

A Genetic Algorithm–based Level Set Curve Evolution for Prostate Segmentation on Pelvic CT and MRI Images

Payel Ghosh

Portland State University, USA

Melanie Mitchell

Portland State University, USA

Santa Fe Institute, USA

James A. Tanyi

Oregon Health and Science University, USA

Oregon State University, USA

Arthur Hung

Oregon Health and Science University, USA

ABSTRACT

A novel genetic algorithm (GA) is presented here that performs level set curve evolution using texture and shape information to automatically segment the prostate on pelvic images in computed tomography and magnetic resonance imaging modalities. Here, the segmenting contour is represented as a level set function. The contours in a typical level set evolution are deformed by minimizing an energy function using the gradient descent method. In these methods, the computational complexity of computing derivatives increases as the number of terms (needed for curve evolution) in the energy function increase. In contrast, a genetic algorithm optimizes the level-set function without the need to compute derivatives, thereby making the introduction of new curve evolution terms straightforward. The GA developed here uses the texture of the prostate gland and its shape derived from manual segmentations to perform curve evolution. Using these high-level features makes automatic segmentation possible.

KEYWORDS

Genetic algorithm, Medical image processing, Level set method, Segmentation, Prostate cancer.

INTRODUCTION

Target-volume and organ-at-risk delineation on medical images such as computed tomography (CT), magnetic resonance imaging (MRI), and ultrasound, is a time-consuming process. In radiation therapy (RT) planning, the growth of conformal techniques, intensity-modulated RT with inverse planning, and adaptive four-dimensional RT has increased the difficulty of organ-delineation tasks. Uncertainty and variability in the definition of tumor margins can result in suboptimal treatment of some patients. The development of automated segmentation tools are, therefore, essential but remains a challenge for several reasons such as the variability of organ shapes and variability of tissue contrast on medical images. Despite the various techniques that have been investigated, most medical image segmentation algorithms are either semi-automatic or require some form of human intervention to perform satisfactorily (He et al., 2008; McInerney & Terzopoulos, 1996). This is mainly due to the fact that these algorithms do not encode the knowledge of the human anatomy that a physician uses to manually segment an image. Automatic segmentation can be accomplished if the knowledge of shapes, relative locations, and textures of organs is incorporated into a single algorithmic framework.

Level set methods have become very popular in the field of medical image segmentation due to their ability to represent dynamic or ill-defined boundaries of objects (Sethian, 1999). In this method, a deformable segmenting curve is associated with an energy function. However the complexity of the energy function increases with the number of terms, and computing derivatives for determining the direction of curve evolution can become computationally intensive. A genetic algorithm simplifies the level set evolution process by eliminating the energy function. Instead, it uses a fitness function that implicitly encodes the energy function used for curve evolution. The genetic algorithm presented here uses the learned shape and textural properties of a known object to segment images. Each individual in the GA population is a vector of parameters of a level set function. The GA optimizes the parameters of the level set function to produce fit individuals or good segmentations of the given image using the information encoded in its fitness function.

The rest of the chapter is organized as follows: At first a literature review is provided on the level set method of image segmentation specifically emphasizing algorithms applied to medical images. The new method combining level sets and genetic algorithms (LSGA) is then described in detail. A description of the dataset used and the results achieved from applying the GA to prostate segmentation are then discussed followed by a discussion on the significance of the research work.

LITERATURE REVIEW

Segmentation is the process of demarcating an object of interest on an image. Before segmentation can be performed properties of the object that set it apart from the rest of the image must be determined. These properties can be image pixel-based properties such

as edges, texture, pixel intensity variation inside the object, or object-level properties such as shape, size, orientation, or location with respect to other objects. The pixel-based features are referred to as *low-level features* because they can be inferred using simple image processing routines on an image. For example, edges of an image can be derived using a gradient operator on the image pixel values. The object-level features on the other hand, are so-called *high-level features* because they involve an extra step of finding an appropriate concept to describe a particular feature. For example, “size” of an object can be determined using the distance between two pixels located in the opposite extremities of the object or by the diameter of a circle enclosing the object.

Pixel based operations can be used on simple problems where the object has a prominent edge and markedly different pixel intensity values inside and outside the object. However such techniques alone are not suitable for medical image segmentation, and when used are usually performed semi-automatically with considerable manual intervention. Most medical images have significant amounts of noise and artifacts that appear during the image acquisition process. They also have low contrast and broken/diffuse edges around regions of interest. This makes segmentation on medical images a challenging problem. Therefore, object-based approaches or a combination of pixel and object-based techniques are often more suitable for medical image segmentation. Segmentation using object level-features involves quantifying object characteristics such as shape, pose, and relative position with respect to objects as well as region-based properties of the object. This quantification process transforms the object into a series of feature vectors (Costa & Caesar, 2001). Deriving these transformations and combining them is the main challenge of object-based segmentation.

The shape of an object is an abstract concept since the same shape can be described differently by people based on their varied perceptions. Shapes are generally represented using contours, transforms, or regions. Contour-based methods represent the shape outline either using a set of points on the contour or approximate the curve using a function such as the level set function. Region-based methods may partition a shape into simpler forms (such as polygons), approximate the shape using a bounding region (such as a bounding rectangle or convex hull), or represent internal features of the shape (e.g., a skeleton). Transform-based representations decompose a shape into one-dimensional (1D) or two-dimensional (2D) signals (for example Fourier transform and wavelet transforms are linear transforms, while the Hough transform is a nonlinear transform). Transform-domain descriptors of shape can be transform coefficients or transform energy. Here, the shape of the prostate has been represented using a contour because it can be easily deformed to represent a flexible boundary.

When an object is enlarged, rotated or moved it is still recognizable by a human. This property of an object is called pose invariance; that is, the object is identifiable from a different angle or position, or at a different scale. The pose of an object in an image can be changed using an affine transform (section 4). Pose is a relative concept and it is usually calculated with respect to the pose of another similar object on an image. It can be estimated by deriving the parameters of the affine transform needed to match the two shapes (Ünsalan, 2007).

