

Segmentation of Medical Images Using a Genetic Algorithm

Payel Ghosh
Dept. of ECE
Portland State University
Portland, OR 97207
payel@pdx.edu

Melanie Mitchell
Dept. of Computer Science
Portland State University
Portland, OR 97207
mm@cs.pdx.edu

ABSTRACT

Segmentation of medical images is challenging due to poor image contrast and artifacts that result in missing or diffuse organ/tissue boundaries. Consequently, this task involves incorporating as much prior information as possible (e.g., texture, shape, and spatial location of organs) into a single framework. In this paper, we present a genetic algorithm for automating the segmentation of the prostate on two-dimensional slices of pelvic computed tomography (CT) images. In this approach the segmenting curve is represented using a level set function, which is evolved using a genetic algorithm (GA). Shape and textural priors derived from manually segmented images are used to constrain the evolution of the segmenting curve over successive generations.

We review some of the existing medical image segmentation techniques. We also compare the results of our algorithm with those of a simple texture extraction algorithm (Laws' texture measures) as well as with another GA-based segmentation tool called GENIE. Our preliminary tests on a small population of segmenting contours show promise by converging on the prostate region. We expect that further improvements can be achieved by incorporating spatial relationships between anatomical landmarks, and extending the methodology to three dimensions.

Categories and Subject Descriptors

I.4 [Image Processing and Computer Vision]: Segmentation – *pixel classification, edge and feature detection.*

General Terms: Algorithms, Experimentation.

Keywords

Level Set Methods, Texture Segmentation, Genetic Algorithms.

1. INTRODUCTION

Identifying specific organs or other features in medical images requires a considerable amount of expertise concerning the shapes and locations of anatomical features. Such segmentation is typically performed manually by expert physicians as part of

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

GECCO'06, July 8–12, 2006, Seattle, Washington, USA.
Copyright 2006 ACM 1-59593-186-4/06/0007...\$5.00.

treatment planning and diagnosis. Due to the increasing amount of available data and the complexity of features of interest, it is becoming essential to develop automated segmentation methods to assist and speed-up image-understanding tasks.

Medical imaging is performed in various modalities, such as magnetic resonance imaging (MRI), computed tomography (CT), ultrasound, etc. Several automated methods have been developed to process the acquired images and identify features of interest, including intensity-based methods, region-growing methods and deformable contour models [17]. Intensity-based methods identify local features such as edges and texture in order to extract regions of interest. Region-growing methods start from a seed-point (usually placed manually) on the image and perform the segmentation task by clustering neighborhood pixels using a similarity criterion. Deformable contour models are shape-based feature search procedures in which a closed contour deforms until a balance is reached between its internal energy (smoothness of the curve) and external energy (local region statistics such as first and second order moments of pixel intensity). Such methods are typically based on only one image feature, such as texture, shape, pixel intensity, etc. However, due to the low contrast information in medical images, an effective segmentation often requires extraction of a combination of features such as shape and texture or pixel intensity and shape. This paper describes our attempts to develop a segmentation algorithm that incorporates both shape and textural information to delineate a desired object in an image. In particular, motivated by the work of Harvey et al. [7] and Tsai et al. [20][21], we developed a genetic algorithm for medical image segmentation. The genetic algorithm framework [14] brings considerable flexibility into the segmentation procedure by incorporating both shape and texture information. In the following sections we describe our algorithm in depth and relate our methodology to previous work in this area. We start by reviewing the active shape modeling approach for image segmentation, specifically the level set method of shape characterization. We then describe texture-based segmentation methods such as Laws' textural feature extraction method. We also describe our studies of the GENIE system for multi-spectral feature extraction. Finally, we present the results of our method on segmenting the prostate based on a small training set of pelvic CT images. We compare these results with those from similar runs on GENIE.

