

# Medial-Based Deformable Models in Nonconvex Shape-Spaces for Medical Image Segmentation

Chris McIntosh\*, *Student Member, IEEE*, and Ghassan Hamarneh, *Senior Member, IEEE*

**Abstract**—We explore the application of genetic algorithms (GA) to deformable models through the proposition of a novel method for medical image segmentation that combines GA with nonconvex, localized, medial-based shape statistics. We replace the more typical gradient descent optimizer used in deformable models with GA, and the convex, implicit, global shape statistics with nonconvex, explicit, localized ones. Specifically, we propose GA to reduce typical deformable model weaknesses pertaining to model initialization, pose estimation and 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 are driven in terms of localized principal modes of variation, instead of modes of variation across the entire shape that often fail to capture localized shape changes. Although GA are not guaranteed to achieve the global optima, our method compares favorably to the prevalent optimization techniques, convex/nonconvex gradient-based optimizers and to globally optimal graph-theoretic combinatorial optimization techniques, when applied to the task of corpus callosum segmentation in 50 mid-sagittal brain magnetic resonance images.

**Index Terms**—Deformable models, evolutionary computing, genetic algorithms, medical image segmentation, medial-shape representation.

## I. INTRODUCTION

**M**EDICAL image segmentation has important applications in computer-aided diagnosis, statistical shape analysis, and medical image visualization. Deformable models are a popular class of segmentation techniques that have been widely researched and adopted to medical image analysis (MIA) problems [1], [2]. This adoption is due to their inherent smoothness properties and ability to fit to missing boundaries.

At a high level, deformable models work by deforming a user provided initial shape to fit to a target structure in a medical image. Shape-changing deformations result from the minimization, with respect to the shape, of an energy function measuring

how plausible the shape model is and how well it aligns with the boundaries of the target anatomy in the image. The energy function is typically formulated as a weighted sum of internal and external terms. Internal terms favor plausible shapes regardless of image data. External terms favor shape models that align to image features, such as image boundaries with high gradient magnitude, and are primarily pixel-based. Being pixel-based, external terms are susceptible to noise that can obscure structures and lead to false boundaries, thereby necessitating the use of internal terms in the form of shape priors [3]–[10]. Shape priors provide resilience to false boundaries by heavily penalizing the implausible shape configurations that the false boundaries imply. Since the shape itself is most commonly represented by a function, the cost function is often termed an energy functional and its gradient is derived using methods from variational calculus. The shape deformations are therefore typically simulated by solving an initial value problem using gradient descent optimization algorithms [11].

The primary goal of this work is to address two open and related problems in deformable-model-based medical image segmentation.

- 1) **Shape model fidelity:** The ideal shape model must capture the anatomical variability of the target anatomy without representing invalid shapes. However, the better the shape model the harder it becomes to formalize and optimize in a segmentation framework. For example, a convex shape-model is readily formulated and optimized, but inaccurately fits a nonconvex shape-space; nonconvex shape-spaces are more accurate, but are harder to formalize and optimize. By the same token, localized, medial-axis-based shape statistics better respect anatomical variability than global statistics [12], but their inherent nonconvexity makes them challenging to optimize. There is a trade-off: one gains optimizability but loses fidelity.
- 2) **Optimizability:** As the objective function of a deformable model more faithfully encodes the desired properties of the target anatomy, the energy landscape becomes more complicated and problems with initializations and local minima arise. For example, adding pose to even convex shape-model-based deformable models is well known to cause nonconvexity [9]. At the expense of the desired faithfulness, recent trends have focused on ensuring global optima by simplifying the objective function through the use of convex, or submodular, approximations and shape spaces [9], [13]–[17]. Thus trading fidelity for optimizability. Where this sacrifice is not practical, a different solution is required. Genetic algorithms (GA) are a powerful alternative optimization-technique that place very little requirement on the objective function, but to the best

Manuscript received June 05, 2011; accepted July 06, 2011. Date of publication July 22, 2011; date of current version December 30, 2011. The work of C. McIntosh was supported by the National Sciences Research and Engineering Council of Canada, and the Michael Smith Foundation for Health Research. *Asterisk indicates corresponding author.*

\*C. McIntosh is with the Medical Image Analysis Lab, School of Computing Science, Simon Fraser University, Burnaby, BC, V5A 1S6, Canada (e-mail: cmcintos@sfu.ca).

G. Hamarneh is with the Medical Image Analysis Lab, School of Computing Science, Simon Fraser University, Burnaby, BC, V5A 1S6, Canada (e-mail: hamarneh@sfu.ca).

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TMI.2011.2162528

of our knowledge they remain entirely unexplored in their application to localized shape-based image segmentation. There are, however, some existing applications of GA to deformable models (see Section I-C for details).

In what follows, we elaborate on these open problems, and propose our solution that combines GA with localized shape training.<sup>1</sup> There is a trade-off in 1) and 2) above. Shape model fidelity is desirable, but lowers optimizability. Optimizability is desirable, but lowers fidelity. We will use GA to mitigate the trade-off that is usually required between between 1) and 2) above, allowing complex shape models while reducing problems with initializations and still being able to attain solutions close to the global optimum.

#### A. Shape Model Fidelity

In many applications, prior knowledge about object shape variability is available or can be obtained by studying a training set of shape examples. This knowledge restricts the space of allowable deformations to a learned shape-space that approximates the space of anatomically feasible shapes [3]–[6], [20]. One of the most notable works in this area is that of Cootes *et al.*, where they introduced and refined active shape models (ASM) [3], [21], [22]. In ASM, principal component analysis (PCA) is calculated over a set of landmark points extracted from training shapes. The resulting principal components are used to construct a point distribution model (PDM) and an allowable shape domain (ASD). In a natural extension to their previous work, Cootes *et al.* modify their method to include image intensity statistics [4]. Staib and Duncan constrained the deformable models in Fourier space by conforming to probability distributions of the parameters of an elliptic Fourier decomposition of the boundary [23]. Statistical prior shape knowledge was also incorporated in implicit, level-set-based deformable models. Leventon *et al.* introduced statistical shape priors by using PCA to capture the main modes of variation of the level set representation [5]. However, as Pohl *et al.* point out, level-sets do not form a vector space and hence more accurate shape statistics could be captured by transforming the shapes into a vector space using the logarithm of odds before performing PCA [24].

Though simpler to optimize than their nonlinear counterparts, linear models of shape deformation may not always adequately represent the variance observed in real data. Linear shape models assume the data lies on a linear manifold, but shapes often lie on nonlinear manifolds where the manifold's properties are not accurately captured by linear statistics [25]. Picture fitting an ellipse to an “S”-like shape space. In order to include the entire letter, extraneous white-space (non-valid shapes) must also be included. Nonlinear shape models have been introduced to address this problem [7], [25]–[31].

We argue that the shape model fidelity problem, as described above, is not necessarily due to the application of a linear model to nonlinear data but rather because of the implicit nature in which the statistics are applied. By *implicit*, we mean the statistics attempt to model variation in the shape, rather than variation in the parameters governing the deformations themselves. Note that we are not referring to the shape representation being implicit or explicit, but instead whether the deformations are




Fig. 1. Shape deformations of the boundary of the corpus callosum (CC) under its global modes of variation using a boundary-based model (a) and a medial-axis-based model (b). In (a) the mean shape is shown in blue. The red and green contours indicate a positive and negative deformation of the CC along the first principal component. Notice the global and linear statistics in (a) enforce a very unnatural deformation of the shape compared to the deformations in (b).

implicitly or explicitly studied. Implicit shape statistics result from the majority of previous deformable model approaches adopting a boundary-based shape representation, aside from a few exceptions [12]. As a consequence, the statistics are calculated using boundary models of the shape instead of models representing the interior and skeletal topology of the structures. Studying the underlying structural changes of a shape allows deformations that were previously nonconvex to be decomposed into linear models. We refer to these as *explicit* shape statistics since they are calculated over the very parameters responsible for varying the object's shape. Consider an object represented by a single pixel. Different images of the object show the pixel moving around in a circle. A nonlinear function is required to describe the pixels motion and hence no linear statistics can capture the motion adequately as long as it is the object's  $x, y$  position being studied. However, once decomposed into a function of  $\sin$  and  $\cos$ , the underlying parameter that controls the objects variability is linear in its variation, and hence linear statistics will have greatly improved shape model fidelity. The same argument carries forward, albeit more complexly, to a more complex object. A simple bending of a shape's medial-axis is a linear deformation under the appropriate representation, as it's simply a rotation of some of the medial-nodes (Fig. 1). However, the bending is a highly nonconvex deformation once embedded in the image domain, as either an implicit shape [24] or an explicit-boundary based model [3].

In addition to being implicit in nature, the statistics are often global in nature, whereby *global* we mean the statistics attempt to model variation in the entire shape. In other words, each shape is a single point in some high-dimensional space, and the statistics, linear or not, describe some restricted set of that space. Global shape statistics are unable to restrict deformations to particular locations of the shape model or anatomy. Furthermore, as most dimensionality reduction techniques are designed to preserve variance, small changes in shape are discarded to preserve larger ones. Though small in scale, those changes can represent true signal changes (i.e., not due to noise) in a localized region of a shape [32], [33]. Global deformations can not adapt to localized variations in shape, which are often of high interest, e.g., regions of pathology.

Decomposing implicitly complex deformations, like bending an object, into simpler linear parameters is made possible through localized shape statistics defined over a medial-based shape representation, which allows deformations to be further quantified into specific types (bending, bulging, stretching) that can be linearly measured. Specifically, medial-axis-based

<sup>1</sup>This work builds on that presented in [18] and extends [19].

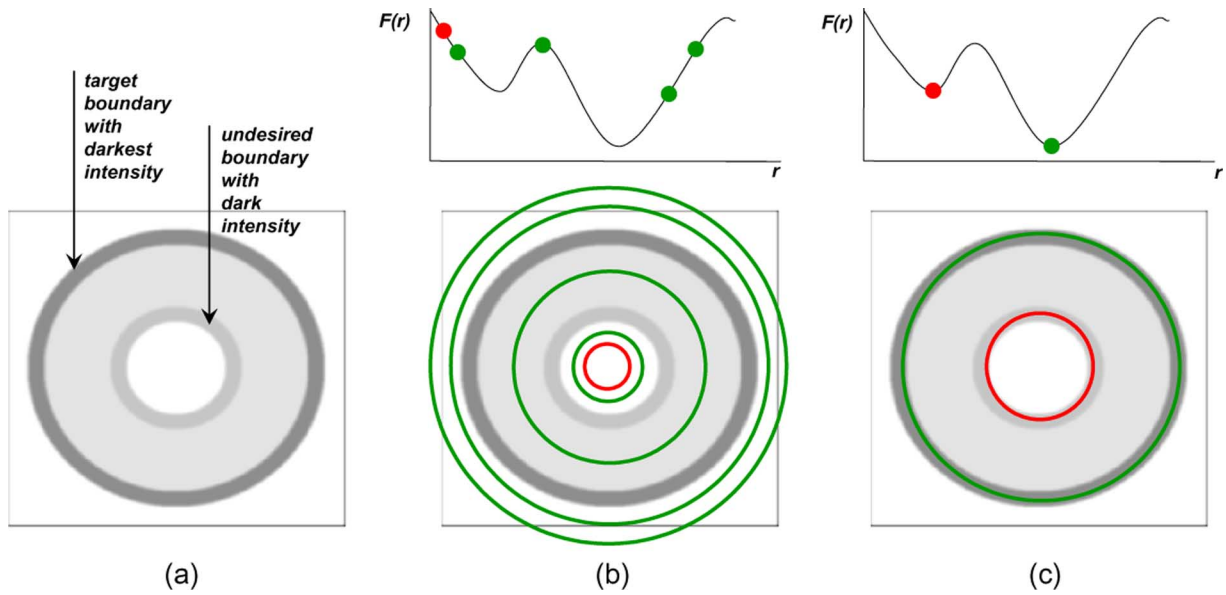


Fig. 2. Synthetic example of a single parameter deformable model with local minima. The circular deformable model’s only parameter is its radius  $r$ . The energy function  $F(r)$  is minimal for the circle with darkest average intensity. The input image is shown in (a) with the darkest outmost boundary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

2-D 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 object’s boundary. Rotating part of the medial axis produces bending deformations, increasing the thickness along part of the axis creates bulging, and lengthening the medial-axis produces a stretching of the object. Medial-based shape representations have been emerging as a powerful alternative to boundary-based techniques [18], [34]–[39]. In addition to decomposing seemingly complex deformations into their linear variables, medial-axis representations also avoid having to define absolute positions of boundary points. With medial-axis representations, deformations like bending the object are performed directly on the medial-axis, which implicitly moves the boundary and ensures the affected boundary points move together. By comparison, with a boundary representation, one has to carefully move all the connected points to ensure boundary continuity. Statistics on the nonlinear manifolds of medial representations have also been proposed [25]. However, there has been little exploration of building deformable models using these shape representations and their corresponding shape statistics [12], [19]. In this work, we explore this very notion by building a medial-based deformable model with deformations governed by explicit, localized shape statistics.

