

Active shape model segmentation using a non-linear appearance model: application to 3D AAA segmentation

Marleen de Bruijne, Bram van Ginneken, Wiro J. Niessen, and Max A. Viergever

ABSTRACT

Active Shape Models (ASM) have proven to be an effective approach for image segmentation. In some applications, however, the linear model of gray level appearance around a contour that is used in ASM is not sufficient for accurate boundary localization. This report concerns one of those applications: segmenting the thrombus in CTA images of abdominal aortic aneurysms (AAA).

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 is trained 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 AAA datasets. Segmentation using the k NN appearance model significantly outperforms the original ASM scheme; average volume errors are 5.1% and 45% respectively.

1. INTRODUCTION

Segmentation methods that are trained on examples are becoming increasingly popular in medical image analysis. Of particular interest are the deformable template techniques¹ that model both the shape and the gray level appearance of an object, such as Active Shape Models (ASM),² Active Appearance models (AAM),³ or Active Blobs.⁴ These techniques can produce correct results even in the case of missing or confusing boundary evidence. In this work we shall concentrate on the frequently used ASMs, which consist of a landmark-based shape model, gray value appearance models around the landmarks, and an iterative optimization scheme. ASMs have been applied to a variety of segmentation tasks in medical imaging.⁵⁻¹⁶

Conventional ASMs use a linear model of boundary appearance to fit the shape model to the image. Separate models are constructed for each landmark. This works well in applications where gray level appearance is fairly consistent along the object boundary, or if appearance variation is strictly correlated with landmark position. However, in many medical image segmentation tasks the object to segment lies within variable anatomy. A given landmark can then be next to different tissue types, causing boundary appearance to vary largely. This is the case for most vascular images. In such circumstances, a linear model of gray value appearance may produce unreliable results.

In this paper we develop a non-parametric appearance model that can deal with highly varying gray values. In contrast with the conventional ASM approach, the new appearance model is trained on both true and false examples of boundary profiles, thus more effectively employing prior information on gray level structure around the object contour. The probability that a given image profile belongs to the boundary is obtained using k nearest neighbor (k NN) probability density estimation.

This work was motivated by the problem of segmenting the aneurysm outer boundary in CTA scans of abdominal aortic aneurysms (AAA). This so-called thrombus segmentation is needed for volume measurements, a key component of risk analysis for AAA patients.¹⁷⁻¹⁹ Most publications on computerized AAA segmentation have concentrated on segmenting the contrast-filled lumen.²⁰⁻³¹ Thrombus segmentation is a more difficult problem, complicated by regions of low boundary contrast and by many neighboring structures in close proximity to the aneurysm wall. Previous approaches to thrombus segmentation have used deformable models with local smoothness constraints.³²⁻³⁴ None of these methods produced results that are sufficiently accurate for clinical use.

If boundary evidence is incomplete, which is often the case in AAA images, local smoothness constraints are not sufficient to prevent the model from "leaking" into neighboring tissue. More specific, global, shape constraints are required. Strong, but realistic, constraints can be derived from a training set of segmented examples, as is done in ASM. We previously showed that the linear boundary appearance model that is originally used in ASM does not perform well in AAA boundary localization.³⁵ We adopted a slice-by-slice segmentation approach in which the similarity between adjacent image slices can be exploited to steer shape model deformation. Performance of this scheme deteriorates with distance from the reference contour, and extensive re-initialization was needed.

In this paper we present a three-dimensional ASM based approach. A model of cylindrical shape variations is used, restricting the deformation to in-slice landmark displacements. The shape model contains both statistical and synthetic deformation modes. The original ASM appearance model is replaced by the new multi-class k NN model. The original ASM segmentation scheme is briefly described in Section 2. The new appearance model is presented in Section 3. Several adaptations to the ASM scheme for AAA segmentation, including the previously described tubular shape model, robust optimization methods, and a constrained fit based on manual model initialization, are discussed in Section 4. Section 5 presents the results of a series of leave-one-out experiments on 23 datasets. Discussion and conclusions are given in Sections 6 and 7.