1.1 Active Shape Modeling

Since the pioneering work by Kass et al. [9], much work has been done on active-contour approaches for image segmentation [2][4][12]. Active-contour segmentation algorithms automatically construct one or more contours that segment a particular structure

or a set of structures in the image. These algorithms associate the segmenting contour with an energy (cost) function usually defined by curvature or image gradient. The curve representing the segmenting contour is deformed by minimizing the internal and external energy of the curve. The internal energy is defined as an intrinsic property of the curve itself, such as the smoothness of the curve (a curve segment is defined as smooth if the derivative of the function defining the curve exists, and is nonzero at all points on that segment of the curve). The external energy is defined using extrinsic properties (i.e., properties of the image and not the curve) such as the image gradient or pixel intensity. The energy function is a weighted sum of internal and external energy terms. Minimizing this energy function attracts the contour towards the object.

The level set approach for active contour modeling was proposed by Malladi et al. [13]. This methodology became very popular due to its ability to automatically represent changes in the topology of dynamic curves, such as the boundaries of soap bubbles, flames and other physical phenomena whose shape changes with time. In this approach the evolving boundary (interface) is built into a surface by adding another dimension to the curve evolution coordinate system. This level set function is defined in terms of the signed distance function. The signed distance function takes any pixel in the image and returns as its output the Euclidean distance between the pixel and the closest point on the interface. Pixels outside the interface have positive distance while pixels inside have negative distance values assigned to them. The “zero” level set is defined as the interface itself, i.e., the set of all points whose distance to the interface is zero. Figure 1 shows a square-shaped object in a binary image and illustrates how the signed distance map is computed from its pixel values.

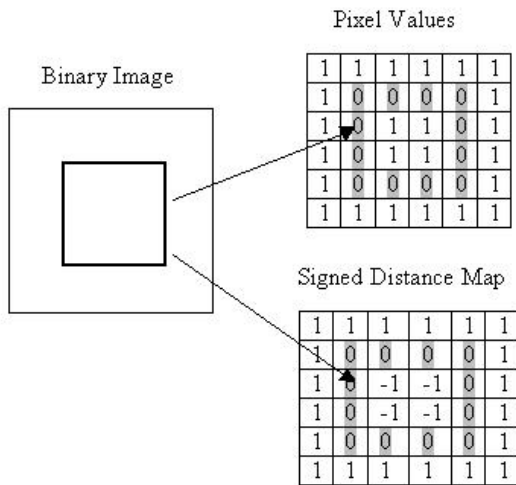


Figure 1. A square shaped object in a binary image. The signed distance values are computed using the Euclidean distance of each pixel from the closest point on the contour.

The level set approach is a powerful and general technique for image segmentation. In this framework, a three-dimensional surface is created from the signed distance representation of the contour with negative values representing regions deeper than or below the zero level and positive values representing regions

protruding above the zero level. As the contour deforms, the zero-valued pixels of the signed distance map move along this three dimensional surface. Thus, this third dimension depicts the time dimension of the contour deformation. This representation of shape is tolerant to slight misalignments of object features and does not require finding correspondences between the pixel coordinates of the original and the deformed contours. Shape statistics such as mean shape and variance can be computed directly from the signed distance maps instead of averaging over pixel coordinates of different contours.

Level set methods have been used by Leventon et al. [12], and Chan and Vese [3] for medical image segmentation. Leventon et al. introduced the concept of shape representation by principal component analysis (PCA) on signed distance functions. They also incorporated statistical shape priors (i.e., shape information from training examples) into their geodesic active contour model in order to generate a posteriori estimates of pose and shape. (Pose is defined as a representation of the position, size, and orientation of an object in an image). Vese and Chan [22] introduced a region-based energy function in order to detect features with diffuse boundaries.

A shape-based level set function was derived by Tsai et al. [20][21] and has been incorporated in this paper. Tsai et al.'s goal was to find the parameters of the level set function that produce a good model of the object shape based on priors from the training data. Tsai et al. derived these parameters via an optimization procedure that used statistics based on the pixel intensities of local regions in a set of training images.

1.2 Texture-Based Segmentation

Texture is defined as a quantitative measure of the variation in intensity of a surface. Texture-based segmentation algorithms are aimed at finding similarity measures to group image pixels. Various approaches for textural feature extraction have been developed to date [19] including co-occurrence matrices, wavelet-based methods, Fourier transform methods, and intensity histogram methods, to name a few.

