Boundary Finding with Correspondence Using Statistical Shape Models

Yongmei Wang* and Lawrence H. Staib*+

Departments of Electrical Engineering* and Diagnostic Radiology+
Yale University, P.O. Box 208042, New Haven, CT 06520-8042
Email: wang@noodle.med.yale.edu, lawrence.staib@yale.edu

Abstract

We propose an approach for boundary finding where the correspondence of a subset of boundary points to a model is simultaneously determined. Global shape parameters derived from the statistical variation of object boundary points in a training set are used to model the object. A Bayesian formulation, based on this prior knowledge and the edge information of the input image, is employed to find the object boundary with its subset points in correspondence with boundaries in the training set or the mean boundary. We compared the use of a generic smoothness prior and a uniform independent prior with the training set prior in order to demonstrate the power of this statistical information. A number of experiments were performed on both synthetic and real medical images of the brain and heart to evaluate the approach, including the validation of the dependence of the method on image quality, different initialization and prior information.

1 Introduction

Locating the boundary of structures in an image is of great importance in a variety of image analysis and computer vision applications including robot vision, pattern recognition and biomedical image processing. Numerous boundary finding methods have been proposed [2, 3, 4, 5, 7, 10, 11]. However, these methods, except, to some extent, [3, 4], do not provide any notion of correspondence. Correspondence is a key step in a number of computer vision applications such as stereo disparity, object recognition, motion estimation and non-rigid registration. There has been much work on determining boundary correspondence using local shape features [12, 8]. In this paper, however, we are interested in both determining an object's boundary and simultaneously determining spatial correspondence between similar structures over different subjects. Prior shape information is quite helpful in delineating the object boundary and solving the non-rigid correspondence problem. The goal of our work is to use prior shape models with point based global shape parameters to find the object boundary and the spatial correspondence. We use a Bayesian objective function based on this model and image derived information.

One of the most generic and popular methods of detecting whole boundaries using deformable models is the active contour approach (snakes) of Kass *et al.* [7]. A snake is a continuously deformable curve used to locate features in an image controlled by internal smoothness forces and external image forces. There have been many refinements, however, the parameters are still free to take almost any smooth boundary with no constraints on the overall shape.

Staib and Duncan [10] used elliptic Fourier descriptors as model parameters to represent open and closed boundaries. The Fourier coefficients are used to bias toward a range of shapes about a mean by using a Gaussian distribution on the parameters as a prior probability. A Bayesian approach is then used to obtain the maximum *a posteriori* estimate of the boundary. Chakraborty *et al.* [2] extended this approach to incorporate region homogeneity. Fourier descriptors are somewhat limited because they are not suitable for describing some shapes, such as those with convolutions or corners. In addition, these methods are designed for boundary finding, without regard to correspondence, which is one of our goals here.

Cootes *et al.* [3] combined deformable shape descriptors with statistical modal analysis built from a training set of annotated images. Object shapes are represented by a subset of boundary points, and a correspondence is established between these points from the different images of the training set. The deformations are modeled using linear combinations of the eigenvectors of the variations from the mean shape, thus defining the characteristic pattern of a shape class

and allowing deformation reflecting the variations in the training set. A similar Modal Analysis scheme is proposed by Pentland and Sclaroff [9] which gives a set of linear deformations of the shape equivalent to the modes of vibration of the original shape. However, the modes are somewhat arbitrary and may not be representative of the real variations which occur in a class of shapes. In Cootes' image search algorithm, the model is adjusted by searching a region of the image around each model point for an improved displacement. These local deformations are transformed into adjustments to the pose and shape parameters of the point model. By projecting the shape onto the shape parameters and enforcing limits, global shape constraints are applied ensuring that the current shape remains similar to the training set. The estimation of the displacement, however, is determined by a search only in the normal direction toward the strongest image edge [3]. In some situations, the adjustment of the pose, scale and shape parameters can not accommodate the inaccuracies this causes resulting in a misadjusted boundary. Shape parameters are accepted unless they are beyond three standard deviations. A more continuous penalty is more appropriate. In addition, the influence of the training set on the results is always fixed, and it can not be adjusted according to the image quality.

In this work, we use statistical point models with shape and shape variation generated from sets of examples using principal component analysis of the covariance matrix. This idea has been used successfully for boundary determination [3, 11]. Also, the Bayesian formulation of the objective function is akin to the work of Staib and Duncan [10]. The primary contribution of this paper is that it extends these ideas to provide a more robust method for both boundary and correspondence finding. Moreover, in order to show the important role of the prior shape model, we also compare with two other kinds of generic prior information. Specifically, the use of a Bayesian framework allows us to adjust the weighting between the statistical prior knowledge and the image information based on the image quality and the reliability of the training set. During the optimization, we search the shape and pose parameters together, resulting in a better optimum. Continuous penalizing criteria based on a training set derived Gaussian distribution for the prior shape and pose parameters are used during the searching. In addition, the optimization algorithm is faster and more robust than the image searching of Cootes *et al.* [3].

