

Model-based Segmentation of Abdominal Aortic Aneurysms in CTA Images

Marleen de Bruijne, Bram van Ginneken, Wiro J. Niessen, Marco Loog, Max A. Viergever
Image Sciences Institute, University Medical Center Utrecht, Utrecht, The Netherlands

ABSTRACT

Segmentation of thrombus in abdominal aortic aneurysms is complicated by regions of low boundary contrast and by the presence of many neighboring structures in close proximity to the aneurysm wall. This paper presents an automated method that is similar to the well known Active Shape Models (ASM), which combine a three-dimensional shape model with a one-dimensional boundary appearance model.

Our contribution is twofold: First, we show how the generalizability of a shape model of curvilinear objects can be improved by modeling the objects axis deformation independent of its cross-sectional deformation. Second, a non-parametric appearance modeling scheme that effectively deals with a highly varying background is presented. In contrast with the conventional ASM approach, the new appearance model trains on both true and false examples of boundary profiles. The probability that a given image profile belongs to the boundary is obtained using k nearest neighbor (k NN) probability density estimation. The performance of this scheme is compared to that of original ASMs, which minimize the Mahalanobis distance to the average true profile in the training set.

A set of leave-one-out experiments is performed on 23 datasets. Modeling the axis and cross-section separately reduces the shape reconstruction error in all cases. The average reconstruction error was reduced from 2.2 to 1.6 mm. Segmentation using the k NN appearance model significantly outperforms the original ASM scheme; average volume errors are 5.9% and 46% respectively.

Keywords: active shape model, classification, image segmentation, 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 at present 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.¹ Patients are scanned every three to twelve months, depending on the state of the aneurysm.

Changes in aneurysm volume form an important indicator for the risk of rupture.²⁻⁴ The current gold standard for volume assessment is Computed Tomography Angiography (CTA) with subsequent manual aneurysm delineation in each slice.⁵ 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.⁶ reported averaged inter-operator volume errors of 8.3% and intra-operator errors, for two operators, of 3.2% and 5.8%.

To reduce analysis time and to increase reproducibility, automated segmentation would be of great value. However, thrombus segmentation in CTA images is a difficult task. The boundary can be obscured by surrounding tissue of similar gray value, many neighboring structures induce strong edges in close proximity to the aneurysm wall, and the variable aneurysm radius in combination with partial volume effects makes the boundary hard to distinguish even in the absence of neighboring structures. In addition, the size and appearance of aneurysms and neighboring structures 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.

The authors can be contacted by e-mail as {marleen,bram,wiro,marco,max}@isi.uu.nl

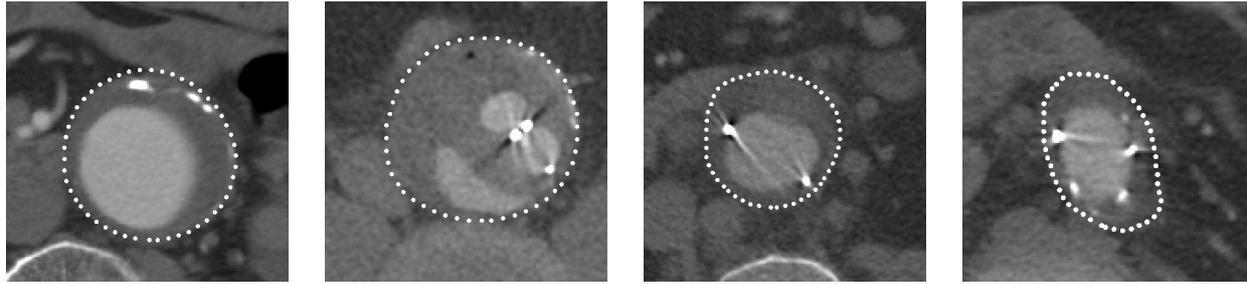


Figure 1. Four CTA slices of different patients, with manually delineated aneurysms. The left image is taken before surgery, the others are post-operative. The interior of the contour contains thrombus, contrast enhanced blood in the lumen and possibly endoleaks, calcification, and — post-operatively — metal graft markers and attachment systems. Neighboring structures are the spinal column, at the bottom of the images, blood vessels and intestines. All CTA images in this paper have been contrast stretched for better visibility of the thrombus.

Most publications on computerized AAA segmentation have concentrated on segmentation of the contrast-filled lumen, either in pre-operative^{7–15} or in post-operative CTA scans.¹⁶ The more difficult problem of thrombus segmentation has less frequently been addressed. Wilson et al.¹⁷ claimed successful thrombus segmentation using an active surface with the central lumen axis as a seed, but did not present an evaluation study. Subasic et al.¹⁸ segmented the aortic lumen using a three-dimensional level-set approach, but found that the method could not cope with the large regions lacking boundary evidence in thrombus segmentation. They solved this problem using a combination of thresholding, edge detection, morphological operations, and shape restrictions in the form of a distance map, favoring convex cross-sections. A level-set was evolved in the resulting binary image to ensure a smooth object. The average segmentation error per slice was 19.8%, which makes the method suitable for visualization purposes but not for accurate volume measurements.

Since one of the main problems in AAA segmentation is the incompleteness of boundary evidence, we believe that the use of shape priors can achieve considerable improvement with respect to the approaches mentioned above. Several authors have combined deformable models with statistical analysis of the shape variables in a training set.^{19–21} Our approach is based on the Active Shape Model (ASM) segmentation scheme as put forward by Cootes and Taylor,^{21,22} which combines a landmark based statistical shape model with local appearance models near object contours. ASMs have been successfully applied to various segmentation tasks in medical imaging.^{23–30}

Conventional ASMs use a linear model generated from gray value profiles in training images to fit the shape model to the image. Since AAAs have a diverse range of possible surrounding structures, which vary in location, shape, intensity and texture, a linear model of gray level structure about the contours is not always able to find the true contour in a new image. We have previously reported an interactive approach to AAA segmentation in which a shape model built of two adjacent image slices is fitted progressively to successive slices, using gray value correlation between adjacent slices as a fit measure.³¹ This approach is easily integrated with the current manual slice-by-slice segmentation. However, the use of slice similarity features and a two-slice shape model results in propagation of errors and performance deteriorates with distance from the reference contour.

