

Genetic Algorithm Driven Statistically Deformed Models for Medical Image Segmentation

Chris McIntosh
cmcintos@sfu.ca

Ghassan Hamarneh
hamarneh@sfu.ca

Medical Image Analysis Lab
School of Computing Science
Simon Fraser University, Burnaby, BC

ABSTRACT

We present a novel evolutionary computing based approach to medical image segmentation. Our method complements the image-pixel integration power of deformable shape models with the high-level control mechanisms of genetic algorithms (GA). Specifically, the GAs alleviate typical deformable model weaknesses pertaining to model initialization, deformation parameter selection, and energy functional local minima through the simultaneous evolution of a large number of models. Furthermore, we constrain the evolution, and thus reduce the size of the search-space, by using statistically-based deformable models whose deformations are intuitive (stretch, bulge, bend) and driven in-terms of principal modes of variation of a learned mean shape. We demonstrate our work through the application to segmentation of the corpus callosum (CC) in mid-sagittal brain magnetic resonance images (MRI).

Categories and Subject Descriptors

I.4.6 [Image Processing and Computer Vision]: Segmentation

General Terms

Algorithms

Keywords

Genetic algorithms, deformable models, segmentation, medical imaging

1. INTRODUCTION

Medical image segmentation remains a daunting task, but one whose solution will allow for the automatic extraction of important structures, organs and diagnostic features, with applications to computer-aided diagnosis, statistical shape analysis, and medical image visualization. Several classifications of segmentation techniques exist including edge, pixel,

and region-based techniques, in addition to clustering, graph theoretic, and model correlation-based approaches [32, 24, 28]. However, the unreliability of traditional purely pixel based methods in the face of shape variation and noise has caused recent trends to focus on incorporating prior-knowledge about the location, intensity, and shape of the target anatomy [13]. One method that has been of particular interest to meeting these requirements is deformable models due to their inherent smoothness properties and ability to fit ill-defined boundaries, compared to say region-growing approaches.

Deformable models for medical image segmentation gained popularity since the 1987 introduction of snakes by Terzopoulos *et al* [29, 16]. In addition to physics-based explicit deformable models [20, 21], geometry-based implicit implementations also attracted attention [22, 3, 26]. Several techniques were proposed to improve segmentation results by controlling model deformations [4, 31]. However, with only smoothness and image-based constraints on their deformations these models were highly susceptible to significant gradient noise, and local minima (Figure 1).

In many applications, prior knowledge about object shape variability is available or can be obtained by studying a training set of shape examples. This allows only feasible deformations to be produced through the incorporation of shape knowledge [7, 8, 17, 30, 11]. However, these methods often used globally calculated statistics which don't generalize to specific locations, and scales on the model. Consequently, the deformations are constrained to be global in nature and, hence, can not adapt to local variations in shape which are often of high interest. Moreover, the deformations themselves are un-intuitively defined and as such it is not clear what deforming along a particular mode of global variation will accomplish (whether it will bulge, stretch, or bend the shape).

However, there are many main issues common to these deformable model-based techniques. The first is that the model still needs to be initialized to some optimal target area of an image, with some shape, orientation, and scale. Secondly, the deformations they exhibit are boundary based, non-intuitive, do not follow the geometry of the object, are not properly spatially constrained, and their extent is not spatially localized (to a location or scale). These problems motivate the use of the hierarchical regional principal component analysis [10] (HRPCA) of a medial based shape representation, in combination with GA driven deformations to obtain: intuitive deformations (via medial-based shape rep.), that are statistically-based and spatial localized (via HRPCA), with automatic parameter setting and robust

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, to 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.

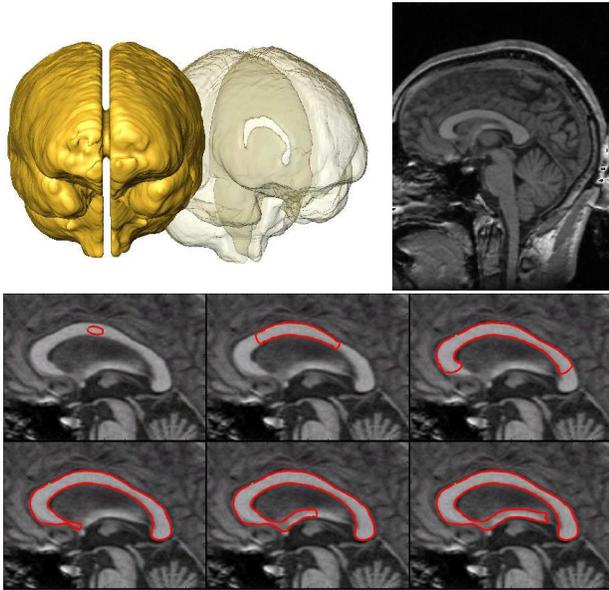


Figure 1: (top) The corpus callosum is the band of nerve fiber tissue connecting the left and right hemispheres of the brain (bottom) Incorrect progress of a snake (red) segmenting the CC in an MRI image. Leaking occurs due to weak edge strength and incorrect parameter setting. Both of these problems are addressed by our proposed method.

model initialization (via GA).

The fitting of deformable models to image data is typically performed through the minimization of a particular energy functional or some higher-level, for example user driven, control process [13].