2 Statistical Point Models

Suppose we want to derive a model to represent the shapes whose boundaries are shown in Figure 1(b). We can represent each example shape as a set of hand-labeled landmark points. The model is then based on the mean positions of the points on the aligned shapes and the main variation of the points from the mean.

2.1 Capturing the Statistics

The critical points on the boundary are usually easily identified features, such as high curvature points, sharp corners, etc. Equally spaced points are interpolated between the critical points (Figure 1(a)). The training set points are first aligned to minimize a weighted sum of squares of distances between equivalent points on different shapes. During the alignment, we calculate the standard deviations of the pose (scale, rotation, translation).

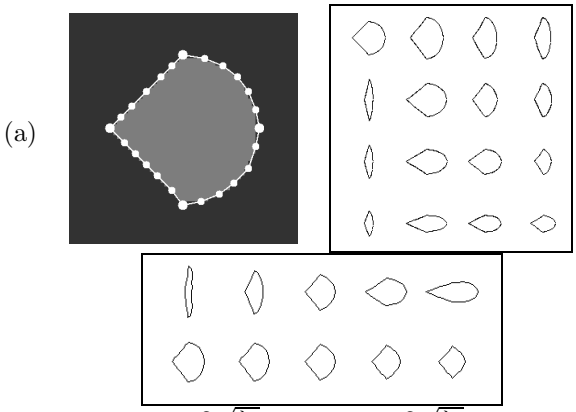

Figure 1: Synthetic shape model. (a): synthetic image (64×64) with its 24 point model of the boundary (4 critical points, (large dots), shown with interpolated points (small dots)); (b): 16 examples of synthetic shapes from a training set; (c): effects of varying each of the first two shape parameters of the synthetic model.

The following 2D formulation is similar to that of Cootes *et al.* [3]. Given m aligned examples and each example of a set of N aligned labeled points, $\mathbf{L}_i = (x_i(1), y_i(1), x_i(2), y_i(2), ..., x_i(N), y_i(N))^T$ $(i = 1, ..., m)$, we calculate the mean shape, $\bar{\mathbf{L}}$, and the covariance about the mean, $C_{training}$. It can be shown that by principal component analysis, the eigenvectors of the covariance matrix, $C_{training}$, corresponding to the largest eigenvalues describe the most significant modes of variation in the variables used to derive the covariance matrix, and that the proportion of the total variance explained by each eigenvector is equal to

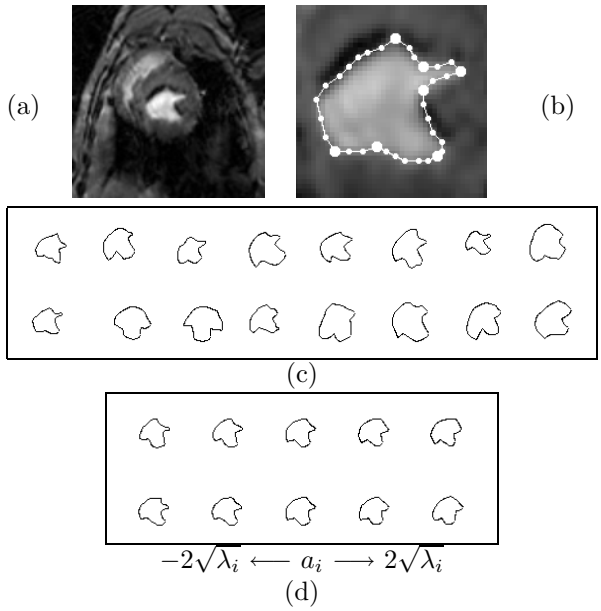

Figure 2: Heart shape model. (a): MR heart image (150 × 150); (b): 34 point model of the endocardium boundary (7 critical points, large dots) (cropped); (c): 16 examples of heart shapes from a training set; (d): effects of varying each of the first two shape parameters of the heart model.

the corresponding eigenvalue λ_k [6]. Most of the variation can be explained by a small number of modes, t ($< 2N$). Any shape in the training set can be approximated using the mean shape and a weighted sum of deviations obtained from the first t modes: $\mathbf{L} = \bar{\mathbf{L}} + Q\mathbf{a}$, where $Q = (\mathbf{q}_1 \mid \mathbf{q}_2 \mid ... \mid \mathbf{q}_t)$ is the matrix of the first t eigenvectors, and $\mathbf{a} = (a_1 a_2 ... a_t)^T$ is a vector of weights, which is also the set of t shape parameters to be optimized later. This equation allow us to generate new examples of the shapes by varying the parameter $\mathbf{a}$ within suitable limits, for example, $-3\sqrt{\lambda_k} \le a_k \le 3\sqrt{\lambda_k}$. We have used the techniques described above to generate statistical point models for both synthetic objects (Figure 1) and biological objects (Figure 2 for heart endocardium).

