12

DEFORMABLE ORGANISMS FOR MEDICAL IMAGE ANALYSIS

Ghassan Hamarneh and Chris McIntosh

School of Computing Science, Simon Fraser University
Burnaby, British Columbia, Canada

In medical image analysis strategies based on deformable models, controlling the deformations of models is a desirable goal in order to produce proper segmentations. In Chapter 11 — "Physically and Statistically Based Deformable Models for Medical Image Analysis" — a number of extension were demonstrated to achieve this, including user interaction, global-to-local deformations, shape statistics, setting low-level parameters, and incorporating new forces or energy terms. However, incorporating expert knowledge to automatically guide deformations can not be easily and elegantly achieved using the classical deformable model low-level energy-based fitting mechanisms. In this chapter we review Deformable Organisms, a decision-making framework for medical image analysis that complements bottom–up, data-driven deformable models with top–down, knowledge-driven mode-fitting strategies in a layered fashion inspired by artificial life modeling concepts. Intuitive and controlled geometrically and physically based deformations are carried out through behaviors. Sensory input from image data and contextual knowledge about the analysis problem govern these different behaviors. Different deformable organisms for segmentation and labeling of various anatomical structures from medical images are also presented in this chapter.

1. INTRODUCTION AND MOTIVATION

Due to the important role of medical imaging in the understanding, diagnosis, and treatment of disease, potentially overwhelming amounts of medical image data

Address all correspondence to: Dr. Ghassan Hamarneh, School of Computing Science, Simon Fraser University, 8888 University Drive, Burnaby, BC, V5A 1S6, Canada. Phone: +1.604.291.3007, Fax: +1.604.291.3045, hamarneh@cs.sfu.ca.

are continuously being acquired. This is creating an increasing demand for medical image analysis (MIA) tools that are not only robust and highly automated, but also intuitive for the user and flexible to adapt to different applications. Medical image segmentation in particular remains one of the key tasks indispensable to a wide array of subsequent quantification and visualization goals in medicine, including computer-aided diagnosis and statistical shape analysis applications. However, the automatic segmentation and labeling of anatomical structures in medical images is a persistent problem that continues to defy solution. Several classifications of segmentation techniques exist, including edge, pixel, and region-based techniques, clustering, graph theoretic, and model correlation approaches [1–5]. However, no one method can yet handle the most general case with sufficient accuracy.

It is important to note that a substantial amount of knowledge is often available about anatomical structures of interest — shape, position, orientation, symmetry, relationship to neighboring structures, landmarks, etc. — as well as about the associated image intensity characteristics. However, medical image analysis researchers have struggled to develop segmentation techniques that can take full advantage of such knowledge.

The development of general-purpose automatic segmentation algorithms will require not only powerful bottom–up, data-driven processes, but also equally powerful top–down, knowledge-driven processes within a robust decision-making framework that operates across multiple levels of abstraction. A flexible framework is needed that can operate at the appropriate level of abstraction and is capable of incorporating and applying all available knowledge effectively (i.e., at the correct time and location during image analysis). A critical element in any such viable highly automated solution is the decision-making framework itself. Top–down, hierarchically organized models that shift their focus from structures associated with stable image features to those associated with less stable features are promising examples [6, 7]. Although Knowledge-based [8–14] and agent-based segmentation techniques [15–19] have been proposed in the past, their use of high-level contextual knowledge remains largely ineffective because it is intertwined much too closely with the low-level optimization-based mechanisms. We revisit ideas for incorporating knowledge that were explored in earlier systems and develop a new frameworks that focuses on top–down reasoning strategies that may best leverage the powerful bottom–up feature detection and integration abilities of deformable models and other modern model-based medical image analysis techniques.

Deformable models, one of the most actively researched model-based segmentation techniques [20], feature a potent bottom–up component founded in estimation theory, optimization, and physics-based dynamical systems, but their top–down processes have traditionally relied on interactive initialization and guidance by knowledgeable users (see Section 1.4, Snakes Drawbacks, in Chapter 11). Since their introduction by Terzopoulos et al. [83, 84], deformable models for medical image segmentation have gained increasing popularity. In addition to physics-based explicit deformable models [20, 21], geometry-based implicit implementations have also attracted attention [22–24].

Although most deformable shape models are capable of deforming into a variety of shapes, they lack the ability to undergo intuitive deformations such as bending, stretching, and other global motions such as sliding and backing up. The reason is that at the geometric level deformable models are typically boundary-based and are not designed with intuitive, multiscale, multi-location deformation controllers or deformation handles. Their inability to perform controlled and intuitive deformations makes it difficult to develop reasoning and planning knowledge-driven strategies for model-to-data fitting at the appropriate level of abstraction.

In more sophisticated deformable models, prior information in the form of measured statistical variation is used to constrain model shape and appearance [25–27]. However, these models have no explicit awareness of where they or their neighbors are, and the effectiveness of these constraints is consequently dependent upon model initialization conditions. The lack of awareness also prevents the models from taking proper advantage of neighborhood information via model interaction and prevents them from knowing when to trust the image feature information and ignore the constraint information and vice versa. The constraint information is therefore applied arbitrarily. Furthermore, because there is no active, explicit search for stable image features, the models are prone to latching onto incorrect features [26], simply due to their myopic decision-making abilities and the proximity of spurious features. Once this latching occurs, the lack of explicit control of the fitting procedure prevents the model from correcting such missteps. The result is that the local decisions that are made do not add up to intelligent global behavior.

For example, when segmenting the corpus callosum (CC) in 2D midsagittal images[1], the vocabulary that one uses should preferably contain words that describe principal anatomical features of the CC, such as the genu, splenium, rostrum, fornix, and body (Figure 1), rather than pixels and edges. The deformable model should match this natural descriptiveness by grouping intuitive model parameters at different scales and locations within it, rather than providing localized boundary-based parameters only.

