# Adapting Active Shape Models for 3D Segmentation of Tubular Structures in Medical Images

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

Image Sciences Institute, University Medical Center Utrecht, The Netherlands

**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. Furthermore, the statistical shape model may be too restricted if the training set is limited.

This paper describes modifications to both the shape and the appearance model of the original ASM formulation. Shape model flexibility is increased, for tubular objects, by modeling the axis deformation independent of the cross-sectional deformation, and by adding supplementary cylindrical deformation modes. Furthermore, a novel appearance modeling scheme that effectively deals with a highly varying background is developed. In contrast with the conventional ASM approach, the new appearance model is trained on both boundary and non-boundary points, and the probability that a given point belongs to the boundary is estimated non-parametrically.

The methods are evaluated on the complex task of segmenting thrombus in abdominal aortic aneurysms (AAA). Shape approximation errors were successfully reduced using the two shape model extensions. Segmentation using the new appearance model significantly outperformed 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. The techniques that model both the shape and the gray level appearance of the object, such as Active Shape Models (ASM) [1], Active Appearance models [2], and M-Reps [3], can produce correct results even in the case of missing or confusing boundary evidence. In this paper we shall concentrate on the frequently used ASMs, which consist of a landmark-based linear shape model, linear gray value appearance models around the landmarks, and an iterative optimization scheme. ASMs have been applied to various segmentation tasks in medical imaging [4–8], most successfully in 2D segmentation of objects with fairly consistent shape and gray level appearance.

However, many segmentation problems in medical imaging are 3D, and gray levels may be variable. Often not enough training data is available to build a correct 3D model. The model will be over-constrained and hence does not generalize well to new shapes of the same class. Furthermore, if the object to segment lies within variable anatomy, such that a given landmark can be next to different tissue types, boundary appearance may vary largely. In that case, a linear model of gray value appearance may produce unreliable results. We show how ASMs can be adapted to deal with these problems. We focus on the segmentation of tubular structures, but some of the adaptations presented are more generally applicable.

We propose three main modifications to conventional ASMs. First, elongated structures are modeled more flexible by modeling the axis and cross-sectional shape deformation separately, thus assuming both types of variation to be uncorrelated. The two models are combined into one model describing both deformations, that is fitted using the regular ASM optimization scheme. Second, supplementary smooth deformation is introduced by adding synthetic covariance. Our approach is similar to that of Wang and Staib [9], but differs in that we decouple the smooth deformation in x, y and zdirections, which makes the approach feasible in 3D. Third, the linear one-class gray value model that is used in ASM is replaced by a novel non-parametric multi-class model that can deal with arbitrary gray value distributions and makes more effective use of the prior information on gray level structure around the object contour.

We have evaluated our method on segmentation of thrombus in abdominal aortic aneurysms (AAA) in CTA data. Most publications on computerized AAA segmentation have concentrated on segmentation of the contrast-filled lumen. 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. Previously reported approaches yield inaccurate results [10] or need extensive user interaction [7].

Leave-one-out experiments were performed on 23 routinely acquired CTA scans of different patients, to compare the proposed modifications with the conventional ASM approach. All images were segmented manually by an expert.

#### 2 Methods

The original ASM segmentation scheme is briefly described in Section 2.1. Section 2.2 presents several extensions to the shape model. The new appearance model is discussed in Section 2.3.

#### 2.1 Active Shape Models

In ASMs [1], shape variations in a training set are described using a Point Distribution Model (PDM). The shape model is used to generate new shapes, similar to those found in the training set, which are fitted to the data using a model of local gray value structure.

**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 **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 [1].

Principal Component Analysis (PCA) is applied to the aligned shape vectors. To this end, the mean shape  $\overline{\mathbf{x}}$ , the covariance matrix  $\mathbf{S}$ , and the eigensystem of  $\mathbf{S}$  are computed. The eigenvectors  $\phi_i$  of  $\mathbf{S}$  provide the modes of shape variation present in

the data. The eigenvectors corresponding to the largest eigenvalues  $\lambda_i$  account for the largest variation; a small number of modes usually explains most of the variation. Each shape **x** in the set can then be approximated using  $\mathbf{x} \approx \overline{\mathbf{x}} + \Phi \mathbf{b}$ , where  $\Phi$  consists of the eigenvectors corresponding to the *t* largest eigenvalues,  $\Phi = (\phi_1 | \phi_2 | \dots | \phi_t)$ , and **b** is the model parameter vector that weighs the contribution of each of the modes.