For this project we compared our system with a simple texture feature extraction method called Laws' texture energy measures [10]. The basic 1-D convolution kernels derived by Laws stand for level (L), edge (E), spot (S), wave (W) and ripple (R) texture types respectively. Justification for the convolution kernels can be found in [11]. Two-dimensional masks are generated from these kernel vectors by convolving each vector with the transpose of the other. The textural energy features are obtained by convolving an image with these two-dimensional integer coefficient masks (usually 5x5) followed by a non-linear windowing operation.

1.3 GENIE

Genetic algorithms have been applied to many image processing problems, such as edge detection [6], image segmentation [18], image compression [15], and feature extraction from remotely sensed and medical images [5]. A general-purpose image-segmentation system called GENIE (“Genetic Imagery Exploration”) [7][8][16] was developed at the Los Alamos National Laboratory. GENIE uses a genetic algorithm to evolve image-processing “pipelines”: sequences of elementary image processing operations, including morphological, arithmetic and

point operators, and filters and edge detectors, among others. Each pipeline performs a segmentation of the image by classifying each pixel as being a positive or negative instance of a desired “feature” such as water, clouds, snow, etc.

The genetic algorithm, starting with a population of random pipelines, evaluates the fitness of pipelines in the population, selects the fittest to produce the next generation, using crossover and mutation to produce offspring. The fitness of each pipeline in the population is computed by comparing its final classification output with a set of training images, in which positive and negative examples of the desired feature have been manually highlighted. At the end of a run of GENIE, the “fittest” pipeline in the population is used in conjunction with a linear classifier to segment the desired feature in new images by labeling each pixel as positive or negative.

Harvey et al. [7] applied GENIE to a medical feature-extraction problem using multi-spectral histopathology images. Their specific aim was to identify cancerous cells on images of breast cancer tissue. Their method was able to discriminate between benign and malignant cells from a variety of samples.

1.4 Overview of Our Work

In this paper we describe our method to combine high-level textural and shape information for image segmentation. Our system uses a set of training images in each of which a segmenting contour surrounding a particular object (e.g., the prostate in a two-dimensional CT image) is drawn by hand. In our system a segmenting contour is represented by a level set function. Each segmenting contour has a unique shape and pose (i.e., size, position, and orientation).

We also have a set of “test images”, not included in the training set, for which a human has provided segmenting contours. Given a new image containing an object of the desired class, the goal here is to evolve a contour that segments that object in the new image, such that the contour obeys shape constraints learned in the training images and also encloses a region whose texture is a good match for textures learned in the training images. Several candidate contours form the individuals of a GA population. The GA is iterated until a fitness value greater than a certain threshold is achieved or the number of generations equals 1000.

Earlier work on segmentation based on level set methods typically derived a curve evolution equation or used gradient descent procedures to search for a contour that minimizes internal and external energy. However, only first and second order statistics such as pixel intensity or variance have been used in these methods because they can be easily incorporated in an implicit representation of the curve. The derived gradient for this implicit function determines the direction of curve evolution.

Cagnoni et al. [1] used a GA for segmenting medical images. The GA optimized the parameters of an elastic contour model using edge information (first-order statistics) from the images. In contrast, the GA framework here allows the use of any kind of high-level textural features for performing segmentation. The fitness function based on textural priors gives a fitness score that is used to rank good candidate solutions and propagate them to future generations. This eliminates the need to derive gradients of energy functions unlike other active shape contour model based segmentation algorithms.

2. PROCEDURE

A two-stage approach is proposed here for image segmentation using a genetic algorithm: the training stage and the segmentation stage. The data for the training stage is obtained from a set of n training images on which a human has outlined the object to be segmented by drawing a contour around it (e.g., the prostate in a 2D slice of a pelvic CT image). The “shape prior” of a training set is defined as a representation of the mean shape over all these manually drawn contours, together with the average deviations from that mean. The textural properties of the object of interest are also derived from the same set of training data. The segmentation phase consists of the genetic algorithm evolving candidate solutions (i.e., candidate contours for segmenting the desired object in a new image), iterating over successive generations until a stopping criterion is satisfied.