2. ACTIVE SHAPE MODELS

In ASMs,² shape variations in a training set are described using a Point Distribution Model (PDM). The shape model is fitted to new images using a model of boundary appearance and an iterative optimization scheme. This section briefly reviews the three parts of ASM: the shape model, the appearance model, and the optimization scheme.

2.1. Point distribution models

A statistical model of object shape and shape variation is derived from a set of s training examples. Each training example is described by a shape vector \mathbf{x} containing the coordinates of n 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. The shape vectors are aligned using Procrustes Analysis and transformed into the tangent space to the mean shape.²

Principal Component Analysis (PCA) is applied to the aligned shape vectors. To this end, the mean shape $\bar{\mathbf{x}}$, the covariance matrix \mathbf{S} , and the eigensystem of \mathbf{S} are computed. 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 variation; a small number of modes usually explains most of the variation. Each shape \mathbf{x} in the set can then be approximated by

$$\mathbf{x} \approx \bar{\mathbf{x}} + \Phi \mathbf{b} \quad (1)$$

where Φ consists of the eigenvectors corresponding to the t largest eigenvalues, $\Phi = (\phi_1 | \phi_2 | \dots | \phi_t)$, and \mathbf{b} is the model parameter vector that weighs the contribution of each of the modes.

2.2. Appearance model

Fitting the shape model to a new image requires a measure of probability that an image point belongs to the boundary. The object's boundary appearance can be learned from the training set.

In the original ASM formulation, a linear model is built from gray value profiles of length k that are sampled around the landmarks from the training set, perpendicular to the object contour. 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:

$$f(\mathbf{g}_s) = (\mathbf{g}_s - \bar{\mathbf{g}})^T \mathbf{S}_g^{-1} (\mathbf{g}_s - \bar{\mathbf{g}}). \quad (2)$$

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

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 voxel and sampled with a spacing of two voxels, and subsequent resolutions are obtained by doubling both the kernel width and the sample spacing. Because of the large anisotropy, the convolution with a Gaussian is performed in the x and y directions only. In Section 3 we will propose an alternative to this appearance model.

2.3. Optimization

In ASMs, the shape model is fitted to new images using a fast deterministic optimization scheme. The process initializes with a plausible shape, usually the mean. The initial position can be pointed out by a user, it can be the mean position in the training images, or it can be estimated using e.g a Hough transform approach.^{36,37}

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 the gray value models that was discussed in Section 2.2. Iteratively, the current shape is translated, rotated, scaled, and deformed such that the sum of squared distances between the landmarks and the optimal positions is minimized. The optimal shape parameters \mathbf{b} are computed using

$$\mathbf{b} = \Phi^T(\mathbf{x} - \bar{\mathbf{x}}). \quad (3)$$

Hard limits are applied to constrain the elements of \mathbf{b} to plausible values:

$$|b_i| \leq f_c \sqrt{\lambda_i}. \quad (4)$$

This process of adjusting landmark positions and shape parameters is repeated a fixed number of N times, whereupon it is repeated at the next level of resolution.

3. A NON-LINEAR APPEARANCE MODEL

We previously showed that the linear boundary appearance model that was originally used in ASM does not perform well in AAA boundary localization for slice-by-slice segmentation.³⁵ A shortcoming of this gray value model is that only the appearance of the correct boundary is learned from the training set. If a profile is similar to many profiles belonging to the background in the training set, that has no effect on the obtained quality of fit. Furthermore, the underlying assumption of a normal profile distribution often does not hold.

To deal with a non-linear profile distribution, Bosch et al.³⁸ performed a non-linear normalization to transform the distribution into a Gaussian. This method was developed specifically for the asymmetric, but unimodal distribution of profiles that is present in echocardiograms. Brejl and Sonka³⁷ applied fuzzy c-means clustering to feature vectors derived from intensity profiles, allowing for a (known) number of separate normal distributions. Van Ginneken and co-authors¹³ did not use intensity profiles, but extracted local texture features on separate grid points around the landmarks. The grid points are labeled as inside or outside the object, an optimal set of features is selected, and a k NN classifier is trained on the labeled points. The optimal landmark position 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 can deal with arbitrary distributions as long as the texture of object and background are different. In medical image segmentation tasks, surrounding structures are often similar to the object in gray value and texture, and the ordering of gray values along the profile can become important. In addition, in small or irregularly shaped objects, the extreme "inside" part of the profile may in fact be outside the object on the other side, which makes the method proposed in¹³ less reliable.