Energy-functional minimization can be carried out in a variety of ways. One solution is to perform explicit differentiation using the Euler-Lagrange where each new application requiring a modified energy functional must be accompanied by one such derivation. Subsequently, the roots of the derivative (minima/maxima of the functional) can be located using Newton’s method, or some modification thereof. However, as the number of variables increases these methods become to expensive to compute and must, therefore, be approximated. One such technique based on a first-order derivative estimate is gradient descent. Common amongst these root-finding methods is that as the number of dependant variables (shape, location, scale, orientation, etc.) increases as does the complexity of the search space, which often increases the amount of local minima and requires the calculation of an increasingly large number of derivatives (Figure 2).

What is, therefore, required is a method that avoids calculating or estimating the energy functionals derivative, while allowing the exploration of the search space in a manner that still converges towards an optimal solution. Thus retaining speed by avoiding gradient calculations, allowing the exploration to be carried out from a variety of initial locations, and enabling it be done in a way that reflects the learned variations of shape in terms of bends, bulges and stretches. GA are an example of one such method.

Our proposed method utilizes prior knowledge to produce

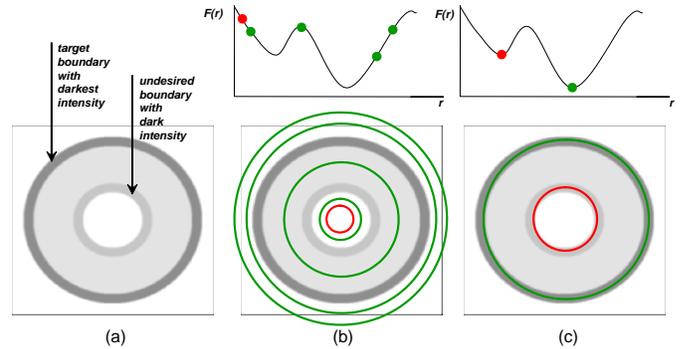


Figure 2: Example of single parameter deformable model with local minima. The circular deformable model’s only parameter is its radius r . The energy function $F(r)$ is minimal for the circle with darkest average intensity. The input image is shown in (a) with the darkest outmost boundary representing the global minima. In (b) traditional deformable models are initialized at a single point shown in red while GA-based statistical deformable models are initialized at multiple locations (green) and perform mutations on r . In (c) the GA converges to the global minima (darkest), while the deformable model gets stuck in a less dark local minima.

feasible deformations while also controlling the scale and location of these deformations. Moreover, through the use of GAs we solve the initialization problem, improve resistivity to local minima, and allow the optimization of highly customizable energy functionals while avoiding costly derivative estimations. Thus allowing us to explore a high-degree of simultaneous solutions in a highly-multivariate search-space, while decomposing the shape deformations into intuitive and localized constituents (bulge, bend, and stretch) that render the results more interpretable by clinicians (e.g. how much bend was needed in a particular part to fit to the new patient anatomy).

Other works have used GAs to drive traditional deformable models [1, 2, 18]. However, their lack of statistical constraints forces the reliance on traditional energy functionals that incorporate internal smoothness constraints. In [15] GAs were used with statistically-based Active Shape Models (ASMs) (section 1.2), where the parameter space consists of possible ranges of values for the pose and shape parameters of the model. The objective function to be maximized reflects the similarity between the gray levels related to the object in the search stage and those found from training.

In what follows we develop our method for medical image segmentation using GAs to drive statistically-based, controlled (location/scale) and intuitive (bend, bulge, and stretch), deformations towards optimal solutions of a problem-specific energy functional. We begin with a summary of how classical deformable models optimize their energy functionals (section 1.1) then describe how statistical shape constraints can be added to restrict segmentation (section 1.2). We describe how to obtain statistically constrained and intuitive deformations (section 1.3), resulting in a shape representation that allows for statistically constrained deformation types at specific locations and scales. In section 1.4 we present an overview of GAs. Finally, we detail how we

employ GAs to drive our statistically based deformations in section 2, present results in section 3, and conclude in section 4.

1.1 Classical Deformable Shape Models

Classical deformable shape models [29], are represented as a dynamic 2D parametric contour $\mathbf{v}(s, t) = (x(s, t), y(s, t))$, where $s \in [0, 1]$ traverses the contour, t denotes time, and \mathbf{v} is deformed to fit to image data by minimizing an energy term ξ ,

$$\xi(\mathbf{v}(s, t)) = \alpha(\mathbf{v}(s, t)) + \beta(\mathbf{v}(s, t)). \quad (1)$$

that depends on both the shape of the contour and the image data $I(x, y)$ reflected via the internal and external energy terms, $\alpha(\mathbf{v}(s, t))$ and $\beta(\mathbf{v}(s, t))$, respectively.

The internal energy term is given as

$$\alpha(\mathbf{v}(s, t)) = \int_0^1 w_1(s) \left| \frac{\partial \mathbf{v}(s, t)}{\partial s} \right|^2 + w_2(s) \left| \frac{\partial^2 \mathbf{v}(s, t)}{\partial s^2} \right|^2 ds. \quad (2)$$

Whereas the external energy term is given as

$$\beta(\mathbf{v}(s, t)) = \int_0^1 w_3(s) P(\mathbf{v}(s, t)) ds. \quad (3)$$

The weighting functions w_1 and w_2 control the tension and flexibility of the contour, respectively, and w_3 controls the influence of image data. w_i 's can depend on s but are typically set to different constants. For the contour to be attracted to image features, the function $P(\mathbf{v}(s, t))$ is designed such that it has minima where the features have maxima. For example, for the contour to be attracted to high intensity changes (high gradients) we can choose

$$P(\mathbf{v}(s, t)) = P(x(s, t), y(s, t)) = -\|\nabla [G_\sigma * I(x(s, t), y(s, t))]\| \quad (4)$$

where $G_\sigma * I$ denotes the image convolved with a smoothing (e.g. Gaussian) filter with a parameter σ controlling the extent of the smoothing (e.g. variance of Gaussian).

