

Three-dimensional point distribution models for tubular objects

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

ABSTRACT

Landmark based statistical shape models are becoming increasingly popular in medical image analysis. These models require a set of training shapes in which corresponding landmarks are indicated. If a limited number of training shapes is available while a large number of landmarks is needed to describe a shape properly, a model based on observations in the training set alone may be too restricted. This will often be the case in three-dimensional shape description.

This paper describes the construction of a three-dimensional shape model of a tubular object. Two approaches for generalizing this model are proposed. First, the model flexibility is increased by modeling the axis deformation independent of the cross-sectional deformation. Second, supplementary cylindrical deformation modes are added to the model.

The methods are evaluated on model construction of an abdominal aortic aneurysm from segmented CTA data. In leave-one-out experiments on 23 datasets, shape approximation errors are successfully reduced using the two shape model extensions. The combination of statistical and synthetic deformation modes performs better than either of the two models alone.

1. INTRODUCTION

Segmentation methods that are trained on examples have proven to be an effective approach for medical image segmentation. 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 (AAM),² and M-Reps,³ can produce correct results even if neighboring objects induce strong edges or if part of the boundary is not discernible.

A vital part of these methods is the underlying shape model. A good shape model should be *general*, i.e. it should be able to approximate any instance of the object class, not just those seen in the training set. Furthermore, it should be *specific*, i.e. it should generate only valid shapes, and, to reduce computation time, and it should be *compact*, i.e. it should explain the observed variation with the smallest number of modes possible.

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 relatively large number of parameters is needed to describe a shape properly. Several authors have therefore developed segmentation schemes that use a statistical shape model merely for object localization. In the final stage, the contour is allowed to refine outside the model space. For instance, Brejl and Sonka⁴ use a statistical shape model to localize the object, whereupon snake optimization or dynamic programming is applied for accurate border detection. Kervrann and Heitz propose a hierarchical framework in which global deformations are modeled using PCA on a training set and local deformations are modeled by a Markov process.⁵ Cootes and Taylor combined their global, statistical Active Appearance Models (AAM) with a local optical-flow-like registration technique that refines the model after AAM convergence.⁶ An alternative approach, which we investigate in this paper, is to extend the shape model itself by incorporating variation that is not observed in the training set.

We concentrate on the frequently used Point Distribution Models (PDM),¹ which describe shape variation across a training set by variation in positions of a set of corresponding landmarks. The 'modes of shape variation' are obtained through Principal Component Analysis (PCA) of the covariance matrix of landmark coordinates. PDMs are the shape models that are used in ASM and AAM. We describe the construction of a three-dimensional PDM of an abdominal aortic aneurysm (AAA), a more or less tubular structure. We propose two different ways to increase the generalizability of this model if the training set is small. The eventual goal in this application is automated AAA segmentation, but the techniques proposed can be used for segmenting other tubular structures as well.

Most research concerning three-dimensional PDMs has been on automated model construction. The main question is how to establish a correct correspondence between landmarks, such that a landmark represents the same anatomical point in different shapes.⁷⁻¹⁴ On the other hand, most papers on three-dimensional application of PDMs ignore the issue of optimizing actual anatomical correspondence and construct a PDM simply by interpolating an equal number of slices for each object and placing the landmarks equidistantly around the object contours in these slices.¹⁵⁻¹⁷ This seems to be a reasonable approximation in many applications, especially in the case of anisotropic data, where manual segmentations are also performed slice-by-slice, and if true anatomical landmarks are sparse or altogether absent. We shall follow this slice-based approach.

To overcome the problem of overconstraining the model, two different approaches can be distinguished. First, the number of training shapes can practically be extended by making simple assumptions about shapes that can be expected, given the set of training shapes. For instance, Hill et al.¹⁵ extended a three-dimensional PDM of brain structures in MR images by including for each training example a second example in which the left and right sides of the brain were reflected through the mid-plane. In this work we follow a similar approach. To extend the model of a tubular object, we assume that each cross-sectional shape found in the training set can occur with each observed axis. Axis and cross-sectional deformation are then modeled separately. The two models are combined into one model describing both deformations, with twice the number of modes that a normal PDM would have. This is equivalent to performing PCA on $s \times s$ shapes with all possible combinations of axes and cross-sections.

Second, the model can be extended by adding plausible synthetic deformation, for instance in the form of physically based deformation modes. An effective approach to physically based shape modeling, using modal analysis of Finite Element Method (FEM) equations, was provided by Pentland and Sclaroff.¹⁸ The object deformation is described in terms of FEM eigenmodes, which provide a frequency-ordered orthonormal basis of the object's free vibration modes if the object were made of elastic material. Several authors have combined FEM eigenmodes with statistical deformation modes obtained from a training set. Martin et al.¹⁹ used a shape model that is a linear combination of the FEM deformation modes of the object under study and the mean shape and statistical deformation modes of a training set for shape analysis of brain structures. Note that there will be redundancy in this combined model since the two sets of modes are not independent. To make it suitable for shape fitting, which requires linear projection onto the (reduced) model space, a second PCA would be needed to remove the correlation between the different modes. Garrido and Pérez de la Blanca²⁰ express the training shapes in the FEM basis and then perform PCA. Cootes and Taylor²¹ computed the covariance matrices associated with the FEM deformation of each training shape and added those, multiplied by a weighting factor, to the covariance of the training shapes about the mean. The weighting factor decreases with increasing size of the training set. The combined model thus reduces to a FEM model if only one training example is present, and tends to a normal PDM if the training set is large.