2.2 Identity Covariance Matrix

Consider the use of a $2N \times 2N$ identity covariance matrix $C_{identity}$ instead of the covariance derived from the training set. This means that all points (x and y coordinates) are independent. The eigenvectors of $C_{identity}$ $\mathbf{q}_k$ (size $2N$), for $k = 1, 2, \cdots, 2N$, are of the format

$$\begin{array}{rcl} \mathbf{q}_1 & = & (1,0,0,0,\cdots,0,0,0)^T \\ \mathbf{q}_2 & = & (0,1,0,0,\cdots,0,0,0)^T \\ \vdots & \vdots & \vdots \\ \mathbf{q}_{2N} & = & (0,0,0,0,\cdots,0,0,1)^T \end{array} \quad (1)$$

with the corresponding eigenvalues $\lambda_k = 1$. If k is even, the k'th eigenvector moves point $k/2$ in the y direction; if k is odd, the k'th eigenvector moves point $(k+1)/2$ in the x direction. Combinations of vectors, one for each mode, can move the modeled landmark points anywhere in the image. Any shape can also be approximated using the mean shape and a weighted sum of deviations obtained from the $2N$ modes. Since the eigenvalues here are very small, ten times the original range is used ($-20\sqrt{\lambda_k} \le a_k \le 20\sqrt{\lambda_k}$) to demonstrate the generation of new shapes, as shown in Figure 3.

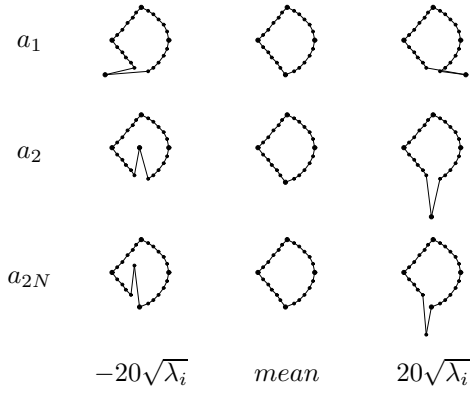

Figure 3: Effects of varying each of the 1st, the 2nd and the $2N$th shape parameters of the synthetic model with the identity covariance matrix.

2.3 Smoothness Covariance Matrix

Since the identity covariance matrix is too underconstrained, consider the incorporation of a type of smoothness constraint into the covariance matrix where neighboring points are correlated. That is

$$C_{smooth} = \begin{bmatrix} 1 & 0 & 0.5 & 0 & 0 & \cdots & 0.5 & 0 \\ 0 & 1 & 0 & 0.5 & 0 & \cdots & 0 & 0.5 \\ 0.5 & 0 & 1 & 0 & 0.5 & \cdots & 0 & 0 \\ 0 & 0.5 & 0 & 1 & 0 & \cdots & 0 & 0 \\ 0 & 0 & 0.5 & 0 & 1 & \cdots & 0 & 0 \\ \vdots & \vdots & \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ 0.5 & 0 & 0 & 0 & 0 & \cdots & 1 & 0 \\ 0 & 0.5 & 0 & 0 & 0 & \cdots & 0 & 1 \end{bmatrix} \quad (2)$$

where C_{smooth} is a $2N \times 2N$ matrix. Here, neighboring points are more likely to move together than if they were independent. The first two eigenvectors (with equal eigenvalues) allow for a rigid translation. Combinations of these two eigenvectors and the other eigenvectors, one for each mode, can move the modeled landmark points to anywhere in the image with the neighboring points moving together. The shapes generated (Figure 4) are smoother than those using the identity covariance matrix. Shapes are not restricted by a training set and thus the model is not specific but allows variability. Parameters corresponding to higher frequency variation have lower eigenvalues and thus have less influence. Also, the critical points no longer correspond. Note, the degree or scale of smoothing could be controlled by changing the coefficients along the diagonals. The motivation or goal of the use of this smoothness covariance matrix is actually similar to that of the internal smoothness forces in Kass's snakes [7]. However, it is expected that our image searching algorithm with this smoothness constraint would be more efficient than Kass's because we select the eigenvectors of C_{smooth} corresponding to the largest eigenvalues and thus reduce the dimensionality of the search space.