Traditionally, the dynamic contour \mathbf{v} that minimizes the energy ξ must, according to the calculus of variations [9], satisfy the vector-valued partial differential (Euler-Lagrange) equation

$$\mu(s) \frac{\partial^2 \mathbf{v}}{\partial t^2} + \gamma(s) \frac{\partial \mathbf{v}}{\partial t} - \frac{\partial}{\partial s} \left(w_1 \frac{\partial \mathbf{v}}{\partial s} \right) + \frac{\partial^2}{\partial s^2} \left(w_2 \frac{\partial^2 \mathbf{v}}{\partial s^2} \right) + w_3 \nabla P(\mathbf{v}(s, t)) = \mathbf{0} \quad (5)$$

where $\mu(s)$ and $\gamma(s)$ are mass and damping densities, respectively. Solving for $\frac{\partial \mathbf{v}}{\partial t}$ yields a first-order iterative optimization method. Though other optimizations have been explored using simulated physical dynamics [12], or more recently GAs [1, 2, 18].

1.2 Statistically Constrained Deformations

In order to constrain the deformations according to some learned shape variation, a training set of shapes in different configurations must be collected. The training set is typically created by labelling corresponding landmark points in each shape example. A classical approach is to perform principal component analysis (PCA) on the boundary points of a training set to obtain a point distribution model (PDM)[7].

The PDM describes the average positions of the landmarks, the main modes of variation of landmark positions,

and the amount of variation each mode explains. To emphasize how our method differs from the traditional statistical methods used with GAs in [15] we summarize the steps involved in generating a PDM and the use of ASMs for image segmentation.

First to construct the PDM L landmarks are chosen to describe the training shapes [5]. After global PCA an object shape is represented by the sum of a mean shape and a linear combination of, say t , principal components, i.e.,

$$\mathbf{x} = \bar{\mathbf{x}} + \mathbf{P}\mathbf{b} \quad (6)$$

where \mathbf{x} is the vector of landmark coordinates, $\bar{\mathbf{x}}$ is the mean shape, \mathbf{P} is the matrix of principal components, \mathbf{b} is a vector of weighting parameters, also called shape parameters. \mathbf{x} and $\bar{\mathbf{x}}$ are each of length $2L$. \mathbf{P} is $2L \times t$ and \mathbf{b} is a vector of length t .

Constraining the entries of \mathbf{b} to \pm a few standard deviations along each principle component ensures deformed shapes lie in the Allowable Shape Domain (ASD).

ASMs find the proposed movement of the landmarks of a current shape estimate to new and better locations, which requires a model of gray level information (image data). The model can be obtained by examining the intensity profiles at each landmark and normal to the boundary created by the landmark and its neighbors. Then the intensity profiles are used to derive a normalized intensity difference (gradient, or derivative) profile giving invariance to the offsets and uniform scaling of the gray levels [6]. With L landmarks representing each shape, N training shapes, and N training images, they derive N profiles for each landmark, one from each image, and calculate the mean profile for each landmark using

$$\bar{\mathbf{y}}_j = \frac{1}{N} \sum_{i=1}^N \mathbf{y}_{ij} \quad (7)$$

where \mathbf{y}_{ij} is the normalized derivative profile for the j^{th} landmark in the i^{th} image and $\bar{\mathbf{y}}_j$ is the mean normalized derivative profile for the j^{th} landmark.

Given a new image, the basic idea is to start with an initial estimate, then examine the neighborhood of the landmarks aiming at finding better locations for the landmarks. The shape and the pose of the current estimate are hence changed to better fit the new locations of the landmarks while producing in the process a new acceptable or allowable shape.

The pose parameters are first found by aligning the current estimate to the new proposed shape. The remaining landmark position modifications generally span $2L$ -dimensions, whereas the shape variations obtained from the model are only t -dimensional. A least-squares solution can be used to solve the following equation for the changes in shape parameters $d\mathbf{b}$ (with orthonormal column of \mathbf{P} we have $\mathbf{P}^T \mathbf{P} = \mathbf{I}$)

$$d\mathbf{x} = \mathbf{P}d\mathbf{b} \Rightarrow d\mathbf{b} = (\mathbf{P}^T \mathbf{P})^{-1} \mathbf{P}^T d\mathbf{x} = \mathbf{P}^T d\mathbf{x} \quad (8)$$

where $d\mathbf{x}$ is a vector containing the remaining landmark position modifications, $d\mathbf{b}$ is a vector of changes in the shape parameters, and \mathbf{P} is the matrix of principal components.

Finally, the shape variations are limited to obtain an acceptable or allowable shape within the ASD by applying the constraints on the shape parameters. They obtain new estimates and re-iterate until approximate convergence (when the parameter changes are insignificant).

However, a significant drawback of this approach is that the result of varying the weight of a single variation mode generally causes all the landmark positions to change. In other words, although the original PDM model produces feasible shape deformations only, a desirable trait, it generally produces global deformations over the entire object. In the next section we detail how HRPCA can be used to obtain localized deformations describing variation of specific anatomical regions, and allowing finer control over the deformation process.