Attempts to fully automate deformable model segmentation methods have so far been less than successful at coping with the enormous variation in anatomical structures of interest, the significant variability of image data, the need for intelligent initialization conditions, etc. It is difficult to obtain intelligent, global (i.e., over the whole image) model behavior throughout the segmentation process from fundamentally local decisions. In essence, current deformable models have no explicit awareness of where they are in the image, how their parts are arranged, or what they or any neighboring deformable models are seeking at any time during the optimization process.

An explicit search for anatomical landmarks requires powerful, flexible, and intuitive model deformation control coupled with appropriate feature detection

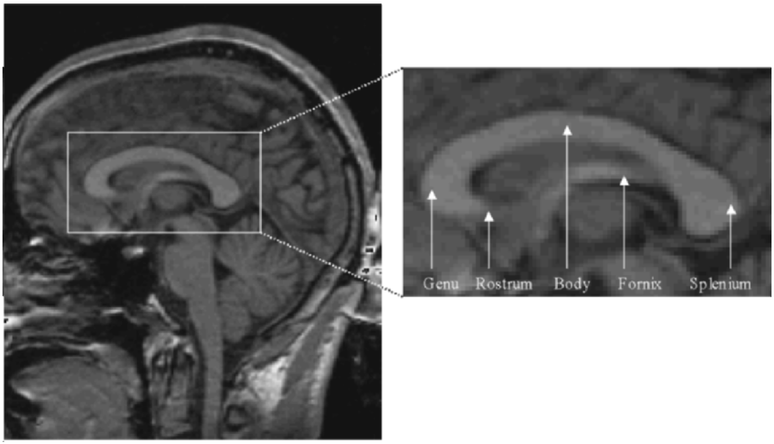

Figure 1. Corpus callosum anatomy overlain on a midsagittal MRI brain slice.

capabilities. Consequently, controlling the deformations of an object's shape in a way that is based on the natural geometry of the object is highly desirable in medical image segmentation and interpretation. This intuitive deformation ability reflects the flexibility of clay to be shaped in a sculptor's hands and naturally lends itself to guidance by high-level controllers. Furthermore, the performance of the controllers can be greatly enhanced by keeping the deformations consistent with prior knowledge about the possible object shape variations. Several techniques exist that provide some control of model deformations by incorporating new energy terms [29, 30] or by allowing only feasible deformations to be produced through the incorporation of prior shape knowledge [25, 31–34]. However, the deformation control these earlier methods provide is tightly linked to low-level cost or energy terms, in which high-level anatomical knowledge is difficult to encode.

We facilitate the merger of higher-level control mechanisms with powerful controlled deformations through a layered architecture, where the high-level reasoning layer has knowledge about and control over the low-level model (or models) at all times. The reasoning layer should apply an active, explicit search strategy that first looks for the most stable image features before proceeding to less stable image features, and so on. It should utilize contextual knowledge to resolve ambiguity in regions where there is a deficiency of image feature information. By constructing the framework in a layered fashion, we are able to separate and complement the data-driven, local image feature integration functionality with the knowledge-driven, model-fitting control functionality, exploiting both for maximal effectiveness. This separation allows us to construct a model-fitting controller from an extensible set of standardized subroutines. The subroutines are defined in terms of deformation maneuvers and high-level anatomical landmark detectors rather than low-level image features. Different subroutines are invoked based on high-level model-fitting decisions by integrating image features, prior contextual anatomical knowledge, and a pre-stored segmentation plan. Furthermore, by combining a layered architecture with a set of standard subroutines, powerful and flexible "custom-tailored" models can be rapidly constructed, thus providing general-purpose tools for automated medical image segmentation and labeling.

2. DEFORMABLE ORGANISMS: AN ARTIFICIAL LIFE MODELING PARADIGM FOR MEDICAL IMAGE ANALYSIS

To realize the ideas and achieve the goals mentioned above in Section 1, we introduced a new paradigm for automatic medical image analysis that adopts concepts from the emerging field of artificial life (AL)[2]. In particular, we developed *deformable organisms*, autonomous agents whose objective is the segmentation and analysis of anatomical structures in medical images [38]. The AL modeling-based approach provides us with the required flexibility to adhere to an active, explicit search strategy that takes advantage of contextual and prior knowledge of anatomy. The organisms are aware of the progress of the segmentation process and of each other, allowing them to effectively and selectively apply knowledge of the target objects throughout their development.

Viewed in the context of the AL modeling hierarchy (Figure 2a), current *automatic* deformable model-based approaches to medical image analysis utilize geometric and physical modeling layers only (Figure 2b). In *interactive* deformable models, such as snakes, the human operator is relied upon to provide suitable behavioral and cognitive level support (Figure 2c). At the physical level, deformable models interpret image data by simulating dynamics or minimizing energy terms, but the models themselves do not monitor or control this optimization process except in a most primitive way.

To overcome the aforementioned deficiencies while retaining the core strengths of the deformable model approach, we add high-level controller layers (a "brain") on top of the geometric and physical layers to produce an autonomous deformable organism (Figure 2d). The planned activation of these lower deformation layers allows us to control the fitting or optimization procedure. The layered architecture approach allows the deformable organism to make deformation decisions at the correct level of abstraction utilizing prior knowledge, memorized information, sensed image features, and inter-organism interaction.

Specifically, a deformable organism is structured as a "muscle"-actuated "body" whose "behaviors" are controlled by a "brain" (Figure 3) that makes decisions based on perceived image data and extraneous knowledge. The brain is the organism's cognitive layer, which activates behavior routines (e.g., for a CC organism: find-splenium, find-genu, find-upper-boundary, etc. (Figure 1)) according to a plan or schedule (Figure 3).