We propose to treat the position evaluation step in the ASM optimization as a classification of boundary profiles. Like in the original ASM formulation, gray value profiles are sampled from the training set, but now a classifier is trained on both correct and incorrect boundary profiles. Raw intensity profiles are used instead of the normalized derivative profiles of the linear model. 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

$$P(\text{boundary}|g_s) = \frac{n_{\text{true}}}{k_{nn}}, \quad (5)$$

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

By analogy with the original ASM formulation, separate models are built for different resolutions. When the profile sample spacing is doubled, the spacing d_{shift} between the shifted examples is doubled as well.

In vascular images, there is no true anatomical correspondence between the landmarks of different shapes. Therefore, we use one appearance model for all landmarks together, instead of building separate models as is more commonly done in ASM. Pilot experiments on image slices have shown that this approach gives slightly better results for both the linear and the k NN model, even if many training examples are available.

4. ADAPTATIONS OF ASM

Apart from the appearance model, several other aspects of the original ASM formulation are altered to enable automated AAA segmentation. Section 4.1 briefly reviews the flexible PDM for tubular structures that was presented in. Several robust optimization methods are discussed in Section 4.2, and a method to constrain the model fit on basis of the manual initialization is presented in Section 4.3.

4.1. PDMs for tubular structures

The model is constructed from contours in slices through the object. The original CT-slices are used since they are perpendicular to the body axis and give approximately perpendicular cross-sectional views of the aorta. A fixed number of slices n_z is interpolated between beginning and end of the aneurysm. The aneurysm outline is drawn manually by an expert. An equal number of landmarks n_{xy} is placed in each slice, equidistantly along the contour. The starting point of a contour is the posterior point with the same y -coordinate as the center of mass.

We model 3D cylindrical shape variations, restricting the deformation to in-slice landmark displacements. Before the model is fitted to a new image, the user indicates the beginning and end of the desired segmentation, thus removing the need for scaling in the z direction. As a consequence, the shape vectors contain only x and y coordinates.

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. We have applied two generalizations for PDMs of tubular structures that were discussed in more detail in.

First, the cross-sectional and axis shape variations are modeled independently. This results in two shape models Φ_{cross} and Φ_{axis} , both consisting of the eigenvectors that correspond to the t largest eigenvalues. The two models are combined into one model by computing the principal components of

$$(\Phi_{\text{cross}} \mathbf{W}_{\text{cross}} | \Phi_{\text{axis}} \mathbf{W}_{\text{axis}}) \quad (6)$$

where $\mathbf{W}_{\text{cross}}$ and \mathbf{W}_{axis} are diagonal weight matrices of the corresponding $\sqrt{\lambda_i}$. This results in a model containing twice the number of modes of a normal PDM, provided that $2(s-1) < D$, where D is the dimension of the shape vectors.

Second, synthetic deformation modes, obtained from a smoothness prior on deformation in x , y and z -directions, are added. The modes are given by the eigenvectors of a *smoothness matrix* with elements

$$e^{-\left(\frac{d_{i,j}}{2\sigma}\right)^2}, \quad (7)$$

where i and j are the row and column indices, and $d_{i,j}$ are given by

$$d_{i,j} = \text{Min}\{|i - j|, |i - j + n_c|, |i - j - n_c|\} \quad (8)$$

for a cyclic sequence of n_c coordinates (x and y deformation in our tubular object), and

$$d_{i,j} = |i - j|, \quad (9)$$

for a non-cyclic sequence of n_z coordinates (z deformation). Subsequently, the separate deformation modes are combined into shape vectors describing full three-dimensional deformation, resulting in an orthonormal set of $2n$ vectors describing smooth cylindrical deformations. In practice, a much smaller number of harmonics is chosen, such that only low-frequency deformations remain. The synthetic model is combined with the statistical model in the same way as the axis and cross-section models are combined (see Equation 6).