1.3 Statistically Constrained Localized and Intuitive Deformations using HRPCA

We use our multi-scale (hierarchical) and multi-location (regional) PCA method introduced in [10] on our publicly available training set of medial shape profiles computed using 51 mid-sagittal CC images provided in citeshenton1992. We will first give an overview of medial shape profiles and then proceed to describe how HRPCA can be applied.

Medial-axis based 2D shape representations enable such deformations by describing the object's shapes in terms of an axis positioned in the middle of the object along with thickness values assigned to each point on the axis that imply the shape of the boundary. We therefore describe the shapes as a mapping $x : \mathbb{R} \rightarrow \mathbb{R}^4$. The domain of which is a parameter m that traverses the medial axis. We use a single primary medial axis. Though secondary medial axes are needed to represent more complex structures. The range of the mapping consists of 4 scalar values for each m , called medial profiles. These are a length profile $L(m)$, an orientation profile $R(m)$, a left (with respect to the medial axis) thickness profile $T^l(m)$, and a right thickness profile $T^r(m)$, where $m = 1, 2, \dots, N$, N is the number of medial nodes, and nodes 1 and N are the terminal nodes. The length profile represents the distances between consecutive pairs of medial nodes, and the orientation profile represents the angles between segments connecting consecutive pairs of medial nodes. The thickness profiles represent the distances between medial nodes and their corresponding boundary points on both sides of the medial axis (Figure 3-bottom). Corresponding boundary points are calculated by computing the intersection of a line passing through each medial node in a direction normal to the medial axis, with the boundary representation of the object. Example medial profiles are shown in Figure 3-top.

These profiles are rotation- and translation-invariant and capture intuitive measure of shape: length, orientation, and thickness. Altering these profiles produces intuitive, controlled deformations: stretching, bending, and bulging, respectively.

Here the principal component analysis is a function of the location, scale, and type of shape profile (length, orientation, or thickness) (Figure 4). Hence we obtain an average medial sub-profile, the main modes of variation, and the amount of variation each mode explains for each location, scale, and shape profile type.

Global PCA becomes a special case of HRPCA by specifying $loc = 1$ and $scl = N$, where N is the number of shape variables covering the whole extent of the object. Hence we obtain N modes of variation of length N . Whereas, $loc = l$ and $scl = s$ will produce $s \times l$ values, say thickness values for the T_r profile, and, as such, result in an $s \times s$ covariance matrix for the s modes of variation of length s . Consequently,

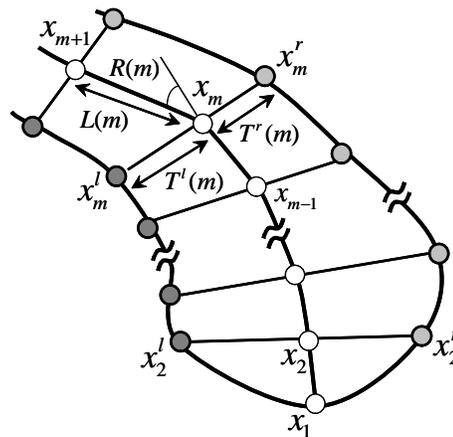
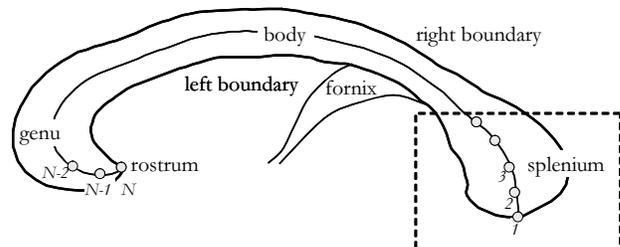
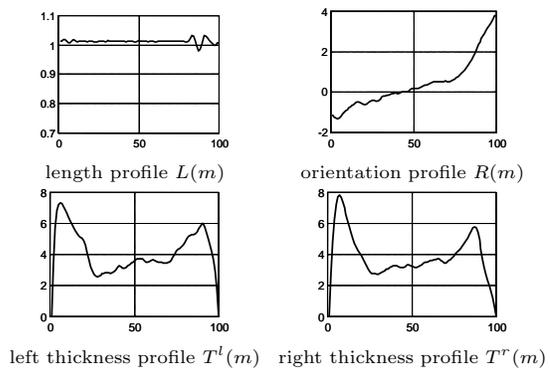


Figure 3: Top: Example medial shape profiles used to reconstruct the CC. Middle: Anatomically labeled CC shape reconstruction resulting from the medial profiles. Bottom: Details of outlined reconstruction showing medial profiles shape representation. Medial nodes shown in black, left and right boundary nodes shown in dark and light gray, respectively. x_m , x_m^l and x_m^r are the m^{th} medial, left boundary and right boundary nodes, respectively. $L(m)$, $R(m)$, $T^l(m)$ and $T^r(m)$ are the length, orientation, left and right thickness profile values, respectively (adapted from [10]).

we can now generate a statistically feasible stretch, bend, or bulge deformation at a specific location and scale in terms of the corresponding main modes of variation.