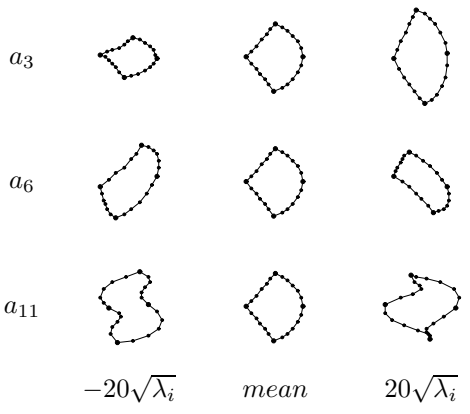

Figure 4: Effects of varying each of the 3rd, the 6th and 11th shape parameters of the synthetic model with the smoothness covariance matrix.

3 Bayesian Objective Function

Given the statistical models, our aim is to use them to model particular examples of structure in individual images, and then to find the shape parameters $\mathbf{a} = (a_1 a_2 ... a_t)^T$, and pose parameters: scale s, rotation θ, translation T_x, T_y. The combined pose and shape parameter vector to be determined is represented by $\mathbf{p} = (s, \theta, T_x, T_y, a_1, a_2, ..., a_t)^T$. The point representation of the nth boundary point $(n = 0, 1, ..., N-1)$ is

$$\begin{aligned} x(\mathbf{p}, n) &= s\cos\theta\{\bar{x}(n) + \textstyle\sum_{j=1}^{t} Q_{2n,j} a_j - \frac{S_x}{2}\} - \\ &\quad s\sin\theta\{\bar{y}(n) + \textstyle\sum_{j=1}^{t} Q_{2n+1,j} a_j - \frac{S_y}{2}\} + T_x + \frac{S_x}{2} \\ y(\mathbf{p}, n) &= s\sin\theta\{\bar{x}(n) + \textstyle\sum_{j=1}^{t} Q_{2n,j} a_j - \frac{S_x}{2}\} + \\ &\quad s\cos\theta\{\bar{y}(n) + \textstyle\sum_{j=1}^{t} Q_{2n+1,j} a_j - \frac{S_y}{2}\} + T_y + \frac{S_y}{2} \end{aligned} \quad (3)$$

where $\bar{x}(n)$ and $\bar{y}(n)$ are the mean shape of the nth point, and S_x, S_y are the image size in x, y directions respectively. In order to apply the prior knowledge of the shape model to the problem of boundary and correspondence determination, we pose the problem in a maximum *a posteriori* Bayesian formulation.

3.1 Prior Probability Density

Prior information can bias the boundary finder to search for a particular range of shapes and poses. We model the prior by using a zero mean multivariate Gaussian density $\Pr(\mathbf{p})$ for the shape and pose parameters (as in [10]). The variance for the shape parameters is the eigenvector's corresponding eigenvalue. For the pose parameters, the variance can be calculated from the training set alignment.

3.2 Likelihood

The likelihood is a measure of the similarity between the deformed template and the object present in the image. The likelihood we propose only uses the edge information in the input image, which is denoted as E here. The edge image E is assumed to consist of one of the deformed templates, $t_{\mathbf{p}}$, corrupted by additive white zero mean Gaussian noise with standard deviation σ_n, i.e. $E = t_{\mathbf{p}} + n$. This leads to (as in [10]):

$$\Pr(E \mid \mathbf{p}) = \prod_{\mathcal{A}} \frac{1}{\sqrt{2\pi}\sigma_n} e^{-\frac{(E(x,y) - t_{\mathbf{p}}(x,y))^2}{2\sigma_n^2}} \quad (4)$$

where $\mathcal{A}$ is the whole image area.

3.3 Posterior Probability Density

By using Bayes rule, the *a posteriori* probability density of the deformed template given the input edge image can be expressed as:

$$\Pr(\mathbf{p} \mid E) = \frac{\Pr(E \mid \mathbf{p}) \Pr(\mathbf{p})}{\Pr(E)} \quad (5)$$

Our objective is to maximize the *a posteriori* density in Eq.(5) with respect to $\mathbf{p}$. This can be simplified to maximize (as in [10]):

$$M(\mathbf{p}) = \sum_{i=1}^{t+4} \left[-\frac{(p_i - m_i)^2}{2\sigma_i^2} \right] + \frac{1}{\sigma_n^2} \sum_{n=1}^{N} E(x(\mathbf{p}, n), y(\mathbf{p}, n)) \quad (6)$$

where m_i is the mean of p_i and σ_i is the standard deviation for each of the parameters. This equation is the maximum *a posteriori* objective incorporating a prior bias to likely shapes and poses (first term) and match to the edges in the image (second term).

4 Experimental Results

We optimize the objective function $M(\mathbf{p})$ using the conjugate gradient method because we can efficiently compute the gradient from an analytic formulation.

For comparison, we also implemented the image searching of Cootes *et al.* [3]. On an SGI Indy 133MHZ, the convergence time of our boundary finding is about 5 seconds compared with 50 seconds for theirs for an average of 36 points, a speed up by a factor of ten.