In this work we present a three-dimensional approach using a statistical model of cylindrical shape variations, restricting the deformation to in-slice landmark displacements. We discuss how an overconstrained model of a curvilinear object can be generalized using the assumption that cross-section and axis deformations are uncorrelated, and we developed a novel non-parametric appearance modeling scheme that effectively deals with a highly varying background. In contrast with conventional ASMs, the proposed appearance modeling scheme extracts both boundary and non-boundary profiles from the training set. The quality of fit of a given landmark to the local image structure, the probability that a given profile corresponds to the aneurysm boundary, is determined using k nearest neighbors (k NN) probability density estimation.

The shape model is briefly discussed in Section 2.1. Section 2.2 describes the appearance model of ASM and our augmented model, and Section 2.3 explains how the shape model is fitted to the image using these appearance models. Details on how we apply the method to AAA segmentation are given in Section 2.4. Section 3 presents the results of a series of leave-one-out experiments on 23 datasets. Discussion and conclusions are given in Sections 4 and 5.

Table 1. Parameters of the segmentation scheme.

Shape model (Section 2.1)		
n	1500	Number of landmark points
f_v	0.99	Part of shape variance to be explained by the model, controlling the number of modes t (Eq. 6)
f_c	3	Bound on eigenvalues λ_i (Eq. 9)
Gray value model (Section 2.2)		
k	7	Patch length
n_{shift}	2	Number of shifted examples on both sides of the contour (k NN model)
d_{shift}	2	Amount of shift (k NN model)
k_{nn}	80	Number of neighbors to evaluate in k NN search (k NN model)
Fitting algorithm (Section 2.3)		
n_s	5	Number of new landmark positions to examine per iteration on either side of the current position
L	4	Number of resolution levels
N	5	Number of iterations per resolution level

2. METHOD

This section describes the segmentation scheme. Shape variations in the training set are described using a Point Distribution Model (PDM).²² 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 gray value structure.

2.1. Point Distribution Models

In PDMs, a statistical model of object shape and shape variation is derived from a set of s training examples. Each training example is described by n manually or automatically identified landmark points that correspond between shapes. Variations in the coordinates of these landmark points describe the variation in shape and pose across the training set.

A shape is described by its shape vector \mathbf{x} containing all landmark coordinates, in three dimensions:

$$\mathbf{x} = (x_1, y_1, z_1, x_2, y_2, z_2, \dots, x_n, y_n, z_n). \quad (1)$$

To maximize the specificity of the model and to reduce non-linearities in the shape distribution, the shapes are aligned by rotation, translation, and scaling using Procrustes Analysis,^{32, 33} and transformed into the tangent space to the mean shape.²²

Principal Component Analysis (PCA) is applied to the aligned shape vectors. Hereto, the mean shape $\bar{\mathbf{x}}$, the covariance matrix \mathbf{S} , and the eigensystem of \mathbf{S} are computed:

$$\bar{\mathbf{x}} = \frac{1}{s} \sum_{i=1}^s \mathbf{x}_i \quad (2)$$

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

The eigenvectors ϕ_i of \mathbf{S} provide the modes of shape variation present in the data. The eigenvectors corresponding to the largest eigenvalues λ_i account for the largest shape variation; a small number of modes usually explains most of the variation. Each shape \mathbf{x} in the set can thus be approximated by

$$\mathbf{x} \approx \bar{\mathbf{x}} + \Phi \mathbf{b} \quad (4)$$

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

$$\mathbf{b} = \Phi^T (\mathbf{x} - \bar{\mathbf{x}}). \quad (5)$$

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:

$$\sum_{i=1}^t \lambda_i \geq f_v \sum_i \lambda_i. \quad (6)$$

A common problem in statistical shape modeling is that the model can be too specific to fit to new shapes properly, owing to a limited amount of training data. This is often the case with three-dimensional models, where a large number of landmarks is needed to describe a shape properly. We propose to generalize models of curvilinear structures like AAA using the assumption that the cross-sectional shape is independent of the object axis. Both the cross-sections and the axes are modeled separately and then combined into one model describing both deformations. Details are given in Section 2.4.

2.2. Gray value model

Fitting the shape model to a new image requires a measure of probability that a point belongs to the boundary. The objects boundary appearance can be learned from the training set. To this end, gray value profiles of k pixels are sampled around each landmark, perpendicular to the object contour.

In the original ASM formulation, a linear model is built from these profiles, similar to the shape model. The effect of global intensity changes is reduced by sampling the first derivative and normalizing the profile. The normalized samples are assumed to be distributed as a multivariate Gaussian, and the mean $\bar{\mathbf{g}}$ and covariance matrix \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 squared Mahalanobis distance $f(\mathbf{g}_s)$ from the sample to the model mean:

$$f(\mathbf{g}_s) = (\mathbf{g}_s - \bar{\mathbf{g}})^T \mathbf{S}_g^{-1} (\mathbf{g}_s - \bar{\mathbf{g}}). \quad (7)$$

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

We previously showed that the Mahalanobis distance to the average training profile does not perform well in AAA boundary localization for slice-by-slice segmentation.³¹ One of the shortcomings of this gray value model is that only the appearance of the correct boundary is learned from the training set. Boundaries of neighboring structures may have similar appearance, but that has no effect on the obtained quality of fit. Furthermore, the underlying assumption of a normal profile distribution does not hold in the presence of distinct background structures of varying shape and brightness.