### B. Optimizability

Energy-functional minimization can be carried out in a variety of ways through various forms of optimization. One solution is to perform explicit differentiation under the Euler-Lagrange constraints, where each new image segmentation application that requires a modified energy functional must be accompanied by one such derivation. The result is a set of constraints, that guarantees a stationary point of the energy func-

tional. The solution is then obtained through a gradient descent process where the change of the shape model (with respect to an artificial time variable) is equated to the Euler-Lagrange constraints. When applied to typical deformable models, which are represented by nonconvex energy functionals, the optimization process still requires initialization at some optimal target location of an image, with some shape, orientation, and scale. Together, the location, orientation, and perhaps the scale of the shape model are referred to as the pose. Depending on the initialization, a different segmentation is typically obtained, and different initializations can therefore produce different amounts of segmentation error (Fig. 2). It is, as such, common for deformable models to be described as having problems with local minima.

As a direct result of the commonly occurring problems with local minima, there has been a number of recent approaches to obtaining global optima to energy functionals [9], [13]–[17]. Graph cuts were demonstrated as a global minimization technique for a popular energy-functional [40], as a special case of computing a geodesic on a Riemannian space whose metric is computed from the image [13]. However, graph cuts have been shown to apply only to a restricted class of energy functionals that are submodular [41], and their solutions are discrete approximations to the continuous formulations whose accuracy is dependent on the resolution of the approximating graph [13]. Naturally, as that resolution increases so does their running time. Random walkers were developed in a similar nature, solving image segmentation as a graph problem wherein the global optimum is obtained to a particular cost function [15]. In fact, graph cuts and random walkers have been shown to be specific instantiations of a single framework [42]. Another line of work has come from the relaxation of the underlying shape model from a nonconvex space to a convex one; thereby defining convex energy functionals which can then be mini-

mized instead of their nonconvex counterparts. This convex relaxation work, which began in 2004 with a simple restricted class of functionals [16], was later extended to a broader class in [17], and then a similar work appeared in 2008 with the addition of a shape prior [9]. However, restrictions still exist in that the functionals and the shape spaces they are optimized over must be convex when defined over the relaxed space, and that the relaxed shape-space must itself be convex. Though they can be globally optimized, convex energies and shape spaces may not be expressive enough to capture the characteristics of the segmentation problem at hand as accurately as their nonconvex counterparts. In other words, they may not faithfully encode the desired properties of the target anatomy. This is especially true for the shape spaces, as our results demonstrate later in the paper (Section III). Similarly, not all energy functionals and shape spaces can be represented convexly. There is thus a trade-off: gain global optimality but likely lose accuracy. As a consequence, these works do not provide a definitive solution to the problem. Another solution present in the optimization literature when initialization and local minima are problematic is GA, a method which is largely unexplored in relation to shape-model-based deformable models.

To reemphasize the argument of the previous paragraph, we pose the following question: Is it worth sacrificing the complexity of the energy functional, and thereby reducing its ability to address image variability, in exchange for global optimality, or can more complex functions and shape spaces be utilized with an approximate optimization method and achieve superior accuracy? Our hypothesis is that while convex approximations to nonconvex functions and shape spaces can be globally optimized they sacrifice fidelity to the data and, in doing so, are less accurate than nonconvex spaces and functions optimized using GA. We propose a novel method that uses GA to overcome the initialization and local-minima entrapment typically encountered in nonconvex problems. While we are not guaranteed global optima, our results show that using nonconvex shape spaces our method can outperform convex shape spaces under the same energy function (Section III), i.e., the energy function remains fixed, while only the shape constraints and solvers are varied.

### C. Contributions and Related Work

To summarize the previous sections, in order to be applicable to medical image segmentation, deformable models require learned, intuitive shape models, and a level of robustness to initialization. It is well known that implicit shape statistics, though easily represented and applied to deformable models, provide unintuitive deformations of shape and often fail to adequately represent nonconvex, localized deformations [12]. Localized shape statistics are more intuitive to clinicians, and better represent local variations in shape, but the proposed methods remain largely unapplied to deformable models. Robustness to initialization can be handled through convex optimization methods, but there are drawbacks in that: not all energy functions are convex; convex functionals may not faithfully represent the segmentation problem; and convex shape-spaces may not accurately summarize the variability of

the target anatomical structure. Therefore, what is desirable is a method that generalizes to a larger class of functionals (convex or not) while allowing the exploration of the search space in a manner that still converges towards an optimal solution. The method must allow the exploration to be carried out from a variety of initial locations, and enable it to be done in a way that reflects the nonconvex, localized variations of shape in terms of bends, bulges, and stretches. We propose one such method.

In this paper, we explore the application of GA to deformable models as well as propose a novel deformable-model based on localized shape statistics using hierarchical regional principle component analysis (HRPCA). Our main contribution in this paper is a novel method for fitting complex, medial-based shape models to medical images under arbitrary energy functionals. In a narrower sense, our new segmentation method replaces gradient descent with GA, and convex, implicit, global shape statistics with nonconvex, explicit, localized ones. GA address the typical initialization, local minima problems associated with traditional energy-minimization techniques. Our unique application of a medial-based deformable model with HRPCA maintains a statistically feasible shape model. It also reduces the size of the search space by constraining the evolution to deformations that are intuitive (stretch, bulge, bend) and are driven in terms of localized principal modes of variation, instead of modes of variation across the entire shape. Though not expressly guaranteed to find the global optimum, our method compares favorably to leading segmentation techniques, including convex/nonconvex gradient based optimizers and globally optimal graph-theoretic combinatorial optimization techniques, on a data set of 50 mid-sagittal brain magnetic resonance images.

Though prior works have used GA for medical image segmentation, here we limit our discussion to those pertaining to deformable models. For a comprehensive survey of the use of GA for medical image segmentation, the reader is referred to [43]. A few methods have used GA to minimize traditional deformable models [44]–[49]. Ballerini extends the classical active contour models [50] by using GA to directly minimize the standard energy functional [44]. Members of the GA population are hypothetical shape configurations, represented by their explicit contour locations. The method was later extended to color images by using one image term per color channel [45]. MacEachern and Manku presented a similar method using a binary representation of the contour [46]. Similarly, Tohka presented simplex meshes paired with image-based energies, minimized via a hybrid GA-greedy approach, and applied the technique to the segmentation of 3-D medical images [47]. Fan *et al.* also develop a GA method for an explicit active contour, but describe their method using Fourier descriptors and employ parallel GAs to speed up minimization [48]. A different shape representation, known as topological active nets, is used by Ibáñez *et al.* to enable the segmentation of objects with unknown topologies, or even multiple objects in the same scene [49]. However, aside from simple boundary smoothness constraints, all of these methods are based on classical active contour models or their variants without incorporating prior shape knowledge, making them prone to latching to erroneous edges and ill-equipped to handle gaps in object boundaries. In [51],

GA were used with statistically-based ASMs, 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. Additional works use convex, implicit, global shape statistics assuming a Gaussian distribution around a mean shape [52]–[54]. Mignotte and Meunier [54] incorporate prior shape information by defining the mean as a circular deformable template, while [53] uses a PDM for occluded shape reconstruction, and [52] uses a level set shape representation and a learned mean from training data. Although these techniques apply GA to produce generations of plausible populations of shapes, the statistically-based deformations are global (over the whole shape) and their convexity may not offer the required flexibility to accommodate for shape variations that are restricted to particular locations, nor are they intuitively defined (bulge, bend, stretch) deformations. In summary, though GA have been used to optimize both classical, and shape-prior-based deformable models, they have not been applied to deformable model methods incorporating localized, medial-based shape statistics. In Section III, we show superior accuracy to a globally optimized method that makes use of convex, implicit, global PCA-based shape statistics [9].

Other works have set out to automatically fit medial-based shape models to image data [55], [56]. These approaches build on the m-reps shape representation of Pizer *et al.* [57], in which shapes are represented using a global-to-local scale approach as objects, “figures” (anatomical subregions or parts), and “atoms” (individual medial nodes). Once initialized by registration to manually specified points, models are fitted to the data using a hierarchical local-search method, i.e., starting first with each object, then each figure, and finally each atom. At each scale, models are fitted to the image data using an iterative conditional modes (ICM) strategy across the scale’s entities (objects, figures, or atoms), with the conjugate gradient method being used to optimize each entry (object, figure, or atom) individually [58]. Our HRPCA-based approach complements the aforementioned works [55], [56] by computing statistics across a spectrum of shape scales and locations; ranging from an individual medial node (or atom), and adding increasingly more atoms until covering a whole anatomical part (or figure). This is achieved through our HRPCA in which multilocation, multiscale HRPCA statistics are gathered for different deformation profiles. This offers greater deformation flexibility (i.e., at multiple locations, sizes, and deformation types) during the model-to-image fitting and, at the same time, remains faithful to the underlying variability in the training data. Inspired by the approach in [55], [56], we also seek to fit our shape model to the image in a global to local manner. However, in our case we employ genetic algorithms in order to alleviate the initialization requirement of the gradient descent-like ICM method, which is known to be heavily sensitive to its initial configuration [59, Sec. 9.3.1, p. 247]). Furthermore, we present comparison to numerous convex and nonconvex methods for medical image segmentation, exploring the aforementioned questions relating to energy function and shape space fidelity versus optimizability (Section III).

To the best of our knowledge, none of the existing image segmentation techniques evolve a population of shapes using intuitive, spatially constrained, and plausible deformations, nor have they enabled such deformations under general energy functionals that are customizable for the problem domain.

## II. METHODS

In this work, we build a deformable model that uses a medial shape representation to provide a way to synthesize intuitive deformations and localized shape statistics (HRPCA). Our use of HRPCA captures regional medial-based shape information, thereby providing controlled localized fitting of specific anatomical regions of a shape [Fig. 1(b)]. We solve our resulting optimization problem using GA, demonstrating a generic way to fit complex shape models to medical images.

In order for the descriptions that follow to lend more readily to standard descriptions of GA, from here on in the paper, we consider the problem of fitting an HRPCA-based deformable model to an image as a fitness maximization problem, rather than an energy functional minimization problem. This is done without loss of generality.

We begin with an overview of how we use HRPCA to build localized shape statistics, where some of the key differences between what we have referred to as global and local statistics should become clearer. We also provide a short introduction to GA for the unfamiliar reader. We then use a motivating example to describe how GA can be used for deformable models, before describing our use of GA in detail.

### A. Statistically-Constrained Localized and Intuitive Deformations Using HRPCA

We use the multiscale (hierarchical) and multilocation (regional) PCA method introduced in [18] on a training set of medial shape profiles. We will first give an overview of medial shape profiles and then proceed to describe how HRPCA is applied in this work.

