

Prostate Segmentation on Pelvic CT Images Using a Genetic Algorithm

Payel Ghosh^a, Melanie Mitchell^b

^aDept. of Electrical Engineering, Portland State University, Portland, OR, USA;

^bDept. of Computer Science, Portland State University, Portland, OR, USA and
Santa Fe Institute, Santa Fe, NM, USA

ABSTRACT

A genetic algorithm (GA) for automating the segmentation of the prostate on pelvic computed tomography (CT) images is presented here. The images consist of slices from three-dimensional CT scans. Segmentation is typically performed manually on these images for treatment planning by an expert physician, who uses the “learned” knowledge of organ shapes, textures and locations to draw a contour around the prostate. Using a GA brings the flexibility to incorporate new “learned” information into the segmentation process without modifying the fitness function that is used to train the GA. Currently the GA uses prior knowledge in the form of texture and shape of the prostate for segmentation. We compare and contrast our algorithm with a level-set based segmentation algorithm, thereby providing justification for using a GA.

Each individual of the GA population represents a segmenting contour. Shape variability of the prostate derived from manually segmented images is used to form a shape representation from which an individual of the GA population is randomly generated. The fitness of each individual is evaluated based on the texture of the region it encloses. The segmenting contour that encloses the prostate region is considered more fit than others and is more likely to be selected to produce an offspring over successive generations of the GA run. This process of selection, crossover and mutation is iterated until the desired region is segmented. Results of 2D and 3D segmentation are presented and future work is also discussed here.

Keywords: Segmentation, shape, texture, deformable geometry

1. INTRODUCTION

This study is aimed at segmenting the prostate, to help speed up the treatment and diagnosis of prostate diseases. The data for this project has been derived from a database of 2700 pelvic CT scans, acquired through collaboration with Oregon Health & Science University (OHSU). About 200 CT scans from this database have been manually segmented by an expert physician. Manual segmentation is a cumbersome process and has to be performed for each slice of the CT scan. 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 such image-understanding tasks. Figure 1 shows a slice from a typical pelvic CT scan and the corresponding manually segmented image depicting the prostate area. Prostate segmentation on CT scans is challenging because the shape and size of the prostate varies considerably across patients. Also, there are no significant edges to make the prostate visible on these images. So the interface between the prostate and the bladder or the rectum is typically not clearly defined. An expert uses prior knowledge of organ shapes and the information of the relative positions of various anatomical landmarks to approximately and intelligently demarcate the prostate on these images. Thus, developing an efficient automated method involves coding the knowledge of the expert into an algorithm.

Motivated by the work of Harvey et al.¹ and Tsai et al.², we developed a genetic algorithm for prostate segmentation. The genetic algorithm makes it possible to combine shape and texture features extracted from training images for automatic segmentation without the need for defining an energy function thereby simplifying the optimization procedure. The longer-term goal of the project is to make this algorithm more accurate by incorporating information about the relative spatial locations of organs.

The sections in this paper are organized as follows: the significance of using a genetic algorithm is described in section 2 followed by a brief description of the genetic algorithm for 2D and 3D segmentation is presented in section 3. Discussion of the segmentation results and comparison with a level-set based algorithm are presented in section 4. Conclusion and future work are discussed in section 5.




Fig. 1. A typical 2-D pelvic CT scan (Left). Manually segmented prostate region marked by a radiologist (Right).

