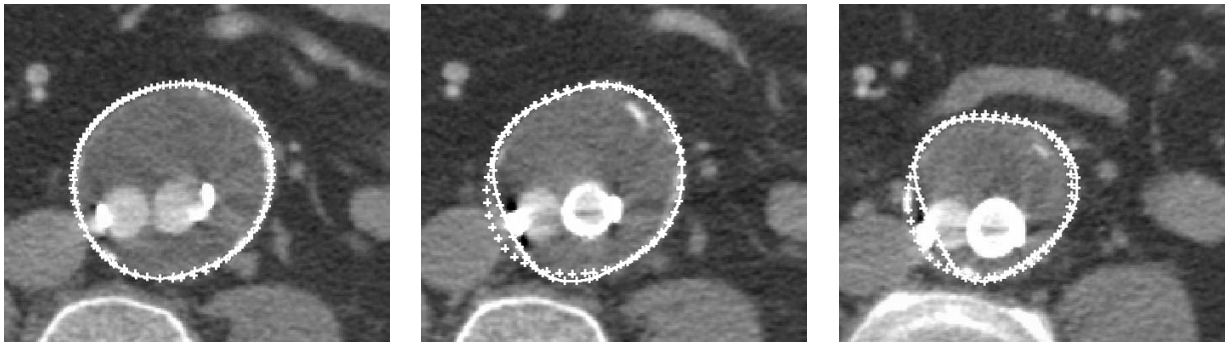

Figure 9. An example where the one-slice model performs better. These images are taken from dataset IV, from left to right five, seven and fourteen slices below the initialisation slice. The contour obtained by the one-slice model is depicted with pluses, the two-slice contour with pluses and a solid line. In one slice a slightly flattened shape is fitted (left). In successive slices (middle and right) this shape is extrapolated by the two-slice model, while the one-slice model recovers the true aneurysm shape.

error for the two-slice model fit still decreases with new data added to the training set. A larger training set would make the model more general and may solve the problems we observed. The 3D nature of the method could be further improved by adding more slices to the model, or by e.g. adjusting the allowed size of the model with distance from the initialisation slice. Then, certainly more training data will be needed.

The used similarity measure, *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.

Moreover, original ASMs on the one hand and the schemes employing slice correlation on the other hand are extreme cases, where in the first no information of neighbouring slices is used, while in both modified schemes no model information is used. As a consequence, segmentation errors are propagated through the dataset in the slice correlation schemes. A hybrid scheme combining both model information and information of neighbouring slices could improve overall performance. It was shown that the original ASM scheme, using local profile models, does not perform well on these images. This may be owing to the underlying assumption of a normal profile distribution. In the presence of distinct background structures of varying shape and brightness this assumption does not hold. The use of non-linear models[22] or a non-linear classifier[21] could be more suitable.

The dependency on image slicing and orientation restricts the use of this method to segmentation tasks where the shape change between slices is not too large and the object is always imaged in approximately the same direction. However, many medical images are made using fixed scan protocols which satisfy these conditions.

6. CONCLUSIONS

A new approach to the semi-automatic segmentation of thrombus in abdominal aortic aneurysms, based on active shape model fitting in sequential slices, is reported. We investigated the use of a one-slice model, modelling all slices independently, as well as a two-slice model, where the shape change between slices is restricted. 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. The proposed scheme was extended with a penalty function for inclusion of low-intensity tissue and a refinement step to locally adjust the position of the landmarks to points with maximum gradient.

In a leave-one-out experiment with six datasets, comprising a total number of 274 aneurysm slices, fitting on slice correlation alone outperformed the conventional ASM, and the proposed extension to the slice correlation scheme still improved the results significantly. The results obtained by the one-slice model were better than those of the two-slice model (all with $p < 0.0001$ in a paired t-test).

Using the extended slice correlation scheme and the one-slice model on average 29 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. For the basic slice correlation scheme this number was 18, and 14 for the original ASM.

These results show that shape model based segmentation of AAA will be valuable in speeding up the current manual segmentation process.

ACKNOWLEDGMENTS

This research was funded by the Netherlands Organisation for Scientific Research (NWO).
We would like to thank Drs. M. Prinssen, and Dr. J.D. Blankensteijn from the Department of Vascular Surgery for providing the datasets and expert segmentations.

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* **61**(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* **14**(2), pp. 84–90, 1997.

3. J. J. Wever, J. D. Blankensteijn, W. P. T. M. 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* **20**(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* **12**(2), pp. 182–188, 1996.
5. J. J. Wever, J. D. Blankensteijn, J. C. van Rijn, I. A. M. J. 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* **175**(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 *Medical Imaging: Image Processing*, *Proceedings SPIE* **3034**, pp. 86–94, 1997.
7. A. J. Bulpitt and E. Berry, "Spiral CT of abdominal aneurysms: comparison of segmentation with an automatic 3D deformable model and interactive segmentation," in *Medical Imaging: Image Processing*, *Proceedings SPIE* **3338**, pp. 938–46, 1998.
8. G. Rubin, D. Paik, P. Johnston, and S. Napel, "Measurement of the aorta and its branches with helical CT," *Radiology* **206(3)**, pp. 823–829, Mar 1998.
9. 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* **19**(4), pp. 337–346, 2000.
10. 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 *Medical Imaging: Image Processing*, *Proceedings SPIE* **4322**, pp. 395–406, 2001.
11. M. Subasic, S. Loncaric, and E. Sorantin, "3D image analysis of abdominal aortic aneurysm," in *Medical Imaging: Image Processing*, *Proceedings SPIE* **4322**, pp. 388–394, 2001.
12. T. F. Cootes and C. J. Taylor, "Statistical models of appearance for computer vision," tech. rep., University of Manchester, 2000.
13. T. F. Cootes and C. J. Taylor, "Statistical models of appearance for medical image analysis and computer vision," in *Medical Imaging: Image Processing*, *Proceedings SPIE* **4322**, pp. 236–248, 2001.
14. A. Hill, A. Thornham, and C. J. Taylor, "Model-based interpretation of 3D medical images," in *British Machine Vision Conference*, pp. 339–348, 1993.
15. 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* **12**(6), pp. 355–366, 1994.
16. S. Solloway, C. J. Taylor, C. E. 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.
17. N. Duta and M. Sonka, "Segmentation and interpretation of MR brain images: An improved active shape model," *IEEE Transactions on Medical Imaging* **17**(6), pp. 1049–1067, 1998.
18. 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.
19. G. Hamarneh and T. Gustavsson, "Combining snakes and active shape models for segmenting the human left ventricle in echocardiographic images," *IEEE Computers in Cardiology* **27**, pp. 115–118, 2000.
20. 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* **20**, pp. 415–423, May 2001.
21. B. van Ginneken, A. F. Frangi, J. J. Staal, B. M. ter Haar Romeny, and M. A. Viergever, "A non-linear gray-level appearance model improves active shape model segmentation," in *Mathematical Methods in Biomedical Image Analysis*, 2001.
22. T. F. Cootes and C. J. Taylor, "A mixture model for representing shape variation," *Image and Vision Computing* **17**(8), pp. 567–573, 1999.