Medial-axis-based 2-D shape representations describe 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 shape as a mapping from  $\mathbb{R}$  to  $\mathbb{R}^4$ , the domain of which is a parameter  $m$  that traverses the medial axis. We use a single primary medial axis, though secondary medial axes or subfigures would be needed to represent more complex structures. The range of the mapping consists of four scalar values for each  $m$ , forming 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  $m$  is now discretized,  $m = 1, 2, \dots, N$ ,  $N$  is the number of medial nodes, and nodes 1 and  $N$  are the terminal nodes. For all examples here and all results presented later,  $N = 100$ . The length profile represents the distances between consecutive pairs of medial nodes. The orientation profile represents the angles between segments connecting consecutive pairs of medial nodes. Finally, the thickness profiles represent the distances between medial nodes and their corresponding boundary points on both sides of the medial axis (Fig. 3), and will be different only



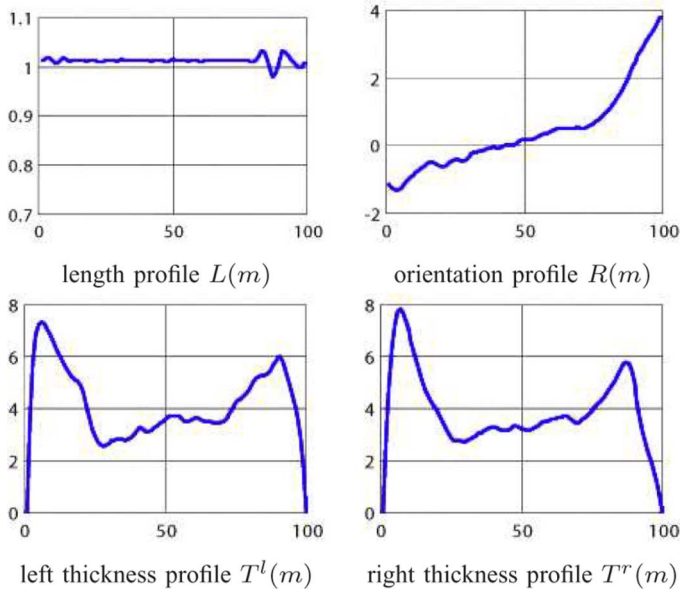


Fig. 3. Example medial shape profiles used to reconstruct a CC shape model. The values of the medial profiles are plotted along the  $y$ -axis with pixels for length and thickness profiles, and radians for the orientation profile. All plots are a function of the medial node number  $m = 1, 2, \dots, 100$ , indicated along the  $x$ -axis. Adapted from [18].

when the shape is asymmetric. In other words, corresponding boundary points  $x_m^l$  and  $x_m^r$  lie at distances  $T^l(m)$  and  $T^r(m)$ , respectively, along a direction normal to the medial axis at  $x_m$  (Fig. 4). As the medial-axis is discretized into  $N$  nodes, implying  $N - 1$  segments, the normal to node  $x_m$  is measured from the medial-segment connecting  $x_m$  and  $x_{m-1}$  (see right-angles in Fig. 4, right). Clearly,  $x_1$  is an exception and in that case the medial-segment connecting  $x_1$  and  $x_2$  is used. Other more elaborate medial representation may be adopted as well (e.g., m-reps [57]).

In order to build energy functions using HRPCA-based shape models, the boundary of the CC shape must be reconstructed from the four medial profiles and a set of affine parameters that describe the objects pose and scale. The affine parameters are: the position,  $(t_x, t_y)$ , describing the  $(x, y)$ -coordinate of a known reference node; the scale,  $(s_x, s_y)$ , describing the scale of the object; the base angle,  $\theta$ , describing the angular offset of the model from the  $x$ -axis. 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 determined using an offset vector  $v$  whose angle is specified by the orientation profile plus the base angle  $\theta$ , and whose length is specified by the stretch profile scaled by  $(s_x, s_y)$ . The corresponding boundary nodes  $x_2^l$  and  $x_2^r$  (Fig. 4, right) 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 complete details see [18]. Finally, with the medial profiles like those shown in Fig. 3 as an input, we can reconstruct the CC structure in Fig. 4. Using polygon rasterization, the reconstructed set of nodes can then be used to create a binary image, of the same size as an input image,  $I$ , with ones representing pixels on or within the




Fig. 4. Representing the CC using medial profiles. (Top) Anatomically labeled CC shape reconstruction resulting from the medial profiles in Fig. 3. (Bottom) The details of the medial profile representation for the splenium region (dashed). Medial nodes are shown in white, whereas the left and right boundary nodes are 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 [18].

shapes' border, and zeros representing pixels not belonging to the shape [60].

The profiles are thus rotation- and translation-invariant, and capture intuitive measures of shape: length, orientation, and thickness. Altering the different profiles produces intuitive, controlled deformations: stretching, bending, and bulging, respectively. For example, setting  $\theta = 0$  and  $\{R(m) = 0\}_{m=1}^N$  would produce a straight-line medial axis, parallel to the  $x$ -axis of the image. Setting  $\{R(m) = \pi/2\}_{m=N/2}^N$  would create a right-angle bend at the midpoint of the medial-axis. Increasing  $\theta$  would simply rotate the object around the reference point. Of course, setting  $R(m)$  randomly may be undesirable and produce unwanted shapes. Next we describe how to use training data to calculate statistics for each profile in order to restrict the deformations to plausible shapes.

In HRPCA, the principal component analysis is a function of the location,  $l$ , scale,  $s$ , and type of shape profile,  $d$  (i.e., length  $L(m)$ ; orientation  $R(m)$ ; left thickness  $T^l(m)$ ; or right thickness  $T^r(m)$  profiles) (Fig. 5). For clarity of presentation, we define  $p_d(m)$  as the  $d$ th shape profile (e.g.,  $p_1(m) = L(m)$ ). Hence, for each location, scale, and shape profile type, we obtain an average medial sub-profile,  $\bar{p}_{dls}$ , the main modes of variation,  $\mathbf{M}_{dls}$ , and the amount of variation each mode explains. The subscript "dls" indicates that the statistics are made specific to a location  $l$  (e.g., to the left of the shape), scale or extent  $s$  (e.g., affecting one tenth of the shape, half of it, or all of it), and the deformation profile type  $d$  (e.g., bending or bulging). Note, of course, that at each location,  $l$ , a maximum scale,  $s$ , of  $N - l + 1$  can be selected. For example,  $p_{L,1.5} = \bar{p}_{L,1.5} + \mathbf{M}_{L,1.5}\alpha_{L,1.5}$  would calculate a new  $L(m) = p_{L,1.5}$ , for  $m = 1 \dots 5$ , since only those 5 positions are affected. Example modes of variation for a global bending deformations are shown in Fig. 1(b), i.e.,

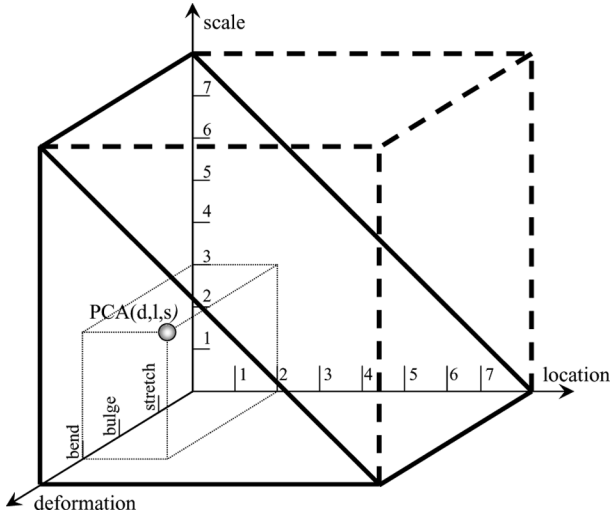


Fig. 5. Hierarchical regional principal component analysis is a function of the deformation ( $d$ ), location ( $l$ ), and scale ( $s$ ). Since all scales are not valid for all locations, only a subset of the full space (the right angle prism drawn with solid black) is valid (see Section II-D2). Adapted from [18].

$d = 2$ ,  $l = 1$  and  $s = 100$ . For each location,  $l$ , with scale  $s$ , there are  $s$  values, say thickness values, for the  $T^r$  profile and, as such, for each profile, location, and scale,  $\mathbf{M}_{dls}$  is an  $s \times s$  covariance matrix for the  $s$  modes of variation of length  $s$ . Consequently, 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 [18].

More specifically, a new statistically feasible profile, made specific to a particular deformation type, location and scale, is calculated by perturbing the mean training profiles by a weighted sum of the covariance of the corresponding training profiles using the following model of medial profile variations

$$p_{dls} = \bar{p}_{dls} + \mathbf{M}_{dls}\alpha_{dls} \quad (1)$$

where  $p_{dls}$  is the resulting shape profile,  $d$  is the deformation profile type ( $L(m)$ ,  $R(m)$ ,  $T^l(m)$ , or  $T^r(m)$ ),  $l$  and  $s$  are the specified location and scale values,  $\bar{p}_{dls}$  is the average medial profile for the specified type ( $L(m)$ ,  $R(m)$ ,  $T^l(m)$ , or  $T^r(m)$ ), the columns of  $\mathbf{M}_{dls}$  encode the main variation modes for the specified profile type, and  $\alpha_{dls}$  are weights of the variation modes and are typically limited to  $\pm 3$  standard deviations.

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

$$p_d = \bar{p}_d + \sum_{l=1}^N \sum_{s=1}^{N-l+1} \mathbf{M}_{dls}\alpha_{dls}. \quad (2)$$

The upper limit of the inner summation of (2) indicates that all locations have PCA performed at them for all scales as long as the combination of locations and scales is valid.

In summary, varying the weights of one or more of the variation modes alters the length, orientation, or thickness profiles and generates, upon shape reconstruction, statistically feasible




Fig. 6. Overlay plot of (a) thickness profiles, measured in pixels, and (b) orientation profiles, measured in radians. Results are from a training set of 50 CC shapes, with both profiles plotted as a function of the medial-node number,  $m = 1 \dots 100$ . Mean profiles appear in bold. Adapted from [18].

stretch, bend, or bulge deformations at specific locations and scales.

Each scale and location provides a unique measurement for the analysis of a shape. Examples of the global (maximal scale) thickness and orientation samples with their corresponding averages are shown in Fig. 6. These examples also serve to illustrate how visualization of these profiles can provide better understanding of the variations present in the data by identifying contributions of different types of deformations along the anatomical shape. Notice, for example, that towards the center of the CC (value range 30–50 on the horizontal axis of the plots) the bulk of the variation lies in the thickness profile, something not readily interpreted from results with classical PCA approaches (i.e., global PCA on boundary, or implicit, shape models, as opposed to localized PCA on a medial-axis based shape model). With classical PCA approaches, one would capture the variation in the CC's shape, but could not readily interpret it as being mostly a change in thickness instead of a bending or stretching of the object, nor could one readily identify exactly where the variance occurs on the model.

We note, however, that although each scale and location provides a unique measurement for analysis, there is a degree of redundancy in the representation when building a statistical model of deformations. Each profile is a real-valued vector of length  $N$ , and thus lies in  $\mathbb{R}^N$ . Any orthogonal set of  $N$  vectors is a basis for  $\mathbb{R}^N$ , and by definition any set larger than  $N$  must have vectors that are linear combinations of other members of the set. As such: 1) any large-scale deformation can be represented as the union of multiple smaller-scale deformations and 2) the entire  $N \times 4 = 400$  dimensional space of deformations can be captured via a global, explicit PCA (over the whole shape, i.e., by setting  $l = 1$  and  $s = N$ ), assuming all the PCA modes are retained in the model and no limits are placed on the amount of variation (typically measured as multiples of standard deviations) along each mode. In other words, with enough modes of variation, any single scale of deformations can capture the entire shape space.