4.1 Evaluation Criteria

The boundary error of each labeled boundary point on the final boundary is calculated by finding the distance to the closest point on the true boundary. We use both average, Eb_a, and maximum, Eb_m, error measures. The correspondence error of each critical point on the final boundary is the distance between this point and its corresponding critical point on the true boundary. The average error of the correspondence is denoted as Ec_a.

4.2 Synthetic Images

The image shown in Figure 5 is a simple synthetic image where the target object (the brightest) belongs to the training set family shown in Figure 1, but is not among that training set. The initial curve position is defined by the mean of the training set. The edge map of the input image is calculated by the Canny edge detector [1]. The final curve position accurately finds the target object. The following are four sets of experiments testing the effect of noise, initial parameters, prior probability bias and the form of prior.

The first experiment, shown in Figure 6, demonstrates the effect of noise on the method, by adding different amounts of zero mean Gaussian noise to the synthetic image shown in Figure 5 and measuring the boundary and correspondence error. Not all the Canny scales we chose here are optimal so that we can test our algorithm's tolerance to noise, to spurious and broken edges. The initial curve position is the mean of the training set for all testing. Signal-to-noise ratio (SNR) is defined here as the ratio of gray-level contrast between the target object and the background to the standard deviation of the Gaussian noise.

The second experiment, shown in Figure 7, examines the effect of the initial parameters. We average three results for each initial parameter using the synthetic image from Figure 5 with SNR of 5.0,

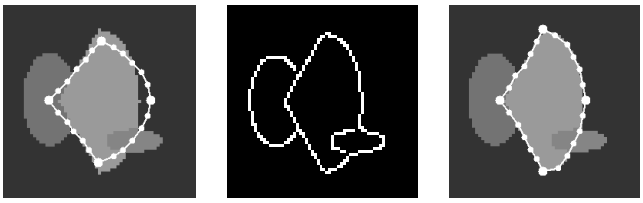

Figure 5: Synthetic image example. Left: initial contour on the synthetic image (64 × 64); Middle: Canny edge image (scale: 1.0); Right: final contour on target shape.

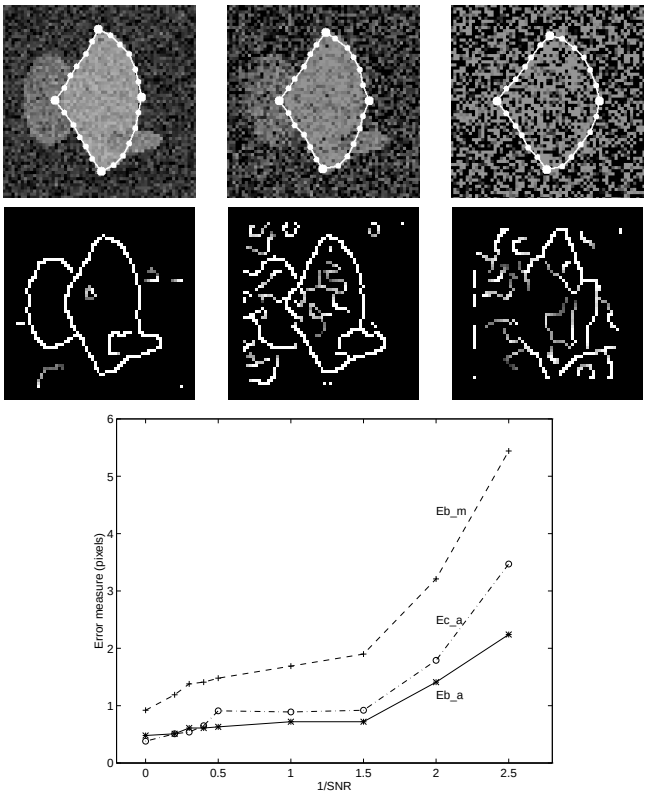

Figure 6: Sensitivity to noise experiment. Top, left to right: image from Fig. 5 with Gaussian noise added with SNR of 5.0, 2.5, 0.5. Each shown with final contour; Middle, left to right: corresponding Canny edge image with scale of 1.2, 1.4, 2.0; Bottom, error measures (Eb_a — boundary average error; Eb_m — boundary maximum error; Ec_a — correspondence average error).