Level set method of segmentation

Deformable contour models or active contour models are shape-based procedures in which a closed contour deforms by minimizing an energy function. This energy function incorporates low-level visual properties of an object such as edges or pixel intensity, and/or object-level features such as curvature of the object and size. In the level set method introduced by Sethian (1999) the evolving boundary (interface) is represented implicitly as the zero iso-contour of some function. For example, the zero iso-contour of $\phi(s) = x^2 + y^2 - 1$, is given by the unit circle $\phi(s) = 0$. In this framework, the equation of motion of the interface is defined using a simple convection equation such as (Osher & Fedkiw, 2002):

$$\frac{d\phi}{dt} + V \cdot \nabla \phi = 0. \quad (1)$$

Here, $V = (u, v, w)$ is the velocity field (u, v, w are components of the velocity field in the x, y and z directions respectively), and ∇ is the spatial gradient operator. This equation is referred to as the level set function. The level set function can be defined in terms of the signed distance function. The signed distance function is an implicit function that 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 the pixels inside have negative distance values assigned to them. The zero level set is defined as the set of all points whose distance to the interface is zero. The level set update equation is derived by discretizing the level set equation using the forward Euler time discretization given by:

$$\frac{\phi(n+1) - \phi(n)}{\Delta t} + \vec{V}(n) \cdot \nabla \phi(n) = 0. \quad (2)$$

The spatial derivative terms in equation (1) can be expanded as:

$$\frac{\phi(n+1) - \phi(n)}{\Delta t} + u(n)\phi_x(n) + v(n)\phi_y(n) + w(n)\phi_z(n) = 0. \quad (3)$$

The upwind differencing used for the spatial derivative terms along with the forward Euler time discretization makes the level set update stable. This guarantees that small approximation errors are not amplified with time.

Level set methods have been used extensively for medical image segmentation. Some of the popular methods are by Leventon et al. (2000), Tsai et al. (2003), and Chan & Vese (2001). Leventon et al. (2000) introduced the principal component analysis (PCA) of shapes in the signed distance function domain. They used Bayesian techniques to generate maximum *a posteriori* estimates of pose and shape from the prior knowledge of shapes. Shape priors have also been used with active contour-based image segmentation by Etyngier et al. (2007). They used diffusion maps to model shapes as a finite-dimensional manifold. Their segmentation results were accurate but the initial contour

was placed manually in the images. Chan & Vese (2001) introduced a region-based energy function based on Mumford-Shah segmentation techniques to detect features with diffuse boundaries. However, their algorithm used intensity average values for detecting objects and they mention that textural features need to be incorporated into the level set framework to perform generalized Mumford-Shah segmentation. Tsai et al. (2003) derived a shape-based level set function as the sum of mean shape and shape variability. Tsai et al. also incorporated pose into their level set function. They optimized the parameters of this function to produce a good model of the object shape based on the knowledge of mean shape and shape deviations from the training data. The level set function derived by Tsai et al. has been adopted by the GA developed here.

Optimization using a GA

Genetic algorithms (Mitchell, 1996) simulate the process of biological evolution using selection, crossover, and mutation. GAs have been used for a variety of image processing applications, such as image segmentation (Poli & Cagnoni, 1997), feature extraction from remotely sensed images (Daida et al., 1996), and medical feature extraction (Harvey et al., 2003). In contrast with traditional optimization methods, a GA uses a stochastic parallel search to reach the optimum solution and so is less likely to be stuck in a local maximum. Individuals of the GA are candidate solutions and are typically encoded as bit strings or vectors defined based on the application. The GA searches the space of candidate solution to identify the best solution for the problem at hand. Individuals of the GA are candidate solutions and are typically encoded as bit strings or vectors whose interpretation depends on the application. The GA searches the space of candidate solutions to identify the best (or at least an adequate solution) solution for the problem at hand. A fitness function is used to evaluate individuals and compare them based on a fitness score. This fitness score is used in the selection process to determine which individuals get to produce an offspring and propagate their “genes” (bits/vector elements) to future generations.

Selection can be performed in a number of different ways. Some of the popular methods are rank selection, fitness proportionate selection, and tournament selection. In rank selection, candidate solutions are sorted according to their fitness score and higher ranked individuals are more likely to be chosen for crossover than lower ranked individuals. In fitness proportionate selection, the probability of an individual for being selected is given by the ratio of its fitness to the fitness of other members of the population. In tournament selection, two individuals are first chosen randomly from the current population. One of the two individuals is then selected probabilistically, based on fitness.

Genetic algorithms often suffer from premature convergence. This occurs when some individuals in the population are much more fit than others and are the only ones selected for producing the future generations thereby resulting in the reduction of diversity of the population over successive generations of the GA. This can slow the performance of the GA. Selection procedures such as tournament selection and rank selection can be used instead of fitness proportionate selection to avoid premature convergence. After selecting

two individuals from the current population, the crossover operator is applied to produce two new offsprings that are members of the next generation. Crossover is performed using a crossover mask which swaps same length segments of genes between two parents to produce the offspring. The mutation operator chooses a single gene at random and changes its value. The new population thus generated is evaluated using the fitness function and the process of selection, crossover, and mutation is repeated until an offspring with an acceptable fitness value is produced.

Genetic algorithms have been used for segmenting medical images by Cagnoni et al. (1999). The GA developed by Cagnoni et al. optimized the parameters of an elastic contour model called “snakes” (Kass et al., 1988) and minimized an energy function for curve evolution. The “snakes” algorithm used an energy function based on low-level features like smoothness of the curve, curvature and image gradient. In contrast, the GA framework here allows the use of different kinds of features such as texture and shape for deriving the fitness of GA individuals. Another recent work using GAs for medical image segmentation is by Chabrier et al. (2008). They use a GA to find the optimal combination of information extracted from several different segmentation algorithms.

PROBLEM DESCRIPTION: PROSTATE SEGMENTATION

The prostate gland is a male reproductive organ located below the bladder and in front of the rectum and is about 3 cm in length along the height of the body. Prostate cancer is the most commonly diagnosed malignancy in men over the age of 50. When diagnosed at an early stage, the disease is curable, and even at later stages treatment can be effective. Nevertheless, treatment options vary depending on the extent of the cancer, and prognosis worsens when diagnosis occurs at an advanced stage. External beam radiotherapy is a well established standard treatment of prostate cancer. Intensity-modulated radiotherapy is used as a technique to improve tumor control and reduce radiotherapy-associated toxicity. Margin reduction around the clinical target volume (CTV) is essential to reduce the irradiated volumes of the organs at risk (bladder, rectum) and reduce toxicity. Image-guided radiotherapy (IGRT) allows margin reduction due to prostate localization before each radiotherapy fraction. A key ingredient for optimal treatment of prostate cancer is target segmentation on medical images used for treatment planning. Traditionally images are manually segmented by a radiologist or radiation oncologist to localize the prostate gland within the pelvic anatomy prior to treatment planning. Nevertheless, manual segmentation has its limitations particularly due to inter-observer and intra-observer segmentation variability. Furthermore, with the rise in the implementation of adaptive radiotherapy, manual segmentation has become a tedious process, necessitating a reliable automated process for expediting prostate localization on CT and/or MRI images. The GA developed here aims at automating the segmentation process by incorporating shape and textural priors into a single framework.