The original model is applied to obtain an initial estimate, up to the second highest resolution. The fit is then refined using the extended model on the smallest scale.

4.2. Robust optimization

The objective function that is minimized during ASM optimization, the sum of squared distances, is sensitive to outliers. This can cause problems in segmenting objects with noisy or highly variable edge evidence. Modifications to the fitting algorithm have been proposed to make it more outlier resistant. Rogers and Graham¹⁵ 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 Sonka⁹ 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,¹¹ Behiels et al.,¹⁴ and Mitchell et al.³⁹ proposed to use dynamic programming to favor landmark displacements that are smoothly varying between neighbors.

In our experiments we have applied dynamic programming regularization^{11, 14, 39} 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 in the anisotropic CTA data. The effect of outliers is reduced further by applying a weighted least squares fit¹⁵ on the basis of image information. Equation 3 then becomes

$$\mathbf{b} = \mathbf{K}(\mathbf{x} - \bar{\mathbf{x}}), \quad (10)$$

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

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 this reads

$$P(\text{boundary}|g_s) = ce^{\frac{-f(g_s)}{2}}, \quad (11)$$

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 5.

4.3. Initialization and constrained optimization

The complexity of the images and the local nature of ASM optimization require an accurate initialization. In our segmentation system, the user draws 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^{24, 31} or a—more easily automated—lumen segmentation^{28, 31, 33, 34} could be used for initialization.

During segmentation, the two manually drawn slices constrain the fit in the same way as in the initialization.

Table 1. Parameters of the segmentation scheme.

PDM (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
f_c	3	Bound on eigenvalues λ_i (Equation 4)
Synthetic deformation		
α	0.1	Weight of synthetic model
σ	4	Smoothness scale of synthetic deformation (Equation 7)
t_{xy}	13	Number of xy smoothing modes
t_z	12	Number of z smoothing modes
Appearance model (Sections 2.2 and 3)		
k	7	Patch length
n_{shift}	2	Number of shifted examples on both sides of the contour (k NN model)
d_{shift}	2	Magnitude 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 evaluate on either side of the current position
L	4	Number of resolution levels
N	5	Number of iterations per resolution level

5. 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.

5.1. Appearance model

First, the performance of the two appearance models is compared independent of the rest of the segmentation process. Hereto, the optimal landmark positions according to the two gray value models are determined, in a search region symmetric around the correct landmark positions as provided by the manual tracings. The distances between the ‘optimal’ and the correct positions are measured. The resulting root mean squared distances as a function of the size of the search region are shown in Figure 1. 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, which suggests that an accurate initialization and a restrictive shape model may be needed.

5.2. Segmentation

Given the initialization described in 4.3, the segmentation method using the extended shape model and the k NN gray value model converged successfully in 21 out of 23 cases. Examples of segmented slices, randomly chosen from these 21 datasets, are shown in Figure 2.