Independent of the shapes in the training set, the model can be extended by adding synthetic variance and covariance directly to the covariance matrix S obtained from the training data. This method was proposed by Cootes and Taylor,²² and was described in more detail and integrated in a Bayesian framework by Wang and Staib.^{23,24} If, for instance, the identity matrix is added to S , all landmarks are allowed to move freely in any direction, independent of the other landmarks. A smoothness constraint can be incorporated by adding a positive number to the off-diagonal elements that represent neighboring points in the covariance matrix, so that neighboring points are more likely to vary together. A model combining synthetic and statistical deformation is computed by an eigenvector decomposition of the combined covariance matrix.

A disadvantage of the above procedures for adding synthetic deformation is that the eigenvectors of the full $D \times D$ covariance matrix have to be computed, where D is the dimensionality of the shape vectors. In the case that the number of samples s is smaller than D , construction of a normal PDM would require only the eigenvectors of a $s \times s$ matrix. Eigenvector decomposition is an $\mathcal{O}(D^3)$ problem and becomes impractical for high dimensions.

To avoid eigenvector decomposition, Nastar and Ayache²⁵ derived analytic expressions for the FEM eigenmodes of curves and surfaces of different topology. We present an approach that is similar to the shape independent approach of Wang and Staib,²³ but the computation of the eigenvectors of the full covariance matrix

is circumvented by decoupling the deformation in the x , y , and z directions, which makes the approach feasible for high dimensional shapes.

Both methods we propose to extend shape models in case of a small number of training images are evaluated with respect to generality, specificity and compactness. The generalizability of a model is easily tested using cross validation experiments, wherein part of the available shapes are used to train a model that is tested on the remaining shapes. The specificity of a model is probably best measured with respect to a given object recognition task. To evaluate model specificity independent of a segmentation method, one can evaluate the compactness as a substitute for specificity; the simplest model, using the fewest parameters, is least likely to generate invalid shapes.^{26, 27} In this report we will study the compactness of the PDM and the proposed extensions on a set of 23 binary segmentations of AAA in CTA scans. The generalizability is evaluated in leave-one-out experiments.

The next section describes basic PDMs. Two extensions to PDMs are presented in Section 3, and evaluated in Section 4.

2. 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:

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

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,^{28, 29} and transformed into the tangent space to the mean shape.³⁰

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:

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

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

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 explain 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 model parameter vector \mathbf{b} weighs the contribution of each of the modes,

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

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:

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

3. PDM FOR TUBULAR STRUCTURES

This section describes the building of a PDM of a tubular object. First, the automatic landmarking is explained in Section 3.1. Sections 3.2 and 3.3 present two approaches for generalizing overconstrained models. The first practically extends the statistical model using the assumption that the cross-sectional shape is independent of the object axis. The second extends the model with synthetic smooth deformation modes.

3.1. Landmarking strategy

The input for automatic shape model construction is a binary segmentation, usually made by manual slice-by-slice delineation. A fixed number of slices n_z is interpolated between beginning and end of the object, using nearest neighbor interpolation. 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 the object boundary. The starting point of a contour is the posterior point with the same y -coordinate as the center of mass.

We model 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.

3.2. Modeling axis and cross-sections separately

The generalizability of the model is 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.

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 \mathbf{x} , the landmarks in each slice of \mathbf{x} are translated such that their centroid coincides with the deformed axis. In general, the two models will not be linearly independent, therefore a second PCA is performed to remove any correlation between the axis modes and the cross-sectional modes. The modes of shape variation of the combined model are then given by the principal components of

$$(\Phi_{\text{cross}} \mathbf{W}_{\text{cross}} | \Phi_{\text{axis}} \mathbf{W}_{\text{axis}}) \quad (7)$$

where Φ_{cross} and Φ_{axis} are concatenations of mode vectors and $\mathbf{W}_{\text{cross}}$ and \mathbf{W}_{axis} are diagonal weight matrices of the corresponding $\sqrt{\lambda_i}$. In the remainder of this thesis, we will refer to this combined model as tubular PDM (TPDM).

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.

3.3. Synthetic deformation

The model is further extended by adding synthetic variance and covariance to the model, similar to approaches for relaxing shape models of closed two-dimensional contours, as were presented by Cootes and Taylor²² and Wang and Staib.^{23,24} They added a smoothness covariance matrix C_{smooth} to the covariance matrix \mathbf{S} obtained from the training data, and a combined model of statistical and synthetic deformation is obtained by computing the eigenvectors of the new covariance matrix. C_{smooth} consists of positive numbers on the diagonal and on the off-diagonal elements representing neighboring points, so that each point is allowed more variation, and