The almond-shaped prostate gland can be deformed by bladder and rectal filling. In addition to this, the size of prostate can vary considerably across patients making automatic segmentation a challenging problem. Figure 1 shows single slices of a pelvic

CT scan and MRI of two patients. The contour in the center was marked by a radiologist as the prostate. The organ just below the prostate contour is the rectum. The large structures around the prostate (white on CT images shown in the top panels of figure 1) are the bones. One can see that the edges near the boundary of the prostate that was marked by the radiologist are not prominent. On CT images, the radiologist adjusts the contrast of the images before performing manual segmentation on them. The MRI images have better soft-tissue contrast therefore no contrast adjustment is needed on these images. These contours are stacked on top of one another to create the three dimensional (3D) shape of the prostate.




Figure 1. The 2D pelvic CT scan (upper left). The white contour (upper right) in the center is the prostate outlined by an expert. The black region just below the prostate is the rectum. The white structures surrounding the prostate are the bones. The 2D pelvic MRI scan (lower left) and manually segmented contour (upper right).

SEGMENTATION USING THE GENETIC ALGORITHM

The GA performs segmentation in two stages. In the first stage, termed training stage, the texture, mean position and shape variability of the prostate are derived from training

images. The latter are obtained from manual segmentation of the dataset. In the second stage, termed segmentation, the GA generates candidate segmenting contours and evaluates their performance using the information derived in the training stage. Better solutions are allowed to propagate to future generations using selection, crossover and mutation to produce a new generation. The process is iterated until a stopping criterion is satisfied. The process of selection crossover and mutation is repeated to get the final segmentation result. The flowchart in figure 2 gives an overview of the algorithmic framework.




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

Deriving priors

Shape representation using a level set function

To derive information from the training data, the manually drawn contours/shapes are first aligned. Then the mean shape and shape variability is derived from this dataset to create a level set-based shape representation (adopted from Tsai et al., 2001) for representing segmenting contours. They are derived using principal component analysis (PCA) on the manually segmented images as follows. The manually drawn contours from the training data are first converted into signed distance functions, ψ_i ($i = 1$ to n , where n denotes the number of training contours). The level set function is a linear combination of the mean shape and weighted shape variances in the signed-distance domain. The mean shape is defined for n contours as:

$$\overline{\Phi} = \left(\frac{1}{n}\right) \sum_{i=1}^n \psi_i. \quad (4)$$

The mean offset functions are then derived by subtracting the mean from each of the training contours in signed-distance domain ($\tilde{\psi}_i = \psi_i - \overline{\Phi}$). The columns of the mean offset functions (size $N = N_1 \times N_2$ same as the training images) are then serially stacked to form one column vector (β_i) of size $1 \times N$. The shape variability matrix S (size $N \times n$) is obtained from n such column vectors

$$S = [\beta_1, \beta_2, \dots, \beta_n]. \quad (5)$$

The variance in shape is then computed by an eigen value decomposition on this shape variability matrix

$$\frac{1}{n} S S^T = U \Sigma U^T. \quad (6)$$

Here U is an $N \times n$ matrix whose columns represent n orthogonal modes of shape variation and Σ is an $n \times n$ diagonal matrix of eigen values. By rearranging the columns of U to form an $N_1 \times N_2$ structure, the n different eigen shapes can be obtained $\{\Phi_1, \Phi_2, \dots, \Phi_n\}$. The mean shape and shape variability derived from the training phase are used to define a level set function that implicitly represents the segmenting curve,

$$\Phi[\mathbf{w}] = \overline{\Phi} + \sum_{j=1}^k w_j \Phi_j. \quad (7)$$

Pose is inserted into this function using an affine transform. The affine transform is the product of three matrices, the translation matrix, the scaling matrix and the rotation matrix respectively

$$\begin{bmatrix} \tilde{x} \\ \tilde{y} \\ 1 \end{bmatrix} = \begin{bmatrix} 1 & 0 & a \\ 0 & 1 & b \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} h & 0 & 0 \\ 0 & h & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} \cos \theta & -\sin \theta & 0 \\ \sin \theta & \cos \theta & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ 1 \end{bmatrix}. \quad (8)$$

Here, x and y are the pixel coordinates of the input image and $\tilde{x}, \tilde{y}$ are the pixel coordinates of the affine transformed image. The new level set function is given by $\Phi[\mathbf{w}, \mathbf{p}]$ where $\mathbf{p} = [a, b, h, \theta]$ are the x, y translation (a, b), scale (h) and rotation (θ) parameters respectively. The zero-level of this level set function gives the segmenting contour and its parameters $[\mathbf{w}, \mathbf{p}]$ are evolved by the GA

$$\Phi[\mathbf{w}, \mathbf{p}](x, y) = \overline{\Phi}(\tilde{x}, \tilde{y}) + \sum_{j=1}^k w_j \Phi_j(\tilde{x}, \tilde{y}). \quad (9)$$

Before deriving the mean shapes and shape variance from the training data, the images are aligned for pose. Gradient descent is used to minimize the difference between a pair

of binary images with respect to their pose parameters. The transformed image based on pose is given by:

$$\tilde{I} = T[p] * I. \quad (10)$$

where $T[p]$ is the two-dimensional transformation matrix i.e., the product of the translation, scaling and rotation matrices of equation (8). The energy functional used to minimize the difference between two images is given by:

$$E_{align} = \frac{\iint_{\Omega} (I^i - I^j)^2}{\iint_{\Omega} (I^i + I^j)^2}, i \neq j. \quad (11)$$