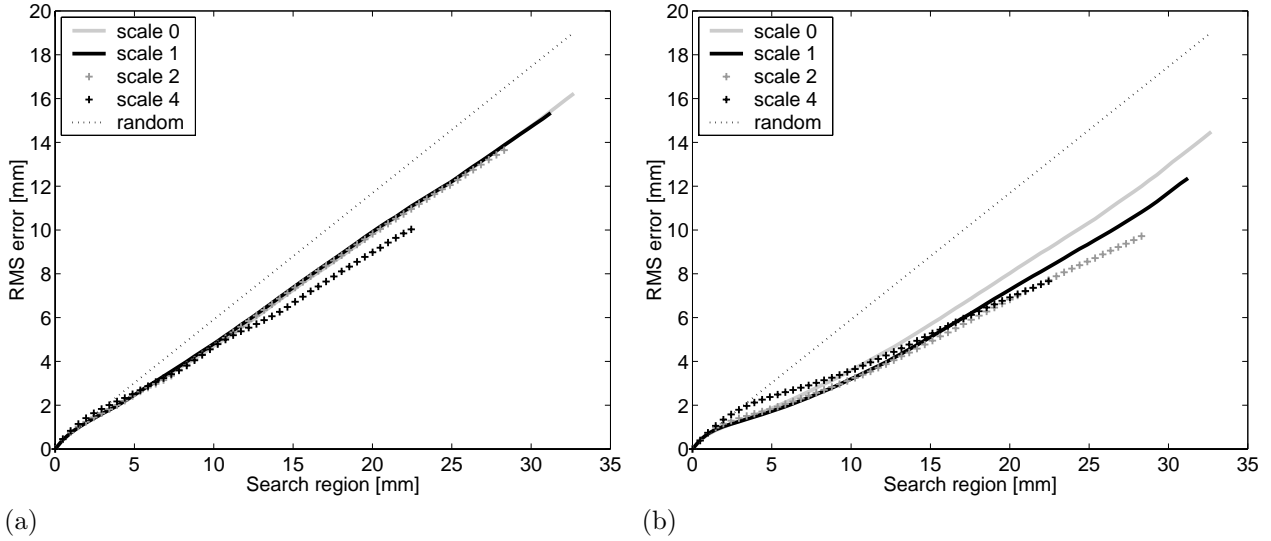


Figure 1. Root mean squared error of landmark positioning as a function of the length of the search region on either side of the contour, for (a) the Mahalanobis distance, and (b) the k NN gray value model. The dotted lines denote the expected error for random positioning.

Figure 3 shows the segmentation errors 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 1.9 and 8.1 mm (3.9 and 17 voxels). The relative volumes of overlap are 95% and 64%, and average volume errors are 5.1% and 45%. 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 4. If these two problematic datasets are left out of consideration, the average error of the remaining 21 datasets is 1.4 mm. The corresponding volume of overlap is 96% and the relative volume error 2.8%. 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 the literature⁴⁰ are expressed relative to the pure thrombus volume, excluding the passable 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.³¹ Under these assumptions, the thrombus volume error is 4.1% on average for the 21 successful segmentations using the k NN model. In a reproducibility study on manual thrombus segmentation, Wever et al.⁴⁰ reported a mean signed inter-observer error of 1.89% with a standard deviation of 4.22%, leading to a repeatability coefficient of 8.30%. Assuming a Gaussian distribution of errors, the mean unsigned volume error in that study would be 3.92%, comparable to our 4.1%.

Performance of the different shape models is compared in Figures 5 and 6. In the two problematic datasets, segmentation errors increase if a more flexible shape model is used. For the remaining 21 datasets, segmentation is improved significantly by independent modeling of axis and cross-section ($p = 0.002$ in a paired T-test), and by using the model extended with synthetic modes for refinement ($p < 0.00001$). If the extended model is used at all resolution levels, however, segmentations are significantly worse. Although still accurate at many locations, the extended model easily includes other edges into the segmentation. Examples of errors that can occur using this model are given in Figure 7.

6. DISCUSSION

The results clearly demonstrate that the original ASM appearance model can not be used for AAA segmentation, while the new k NN appearance model obtains accurate results in most cases. 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. Second, we do




Figure 2. Image slices taken randomly from the 21 successful segmentations, showing 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.




Figure 3. Root mean squared point-to-contour segmentation error for all 23 datasets, for the linear one-class model (gray) and the k NN multi-class model (black).




Figure 4. The two datasets in which the k NN method was unable to find an adequate segmentation, showing 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 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 ??). In the second dataset the aneurysm is embedded in other structures with similar gray values 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 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, or the incorporation of information on the enclosed region in the segmentation process.




Figure 5. Segmentation errors obtained using the k NN appearance model and one of the shape models; normal three-dimensional PDM (light gray), independent axis and cross-section TPDM (dark gray), and TPDM extended with synthetic smooth deformation modes in the last stage of model fit (black).