**Appearance model** Fitting the shape model to a new image requires a measure of probability that an image point belongs to the boundary. In the original ASM formulation, a linear model is constructed from gray value profiles 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  $\overline{\mathbf{g}}$  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:

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

**Optimization** 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 appearance model determines for each landmark its optimal new position among  $2n_s + 1$  candidate positions along the line perpendicular to the contour,  $n_s$  on both sides. Iteratively, the shape is translated, rotated, scaled, and deformed, such that the squared distance between the landmarks and the optimal positions is minimized. To ensure plausible shapes, the shape parameters  $b_i$  are constrained to lie within  $\pm 3\sqrt{\lambda_i}$ . 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.

#### 2.2 Extending PDMs

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. This section describes two different approaches for generalizing models of curvilinear structures. Hereto, first the automatic landmarking used in our experiments needs to be explained.

Landmarking strategy A fixed number of slices  $n_z$  is interpolated between beginning and end of the object. In AAA segmentation, the original CT-slices are used since they are perpendicular to the body axis and give approximately perpendicular cross-sectional views of the aorta. Alternatively, reformatted slices perpendicular to the object axis could be used. An equal number of landmarks  $n_{xy}$  is placed in each slice, equidistantly along contours that were drawn manually by an expert. The starting point of a contour is the posterior point with the same x-coordinate as the center of mass.



**Fig. 1.** Shape models built from the two input shapes on the top left. The axis of one of the input shapes is straight while the diameter of its cross-section increases towards the vertical center. The other input shape has a constant circular cross-section around a curved axis. A PDM built directly from these two input shapes contains one mode of shape variation, varying between thin shapes curved to the left and fat shapes curved to the right. The combined model (TPDM) finds two modes of variation; the first describes a curving of the object's axis and the second describes an increase or decrease in diameter from the ends towards the center.

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.

**Modeling axis and cross-sections separately** The ability of the model to generalize to unseen shapes can be increased by modeling the axes and cross-sections separately, thus assuming that both types of shape variation are uncorrelated. Subsequently, the two models are combined into one model describing both deformations.

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 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.

To combine the mode vectors of the two models they need to be of equal dimensions. However, while the axis modes have  $2n_z$  coordinates, the straightened modes are of dimension  $2n_z n_{xy}$ . To extend a mode of axis variation to  $2n_z n_{xy}$ , a shape vector is constructed which has the  $n_{xy}$  landmarks in each slice positioned at the axis points. If this deformation is applied to a shape x, the landmarks in each slice of x are translated such that their centroid coincides with the deformed axis.

In general, the two models will not be linearly independent. A second PCA is therefore performed to remove any correlation between the axis and the cross-sectional



Fig.2. Examples of synthetic deformation of a cylinder. The frequency of z-deformation increases from left to right, and that of x-deformation from top to bottom.

modes. The modes of shape variation of the combined model are thus given by the principal components of  $(\Phi_{cross} W_{cross} | \Phi_{axis} W_{axis})$ , where  $\Phi_{cross}$  and  $\Phi_{axis}$  are concatenations of mode vectors and  $W_{cross}$  and  $W_{axis}$  are diagonal weight matrices of the corresponding  $\sqrt{\lambda_i}$ .

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. Figure 1 illustrates the effect of this generalization.

Additional smooth variation Several authors have investigated the combination of statistical deformation modes with synthetic smooth deformation obtained from finite element method (FEM) models of shape instances, or smooth deformation independent of the object's shape. For instance, Wang and Staib [9] apply smoothness constraints in the form of a smoothness covariance matrix C that consists of positive numbers on the diagonal and off-diagonal elements representing neighboring points, so that each point is allowed more variation, and neighboring points are more likely to vary together. C is added to the covariance matrix S obtained from the training data, and an extended shape model is obtained by computing the eigenvectors of the new covariance matrix.

A disadvantage of this procedure is that the eigenvectors of the full  $D \times D$  covariance matrix have to be computed, while in the case that the number of samples s is smaller than the dimensionality of the shape vectors D, PCA requires only the eigenvectors of an  $s \times s$  matrix. Eigenvector decomposition is an  $\mathcal{O}(D^3)$  problem and becomes impractical for high dimensions.