Here, Ω is the image domain. The area normalizing term in the denominator improves the cost function by preventing the images from shrinking. The gradient with respect to the pose parameters p_i for any image i is given by:

$$\nabla_{p_i} E_{align} = \frac{2 \iint_{\Omega} (\tilde{I}^i - \tilde{I}^j) \nabla_{p_i} I^i dA}{\iint_{\Omega} (I^i - I^j)^2 dA} - \frac{2 \iint_{\Omega} (\tilde{I}^i - \tilde{I}^j)^2 dA \iint_{\Omega} (I^i + I^j) \nabla_{p_i} I^i dA}{\left(\iint_{\Omega} (I^i + I^j)^2 dA \right)^2} \quad (12)$$

$\nabla_{p_i} \tilde{I}^i$ is the gradient of the transformed image taken with respect to the pose parameters:

$$\nabla_{p_i} \tilde{I}^i(\tilde{x}, \tilde{y}) = \begin{bmatrix} \frac{\partial I^i(\tilde{x}, \tilde{y})}{\partial x} & \frac{\partial I^i(\tilde{x}, \tilde{y})}{\partial y} & 0 \end{bmatrix} \frac{\partial T[p_i]}{\partial p_i} \begin{bmatrix} x \\ y \\ 1 \end{bmatrix}, \quad (13)$$

where the matrix derivatives of $T[p]$ are taken component-wise. The initial pose parameters of one of the shapes are kept fixed and the pose parameters of the second image are calculated to solve the problem. Note that homogeneous coordinate system is used here. The homogeneous coordinate system maps each point (x', y') in the Euclidean space to $[x, y, w]$ ($w \neq 0$). The mapping is achieved by the relation $x' = x/w$ and $y' = y/w$. Using homogeneous coordinate system allows the translation operation to be represented with a matrix multiplication.

Deriving textures

The textural priors are *high-level* feature vectors derived from the training image using Laws' textural measures and Gabor wavelet transform-based features. The Laws' texture measures are computed by convolving the training images with small integer coefficient masks. The basic one-dimensional convolution kernels (usually 5x5) 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. To generate the texture energy planes, the training images are first convolved with the each of the 25 two-dimensional masks to obtain 25 grayscale

images. When an image is convolved with a mask, the pixel values in the image around 5x5 windows are weighed by the mask parameters thereby enhancing certain features in the image. For example, edges in an image can be enhanced by convolving it with the mask $E5 \times E5^T$ which enhances the value of each edge pixel by differently weighting its neighboring pixels in 5x5 windows. Thus pixel intensity differences in the image are amplified making the edges more prominent. A small (15x15) window is then operated on these grayscale images by summing the absolute values of the 225 neighborhood pixels to produce 25 different texture energy planes/maps. Fisher's linear discriminant (FLD) is used to find the weights to linearly combine the 25 texture energy planes and threshold it to obtain a binary image with the desired classification (Figure 3). These weights are saved and used to derive the texture energy planes on test images. Fisher's linear discriminant is a simple dimensionality reduction approach in which a multi-dimensional data x is projected onto a one-dimensional space y such that

$$y = v^T x \quad (14)$$

Although projecting multi-dimensional data along one dimension leads to a loss of information due to significant overlap, the class separation in one dimension can be maximized by adjusting the weight vector v . This weight vector is a function of the projected class means normalized by the within-class scatter along the direction of v . For, a two-class problem the Fisher criterion is:

$$J(v) = \frac{(m_2 - m_1)^2}{s_1^2 + s_2^2} \quad (15)$$

$$s_k^2 = \sum_{n \in C_k} (y^n - m_k)^2 \quad k = 1, 2$$

where m_k is the class mean and s_k is the within class scatter. Maximizing the function J maximizes the class separation and minimizes the within-class scatter. The weight vector v is the derived texture prior.




Figure 3. Flowchart depicting the steps in the segmentation process using Laws' texture measures

Gabor wavelets are based on the Gabor elementary function given by the modulation of the Gaussian with a complex exponential function (equation 16).

$$h(x,y) = g(x,y) \exp \left[j2\pi \left(U_x + V_y \right) \right]$$

$$g(x,y) = \frac{1}{2\pi\sigma_x\sigma_y} \exp \left\{ -\frac{1}{2} \left[\left(\frac{x}{\sigma_x} \right)^2 + \left(\frac{y}{\sigma_y} \right)^2 \right] \right\}. \quad (16)$$

Gabor wavelets are derived by considering $h(x, y)$ as the mother wavelet and a family of functions is obtained by translations and dilations of this mother wavelet. The method of Gabor wavelets assumes that local texture regions are spatially homogeneous and the mean and standard deviation of the transform coefficients are used to represent regions for classification. The Gabor wavelet is given by equation 17. Here, k is the number of orientations.

$$h_{mn}(x, y) = aH(x', y')$$

$$a > 1$$

$$m, n = \text{integers} \quad (17)$$

$$\text{and } x' = a^{-m} \left(x \cos \frac{n\pi}{k} + y \sin \frac{n\pi}{k} \right)$$

$$y' = a^{-m} \left(-x \sin \frac{n\pi}{k} + y \cos \frac{n\pi}{k} \right)$$

Given an image $I(x, y)$, the Gabor wavelet transform is given by equation (18).

$$W_{mn}(x, y) = \int I(x_1, y_1) h_{mn}^*(x - x_1, y - y_1) dx_1 dy_1 \quad (18)$$

Here, h^* is the complex conjugate of h . The mean and standard deviation of transform coefficients for each known region are the derived texture priors.

Segmentation procedure

An individual (I) of the GA population is defined as a vector of weight ($\pm 2\sigma$ variability of the mean shape) and pose parameters of the level set function

$$I = [w_1, w_2, a, b, h, \theta] \quad (19)$$

Several such individuals form the GA population. When the weights and pose parameters of each individual are substituted in the level set equation (9), a new shape combining the mean shape, the eigen shapes and a new pose is produced. This new shape is the segmenting contour. Thus, each GA individual is a segmenting contour. The fitness of an

individual is measured by comparing the textural difference between the regions inside and outside the contour. The textural difference is computed as follows. At first a binary image is generated using texture segmentation on a test image. In this binary image, pixels labeled '1' represent one texture type (regions texturally similar to the prostate) and pixels labeled '0' represent all other texture types (regions not texturally similar to the prostate on pelvic CT/MR images). The number of '1's and '0's inside and outside the candidate segmenting contour are counted. The fitness (F) is defined as a function of this count and is given by

$$F = 500 \times (A + (1-B)). \quad (20)$$

Here, A denotes the detection rate: the fraction of pixels labeled '1' ('0') that are inside (outside) the candidate contour. Thus, A is the sum of true positives and true negatives. Conversely, B is the sum of false positives and false negatives. B denotes the false alarm rate: the fraction of pixels labeled '1' ('0') that are outside (inside) the candidate contour. A higher fitness score means that more pixels inside the contour belong to the desired texture type that was derived from the training data.