2. SIGNIFICANCE OF THE DEVELOPED ALGORITHM

Several automated/semi-automated methods have been developed to segment medical images. Some of the popular segmentation methods are intensity-based methods, region-growing methods and deformable contour models³. 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 which use gradient descent to minimize an energy (cost) function usually defined by curvature or image gradient^{4, 5}. If $C(q) : [0,1] \rightarrow \mathbb{R}^2$ be the parameterized curve (for every point q on the curve C), then the energy function in the so-called classical “snakes” approach is defined as⁴:

$$E(C) = \int_0^1 \alpha |C'|^2 + \beta |C''|^2 dq - \lambda \int_0^1 |\nabla I(C)| dq. \quad (1)$$

The GA-based segmentation algorithm presented in this paper combines high-level texture and shape information to segment medical images in 2D as well as 3D domains. The signed distance transform-based shape representation derived by Tsai et al.² has been adopted here for shape representation. The goal of the algorithm is to find the parameters of this function that produce a good model of the object shape based on the mean and variability of the shape derived from the training data. Level-set methods have been used by Chan and Vese⁶ to detect features with diffuse boundaries on medical images. They used first and second (low-order) order statistical features (such as pixel intensity and variance) to define a level-set function for curve evolution. In contrast, the genetic algorithm here uses higher-level textural and shape features for performing segmentation. The fitness function is based on textural priors and gives a fitness score that is used to rank good candidate solutions and propagate them to future generations. Segmentation result on the pelvic CT images comparing the results of the genetic algorithm with the level-set algorithm of Chan and Vese⁶ has been presented here.

3. METHODOLOGY

The flowchart below provides a brief description of the genetic algorithm. Shape and texture information are first derived from the training images on which a contour has been drawn by an expert. Following the lead of Tsai et al.² the mean shape and shape variability of the prostate is derived to define the shape representation. The evolution of the segmenting contour is thus constrained by the known average shape and variability. The texture of the prostate region is also derived from the training images. The fitness function of the GA is based on the texture of the region enclosed by the segmenting contour. Selection, crossover and mutation are performed sequentially until the fitness of an individual of a generation exceeds a certain threshold or some specified number of generations is produced.




Fig. 2. Flowchart depicting the sequence of operations of the GA.

3.1 Segmentation in 2D

The manually segmented contours are used for deriving the representation of known shape in 2D. The contours are converted to the signed distance representation in which every pixel is assigned a Euclidean distance value from the nearest point on the contour with negative values inside the contour and positive values outside the contour. The contour itself represents the “zero” level of the signed distance representation. Level-set methods use the same representation for evolving interfaces that segment objects on images.

The mean level-set function is defined (for n contours) as:

$$\bar{\Phi} = \left(\frac{1}{n}\right) \sum_{i=1}^n \psi_i. \quad (2)$$

The shape variance is computed using principal component analysis (PCA). In this approach, shape offsets are computed by subtracting the mean from the signed distance representation of the training contours ($\tilde{\psi}_i = \psi_i - \bar{\Phi}$). The columns of the mean offset functions are then successively stacked on top of another to form one large column vector. A new matrix S , called the shape variability matrix, is formed from n such column vectors. The variance in shape is then computed by an eigen value decomposition on this shape variability matrix. The eigen value decomposition is given by:

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

The columns of U represent n orthogonal modes of shape variation (or eigen shapes) Φ_i , and Σ is a $n \times n$ diagonal matrix of eigen values.

The texture of the prostate region is derived using the Laws' textural measures. These texture measures are computed by convolving the training images with small integer coefficient masks⁷. The basic one dimensional convolution kernels derived by Laws stand for level (L), edge (E), spot (S), wave (W) and ripple (R) texture types respectively. Two-dimensional masks are generated from these vectors by convolving each vector with the transpose of another vector. To

generate the texture energy planes, the training images are first convolved with 2D convolution kernels, followed by a windowing operation which sums the absolute pixel values in the neighborhood. A Fisher linear discriminant is then used to find the weights to linearly combine the texture energy planes such that maximum separation of textural features is achieved between pixels inside and outside the contour on the training images. These weights are saved and used to derive the texture energy plane on test images.