Our approach is similar to that of Wang and Staib [9], but we circumvent the computation of the eigenvectors of the full covariance matrix by decoupling the deformation in the x, y, and z directions. The 3D deformation modes of a cylindrical object are thus built up of smooth deformations of cyclic sequences of x and y-coordinates and a noncyclic sequence of z-coordinates. For the cyclic sequences, C is circulant and therefore has sinusoidal eigenvectors. The first eigenvector is a constant, corresponding to a translation of the entire object, and subsequent eigenvectors correspond to sine-cosine pairs with an increasing number of full periods. For the non-cyclic sequence, the first eigenvector approximates a half period of a sine. Subsequent eigenvectors correspond to approximate sines with an increasing number of half periods.

We set the elements of the synthetic covariance matrix according to a Gaussian. The x and y deformation are then given by the eigensystem of the circulant  $n_{xy} \times n_{xy}$  matrix with elements  $e^{-(\frac{d_{i,j}}{2\sigma})^2}$ , where i and j are the matrix row and column indices,  $n_{xy}$  is the number of landmarks in one slice, and  $d_{i,j} = Min\{|i-j|, |i-j+n_{xy}|, |i-j-n_{xy}|\}$ . The z deformation is given by the eigensystem of a similar but non-circulant  $n_z \times n_z$  matrix, with  $n_z$  the number of slices in the model and  $d_{i,j} = |i-j|$ .

In the following we denote the eigenvectors of the xy and z deformation by  $\chi_i$  and  $\psi_i$  respectively. The deformations in the xy plane of the entire shape are now given by  $2n_{xy}$  shape vectors where the elements corresponding to x-coordinates in each slice are set according to one of the smooth x deformation modes, while the y-elements are zero, or the other way around. To include all possible variations along the z-axis, each of the xy-modes is combined with each of the z-modes by multiplying the elements in a slice of the xy-mode by the corresponding element of the z-mode:

$$\begin{cases} x_{i,j} = \psi(i) \cdot \chi(j) \\ y_{i,j} = 0 \end{cases} \quad \text{or} \quad \begin{cases} x_{i,j} = 0 \\ y_{i,j} = \psi(i) \cdot \chi(j) \end{cases}$$
(2)

where *i* is the slice index and *j* is the number of the landmark in the slice. The resulting deformation vectors are centered around the origin and normalized to unit length. The eigenvalues, used for weighting of the modes, are obtained through multiplication of the eigenvalues that correspond to the original xy and z modes.

The result is 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 eigenvalues are multiplied by a weight factor  $\alpha$  and the model is combined with the statistical model in the same way as the axis and cross-section models are combined in the previous subsection. Figure 2.2 shows several examples of smooth deformation modes applied to a cylinder.

The parameters involved in this augmented model are the smoothing scale  $\sigma$ , the number of synthetic modes retained, and the weight factor  $\alpha$ . The scale  $\sigma$  mainly weighs the modes of different frequencies; a larger  $\sigma$  increases the eigenvalues for low frequencies and decreases the eigenvalues of high frequency variation, thus favoring smoother deformation. The weight factor  $\alpha$  weighs the synthetic model with respect to the statistical model and should decrease if more training shapes are added. These parameters can for instance be selected by defining a threshold on the maximum reconstruction error allowed in leave-one-out experiments on the training data.

#### 2.3 A Nonlinear Appearance Model

We previously showed that the Mahalanobis distance to the average training profile does not perform well in AAA boundary localization [7]. A shortcoming of this gray value model is that only the appearance of the correct boundary is learned from the training set. Furthermore, the underlying assumption of a normal profile distribution often does not hold. To deal with a non-linear profile distribution, Bosch et al. [11] performed a non-linear normalization to transform an asymmetric, but unimodal distribution into a Gaussian. Brejl and Sonka [12] 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-authors [5] did not use intensity profiles, but applied local texture features and kNN classification to determine the boundary between object and background, hence allowing 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.

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 boundary profile is sampled around the landmark and perpendicular to the contour, and  $2n_{\rm shift}$  non-boundary profiles are sampled in the same direction,  $n_{\rm shift}$  displaced outwards and  $n_{\rm shift}$  displaced inwards.