Individuals with a higher fitness score are selected for crossover and mutation to propagate their genes to future generations. Crossover is performed by swapping segments of genes between two individuals. For example, let $I_1 = [w_1, w_2, a, b, h, \theta]$ and $I_2 = [w'_1, w'_2, a', b', h', \theta']$. The following two offsprings are produced by crossover at midpoint between I_1 and I_2 : $[w_1, w_2, a, b', h', \theta']$ and $[w'_1, w'_2, a', b, h, \theta]$. Mutation is performed by randomly changing the value of a gene from a specified range of values (For example, the gene θ can take on values between 0-360°). A crossover/mutation rate (usually specified as a percentage) determines the probability of crossover/mutation between two individuals. Selection, crossover and mutation are performed iteratively for successive generations until an individual with a specified fitness value is attained or the GA runs for a specified number of generations.

Extension to three dimensions

The GA is extended to three-dimensions by using 3D pose parameters; x, y, z translation (a, b, c), scale (h), yaw (α), pitch (β) and roll (θ).

$$\begin{bmatrix} \tilde{x} \\ \tilde{y} \\ \tilde{z} \\ 1 \end{bmatrix} = \begin{bmatrix} 1 & 0 & 0 & a \\ 0 & 1 & 0 & b \\ 0 & 0 & 1 & c \\ 0 & 0 & 0 & 1 \end{bmatrix} * \begin{bmatrix} h & 0 & 0 & 0 \\ 0 & h & 0 & 0 \\ 0 & 0 & h & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} * R_x * R_y * R_z * \begin{bmatrix} x \\ y \\ z \\ 1 \end{bmatrix} \quad (21)$$

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

$$R_x = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & \cos(\alpha) & -\sin(\alpha) & 0 \\ 0 & \sin(\alpha) & \cos(\alpha) & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \quad (22)$$

$$R_y = \begin{bmatrix} \cos(\beta) & 0 & \sin(\beta) & 0 \\ 0 & 1 & 0 & 0 \\ -\sin(\beta) & 0 & \cos(\beta) & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \quad (23)$$

$$R_z = \begin{bmatrix} \cos(\theta) & -\sin(\theta) & 0 & 0 \\ \sin(\theta) & \cos(\theta) & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \quad (24)$$

The individuals of the 3D-GA 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 the slices of the CT/MR scans from the training data. The 3D segmenting contours generated by this GA segments all the slices of the test image at once.

Performance Evaluation

The performance of the algorithm has been evaluated using two measures taken from Udupa et al. (2006). These measures are precision and validity. Precision is defined as the repeatability of the segmentation algorithm. Validity is defined as the closeness of the segmentation outcome with the truth/manual segmentation. Pixel-based (texture feature finding) segmentation methods with which the LSGA has been compared are the Laws' texture method and the Gabor wavelet method. The contour-based segmentation method of Chan & Vese (2001) has also been compared with the LSGA. The application domain for performing the evaluation is given by a pair of task and segmentation protocols, $\langle T, P \rangle$ used to solve the task. Here, the possible values for T and P are:

1. T : Segmenting the prostate on pelvic CT/MRI images.
2. P : LSGA/ Gabor Wavelets/ Law's texture method/ Chan & Vese method

The ground truth is derived from the manual segmentations. Manual segmentation is variable due to intra- and inter-operator variability. The ground truth is therefore obtained from manual delineations by averaging over multiple manual delineations. The precision (repeatability) of a given segmentation method is determined by applying the same algorithm twice on the same image and then comparing the binary outcomes. The $(\cap)$ operator signifies the region common to the two binary outcomes. The $(\cup)$ operator signifies the union of the two binary outcomes

$$\text{Precision(PR)} = \frac{\text{Outcome1} \cap \text{Outcome2}}{\text{Outcome1} \cup \text{Outcome2}}. \quad (25)$$

The validity of the segmentation result is obtained by comparing the segmentation outcome with the ground truth. Four features are derived this way; True Positives (TP), True Negatives (TN), False Positives (FP), False Negatives (FN) as

$$\text{True Positive (TP)} = \frac{|\text{'1's in the binary outcome matching the '1's in manual segmentation}|}{|\text{'1's in the manual segmentation}|} \quad (26)$$

$$\text{True Negative (TN)} = \frac{|\text{'0's in the binary outcome matching the '0's in manual segmentation}|}{|\text{'0's in the manual segmentation}|} \quad (27)$$

$$\text{False Positive (FP)} = \frac{|\text{'1's in the binary outcome matching '0's in the manual segmentation}|}{|\text{'0's in the manual segmentation}|} \quad (28)$$

$$\text{False Negative (FN)} = \frac{|\text{'0's in the binary outcome matching the '1's in the manual segmentation}|}{|\text{'1's in the manual segmentation}|} \quad (29)$$

RESULTS AND DISCUSSION

To perform this analysis, images were obtained from a database of 2700 pelvic CT and MRI scans, acquired through collaboration with Oregon Health & Science University (OHSU). CT and MRI images of 10 patients (for each modality) from this database were manually segmented by Dr. Arthur Hung and Dr. James Tanyi, (Dept. of Radiation Medicine, OHSU). Each scan for a patient contains 15-20 slices of 2D images. Some patients have multiple CT/MRI scans. 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 number of slices on which the prostate is visible depends on the resolution of the scans which varies from one patient to another, and also dependent on the technique used for image acquisition. The prostate has been manually delineated twice on the same set of images by the experts. This provides a database for intra-operator variability. The manually segmented contours derived from the scans of five patients (~ 50x2 =100 images) have been used as the training data for this analysis. The images from the other five patients (for which the ground truth was available) were used as test images. Note that the CT and MR images shown here are taken from two different sets of patients.