To summarize, the steps in this procedure are as follows:

1. From a set of n manually segmented training images, derive a representation of the shape prior, that is, the mean shape and variability of the n segmenting contours.
2. From the same training images, derive a representation of the mean texture of the segmented objects.
3. Given a new image not in the training set, use the GA to evolve a segmenting contour for delineating the desired object in this image, as follows:
 - (i) Start with an initial population of randomly generated shapes, constrained by the shape prior from step (1).
 - (ii) The fitness of a given shape is determined by the match between the texture of its enclosed region and the mean texture from step (2).
 - (iii) Perform selection, crossover, and mutation, as will be described below, to form a new population.
 - (iv) Repeat until a fitness score above a certain threshold is achieved or the number of iterations exceeds 1000.
4. Calculate the “goodness of fit” of the fittest individual from the final generation, as described below.

2.1 The Training Phase

2.1.1 Deriving Shape Priors

To derive the shape prior, each contour from the training data is represented as the zero level set of the signed distance function $\psi_i(x,y)$ (where (x,y) are the pixel coordinates and $i=1$ to n , the number of training contours used to find the shape variability). The mean shape and shape variability of the contours, obtained from the training images, are computed, using the methodology described in [20]. The mean level set function is defined (for n contours) as:

$$\bar{\Phi}(x, y) = \frac{1}{n} \sum_{i=1}^n \psi_i(x, y) \quad (1)$$

Mean offset functions are then derived by subtracting the mean from the signed distance representations of the training contours ($\tilde{\psi}_i = \psi_i - \bar{\Phi}$). Assume the image is of size $N = N_1 \times N_2$. Let β_i be the size $N \times 1$ column vector consisting of the N_2 columns of mean offset image $\tilde{\psi}_i$ stacked to make a single column vector. A

new matrix \mathbf{S} (size $N \times n$), called the shape variability matrix, is formed from n such column vectors, one for each training image

$$\mathbf{S} = [\beta_1, \beta_2, \dots, \beta_n] \quad (2)$$

The variance in shape is then computed by the eigenvalue decomposition of this shape variability matrix.

$$\frac{1}{n} \mathbf{S} \mathbf{S}^T = \mathbf{U} \mathbf{\Sigma} \mathbf{U}^T \quad (3)$$

Here \mathbf{U} is an $N \times n$ matrix whose columns represent n orthogonal modes of shape variation and $\mathbf{\Sigma}$ is an $n \times n$ diagonal matrix of eigenvalues. By rearranging the columns of \mathbf{U} to form an $N_1 \times N_2$ structure, the n different eigenshapes can be obtained $\{\Phi_1, \Phi_2, \dots, \Phi_n\}$. (Details for this procedure can be found in [20].)

2.1.2 Deriving Textural Priors

We define textural priors as the high-level feature vectors derived for each pixel in the training image. Two approaches to deriving high-level texture features are Laws' textural measures and GENIE. We acquired an open source release version of the GENIE software from Los Alamos National Laboratory and tested it using our training images to obtain the high-level textural features. In section 3 we compare the results of our GA using these textural features with the results of GENIE alone as well as Laws' textural measures alone, as applied to the pelvic CT images.