The mean shape and shape variability derived from the training images are used to define a level-set function which is (the segmenting curve) the sum of the mean image and weighted deviations from the mean image:

$$\Phi[w] = \bar{\Phi} + \sum_{i=1}^k w_i \Phi_i. \quad (4)$$

Pose parameters are incorporated into this framework using an affine transform. The affine transform is the product of three matrices in 2D: the x-y translation matrix, the scaling matrix and the rotation matrix respectively,

$$\begin{bmatrix} \tilde{x} \\ \tilde{y} \\ 1 \end{bmatrix} = \begin{bmatrix} 1 & 0 & a \\ 0 & 1 & b \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} h & 0 & 0 \\ 0 & h & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} \cos \theta & -\sin \theta & 0 \\ \sin \theta & \cos \theta & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ 1 \end{bmatrix}. \quad (5)$$

Incorporating pose parameters into the shape representation creates segmenting contours of different sizes and orientation.

$$\Phi[w, p](x, y) = \bar{\Phi}(\tilde{x}, \tilde{y}) + \sum_{i=1}^k w_i \Phi_i(\tilde{x}, \tilde{y}). \quad (6)$$

The individuals ($\mathbf{I}$) of the GA population are defined as a vector of these weights (w) and pose parameters (p). Each value of this vector represents a gene of an individual:

$$\mathbf{I} = [w, p] \quad \text{where, } w = [w_1, w_2, \dots, w_k] \text{ and } p = [a, b, h, \theta]. \quad (7)$$

To form the initial population of the GA, the weights for the k -principal eigen shapes (w) are chosen randomly from the space of $[0, \sigma_i]$ (where σ_i^2 are the eigen values corresponding to the k eigen shapes). The pose parameters, p are chosen randomly from the space of uniformly distributed real numbers: a , and b are chosen from the range $[0-10]$; θ from the range $[-30, +30]$ degrees; and h from the range $[0.5-2.0]$.

The fitness of each individual (i.e., segmenting contour for a test image) is measured by analyzing the texture of the segmented region and verifying whether it encloses the prostate region. The analysis is performed by generating the textural feature planes for the test image (a new image not in the training set). Each pixel of the test image is then classified as “true” (prostate) or “false” (otherwise) by using the saved weights of the Fisher linear discriminant. This binary image is compared with the binary segmentation result (ones inside the segmenting contour and zeros outside) to derive a fitness score for the GA individual. The fitness is a function of the detection rate (A) and the false alarm rate (B) as ¹:

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

A fitness score of 1000, therefore, represents a perfect segmentation result. Once the fitness score is derived for all individuals of the population, they are ranked based on their fitness score and chosen for producing children using crossover and mutation. Crossover is performed by swapping fixed length portions of the genes between individuals. Mutation is performed by randomly changing the value of a gene. The process of GA evolution: selection, crossover, and mutation to create a new generation is iterated until the optimum fitness is attained or after a specified number of generations have been produced.

The goodness of fit (G) of the final segmenting curve is found using the Hamming distance (H) between the final segmenting contour and the manually segmented contour corresponding to the test image. The Hamming distance is defined as the total number of pixels that are classified differently (wrongly) from the manually segmented contour. 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 manual segmentation. G is derived on the segmentation results from the GA to test the accuracy of the automated segmentation procedure.

3.2 Extension to 3D

The above procedure is extended to 3D by using 3D pose parameters; x-y-z translation, scale, yaw, pitch and roll,

$$\begin{bmatrix} \tilde{x} \\ \tilde{y} \\ \tilde{z} \\ 1 \end{bmatrix} = \begin{bmatrix} 1 & 0 & 0 & a \\ 0 & 1 & 0 & b \\ 0 & 0 & 1 & c \\ 0 & 0 & 0 & 1 \end{bmatrix} * \begin{bmatrix} h & 0 & 0 & 0 \\ 0 & h & 0 & 0 \\ 0 & 0 & h & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} * R_x * R_y * R_z * \begin{bmatrix} x \\ y \\ z \\ 1 \end{bmatrix}. \quad (10)$$