By design, global PCA seeks a subset of vectors that maximally represents the variability in the data, assuming the data follows a Gaussian distribution. A significant drawback of this

global approach is that varying the weight of a single variation mode generally causes all the landmark positions to change (Fig. 1). In other words, although the global model produces feasible shape deformations only, a desirable trait, it generally produces global deformations over the entire object, making it difficult to guarantee fitting one part of the shape without unfitting another part. In contrast, setting  $s = 37$ , for example, produces 64 possible locations (from location  $l = 1$  to  $l = N - s + 1 = 64$ ). Each location has 37 modes of variation for a total of  $64 \times 37 = 2368$  variables. Clearly only  $N = 100$  of these vectors can be orthogonal, but no two vectors need to point in the same direction. Hence each vector can represent a unique and statistically valid localized deformation of the shape model for a particular anatomical region.

These statistically valid and localized deformations can be used effectively during the segmentation process, as will be better detailed in Section III. For example, suppose our shape model fits perfectly over all the corpus callosum (CC) except for a misalignment around the splenium region of the CC (Fig. 4). In this scenario, the shape of the model must be adjusted only around the splenium without affecting any of the correctly aligned parts. This is simply accomplished by setting  $l = 1$  and  $s = 37$  to localize the deformations around the splenium part only and ensure plausible deformations respecting the variations observed in the training set in that region only. A compact yet global PCA ( $s = 100$ ) is unlikely to include a mode of variation that directly accomplishes this task. Instead some complex combinations of several global modes must be used to ensure the desired localized effect.

The trade off is clear, finer scale deformations mean more variables that will need to be optimized, but the statistics will better represent localized changes in shape. Larger scale deformations have fewer variables but sacrifice localized changes to preserve global ones, and thus it is difficult to fit one part of the model without unfitting another part. We can exploit the advantages of both options during our segmentation procedure by starting with larger scale deformations, and then moving to increasingly finer scale deformations (Section II-D for details). By starting with large scale deformations, we can find the best global fit of the model to the image data without wasting time on localized changes that represent a small portion of the overall error. Once the global model is fitted, localized deformations can refine anatomical subregions and thus give a better fit to localized variability in the data (results of this can be seen in Fig. 10).

The preceding discussion outlined how to learn our shape model through examples. Additional details on HRPCA are available in [18] for the interested reader. Next we describe how to build a fitness function that can be used to characterize how well the shape model has been fitted to the image. The fitting procedure itself, performed using GA, is described in Section II-D. Essentially we need to discover how to deform the shape (what pose parameters and modes of variation,  $\alpha_{dls}$  to activate) in order to fit the shape to the image. The GA optimization explores the space of these deformations and finds the ones that best fit the model to the data (those that score highly according to the objective function).

## B. Building a Fitness Function

The fitness function is used for ranking the population in order to determine which members should reproduce and which should survive. It is akin to the energy functional in typical deformable models, in that it is the objective function being optimized. As noted in Section I, the use of GA allows us to easily adapt the function to any given task including both convex and nonconvex prior shape and image-based knowledge; something globally-optimal energy-minimizing models do not generally allow for [41]. Rather than focusing on easily described and optimized terms, fitness functions can be constructed that directly reflect the characteristics thought to be exhibited by the target structure. The general form of our energy function is then a weighted combination of the desired terms. For example, we can adopt the fitness function  $\text{Fit}_1(i)$ , where  $i$  is a potential segmentation, to consider the mean and standard deviation of the image intensity enclosed by the shape's boundary, the average edge strength along said boundary, and the learned size of the CC

$$\begin{aligned} \text{Fit}_1(i) = & w_1 \exp\left(\frac{-F_1(i)}{2\sigma_1^2}\right) + w_2 \left(1 - \exp\left(\frac{-F_2(i)}{2\sigma_2}\right)\right) \\ & + w_3 \left(1 - \exp\left(\frac{-F_3(i)}{2\sigma_3}\right)\right) \\ & + w_4 \exp\left(\frac{-F_4(i)}{2\sigma_4}\right) \end{aligned} \quad (3)$$

where

$$\begin{aligned} F_1(i) &= \left(\mu - \sum_{x=1}^{\mathcal{N}} q(x)\right)^2 \\ F_2(i) &= \frac{\sum_{x=1}^{\mathcal{N}} \|\nabla I(x)\| \|\nabla q(x)\|}{\sum_{x=1}^{\mathcal{N}} \|\nabla q(x)\|} \\ F_3(i) &= \frac{\sum_{x=1}^{\mathcal{N}} I(x)q(x)}{\sum_{x=1}^{\mathcal{N}} q(x)} \\ F_4(i) &= \sqrt{\frac{\sum_{x=1}^{\mathcal{N}} (I(x) - F_3(i))^2}{\sum_{x=1}^{\mathcal{N}} q(x)}} \end{aligned} \quad (4)$$

$x_1 \dots x_{\mathcal{N}}$  are the indices of the  $\mathcal{N}$  pixels,  $q(x)$  is a binary image of the reconstructed shape (Section II-A), and  $\mu$  and  $\sigma_1$  are the average and standard deviation, respectively, of the sizes of the CCs learned from the training set.  $I$  is the image, and  $\sigma_2, \sigma_3, \sigma_4$  are the learned averages in edge strength, mean image intensity, and image intensity standard deviation. Hence,  $F_1(i)$  represents the conformance to the learned area of shape  $i$ ,  $F_2(i)$  the average gradient magnitude along the shape's boundary,  $F_3(i)$  the mean image intensity enclosed by the shape's boundary, and  $F_4(i)$  the standard deviation of the intensity enclosed by the shape's



boundary. Finally,  $w_1, w_2, w_3$  and  $w_4$  are scalar weights controlling the importance of each term in the segmentation process.

Alternatively, we can adopt a more traditional energy functional commonly used in the deformable model community

$$\text{Fit}_2(i) = - \sum_{x=1}^{\mathcal{N}} (w_1(f(x)q(x) + g(x)(1 - q(x))) + w_2h(x)|\nabla q(x)|) + w_3F_1(i) \quad (5)$$

where  $f = -\log(P_{\text{obj}}(I))$ ,  $g = -\log(P_{\text{bk}}(I))$ , for object and background likelihoods  $P_{\text{obj}}$  and  $P_{\text{bk}}$ ,  $h = 1/1+|\nabla I|^p$ , and  $F_1$  is the shape term as defined in (4) (see [9] for details). This functional has recently been shown to be convex when defined over convex shape-spaces [9], and can therefore be globally optimized in that scenario. Using this functional is therefore a good way to investigate the fidelity trade-off between convex and nonconvex shape-spaces, since we can compare the accuracy of the guaranteed global optima under convex shape-spaces to the potentially global optima found using the proposed nonconvex shape-space. Since, with GA, we seek to define a maximization problem, the function of [9] has been negated, thus making it a concave optimization problem over a nonconvex shape-space.

Depending on the image segmentation task at hand, one functional will likely perform better than the other. We simply illustrate here that novel, as well as traditional, functionals can be optimized with GA.

### C. Genetic Algorithms

GA are an optimization technique that models our understanding of evolution. In essence, a number of simultaneous potential solutions (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 regions of the search space with exceptionally ‘‘fit’’ members (good solutions). A few important choices are made during any application of GA: 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, etc.), how to perform those crossovers (uniform, single-point, multipoint, etc.), and finally what fitness function to use for evaluation. Though these choices seem complex, we found through experimentation that different genetic operators yielded negligible impact on our results. The interested reader is referred to [61] for an in-depth review of GA.

### D. Genetic Algorithms for HRPCA

An overview of our GA is given in Algorithm 1 including a summary of all the free parameters and the associated settings that were used for all results presented in Section III. Fixed values were set by observing results on the training data. For example, the maximum number of iterations,  $\text{MaxIter}$ , was set by noting that the algorithm always converged within 500 iterations.

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

1) *Population Representation*: In medical image segmentation using GA, the individuals forming the population represent potential shapes of the target structure, each having some level of accuracy measured by the fitness function (Section II-B). We require a shape representation 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-node position, and scale (Section II-A).

2) *Encoding the Shape Parameters of the Population for GA*: We use chromosomes to represent the set of all the weights of the principal components as obtained from the HRPCA, where each gene represents a weight (as a floating point number) for a particular deformation, location, scale, and mode of variation (Fig. 7, ). In total, there are at most  $\sum_1^d \sum_{l=1}^{\mathcal{N}} \sum_{s=1}^{\mathcal{N}-l+1} s = d*(1/6)(2\mathcal{N}+3\mathcal{N}^2+\mathcal{N}^3)$  weights available for mutation since, for each of the four deformations,  $d = 4$ , we have  $\mathcal{N} = 100$  different locations, but for each location,  $l$ , we can only have up to  $\mathcal{N} - l + 1$  scales,  $s$ , each of which has at most  $s$  weights for the  $s$  principal components (Fig. 5). In our application, this would add up to 686, 800 dependent variables for our model, which motivates the need for GA with a coarse-to-fine approach to search the highly multivariate space. The large number of variables is based on the assumption that for each deformation type, location, and scale, all modes of variation are used. However, in practice we found that only a maximum of five modes of variation for each deformation type at each location and scale are needed to account for at least 90% of the variance, with at most  $s$  modes for  $s \leq 5$ . As previously discussed, there is a level of redundancy in the HRPCA space whereby any large scale deformation can be achieved by a combination of smaller scale deformations. Although smaller scale deformations have more valid locations and therefore more variables to optimize, they are desirable in that they better represent smaller, localized changes in the model. To avoid optimizing a needlessly high number of variables early on, we slowly decrease the scale of the deformation in a coarse-to-fine manner, and never optimize over more than one scale at a time (see Section II-D3 for details).

3) *Mutations and Crossovers*: As previously discussed, GA use mutations to walk randomly around the search space and crossovers to initialize the search in new locations that result from combinations of pairs of current solutions.

With regards to mutations, in order to steer the evolution in a coarse-to-fine manner and, thereby, facilitate faster initial convergence, we employ a coarse-to-fine approach. Initially we constrain the mutations to the affine transformation parameters: base-node position (translation values)  $(t_x, t_y)$ , model orientation angle  $\theta$ , and scale values  $(s_x, s_y)$  (Fig. 7, top). Since our initial shape is the mean CC (obtained by setting all weights,  $\alpha_{\text{dis}}$  in (2), 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




Fig. 7. Segmenting an anatomical structure amounts to finding the optimal set of shape parameters. (Top) In our GA implementation (Section II-D1), we represent each shape as a chromosome with genes encoding affine and statistical shape deformation parameters. Affine transformation parameters are those encoding the global rotation, scale, and position of the shape. Statistical shape deformation parameters are those representing weights for a particular deformation  $d$ , location  $l$ , and scale  $s$ . As noted in Fig. 5, Section II-D2, endmost locations have only a restricted scale. (Middle) Mutation (Section II-D3) is performed by altering some weights of the HRPCHA. (Bottom) Crossover amounts to swapping a set of weights between two individuals.

score for a good location, scale and orientation, simply because of a bad random shape mutation.

With an adequate location, scale, and orientation obtained, we allow the mutations to begin including shape deformations (Fig. 7, middle). Dynamic mutation of a single gene amounts to altering the corresponding weight by sampling it from a uniform random distribution under the constraint that the total weight lies within  $\pm 3$  standard deviations of the corresponding mode of variation (square root of the explained variance obtained in HRPCHA). Modifying a weight will change the medial profiles and hence the reconstructed shape boundary (Section II-A).

During the initial phases of the evolution, every member of the population undergoes a random deformation with a 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 Fig. 4) will not help until an acceptable global fit is obtained. Consequently, after no noticeable change in population fitness occurs, 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), and then smaller deformations which amount to small variations in local regions (Fig. 8). This process is outlined in algorithm Algorithm 1, where the *Evolve* function is sequentially called with decreasing scale.

In essence, GA 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 where under both parents have an equal probability of making the contribution (Fig. 7, bottom and Fig. 9).

