

Statistical Shape Influence in Geodesic Active Contours

Michael E. Leventon¹ W. Eric L. Grimson¹ Olivier Faugeras^{1,2}

¹ MIT AI Lab
Cambridge, MA 02139
leventon, welg@ai.mit.edu

² INRIA
Sophia Antipolis, France
Olivier.Faugeras@sophia.inria.fr

Abstract

A novel method of incorporating shape information into the image segmentation process is presented. We introduce a representation for deformable shapes and define a probability distribution over the variances of a set of training shapes. The segmentation process embeds an initial curve as the zero level set of a higher dimensional surface, and evolves the surface such that the zero level set converges on the boundary of the object to be segmented. At each step of the surface evolution, we estimate the maximum a posteriori (MAP) position and shape of the object in the image, based on the prior shape information and the image information. We then evolve the surface globally, towards the MAP estimate, and locally, based on image gradients and curvature. Results are demonstrated on synthetic data and medical imagery, in 2D and 3D.

1 Introduction

The anatomical structures that appear in magnetic resonance (MR) or computed tomography (CT) scans are often explicitly extracted or segmented from the image for use in surgical planning, navigation, simulation, diagnosis, and therapy evaluation. By segmentation, we refer to the process of labeling individual voxels in the volumetric scan by tissue type, based on properties of the observed intensities as well as anatomical knowledge about normal subjects.

Segmentation is often performed using automated techniques and semi-automated techniques. With CT data, segmentation of some structures can be performed just using intensity thresholding or other low-level image processing. In general, however, segmentation is challenging and requires more sophisticated algorithms and significant human input. For example, the distribution of intensity values corresponding to one structure may overlap those of another structure, defeating intensity-based segmentation techniques.




Figure 1. Level sets of an embedding function u for a closed curve C in \mathbb{R}^2 .

Boundary finding segmentation methods such as Snakes [8], are generally local algorithms that require some feature (such as an edge) to be present along the boundary of the object, and gravitate toward that feature. These methods may be sensitive to the starting position and may “leak” through the boundary of the object if the edge feature is not salient enough in a certain region in the image.

Level set based segmentation, introduced in [1, 9], involves solving the energy-based active contours minimization problem by the computation of geodesics or minimal distance curves [11, 13]. In this approach, a curve is embedded as a zero level set of a higher dimensional surface (Figure 1). The entire surface is evolved to minimize a metric defined by the curvature and image gradient. Recent work based on level sets include extensions such as codimension-two regularization [10], texture models [12], global intensity statistics [17], and pedal curve evolution [5].

When segmenting or localizing an anatomical structure, having prior information about the expected shape of that structure can significantly aid in the segmentation process. Both Cootes, *et al.* [4] and Wang and Staib [16] find a set of corresponding points across a set of training images and construct a statistical model of shape variation that is then used in the localization of the boundary. Staib and Duncan [14] incorporate global shape information into the segmentation task by using an elliptic Fourier decomposition of the boundary and placing a Gaussian prior on the Fourier




Figure 2. Corpus callosum outlines for 6 out of 51 patients in the training set embedded as the zero level set of a higher dimensional signed distance surface.

coefficients.

Our approach to object segmentation extends geodesic active contours [1, 9] by incorporating shape information into the evolution process. We first compute a statistical shape model over a training set of curves. To segment a structure from an image, we evolve an active contour both locally, based on image gradients and curvature, and globally to a maximum *a posteriori* estimate of shape and pose.

2 Probability distribution on shapes

To incorporate shape information into the process of segmenting an object in an image, we consider a probabilistic approach, and compute a prior on shape variation given a set of training instances. To build the shape model, we choose a representation of curves, and then define a probability density function over the parameters of the representation.