2.5 and 1.0. The parameters individually varied were the first shape parameter, scale, horizontal translation and rotation, holding the other parameters constant. All results are good within reasonable ranges. Note, some error curves are not symmetric due to the non-symmetric test image. For all the initial parameters, when they are too far away from the true boundary,

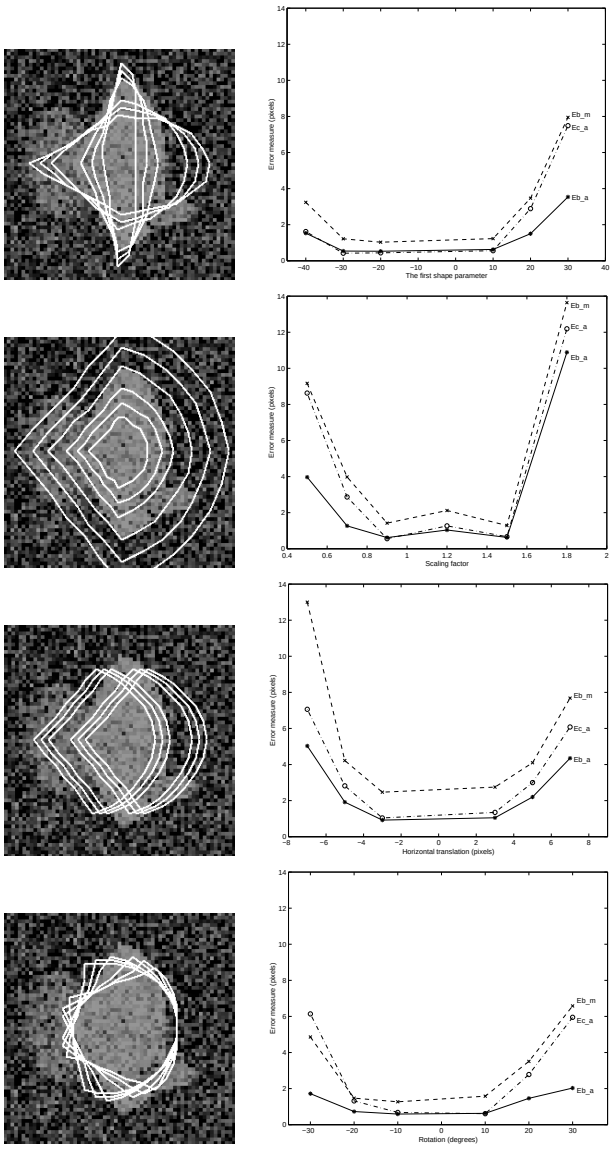

Figure 7: Sensitivity to initial parameters experiment. Left, top to bottom: image from Fig. 5 with Gaussian noise added (SNR = 2.5) shown with initial contours for range of the first shape parameter tested, range of scale tested, range of horizontal translation tested, range of rotation tested; Right: corresponding error measures respectively (Eb_a — boundary average error; Eb_m — boundary maximum error; Ec_a — correspondence average error).

the optimization may be trapped by local minima corresponding to nearby edges.

In the third experiment, we demonstrate the effect of different prior probability densities. The shape model for the training set is shown in the top of Figure 8. A synthetic image, shown in the bottom of Figure 8, was designed containing two objects. The light object corresponds to the shape with the first shape parameter at two standard deviations. The dark object underneath it is the shape with the second shape parameter at two standard deviations. Using densities with the same mean, but different variances, we can demonstrate different results that are completely due to the prior bias density. The prior can be biased towards finding the light object by having a wide distribution on the first shape parameter and narrow distribution on the second parameter and vice versa (the bottom of Figure 8). This would not be possible without the continuous bias of the prior term.

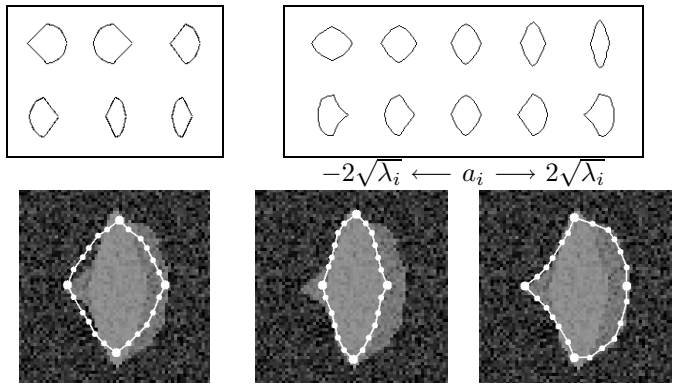

Figure 8: Bias experiment. Top left: 6 examples of synthetic shapes from a training set, each containing 24 points; Top right: effects of varying each of the first two shape parameters of the synthetic model; Bottom left: initial contour on the synthetic image (64 × 64); Bottom middle: final contour, biased to the brighter target shape; Bottom right: final contour, biased to the darker target shape.