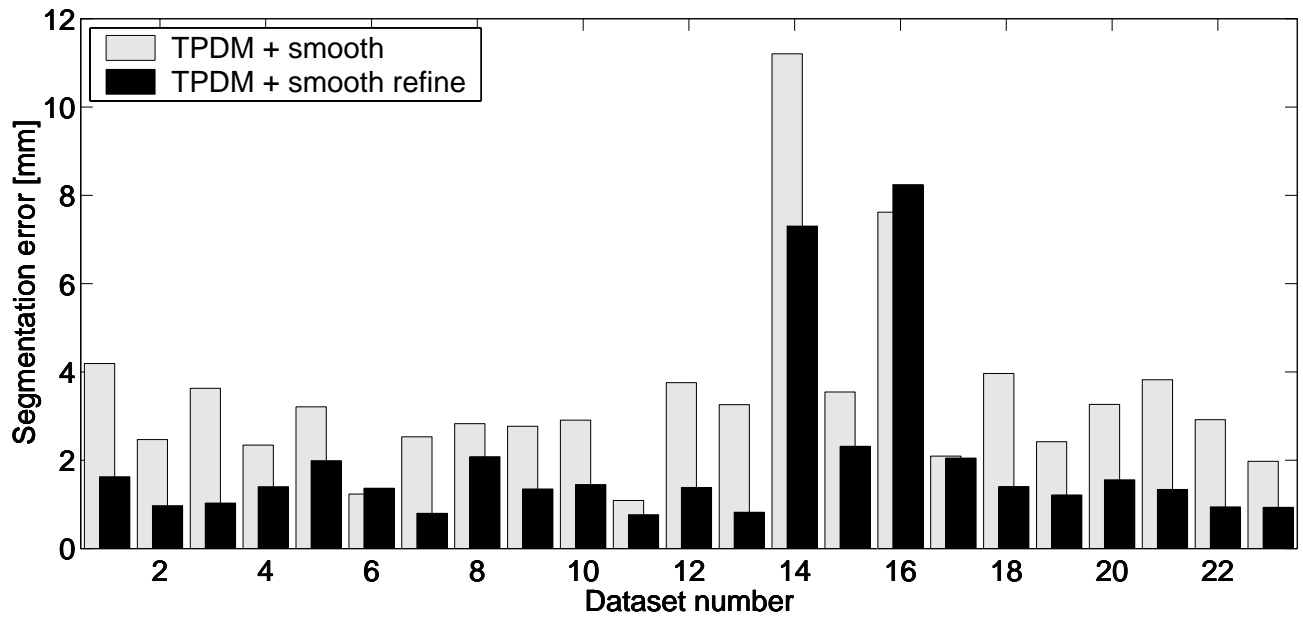


Figure 6. Performance of the TPDM extended with synthetic deformation, used at all levels of resolution (gray), or only as a refinement step at the highest level (black).




Figure 7. Erroneous segmentations obtained when the extended PDM is used at all resolution levels (continuous line). The manual segmentation is denoted with dots. The flexible model fits the correct boundary in large parts of the image, but easily draws to wrong boundaries, such as the spine (left), the vena cava (middle), or the aortic lumen and metal artifacts (right).

not assume a Gaussian intensity profile distribution but estimate the distribution non-parametrically, which can be important if either the object or the background varies largely in gray value or texture.

This model may improve results of contour-based segmentation in many applications that suffer from large appearance variations. Even if the appearance is similar for a given landmark in different images, but varies between landmarks, in small training sets the use of one non-parametric model for all landmarks together may improve upon separate linear models for each landmark. However, if the boundary appearance at one landmark is similar to the neighborhood appearance of another landmark, mixing the model at different landmarks produces unreliable results. In vascular images, where the model landmarks usually do not correspond to anatomical landmarks, boundary appearance is not directly related to the landmark number and the construction of separate models is of little use. One could think of a compromise between the two, where one model is constructed for all landmarks, and an additional location parameter in the feature vector biases the collection of neighbors visited towards profiles that were sampled at a similar location. This location parameter could denote, for instance, the length along the vessel or the angle of the given landmark with respect to the starting point.