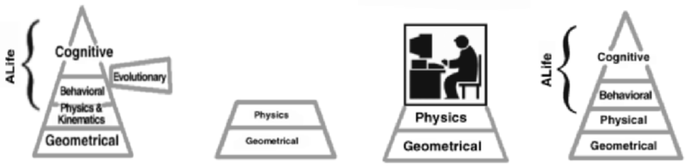

(a) (b) (c) (d)

Figure 2. AL, Deformable Models, and Deformable Organisms. (a) AL modeling pyramid. Adapted with permission from [39]. Copyright ©1999, ACM. (b) Automatic deformable models (incorporating geometry and physics layers only). (c) Deformable models guided by an expert human operator. (d) Intelligent deformable models (deformable organisms) provide a model of the cognitive abilities of human operators (by including higher cognitive layers).

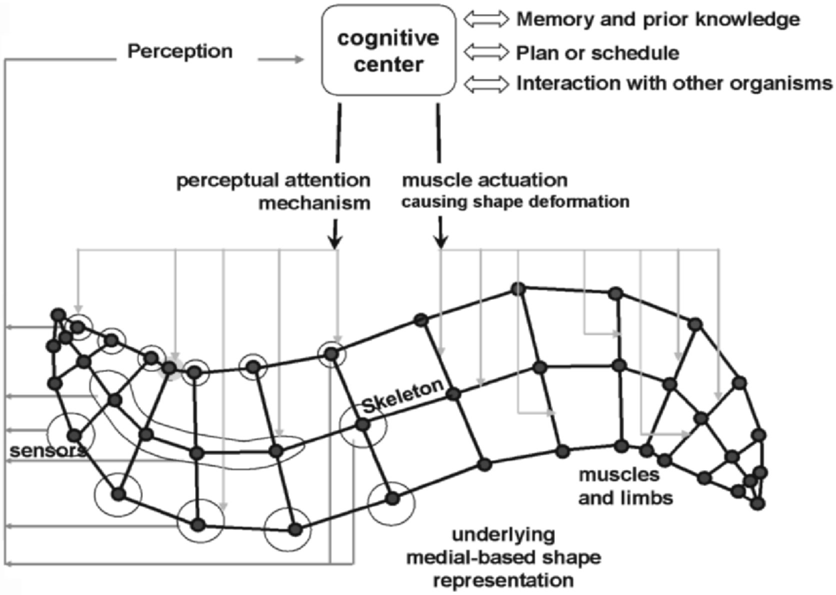

Figure 3. A Deformable Organism. The brain issues "muscle" actuation and perceptual attention commands. The organism deforms and senses image features, whose characteristics are conveyed to the brain. The brain makes decisions based on sensory input, memorized information and prior knowledge, and a pre-stored plan, which may involve interaction with other organisms. Reprinted with permission from [85]. Copyright ©2002, Elsevier.

Carrying out a behavior routine requires image information in order for the proper shape deformation to take place toward achieving the goal of the current behavior. The deformable organism perception system is responsible for gathering image information and comprises a set of sensors that are adaptively tuned to specific image features (edge strength, texture, color, etc) in a task-specific way. Hence, the organism can disregard sensory information superfluous to its current behavioral needs.

The organism carries out a sequence of active, explicit searches for stable anatomical features, beginning with the most stable anatomical feature and then proceeding to the next best feature. This allows the organism to be "self-aware" (i.e., knows where it and its parts are and what it is seeking at every stage) and is therefore able to perform these searches intelligently and effectively by utilizing a conflux of contextual knowledge, perceived sensory data, an internal mental state, memorized knowledge, and a cognitive plan. For example, it need not be satisfied with the nearest matching feature, but can look further within a region to find the best match, thereby avoiding globally suboptimal solutions. The plan (or plans) can be generated with the aid of a human expert, since the behavior routines are defined using familiar anatomical terminology.

An organism may "interact" with other organisms to determine optimal initial conditions. Once stable features are found and labeled, an organism can selectively use prior knowledge or information from the neighbor organisms to determine the object boundary in regions known to offer little or no feature information. Interaction among organisms may be as simple as collision detection and avoidance, or one or several organisms supplying intelligent initial conditions to another, or the use of inter-organism statistical shape/image appearance constraint information.

Furthermore, by carrying out explicit searches for features, correct correspondences between the organism and the data are more readily assured. If a feature cannot be found, an organism may "flag" this situation (Figure 4b). If multiple plans exist, another plan can be selected and/or the search for the missing feature postponed until further information is available (e.g., from a neighboring organism). Alternatively, the organism can retrace its steps and return to a known state and then inform the user of the failure. A human expert can intervene and put the organism back on course by manually identifying the feature. This strategy is possible because of the sequential and spatially localized nature of the model fitting process.

Customized behavioral routines and explicit feature search requires powerful, flexible and intuitive model deformation control. The behavior routines activate "motor" (i.e., deformation) controller routines or growth controller routines, enabling the organism to fulfill its goal of object segmentation. An organism may begin in an "embryonic" state with a simple proto-shape, and then undergo controlled growth as it develops into an "adult," proceeding from one stable object feature to the next. Alternatively, an organism may begin in a fully developed