In the fourth experiment, we compare different prior models to illustrate the appropriateness of our prior model. If there is no prior information, using an identity covariance matrix, each point on the boundary can move independently. The objective function only includes the likelihood term. The result will try to match to edges without regarding to shape or correspondence. With the smoothness covariance of Eq.(2), neighboring points on the boundary will tend to move together. The resulting boundary points will try to match edges while maintaining the smoothness, but the correspondence of the points may not be maintained. Here, we used the same synthetic image with SNR= 2.5 and the same Canny edge image at scale 1.4 (the middle of Figure 6). The results with different prior models are shown in Figure 9 and Table 1. The training set covariance $C_{training}$ works best,

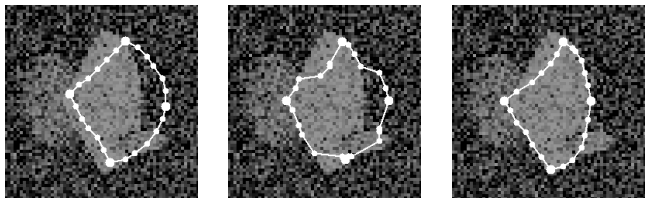

Figure 9: Prior experiment for synthetic image. Left: initial contour; Middle: final contour on target shape using identity covariance matrix $C_{identity}$; Right: final contour on target shape using smoothness covariance matrix C_{smooth}.

as expected, while the smoothness covariance may be suitable in situations where there is no training set and correspondence is not needed.

	$C_{identity}$	C_{smooth}	$C_{training}$
Eb_a	2.58	1.45	0.57
Eb_m	7.23	5.90	1.52
Ec_a	5.04	1.86	0.36

Table 1: Error measure for the synthetic image in the prior experiment with different covariance matrices: Eb_a — boundary average error; Eb_m — boundary maximum error; Ec_a — correspondence average error.

4.3 Real Images

Results of the method applied to sagittal and axial magnetic resonance (MR) images of the human brain are shown in Figures 10 and 11 respectively. For the sagittal corpus callosum image, we used a 49 point model derived from a set of 12 corpus callosum shapes. For the axial brain image, a 93 point model derived from a set of 12 basal ganglia and ventricle boundaries is used. Not only are the final contours delineated successfully, but also the correspondence of the points is established accurately. Figure 12 shows the method applied to the endocardium in an MR image of a dog heart using the heart shape model in Figure 2. Both the boundary and correspondences are found correctly.

We also compared the three different prior models for the dog heart image (Figure 12) with the same initial position as before. With the identity covariance, the boundary cannot be found since the shape is complicated and the points move independently. With the smoothness covariance Eq.(2), although the resulting boundary is similar to the true boundary, the correspondence of the points is not determined. The results in Figure 13 and Table 2 show quantitative comparison with an expert drawn endocardium. For the heart

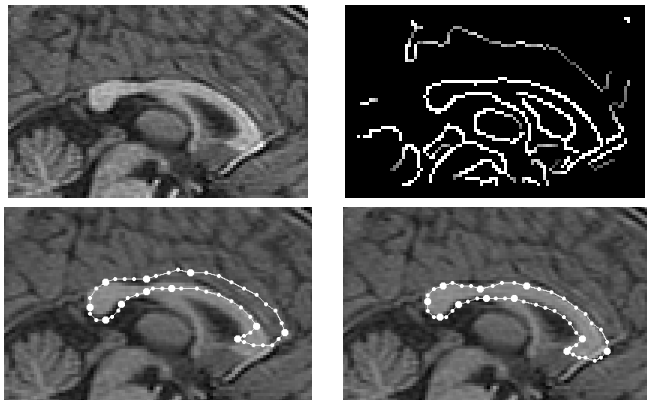

Figure 10: MR sagittal brain example. Top left: original MR image (100×64); Top right: Canny edge image (scale: 1.2); Bottom left: initial contour (mean curve); Bottom right: final contour on corpus callosum.

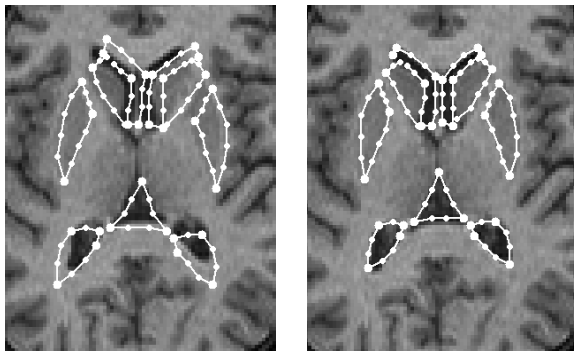

Figure 11: MR axial brain example. Left: initial contour (mean curve) on the original image (80×100); Right: final contour on basal ganglia and ventricle boundaries.

image here, a shape model is a necessity for finding the endocardium and its critical points.