2.2 The Segmentation Phase

Each individual in the GA population consists of a fixed-length string of real-valued "genes". Each such string represents a vector of shape and pose parameters defined as follows. Pose parameters are incorporated into this framework using an affine transform. The affine transform is the product of three matrices (equation 4): the translation matrix, the scaling matrix and the rotation matrix respectively. If x and y are the pixel coordinates of the input image then \tilde{x}, \tilde{y} are the pixel coordinates of the affine-transformed image given by equation (4).

$$\begin{bmatrix} \tilde{x} \\ \tilde{y} \\ 1 \end{bmatrix} = \begin{bmatrix} 1 & 0 & a \\ 0 & 1 & b \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} h & 0 & 0 \\ 0 & h & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} \cos \theta & -\sin \theta & 0 \\ \sin \theta & \cos \theta & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ 1 \end{bmatrix} \quad (4)$$

Here a, b are translation parameters, h is the scaling factor, and θ is the degree of rotation of an object. Following the lead of [20], the mean shape and shape variability derived from the training phase are used to define a level set function (equation 5) that implicitly represents the segmenting curve.

$$\Phi[x, y] = \overline{\Phi}(x, y) + \sum_{i=1}^k w_i \Phi_i(x, y) \quad (5)$$

Each individual \mathbf{I} in the population, represents a segmenting curve defined by the weighted eigenshapes and pose parameters:

$$\mathbf{I} = [\mathbf{W}, \mathbf{P}] \quad (6)$$

where, $\mathbf{W} = [w_1, w_2, \dots, w_k]$, and $\mathbf{P} = [a, b, h, \theta]$. Thus the number of real-valued genes on a GA chromosome is $k+4$, where k is the number of principal eigenshapes. For the experiments reported here, we set k to 6. To form an initial population of individuals for the GA, the weights for the k principal eigenshapes (w_i) are chosen randomly from the space of $[0, \sigma_i^2]$ (where σ_i^2 are the

eigenvalues corresponding to the k eigenshapes). The pose parameters, \mathbf{P} are chosen randomly from the range of values specified in Table 1.

The segmenting contour represented by the GA individual can be expressed as:

$$\Phi(\tilde{x}, \tilde{y}) = \overline{\Phi}(\tilde{x}, \tilde{y}) + \sum_{i=1}^k w_i \Phi_i(\tilde{x}, \tilde{y}) \quad (7)$$

The fitness of each individual is measured by comparing the textural properties of the region segmented by that individual to the desired texture derived from the training images. First, the textural feature planes (formed from the textural feature vectors generated for each pixel of a image) are generated for the test image (a new image not in the training set). Each pixel of the test image is then classified as "True" (desired texture) or "False" (otherwise) by using a Fisher linear discriminant. A Fisher linear discriminant finds an optimal linear combination of the feature planes and maximizes the separation between the desired and undesired texture features (i.e., maximizes the fitness). The fitness function is similar to the one used in GENIE [7][16]:

$$F = 500(A + (1-B)) \quad (8)$$

Here, A denotes the detection rate: i.e., the fraction of pixels inside the segmenting contour that are labeled "True" (i.e., texturally similar to the prostate). B denotes the false alarm rate: the fraction of pixels outside the segmenting contour that are labeled "True". An increase in fitness means that more pixels inside (and fewer pixels outside) the contour are labeled as "True". A fitness score of 1000, therefore, represents a perfect segmentation result.

Our GA uses rank selection, single-point crossover, and mutation. Rank selection is implemented by assigning a numerical rank to each individual, based on its fitness value, and by making higher ranked individuals more likely to be selected to produce offspring. Fixed-length individuals are used here, and single-point crossover is implemented (here, with probability 1 per pair of parents) by swapping same length segments of genes between two individuals. For each pair, a single crossover point is chosen randomly with uniform probability over genes in the chromosome. Mutation is performed by randomly changing the value of a gene (one of the w_i or one of the a, b, h, θ values that make up an individual) based on the ranges specified in Table 1.

Table 1. GA Parameters

Population Size	25
Mutation Probability	0.02 per gene
Crossover Probability	1.0, Single Point
Selection Criteria	Rank Selection
k, No. of principal eigenshapes	6
Pose parameters a, b	Integer (0-10)
Pose parameter θ	0°-360°
Pose parameter h	(0.5-2.0)

The GA is iterated until the optimum fitness is attained or after a specified number of generations has been produced. Following [8], we define the goodness of fit, G , as a means to evaluate the closeness of a candidate segmenting contour to the human-drawn contour for a given test image. The fitness function defined in equation (8) is not used as a goodness of fit because it is defined based on the textural prior obtained from the training images.