Generally with HRPCA, for a single deformation, location, and scale specific PCA we obtain the following model of medial profile variations,

$$p_{dls} = \bar{p}_{dls} + M_{dls} w_{dls} \quad (9)$$

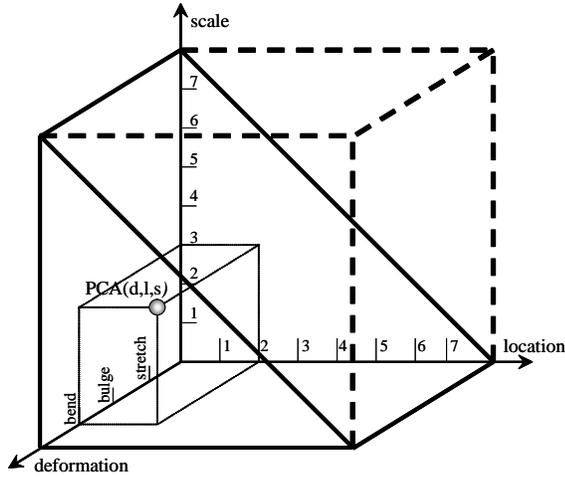


Figure 4: Hierarchical regional principal component analysis is a function of the deformation (d), location (l), and scale (s) (Adapted from [10]).

where p is the shape profile, d is the deformation profile type, l and s are the location and scale values, \bar{p}_{dls} is the average medial profile, M_{dls} describes the main variation modes, and w_{dls} are weights of the variation modes and are typically limited to ± 3 standard deviations.

Note that for any shape profile type multiple variation modes can be activated by setting the corresponding weighting factors to non-zero values. Since each variation mode acts at a certain location and scale we obtain

$$p_d = \bar{p}_d + \sum_l \sum_s M_{dls} w_{dls}. \quad (10)$$

In summary, varying the weights of one or more of the variation modes alters the length, orientation, or thickness profiles and generates, upon reconstruction, statistically feasible stretch, bend, or bulge deformations at specific locations and scales.

1.4 Genetic Algorithms

GAs are a special form of local search that models our own understanding of evolution. In essence a number of simultaneous agents (the population) each having an encoded state (the chromosome) perform a random walk (mutations) around the search space, while forming new solutions from combinations of existing solutions (crossover) and, thus, adjusting and refocusing the efforts of the search on exceptionally good areas once located. A few important choices are made during any application of genetic algorithms, involving how to encode the population (binary, integer, decimal, etc), how to mutate the population (mutate all genes, some genes, etc), how to select the parents for crossovers (roulette wheel, tournament selection), how to perform those crossovers (uniform, single-point), and finally what fitness function to use for evaluation. Though these choices seem complex, in situations where the energy functional has hundreds or even thousands of dependent variables and parameters these few choices can yield nearly-optimal values for all variables and parameters concerned.

2. METHODS

In this work we use GAs to solve the typical initialization, local minima and parameter sensitivity problems associated with traditional energy-minimization techniques. Moreover, the medial shape representation provides an intuitive way to synthesize/control deformations, while HRPCHA enables localized statistics thereby localizing the variations and deformations to specific anatomical regions of a shape (Figure 3-middle).

Here we describe our representation of individuals, our encoding of the model into chromosomes (deformation weights) to be optimized, our method of mutating (deforming) the model, our selection and crossover methods, and our fitness function (energy functional to minimize).

2.1 Population Representation

In medical image segmentation using GAs, the individuals forming the population represent potential shapes of the target structure, each having some level of accuracy measured by the fitness function (section 2.6). Consequently, we require a shape representation consistent across models and capable of intuitive mutations. One straightforward way of representing shapes is using boundary nodes (section 1.1). However, intuitive aspects of shape variation (such as bending, thickness, and elongation) can not be easily captured and, therefore, not properly represented in the mutations, selection and crossover phases. We require a shape representations that allows us to describe and control the shape deformations intuitively and in-terms of our calculated shape statistics. Consequently, we represent each individual by its associated stretching, bending and thickness profiles along with its global orientation, base location, and scale (section 1.3).

2.2 Encoding the Weights for GA

We use chromosomes to represent a set of all the weights of the principal components as obtained from the HRPCHA, where each gene represents a weight for a particular deformation, location, scale and mode of variation (Figure 5-top).

In total there are $\sum_{d=1}^4 \sum_{lc=1}^N \sum_{sc=1}^{N-lc+1} \sum_{w=1}^{sc} 1$ weights available for mutation since for each of the four deformations, d , we have $N = 27$ different locations, but for each location, lc , we can only have up to $N - lc + 1$ scales, sc , each of which has sc weights for the sc principle components. (Figure 4). In our application this adds up to 14616 dependent variables for our model, which motivates the use of GAs to search the highly-multivariate space.

2.3 Mutations

In order to somewhat guide the solution in the right direction, and, thereby, facilitate faster initial convergence, we start by constraining the mutations to the affine transformation parameters: base-node position (translation values) (t_x, t_y) , model orientation angle θ , and scale values (s_x, s_y) (Figure 5-top). Since our initial shape is the mean CC (obtained by setting all weights to zero) it can be expected to provide a reasonably strong fitness value when an acceptable position, orientation, and scale are set. In essence, we eliminate the possibility of getting a low score for a good location, scale and orientation, simply because of a bad random shape mutation.