5 Conclusions and Future Work

This work presents a systematic approach to determine an object's boundary, as well as the correspondence of boundary points to a model. The statistical point models derived from a training set by principal component analysis are used as the prior probability in a Bayesian scheme, capturing prior knowledge of the shape. The structure is delineated and the spatial correspondence of these points to the model is established when the *a posteriori* probability is maximized using conjugate gradient optimization. From experimental results, it was found that this method performs well at both boundary finding and determining correspondence, and it is also relatively insensitive to noise and initialization. The prior models testing

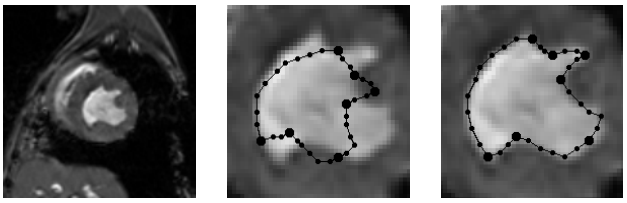

Figure 12: MR heart image example. Left: original MR image (150 × 150); Middle: initial contour (mean curve) on endocardium (cropped); Right: final contour.

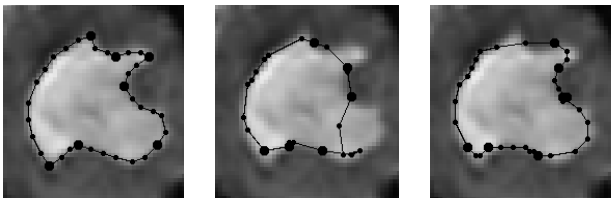

Figure 13: Prior experiment for MR heart image (cropped). Left: contour on endocardium drawn by an expert; Middle: final contour using identity covariance matrix $C_{identity}$; Right: final contour using smoothness covariance matrix C_{smooth}.

further demonstrated that the statistical shape models used are crucial for both boundary and correspondence finding.

	$C_{identity}$	C_{smooth}	$C_{training}$
Eb_m	6.12	1.79	1.36
Ec_a	5.57	7.19	0.89

Table 2: Error measure for the heart image in prior experiment with different covariance matrices: Eb_m — boundary maximum error; Ec_a — correspondence average error.

These methods generalize directly to 3D. Other future directions include testing the smooth covariance matrix C_{smooth} at different smoothing scales, combining the training set approach with the smoothness covariance matrix when few training set examples are available, and using the resulting boundary points as landmarks for non-rigid registration.

Acknowledgments

This work was supported in part by a grant from the Whitaker Foundation. We would like to thank the reviewers for their valuable comments.

References

[1] J. F. Canny, "A computational approach to edge-detection," *IEEE Trans. Pattern Analysis and Machine Intelligence*, Vol. 8, no. 6, pp. 679-698, 1986.

[2] A. Chakraborty, L. H. Staib and J. S. Duncan, "An integrated approach to boundary finding in medical images," *IEEE Workshop Biomedical Image Analysis*, pp. 13-22, 1994.

[3] T. F. Cootes, C. J. Taylor, D. H. Cooper and J. Graham, "Active shape models – their training and application," *Computer Vision and Image Understanding*, Vol. 61, no. 1, pp. 38-59, 1995.

[4] N. Duta and M. Sonka, "Segmentation and interpretation of MR brain images using an improved knowledge-based active shape models," *Information Processing in Medical Imaging*, pp. 375-380, Spring-Verlag Berlin Heidelberg New York, 1997.

[5] A. K. Jain, Y. Zhong, S. Lakshmanan, "Object matching using deformable templates," *IEEE Trans. Pattern Analysis and Machine Intelligence*, Vol. 18, no. 3, pp. 267-278, 1996.

[6] R. A. Johnson and D. W. Wichern, *Multivariate statistics, a practical approach,* Chapman & Hall, 1988.

[7] M. Kass, A. Witkin and D. Terzopoulos, "Snakes: Active contour models," *Int. Journal of Computer Vision*, Vol. 1, no. 4, pp. 312-331, 1988.

[8] J. Mceachen, A. Nehorai and J. Duncan, "A recursive filter for temporal analysis of cardiac motion," *IEEE Workshop on Biomedical Image Analysis*, pp. 124-133, 1994.

[9] A. P. Pentland and S. Sclaroff, "Closed-form solutions for physically based shape modeling and recognition," *IEEE Trans. Pattern Analysis and Machine Intelligence*, Vol. 13, no. 7, pp. 715-729, 1991.

[10] L. H. Staib and J. S. Duncan, "Boundary finding with parametrically deformable models," *IEEE Trans. Pattern Analysis and Machine Intelligence*, Vol. 14, no. 11, pp. 1061-1075, 1992.

[11] G. Szekely, A. Keleman, C. Brechbuler and G. Gerig, "Segmentation of 2-D and 3-D objects from MRI volume data using constrained elastic deformations of flexible Fourier contours and surface models," *Medical Image Analysis*, Vol. 1, no. 1, pp. 19-34, 1996.

[12] H. Tagare, D. O'Shea and A. Rangarajan, "A geometric criterion for shape-based non-rigid correspondence," *Proc. Fifth Int. Conf. Comp. Vision*, pp. 434-439, 1995.