We propose to treat the position evaluation step in the ASM fitting as a classification problem. Again, gray value profiles are sampled from the training set, but now a classifier is trained on both true and false examples of aneurysm contours. For each landmark, one true example is sampled around the landmark and perpendicular to the contour, and $2n_{\text{shift}}$ false examples are sampled in the same direction, n_{shift} displaced outwards and n_{shift} displaced inwards. The first false example in one direction is sampled at a distance d_{shift} from the landmark, and subsequent examples are obtained by each time shifting d_{shift} with respect to the previous sample.

In a new image, the probability that a given profile lies on the aneurysm boundary is given by the posterior probability from the classifier for that profile. In this work, a k NN classifier is used and the posterior probability is given by

$$P(\text{boundary} | g_s) = \frac{n_{\text{true}}}{k_{nn}}, \quad (8)$$

where n_{true} is the number of boundary samples among the k_{nn} nearest neighbors.

For the original ASM gray value model as well as for the new k NN gray value model the profiles are sampled at multiple resolutions, to enable coarse-to-fine fitting. The finest resolution uses the original image and a sample spacing of one voxel, the next resolution is the image convolved with a Gaussian kernel of width one and sampled with a spacing of two voxels, and subsequent resolutions are obtained by doubling both the kernel width and the sample spacing. When the sample spacing is doubled, the amount of shift d_{shift} of the k NN model is doubled as well.

2.3. Fitting procedure

In ASMs, the shape model is fitted to new images using a fast deterministic optimization scheme. The process initializes with some plausible shape, usually the mean. For all landmarks, n_s possible new positions are evaluated along the line perpendicular to the contour, on both sides. The distance between the $2n_s + 1$ positions to evaluate depends on the level of resolution and is equal to the sample spacing of the corresponding gray level model: at the finest resolution the distance is 1 voxel, at each subsequent resolution it is doubled. The optimal position is determined by one of the gray value models that were discussed in Section 2.2. The current shape is translated, rotated and scaled such that the squared distance between the landmarks and the optimal positions are minimized. Next, the optimal shape parameter vector \mathbf{b} is computed using Equation 5. Hard limits are applied to constrain the elements b_i of \mathbf{b} to plausible values:

$$|b_i| \leq f_c \sqrt{\lambda_i}. \quad (9)$$

This process is repeated a fixed number of N times, whereupon it is repeated at the next level of resolution.

This fitting procedure can encounter problems when segmenting objects with noisy or highly variable edge evidence, since the objective function to minimize, the sum of squared distances, is sensitive to outliers. Modifications to the fitting algorithm have been proposed to make it more outlier resistant. Rogers and Graham³⁴ have compared various robust estimation techniques in the context of ASM search, including least median of squares, random sample consensus, and weighted least squares fitting, where the weights can be determined either by image information or by the shape residual distribution. Duta and Sonka²⁶ suggested to detect outliers as points that induce an exceptionally large part of the total amount of variation, and move them to the mean position. Hamarneh and Gustavsson,²⁷ Behiels et al.,³⁰ and Mitchel et al.²⁸ proposed to use dynamic programming to favor landmark displacements that are smoothly varying between neighbors.

In our experiments we have applied dynamic programming regularization followed by a weighted least squares fit. Details are given in Section 2.4.

2.4. Application to AAA segmentation

Landmarking strategy

The original CT-slices are perpendicular to the body axis and therefore give approximately perpendicular cross-sectional views of the aorta. In our datasets, an aneurysm extends over 34 to 65 slices. We model 3D cylindrical shape variations, restricting the deformation to in-slice landmark displacements. To build the shape model, a fixed number of slices n_z is interpolated between the beginning and end slice. An equal number of landmarks n_{xy} is placed in each slice, along contours that were drawn manually by an expert. In the absence of well-defined anatomical landmark points, landmarks are equidistantly placed along the object contours. The starting point of a contour is the posterior point with the same y -coordinate as the center of mass. Before the model is fitted to a new image, the user indicates the first and last aneurysm slice, thus removing the need for scaling in the z direction. As a consequence, the shape vector contains only x and y coordinates.

Handling small training sets

The generalizability of the model is increased by first modeling the axes and cross-sections separately and then combining the two models, thus assuming that both types of shape variation are uncorrelated. To this end, s central axes and s straightened shapes are extracted from the s aligned training shapes. Each axis contains one landmark per slice, defined by the centroid of the contour in that slice. The straightened shapes have the same number of landmarks as the original shapes and are formed by translating each contour such that its centroid is in the origin. PDMs are derived for both shape distributions as described in Section 2.1.

The axis PDM now describes the shape variation with respect to the mean axis, while the straightened PDM describes the variation with respect to the mean cross-sectional shape. To be able to combine the two models they need to be of equal dimensions ($2n$). The $2n$ -dimensional modes of axis variation are obtained by translating the slices of the mean cross-sectional shape such that their origin coincides with the landmarks of the deformed axes, and the total mean shape is determined similarly by translating the mean cross-sectional shape to the landmarks of the mean axis. The two models will in general not be linearly independent, therefore a second PCA is performed to remove any correlation between the axis

and cross-sectional modes. The combined model is thus obtained by computing the principal components of the column vectors of

$$(\Phi_{\text{cross}} \mathbf{W}_{\text{cross}} | \Phi_{\text{axis}} \mathbf{W}_{\text{axis}}) \quad (10)$$

where the Φ are concatenations of eigenvectors and \mathbf{W} are diagonal weight matrices of the corresponding eigenvalues.

The resulting model contains at maximum $2(s - 1)$ modes, provided that $s - 1 < n$, while a model built from all shapes directly would contain only $s - 1$ modes.