With an adequate location, scale, and orientation ob-

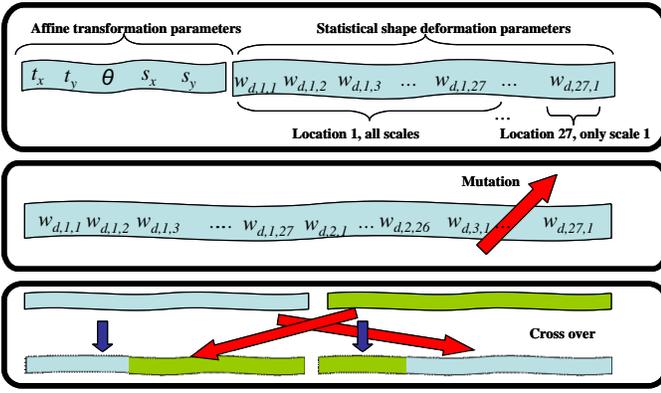


Figure 5: Segmenting an anatomical structure amounts to finding the optimal set of shape parameters. In our GA implementation (section 2.1) we represent each shape as a chromosome with genes encoding affine and statistical shape deformation parameters (top). Mutation (section 2.3) is performed by altering the weights of the HRPCA (middle). Crossover amounts to swapping a set of weights between two individuals (bottom).

tained we allow the mutations to begin including shape deformations (Figure 5-mid). Dynamic mutation of a single gene amounts to altering the corresponding weight by sampling it from a uniform random variable under the constraint that the total weight lie within ± 2 standard deviations of the corresponding mode of variation (square root of the explained variance obtained in HRPCA). Modifying a weight will change the medial profiles and hence the reconstructed shape boundary (section 2.6.1).

During the initial phases of the evolution each member of the population undergoes a random deformation with global scale, and at random amplitudes set in multiples of the corresponding standard deviations. Thus resulting in a new shape. The initial constraint to global deformations is well-suited for our statistical deformations as localized deformations (say bulging the splenium in Figure 3-middle) will not help until an acceptable global fit is obtained. Consequently, after a particular number of generations has passed, we allow the deformations to begin varying in both position and scale to include at first larger deformations (those corresponding to an entire anatomical region, and, hence, a primary area of variation) then smaller deformations which surmount to small variations in local regions.

2.4 Selection

Genetic algorithms require a method of determining which members of a generation will reproduce, and which will survive. We use roulette wheel selection, where each member i of a population P has a probability of selection equal to

$$P_{selection}(i) = Fit(i) / \sum_{j \in P} Fit(j) \quad (11)$$

to randomly select members for reproduction, where $Fit(i)$ is the fitness function (Section 2.6). We also employ an “elitist” strategy under-which the best member of the population is always maintained, and the weakest are replaced by the new individuals resulting from the crossover operation.

2.5 Crossover

Genetic algorithms use crossover to combine the information from two existing “parents” into a single “offspring”, that contains genes from each parent. We used uniform crossover, which makes an independent random decision for each gene whereunder both parents have an equal probability of making the contribution (Figure 5-bottom).

2.6 Fitness

Our fitness function is specifically designed for segmentation of the corpus callosum, though as noted in section 1 the use of GAs allow us to easily adapt the function to any given task including both prior shape and image-based knowledge; something traditional deformable models do not allow for. For example, we have adopted the fitness function $Fit(i)$ to consider mean image intensity, edge strength, standard deviation, and anatomical size of the CC.

$$Fit(i) = \alpha S(i) + \beta \left(1 - e^{\left(\frac{-E(i)}{\chi} \right)} \right) + \phi \left(1 - e^{\left(\frac{-\mu(i)}{\varphi} \right)} \right) + \delta e^{\left(\frac{-\sigma(i)}{\epsilon} \right)} \quad (12)$$

where

$$S(i) = \frac{|\Omega_{internal}|}{\varpi} \quad (13)$$

$$E(i) = \frac{1}{|\Omega_{contour}|} \int_{\Omega_{contour}} \|\nabla I\| \quad (14)$$

$$\mu(i) = \frac{1}{|\Omega_{internal}|} \int_{\Omega_{internal}} I \quad (15)$$

$$\sigma(i) = \sqrt{\frac{1}{|\Omega_{internal}|} \int_{\Omega_{internal}} (I - \mu(i))^2} \quad (16)$$

$|\Omega_{internal}|$ is the area enclosed by the reconstructed boundary (section 2.6.1), $|\Omega_{contour}|$ is the length of the boundary, and ϖ, η are the average size and standard deviation of the CC learned from the training set, respectively. I is the image, and χ, ϵ, φ are learned edge strength, standard deviation, and mean intensity. Hence, $S(i)$ represents the shape’s area, $E(i)$ the average gradient magnitude at the shape’s boundary, $\mu(i)$ the image intensity enclosed by the shape’s boundary, $\sigma(i)$ standard deviation. Finally, α, β, ϕ , and δ are scalar weights controlling the importance of each term in the segmentation process. For all experiments in this paper $\alpha = 0.1, \beta = 0.05, \phi = 0.80$, and $\delta = 0.05$, and parameter learning is performed using leave-one-out validation.