2.1 Curve representation

Each curve in the training dataset is embedded as the zero level set of a higher dimensional surface, u , whose height is sampled at regular intervals (say N^d samples, where d is the number of dimensions). The embedding function chosen is the commonly used signed distance function [13], where each sample encodes the distance to the nearest point on the curve, with negative values inside the curve. Each such surface (distance map) can be considered a point in a high dimensional space ($u \in \mathbb{R}^{N^d}$). The training set, \mathcal{T} , consists of a set of surfaces $\mathcal{T} = \{u_1, u_2, \dots, u_n\}$. Our goal is to build a shape model over this distribution of surfaces. Since a signed distance map is uniquely determined from the zero level set, each distance map has a large amount of redundancy. Furthermore, the collection of curves in the training set presumably has some dependence, as they are shapes of the same class of object, introducing

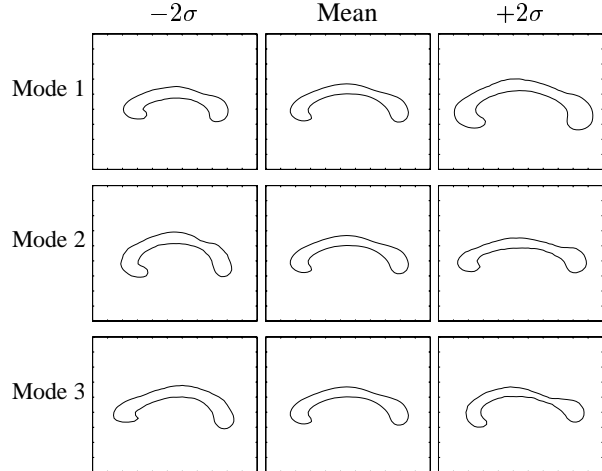


Figure 3. The three primary modes of variance of the corpus callosum training dataset.

more redundancy in the training set. The cloud of points corresponding to the training set is approximated to have a Gaussian distribution, where most of the dimensions of the Gaussian collapse, leaving the principal modes of shape variation.

The mean surface, μ , is computed by taking the mean of the signed distance functions, $\mu = \frac{1}{n} \sum u_i$. The variance in shape is computed using Principal Component Analysis (PCA). The mean shape, μ , is subtracted from each u_i to create an mean-offset map, \hat{u}_i . Each such map, \hat{u}_i , is placed as a column vector in an $N^d \times n$ -dimensional matrix M . Using Singular Value Decomposition (SVD), the covariance matrix $\frac{1}{n} M M^T$ is decomposed as:

$$U \Sigma U^T = \frac{1}{n} M M^T \quad (1)$$

where U is a matrix whose column vectors represent the set of orthogonal modes of shape variation and Σ is a diagonal matrix of corresponding singular values. An estimate of a novel shape, u , of the same class of object can be represented by k principal components in a k -dimensional vector of coefficients, α .

$$\alpha = U_k^T (u - \mu) \quad (2)$$

where U_k is a matrix consisting of the first k columns of U that is used to project a surface into the eigen-space. Given the coefficients α , an estimate of the shape u , namely \tilde{u} , is reconstructed from U_k and μ .

$$\tilde{u} = U_k \alpha + \mu \quad (3)$$

Note that in general \tilde{u} will not be a true distance function, since convex linear combinations of distance maps do not produce distance maps. However, the surfaces generally still have advantageous properties of smoothness, local dependence, and zero level sets consistent with the combination of original curves.




Figure 4. Top: Three-dimensional models of seven thoracic vertebrae (T3-T9) used as training data. Bottom left and right: Extracted zero level set of first and second largest mode of variation respectively.

Under the assumption of a Gaussian distribution of shape represented by α , we can compute the probability of a certain curve as:

$$P(\alpha) = \frac{1}{\sqrt{(2\pi)^k |\Sigma_k|}} \exp\left(-\frac{1}{2} \alpha^T \Sigma_k^{-1} \alpha\right) \quad (4)$$

where Σ_k contains the first k rows and columns of Σ . We also have considered modeling the distribution of shapes as a mixture of Gaussians or using a Parzen window density estimator, but keep to a Gaussian prior in this work.

Figure 2 shows a few of the 51 training curves used to define the shape models of the corpus callosum. The original segmentations of the images are shown as white curves. The outlines are overlaid on the signed-distance map. Before computing and combining the distance maps of these training shapes, the curves were aligned using centroids and second moments to approximate the correspondence. Figure 3 illustrates zero level sets corresponding to the means and three primary modes of variance of the shape distribution of the corpus callosum. Figure 4 shows the zero level set (as a triangle surface model) of seven rigidly aligned vertebrae of one patient used as training data. The zero level sets of the two primary modes are shown as well. Note that for both the corpus and the vertebrae, the mean shapes and primary modes appear to be reasonable representative shapes of the classes of objects being learned. In the case of the corpus callosum, the first mode seems to capture size, while the second mode roughly captures the degree of curvature of the corpus. The third mode appears to represent the shifting of the bulk of the corpus from front to back.