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 to randomly select members for reproduction. Our goal is to ensure that the




Fig. 8. Examples of three sequentially performed mutations on four different members of the population. The scale of the mutation is decreased, progressing from left-to-right top-to-bottom, from affine, to global scale, to medium scale, and finally small scale. Note that the magnitude of each deformation is entirely random, and only the range over which the deformation operates is decreasing. In all images, the red shape was mutated into the green, and then green was mutated into blue.




Fig. 9. Example crossover operations. Each image shows an overlay plot of the two parents and the resulting offspring. The figure shows parents in red and blue, with offspring in green. Note how the offspring carries shape properties from both parents and thus appears somewhat “in between” the two parents.

fittest members of the population have the highest probability of selection in proportion to how much their fitness values dominate the rest of populous. Each member  $i$  of a population  $\mathcal{P}$  has a probability of selection equal to

$$P_{\text{selection}}(i) = \frac{\left( \text{Fit}(i) - \min_{k \in \mathcal{P}} \text{Fit}(k) \right)}{\sum_{j \in \mathcal{P}} \left( \text{Fit}(j) - \min_{k \in \mathcal{P}} \text{Fit}(k) \right)} \quad (6)$$

where  $\text{Fit}(i)$  is the fitness function of the member  $i$  (Section II-B), and  $\min_{k \in \mathcal{P}} \text{Fit}(k)$  is the minimum fitness value of the population. The normalization (subtracting of the minimum fitness value and dividing by the summation) in (6) is performed to ensure valid probabilities for even negative valued fitness functions. We experimented with different methods for determining probabilities from negative fitness functions, and noted only negligible impact on our results. We also employ an “elitist” strategy under-which the best member of the population is always maintained, and the weakest of the set are replaced by the new individuals resulting from the crossover operation. A set number of crossover operations are performed during each iteration of the genetic algorithm. We empirically set the number of crossover operations per iteration to 9 for all results, with different experimental values yielding negligible impact.

### III. RESULTS AND DISCUSSIONS

We demonstrate our work through its application to CC segmentation in mid-sagittal brain magnetic resonance images (MRI). The CC is a bundle of white matter fibre connecting the




Fig. 10. Example segmentations using GA-HRPCA. (left) Finding the pose for one example image, showing ten evenly sampled members from the population, in terms of fitness. Red indicates the fittest model moving linearly to green as the worst of the ten. (right) Three example segmentation (six rows) results progressing left to right, showing fittest individual after population initialization (left), global deformations (next two), and local deformations (last two). Odd image rows show a cropped plot overlaid on the MRI, with the medial axis highlighted in red. Cropping is only for visualization purposes, and all results are produced on full images. Even image rows show color images with the red channel set equal to the ground truth segmentation, the green channel the GA-HRPCA segmentation, and all zeros for the blue channel, i.e., green for true-negatives, yellow for true-positives, red for false-negatives, and green for false-positives. Below each example segmentation, error  $\varepsilon$  is shown.




Fig. 11. Plot of fitness of the best individual versus generation number for the three sample MRI images from Fig. 10, for 24 different runs per image. The slope of the curve demonstrates a rapid increase in fitness as the CC is located and the shape model heavily deformed (first 100 generations) and a more gradual increase as the shape model is manipulated to an increasingly fine degree.

left and right brain hemispheres, and whose cross sectional area and shape correlate with neurological diseases such as multiple sclerosis and schizophrenia [62]–[64]. The vast amount of data acquired in clinical studies of these diseases motivates the necessity of a reliable CC segmentation algorithm.

We conduct all of our experiments on a set of 50 mid-sagittal magnetic resonance images [65], presenting both qualitative and quantitative results of the fully-automatic segmentation of the CC using our GA driven, statistically-constrained deformable models. All experiments are performed using leave-one-out validation for training the shape and intensity priors. In order to discount pose problems for some of the methods tested, all of the data has been affinely registered. We do, however, still search over pose in our GA-HRPCA method. As our tested energy functions are all affine invariant, it should be clear that including or not including the affine registration step for novel images does not change the results for GA-HRPCA. The affine registration step does, however, allow us to obtain globally optimal solutions to the method in [9], which we will use for comparison.

For clarity, we divide our experiments into two sets: those dealing with the strict validation of GA-HRPCA in its ability to segment the CC; and those exploring the aforementioned

trade-off between energy function and shape space fidelity, and convexity. All error results in each set are computed using the Jaccard distance:  $\varepsilon = 1 - (|Auto \cap Manual|) / |Auto \cup Manual|$ , where  $\varepsilon = 0$  represents the ideal segmentation,  $|z|$  denotes the area of  $z$ , and *Auto* and *Manual* denote the automatic segmentation and the manual expert delineation, respectively [66].

#### A. Experiment Set A: Validating GA-HRPCA

We run a total of 500 evolution iterations per image, which takes about 8.5 min per image on a 2 GHz AMD Opteron, implemented in MATLAB (2010a, The MathWorks, Natick, MA). The most expensive operation is the crossover operation since, as previously mentioned, there are a larger number of dependent variables that must be copied between the models. As already noted, in order to speed up convergence, we employ a coarse-to-fine approach where deformations progress from larger to increasingly smaller scales. Again, the basic idea is that initially there is little need for fine, small scale model adjustments since the model may not be in the right position at all. This process is depicted in Figs. 10 and 11. Notice that although the finer scale deformations represent a small amount of the global shape variance, they allow a significant reduction in




Fig. 12. Example segmentations performed using GA-HRPCA  $\text{Fit}_2 - B$ . Six different images are shown, progressing from least-to-most segmentation error under evenly spaced samples. Note how even the worst segmentation is still a plausible CC segmentation. Top row shows segmentations overlaid on cropped MRI, with  $\epsilon$  underneath. Bottom row shows color images with the red channel set equal to the ground truth segmentation, the green channel the GA-HRPCA segmentation, and all zeros for the blue channel, i.e., green for true-negatives, yellow for true-positives, red for false-negatives, and blue for false-positives.

TABLE I  
ERROR RESULTS FOR OUR PROPOSED GA-HRPCA. RESULTS FOR ALL BUT THE LAST THREE COLUMNS ARE FIRST AVERAGED ACROSS ALL RUNS OF EACH IMAGE, AND THEN STATISTICS ARE COMPUTED ACROSS THOSE VALUES FOR 50 IMAGES

Method	mean $\epsilon$	median $\epsilon$	min $\epsilon$	max $\epsilon$	std. dev. $\epsilon$	mean std. dev. $\epsilon$	max std. dev. $\epsilon$	std. dev. of std. dev. $\epsilon$
GA-HRPCA $\text{Fit}_1$	0.176	0.161	0.081	0.360	0.060	0.022	0.073	0.012
GA-HRPCA $\text{Fit}_2 - A$	0.211	0.197	0.093	0.495	0.071	0.023	0.118	0.023
GA-HRPCA $\text{Fit}_2 - B$	0.202	0.192	0.116	0.347	0.055	N/A	N/A	N/A
GA-HRPCA $\text{Fit} = -\epsilon$	0.031	0.029	0.018	0.052	0.008	0.006	0.010	0.002

error by reducing the localized errors in the fitting that remained after the global fit (second image in each row in Fig. 10).

We conducted three main experiments with GA-HRPCA. Since GA are randomized algorithms and may not always produce the same result, we run the algorithm 24 times per image. Our first experiment was performed using  $\text{Fit}_1$ , (3). We learn the optimal  $\mathbf{w}$  for  $\text{Fit}_1$  and  $\text{Fit}_2$  using a method based on [67], [68]. The optimal weights for  $\text{Fit}_1$  where  $w_1 = 0.0090$ ,  $w_2 = 0.1101$ ,  $w_3 = 0.8809$ ,  $w_4 = 0$ . Note that  $F_1, \dots, F_4$  are strictly positive for positive valued images, and therefore each term in  $\text{Fit}_1$  is normalized between 0 and 1. Setting  $w_4 = 0$  means that through training we learned  $F_4$  was misleading the fitness function, and was unable to accurately predict the standard deviation of the CC. Our second experiment was performed using  $\text{Fit}_2$  in (5), where the optimal weights for  $\text{Fit}_2$  were:  $w_1 = 0.0145$ ,  $w_2 = 0.9707$ ,  $w_3 = 0.0148$ . We note that for  $\text{Fit}_2$ ,  $w_2 \gg w_1$  does not imply that one term has vastly more influence than the other because their scales are not normalized. From here on we refer to this experiment with  $\text{Fit}_2$  as  $\text{Fit}_2 - A$ . Our third experiment was also performed using  $\text{Fit}_2$ , but we used the fittest result from all 24 runs for each image. In essence, we increase the population size of the GA by a factor of 24. We refer to this experiment with  $\text{Fit}_2$  as  $\text{Fit}_2 - B$ . The same experiment was conducted using  $\text{Fit}_1$  without any noticeable improvement, and we defer discussion until the end of this section. To give a visual impression of the segmentation error we present some results of  $\text{Fit}_2 - B$  in Fig. 12, evenly sampled from least to most error. We note that even the worst result, far-right, is still a very plausible CC segmentation.

For point of interest, we also conducted a fourth experiment to validate that our GA-HRPCA shape model can fit each ground truth segmentation without error. Simply put, we segmented the ground truth segmentations directly, using  $\text{Fit} = -\epsilon$ . The average error was indeed  $\epsilon \approx 0$ , validating that the GA-HRPCA model can accurately segment CC data, given an ideal energy function.

Results for our three main experiments are shown in Table I. Since GA are randomized, and we segment each image 24 times, we first compute expected error values for each image by averaging the error results from all 24 runs. The 50 resulting expected error values can then be used to compute statistics in expected error across images, as is done in Table I. However, in the last three columns of Table I, the standard deviation (std. dev.) is calculated across all 24 runs, and then statistics (mean, max, and std. dev.) are calculated over the 50 images to demonstrate that even though GA are a randomized method, the average std. dev. is still low (i.e., for a given image, there is little change in its error between runs). These numbers cannot be computed for  $\text{Fit}_2 - B$  because there is effectively only one run of that method. Notice that the std. dev. is about twice as large across images than within each image, for example 0.060 versus 0.022 for  $\text{Fit}_1$ . This indicates that the variation in error as a result of GA-HRPCA being a randomized algorithm is small in comparison to the variation caused by how accurately the energy function itself reflects the desired image properties. From that we conclude that in order to improve segmentation results, a new energy function should be sought as opposed to attempting to improve the optimizer or, in our case, *tweak* the GA. Note that we specifically say energy function here, not shape space, because we have already shown the HRPCA shape model can fit any of the ground truth CCs ( $\text{Fit} = -\epsilon$ ).

There is, however, an outlying case where the std. dev. across numerous runs of the GA for a specific image is larger than normal for  $\text{Fit}_2 - A$  (at 0.118 it lies at over 4 std. dev. from the mean of 0.023). This explains the significant reduction in maximal error obtained when moving from  $\text{Fit}_2 - A$  to  $\text{Fit}_2 - B$ . The increase in accuracy indicates that some local optima were obtained in  $\text{Fit}_2 - A$ , but can be surpassed with a more exhaustive search. This raises the question of whether we are better off with guaranteed global optima from convex functions, or possibly local optima from our GA-HRPCA method. In Section IV we conduct a few experiments with global and nonglobal methods to explore this question.






Fig. 13. Visualization of Geocuts and RW seeding percentages. Images demonstrate the increasing degree of constraint, in the form of additional seed points, required to prevent these simple functionals from producing highly erroneous segmentations. Top row, seeds with seeding percentage listed (See text for details). Middle row (a)–(d), Geocuts segmentations for corresponding seed percentages. Bottom row (e)–(h), RW segmentations. Notice leakage into fornix dip (below the CC) in both methods, due to lack of shape information. As the constraints get tighter, the possible set of segmentations shrinks to exclude the fornix.

It is interesting to note that the minimal error increases from  $\text{Fit}_2 - A$  to  $\text{Fit}_2 - B$ . This is because in one case improving fitness leads to increased segmentation error,  $\varepsilon$ , implying the fitness does not directly correlate with accuracy. Though it may seem an odd result, this is not uncommon for both convex or nonconvex functions used in practice for segmentation.