2.6.1 Shape Reconstruction for Fitness Calculation

In order to evaluate the fitness, the boundary of the CC shape must be reconstructed from the set of affine parameters and medial profiles specified by the shape weights. To reconstruct the object’s shape given its set of medial profiles, we calculate the positions of the medial and boundary nodes from a known reference node at location $x_1 = (t_x, t_y)$. The next node at position $x_2 = x_1 + v_1$ is specified using an offset vector v whose angle is specified by the orientation

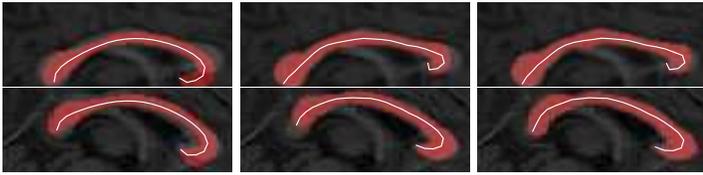


Figure 6: Two example segmentation results progressing left to right, showing fittest individual after automatic initialization (left), global deformations (middle), and local deformations (right).

profile plus the base angle θ , and length is specified by the stretch profile scaled by (s_x, s_y) . The corresponding boundary nodes x_2^l and x_2^r (Figure 3) are then orthogonal to the medial axis, at a distance specified by the thickness profile scaled by (s_x, s_y) . This process is repeated recursively, generating $x_3 = x_2 + v_2$, and so on. For details see [10]. Finally, with the medial profiles of Figure 3-top, for example, as an input we reconstruct the corpus callosum (CC) structure in Figure 3-middle.

3. RESULTS

We validate our method through the segmentation of the corpus callosum, which is the largest white-matter tract in the human brain. Specifically, it serves as the primary means of communication between the two cerebral hemispheres and mediates the integration of cortical processes from opposite sides of the brain. The presence of morphological differences in the corpus callosum in schizophrenics has been the subject of intense investigation [25]. The corpus callosum may also be involved in Alzheimer’s disease [23], mental retardation [19], and other disorders.

Table 1: Error comparison between statistical-based deformations with a hand-crafted schedule, and our GA based models. Error $\varepsilon = (SUM - S \cap M)/M$ is used, where S and M denote the area enclosed within the result of the automatic segmentation and the manual expert delineation, respectively.

Error	mean	median	min	max	std
Hand-crafted	0.1834	0.1706	0.1095	0.4526	0.0576
GA	0.1719	0.1501	0.0732	0.5464	0.0868

We present qualitative as well as quantitative results of the fully-automatic segmentation of 46 corpus callosum in mid-sagittal magnetic resonance images [27] using our GA driven, statistically-constrained deformable models (Figures 6, 7). We compare our results to those previously obtained in [14], where statistically-constrained, physically-based deformations are controlled by a CC specific hand-crafted schedule and initialization method (Table 1). Here our use of GAs enabled us to obtain superior accuracy without depending on an application-specific schedule, which strongly emphasizes the extendability of our method.

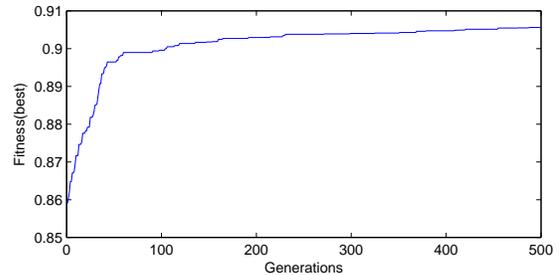


Figure 7: Plot of generation number versus fitness of the best individual.

4. CONCLUSIONS

We have developed a novel segmentation technique by addressing the main concerns associated with both traditional and statistical-based deformable models. Firstly, by using GAs to solve the initialization, local minima and parameter sensitivity problems associated with traditional energy-minimization techniques. Secondly, our medial shape representation provides a powerful way to synthesize and analyze deformations thus decomposing deformations into different types that are intuitively controlled and are more easily communicated to medical experts than boundary based deformations. Finally, our use of HRPCHA enables localized statistics thereby localizing the variations and deformations to specific anatomical regions that render the results more interpretable by clinicians and enable regional statistical analysis.

Furthermore, we have demonstrated how GAs can be combined with constrained shape deformations to effectively explore the search space of a complex energy functional, thereby incorporating prior-knowledge into the solution while retaining multiple simultaneous searches of the space. In essence, we have constrained the random-walks of the GA to lie within the allowable shape domain thus greatly reducing the search space traditionally associated with deformable models.

Our method is also extendible to other segmentation problems. Specifically, given a training set of shapes for a different anatomical structure, one can perform skeletonization followed by medial profile extraction and, subsequently, HRPCHA. Further, the components of the fitness functions presented here can apply to other anatomical structures, with possible minor modifications as the application warrants them. Nevertheless other terms can easily be added related to texture, colour or other image features. Finally, we are working on extending these ideas to 3D, where the genes become the weights of 3D shape representation parameters.

Though other works have used GAs to drive deformable models [1, 2, 18, 15]. To the best of our knowledge, no works have combined GAs with statistical-based deformations in a way that yields intuitively constrained deformations, nor have they employed fitness functions as well-suited to the problem domain. Furthermore, by comparison our method retains speed by avoiding gradient calculations, allows search space exploration to be carried out from a variety of initial locations, and enables it to be done in a way that intuitively reflects the learned variations of shape (bends, bulges and stretches).

5. ACKNOWLEDGMENTS

We wish to thank Dr. Martha Shenton of the Harvard Medical School for providing the MRI data, and Peter Plett for assisting with code development.