The inclusion of examples of shifted profiles raises the question how many and which false examples should be taken into account. In our application, the first two shifted profiles (one on each side of the contour) made a large difference, while adding more examples increased boundary localization accuracy only marginally.

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 in various ways. Possible approaches include using an approximate k NN classification,⁴¹ using less shifted examples, selecting a random subset of training profiles or pruning the k NN tree, or 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.

Two extensions to PDM, both relaxing the shape constraints in an over-specific model, were tested on AAA segmentation. Segmentations were significantly improved by modeling the object axis and cross-sections independently. Accuracy can still be increased by extending the model with smooth synthetic deformation modes. However, the extended model lacks the specificity required for accurate object recognition. One of the problems in this combination of a point based optimization method with Fourier modes is that landmarks are allowed to slide along the contour, while the distance between landmarks is not preserved. A model that has a dense sampling on well defined boundaries and sparse sampling in poor contrast region does provide a better fit, but not necessarily a better contour. A segmentation scheme wherein points are redistributed equidistantly after each iteration would probably improve the results for the extended model at large scales. In its current form, the extended model should only be used to refine the fit at high resolution. An elegant solution would be to increase the dimensionality of the shape model at increasing resolution levels in the segmentation process.

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.

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.^{20, 22–28, 30, 31, 33, 34} 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 4). 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 al⁴² 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,⁴³ 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.

7. CONCLUSIONS

A new gray value model for use in contour-based image segmentation has been presented. The model significantly outperforms the ASM gray value model in boundary localization in CTA images of abdominal aortic aneurysms ($p < 0.00001$). Obtained volume errors are comparable to inter-observer errors reported in the literature.

REFERENCES

1. A. Jain, Y. Zhong, and M. Dubuisson-Jolly, "Deformable template models: A review," *Signal Processing*, vol. 71, pp. 109–129, December 1998.
2. T. Cootes, C. Taylor, D. Cooper, and J. Graham, "Active shape models – their training and application," *Computer Vision and Image Understanding*, vol. 61, no. 1, pp. 38–59, 1995.
3. T. Cootes, G. Edwards, and C. Taylor, "Active appearance models," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 23, no. 6, pp. 681–684, 2001.
4. S. Sclaroff and J. Isidoro, "Active blobs," in *Proceedings of the Sixth International Conference on Computer Vision (ICCV'98)*, (Bombay, India), pp. 1146–1153, IEEE Computer Society Press, 1998.
5. 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.
6. 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*, vol. 12, no. 6, pp. 355–366, 1994.
7. P. Smyth, C. Taylor, and J. Adams, "Automatic measurement of vertebral shape using active shape models," *Image and Vision Computing*, vol. 15, no. 8, pp. 575–581, 1997.
8. 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*, vol. 37, no. 6, pp. 943–952, 1997.
9. 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.
10. P. He and J. Zheng, "Segmentation of tibia bone in ultrasound images using active shape models," in *Proceedings of IEEE - Engineering in Medicine and Biology Society Conference*, 2001.
11. 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.
12. M. Kohnen, A. Mahnken, J. Kesten, E. Koepfel, R. Günther, and B. Wein, "A three dimensional knowledge based surface model for segmentation of organic structures," in *Medical Imaging: Image Processing* (M. Sonka and M. Fitzpatrick, eds.), vol. 4684 of *Proceedings of SPIE*, pp. 485–494, SPIE Press, 2002.
13. 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*, vol. 21, no. 8, pp. 924–933, 2002.
14. 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*, vol. 6, no. 1, pp. 47–62, 2002.
15. 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.), vol. 2353 of *Lecture Notes in Computer Science*, pp. 517–530, Springer, 2002.