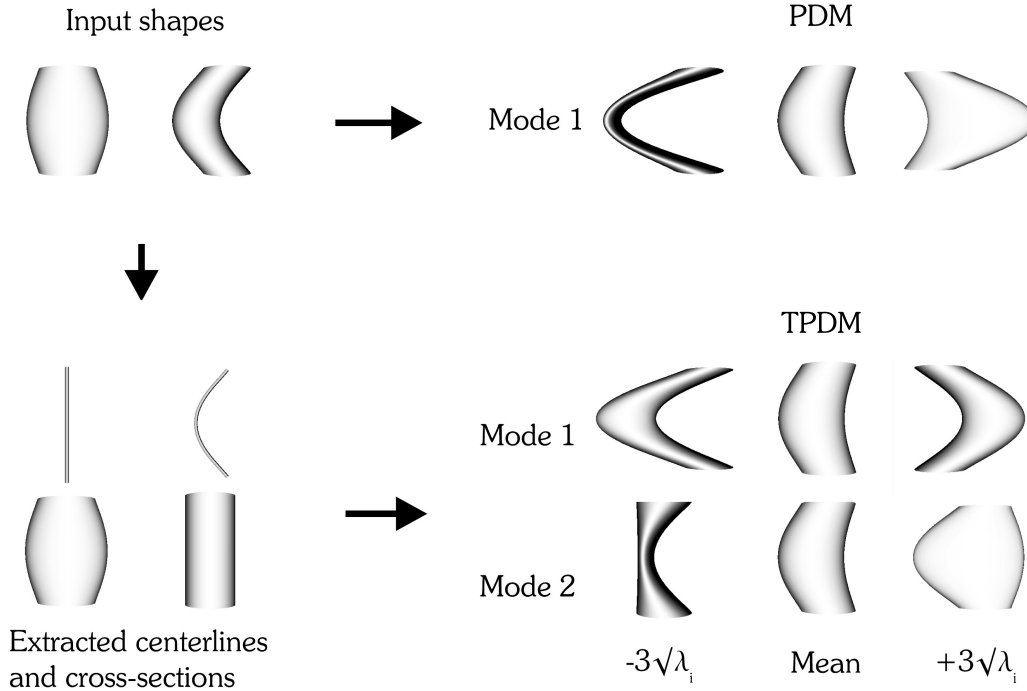


Figure 1. Shape models built from the two input shapes on the top left. The axis of one of the input shapes is straight while the diameter of its cross-section increases towards the vertical center. The other input shape has a constant circular cross-section around a curved axis. A PDM built directly 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.

neighboring points are more likely to vary together. For closed contours, C_{smooth} is circulant and therefore its eigenvectors are sinusoidal.³¹ These are also the basis functions for Fourier parameterization of a closed contour, as were proposed for deformable modeling by Staib and Duncan.^{32, 33}

This procedure becomes impractical for high-dimensional shape vectors, since an eigenvector decomposition of a matrix with the dimension of the shape vectors is required. If the number of samples s is smaller than the dimensionality of the shape vectors D , a normal PDM would require only the eigenvectors of an $s \times s$ matrix.

Our approach is similar, but we avoid the computation of the eigenvectors of the full covariance matrix by decoupling the deformation in the x , y , and z directions. For a cylindrical object, which can be described by a stack of contours with an equal number of landmarks in each contour, the deformation can be decomposed in x and y deformation modes that are equal for each contour, and a set of z deformation modes describing the variation along the axis. The three-dimensional deformation modes are thus built up of smooth deformations of a cyclic sequence of x -coordinates, a cyclic sequence of y -coordinates and a non-cyclic sequence of z -coordinates. For the cyclic sequences, C_{smooth} is circulant and therefore has sinusoidal eigenvectors. The first eigenvector is a constant, corresponding to a translation of the entire contour, 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 smoothness covariances according to a Gaussian. Both the x and the y deformation

are then described by the eigensystem of an $n_{xy} \times n_{xy}$ matrix with elements

$$e^{-\left(\frac{d_{i,j}}{2\sigma}\right)^2} \quad (8)$$

$$d_{i,j} = \text{Min}\{|i-j|, |i-j+n_{xy}|, |i-j-n_{xy}|\},$$

where i and j are the matrix row and column indices, and n_{xy} is the number of landmarks in one slice. The z deformation is given by the eigensystem of an $n_z \times n_z$ matrix with the elements according to Equation 8 and

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

n_z being the number of slices in the model.

In the following we denote the eigenvectors of the xy and z deformations by χ_i and ψ_i respectively. The possible 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 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:

$$\begin{cases} x_{i,j} = \psi(i) \cdot \chi(j) \\ y_{i,j} = 0 \end{cases} \quad \text{or} \quad \begin{cases} x_{i,j} = 0 \\ y_{i,j} = \psi(i) \cdot \chi(j) \end{cases} \quad (10)$$

where i is the slice index and j is the number of the landmark in the slice. The resulting mode vectors are centered around the origin and normalized to unit length. The eigenvalues of the three-dimensional deformation modes are obtained through multiplication of the eigenvalues corresponding 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 of the modes are multiplied by a weight factor c_w and the model is combined with the statistical model in the same way as the axis and cross-section models are combined in Section 3.2 (see Equation 7). We express c_w in terms of the total variance contained in both models and use a weight factor α to control the contribution of the synthetic model:

$$\alpha = c_w \frac{\sum_i \lambda_{i_{synthetic}}}{\sum_i \lambda_{i_{train}}}. \quad (11)$$

A combined model where statistical and synthetic modes are weighted equally is thus given by $\alpha = 1$. Figure 2 shows several examples of smooth deformation modes applied to a cylinder.

The parameters involved in this augmented model are the smoothing scale σ , the number of synthetic modes retained, and the weight factor α . The scale σ determines the degree of smoothness of the synthetic deformation. For the cyclic smoothness matrices, the eigenvectors are always the same, but a larger σ increases the eigenvalues of low frequencies and decreases the eigenvalues of high frequency variation, thus favoring smoother deformation. In the non-cyclic case, a larger σ not only weighs low frequency variation stronger, but also produces smoother eigenvectors. The weight factor α weighs the smooth 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.