2.2 The correspondence problem

When measuring shape variance of a certain part of an object across a population, it is important to compare like parts of the object. For example, when looking at variances

in the shape of the vertebrae, if two training examples are misaligned and a process of one is overlapping a notch of the other, then the model will not be capturing the appropriate anatomical shape variance seen across vertebrae.

One solution to the correspondence problem is to explicitly generate all point-wise correspondences to ensure that comparisons are done consistently. Finding correspondences in general is a difficult problem. Manually localizing corresponding landmarks is tedious, while automatic landmark detection is prone to errors, especially when dealing with 3D objects. In [4], Cootes requires the labeling of 123 corresponding landmarks on each of 72 training instances when building a 2D model of the region of the brain around the ventricles. While this method does an excellent job of performing the model-based matching, in many applications, especially in 3D, careful labeling of such a large training set is infeasible.

Another approach to correspondence is to roughly align the training data before performing the comparison and variance calculation. A rough alignment will not match every part of each training instance perfectly, so one must consider the robustness of the representation to misalignment. Turk and Pentland [15] introduced Eigenfaces as a method of building models for recognition. Each image in a set of face images ($N \times N$ array of intensities) is considered as a point in an N^2 -dimensional space, from which the principal components are computed. The Eigenface method is similar to our method of combining signed distance maps of binary images, with an important distinction. Any slight misalignment in the faces compares the intensity variance between independent objects, while slightly misaligned pixels in a distance map are generally very highly correlated. Smoothing a grayscale or binary image propagates information spatially as well, increasing the correlation between neighboring pixels, but results in loss of information, whereas no information about the binary image is lost by computing its




Figure 5. (a) The curve expected to be found in the image. (b) The image containing the shape to be segmented. (c) The same solid curve appears in three different locations, and the direction of evolution depends on the position of the evolving curve with respect to the object. The dotted lines show a later step in evolution given the curve’s position and shape.

signed distance function.

Using the signed distance map as the representation of shape provides tolerance to slight misalignment of object features, in the attempt to avoid having to solve the general correspondence problem. In the examples presented here, the rough rigid alignment of the training instances resulted in the model capturing the shape variances inherent in the population due to the dependence of nearby pixels in the shape representation. We are exploring the tradeoffs between leaving certain transformation parameters intrinsic in the representation, or extrinsic by aligning the training set under those classes of transformation. Currently, rigid pose is extrinsic, but scale and affine warp are left as intrinsic. We are also extending this method to use correspondences (when they are available) to ensure the comparison of “like” pixels instead of “near” pixels using flow fields as in [7].

3 Shape priors and geodesic active contours

Given a curve representation (the k -dimensional vector α) and a probability distribution on α , the prior shape information can be folded into the segmentation process. This section describes adding a term to the level set evolution equation to pull the surface in the direction of the maximum *a posteriori* shape and position of the final segmentation.

3.1 Geodesic active contours for segmentation

The snake methodology defines an energy function $E(\mathcal{C})$ over a curve \mathcal{C} as the sum of an internal and external energy of the curve, and evolves the curve to minimize the energy [8].

$$E(\mathcal{C}) = \beta \int |\mathcal{C}'(q)|^2 dq - \lambda \int |\nabla I(\mathcal{C}(q))| dq \quad (5)$$




Figure 6. Three steps in the evolution process. The evolving curve is shown solid, superimposed on the image (top row). The curve is matched to the expected curve to obtain a PDF over pose (bottom row). The next evolution step (based on pose and shape) is shown as the dotted line.