Initialization

The complexity of the images requires an accurate initialization. We have asked the user to draw the top and bottom contours of the aneurysm manually. To aid the model in establishing the correct object axis an additional point is placed in the approximate aneurysm center of the central slice. The shape model is iteratively fitted to these points using a constrained ASM scheme, in which after each iteration the landmarks of the manually drawn slices are replaced to their original position and the landmarks of the central slice are translated such that their average position coincides with the manually identified center point. Alternatively, an automatic estimate of the luminal or aneurysmal axis or a — more easily automated — lumen segmentation could be used for initialization.

Subsequently, a fixed number of slices is interpolated from the image, and the shape model is fitted at multiple resolutions to optimally match the local image structure, given the two manually drawn contours. The segmentation process is constrained by keeping the two manually drawn slices fixed.

Regularization

We have applied dynamic programming regularization^{27,28,30} to reduce the effect of outliers within a slice, but the landmark movements from slice to slice are unconstrained, as large shape differences can occur between slices. The effect of outliers is further reduced by applying a weighted least squares fit³⁴ on the basis of image information. Equation 5 then becomes

$$\mathbf{b} = \mathbf{K}^T (\mathbf{x} - \bar{\mathbf{x}}), \quad (11)$$

$$\mathbf{K} = (\Phi^T \mathbf{W}^T \mathbf{W} \Phi)^{-1} \Phi^T \mathbf{W}^T \mathbf{W}$$

where \mathbf{W} is a diagonal weight matrix. Weights w_i describe the certainty that the i th landmark is on its correct position, and are chosen as the probability that a given profile is on the boundary. For the linear model that is

$$P(\text{boundary} | g_s) = c e^{\frac{-f(g_s)}{2}}, \quad (12)$$

with c a normalization constant that can be ignored in this case. For the k NN model the weights are given by the posterior probability as defined in Equation 8.

3. EXPERIMENTS AND RESULTS

A series of leave-one-out experiments is performed on 23 routinely acquired CTA images from 23 different patients, including 3 pre-operative and 8 post-operative scans. The remaining 12 scans are taken at follow-up ranging from 1 to 12 months. The scan resolution is $0.488 \times 0.488 \times 2.0$ mm. Each image consists of circa 125 slices of 512×512 voxels, of which 34 to 63 slices contain aneurysmal tissue.

Unless mentioned otherwise, all parameters are kept fixed at the values given in Table 1.