R_x , R_y , and R_z are the rotation matrices about the x, y and z axes respectively:

$$R_x = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & \cos(\alpha) & -\sin(\alpha) & 0 \\ 0 & \sin(\alpha) & \cos(\alpha) & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}, \quad (11)$$

$$R_y = \begin{bmatrix} \cos(\beta) & 0 & \sin(\beta) & 0 \\ 0 & 1 & 0 & 0 \\ -\sin(\beta) & 0 & \cos(\beta) & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}, \quad (12)$$

$$R_z = \begin{bmatrix} \cos(\theta) & -\sin(\theta) & 0 & 0 \\ \sin(\theta) & \cos(\theta) & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}, \quad (13)$$

The individuals of the 3DGA population are based on the new pose parameters $p = [a, b, c, h, \alpha, \beta, \theta]$. The mean shape and shape variability are derived from the 3D images generated by stacking contours from training images. The 3D segmenting contours generated by this GA segments all the slices of the test image at once. The fitness of the segmenting object is defined similarly and is based on the texture of the region enclosed by it.

3.3 Level-set method of segmentation

In the level-set approach of segmentation, the evolving boundary (interface) is built into a surface by adding another dimension to the curve evolution coordinate system. The level-set approach is a powerful generalized technique for image segmentation, as it does not require finding point correspondences between images. Object shapes can be modeled without finding marker points on the features of interest. In this framework, interfaces corresponding to object shapes in two-dimensions, which we call “contours”, are embedded as zero level-sets of a three-dimensional surface formed by computing the signed distance transform on images. This representation of shape is more tolerant to slight misalignments of object features since the distance values assigned to each pixel after distance transform can be used directly to find shape statistics such as mean shape and variance without the need for finding pixel coordinates on the image.

Level-set methods have been used by Chan and Vese⁶ for medical image segmentation. They introduced a region-based (pixel intensity and variance) energy function in order to detect features with diffuse boundaries on images. The stopping term of their algorithm is based on Mumford-Shah segmentation techniques, which is a minimal partition problem. The algorithm looks for the best approximation of a region-based function which takes only two values, one inside (1) and the other (0) outside the segmenting contour.

4. RESULTS AND DISCUSSION

The training data for this analysis consists of image slices of 3D CT scans from healthy patients on which the prostate is visible. The prostate has been manually segmented two times on the training images by the radiologist to provide the database for intra-operator variability. Figure 3 shows the two different segmentations for the same training image performed manually by the radiologist. The contours are shown stacked on top of each other to create the 3D shape of the prostate. The test data was taken from images of separate subjects. Manual segmentation was also performed on the test images and was used to assess the final segmentation results. The mean shape and shape variability (figure 4) of the prostate on training images was used to derive the shape representation. The initial segmenting contour (initial population of the GA) was placed using the position of the mean image derived from training images. The evolution of the segmenting contour was constrained by the known shape information as the GA run proceeded. Selection, crossover and mutation were performed until 50 generations were produced. The level-set based segmentation was also tried on the same set of test images. In this case too, a small segmenting contour was placed in the center of the test images and the curve was allowed to evolve using gradient descent optimization.

Figure 5 (top left panel) shows a test image with the manually segmented contour and the segmentation result of the level-set algorithm (top right panel). Due to the low contrast of these images, there are no significant differences in the region statistics of the prostate and the nearby organs such as bladder and rectum. Therefore, the low-level texture-based level-set algorithm diverges out of the prostate region. The GA on the contrary uses shape and high-level texture information that constrains the evolution of the segmenting contour to be within known shape bounds and the location of the final contour in the prostate region. The segmentation result of one of the highest fitness individual of the 2DGA on the same test image is shown in figure 5 (bottom left panel). The individual was $[206(w_I), 355(w_2), 832(w_3), 6.72(w_4), 7(a), 10(b), 1.8(h), 9(\theta)]$ and its fitness was 370. Figure 5 (bottom right panel) shows a slice from the 3D segmentation on the same image. The 3D individual was $[1170(w_I), 161(w_2), 9(a), 2(b), 4(c), 1.3(h), 25(\alpha), 13(\beta), 24(\theta)]$ and its fitness was 316. Note that the fitness values are in the range of 0-1000 and they are low because the Law's texture algorithm marks some regions outside the prostate region as texturally similar to the prostate as well. This is again due to the low contrast in these images. Therefore, for achieving better fitness values, the semantic information of relative organ locations must be incorporated into the fitness function. This would be implemented in future.