Comparing our two potential fitness functions,  $\text{Fit}_1$  and  $\text{Fit}_2$ , Table I shows that  $\text{Fit}_1$  obtained superior results to  $\text{Fit}_2 - A$  and  $\text{Fit}_2 - B$ . Interestingly,  $\text{Fit}_1$  did not noticeably benefit from a more exhaustive search. We believe  $\text{Fit}_1$ 's better performance stems from its simpler assumption of an exponential relationship for intensity having fewer parameters to estimate and, therefore, being more resilient to noise than the potentially over-fit learned likelihoods in  $\text{Fit}_2$ . The maximal std. dev. and std. dev. of std. dev. are also much lower for  $\text{Fit}_1$ , and this explains why a similar reduction in error was not observed for  $\text{Fit}_1$  when conducting a more exhaustive search.

### B. Experiment Set B: Fidelity Versus Convexity

We segment the same set of 50 CC images under a few popular frameworks.

- 1) Geodesic active contours (GAC) [40]: Gradient descent performed on a nonconvex level-set-based deformable model with no shape prior. The energy functional itself is convex, and is modeled by the second term in (5),  $h(x) |\nabla q(x)|$ . Nonconvexity is induced by the level-set shape representation (i.e., nonconvex domain).
- 2) Geocuts [13]: A globally optimal solution to the GAC energy using graph cuts.
- 3) Random Walker (RW) [15]: A submodular (convex) graph-based formulation with a globally optimal solution and no shape information.

- 4) Cremers 2008 [9]: A globally optimal segmentation method using  $\text{Fit}_2$ , (5), with a global shape prior<sup>2</sup> implemented via PCA (i.e., convex, implicit, global statistics instead of nonconvex, explicit, localized ones).

We note that although Geocuts and RW do not have explicit shape priors they still impose a specific form of shape constraints; the seed points function as hard constraints in both methods that shrink the set of possible segmentations and thus form an allowable shape space, albeit a user created one. It is important to note that as Geocuts and RW are interactive methods, and our method is fully-automatic, we are not validating against them, but rather using them to explore the results of minimizing convex functions with varying degrees of shape constraints. For exploration, we choose to use random seeds drawn from the ground truth foreground and background, and we report the percentage (%) of pixels used as seeds in each of our experiments (Fig. 13). For example, a 1% seeding means that, on average, one seed is placed in every  $10 \times 10$  pixel region of the image. The user could draw better seeds, i.e., fewer seeds that would also produce less segmentation error. However, this is essentially providing a greater constraint (fewer seeds in number, higher in tightness, i.e., more restricted shape space). We also note that although Geocuts and RW do not have weights between competing terms, they each have an inherit parameter. The Geocuts method has the free parameter  $p$  in  $h = 1/1+|\nabla I|^p$ , and RW has the  $\beta$  parameter which controls the mapping from intensity differences to edge weights (see [15] for details). We ran each method over a large set of parameter values. Thus for each method, we have optimized its free parameters to minimize differences in accuracy stemming from a lack

<sup>2</sup>This changes the third term slightly, but both will end up being Gaussian priors centered on the mean shape, with learned variation from PCA.



TABLE II  
FIDELITY VERSUS OPTIMIZABILITY. SEE TEXT FOR DETAILS

Method	mean	median	min	max	std. dev.
	$\epsilon$	$\epsilon$	$\epsilon$	$\epsilon$	$\epsilon$
GAC [40]	0.478	0.433	0.094	1.000	0.297
Geocuts [13] 0% seeding	0.375	0.370	0.094	0.981	0.232
Geocuts [13] 5% seeding	0.182	0.171	0.089	0.346	0.061
Geocuts [13] 10% seeding	0.160	0.149	0.086	0.309	0.057
RW [15] 0% seeding	0.844	0.933	0.288	0.970	0.170
RW [15] 5% seeding	0.165	0.152	0.106	0.407	0.055
RW [15] 10% seeding	0.132	0.117	0.070	0.343	0.057
Cremers 2008 [9]	0.376	0.319	0.158	0.895	0.186
GA-HRPCA $Fit_1$	0.176	0.161	0.081	0.360	0.060
GA-HRPCA $Fit_2 - A$	0.211	0.197	0.093	0.495	0.071
GA-HRPCA $Fit_2 - B$	0.202	0.192	0.116	0.347	0.055

of parameter tuning, and instead highlight the best results that each method can obtain. To that same end, we initialize GAC with the ground truth segmentation, thereby placing it under optimal conditions. For comparison purposes that will become clear shortly, we also minimize the GAC energy function using GA-HRPCA, i.e.,  $Fit_2$  with  $w_1 = 0$ ,  $w_2 = 1$ ,  $w_3 = 0$ , we denote this method as GAC-GA-HRPCA.

Our reasons for choosing the above methods are to compare and contrast the consequences of each competing method's trade-offs with the trade-offs present in GA-HRPCA. These trade-offs are highlighted in Table II, where the optimizability, shape flexibility, and shape fidelity are measured. Note that for Geocuts and RW we report multiple entries (rows) each using a different seeding percentage. This enables us to explore how the energy functions behave under increasing amounts of constraint.

Shape flexibility measures the method's ability to fit any valid CC. To measure flexibility, we run each method on the ground truth segmentations directly, and measure the average error using  $\epsilon$ . Low values indicate low error, and therefore high flexibility in that the method had no problem segmenting all 50 valid CCs. High values here are due to the shape space being too restrictive to accommodate the shape variability in the target CC. As one expects, methods without shape restrictions (GAC, Geocuts, and RW) have perfect flexibility, since they can segment any shape.

Shape fidelity, measures the methods ability to constrain the shape to only valid members of the anatomy, and is estimated as error, under measurement  $\epsilon$ , in segmenting novel images. We say estimated because, as reported, shape fidelity implicitly includes energy function error, and is thus an approximation. However, when comparing fidelity between equal energy functions under different constraints, it is clear that changes in accuracy must be entirely a result of the solver or the constraints, as the energy function itself remains fixed. A perfect score is  $\epsilon = 0$  in both shape columns.

To balance the comparison, we also include the level of manual constraints required, as otherwise fully-manual segmentation would be the best possible method according to the table. Manual constraints are rated on a scale from 0 to 10, 0 being no manual intervention, and 10 being the most of any method surveyed here. These ratings correspond directly to the percentage of the image that needed to be seeded for each image, as seeding required direct knowledge of the ground truth segmentation. Notice that for manual methods, increasing

fidelity requires increasing manual intervention. In automatic methods, increasing flexibility and fidelity comes at a loss of optimality, as nonconvex shape-spaces are required. We provide a more detailed review of these results below.

Examining the fidelity column of Table II in particular, we notice that the convex methods score poorly unless they are heavily constrained with manual intervention (i.e., high seeding percentage for Geocuts and RW). The high error of GAC, in combination with the fact that we use the ground truth as the initialization, implies directly that a simple convex functional, under few optimization constraints, yields a very high segmentation error. Simply put, there are too many erroneous segmentations with lower energy than the ground truth. This is, of course, expected. Notice, however, that with minimal constraints Geocuts yields almost as high a segmentation error as GAC. Again, this is not surprising; the global optima include many non-CC edges, while the local optima can get stuck before reaching increasingly erroneous segmentations. Had we initialized GAC randomly instead of using the ground truth segmentation, its error would have been far higher. Adding more constraints to Geocuts demonstrates that under appropriate constraints a global optimum to a convex function can yield good results. However, as presented in Fig. 13, the constraints need to get rather tight before erroneous segmentations are eliminated. Naturally, a region-based energy function, as is used in RW, suffers the same fate. We have thus far reached a rather expected result, global optima to simple energy functionals can yield good segmentations, but they require significant constraints (the seeding percentage has to be high). As manual constraints are undesirable, automatic constraints are preferable.

Tight shape constraints can be automated through the use of shape prior information, and statistically shape-constrained energy functionals can be solved, as with Cremers 2008. However, in order to introduce a shape constraint that is convex, i.e., can be globally optimized, the constraint is made simple (global-in-scale PCA). It is especially interesting to compare the results from Cremers 2008 with GA-HRPCA. The Cremers 2008 row of results in Table II compares the same energy function, with the same weights,  $w$ , to our GA-HRPCA method in  $Fit_2 - A$  and  $Fit_2 - B$ . We believe the Cremers 2008 error is high because not only is the shape space unable to address the variability present in the data (i.e., low shape flexibility), but numerous non-CC shapes are still allowable and have lower energy than other, better segmentations (i.e., low shape fidelity). Similar discussions apply when comparing Geocuts to GAC and GAC-GA-HRPCA, since all of the methods use the same energy functional. Under appropriate constraints GAC has low error, and therefore high fidelity, and our method is able to automatically enforce those constraints. We also note that rerunning Cremers 2008 with an increasing number of PCA variation modes had negligible impact on the results, because while using more modes increases shape space flexibility, it also reduces fidelity. More specifically, with more modes of variation, the convex, implicit, global shape space can include the ground truth segmentation and also other more erroneous segmentations, stemming from the fact that the true underlying shape space is not Gaussian (imagine fitting an ellipse to the letter "S," to include

TABLE III  
ERROR COMPARISON BETWEEN TESTED METHODS FOR 50 IMAGES

Method	optimizability	shape flexibility	shape fidelity	manual	automatic
		$\epsilon$	$\epsilon$	constraints	constraints
RW [15] 10% seeding	Global	0	0.132	10	N
Geocuts [13] 10% seeding	Global	0	0.160	10	N
RW [15] 5% seeding	Global	0	0.165	5	N
Geocuts [13] 5% seeding	Global	0	0.182	5	N
RW [15] 0% seeding	Global	0	0.844	0	N
Geocuts [13] 0% seeding	Global	0	0.375	0	N
GAC [40]	Local	0	0.478	0	N
GAC-GA-HRPCA	Local	0.03	0.270	0	Y
Cremers 2008 [9]	Global	0.14	0.376	0	Y
GA-HRPCA $Fit_1$	Local	0.03	0.176	0	Y

the entire letter much of the space outside the letter must also be included).

There is, thus, a trade-off present in all choices. Neither locally optimal nor globally optimal methods without shape information seem well-suited to fully-automatic segmentation of the CC. Shape information can be added, and the solution kept global, but the shape information must be kept convex, which may not accurately fit the desired shape space, without simultaneously allowing erroneous shapes. Alternatively, with HRPCA a more expressive shape constraint can be automatically built, but it cannot be solved optimally and thus necessitates GA in order to minimize sensitivity to initialization. A comprehensive set of results is presented in Table III.

Notice that the standard deviation between images for each tested method, last column of Table III, is in general much greater than the standard deviation across different runs of the same image for our GA-HRPCA method (Table I, column 7, mean std. dev.  $\epsilon$ ). This implies that one can expect a greater variation in accuracy by applying a given method to a new image, than from rerunning our GA on a specific image.

#### IV. CONCLUSION

We have developed a novel segmentation technique by exploring the main concerns associated with both traditional and statistically-based deformable models. By using GA to address the initialization and local minima problems associated with traditional energy-minimization techniques, we propose an alternative to convex optimization methods [9], [13], [15]–[17]. In doing so, we trade guaranteed global optimality for a shape space that more accurately fits the true allowable shape domain (localized, medial-based shape statistics), i.e., higher flexibility and fidelity. Using the same energy functional, we showed a significant improvement in accuracy over the globally optimized, convex, global, implicit PCA shape space, Cremers 2008 method. Consequently, in answer to our question posed in the introduction: Is it worth sacrificing the complexity of the energy functional, and thereby reducing its ability to address image variability, in exchange for global optimality, or can more complex functions and shape spaces be utilized with an approximate optimization method and achieve superior accuracy? In our application the sacrifice is not worth it. It is better to trade convexity with guaranteed global optimality for a more expressive functional and shape-space with GA-based optimization that achieves close to optimal solutions. We therefore believe

research into finding energy functions and shape spaces with high fidelity, with a de-emphasis on convexity, is a very important area of future research. While it is true that global solutions to any such functions would be desired, our results demonstrate that convexly approximating the nonconvex shape-space in the application of CC segmentation causes a higher error than approximating the global optima of the nonconvex space via GA. We imagine the difference will only increase for structures with more complex shape deformations between members of the population, as these structures will form increasingly non-convex shape spaces.