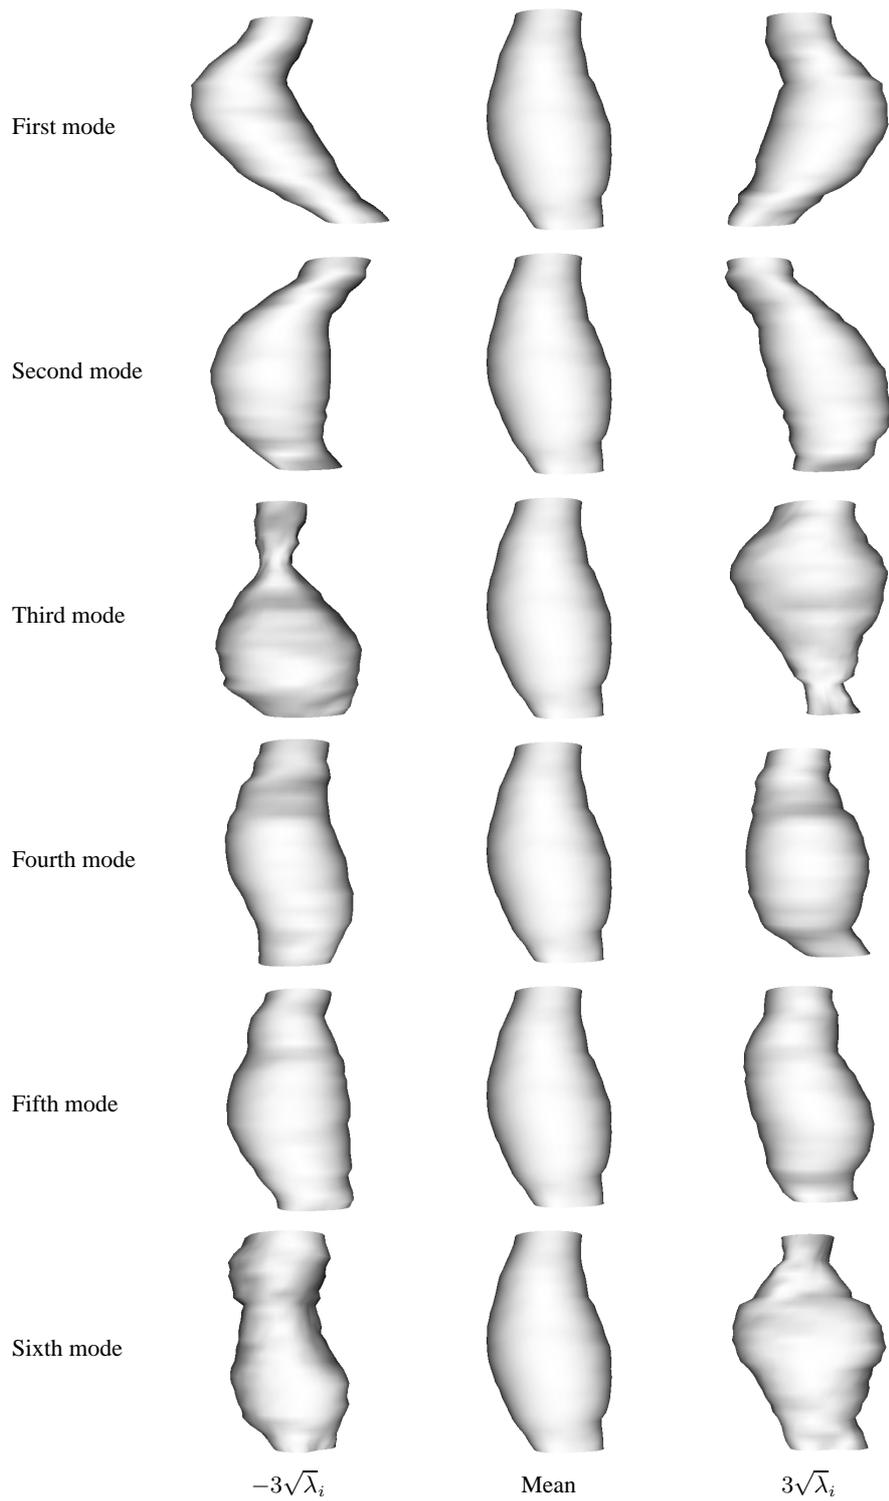


Figure 2. The effect of varying each of the first six shape parameters individually (frontal view). The largest amount of variance across the training set is caused by a curving of the axis, shown in the first two modes. The third mode describes aneurysms with a long or a short proximal neck, while mode four and five reveal a buckling of the axis. Mode six shows some variations in cross-sectional shape along the axis and distinguishes between aneurysms that have the greater part of the volume in the center or spread along the axis.

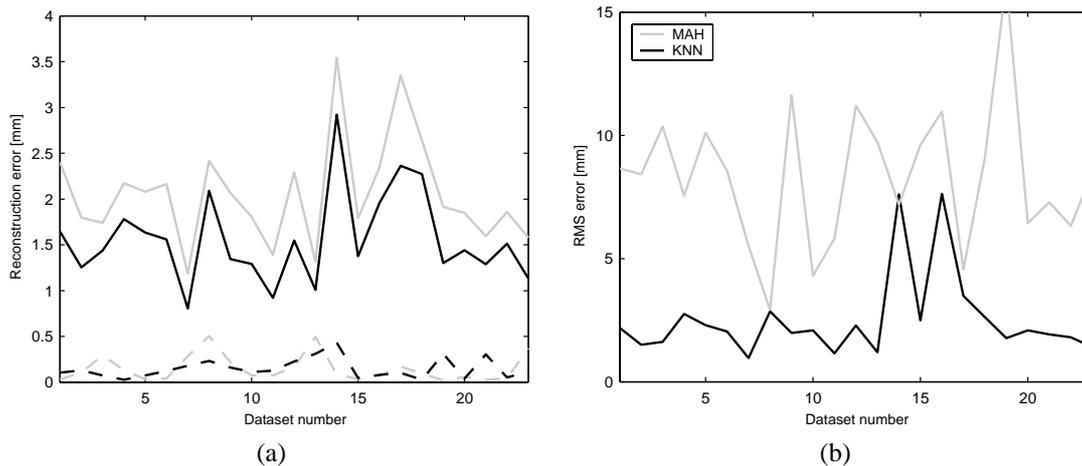


Figure 3. (a) Root mean squared reconstruction error for all 23 datasets. The continuous lines show the leave-one-out error for a normal three-dimensional shape model (gray) and for the combined model of axis and cross-section (black). The dashed lines denote the corresponding error if the test image is added to the training set. Both normal models use 21 modes of shape variation, and the combined models use 21 axis modes and 21 cross-section modes. Shapes are constrained using $f_c = 3$. (b) Root mean squared segmentation error for all 23 datasets, for the Mahalanobis distance based model (gray) and the k NN model (black).

3.1. Shape model evaluation

Figure 2 visualizes the first six modes of shape variation of the combined axis and cross-section model trained with all 23 datasets. The validity of the shape model is tested by fitting the model directly to the manual segmentations, which gives an upper bound for the accuracy that can be obtained when the model is fitted to new image data. Figure 3a shows the root mean squared landmark-to-contour error for all datasets. Modeling the axis and cross-section separately reduced the reconstruction leave-one-out error in all cases; the average error was reduced from 2.2 to 1.6 mm. There is however still a large discrepancy between the leave-one-out errors and the leave-all-in errors, suggesting that a larger training set is needed to span the entire space of possible shapes.

3.2. Gray value model

Starting from the landmarks on the manually drawn contour, the optimal landmark positions according to the two gray value models are determined, for varying n_s . The distances from these points to the landmark positions provided by the manual tracings are measured. The resulting root mean squared distances as a function of the size of the search region are shown in Figure 4, together with the expected error for random landmark selection. For both gray value models, the low-resolution model is more robust in large search regions, yet for accurate boundary localization if the true boundary is within a few millimeters distance, the high resolution models are better suited. The difference in performance for low and high resolution models is more pronounced for the k NN model. The k NN model performs significantly better than the conventional ASM gray value model at all resolutions, but still the errors are quite large, suggesting that an accurate initialization may be needed.

3.3. Segmentation

Figure 3b shows the segmentation results obtained using the two gray value models. The k NN model yields significantly better results than the original ASM model ($p < 0.00001$ in a paired T-test). Average root mean squared errors are 2.5 and 8.3 mm (5.1 and 17 voxels). The relative volumes of overlap are 92% and 62%, and the average volume errors are 5.9% and 46%. There are two datasets in which the error obtained using the k NN model is larger than half a centimeter. An example slice of both images is given in Figure 5. If these two problematic datasets are left out of consideration, the average error of the remaining 21 datasets is 2.0 mm. The corresponding volume of overlap is 94% and the relative volume error 4.2%. Note that the latter is the volume error as a percentage of the total volume enclosed by the aneurysm boundary, while thrombus volume errors reported in literature⁶ are expressed relative to the pure thrombus volume, excluding the passable