Figure 6 (left panel) shows the 3D reconstructed view of the prostate obtained by stacking the segmented contours in 2D (from all slices of one test image) together. Since the correlation of shape and pose information across slices has not been incorporated in 2D, the transition in contour sizes from one slice to another is not smooth. Also the segmentation results are not aligned with respect to each other. The segmentation result in 3D is shown in figure 6 (right panel). In contrast to the 2D segmentation the transition in contour sizes is smooth in this case. Also the slices are aligned with each other and the step-like artifacts of the 2D segmentation are removed. The goodness of fit of the segmenting contour compared with the manually segmented contour on the test image was found to be 866 (on a scale of 0-1000) in the case of 2D segmentation and 951 (on scale of 0-1000) in the case of 3D segmentation.




Fig. 3. Two training images with all the manually segmented contours stacked to form a 3D shape.




Fig. 4. Mean Shape (Left panel) and shape variability (right panel) derived from the training images.




Fig. 5. (Top left panel) Manually segmented test image (white contour in the center). (Top right panel) The contour evolved by the level-set algorithm diverges out of the prostate region. (Bottom left panel) 2D segmentation result on the test image. (Bottom right panel) A slice of the 3D segmentation result generated by the GA. The white contour in the center of the bottom panels shows the segmenting contour evolved by the GA.




Fig. 6. (Left panel) Stacked slices from the 2D segmentation on a test image. (Right panel) 3D Segmentation result of the GA on the same test image.

5. CONCLUSION AND FUTURE WORK

The GA framework developed here implements an image segmentation algorithm that incorporates texture and shape information to extract objects without prominent edges, such as the prostate on pelvic CT images. Representing candidate solutions of the GA as segmenting contours and assessing their performance using a fitness function eliminates the need for defining an energy function (and the associated derivatives) and simplifies the optimization procedure needed for curve evolution. Our test results show that the algorithm converges on the prostate region. Better accuracy of final position and initial contour placement are expected to be achieved when the semantic knowledge of organ locations with respect to the prostate is incorporated into this framework.

ACKNOWLEDGEMENTS

This research has been partly funded by Intel Corporation and the J.S. McDonnell Foundation. The authors would like to thank Dr. Arthur Hung (Radiation Oncologist, OHSU) for providing training data for this analysis. Thanks also to Dr. Xubo Song and Kun Yang for their help in providing data and for valuable discussions on this project.

REFERENCES

- [1] 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, 5032, 557- 566(2003).
- [2] Tsai, A. et al., "A shape-based approach to the segmentation of medical imagery using level sets," IEEE Trans. on Medical Imaging, 22,137-154, (2003).
- [3] Pham, D. L., Xu, C., Prince, J. L., "Current methods in medical image segmentation," Annual Review of Biomedical Eng., Vol. 2, 315-337 (2000).
- [4] Caselles, V., Kimmel, R., and Shapiro, G., "Geodesic active contours," Int. J. Comp. Vis., 22, 61-79, (1998).
- [5] Cohen, L., "On active contour models and balloons," CVGIP: Image Understanding, 53, 211-218, (1991).
- [6] Chan, T., and Vese, L., "Active contours without edges," IEEE Trans. Image Proc., 10, 266-277, Feb. (2001).
- [7] Laws, K. I., "Rapid texture identification," SPIE, Image Processing for Missile Guidance, 238, 376-380, (1980).