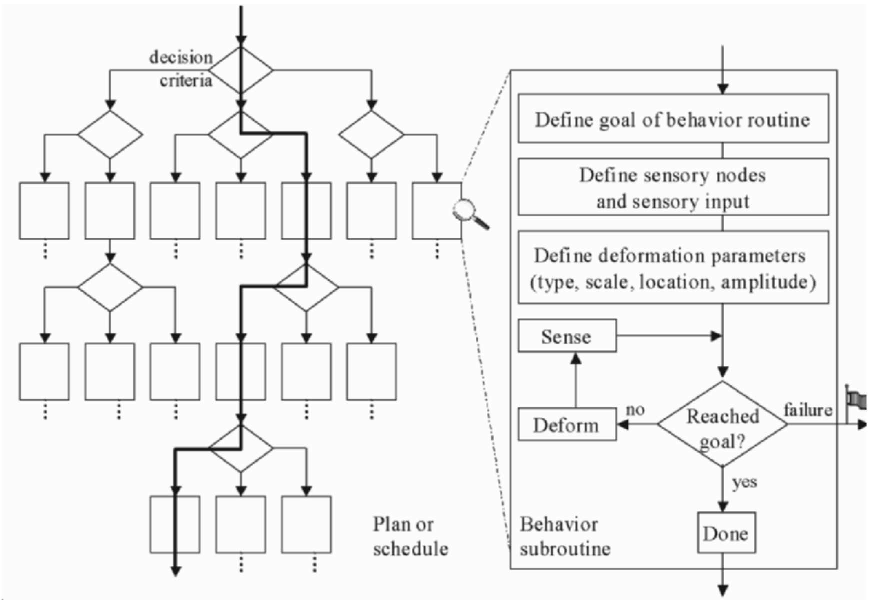

(a) (b)

Figure 4. (a) Procedural representation of a fragment of a deformable organism's plan or schedule. The organism goes through several behavior subroutines (bold path in (a)). (b) Generic example of a behavior routine. Reprinted with permission from [85]. Copyright ©2002, Elsevier.

state and undergo controlled deformations as it carries out its model-fitting plan. The type of organism to use, or whether to use, some sort of hybrid organism, is dependent on the image and shape characteristics of the target anatomical structure (different examples are presented in Section 4).

Deformation controllers are parametrized procedures dedicated to carrying out a complex deformation function, such as successively bending a portion of the organism over some range of angle or stretching part of the organism forward some distance. They translate natural control parameters such as <bend_angle, location, scale> or <stretch_length, location, scale> into detailed deformations.

To summarize, the AL layered architecture provides the needed framework to complement the geometrical and physical layers of classical deformable organisms (the model shape, topology, and deformations) with behavioral and cognitive layers (the high-level controllers that activate the routines) utilizing a perception system (the source of the image data) and contextual-knowledge (knowledge about the anatomy's shape, appearance, neighborhood relationships, etc.).

3. THE LAYERED ARCHITECTURE OF DEFORMABLE ORGANISMS

This section describes the layers of the deformable organism architecture, starting with the lowest layers (geometrical and physical), and working up to the higher layers (behavioral, cognitive, and sensory). Throughout these sections we will provide examples of each layer to solidify the abstract terminology. To this end, we will draw upon particular applications of deformable organisms, each of which is briefly outlined below.

- **Geometrically based deformable CC organism**: This CC organism relies on pure geometry-based shape deformations and was designed to begin as a small "worm," and then grow and expand itself to fit the data (Figure 5a). Details of its deformation module are presented in Section 3.3.1, sensors in Section 3.4, and with behaviors and results in Section 4.1.

- **Physically based deformable CC organism:** Incorporates a true physically based deformation layer (Section 3.3.3) to allow for increased robustness and intuitive interaction (Figure 5b). Example physics, behavioral, and cognitive layers of this organism are presented in Sections 3.3.3, 3.5.1, and 3.6.1, respectively, with results in Section 4.4.

- **2D vessel crawler:** The 2D vessel crawler is a worm-like deformable organism (Figure 5c) that grows to segment vessels in 2D angiograms by continuously sensing which direction to grow in, then expanding and fitting to the walls. Its distinguishing feature is its ability to discriminate between bifurcations and overlaps (Figure 41) using an off-board sensor (Section 3.4).

- **3D vessel crawler:** This organism extends the 2D vessel crawler to 3D geometry, incorporates a true physical deformation layer (Section 3.3.4), and is equipped with new behavioral and sensory modules (Section 3.4) [86]. It crawls along vasculature in 3D images (Figure 5d), accurately segmenting vessel boundaries, detecting and exploring bifurcations, and providing sophisticated, clinically relevant structural analysis. Results are presented in Section 4.6.

3.1. The Geometrical Layer (Shape and Topology)

The geometrical layer of the deformable organism houses the shape representation and morphological and topological constraints of the organism. The deformable organisms we developed to date are based on medial-axis or medial-sheet deformable shape representations. The medial-based shape representation allows for a variety of intuitive and controlled changes in the model's shape that are described relative to the natural geometry of the object rather than as displacements