For the same set of test images the segmentation outcomes using the four different methods were computed. At first the level set based segmentation algorithm of Chan & Vese (2001) was tried on the test images. The initial contour was placed in the center of each test image (left panels of figure 4). The optimization method used for curve evolution in this method is the gradient descent method. The upper right panel on figure 4 shows the outcome of the algorithm for a CT image. The CT image was contrast-stretched because the original image had very low contrast. Figure 4 (lower right panel)

shows the result of the algorithm on a MRI image. In both the cases the algorithm found boundaries between the dark and light regions. The result that was obtained was expected because the algorithm is designed to find regions with markedly different pixel intensity values inside and outside the contour. The result obtained in this case is not a binary image. It is the original image superimposed by the evolved curve. Since, the original curve divides into many curves surrounding several regions, it is not possible to obtain a binary image from this outcome. Since the evaluation of the algorithm performance requires a binary image, the evaluation cannot be performed for the outcome of this algorithm. Therefore, NA or (Not Applicable) is mentioned in the evaluation results tables 2 and 3 for the corresponding values.




Figure 4. (Upper left panel) Initial contour placed on top of a test CT image. (Upper right panel) Outcome of the level set algorithm of Chan & Vese. (Lower left panel) Initial contour placed on top of a test MR image. (Lower right panel) Outcome of the level set algorithm of Chan & Vese.

Figure 5 (middle panels) shows the segmentation output generated by the Laws' texture segmentation method on a test CT and MR image. For the CT images the result was again computed on contrast-stretched CT images. The figure shows the classification of pixels as similar to prostate (white) or otherwise (black). Figures 5 (lower panels) shows the outcome of the Gabor wavelet-based segmentation algorithm applied on the same test CT (left) and MR image (right panel). Both of these methods found regions texturally similar to the prostate, but also marked some regions outside the prostate as similar to the prostate. The GA used this texture based classification for its fitness

determination. It used the mean position information derived from training images to place the segmenting contour in the prostate region and then optimized its pose and location in the prostate region using the texture information derived using these methods.




Figure 5. Test CT and MR images (upper left and right panels). Laws' texture classification method applied to the test CT image (middle left panel). Laws' texture classification method applied to the test MR image (middle right panel). Gabor wavelet segmentation algorithm applied to the test CT (lower left) and MR (lower right) images.

The parameters used by the GA are shown in table 1. Before computing PCA on the training contours, they were aligned with respect to pose with one training contour. Figure 6 shows the mean shape and shape variability of the prostate derived in 2D from the training images. The initial population of the GA was generated using values chosen randomly from the range of parameter values specified in table 1. These parameter values when substituted into the level set equation 9, generated segmenting contours as a sum of the mean shape and shape variability information derived from the training images. The evolution of the curve (selection, cross-over and mutation over 100 generations) was guided by the fitness derived from textural priors generated by Laws' texture method since it narrowed the region of interest more than the Gabor wavelet-based method. The GA used the textural classification map (around the mean location of the prostate) on a test image, to place the final the segmenting contour.

Population Size	50
Mutation Rate	10 % per gene
Crossover Rate	50% single-point
Selection Criteria	Rank Selection
Weights for eigen shapes, w	$(0 \pm 2\sigma)$
Translation parameters a, b	Integer (0-30)
Rotation parameter θ	-90° to $+90^\circ$
Scale parameter h	(0.5-2)

Table 1. LSGA Parameters




Figure 6. Mean shape of the prostate derived from training images (left panel). Shape variability of the prostate derived from training images (right panel).

The segmentation result from the LSGA in 3D is shown in figure 7 (right panel). Figure 7 (left panel) shows a slice from the 3D segmentation result obtained on the test CT image. The highest-fitness 3D individual evolved by the GA was found to be $[1170(w_1), 161(w_2), 9(a), 2(b), 4(c), 1.3(h), 25(\alpha), 13(\beta), 24(\theta)]$ and its fitness was 316. Figure 8 shows a slice from the 3D segmentation result obtained on a test MR image. The highest-fitness 3D individual evolved by the GA was found to be $[3.99e+003(w_1), 4.7e+003(w_2), 8(a), 27(b), 2(c), 1.9(h), 4(\alpha), 14(\beta), 6(\theta)]$ and its fitness was 526. There were only five slices in the 3D MR image as compared to ten slices in the 3D CT image, therefore the 3D structure of the shapes look different. Note that the fitness values found are low because it was derived from the texture segmentation of the entire image and not just the prostate region. For achieving more accurate fitness values only the textural segmentation around the prostate region could be used along with the information of the relative position of the prostate with respect to the bladder and the rectum. This will be the subject of future work.

Tables 2 and 3 show the performance evaluation by comparing the binary results obtained from the four different methods. The abbreviations LSGA, LTM, GW, and CV are for the genetic algorithm developed here, the Laws' texture method, Gabor Wavelet based segmentation method and the Chan & Vese method respectively. As mentioned before the outcome of the Chan & Vese algorithm is not a binary image and therefore it cannot be evaluated using the measures used here. For the other methods, the precision, and validity measures (TP, TN, FP, FN) were derived for all the test images. The segmentation result of the GA in 3D has been used for evaluation purposes.

Protocol	Precision	TP	FP	TN	FN
LSGA	0.42	0.67	0.007	0.99	0.32
LTM	1	0.5	0.07	0.93	0.5
GW	1	0.44	0.08	0.92	0.56
CV	NA	NA	NA	NA	NA

Table 2. Performance evaluation of the four protocols for segmenting the prostate on pelvic CT images.

Protocol	Precision	TP	FP	TN	FN
LSGA	0.4	0.73	0.005	0.99	0.28
LTM	1	0.44	0.24	0.76	0.56
GW	1	0.79	0.42	0.58	0.21
CV	NA	NA	NA	NA	NA

Table 3. Performance evaluation of the four protocols for segmenting the prostate on pelvic MR images.




Figure 7. 3D segmentation result of the LSGA on a test CT image (right panel). A slice of the 3D segmentation generated by the GA (left panel).




Figure 8. 3D segmentation result of the LSGA on a test MR image (right panel). A slice of the 3D segmentation generated by the GA (left panel). There are only 5 slices in this figure as compared to 10 in the previous figure, therefore the 3D structures look different.