4. EXPERIMENTS AND RESULTS

In this section, we evaluate the compactness and flexibility of PDMs, TPDMs, and the synthetic models. Compactness is studied by comparing the contribution of the separate modes to the total variance. The ability of the models to fit to unseen shapes is assessed in leave-one-out experiments.

Experiments were performed on 23 routinely acquired CTA images from 23 different patients, segmented manually by an expert. Of these scans, 3 were taken pre-operatively, 8 within two days after surgery, and 12 at

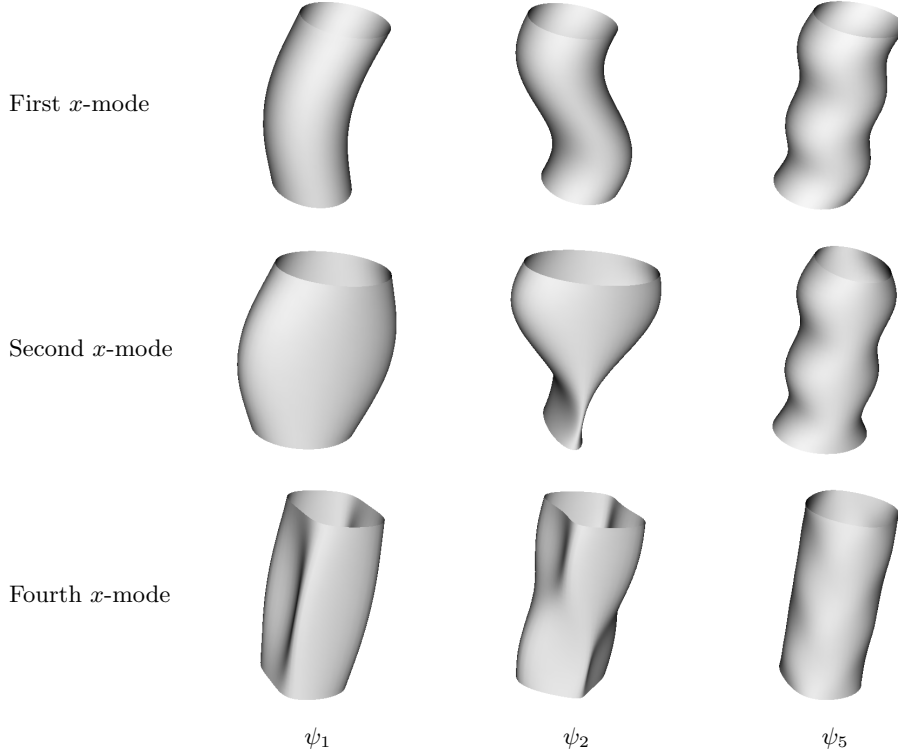


Figure 2. Examples of smooth deformation of a cylinder, the frequency of z -deformation increases from left to right, and that of x -deformation from top to bottom.

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. The shapes are described by $n_z = 30$ slices each containing $n_{xy} = 50$ landmarks; a total of 3000 landmark coordinates.

Figure 3 visualizes the first six modes of shape variation of the combined axis and cross-section model trained with all 23 datasets.

The variance associated with the different modes and models is shown in Figure 4. The total variance of the normal PDM and the combined axis and cross-section PDM is equal, but the normal PDM is more compact; it uses half the number of modes. The total variance of both the normal and the combined PDM is dominated by axis deformation, mainly evident in the first few modes.

Figure 5 illustrates the compactness of the smooth deformation models of varying scale σ . The sum of eigenvalues, associated with the total variance, is the same for all models. With increasing σ the model becomes far more compact, and describes at $\sigma = 16$ most of its variation in the first 10 modes.

The ability of the model to fit to new shapes is tested by fitting a model constructed from 22 shapes directly to the manual segmentation of the remaining image. This gives an upper bound for the accuracy that can be obtained in a segmentation scheme where the model is fitted to new image data. The shape is approximated through linear projection onto the model space (see Equation 5). No constraints are set on the model parameters. Examples of obtained shape approximations are given in Figure 6. A normal PDM is not able to fit to the upper and lower part of the aneurysm accurately. The resulting shape is quite jagged, more than the original. Although the TPDM captures the global shape better, it still shows rippling that is not present in the original shape. Furthermore, the transition from a fairly round cross-section in the upper part of the aneurysm to an elliptical cross-section in the lower part is not accurately described. The TPDM combined with synthetic deformation has approximated to the data in a smooth way, while retaining most of the variation between slices. The purely synthetic model gives a very smooth approximation.