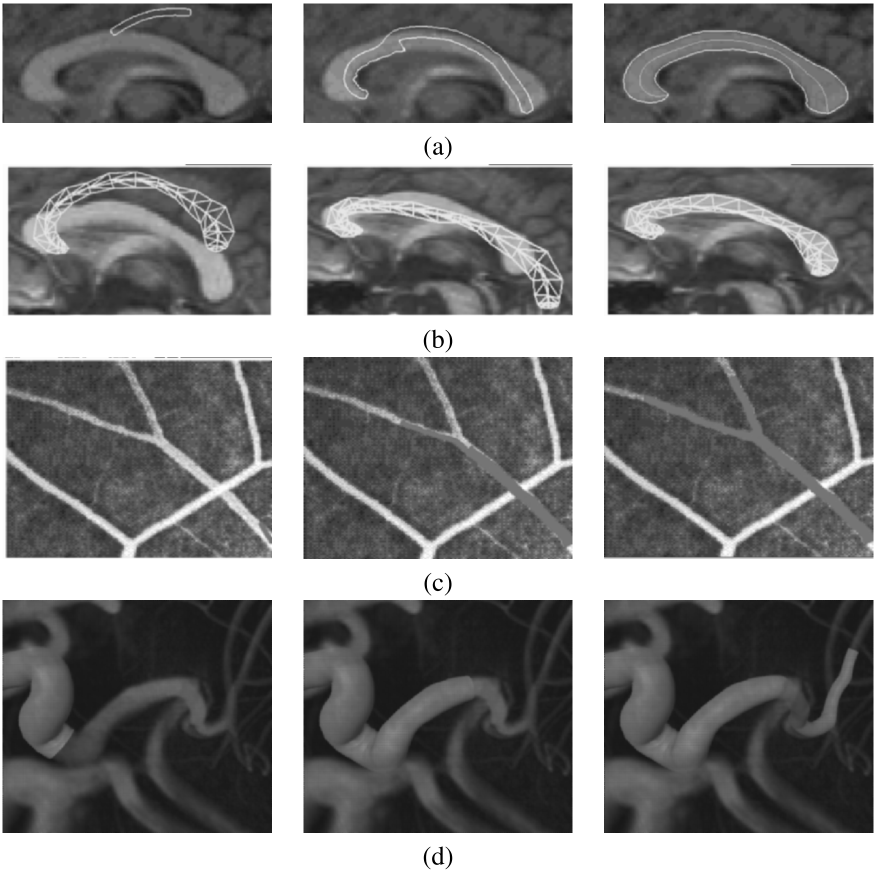

Figure 5. Example deformable organisms for medical image analysis. (a) Geometrically based and (b) physically based deformable CC organisms. (c) 2D and (d) 3D vessel crawler. Progress of segmentation is shown from left to right. See attached CD for color version.

to its boundary. These changes in shape include stretching, thickening, or bending and are realized as either pure geometric deformations or a result of deformations simulated at the physics layer of the organism. In Section 3.3 we present examples of different shape representations that allow such controlled deformations.

3.2. The Physical Layer (Motor System and Deformations)

The physical layer is responsible for deforming the shape of the organism and changing its geometry. Deformations carried out by the physical layer of the organism are divided into two classes: simple deformation capabilities and more complex locomotion maneuvers. The simple deformation capabilities constitute the low-level motor skills or the basic shape deformation actuators of the organism. Simple deformations include global deformations such as affine transformation (translate, rotate, scale) or localized deformations such as a simple bulge or a stretch. The complex locomotion maneuvers are the high-level motor skills of the organism constructed from several basic motor skills. These are parametrized procedures that carry out complex deformation functions such as sweeping over a range of rigid transformation parameters, sweeping over a range of stretch/bend/thickness parameters, bending at increasing scales, moving a bulge on the boundary, etc. Other high-level deformation capabilities include smoothing the medial or the boundary of the model, or moving the medial axis or medial sheet to a position exactly midway between the boundaries on both sides of the medial. In the following section (3.3) we present different examples of deformation enabled by a number of medial-based geometrical representations.

3.3. Controlling Shape Deformation

Top–down, knowledge-driven model fitting strategies require underlying shape representations responsive to high-level controlled shape deformation commands. As the organism's cognitive layer decides on a behavior, the behavior routine will involve carrying out shape deformations realized through high-level motor skills, which in turn are realized through simpler low-level motor skills. The goal here is the ability to intelligently control the different types and extent of model deformations during the model-to-data fitting process in an effort to focus on the extraction of stable image features before proceeding to object regions with less well-defined features.

The choice of shape representation is undeniably crucial for segmentation, recognition, and interpretation of medical images. The study of shape is unsurprisingly attracting a great deal of attention within the medical image analysis community [40–42]. A desirable trait in a shape representation is its ability to control non-rigid object deformations at multiple locations and scales in an interactive and intuitive manner. Deformable shape models [20], in particular, are mostly boundary based, and therefore multiscale deformation control is constructed upon arbitrary boundary point sets and not upon object-relative geometry. Although they provide excellent local shape control, they lack the ability to undergo intuitive global deformation and decompose shape variability into intuitive deformations. As a result, it is difficult to incorporate intelligent deformation control operating at the right level of abstraction into the typical deformable model framework of energy minimization. Consequently, these models remain sensitive to initial conditions and spurious image features in image interpretation tasks.

Hierarchical boundary-based shape models [43–46] and volume-based shape representations or free-form deformation mechanisms have been proposed [47–51]; however, as their deformation is not based on object-relative geometry, they are limited either by the type of objects they can model, or by the type and intuitiveness of the deformations they can carry out. They are also typically not defined in terms of the object but rather the object is unnaturally defined (or deformed) in terms of the representation or deformation mechanism. Other 3D shape representations with similar drawbacks include spherical harmonics, FEM, NURBS, and wavelet-based representations [52–55].

Shape models founded upon the use of the medial-axis transform [56] are emerging as a powerful alternative to the earlier boundary-based and volume-based techniques [57–68]. Medial representations provide both a local and global description of shape. Deformations defined in terms of a medial axis are natural and intuitive and can be limited to a particular scale and location along the axis, while inherently handling smoothness and continuity constraints.

Statistical models of shape variability have been used for medical image interpretation [25, 27, 31, 32, 69]. These typically rely on principal component analysis (PCA) and hence are only capable of capturing global shape variation modes. Statistical analysis of medial-based shape representation has been the focus of much recent research [70–74].

In the following subsections we present the details of a variety of shape representation and controlled deformation techniques that are crucial to the operation of deformable organisms and modeling their lower geometrical and physical layers. These include 2D medial profiles (Section 3.3.1), 3D shape medial patches (3.3.2), 2D spring-mass systems with physics-based deformations (3.3.3), and their 3D extension (3.3.4).

3.3.1. 2D Shape Representation and Deformation with Medial Profiles

We describe a multiscale, medial-based approach to shape representation and controlled deformation in an effort to meet the requirements outlined above in the previous sections. In this shape representation and deformation scheme, an anatomical structure (e.g., the CC) is described with four shape profiles derived from the primary medial axis of an organism's boundary contour (e.g., CC boundary contour). The medial profiles describe the geometry of the structure in a natural way and provide general, intuitive, and independent shape measures. These profiles are: length, orientation, left (with respect to the medial axis) thickness, and right thickness. Once the profiles are constructed, various deformation functions or *operators* can be applied at certain locations and scales on a profile, producing intuitive, controlled deformations: stretching, bending, and bulging. In addition to the general deformation operators, we wish to use as much knowledge as possible about the object itself and to generate statistically proven feasible deformations from a training set. We would like to control these statistical deformations locally along the medial shape profiles to support our goal of intelligent deformation scheduling. Since general statistically derived shape models only produce global shape variation modes [25, 27], we present spatially localized feasible deformations at desired scales by utilizing hierarchical (multiscale) and regional (multi-location) principal component analysis to capture shape variation statistics.