To calculate G , we generate two binary images corresponding to the human-drawn contour and the contour derived from an evolved individual: in each, the pixels inside the segmenting contour are set to 1 and outside are set to 0. Define H as the Hamming distance between these two binary images—that is, the number of pixels that are classified differently (wrongly) in the evolved individual's segmentation from corresponding pixels in the manually segmented binary image. The goodness of fit is numerically defined as:

$$G = (1 - (H/N)) \times 1000 \quad (9)$$

where N is the total number of pixels in the image. A score of 1000 represents a perfect match with the training data.

3. RESULTS




Figure 2. A typical 2D pelvic CT scan before manual segmentation by a radiologist.




Figure 3. The 2D pelvic CT scan of Figure 2 manually segmented by a radiologist. The white contour labeled "prostate" was drawn by hand.




Figure 4. A different slice from the CT scan of the same patient.




Figure 5. The prostate region marked the CT scan for the same patient. The white contour was marked by a radiologist.

We tested our system on images taken from a database of 2700 pelvic CT scans, acquired through collaborations with radiologists at Oregon Health & Science University (OHSU). About 100 CT scans from this database have been manually segmented by Arthur Hung, M.D. (Dept of Radiation Oncology, OHSU).

A 3D CT scan for each patient consists of 15-20 slices of 2D images stacked together. The prostate is visible in about 10-12 of these slices; the rest display other organs in the pelvic region such as the bladder and the rectum. The prostate is located between the bladder and the rectum and is about 3 cm in size along the height of the body. The bladder and the rectum are more texturally prominent on the CT scans and are used by the radiologist to help locate the prostate on these images. The prostate has been manually delineated three times on the same set of images by the radiologist. This provides a database for intra-operator variability.

Figure 2 shows a typical pelvic CT scan and Figure 3 shows the same image with a manually placed contour depicting the prostate area. Figure 4 shows a different slice of the same patient's CT scan. Figure 5 shows the manually placed prostate contour on this new slice. Note how the shape of the prostate changes from slice to slice.

The manually segmented contours derived from the CT scans of one patient (10 x 3= 30 images) have been used as the training data for this analysis. A set of 10 2D CT scan slices from another patient has been used as test images. Figure 6 shows a test image used for experiments here.

Prostate segmentation is challenging because the shape and size of the prostate varies considerably across patients. Also there are neither significant edges nor distinct textural differences to make the prostate visible in these images. Thus the interface between the prostate and the bladder or the rectum is typically not clearly defined. An expert radiologist uses prior knowledge of organ shapes and the relative positions of various anatomical landmarks to approximately and intelligently “guess” the location of the prostate on these images. A longer-term goal of our project is to simulate this procedure accurately by incorporating prior information about the relative spatial locations of organs.

We compared our system with other methods for segmenting the prostate on the pelvic CT images. First, we tried to segment the images using Laws’ textural priors alone. The textural energy measures were derived for every pixel on the training images to produce texture feature planes [10]. An optimal linear combination of these feature planes was then derived such that the separation between feature pixels (pixels in the prostate region) and non-feature pixels was maximized. This Fisher linear discriminant was used on test images to classify pixels as feature or non-feature based on the computed texture energy planes for the test images. We found that Laws’ textural priors could only differentiate soft-tissue (regions marked white on Figure 7) from body cavity (regions marked black on Figure 7) due to the relatively low contrast information in the images.

We then trained GENIE to generate an image processing pipeline to discriminate the prostate region on the same set of training images. This image processing pipeline was applied to the test images to derive regions on the pelvic CT scans with high-level texture similar to the prostate region. A sample segmentation result of a GENIE run on a test image is shown in Figure 8. The pixels classified by GENIE as texturally similar to the prostate are marked white and the other pixels are marked black. Not surprisingly, because GENIE does not create contour shape descriptions, it was not successful in identifying the prostate.