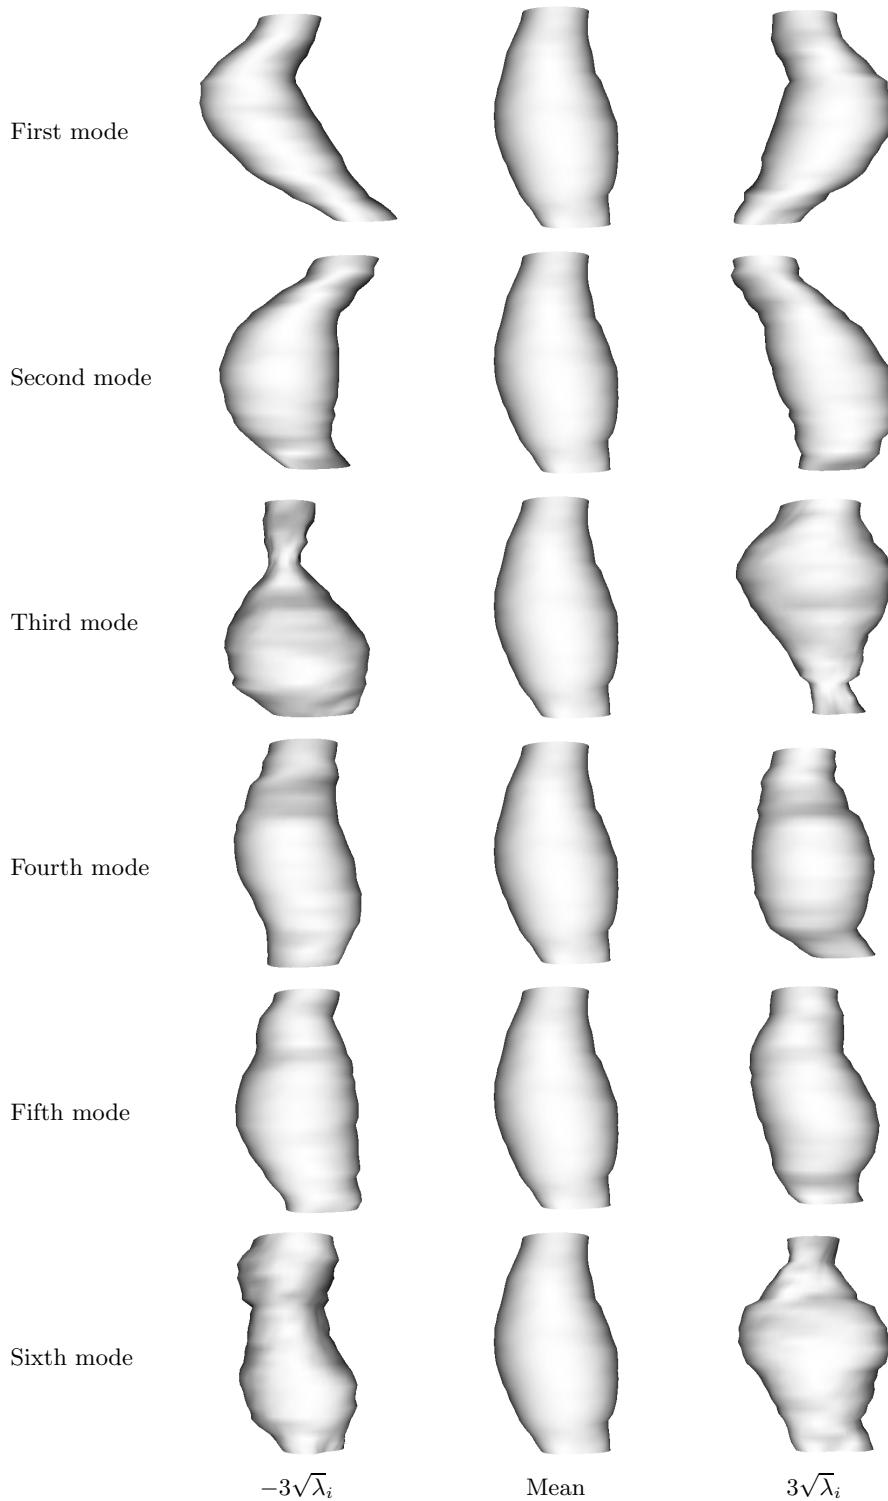


Figure 3. The effect of varying each of the first six shape parameters individually (frontal view). The largest amount of variance across the training set is caused by a curving of the axis, shown in the first two modes. The third mode describes aneurysms with a long or a short proximal neck, while mode 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.