In the following sections, we demonstrate the ability to produce controlled shape deformations by applying them to 2D medial-based representations of the CC, derived from 2D midsagittal MRI slices of the brain. We begin by describing the generation and use of medial-based profiles for shape representation and describe a set of general operators that act on the medial shape profiles to produce controlled shape deformations. 3.3.1). We then present a technique for performing a multiscale multi-location statistical analysis of the shape profiles and describe statistics-based deformations based on this analysis. We present a simple application of the controlled shape deformations and demonstrate their use in an automatic medical image analysis system (Section 4.1).

Medial Profiles for Shape Representation We use a boundary representation of an object to generate the medial-based profiles. Generation of the profiles begins with extraction of a sampled pruned skeleton of the object to obtain a set of medial nodes. Four medial profiles are constructed: a length profile $L(m)$, an orientation profile $O(m)$, a left (with respect to the medial axis) thickness profile $T^l(m)$, and a right thickness profile $T^r(m)$, where $m = 1, 2, \ldots, N$, N being the number of medial nodes, with nodes 1 and N being the terminal nodes. The length profile represents the distances between consecutive pairs of medial nodes, and the orientation profile represents the angles of the between edges connecting consecutive pairs of medial nodes. The thickness profiles represent the distances between medial nodes and their corresponding boundary points on both sides of the medial axis (Figure 6). Corresponding boundary points are calculated by computing the intersection of a line passing through each medial node in a direction normal to the medial axis, with the boundary representation of the object. Example medial profiles are shown in Figure 7.

Shape Reconstruction from Medial Profiles To reconstruct the object's shape given its set of medial profiles, we calculate the positions of the medial and boundary nodes by following these steps:

1. Specify affine transformation parameters: orientation angle θ, translation values (t_x, t_y), and scale (s_x, s_y).

2. Using medial node 1 as the base or reference node, place it at location $x_1 = (t_x, t_y)$.

3. Repeat steps 4 and 5 for $m = 1, 2, \ldots, N$.

4. Compute locations x_m^l and x_m^r of boundary points l and r at either side of the mth medial node (Figure 6) as

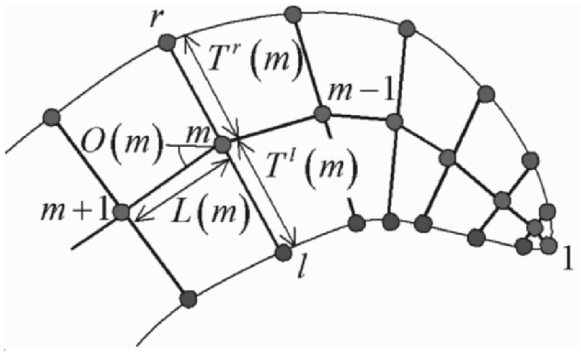

Figure 6. Diagram of shape representation. Reprinted with permission from [85]. Copyright ©2002, Elsevier.

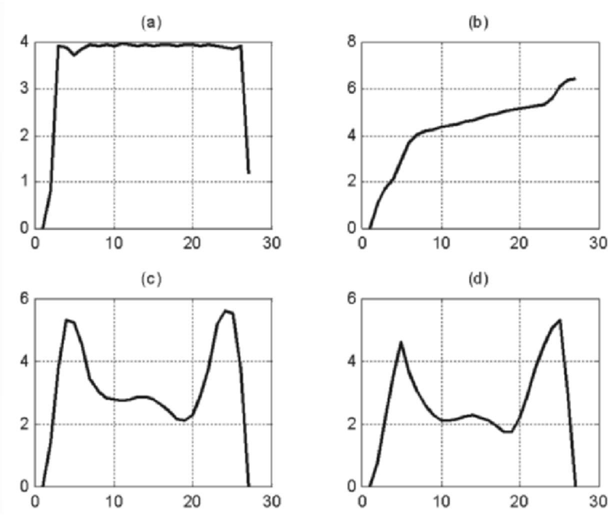

Figure 7. Example medial shape profiles: (a) length profile $L(m)$, (b) orientation profile $O(m)$, (c) left thickness profile $T^l(m)$, and (d) right thickness profile $T^r(m)$. Reprinted with permission from [85]. Copyright ©2002, Elsevier.

$$x_m^l = x_m + T^l(m)\begin{pmatrix} s_x \cos\left(\theta + O(m) + \frac{\pi}{2}\right) \\ s_y \sin\left(\theta + O(m) + \frac{\pi}{2}\right) \end{pmatrix} \tag{1}$$

and, similarly,

$$x_m^r = x_m + T^r(m)\begin{pmatrix} s_x \cos\left(\theta + O(m) - \frac{\pi}{2}\right) \\ s_y \sin\left(\theta + O(m) - \frac{\pi}{2}\right) \end{pmatrix}. \tag{2}$$

5. If $m < N$, compute location x_{m+1} of the next medial node $m+1$ as

$$x_{m+1} = x_m + L(m)\begin{pmatrix} s_x \cos(\theta + O(m)) \\ s_y \sin(\theta + O(m)) \end{pmatrix}. \tag{3}$$

An example shape reconstruction is shown in Figure 8.