In our next experiment we implemented our genetic algorithm with the parameters specified in Table 1. The initial contour was generated using mean shape and shape variability information derived from the training images and was placed randomly on a test image. The evolution of the curve (selection after cross-over and mutation) was guided by the fitness function derived from textural priors evolved by GENIE. We found that the curve evolution process was able to converge on the approximate area of the prostate. Also, the shape priors constrained the growth of the curve to within the expected shape of the prostate. Figure 9 shows the segmentation result of our algorithm on a single slice of the pelvic CT image. It took about 20 generations for the GA to converge on a contour that gives a reasonably good segmentation of the prostate area.




Figure 6. A slice from the CT scan of another patient (test image)



Figure 7. Segmentation result (on the test image) using Laws’ textural measures. The white regions marked are classified texturally similar to prostate (positive). The black region is classified as negative.




Figure 8. Segmentation result using GENIE: white regions (positive classifications), black regions (negative classifications)




Figure 9. Segmentation result using our GA with shape and textural priors, on one test image. The white contour shows the final segmenting contour (labeled as prostate).




Figure 10. The test image segmented by a radiologist. This contour was used to find the goodness of fit of the segmenting contour in Figure 9.

Not surprisingly, we found that our algorithm performed significantly better in segmenting the prostate on the images than the other methods that use textural information alone. Table 2 shows the goodness of fit of the final segmentation results for each of the methods implemented here. The values given are averages of G over all the slices for the given patient (slices from one patient were used for the training data and slices from a second patient were used for the test data).

Table 2. Goodness of fit of the final segmentation obtained from the three different methods. The values given are averages of G over all the slices for the given patient.

Classifier	G : Training Data	G : Test Data
Our GA	985	991
GENIE	950	708
Laws' Texture Measures	850	580

4. CONCLUSIONS AND FUTURE WORK

The algorithm developed here evolves a segmenting contour by incorporating both texture and shape information to extract objects without prominent edges, such as the prostate on pelvic CT images. Representing the shape of the contours as level sets and encoding candidate solutions of the GA as segmenting contours eliminates the need for deriving the gradients of energy functions for shape evolution and simplifies the optimization procedure. Our experiments using a small training set and a small population of candidate segmentation contours shows promise by converging on the prostate area.

The following enhancements to the above framework are proposed for improving the segmentation results.

1. Incorporating position information: The relative position of the various organs, if incorporated, can be used for initial placement of the segmenting curve (which is random at present). This has the potential to significantly improve the segmentation results
2. Extension to 3-D: Following the lead of [20], the pose parameters can be extended to represent the 3D pose of an object. The above framework can be used to evolve a surface instead of a curve in a 3-D domain. Thus information from all the slices of a CT scan can be used simultaneously for 3-D segmentation. We would also like to compare our results with the shape-based segmentation procedure implemented by Tsai et al.

5. ACKNOWLEDGMENTS

We are grateful to the Intel Corporation and the J. S. McDonnell Foundation for supporting this research. Thanks also to Dr. Arthur Hung for providing training data for this analysis and to Dr. Xubo Song and Kun Yang for their help in providing data and for valuable discussions on this project.

6. REFERENCES

- [1] Cagnoni, A., Dobrzeniecki, A. B., Poli, R., and Yanch, J. C. Genetic algorithm-based interactive segmentation of 3D medical images. *Image and Vision Computing*, 17(12), 881-895, 1999.
- [2] Caselles, V., Kimmel, R., and Shapiro, G. Geodesic active contours. *Int. J. Comput. Vis.*, 22, 61-79, 1998.
- [3] Chan, T., and Vese, L. Active contours without edges. *IEEE Trans. Image Proc.*, 10, 266-277, 2001.
- [4] Cohen, L. On active contour models and balloons. *CVGIP: Image Understanding*, 53, 211-218, 1991.
- [5] Daida, J.M., Hommes, J. D., Bersano-Begey, T.F., Ross, S. J, and Vesecky, J. F. Algorithm discovery using the genetic programming paradigm: Extracting low-contrast curvilinear features from SAR images of Arctic ice. In *Advances in Genetic Programming 2*, P. J. Angeline, and K. E. Kinnear, Jr., editors, chap. 21, MIT, Cambridge, 1996.
- [6] Harris, C., and Buxton, B. Evolving edge detectors with genetic programming. In Koza, J. R., Goldberg, D. E., Fogel, D. B., and Riolo, R. L., editors, *Genetic Programming 1996: Proceedings of the First Annual Conference*, pp. 309-315. MIT Press, 1996.
- [7] Harvey, N., Levenson, R. M., Rimm, D. L. Investigation of automated feature extraction techniques for applications in