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 kNN classifier is used and the posterior probability is given by

$$P(\text{boundary}|g_s) = \frac{n_{\text{boundary}}}{k},\tag{3}$$

where  $n_{\text{boundary}}$  is the number of boundary samples among the k nearest neighbors.

Like in the original ASM formulation, separate models are built for different resolutions.

## **3** Experiments and Results

A series of leave-one-out experiments is performed on 23 routinely acquired CTA images including pre-operative as well as post-operative scans. The scan resolution is  $0.488 \times 0.488 \times 2.0$  mm.

**Parameter settings** We have selected a set of parameters on the basis of pilot experiments, and keep those settings fixed throughout the experiments. The shapes are described by  $n_z = 30$  slices each containing  $n_{xy} = 50$  landmarks; a total of 3000 landmark coordinates. The number of modes of the axis and cross-section models is chosen such that both models describe at least 99% of the total variance across the training set. The smooth deformation modes are built of a smoothness matrix with scale  $\sigma = 4$ , and the 26 strongest xy deformation modes and 12 strongest z deformation modes are selected, thus allowing up to 6 sine periods in all directions. The weight factor  $\alpha$  is set such that the contribution of the synthetic model to the total variance is 10% of that of the statistical model. The statistical shape model is applied to obtain an initial estimate,



**Fig. 3.** (a) Root mean squared reconstruction error for all 23 datasets using all modes of variation, for a normal three-dimensional PDM (white), the combined model of axis and cross-sections (gray), and the combined model with additional smooth deformation (black). (b) Root mean squared segmentation error for all 23 datasets, for the linear model (gray) and the kNN model (black).

up to the second highest resolution. The fit is then refined on the smallest scale using the model extended with synthetic deformation.

The profiles of the gray value models consist of 7 samples. The *k*NN appearance model contains, in addition to a correct boundary profile, examples shifted  $d_{\text{shift}} = 2$  voxels inwards and outwards, for each landmark. The number of neighbors considered in the *k*NN probability estimation, *k*, is 80. The fitting algorithm evaluates  $n_s = 5$  possible new positions on both sides of the present landmark position, and performs 5 iterations at each of 4 resolution levels.

**Shape model evaluation** 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 3.a 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. The average error decreases to 0.74 mm if smooth deformation modes are added to the model.

Appearance model evaluation 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. Figure 4 shows the error in optimal position selection as a function of the size of search region  $n_s$  around the manual contour. The *k*NN



**Fig. 4.** Root mean squared error of landmark positioning, without fitting the shape model, as a function of the length of the search region on either side of the contour, for (a) the original ASM gray value model, and (b) the kNN gray value model. The dotted line corresponds to the expected error for random landmark selection.

model performs significantly better than the conventional ASM gray value model at all resolutions.

**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. An initial estimate is obtained by iteratively fitting the shape model to these points. 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. To make the fitting process more resistant to outliers, we have applied dynamic programming regularization [6] followed by a weighted least squares fit [8], in which the weights are given by the posterior probability obtained from the gray value model.

Segmentation results Given this initialization and the constrained optimization scheme, the segmentation method using the extended shape model and the kNN gray value



**Fig. 5.** 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 kNN gray value model (continuous line). The kNN model obtains a segmentation near the manual contour in all four cases, while the original ASM gray value model finds a satisfactory segmentation only in the third image.

model converged successfully in 21 out of 23 cases. Examples of segmented slices, randomly chosen from these 21 datasets, are shown in Figure 5.

Figure 3.b 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. One of these combines an extremely wide aneurysm with calcifications, which are usually found only at the boundary, inside the aneurysm; in the other 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 forms a plausible aneurysm shape. 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%.

Wever et al. [13] reported an inter-observer reproducibility coefficient (RC) of 8.3% and intra-observer RC of 3.2% and 5.7% for measurement of the total aneurysm volume in CTA. RC is, according to Bland and Altman [14], defined as 1.96 times the standard deviation of the differences in measurements. The RC of the automated measurements with respect to the expert measurements is 4.7%. Automated segmentations initialized by a second observer yield RC=5.2% with respect to the manual segmentations, and RC=1.7% as compared to the first set of automated measurements.