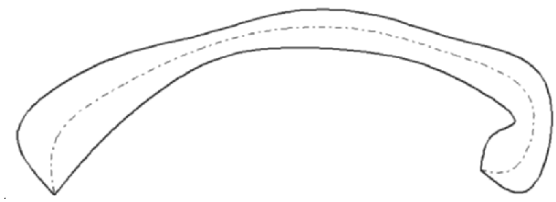

Figure 8. Object reconstruction resulting from the shape profiles in Figure 7.

Shape Deformations Using Medial-Based Operators Once the shape profiles have been generated, we can construct deformation operators and apply these operators to the shape profiles. This results in intuitive deformations of the object upon reconstruction. That is, by applying an operator to the length, orientation, or thickness shape profile, we obtain a stretch, bend, or bulge deformation, respectively. Each deformation operator is implemented by defining a medial-based operator profile, $k(m)$, of a particular type (Figure 9) and specifying an amplitude, location, and scale.

The operator profile is then added to (or blended with) the medial shape profile corresponding to the desired deformation. For example, to introduce a bulge on the right boundary, an operator profile with a specific amplitude, type, location, and scale is generated and added to the right thickness medial profile $T^r(m)$ to obtain $T^r(m) + k(m)$ (Figure 10).

In general the application of a deformation operator $k(m)$ alters the desired shape profile according to

$$p_d(m) = \bar{p}_d(m) + \alpha_{dlst} k_{dlst}(m), \tag{4}$$

where p is the shape profile, d is the deformation type (stretch, bend, left/right bulge), i.e., $p_d(m) : \{L(m), O(m), T^l(m), T^r(m)\}$, $\bar{p}$ is the average shape

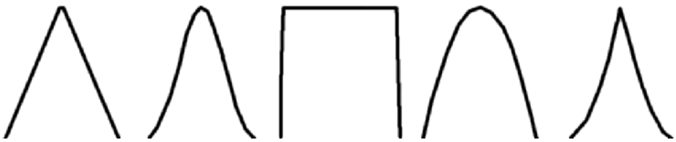

Figure 9. Examples of operator types: (left to right) triangular, Gaussian, flat, bell, and cusp [75]. Reprinted with permission from [63]. Copyright ©2004, World Scientific.

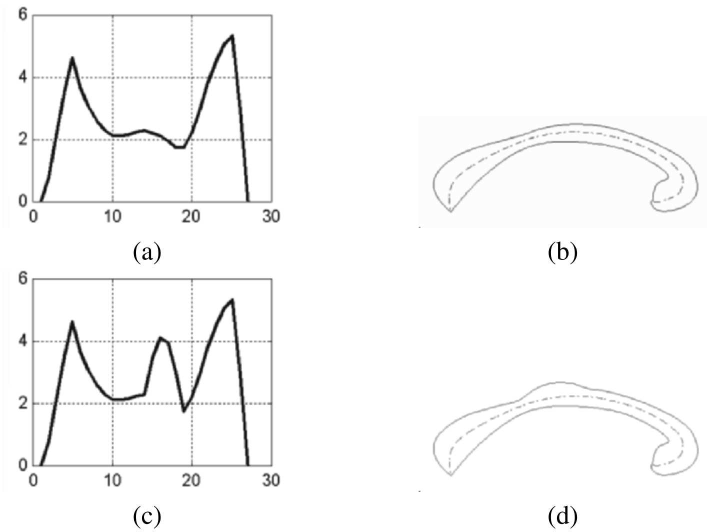

Figure 10. Introducing a bulge on the right boundary by applying a deformation operator on the right thickness profile: (a) $T^r\,(m)$ before and (c) after applying the operator. (b) Reconstructed shape before and (d) after the operator. Reprinted with permission from [85]. Copyright ©2002, Elsevier.

profile, k is the operator profile (with unity amplitude), l is the location, s is the scale, t is the operator type (Gaussian, triangular, ..., etc.), and α is operator amplitude.

Altering one shape profile only affects the shape property associated with that profile and does not affect any other object shape properties. For example, applying an operator to the orientation profile results in a bend deformation only and does not result in a stretch or bulge. This implies the ability to perform successive operator-based object deformations of varying amplitudes, types, locations or scales, which can be expressed as

$$p_d\,(m) = \bar{p}_d\,(m) + \sum_l \sum_s \sum_t \alpha_{dlst} k_{dlst}\,(m). \tag{5}$$

Examples of operator-based deformations are shown in Figure 11a–d.

Statistical Shape Analysis by Hierarchical Regional PCA In many applications, prior knowledge about object shape variability is available or can be obtained by studying a training set of shape examples. The training set is typically created by labeling corresponding landmark points in each shape example. Principal Component Analysis (PCA) is then applied to the training set, resulting in a point distribution model (PDM) [25]. The PDM describes the main modes of variation of the landmark positions and the amount of variation each mode explains. A drawback of this original approach is that the result of varying the weight of a single variation mode generally causes all the landmark positions to change. In other words, although the original PDM model produces only feasible shape deformations, a desirable trait, it generally produces global deformations over the entire object. Our goal is to utilize prior knowledge and produce feasible deformations, while also controlling the scale and location of these deformations. Toward this end, we perform a multiscale (hierarchical) multi-location (regional) PCA on a training set of medial shape profiles. To achieve this, we collect spatially corresponding subprofiles from the shape profiles. The length of a subprofile reflects the scale over which the analysis is performed. The principal component analysis is now a function of the location, scale, and type of shape profile (length, orientation, or thickness). Thus, for each location, scale, and shape profile type, we obtain an average medial subprofile, the main modes of variation, and the amount of variation each mode explains. The result is that we can now generate a feasible stretch, bend, or bulge deformation at a specific location and scale.

A shape profile can now be written as the sum of the average profile and the weighted modes of variation as follows:

$$p_d\left(m\right) = \bar{p}_d\left(m\right) + M_{dls} w_{dls}, \tag{6}$$

where p, d, $\bar{p}$, $p_d(m)$, l, and s are defined in (4); M_{dls} are variation modes (columns of M) for specific d, l, and s; and w_{dls} are weights of the variation modes, where the weights are typically set such that the variation is within three standard deviations.