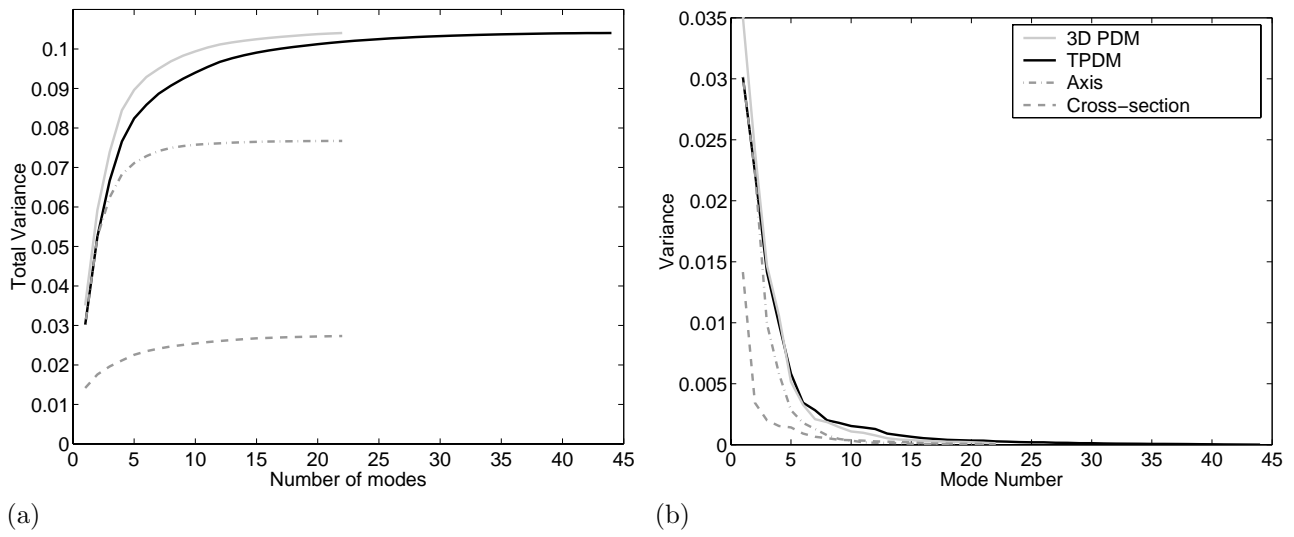


Figure 4. (a) Total variance of the model as a function of the number of modes retained. (b) Contribution to the variance of each mode.




Figure 5. Compactness of the smooth deformation model for varying scale σ .




Figure 6. Example of shape reconstruction by the different models. (a) Manual segmentation (b) Approximation by a normal PDM, all 21 modes (rms distance 4.09 mm; max distance 11.0 mm) (c) Approximation by TPDM, all 42 modes (2.88 mm; 7.91 mm) (d) Approximation by TPDM + synthetic deformation of $\sigma = 8$, 200 modes (1.30 mm; 3.58 mm). (e) Approximation by synthetic model only, $\sigma = 8$, 200 modes (1.47 mm; 4.07 mm).


Figure 7a shows the shape approximation error averaged over all datasets as a function of the number of modes retained in the model. The normal PDM and the TPDM give similar results at first. From the fourteenth mode, the extended model performs better. The difference in performance at $t=21$ modes is significant with $p < 0.0001$ in a paired T-test. The minimum root mean squared reconstruction error for the normal PDM is 2.54 mm, and 1.85 mm for TPDM. The corresponding maximum errors are 13.3 mm and 9.22 mm.

Figure 7b shows the shape approximation errors for the synthetic models, without using any information from the training set. The reconstruction performance of a model of a limited number of modes depends strongly on the smoothing scale σ . A larger σ provides smoother z -deformation and a different ordering of eigenvectors, which improves results up to $\sigma = 8$. All synthetic models need many more modes of shape variation to reach the accuracy of one of the statistical models. To achieve the reconstruction ability of one of the statistical models with 10 modes, the best synthetic model, that of $\sigma = 8$, needs 30 modes. However, the number of synthetic deformation modes is not limited by the size of the training set, and the approximation error will eventually reach zero for 3000 modes.


Figure 8a gives the shape approximation errors for PDM and TPDM combined with modes of the best synthetic model. For comparison, the curves of PDM, TPDM and the synthetic model without statistical modes are given as well. The weight is chosen such that both models contain equal total variance, $\alpha = 1$. The model is allowed to occupy the entire parameter space; no limits are applied to b_i . Therefore, α influences only the ordering of statistical and synthetic modes and has no effect on the final reconstruction error obtained. The reconstruction error equals that of a statistical model for a combined model using up to 21 modes for PDM or 42 modes for TPDM, but inclusion of more modes increases accuracy. For a model containing 200 modes, TPDM combined with synthetic deformation still performs better than PDM, and better than a synthetic model alone.

If the model parameters are constrained to lie within a certain range, as is commonly done to ensure plausible shapes, the weight factor α becomes more important. If α is very small, the synthetic deformation has a small amplitude, while a large α would allow a large amount of synthetic deformation, which probably leads to illegal shapes. The effect of α is illustrated in Figure 8b, which shows the reconstruction error if the parameters b_i are constrained to lie within $\pm 3\sqrt{\lambda_i}$. For a large range of parameters, the combined model is equal to the statistical model as long as the number of modes is smaller than the number of statistical modes available. Naturally, for the models that contain a large number of modes, a larger α — allowing more deformation — provides a better fit. At $\alpha = 8$, where the variance in the synthetic model is eight times larger than in the statistical model, the reconstruction error for a model of 200 modes is nearly as small reconstruction error without constraints on the parameters that was given in Figure 8a.

Finally, Figure 9 shows the shape approximation error for each dataset. Both the separate modeling of axis and cross-section and the addition of synthetic deformation modes improved the shape approximation accuracy in all cases.

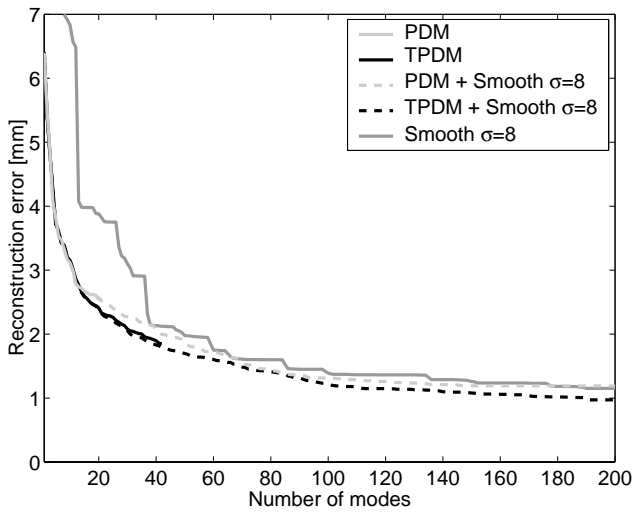


(a)




(b)

Figure 7. (a) Root mean squared point-to-point shape approximation error as a function of the number of modes, obtained using the normal PDM (gray) and the tubular PDM with independent axis and cross-section deformation (TPDM, black). (b) Shape approximation error as a function of the number of modes, for synthetic smooth deformation models of varying σ .