In [1], Caselles, *et al.* derive the equivalence of geodesic active contours to the traditional energy-based active contours (snakes) framework by first reducing the minimization problem to the following form:

$$\min_{\mathcal{C}(q)} \int g(|\nabla I(\mathcal{C}(q))|) |\mathcal{C}'(q)| dq \quad (6)$$

where g is a function of the image gradient (usually of the form $\frac{1}{1+|\nabla I|^2}$). Using Euler-Lagrange, the following curve evolution equation is derived [1]

$$\frac{\partial \mathcal{C}(t)}{\partial t} = g\kappa\mathcal{N} - (\nabla g \cdot \mathcal{N})\mathcal{N} \quad (7)$$

where κ is the curvature and \mathcal{N} is the unit normal. By defining an embedding function u of the curve \mathcal{C} , the update equation for a higher dimensional surface is computed.

$$\frac{\partial u}{\partial t} = g(c + \kappa) |\nabla u| + \nabla u \cdot \nabla g \quad (8)$$

where c is an image-dependent balloon force added to force the contour to flow outward [1, 3]. In this level set framework, the surface, u , evolves at every point perpendicular to the level sets as a function of the curvature at that point and the image gradient.

3.2 Estimation of pose and shape

In addition to evolving the level set based on the curvature and the image term, we include a term that incorporates information about the shape of the object being segmented. To add such a global shape force to the evolution, the pose of the evolving curve with respect to the shape model must




Figure 7. The relationship between the distance map and the image gradient.

be known (see Figures 5 and 6). Without an estimate of the pose, the shape model cannot adequately constrain or direct the evolution. Therefore, at each step of the curve evolution, we seek to estimate the shape parameters, α , and the rigid pose parameters, p , of the *final* curve using a maximum *a posteriori* approach.

$$\langle \alpha_{\text{MAP}}, p_{\text{MAP}} \rangle = \underset{\alpha, p}{\operatorname{argmax}} P(\alpha, p \mid u, \nabla I) \quad (9)$$

In this equation, u is the evolving surface at some point in time, whose zero level set is the curve that is segmenting the object. The term ∇I is the gradient of the image containing the object to be segmented. By our definition of shape and pose, the final segmentation curve is completely determined by α and p . Let u^* be the estimated final curve, which can be computed from α and p . Therefore, we also have

$$u_{\text{MAP}}^* = \underset{u^*}{\operatorname{argmax}} P(u^* \mid u, \nabla I) \quad (10)$$

To compute the maximum *a posteriori* final curve, we expand the terms from Eq. 9 using Bayes' Rule.

$$\begin{aligned} P(\alpha, p \mid u, \nabla I) &= \frac{P(u, \nabla I \mid \alpha, p) P(\alpha, p)}{P(u, \nabla I)} \quad (11) \\ &= \frac{P(u \mid \alpha, p) P(\nabla I \mid \alpha, p, u) P(\alpha) P(p)}{P(u, \nabla I)} \end{aligned}$$

Note that the preceding step assumes that shape is independent from pose. Since our current model does not attempt to capture the relationships between these two quantities, this is reasonable. Future work may incorporate positional priors (and relative positional priors between objects) into our shape model. We proceed by defining each term of Eq. 11 in turn. We discard the normalization term in the denominator as it does not depend on shape or pose.

Inside Term. The first term in Eq. 11 computes the probability of a certain evolving curve, u , given the shape and pose of the final curve, u^* (or $\langle \alpha, p \rangle$). Notice that this term does not include any image information whatsoever. Given our method of initializing the curve with a point *inside* the object, it is reasonable to assume that the curve should remain inside the object throughout the evolution. Therefore,




Figure 8. Illustration of the various terms in the evolution of the surface, u . The surface u^* , is the maximum *a posteriori* final shape. To update u , we combine the standard gradient and curvature update term, v , and the direction of the MAP final shape, $u^* - u$.

if the evolving curve lies completely inside the final curve, then it is more likely than a curve that lies partially or fully outside the final curve. We model this term as a Laplacian density function over V_{outside} , the volume of the curve u that lies outside the curve u^* .

$$P(u \mid \alpha, p) = \exp(-V_{\text{outside}}) \quad (12)$$

This term assumes that any curve u lying inside u^* is equally likely. Since the initial curve can be located at any point inside the object and the curve can evolve along any path, we do not favor any such curve.