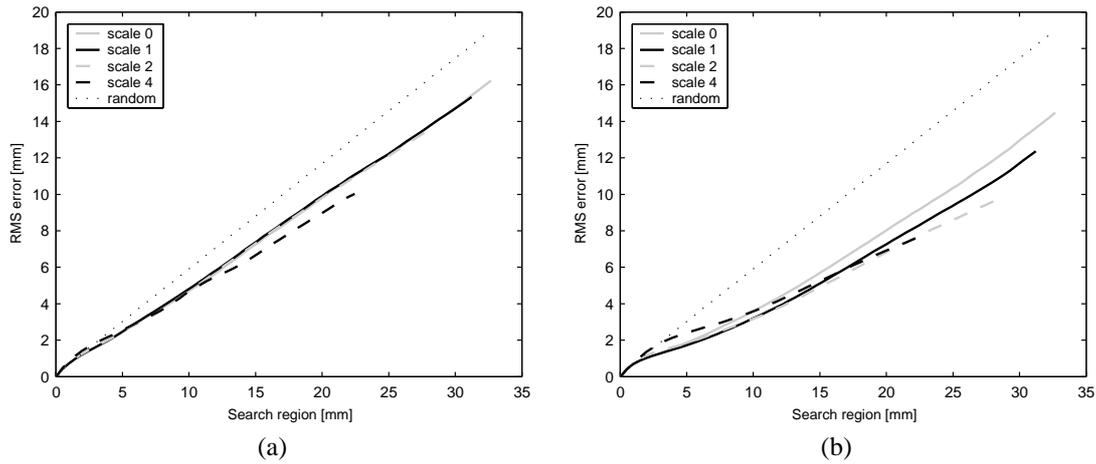


Figure 4. Root mean squared error of landmark positioning as a function of the length of the search region on either side of the contour, for (a) the Mahalanobis distance, and (b) the k NN gray value model.

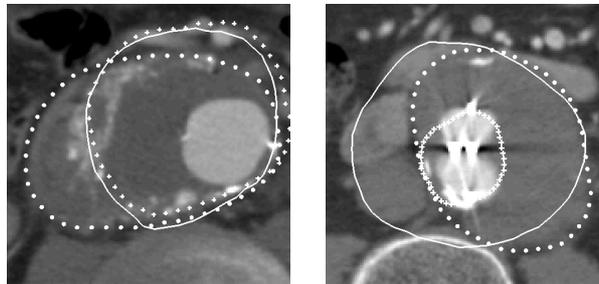


Figure 5. The two datasets in which the k NN method was unable to find an adequate segmentation. The dataset on the left combines an extremely wide aneurysm with calcifications, which are usually found only at the boundary, inside the aneurysm. The reconstruction error for this dataset was also large (dataset number 14 in Figure 3a). In the second dataset the aneurysm is embedded in other structures with similar gray value for over 10 adjacent slices, while the total region comprised by the aneurysm and its surrounding structures would be a plausible aneurysm shape. The first problem may be solved if more training shapes are available, the second case is clearly one that this method, using only shape constraints and the gray values near the contour, cannot handle. This image requires very accurate initialization or user interaction.

lumen. We estimate the relative thrombus volume errors of the obtained segmentations by subtracting the lumen volume, obtained through thresholded 3D volume growing under expert supervision, and allowing for a 3.5% volume error in this lumen segmentation.¹⁶ Under these assumptions, the thrombus volume error is 5.8% on average for the 21 successful segmentations using the k NN model, which is comparable to the intra-operator errors reported by Wever et al.⁶

Examples of segmented slices, randomly chosen from the 21 datasets that were successfully segmented, are shown in Figure 6.

4. DISCUSSION

We have restricted the model deformation to in-slice landmark displacements and consider only in-slice gray value profiles, while a full three-dimensional approach — allowing gray value modeling perpendicular to the object's surface — could be more appropriate in some applications. In the case of CTA images, which are in general highly anisotropic (in our images the voxels are over 4 times larger in the z -direction), we do not expect a significant improvement in boundary localization if the profiles would be sampled in three dimensions.

The combination of axis and cross-section models effectively increases the ability of the model to generalize to new images. However, reconstruction errors as shown in Figure 3a are still quite large. Our current work focuses on extending the model

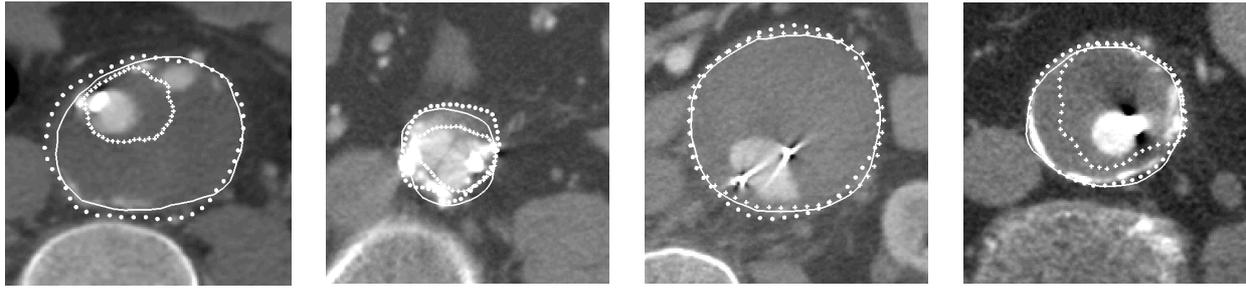


Figure 6. Image slices taken randomly from the 21 successful segmentations, with the manually drawn contour (dots), the segmentation obtained using original ASM (pluses) and the segmentation obtained with the k NN gray value model (continuous line). The k NN model obtained a segmentation near the manual contour in all four cases. The original ASM gray value model tends to draw to the lumen boundary, and finds a satisfactory segmentation only in the third image.

to incorporate other smooth deformations not observed in the training set, for example by adding artificial covariance as was suggested by Wang and Staib.³⁵

The improvement of the presented gray value model over the original ASM gray value model is twofold. First, not only the appearance of the boundary but also the appearance of points near the boundary is learned from the training set. This raises the question how many and which false examples should be taken into account. In our application, the first two shifted examples (one on each side of the contour) made a large difference, while adding more examples increases boundary localization accuracy only marginally. Second, we do not assume a Gaussian intensity profile distribution but estimate the distribution non-parametrically with k NN probability density estimation. This can be important in the presence of distinct background structures of varying brightness. It is however computationally more expensive and it can be less accurate if the number of training profiles is small.