(a)



(b)

Figure 8. (a) Shape approximation error as a function of the number of modes, for the two statistical models combined with a synthetic model. (b) Effect of varying α . Shape approximation error as a function of the number of modes, for TPDM combined with the synthetic model of $\sigma = 8$. Shape parameters are constrained within $\pm 3\sqrt{\lambda}$.

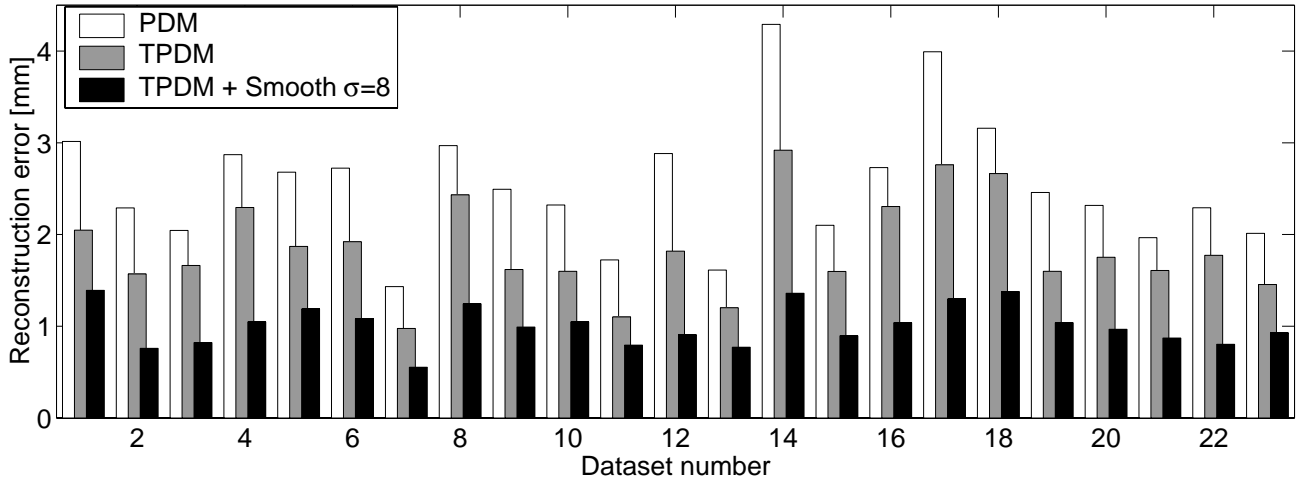


Figure 9. Leave-one-out reconstruction error for all 23 datasets. The PDM and TPDM use the maximum of 21 and 42 modes respectively, and the combined model contains 200 modes.

5. DISCUSSION AND CONCLUSION

Statistical shape models, such as PDM, form a powerful tool for the interpretation of medical images. In many cases however, especially in three or more dimensions, a shape model based purely on the statistics of a small training set can be too specific.

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, such as were constructed for instance of the left and right ventricle of the heart¹³ and of brain structures,³⁴ 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.

To extend the model further, an orthonormal basis of smooth synthetic deformation modes was constructed. The approach is similar to the smoothness prior used by Wang and Staib^{23,24} and Cootes and Taylor,²² but we greatly reduce the computational burden by modeling x , y , and z -deformations separately. The approach presented is valid for tubular objects, where the shape can be described by a stack of contours with an equal number of landmarks in each contour. In arbitrary shapes in three dimensions, 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 considerably reducing computation time.

In the construction of an AAA model, 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 AAA scans the voxels are over 4 times larger in the z -direction). The expert segmentations are obtained by delineation in the original slices as well. However, the presented methods can also be applied to reformatted slices perpendicular to the object axis.

In leave-one-out experiments on 23 datasets, the shape approximation error was successfully reduced by modeling axis and cross-section deformation modes independently, and by adding supplementary smooth deformation modes. It was shown that synthetic deformation models need more modes to describe the shapes with the same accuracy as a statistical model. Decoupling axis and cross-section shape variation improved the ability of the shape model to generalize to unseen shapes, even if a large number of synthetic deformation modes was added. For a large range of weights, a model combining synthetic and statistical shape variation tends to a normal

statistical model if the number of modes is smaller than the maximum number of modes that can be obtained from the training set.

We have discussed ways to relax shape constraints in an overly specific model. However, a model with good generalizing properties does not always give better results in image interpretation. Such a model may be able to fit to any data and thus produce physically impossible shapes. For instance, the Fourier modes that are applied can easily create self intersecting contours. A common approach is to use a restricted model for coarse object localization, and then refine the fit using a more flexible model.

With a fixed number of modes, the model with the best shape approximation ability will yield the best segmentations. We can conclude that the combined axis and cross-section PDM is a better choice for description of AAA than a normal PDM, and that the combination of statistical models with synthetic deformation is better than using only a synthetic model. The improvement in shape approximation using synthetic deformation is encouraging, but the true benefit of synthetic deformation modes in model-based AAA segmentation still needs to be established.