# 4 Discussion and Conclusions

Segmentation methods based on linear statistical models, such as ASM, obtain good results in many applications. In some cases however, a shape model based on statistics alone is too specific, and a linear gray value model is not always able to find the correct contour. We have presented an application—AAA segmentation—in which conventional ASM can not find a suitable segmentation.

We have shown how shape models of elongated objects can be made more flexible by modeling the object axis and cross-sections independently. The idea of decoupling different types of shape variation and treating them as independent can be applied more generally. For instance, when modeling vascular trees, different segments could be modeled separately. In multiple object models, each object can be modeled separately whereafter the objects are joined into one combined model. The general relations between different objects are then retained, while the correlation between shape variation in different objects is removed. Such a model is more flexible but also increases the risk of producing invalid shapes, like overlapping objects.

An orthonormal basis of smooth deformation modes was constructed using the eigenvectors of small matrices. The approach presented is valid for tubular objects, where the shape can be described by a stack of contours with an equal number of land-marks in each contour. In arbitrary shapes, decoupling the deformation in x, y, and z-direction would require computation of the eigenvectors of an  $n \times n$  instead of an  $3n \times 3n$  matrix, still greatly reducing computation time.

In AAA segmentation, we used contours in the original CT slices to build the model, and deformation is restricted to in-slice landmark displacements. We believe this approach is valid in the case of CTA images, which are in general highly anisotropic (in the images used in this study the voxels are over 4 times larger in the z-direction). However, the presented methods can also be applied to reformatted slices perpendicular to the object axis.

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 not assume a Gaussian intensity profile distribution but estimate the distribution non-parametrically with kNN probability density estimation. The latter is responsible for a dramatic increase in computation time; a full segmentation took on average 25 seconds on a 1.7 GHz Pentium PC when the original ASM gray value model was used and 450 seconds using the kNN model. If computation time is an issue the method could be sped up by using fewer shifted examples and pruning the kNN tree, or by using other classifiers [15]. For instance, a quadratic discriminant classifier could be used, which is equivalent to extending the original ASM gray value modeling scheme to more classes.

In leave-one-out experiments on 23 datasets, the shape approximation error was successfully reduced by modeling axis and cross-section deformation independently, and by adding supplementary smooth deformation modes. The *k*NN appearance model significantly outperforms the original one-class linear gray value model (p < 0.00001). Obtained volume errors with respect to expert segmentations are comparable to inter-observer errors reported in the literature, while the inter-observer agreement for automated segmentation initialized by two different observers is better than for manual segmentation.

#### Acknowledgments

This research was funded by the Netherlands Organization for Scientific Research (NWO). We would like to thank our colleagues M. Prinssen, M.J. van der Laan, and J.D. Blanken-

steijn from the Department of Vascular Surgery for providing the datasets and expert segmentations.

# References

- 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.
- T. Cootes, G. Edwards, and C. Taylor, "Active appearance models," *IEEE Transactions on Pattern Analysis and Machine Intelligence* 23(6), pp. 681–684, 2001.
- S. Joshi, S. Pizer, P. Fletcher, P. Yushkevich, A. Thall, and J. Marron, "Multiscale deformable model segmentation and statistical shape analysis using medial descriptions," *IEEE Transactions on Medical Imaging* 21(5), pp. 538–550, 2002.
- 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.
- 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.
- 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.
- 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.
- 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, 2002.
- 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.
- 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.
- 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., *Proceedings of SPIE* 4322, SPIE Press, 2001.
- M. Brejl and M. Sonka, "Object localization and border detection criteria design in edgebased image segmentation: automated learning from examples," *IEEE Transactions on Medical Imaging* 19(10), pp. 973–985, 2000.
- J. Wever, J. Blankensteijn, J. van Rijn, I. Broeders, B. Eikelboom, and W. Mali, "Interand intra-observer variability of CTA measurements obtained after endovascular repair of abdominal aortic aneurysms," *American Journal of Roentgenology* 175(5), pp. 1297–1282, 2000.
- 14. J. Bland and D. Altman, "Statistical methods for assessing agreement between two methods of clinical measurement," *Lancet* (1), pp. 307–310, 1986.
- 15. A. Jain, R. Duin, and J. Mao, "Statistical pattern recognition: a review," *IEEE Transactions* on *Pattern Analysis and Machine Intelligence* **22**(1), pp. 4–37, 2000.