Gradient Term. The second term in Eq. 11 computes the probability of seeing certain image gradients given the current and final curves. Consider the relationship between u^* and $|\nabla I|$ when u^* correctly outlines the boundary of the object (see Figure 7). Notice that the distance map u^* is linear along the normal direction of the curve at any boundary point, q , and $u^*(q) = 0$. Furthermore, under the assumption that the object boundary is a smoothed step edge, $|\nabla I|$ approximates a Gaussian along the normal at q . Therefore, we'd expect the relationship between $|\nabla I|$ and u^* to be Gaussian in nature. Figure 7 shows an example scatter plot of these quantities when u^* is aligned with the object boundary. Let $h(u^*)$ be the best fit Gaussian to the samples $(u^*, |\nabla I|)$. We model the gradient probability term as a Laplacian of the goodness of fit of the Gaussian.

$$P(\nabla I \mid u^*, u) = \exp(-|h(u^*) - |\nabla I||^2) \quad (13)$$

Shape and pose priors. The last two terms in Eq. 11 are based on our prior models, as described in Section 2. Our shape prior is a Gaussian model over the shape parameters, α , with shape variance Σ_k .

$$P(\alpha) = \frac{1}{\sqrt{(2\pi)^k |\Sigma_k|}} \exp\left(-\frac{1}{2} \alpha^T \Sigma_k^{-1} \alpha\right) \quad (14)$$

In our current framework, we seek to segment one object from an image, and do not retain prior information on the likelihood of the object appearing in a certain location. Thus, we simply assume a uniform distribution over pose parameters, which can include any type of transformation, depending on application.

$$P(p) = \mathcal{U}(-\infty, \infty) \quad (15)$$

Currently we are modeling translation and rotation. We feel, however, that positional priors could provide a rich source of information to explore in the future, especially when segmenting multiple objects from a single image that may have clear prior relative poses, or when a distribution over pose in a fixed image-based coordinate system is known.

These terms define the maximum *a posteriori* estimator of shape and pose, which estimates the final curve or segmentation of the object. For efficiency, these quantities are computed only in a narrow band around the zero level set of the evolving surface, and the MAP pose and shape are re-estimated at each evolution step using simple gradient ascent on the log probability function in Eq. 11. While each ascent may yield a local maximum, the continuous re-estimation of these parameters as the surface evolves generally results in convergence on the desired maximum. Next, we incorporate this information into the update equation commonly used in level set segmentation.

3.3 Evolving the surface

Initially, the surface, u , is assumed to be defined by at least one point that lies inside the object to be segmented. Given the surface at time t , we seek to compute an evolution step that brings the curve closer to the correct final segmentation based on local gradient and global shape information.

The level set update expression shown in Eq. 8 provides a means of evolving the surface u over time towards the solution to the original curve-minimization problem stated in Eq. 6. Therefore, the shape of the surface at time $t+1$ can be computed from $u(t)$ by:

$$u(t+1) = u(t) + \lambda_1 (g(c + \kappa) |\nabla u(t)| + \nabla u(t) \cdot \nabla g) \quad (16)$$

where λ_1 is a parameter defining the update step size.

By estimating the final surface u^* at a given time t , (Section 3.2), we can also evolve the surface in the direction of the maximum *a posteriori* final surface:

$$u(t+1) = u(t) + \lambda_2 (u^*(t) - u(t)) \quad (17)$$

where $\lambda_2 \in [0, 1]$ is the linear coefficient that determines how much to trust the maximum *a posteriori* estimate. Combining these equations yields the final expression for computing the surface at the next step.

$$u(t+1) = u(t) + \lambda_1 (g(c + \kappa) |\nabla u(t)| + \nabla u(t) \cdot \nabla g) + \lambda_2 (u^*(t) - u(t)) \quad (18)$$

Figure 8 illustrates this evolution. The two parameters λ_1 and λ_2 are used to balance the influence of the shape model and the gradient-curvature model. The parameters also determine the overall step size of the evolution. The tradeoff between shape and image depends on how much faith one has in the shape model and the imagery for a given application. Currently, we set these parameters empirically for a particular segmentation task, given the general image quality and shape properties.