A full segmentation took on average 25 seconds on a 2GHz Pentium PC when the original ASM gray value model was used and 450 seconds using the k NN model. The main reason for the increase in computation time is the time it takes to find the nearest neighbors in a large set of training profiles. We believe that a segmentation time of 450 seconds is acceptable for automatic segmentation, but if computation time is an issue the method could be sped up by using an approximate k NN classification, by using less shifted examples and pruning the k NN tree, or by using other classifiers. For instance, a quadratic discriminant classifier could be used, which is equivalent to extending the original ASM gray value modeling scheme to more classes, such that the probability that a given profile is on the boundary is determined by the Mahalanobis distances to the means of boundary and shifted profiles.

A different approach to applying a non-linear appearance model in ASM segmentation was proposed by van Ginneken and co-workers.²⁹ They extract local texture features at points on a grid around each landmark, label the points as inside or outside the object, perform feature selection and train a k NN classifier on the labeled points. The optimal landmark position in the model fitting step is then the position that maximizes the probability that the points inside the contour indeed belong to the object and those outside are part of the background. This method was shown to perform very well on segmentation problems where the object and background have different texture. In the case of AAA segmentation, where many background structures are similar to the object in gray value and texture, the ordering of different gray values along the profile can become important and our method may be more suitable. In addition, in the top and bottom part where the aneurysm diameter is small, the extreme "inside" part of the profile may in fact be outside the object on the other side, which makes the method proposed in²⁹ less reliable.

In this work we have discussed segmentation of the thrombus outer boundary, while for thrombus volume measurements the passable lumen must be extracted as well. For the lumen segmentation one of several existing automatic or semi-automatic methods can be used.⁷⁻¹⁶ Another option is to apply the proposed method to the lumen as well, thus building a combined shape model of both the inner and the outer boundary. In that case, the more easily detected lumen-thrombus boundary could help the model in locating the thrombus-background boundary. Obviously, this increases the problem of overconstraining the model using a small training set and probably more than the 23 datasets used in this study will be needed to obtain satisfactory segmentation results.

The presented method achieves good results in most cases, but sometimes fits (partly) to a different boundary (see Figure 5). In those cases, an interactive scheme in which a user drags one or a few landmark points to the correct boundary and the model is constrained to go through those points could be helpful. Van Ginneken et al³⁶ show that fixing one or two points can drastically improve segmentation accuracy in two-dimensional ASM segmentation of AAA. Such a scheme could also be used for initialization,³⁷ so that the user has to click only a few points on the boundary of the first and last slice, instead of drawing an entire contour.

5. CONCLUSIONS

We have presented a three-dimensional model based segmentation of thrombus in CTA scans of abdominal aortic aneurysms. The generalizability of a point distribution model is extended by modeling axis and cross-section deformation modes independently. In addition, we propose a new gray value model which significantly outperforms the gray value model that is normally used in ASM ($p < 0.00001$). Obtained volume errors are comparable to intra-observer errors reported in literature.

ACKNOWLEDGMENTS

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

REFERENCES

1. I. Broeders, J. Blankensteijn, A. Gvakharia, J. May, P. Bell, J. Swedenborg, J. Collin, and B. Eikelboom, "The efficacy of trans-femoral 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.
2. J. Wever, J. Blankensteijn, W. Mali, and B. 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.
3. B. Czermak, G. Fraedrich, M. Schocke, I. Steingruber, P. Waldenberger, R. Perkmann, M. Rieger, and W. Jaschke, "Serial CT volume measurements after endovascular aortic aneurysm repair," *Journal of Endovascular Therapy* **8**(4), pp. 380–389, 2001.
4. J. Pollock, S. Travis, S. Whitaker, I. Davidson, R. Gregson, B. Hopkinson, P. Wenham, and S. MacSweeney, "Endovascular AAA repair: Classification of aneurysm sac volumetric change using spiral computed tomographic angiography," *Journal of Endovascular Therapy* **9**(2), 2002.
5. R. Balm, R. Kaatee, J. Blankensteijn, W. Mali, and B. Eikelboom, "CT-Angiography of abdominal aortic aneurysms after trans-femoral endovascular aneurysm management," *European Journal of Vascular and Endovascular Surgery* **12**(2), pp. 182–188, 1996.
6. J. Wever, J. Blankensteijn, J. van Rijn, I. Broeders, B. Eikelboom, and W. 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.
7. M. Fiebich, M. Tomiak, R. Engelmann, J. McGilland, and K. Hoffman, "Computer assisted diagnosis in CT angiography of abdominal aortic aneurysms," in *Medical Imaging: Image Processing*, K. Hanson, ed., *Proceedings of SPIE* **3034**, pp. 86–94, SPIE Press, 1997.
8. A. 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*, K. Hanson, ed., *Proceedings of SPIE* **3338**, pp. 938–946, SPIE Press, 1998.
9. G. Rubin, D. Paik, P. Johnston, and S. Napel, "Measurement of the aorta and its branches with helical CT," *Radiology* **206**, pp. 823–829, Mar 1998.
10. O. Wink, W. Niessen, and M. Viergever, "Fast delineation and visualization of vessels in 3-D angiographic images," *IEEE Transactions on Medical Imaging* **19**(4), pp. 337–346, 2000.
11. H. Tek, D. Comaniciu, and J. Williams, "Vessel detection by mean shift-based ray propagation," in *IEEE Workshop on Mathematical Methods in Biomedical Image Analysis*, L. Staib and A. Rangarajan, eds., IEEE Computer Society Press, 2001.
12. R. Pohle and K. Toennies, "A new approach for model-based adaptive region growing in medical image analysis," in *Computer Analysis of Images and Patterns*, W. Skarbek, ed., *Lecture Notes in Computer Science* **2124**, pp. 238–246, Springer, 2001.
13. J. Weese, M. Kaus, C. Lorenz, S. Lobregt, R. Truyen, and V. Pekar, "Shape constrained deformable models for 3D medical image segmentation," in *Information Processing in Medical Imaging*, M. Insana and R. Leahy, eds., *Lecture Notes in Computer Science* **2082**, pp. 380–387, Springer, 2001.
14. D. Magee, A. Bulpitt, and E. Berry, "Combining 3D deformable models and level set methods for the segmentation of abdominal aortic aneurysms," in *Proceedings of the British Machine Vision Conference*, T. Cootes and C. Taylor, eds., pp. 119–126, The British Machine Vision Association, 2001.