Furthermore, as its well known that pose introduces an inherent nonconvexity in shape-model-based deformable models, for example [9], we provide a method that directly handles the pose problem. In essence, the basic idea with GA is to avoid choosing one pose over another until there is enough evidence suggesting one pose is vastly superior. For example, suppose the energy function had two local optima with different pose but similar shape; a binary image containing a target shape, and a second similar structure with a different pose. Equation (6) ensures that our method would keep approximately half of its population at each pose initially, since population members under each pose would have equal probability of replacement. Our GA would then start shifting the bulk of the population towards whichever object better satisfies the fitness function once shape deformations were allowed. This is exemplified by Fig. 10, as numerous false pose configurations have similar response to the fitness function. However, the pose that corresponds with the CC has a much higher degree of fitness and so it quickly overtakes the focus of the search.

We have demonstrated how GA 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 members of the GA population to lie within the allowable shape domain, thus greatly reducing the search space traditionally associated with deformable models, but without sacrificing fidelity to the target anatomical shape space through global PCA and the inherent convex approximations.

Though other works have used GA to drive deformable models [44]–[46], [48], [51], [52], [54], to the best of our knowledge, no works have combined GA with statistical-based deformations in a way that yields intuitively constrained deformations (bends, bulges, and stretches) that render the results

---

**Algorithm 1** GA for HRPCA
 

---

**Require:** A set of images,  $\mathbf{I}$ , an HRPCA matrix,  $\mathbf{M}$ , a fitness function  $Fit$ , and GA variables:

```

1: Population size,  $PopSize = 32$ .
2: Number of crossovers-per-evolution,  $CrossNum = 9$ .
3: Maximum number of standard deviations to deform,  $MxStdDev = 2$ .
4: Lower- and Upper-bounds on the position of the reference node,  $x_1$ ,  $(Min_x, Min_y)$  and  $(Max_x, Max_y)$ . Always set such that the extent of the mean shape model lies within the image bounds.
5: Number of modes of variation to consider (for each  $d, l, s$ ),  $k = 5$ .
6: Population,  $\mathcal{P}$ , with members denoted  $\mathcal{P}(i)$  and base-node position for  $\mathcal{P}(i)$ :  $(Min_x, Min_y) \leq (\mathcal{P}(i).t_x, \mathcal{P}(i).t_y) \leq (Max_x, Max_y)$ ; model orientation angle  $0 \leq \mathcal{P}(i).\theta \leq \pi$ ; scale  $(0.6, 0.6) \leq (\mathcal{P}(i).s_x, \mathcal{P}(i).s_y) \leq (1.4, 1.4)$ ; and HRPCA weights  $\mathcal{P}(i).\alpha$ .
7: Number of iterations allowed without improvement,  $ChangeIter = 40$ .
8: Maximum number of iterations allowed,  $MaxIter = 500$ .
9: The best population member seen so far,  $best$ , with fitness  $best.fit$ .

10: procedure GA-HRPCA
11:   Randomly initialize  $\mathcal{P}$ , from a uniform distribution with  $\mathcal{P}.\alpha = 0$ .
12:    $\mathcal{P}, best = UpdateBest(\mathcal{P}, best, 0)$ 
13:    $cnt = 1$ 
14:    $\mathcal{P}, best, cnt = Evolve(\mathcal{P}(i).\alpha=0, cnt)$  ▷ Search only for top candidates in position, orientation, and scale.
15:    $\mathcal{P}, best, cnt = Evolve(s = N, cnt)$  ▷ Perform global deformations. There are  $4 \times k \times (N - s + 1) + 5 = 25$  free variables at this stage.
16:    $\mathcal{P}, best, cnt = Evolve(s = 37, cnt)$  ▷ Perform medium scale deformations. There are  $4 \times k \times (N - s + 1) + 5 = 1285$  free variables at this stage.
17:    $\mathcal{P}, best, cnt = Evolve(s = 18, cnt)$  ▷ Perform fine scale deformations. There are now  $4 \times k \times (N - s + 1) + 5 = 1665$  free variables at this stage.
18:    $\mathcal{P}, best, cnt = Evolve(s = 11, cnt)$  ▷ Perform finer scale deformations. There are now  $4 \times k \times (N - s + 1) + 5 = 1805$  free variables at this stage.
19:   return  $best$ 
20: end procedure

21: procedure UPDATEBEST( $\mathcal{P}, best, changed$ )
22:   Sort  $\mathcal{P}(i)$  in ascending order according to  $Fit(i)$ .
23:   if  $Fit(PopSize) \geq best.fit$  then
24:      $best = \mathcal{P}(PopSize)$  and  $best.fit = Fit(PopSize)$ 
25:      $changed = 0$ 
26:   else
27:      $\mathcal{P}(1) = best$  ▷ Replace the worst member of the population with the best.
28:      $changed = changed + 1$  ▷ Keep track of how many iterations since fitness last improved.
29:   end if
30:   return  $\mathcal{P}, best, changed$ 
31: end procedure

32: procedure EVOLVE( $Constraints, cnt$ ) ▷ While  $best.fit$  is improving and  $cnt \leq MaxIter$ , perform crossovers and mutations.
33:    $changed = 0$ 
34:   while  $(changed \leq ChangeIter)$  AND  $(cnt \leq MaxIter)$  do
35:     for  $i = 1$  to  $CrossNum$  do
36:       Select parents  $(a, b)$  with probability of selection  $P_{selection}(a)$  and  $P_{selection}(b)$ 
37:       Replace  $\mathcal{P}(i)$  with crossover on  $\mathcal{P}(a), \mathcal{P}(b)$ . ▷ Replaces the least-fit individuals.
38:     end for
39:      $\mathcal{P}, best, changed = UpdateBest(\mathcal{P}, best, changed)$ 
40:     for  $i = 1$  to  $PopSize$  do
41:       Mutate  $\mathcal{P}(i)$ , only modifying  $\mathcal{P}(i).\alpha$  with scale,  $s$ , described by  $Constraints$ .
42:     end for
43:      $\mathcal{P}, best, changed = UpdateBest(\mathcal{P}, best, changed)$ 
44:      $cnt = cnt + 1$ 
45:   end while
46:   return  $\mathcal{P}, best, cnt$ 
47: end procedure

```

---

more interpretable by clinicians and enable regional statistical analysis [18].

Our method can handle a variety of classical energy terms [9], [40], [69]. Other terms can be easily added related to texture, color or other image features. There is no requisite that the terms be convex or even differentiable; they simply need to properly reflect the characteristics of the appearance of the target object in the image data. Not requiring terms to be differentiable is an interesting side benefit of using GA.

Our method is also extendable to other segmentation problems. Specifically, given a training set of shapes for a different anatomical structure, one can perform skeletonization [70] followed by medial profile extraction and, subsequently, HRPCA. More complex topologies than the single medial-axis structure used in our experiments would be an interesting and

natural extension to our method. The largest challenge would be learning the underlying medial-structure, but once learned we believe there would be very little modification required to the GA-HRPCA algorithm. One potential modification would be to fit the primary medial-axis first, then extend down to the branches in a coarse-to-fine approach similar that used to explore different scales in Algorithm 1.

We are working on extending these ideas to 3-D where the genes of the population become the 3-D shape representation's parameters [71]. However, it remains to be explored how much added complexity must be introduced to address complicated 3-D objects with possibly varying topology using a medial representation, and whether a medial based shape representation will remain the ideal choice for an evolutionary computing strategy with localized deformations in 3-D. Another interesting

avenue is the extension to multilabel segmentation, where GA could be used to model numerous, interacting, simultaneously evolving populations, each representing a different anatomical structure.

#### ACKNOWLEDGMENT

The authors thank the reviewers and editors for their valuable feedback that resulted in an improved paper. The authors would also like to thank Dr. M. Shenton of the Harvard Medical School for providing the MRI data, and P. Plett for assisting with code development.