The original evolution equation (Eq. 16), to which we added the shape influence term, was derived from an energy minimization expression (Eq. 6). We are currently exploring ways of adding a “shape energy” term to the curve integral in Eq. 6, and then deriving the complete evolution equation using Euler-Lagrange, instead of adding the shape force after the fact. In this new framework, both processes will be fused into a single energy minimization. One difficulty is that the curve integral is inherently local and does not require any notion of correspondence, whereas global shape information involves the comparison of a shape to a model, in correspondence.

4 Results

We have tested the segmentation algorithm on synthetic and real shapes, both in 2D and in 3D. For controlled testing, a training set of rhombi of various sizes and aspect ratios was generated to define a shape model. Test images were constructed by embedding the shapes of two random rhombi with the addition of Gaussian speckle noise and a low frequency diagonal bias field. Figure 9 illustrates several steps in the segmentation of the synthetic objects. In the first frame of each trial, the small black circle represents the initialization point. The white curve illustrates the MAP shape and pose at each time step. The final segmentations are shown in the last frames.

Segmentation experiments were performed on 2D slices of MR images of the femur and corpus callosum (Figures 10 and 11). For the femur experiments, the training set consisted of 18 nearby slices of the same femur, leaving out the slice being segmented and its neighbors. In both femur examples, the same initialization point was used to seed the evolution. As the curve evolves, the MAP estimator of shape and pose locks into the shape of the femur slice.

The corpus callosum training set consisted of 49 examples like those in Figure 2. The segmentations of two corpora callosa are shown in Figure 11. Notice that while the MAP shape estimator is initially incorrect, as the curve evolves, the pose and shape parameters converge on the boundary. The segmentations of the femur slices and corpora converged in under a minute on a 550 MHz Pentium.

The vertebrae example illustrates the extension of the algorithm to 3D datasets. Figure 12 illustrates a few steps

K	Corpus 1	Corpus 2	Vertebra
95%	1.3 mm	1.5 mm	2.7 mm
99%	1.6 mm	2.0 mm	4.4 mm

Table 1. Partial Hausdorff distance between our segmentation and the manually-segmented ground truth.

in the segmentation of vertebra T7. The training set in this case consisted of vertebrae T3-T9, with the exception of T7. The initial surface was a small sphere placed in the body of the vertebra. The black contour is a slice through the zero level set of the evolving hyper-surface. The white contour is the MAP pose and shape estimate. Segmenting the vertebra took approximately six minutes.

To validate the segmentation results, we compute the undirected partial Hausdorff distance [6] between the boundary of the computed segmentation and the boundary of the manually-segmented ground truth. The directed partial Hausdorff distance over two point sets A and B is defined as

$$h_K(A, B) = K^{\text{th}} \min_{a \in A, b \in B} \|a - b\|$$

where K is a quantile of the maximum distance. The undirected partial Hausdorff distance is defined as $H_K(A, B) = \max(h_K(A, B), h_K(B, A))$. The results for the corpora and vertebrae shown in Table 1 indicate that virtually all the boundary points lie within one or two voxels of the manual segmentation.

5 Conclusions

This work presents a means of incorporating prior shape information into the geodesic active contour method of medical image segmentation. The shape representation of using PCA on the signed distance map was chosen with the intention of being robust to slight misalignments without requiring exact point-wise correspondences. Our extension to active contours that estimates the model parameters and then evolves the curve based on that estimation provides a method of robustly converging on the boundary even with noisy inputs. The representation and the curve evolution technique merge well together since the evolution requires a distance map of the evolving curve, which is inherent in the shape model. However, these two modules need not be coupled. A different statistical shape model could be tied into the evolution method, or a different method of model-based matching could be used with the proposed shape model.

Acknowledgments

This work was supported by NSF Contract IIS-9610249, NSF Contract DMS-9872228, and NSF ERC (Johns Hopkins University agreement) 8810-274. The MRI data was provided by Martha Shenton, Jane Anderson, and Robert

W. McCarley, Department of Psychiatry and Surgical Planning Lab of Brigham & Women's Hospital, and Brockton VA Medical Center. The authors would like to acknowledge Dr. Shenton's NIMH grants, K02 M-01110 and R01 MH-50747, and Dr. McCarley's grant, R01-40799, and the Brockton Schizophrenia Center for the Department of Veterans Affairs.