6. REFERENCES

- [1] L. Ballerini. Genetic snakes for medical image segmentation. *Proc. SPIE, Mathematical Modeling and Estimation Techniques in Computer Vision*, 3457, 1998.
- [2] L. Ballerini. Genetic snakes for color images segmentation. *Lecture Notes in Computer Science*, 2037:268, 2001.
- [3] V. Caselles, R. Kimmel, and G. Sapiro. Geodesic active contours. In *IJCV*, volume 22(1), pages 61–79, 1997.
- [4] L. D. Cohen. On active contour models and balloons. *CVGIP: Image Underst.*, 53(2):211–218, 1991.
- [5] T. Cootes, C. Taylor, D. Cooper, and J. Graham. Training models of shape from sets of examples. *BMVC*, pages 9–18, 1992.
- [6] T. Cootes, C. Taylor, A. Hill, and J. Halsam. The use of active shape models for locating structures in medical images. *Proceedings of the 13th International Conference on Information Processing in Medical Imaging*, pages 33–47, 1993.
- [7] T. F. Cootes, D. Cooper, C. J. Taylor, and J. Graham. Active shape models- their training and application. In *Computer Vision and Image Understanding*, volume 61, pages 38–59, 1995.
- [8] T. F. Cootes, G. J. Edwards, and C. J. Taylor. Active appearance models. In *IEEE Transactions on Pattern Analysis and Machine Intelligence*, volume 23(1), pages 681–685, 2001.
- [9] L. Elsgolc. *Calculus of Variations*. Pergamon Press Ltd., 1962.
- [10] G. Hamarneh, R. Abu-Gharbieh, and T. McInerney. Medial profiles for modeling and statistical analysis of shape. *International Journal of Shape Modeling*, 10(2):187–209, 2004.
- [11] G. Hamarneh and T. Gustavsson. Statistically constrained snake deformations. In *IEEE International Conference on Systems, Man, and Cybernetics*, volume 3, pages 1610–1615, 2000.
- [12] G. Hamarneh and T. McInerney. Physics-based shape deformations for medical image analysis. volume 5014, pages 354–362. SPIE, 2003.
- [13] G. Hamarneh, T. McInerney, and D. Terzopoulos. Deformable organisms for automatic medical image analysis. In *MICCAI*, pages 66–76, 2001.
- [14] G. Hamarneh and C. McIntosh. Physics-based deformable organisms for medical image analysis. *Proceedings of SPIE Medical Imaging: Image Processing*, 5747:326–335, 2005.
- [15] A. Hill and C. Taylor. Model-based image interpretation using genetic algorithms. *Image and Vision Computing*, 10(5):295–300, 1992.
- [16] M. Kass, A. Witkin, and D. Terzopoulos. Snakes: Active contour models. volume 1(4), pages 321–331, 1987.
- [17] M. Leventon, E. Grimson, and O. Faugeras. Statistical shape influence in geodesic active contours. In *CVPR*, 2000.
- [18] L. MacEachern and T. Manku. Genetic algorithms for active contour optimization. *IEEE Proceedings of the International Symposium on Circuits and Systems*, 4:229–232, 1998.
- [19] E. Marszal, E. Jamrow, J. Pilch, and et. al. Agenesis of corpus callosum: Clinical description and etiology. *Journal of Child Neurology*, 15:401–405, 2000.
- [20] T. McInerney and D. Terzopoulos. Deformable models in medical image analysis: A survey. In *Medical Image Analysis*, volume 1(2), pages 91–108, 1996.
- [21] J. Montagnat, H. Delingette, and N. Ayache. A review of deformable surfaces: topology, geometry and deformation. In *Image and Vision Computing*, volume 19(14), pages 1023–1040, 2001.
- [22] S. Osher and N. Paragios. *Geometric Level Set Methods in Imaging Vision and Graphics*. Springer Verlag, 2003.
- [23] J. Pantel, J. Schroder, M. Jauss, and et. al. Topography of callosal atrophy reflects distribution of regional cerebral volume reduction in alzheimer’s disease. *Psychiatry Research*, 90:180–192, 1999.
- [24] R. A. Robb. *Biomedical Imaging, Visualization, and Analysis*. Wiley-Liss Inc., 2000.
- [25] R. Rosenthal and L. Bigelow. Quantitative brain measurements in chronic schizophrenia. *British Journal of Psychiatry*, 121:259–64, 1972.
- [26] J. A. Sethian. Level set methods; evolving interfaces in geometry, fluid mechanics. In *Computer Vision and Material Sciences*. Cambridge University Press, 1996.
- [27] M. Shenton, R. Kikinis, F. Jolesz, S. Pollak, M. LeMay, C. Wible, H. Hokama, J. Martin, D. Metcalf, M. Coleman, and et al. Abnormalities in the left temporal lobe and thought disorder in schizophrenia: A quantitative magnetic resonance imaging study. In *New England J. Medicine*, volume 327, pages 604–612, 1992.
- [28] M. Sonka and J. Fitzpatrick. *Handbook of Medical Imaging, Volume 2: Medical Image Processing and Analysis*. SPIE-International Society for Optical Engine, 2000.
- [29] D. Terzopoulos. On matching deformable models to images. *Topical Meeting on Machine Vision, Technical Digest Series*, 12:160–167, 1987.
- [30] S. K. Warfield, M. Kaus, F. A. Jolesz, and R. Kikinis. Adaptive, template moderated, spatially varying statistical classification. In *Medical Image Analysis*, volume 4(1), pages 43–55, 2000.
- [31] C. Xu and J. Prince. Snakes, shapes, and gradient vector flow. *IEEE Transactions on Image Processing*, 7(3):359–369, 1998.
- [32] T. S. Yoo. *Insight into Images: Principles and Practice for Segmentation, Registration, and Image Analysis*. A k Peters Ltd., suite 230, 888 Worcester St. Wellesey, MA, USA, 2004.