Significance of this work

This algorithm is a step towards automating the complex task of prostate segmentation. The GA allows multiple types of features to be used for segmentation: texture, shape and mean location of prostate for initial contour placement. The GA provides a means for derivative-free curve evolution for performing segmentation. This brings flexibility to use *high-level* features such as the combination of textures, shapes and location information for curve evolution. The GA developed here is not specific to prostate segmentation and can be used in other domains and applications involving multiple feature based segmentation.

FUTURE RESEARCH DIRECTIONS

There is a considerable scope for improvement in this GA framework. Currently the initial position of the contour is derived from the mean position of the contour from training images. The desired method of automatically placing the initial contour is to encode the information of relative position of the organs in the pelvic region image into

the GA framework. Some algorithms used for finding spatial relations between objects are Bayesian network classifiers (Pham & Smeulders, 2006), graph matching (Bunke, 2000) and histograms of distances and angles (Miyajima & Ralescu, 1994). Relative position information would be incorporated into the GA framework in future not only for initial contour placement but also in the fitness function to determine the fitness of a segmenting contour.

Also, a linear classifier has been used for texture based classification of pixels. There are many other methods of classification such as support vector machines (SVMs), neural networks, etc. that can be tried in future to improve the classification outcome of the texture-based segmentation. Tables 2 and 3 show that the precision of the GA could not achieve high precision on these images. Perfect precision is impossible to achieve for this dataset because there are no clearly visible edges to drive an optimization algorithm to the exact boundary of the prostate. However, using the texture classification outcome only from the prostate region and surrounding organs, for deriving fitness may improve the precision of the segmentation outcome.

Another major shortcoming of the GA based optimization is its speed of execution. Currently the 3D LSGA takes about 15 minutes to a few hours for each set of test images to find the segmentation outcome. The speed depends on a number of factors such as, the number of generations of the GA run, the population size, etc. Methods to speed up the GA such as parallel GA (Luque et al., 2005) would be explored for finding the practical viability of the algorithm.

CONCLUSION

A genetic algorithm framework has been applied to level-set curve evolution to incorporate multiple features for performing the complex segmentation task of prostate segmentation on pelvic CT and MR 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. Complete automation can be achieved when the knowledge of organ locations with respect to the prostate is incorporated into this framework.

ACKNOWLEDGEMENTS

We are thankful to Intel Corporation and the J.S. McDonnell Foundation for funding this research.

REFERENCES

- Ahmadian, A., & Mostafa, A. (2003). An efficient texture classification algorithm using gabor wavelet. In *Proceedings of the 25th Annual International Conference of the IEEE EMBS* (pp. 930-933). Washington, DC: IEEE Press.
- Bunke, H. (2000). Graph matching for visual object recognition. *Spatial Vision*, 13, 335-340.
- Cagnoni, A., Dobrzeniecki, A. B., Poli, R., & Yanch, J. C. (1999). Genetic algorithm-based interactive segmentation of 3D medical images. *Image and Vision Computing*, 17(12), 881-895.
- Chabrier, S., Rosenburger, C., Emile, B., & Laurent, H. (2008). Optimization based image segmentation by genetic algorithms. *EURASIP Journal on Video and Image Processing*, 2008 (February, 2008), 1-23.
- Chan, T., & Vese, L. (2001). Active contours without edges. *IEEE Trans. on Image Proc.*, 10, 266-277.
- Costa, L. F., & Cesar, R. M., Jr. (2001). *Shape analysis and classification theory and practice*. Boca Raton, FL: CRC Press.
- Daida, J.M., Hommes, J. D., Bersano-Begey, T.F., Ross, S. J., & Vesecky, J. F. (1996). Algorithm discovery using the genetic programming paradigm: Extracting low-contrast curvilinear features from SAR images of Arctic ice. In P. J. Angeline & K. E. Kinnear Jr. (Eds.), *Advances in Genetic Programming 2* (pp. 417-442). Cambridge, MA: MIT Press.
- Etyngier, P., Segonne, F., & Keriven R. (2007). Active-contour-based image segmentation using machine learning techniques. In *Proceedings of the Medical Image Computing and Computer-Assisted Intervention Conference* (pp. 891-900). Heidelberg, Germany: Springer.
- Harvey, N., Levenson, R. M., & Rimm, D. L. (2003). Investigation of automated feature extraction techniques for applications in cancer detection from multi-spectral histopathology images. *Proceedings of SPIE*, 5032, 557-566.
- He, L., Peng, Z., Everding, B., Wang, X., Han, C. Y., Weiss, K. L., & Wee, W. G. (2008). A comparative study of deformable contour methods on medical image segmentation. *Image and Vision Computing*, 26(2), 141-163.
- Kass, M., Witkin, A., & Terzopoulos, D. (1988). Snakes: Active contour models. *International Journal of Computer Vision*, 1, 321-331.
- Laws, K. I. (1980). Rapid texture identification. *SPIE Image Processing for Missile Guidance*, 238, 376-380.
- Leventon, M., Grimson, E., & Faugeras, O. (2000). Statistical shape influence in geodesic active contours. In *Proc. of IEEE Conference on Computer Vision and Pattern Recognition* (Vol. 1, pp. 316-323). Washington, DC: IEEE Press.

- Luque, G., Alba, E., Dorronsoro, B. (2005). Parallel genetic algorithms. In E. Alba (Ed.), *Parallel metaheuristics: A new class of algorithms* (pp. 107-126). Hoboken, NJ: John Wiley & Sons Inc.
- McInerney, T., & Terzopoulos, D. (1996). Deformable models in medical image analysis: A survey. *Medical image analysis*, 1(2), 91-108.
- Mitchell, M. (1996). *An introduction to genetic algorithms*. Cambridge, MA: MIT Press.
- Mitchell, T. M. (1997). *Machine learning*. Singapore: McGraw-Hill International Editions.
- Miyajima, K., & Ralescu, A. (1994). Spatial organization in 2D segmented images: Representation and recognition of primitive spatial relations. *Fuzzy Sets and Systems*, 65(2/3), 225-236.
- Osher, S. J., & Fedkiw, R. P. (2002). *Level set methods and dynamic implicit surfaces*. New York: Springer.
- Pham, T. V., & Smeulders, A. W. M. (2006). Learning spatial relationships in object recognition. *Pattern Recognition Letters*, 27, 1673-1684.
- Poli, R., & Cagioni S. (1997). Genetic programming with user-driven selection: Experiments on the evolution of algorithms for image enhancement. In *Proceedings of the 2nd Annual Conference on Genetic Programming* (pp. 269-277). San Francisco, CA: Morgan Kaufmann.
- Sethian, J. A. (1999). *Level set methods and fast marching methods*. New York, NY: Cambridge University Press.
- Tsai, A., Yezzi, A., Wells, W., Tempany, C., Tucker, D., Fan, A., Grimson, E., & Willsky A. (2003). A shape-based approach to the segmentation of medical imagery using level sets. *IEEE Transactions on Medical Imaging*, 22, 137-154.
- Udupa, J. K., LeBlanc, V. R., Zhuge, Y., Imielinska, C., Schmidt, H., Currie, B. E., Hirsch, L. M., & Woodburn, J. (2006). A framework for evaluating image segmentation. *Computerized Medical Imaging and Graphics*, 30, 75-87.