#### REFERENCES

- [1] M. Kass, A. Witkin, and D. Terzopoulos, "Snakes: Active contour models," *Int. J. Comput. Vis.*, vol. 1, no. 4, pp. 321–331, 1987.
- [2] T. McInerney and D. Terzopoulos, "Deformable models in medical image analysis: A survey," *Med. Image Anal.*, vol. 1, no. 2, pp. 91–108, 1996.
- [3] T. F. Cootes, D. Cooper, C. J. Taylor, and J. Graham, "Active shape models—Their training and application," *Comput. Vis. Image Understand.*, vol. 61, pp. 38–59, 1995.
- [4] T. F. Cootes, G. J. Edwards, and C. J. Taylor, "Active appearance models," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 23, no. 6, pp. 681–685, Jun. 2001.
- [5] M. Leventon, W. Grimson, and O. Faugeras, "Statistical shape influence in geodesic active contours," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2000, vol. 1, pp. 316–323.
- [6] S. Warfield, M. Kaus, F. A. Jolesz, and R. Kikinis, "Adaptive, template moderated, spatially varying statistical classification," *Med. Image Anal.*, vol. 4, no. 1, pp. 43–55, 2000.
- [7] D. Cremers, T. Kohlberger, and C. Schnörr, "Nonlinear shape statistics via kernel spaces," in *Pattern Recognition (Proc. DAGM)*, ser. LNCS, B. Radig and S. Florczyk, Eds. Munich, Germany: Springer, 2001, vol. 2191, pp. 269–276.
- [8] N. Paragios, M. Taron, X. Huang, M. Rousson, and D. Metaxas, "On the representation of shapes using implicit functions," in *Statistics and Analysis of Shapes*, ser. Modeling and Simulation in Science, Engineering and Technology, H. Krim and A. Yezzi, Eds. Boston, MA: Birkhäuser, 2006, pp. 167–199.
- [9] D. Cremers, F. Schmidt, and F. Barthel, "Shape priors in variational image segmentation: Convexity, Lipschitz continuity and globally optimal solutions," in *IEEE Conf. Comput. Vis. Pattern Recognit.*, Jun. 2008, pp. 1–6.
- [10] N. Vu and B. Manjunath, "Shape prior segmentation of multiple objects with graph cuts," in *IEEE Conf. Comput. Vis. Pattern Recognit.*, Jun. 2008, pp. 1–8.
- [11] L. Elsgolc, *Calculus of Variations*. Oxford, U.K.: Pergamon, 1962.
- [12] S. M. Pizer, "Guest editorial—Medial and medical: A good match for image analysis," *Int. J. Comput. Vis.*, vol. 55, no. 2–3, pp. 79–84, 2003.
- [13] Y. Boykov and V. Kolmogorov, "Computing geodesics and minimal surfaces via graph cuts," in *Proc. 9th IEEE Int. Conf. Comput. Vis.*, Oct. 2003, vol. 1, pp. 26–33.
- [14] Y. Boykov and G. Funka-Lea, "Graph cuts and efficient N-D image segmentation," *Int. J. Comput. Vis. (IJCV)*, vol. 70, no. 2, pp. 109–131, 2006.
- [15] L. Grady, "Random walks for image segmentation," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 28, no. 11, pp. 1768–1783, Nov. 2006.
- [16] M. Nikolova, S. Esedoglu, and T. F. Chan, "Algorithms for finding global minimizers of image segmentation and denoising models," *SIAM J. Appl. Math.*, vol. 66, no. 5, pp. 1632–1648, 2006.
- [17] X. Bresson, S. Esedoglu, P. Vanderghenst, J.-P. Thiran, and S. Osher, "Fast global minimization of the active contour/snake model," *J. Math. Imag. Vis.*, vol. 28, no. 2, pp. 151–167, Jun. 2007.
- [18] G. Hamarneh, R. Abu-Gharbieh, and T. McInerney, "Medial profiles for modeling and statistical analysis of shape," *Int. J. Shape Model.*, vol. 10, no. 2, pp. 187–209, 2004.
- [19] C. McIntosh and G. Hamarneh, "Genetic algorithm driven statistically deformed models for medical image segmentation," in *ACM Workshop Med. Appl. Genetic Evolutionary Computation Workshop (MedGEC)*, in Conjunction With the Genetic Evolutionary Computat. Conf. (GECCO), 2006.
- [20] G. Hamarneh and T. Gustavsson, "Statistically constrained snake deformations," in *IEEE Int. Conf. Syst., Man, Cybern.*, 2000, vol. 3, pp. 1610–1615.
- [21] T. Cootes, C. Taylor, D. Cooper, and J. Graham, "Training models of shape from sets of examples," in *Br. Mach. Vis. Conf.*, 1992, pp. 9–18.
- [22] T. Cootes, C. Taylor, A. Hill, and J. Halsam, "The use of active shape models for locating structures in medical images," in *Proc. 13th Int. Conf. Inf. Proces. Med. Imag.*, 1993, pp. 33–47.
- [23] L. Staib and J. Duncan, "Boundary finding with parametrically deformable models," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 14, no. 11, pp. 1061–1075, Nov. 1992.
- [24] K. M. Pohl, J. Fisher, S. Bouix, M. Shenton, R. W. McCarley, W. E. L. Grimson, R. Kikinis, and W. M. Wells, "Using the logarithm of odds to define a vector space on probabilistic atlases," in *Med. Image Anal.*, 2007, vol. 11, pp. 465–477.
- [25] P. Fletcher, C. Lu, S. Pizer, and S. Joshi, "Principal geodesic analysis for the study of nonlinear statistics of shape," *IEEE Trans. Med. Imag.*, vol. 23, no. 8, pp. 995–1005, Aug. 2004.
- [26] P. Sozou, T. Cootes, C. Taylor, and E. Di Mauro, "Non-linear point distribution modelling using a multi-layer perceptron," in *Br. Mach. Vis. Conf.*, 1995, pp. 107–116.
- [27] T. Cootes and C. Taylor, "A mixture model for representing shape variation," in *Br. Mach. Vis. Conf.*, 1997, pp. 110–119.
- [28] D. Cremers, T. Kohlberger, and C. Schnörr, A. Heyden, Ed. *et al.*, "Nonlinear shape statistics in Mumford-Shah based segmentation," in *Proc. Comput. Vis. (ECCV 2002)*, Copenhagen, Denmark, May 2002, vol. 2351, pp. 93–108.
- [29] S. Dambreville, Y. Rathi, and A. Tannen, "Shape-based approach to robust image segmentation using kernel PCA," in *IEEE Comput. Soc. Conf. Comput. Vis. Pattern Recognit.*, Jun. 2006, vol. 1, pp. 977–984.
- [30] P. Etyngier, F. Segonne, and R. Keriven, "Shape priors using manifold learning techniques," in *IEEE 11th Int. Conf. Comput. Vis.*, Oct. 2007, pp. 1–8.
- [31] D. Cremers, "Nonlinear dynamical shape priors for level set segmentation," *J. Sci. Comput.*, vol. 35, no. 2–3, pp. 132–143, Jun. 2008.
- [32] A. Ward and G. Hamarneh, "Gmat: The groupwise medial axis transform for fuzzy skeletonization and intelligent pruning," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 32, no. 6, pp. 1084–1096, Jun. 2010.
- [33] B. Ng, R. Abugarbieh, and G. Hamarneh, "Group MRF for fMRI activation detection," in *IEEE Comput. Vis. Pattern Recognit.*, 2010, pp. 2887–2894.
- [34] H. Blum, "Biological shape and visual science," *Theoretical Biol.*, vol. 38, pp. 205–287, 1973.
- [35] P. Dimitrov, J. Damon, and K. Siddiqi, "Flux invariants for shape," in *Proc. IEEE Comput. Soc. Conf. Comput. Vis. Pattern Recognit.*, Jun. 2003, vol. 1, pp. 835–841.
- [36] G. Hamarneh and C. McIntosh, "Physics-based deformable organisms for medical image analysis," *Proc. SPIE Med. Imag.: Image Process.*, vol. 5747, pp. 326–335, 2005.
- [37] S. Pizer, G. Gerig, S. Joshi, and S. Aylward, "Multiscale medial shape-based analysis of image objects," *Proc. IEEE*, vol. 91, no. 10, pp. 1670–1679, Oct. 2003.
- [38] K. Siddiqi, S. Bouix, A. Tannenbaum, and S. Zucker, "Hamilton-Jacobi skeletons," *Int. J. Comput. Vis.*, vol. 48, no. 3, pp. 215–231, 2002.
- [39] T. Sebastian, P. Klein, and B. Kimia, "Recognition of shapes by editing their shock graphs," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 26, no. 5, pp. 550–571, May 2004.
- [40] V. Caselles, R. Kimmel, and G. Sapiro, "Geodesic active contours," *Int. J. Comput. Vis.*, vol. 22, pp. 61–79, 1997.
- [41] V. Kolmogorov and R. Zabini, "What energy functions can be minimized via graph cuts?," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 26, no. 2, pp. 147–159, Feb. 2004.
- [42] A. Sinop and L. Grady, "A seeded image segmentation framework unifying graph cuts and random walker which yields a new algorithm," in *IEEE 11th Int. Conf. Comput. Vis.*, Oct. 2007, pp. 1–8.
- [43] U. Maulik, "Medical image segmentation using genetic algorithms," *IEEE Trans. Inf. Technol. Biomed.*, vol. 13, no. 2, pp. 166–173, Mar. 2009.

- [44] L. Ballerini, "Genetic snakes for medical image segmentation," in *Proc. SPIE Conf. Math. Model. Estim. Techn. Comput. Vis.*, 1998, vol. 3457, pp. 284–295.
- [45] L. Ballerini, "Genetic snakes for color images segmentation," in *Applications of Evolutionary Computing*, ser. Lecture Notes in Computer Science, E. Boers, Ed. Berlin, Germany: Springer, 2001, vol. 2037, pp. 268–277.
- [46] L. MacEachern and T. Manku, "Genetic algorithms for active contour optimization," in *Proc. 1998 IEEE Int. Symp. Circuits Syst.*, May 1998, vol. 4, pp. 229–232.
- [47] J. Tohka, "Global optimization of deformable surface meshes based on genetic algorithms," in *Proc. 11th Int. Conf. Image Anal. Process.*, Sep. 2001, pp. 459–464.
- [48] Y. Fan, T. Jiang, and D. Evans, "Volumetric segmentation of brain images using parallel genetic algorithms," *IEEE Trans. Med. Imag.*, vol. 21, no. 8, pp. 904–909, Aug. 2002.
- [49] O. I. Nez, N. Barreira, J. Santos, and M. Penedo, "Genetic approaches for topological active nets optimization," *Pattern Recognit.*, vol. 42, no. 5, pp. 907–917, 2009.
- [50] D. Terzopoulos, "On matching deformable models to images," in *Topical Meeting Mach. Vis., Tech. Digest Series*, 1987, vol. 12, pp. 160–167.
- [51] A. Hill and C. Taylor, "Model-based image interpretation using genetic algorithms," *Image Vis. Comput.*, vol. 10, no. 5, pp. 295–300, 1992.
- [52] P. Ghosh and M. Mitchell, "Segmentation of medical images using a genetic algorithm," in *Proc. 8th Annu. Conf. Genetic Evolutionary Computat.*, New York, 2006, pp. 1171–1178.
- [53] C. F. Ruff, S. W. Hughes, and D. J. Hawkes, "Volume estimation from sparse planar images using deformable models," *Image Vis. Comput.*, vol. 17, no. 8, pp. 559–565, 1999.
- [54] M. Mignotte and J. Meunier, "Deformable template and distribution mixture-based data modeling for the endocardial contour tracking in an echographic sequence," in *IEEE Comput. Soc. Conf. Comput. Vis. Pattern Recognit.*, 1999, vol. 1, pp. 230–230.
- [55] C. Lu, S. Pizer, S. Joshi, and J.-Y. Jeong, "Statistical multi-object shape models," *Int. J. Comput. Vis.*, vol. 75, pp. 387–404, 2007.
- [56] X. Liu, J.-Y. Jeong, J. Levy, R. Saboo, E. Chaney, and S. Pizer, "A large-to-fine-scale shape prior for probabilistic segmentations using a deformable m-rep," in *IEEE Comput. Soc. Conf. Comput. Vis. Pattern Recognit. Workshops*, Jun. 2008, pp. 1–8.
- [57] S. Pizer, D. Fritsch, P. Yushkevich, V. Johnson, and E. Chaney, "Segmentation, registration, and measurement of shape variation via image object shape," *IEEE Trans. Med. Imag.*, vol. 18, no. 10, pp. 851–865, Oct. 1999.
- [58] S. Kaleem and P. Stephen, *Medial Representations: Mathematics, Algorithms and Applications*. New York: Springer, 2008.
- [59] S. Z. Li, *Markov Random Field Modeling in Image Analysis*, ser. Advances in pattern recognition, 3rd ed. New York: Springer-Verlag, 2009.
- [60] J. D. Foley, A. van Dam, S. K. Feiner, and J. Hughes, *Computer Graphics: Principles and Practice in C*, 2nd ed. Reading, MA: Addison-Wesley, 1995.
- [61] M. Mitchell, *An Introduction to Genetic Algorithms*. Cambridge, MA: MIT Press, 1996.
- [62] E. Marszal *et al.*, "Agenesis of corpus callosum: Clinical description and etiology," *J. Child Neurol.*, vol. 15, pp. 401–405, 2000.
- [63] J. Pantel *et al.*, "Topography of callosal atrophy reflects distribution of regional cerebral volume reduction in alzheimer's disease," *Psychiatry Res.*, vol. 90, pp. 180–192, 1999.
- [64] R. Rosenthal and L. Bigelow, "Quantitative brain measurements in chronic schizophrenia," *Br. J. Psychiatry*, vol. 121, pp. 259–264, 1972.
- [65] M. Shenton *et al.*, "Abnormalities in the left temporal lobe and thought disorder in schizophrenia: A quantitative magnetic resonance imaging study," *New Eng. J. Med.*, vol. 327, pp. 604–612, 1992.
- [66] P. Jaccard, "Étude comparative de la distribution florale dans une portion des alpes et des jura," *Bulletin del la Société Vaudoise des Sciences Naturelles*, vol. 37, pp. 547–579, 1901.
- [67] C. McIntosh and G. Hamarneh, "Optimal weights for convex functionals in medical image segmentation," in *Int. Symp. Vis. Comput.: Special Track on Optimization for Vision, Graphics and Medical Imaging: Theory and Applications (ISVC OVGMI)*, 2009, vol. 5875-I, pp. 1079–1088.
- [68] M. Szummer, P. Kohli, and D. Hoiem, D. Forsyth, P. Torr, and A. Zisserman, Eds., "Learning CRFs using graph cuts," in *Computer Vision—ECCV 2008*, Berlin, Germany, 2008, vol. 5303, pp. 582–595.
- [69] T. Chan and L. Vese, "Active contours without edges," *IEEE Trans. Image Process.*, vol. 10, no. 2, pp. 266–277, Feb. 2001.
- [70] P. Golland, W. Eric, and L. Grimson, "Fixed topology skeletons," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2000, vol. 1, pp. 10–17.
- [71] G. Hamarneh, A. Ward, and R. Frank, "Quantification and visualization of localized and intuitive shape variability using a novel medial-based shape representation," in *IEEE Int. Symp. Biomed. Imag.*, 2007, pp. 1232–1235.