- cancer detection from multi-spectral histopathology images. *Proc. of SPIE 2003*, 5032, 557-566.
- [8] Harvey, N., Perkins, S., Brumby, S. P., Theiler, J., Porter, R. B., Young, A. C., Verghese, A. K., Szymanski J. J., and Bloch, J. J. Finding golf courses: The ultra high-tech approach. *Lecture Notes in Computer Science*, vol. 1803, pp.54-64, 2000.
- [9] Kass, M., Witkin, A., and Terzopoulos, D. Snakes: Active contour models. *Intl. J. Computer Vision*, 1, 321-331, 1988.
- [10] Laws, K. I. Rapid texture identification. *SPIE, Image Processing for Missile Guidance*, 238, 376-380, 1980.
- [11] Laws, K. I. *Texture Image Segmentation*. PhD dissertation, Univ. of Southern Calif., Jan. 1980.
- [12] Leventon, M., Grimson, E., and Faugeras, O. Statistical shape influence in geodesic active contours. In *Proc. IEEE Conf. on Computer Vision and Patt. Recog.*, 1, 316-323, 2000.
- [13] Malladi, R., Sethian, J. A., and Vemuri, B. C. Shape modeling with front propagation: A level set approach. *IEEE Trans. on Patt. Recog. and Machine Intell.*, 17(2), 158-175, 1995.
- [14] Mitchell, M., *An Introduction to Genetic Algorithms*, Cambridge, MA, MIT Press, (1996).
- [15] Nordin, P., and Banzhaf, W. Programmatic compression of images and sound. In *Genetic Programming 1996, Proceedings of the First Annual Conference*, Koza J. R. et al. editors, pp. 345-350, MIT Press, 1996.
- [16] Perkins, S., Theiler, J., Brumby, S. P., Harvey, N. R., Porter, R. B., Szymanski, J. J., and Bloch, J. J. GENIE: A Hybrid Genetic Algorithm for Feature Classification in Multi-Spectral Images. *Proc. SPIE 4120 Applications and Science of Neural Networks, Fuzzy Systems, and Evolutionary Computation III*, pp. 52-62, 2000.
- [17] Pham, D. L., Xu, C., Prince, J. L. Survey of current methods in medical image segmentation. *Annual Review of Biomedical Eng.*, 2, pp. 315-337, 2000.
- [18] Poli, R., and Cagoni, S. Genetic programming with user-driven selection: Experiments on the evolution of algorithms for image enhancement. In *Genetic Programming 1997, Proceedings of the 2nd Annual Conference*, J. R. Koza et al. editors, pp. 269-277, 1997.
- [19] Sonka, M., Hlavac, V., and Boyle, R. *Image Processing, Analysis and Machine Vision*. London, Chapman & Hall, UK, 1994.
- [20] Tsai, A., Yezzi, A., Wells, W., Tempany, C., Tucker, D., Fan, A., Grimson, E., Willsky A. A shape-based approach to the segmentation of medical imagery using level sets. *IEEE Trans. on Medical Imaging*, 22,137-154, 2003.
- [21] Tsai, A., Wells, W., Warfield, S., Willsky, A. An EM algorithm for shape classification based on level sets. *Medical Image Analysis*, Vol. 9, No. 5, 491-502, October 2005.
- [22] Vese L. A., Chan, T. F. A multiphase level set framework for image segmentation using the Mumford and Shah model. *Intl. J. Computer Vision*, 50(3), 271-293, 2002.