REFERENCES

1. 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.
2. 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.
3. 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*, vol. 21, no. 5, pp. 538–550, 2002.
4. 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.
5. C. Kervrann and F. Heitz, "A hierarchical statistical framework for the segmentation of deformable objects in image sequences," in *CVPR'94 (IEEE Conference on Computer Vision and Pattern Recognition)*, pp. 724–728, Computer Society Press, 1994.
6. C. et al, "Combining elastic and statistical models of appearance variation," in *ECCV00*, 2000.
7. M. Fleute and S. Lavallée, "Building a complete surface model from sparse data using statistical shape models: applications to computer assisted knee surgery," in *Medical Imaging Computing & Computer-Assisted Intervention* (W. Wells, A. Colchester, and S. Delp, eds.), vol. 1496 of *Lecture Notes in Computer Science*, pp. 879–887, Springer, 1998.
8. A. Brett and C. Taylor, "A framework for automated landmark generation for automated statistical 3D model construction," in *Information Processing in Medical Imaging* (A. Kuba, M. Sámal, and A. Todd-Pokropek, eds.), vol. 1613 of *Lecture Notes in Computer Science*, pp. 376–81, Springer, 1999.
9. A. Kelemen, G. Székely, and G. Gerig, "Elastic model-based segmentation of 3-D neuroradiological data sets," *IEEE Transactions on Medical Imaging*, vol. 18, no. 10, pp. 828–839, 1999.
10. C. Lorenz and N. Krahnstöver, "Generation of point-based 3D statistical shape models for anatomical objects," *Computer Vision and Image Understanding*, vol. 77, no. 2, pp. 175–191, 2000.
11. Y. Wang, B. Peterson, and L. Staib, "Shape-based 3D surface correspondence using geodesics and local geometry," vol. 2, pp. 644–651, 2000.
12. M. Kaus, V. Pekar, C. Lorenz, R. Truyen, S. Lobregt, J. Richolt, and J. Weese, "Automated 3D PDM construction using deformable models," in *Proceedings of the Eighth International Conference on Computer Vision (ICCV'01)*, pp. 566–572, IEEE Computer Society Press, 2001.
13. A. Frangi, D. Rueckert, J. Schnabel, and W. Niessen, "Automatic construction of multiple-object three-dimensional statistical shape models: Application to cardiac modeling," *IEEE Transactions on Medical Imaging*, vol. 21, no. 9, pp. 1151–1166, 2002.

14. R. Davies, C. Twining, T. Cootes, J. Waterton, and C. Taylor, "3D statistical shape models using direct optimisation of description length," 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. 3–20, Springer, 2002.
15. 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.
16. S. Mitchell, J. Bosch, B. Lelieveldt, R. van der Geest, J. Reiber, and M. Sonka, "3-D active appearance models: Segmentation of cardiac MR and ultrasound images," *IEEE Transactions on Medical Imaging*, vol. 21, no. 9, pp. 1167–1178, 2002.
17. 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.
18. A. Pentland and S. Sclaroff, "Closed-form solutions for physically based shape modeling and recognition," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 13, no. 7, pp. 715–729, 1991.
19. J. Martin, A. Pentland, and R. Kikinis, "Shape analysis of brain structure using physical and experimental modes," in *CVPR'94 (IEEE Conference on Computer Vision and Pattern Recognition)*, Computer Society Press, 1994.
20. 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.
21. T. Cootes and C. Taylor, "Combining point distribution models with shape models based on finite element analysis," *Image and Vision Computing*, vol. 13, no. 5, pp. 403–409, 1995.
22. T. Cootes and C. Taylor, "Data driven refinement of active shape model search," in *Proceedings of the British Machine Vision Conference* (R. Fisher and E. Trucco, eds.), pp. 383–392, The British Machine Vision Association, 1996.
23. Y. Wang and L. Staib, "Statistical shape and smoothness models for boundary finding with correspondence," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 22, no. 7, pp. 738–743, 2000.
24. L. Staib, "Prior shape models for boundary finding," in *Proceedings of the First IEEE International Symposium on Biomedical Imaging: Macro to Nano (ISBI'02)*, pp. 30–33, IEEE Signal Processing Society, 2002.
25. C. Nastar and N. Ayache, "Frequency-based nonrigid motion analysis: Application to four dimensional medical images," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 18, no. 11, pp. 1067–1079, 1996.
26. A. Kotcheff and C. Taylor, "Automatic construction of eigenshape models by direct optimization," *Medical Image Analysis*, vol. 2, pp. 303–314, 2002.
27. R. Davies, C. Twining, T. Cootes, J. Waterton, and C. Taylor, "A minimum description length approach to statistical shape modelling," *IEEE Transactions on Medical Imaging*, vol. 21, no. 5, pp. 525–537, 2002.
28. C. Goodall, "Procrustes methods in the statistical analysis of shape," *Journal of the Royal Statistical Society B*, vol. 53, no. 2, pp. 285–339, 1991.
29. I. Dryden and K. Mardia, *Statistical Shape Analysis*. Wiley Series in Probability and Statistics, 1998.
30. T. Cootes and C. Taylor, "Statistical models of appearance for computer vision," tech. rep., University of Manchester, 2002.
31. J. Davis, *Circulant Matrices*. John Wiley & Sons, 1979.
32. L. Staib and J. Duncan, "Boundary finding with parametrically deformable models," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 14, no. 11, pp. 1061–1075, 1992.
33. L. Staib and J. Duncan, "Model based deformable surface finding for medical images," *IEEE Transactions on Medical Imaging*, vol. 15, no. 5, pp. 720–731, 1996.
34. 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.