16. H. van Assen, R. van der Geest, H. Lamb, M. Danilouchkine, J. Reiber, and B. Lelieveldt, "Three-dimensional active shape model matching for left ventricle segmentation in cardiac CT," in *Medical Imaging: Image Processing* (M. Sonka and M. Fitzpatrick, eds.), vol. 5032 of *Proceedings of SPIE*, SPIE Press, 2003. In Press.
17. 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*, vol. 20, no. 2, pp. 177–182, 2000.
18. 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*, vol. 8, no. 4, pp. 380–389, 2001.
19. 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*, vol. 9, no. 2, 2002.
20. 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.), vol. 3034 of *Proceedings of SPIE*, pp. 86–94, SPIE Press, 1997.
21. 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.), vol. 3338 of *Proceedings of SPIE*, pp. 938–946, SPIE Press, 1998.
22. A. Bulpitt, E. Berry, R. Boyle, D. Scott, and D. Kessel, "A deformable model, incorporating expected structure information, for automatic 3D segmentation of complex anatomical structures," in *Computer Assisted Radiology and Surgery* (H. Lemke, M. Vannier, K. Inamura, A. Farman, and K. Doi, eds.), vol. 1214 of *Excerpta Medica International Congress Series*, pp. 572–577, Elsevier Publishers, 2000.
23. G. Rubin, D. Paik, P. Johnston, and S. Napel, "Measurement of the aorta and its branches with helical CT," *Radiology*, vol. 206, pp. 823–829, Mar 1998.
24. O. Wink, W. Niessen, and M. 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.
25. 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.
26. 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.), vol. 2124 of *Lecture Notes in Computer Science*, pp. 238–246, Springer, 2001.
27. 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.), vol. 2082 of *Lecture Notes in Computer Science*, pp. 380–387, Springer, 2001.
28. 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.
29. K. Subramanyan, D. Smitha, J. Varmab, and S. Chandraa, "Automatic vessel extraction and abdominal aortic stent planning in multi-slice CT," in *Medical Imaging: Image Processing* (M. Sonka and M. Fitzpatrick, eds.), vol. 4684 of *Proceedings of SPIE*, pp. 114–122, SPIE Press, 2002.
30. A. Giachetti, M. Tuveri, and G. Zanetti, "Reconstruction and web distribution of measurable arterial models," *Medical Image Analysis*, vol. 7, no. 1, pp. 79–93, 2002.
31. 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.
32. 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.), vol. 1191 of *Excerpta Medica International Congress Series*, pp. 213–217, Elsevier Publishers, 1999.

33. 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.), vol. 4684 of *Proceedings of SPIE*, pp. 1681–1689, SPIE Press, 2002.
34. K. Subramanyan, M. Steinmiller, D. Sifri, and D. Boll, "Automatic aortic vessel tree extraction and thrombus detection in multi-slice CT," in *Medical Imaging: Image Processing* (M. Sonka and M. Fitzpatrick, eds.), vol. 5032 of *Proceedings of SPIE*, SPIE Press, 2003. In Press.
35. 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.), vol. 4684 of *Proceedings of SPIE*, pp. 463–474, SPIE Press, 2002.
36. A. Garrido and N. de la Blanca, "Physically based active shape models: Initialization and optimization," *Pattern Recognition*, vol. 31, no. 8, pp. 1003–1017, 1998.
37. M. Brejl and M. Sonka, "Object localization and border detection criteria design in edge-based image segmentation: automated learning from examples," *IEEE Transactions on Medical Imaging*, vol. 19, no. 10, pp. 973–985, 2000.
38. H. Bosch, S. Mitchell, B. Lelieveldt, F. Nijland, O. Kamp, M. Sonka, and J. Reiber, "Active appearance-motion models for endocardial contour detection in time sequences of echocardiograms," in *Medical Imaging: Image Processing* (M. Sonka and K. Hanson, eds.), vol. 4322 of *Proceedings of SPIE*, SPIE Press, 2001.
39. 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*, vol. 20, no. 5, pp. 415–423, 2001.
40. 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*, vol. 175, no. 5, pp. 1297–1282, 2000.
41. S. Arya, D. Mount, N. Netanyahu, R. Silverman, and A. Wu, "An optimal algorithm for approximate nearest neighbor searching," *Journal of the ACM*, no. 45, pp. 891–923, 1998.
42. 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.), vol. 5032 of *Proceedings of SPIE*, SPIE Press, 2003. In Press.
43. 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.), vol. 1843 of *Lecture Notes in Computer Science*, pp. 290–306, Springer, 2000.