ADDITIONAL READING

- Bäck, T. (1996). *Evolutionary algorithms in theory and practice*, New York: Oxford University Press.
- Caselles, V., Kimmel, R., & Shapiro, G. (1998). Geodesic active contours. *International Journal of Computer Vision*, 22, 61-79.
- Costa, M. J., Delingette, H., Novellas, S., & Ayache, N. (2007). Automatic Segmentation of Bladder and Prostate Using Coupled 3D Deformable Models. In *Proceedings of the Medical Image Computing and Computer-Assisted Intervention Conference* (Vol. 4791, pp. 252-260). Heidelberg, Germany: Springer.

- Cohen, L. (1991). On active contour models and balloons. *CVGIP: Image Understanding*, 53, 211-218.
- Duncan, J.S., & Ayache, N. (2000). Medical image analysis: progress over two decades and the challenges ahead. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 22(1), 85-106.
- Dunn, D., & Higgins, W. E. (1995). Optimal Gabor filters for texture segmentation. *IEEE Transactions on Image Processing*, 4(7), 947-964.
- Egenhofer, M. (1989). A formal definition of binary topological relationships. In *Proceedings of the 3rd International Conference on Foundations of Data Organization and Algorithms* (Vol. 367, pp. 457-472). Heidelberg, Germany: Springer.
- Ghosh, P., & Mitchell, M. (2008). Prostate segmentation on pelvic CT images using a genetic algorithm. In J. M. Reinhardt & J. P. W. Pluim (Ed.), *Proceedings of the SPIE Medical Imaging Symposium* (Vol. 6914, pp. 691442.1-691442.8). Bellingham, WA: SPIE Press.
- Ghosh, P., & Mitchell, M. (2006). Segmentation of medical images using a genetic algorithm. In *Proceedings of the Genetic and Evolutionary Computation Conference* (pp.1171-1178). New York: ACM.
- Gonzales, R. C., & Woods, R. E. (2002). *Digital image processing*, Upper Saddle River, NJ: Prentice Hall.
- Goldberg, D. E. (1989). *Genetic algorithms in search, optimization, and machine learning*. Reading, MA: Addison-Wesley Publishing Company, Inc.
- Harvey, N.R., Theiler, J., Brumby, S.P., Perkins, S., Szymanski, J. J., Bloch, J. J., Porter, R. B., Galassi, M., & Young, A. C., (2002). Comparison of GENIE and conventional supervised classifiers for multispectral image feature extraction. *IEEE Transactions on Geoscience and Remote Sensing*, 40(2), 393-404.
- Kaiser, G. (1994). *A friendly guide to wavelets*. Cambridge, MA: Birkhäuser Boston.
- Malladi, R., Sethian, J. A., & Vemuri, B. C. (1995). Shape modeling with front propagation: A level set approach. *IEEE Transactions on Pattern Recognition and Machine Intelligence*, 17(2), 158-175.
- Mallat, S. (2008). *Wavelet tour of signal processing* (3rd ed.). St. Louis, MO: Academic Press.
- McBain, C. A., Moore, C. J., Green, M. M. L., Price, G., Sykes, J. S., Amer, A., Khoo, V., & Price, P. (2008). Early clinical evaluation of a novel three-dimensional structure delineation software tool (SCULPTER) for radiotherapy treatment planning. *British Journal of Radiology*, 81(968), 643-652.
- Perkins, S., Theiler, J., Brumby, S. P., Harvey, N. R., Porter, R. B., Szymanski, J. J., & Bloch, J. J. (2000). GENIE: A hybrid genetic algorithm for feature classification in multi-spectral images. In *Proc. of the SPIE 4120 Applications and Science of Neural Networks, Fuzzy Systems, and Evolutionary Computation III* (pp. 52-62).
- Pham, D. L., Xu, C., & Prince, J. L. (2000). Survey of current methods in medical image segmentation. *Annual Review of Biomedical Eng.*, 2, 315-337.

- Popovic, A., Fuente, M., Engelhardt, M., & Radermacher, K. (2007). Statistical validation metric for accuracy assessment in medical image segmentation. *International Journal of Computer Assisted Radiology and Surgery*, 2(3-4), 169-181.
- Sonka, M., Hlavac, V., & Boyle, R. (1993). *Image processing, analysis and machine vision*. London, UK: Chapman & Hall.
- Tremeau, A., & Borel, N. (1997). A region growing and merging algorithm to color segmentation. *Pattern Recognition*, 30(7), 1191-1203.
- Tsai, A., Wells, W., Warfield, S., & Willsky, A. (2005). An EM algorithm for shape classification based on level sets. *Medical Image Analysis*, 9(5), 491-502.
- Ünsalan, C. (2007). A model based approach for pose estimation and rotation invariant object matching. *Pattern Recognition Letters*, 28(1), 49-57.
- Vese, L. A., & Chan, T. F. (2002). A multiphase level set framework for image segmentation using the Mumford and Shah model. *Intl. J. Computer Vision*, 50(3), 271-293.
- Walnut, D. F. (2004). *An introduction to wavelet analysis*. Cambridge, MA: Birkhäuser Boston.
- Wang, Y., & Makedon, F. (2003). R-histogram: quantitative representation of spatial relations for similarity-based image retrieval. In *Proceedings of the eleventh ACM international conference on Multimedia* (pp. 323-326). New York: ACM.