15. A. Giachetti, M. Tuveri, and G. Zanetti, "Reconstruction and web distribution of measurable arterial models," *Medical Image Analysis* **7**(1), pp. 79–93, 2002.
16. M. de Bruijne, W. Niessen, J. Maintz, and M. Viergever, "Localization and segmentation of aortic endografts using marker detection," *IEEE Transactions on Medical Imaging*, 2003. In Press.
17. L. Wilson, S. Brown, J. Young, R. Li, and L. Brandt, "Three-dimensional computer models of abdominal aortic aneurysms by knowledge-based segmentation," in *Computer Assisted Radiology and Surgery*, H. Lemke, M. Vannier, K. Inamura, and A. Farman, eds., *Excerpta Medica International Congress Series* **1191**, pp. 213–217, Elsevier Publishers, 1999.
18. M. Subasic, S. Loncaric, and E. Sorantin, "3D image analysis of abdominal aortic aneurysm," in *Medical Imaging: Image Processing*, M. Sonka and M. Fitzpatrick, eds., *Proceedings of SPIE* **4684**, pp. 1681–1689, SPIE Press, 2002.
19. L. Staib and J. Duncan, "Model based deformable surface finding for medical images," *IEEE Transactions on Medical Imaging* **15**(5), 1996.
20. G. Székely, A. Kelemen, C. Brechbühler, and G. Gerig, "Segmentation of 3D objects from MRI volume data using constrained elastic deformations of flexible Fourier surface models," *Medical Image Analysis* **1**(1), pp. 19–34, 1996.
21. T. Cootes, C. Taylor, D. Cooper, and J. Graham, "Active shape models – their training and application," *Computer Vision and Image Understanding* **61**(1), pp. 38–59, 1995.
22. T. Cootes and C. Taylor, "Statistical models of appearance for computer vision," tech. rep., University of Manchester, 2002.
23. A. Hill, A. Thornham, and C. Taylor, "Model-based interpretation of 3D medical images," in *Proceedings of the British Machine Vision Conference*, J. Illingworth, ed., pp. 339–348, The British Machine Vision Association, 1993.
24. T. Cootes, A. Hill, C. 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.
25. S. Solloway, C. Hutchinson, J. Waterton, and C. Taylor, "The use of active shape models for making thickness measurements of articular cartilage from MR images," *Magnetic Resonance in Medicine* **37**(6), pp. 943–952, 1997.
26. 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.
27. G. Hamarneh and T. Gustavsson, "Deformable spatio-temporal shape models: Extending ASM to 2d+time," in *Proceedings of the British Machine Vision Conference*, T. Cootes and C. Taylor, eds., The British Machine Vision Association, 2001.
28. S. Mitchell, B. Lelieveldt, R. van der Geest, H. Bosch, J. 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**(5), pp. 415–423, 2001.
29. B. van Ginneken, A. Frangi, J. Staal, B. ter Haar Romeny, and M. Viergever, "Active shape model segmentation with optimal features," *IEEE Transactions on Medical Imaging* **21**(8), pp. 924–933, 2002.
30. G. Behiels, F. Maes, D. Vandermeulen, and P. Suetens, "Evaluation of image features and search strategies for segmentation of bone structures in radiographs using active shape models," *Medical Image Analysis* **6**(1), pp. 47–62, 2002.
31. M. de Bruijne, B. van Ginneken, W. Niessen, J. Maintz, and M. Viergever, "Active shape model based segmentation of abdominal aortic aneurysms in CTA images," in *Medical Imaging: Image Processing*, M. Sonka and M. Fitzpatrick, eds., *Proceedings of SPIE* **4684**, pp. 463–474, SPIE Press, 2002.
32. C. Goodall, "Procrustes methods in the statistical analysis of shape," *Journal of the Royal Statistical Society B* **53**(2), pp. 285–339, 1991.
33. I. Dryden and K. Mardia, *Statistical Shape Analysis*, Wiley Series in Probability and Statistics, 1998.
34. M. Rogers and J. Graham, "Robust active shape model search," in *Proceedings of the European Conference on Computer Vision (ECCV'02)*, A. Heyden, G. Sparr, M. Nielsen, and P. Johansen, eds., *Lecture Notes in Computer Science* **2353**, pp. 517–530, Springer, 2000.
35. Y. Wang and L. Staib, "Statistical shape and smoothness models for boundary finding with correspondence," *IEEE Transactions on Pattern Analysis and Machine Intelligence* **22**(7), pp. 738–743, 2000.
36. B. van Ginneken, M. de Bruijne, M. Loog, and M. Viergever, "Interactive shape models," in *Medical Imaging: Image Processing*, M. Sonka and M. Fitzpatrick, eds., *Proceedings of SPIE* **5032**, SPIE Press, 2003. In Press.
37. J. Hug, C. Brechbühler, and G. Székely, "Model-based initialisation for segmentation," in *Proceedings of the European Conference on Computer Vision (ECCV'00), Part II*, D. Vernon, ed., *Lecture Notes in Computer Science* **1843**, pp. 290–306, Springer, 2000.