References

- [1] V. Caselles, R. Kimmel, and G. Sapiro. Geodesic active contours. *Int'l J. Comp. Vision*, 22(1):61–79, 1997.
- [2] V. Caselles, J. Morel, G. Sapiro, and A. Tannenbaum. Introduction to the special issue on PDEs and geometry-driven diffusion in image processing and analysis. *IEEE Trans. on Image Processing*, 7(3):269–274, 1998.
- [3] L. D. Cohen. On active contour models and balloons. *CVGIP: Image Understanding*, 53(2):211–218, 1991.
- [4] T. Cootes, C. Beeston, G. Edwards, and C. Taylor. Unified framework for atlas matching using active appearance models. In *Int'l Conf. Inf. Proc. in Med. Imaging*, pages 322–333. Springer-Verlag, 1999.
- [5] Y. Guo and B. Vemuri. Hybrid geometric active models for shape recovery in medical images. In *Int'l Conf. Inf. Proc. in Med. Imaging*, pages 112–125. Springer-Verlag, 1999.
- [6] D. Huttenlocher, G. Klanderman, and W. Rucklidge. Comparing images using the hausdorff distance. *IEEE Trans. Patt. Analysis and Mach. Intell.*, 15:850–863, 1993.
- [7] M. Jones and T. Poggio. Multidimensional morphable models. In *IEEE Int'l Conf. Comp. Vision*, pages 683–688, 1998.
- [8] M. Kass, A. Witkin, and D. Terzopoulos. Snakes: Active contour models. *Int'l J. Comp. Vision*, 1(4):321–331, 1988.
- [9] A. Kichenassamy, A. Kumar, P. Osher, A. Tannenbaum, and A. Yezzi. Gradient flows and geometric active contour models. In *IEEE Int'l Conf. Comp. Vision*, pages 810–815, 1995.
- [10] L. M. Lorigo, O. Faugeras, W. Grimson, R. Keriven, R. Kikinis, and C.-F. Westin. Co-dimension 2 geodesic active contours for mra segmentation. In *Int'l Conf. Inf. Proc. in Med. Imaging*, pages 126–139. Springer-Verlag, 1999.
- [11] S. Osher and J. Sethian. Fronts propagating with curvature-dependent speed: Algorithms based on hamilton-jacobi formulation. *J. of Comp. Physics*, 79(1):12–49, 1988.
- [12] N. Paragios and R. Deriche. Geodesic active regions for supervised texture segmentation. In *IEEE Int'l Conf. Comp. Vision*, pages 926–932, 1999.
- [13] J. A. Sethian. *Level Set Methods*. Cambridge University Press, 1996.
- [14] L. Staib and J. Duncan. Boundary finding with parametric deformable contour models. *IEEE Trans. Patt. Analysis and Mach. Intell.*, 14(11), 1992.
- [15] M. Turk and A. Pentland. Eigenfaces for recognition. *J. of Cogn. Neuroscience*, 3(1):71–86, 1991.
- [16] Y. Wang and L. Staib. Boundary finding with correspondence using statistical shape models. In *Proc. IEEE Conf. Comp. Vision and Patt. Recog.*, pages 338–345, 1998.
- [17] A. Yezzi, A. Tsai, and A. Willsky. A statistical approach to snakes for bimodal and trimodal imagery. In *IEEE Int'l Conf. Comp. Vision*, pages 898–903, 1999.




Figure 9. Several time steps in the curve evolution process of segmenting two rhombi. The training set for the rhombus consisted of rhombi of various sizes and aspect ratios. The black curve is the zero level set of the evolving surface. The gray curve is the next step in the curve evolution. The white curve is the MAP estimate of the position and shape of the final curve.




Figure 10. Initial, middle, and final steps in the evolution process of segmenting two slices of the femur. The training set consisted of 18 slices of the same femur, leaving out the slice being segmented and its neighbors.




Figure 11. Four steps in the segmentation of two different corpora callosa. The last image in each case shows the final segmentation in black. The dotted contour is the standard evolution *without* the shape influence.




Figure 12. Early, middle, and final steps in the segmentation of the vertebra T7. Three orthogonal slices and the 3D reconstruction are shown for each step. The black contour is a slice through the evolving surface. The white contour is a slice through the inside of the MAP final surface.