For any shape profile type, multiple variation modes can be activated by setting the corresponding weighting factors to nonzero values. Each variation mode acts at a certain location and scale; hence we obtain

$$p_d\left(m\right) = \bar{p}_d\left(m\right) + \sum_l \sum_s M_{dls} w_{dls}. \tag{7}$$

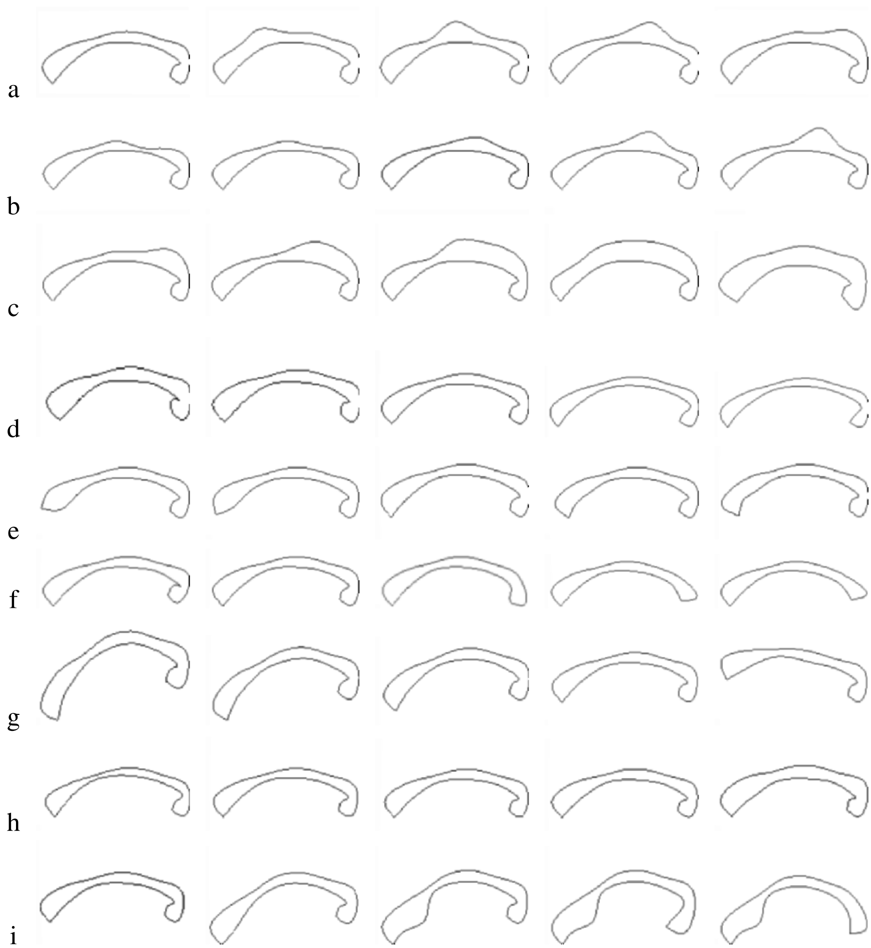

Figure 11. Examples of Controlled Deformations. (a)–(c) Operator-based bulge deformation at varying locations, amplitudes, and scales. (d) Operator-based stretching with varying amplitudes over entire CC. (e)–(g) Statistics-based bending of the left end, the right end, and the left half of the CC. (h) Statistics-based bulge of the left and right thickness over the entire CC. (i) From left to right: (1) mean shape, (2) statistics-based bending of the left half, followed by (3) locally increasing the left thickness using an operator, followed by (4) applying an operator-based stretch and (5) an operator-based bend to the right side of the corpus callosum. Reprinted with permission from [63]. Copyright ©2004, World Scientific.

In summary, varying the weights of one or more of the variation modes alters the length, orientation, or thickness profiles and generates statistically feasible stretch, bend, or bulge deformations at specific locations and scales upon reconstruction. Examples of statistics-based deformations are shown in Figure 11e–h.

Combining Operator- and Statistics-Based Deformations In general, operator- and statistics-based deformations ((5) and (7)) can be combined as

$$p_d = \bar{p}_d + \sum_l \sum_s \left(M_{dls} w_{dls} + \sum_t \alpha_{dlst} k_{dlst} \right). \tag{8}$$

It is worth noting that several deformations, whether operator- or statistics-based, may spatially overlap. Furthermore, adding profiles of different scales, and hence different vector lengths, is possible by padding the profiles with zeros. Figure 11i shows an example of combining operator- and statistics-based deformations.

Hand-Crafted Application To demonstrate the potential of statistics- and operator-based controlled deformations, we handcrafted a deformation schedule for fitting the CC shape model to a midsagittal MRI slice of the brain. Figure 12 shows the resulting medial shape profiles after applying the fitting schedule (compare with the initial profiles in Figure 10). The initial and final CC shapes are shown in Figure 13. The schedule steps are shown in Table 1 and the resulting deformed CC shapes for each step of the schedule are depicted in Figure 14.

3.3.2. 3D Shape Representation and Deformation using Medial Patches

We now extend the idea of the medial profiles method for 2D shape representation and deformation (presented in Section 3.3.1) to 3D. Instead of a single 1D orientation profile and a 1D elongation profile describing a primary medial axis, we now require two 2D orientation patches and a 2D elongation patch that together describe a primary medial sheet. Further, instead of two 1D thickness profiles describing the boundary of a 2D shape, we now require two 2D thickness patches to describe the surface the 3D object. Furthermore, rather than using 1D operators to change the profiles and generate new 2D shapes upon reconstruction, we now require 2D operators to alter the patches and generate new 3D shapes. As before, an operator applied to the orientation, elongation, or thickness components will cause a bend, stretch, or bulge shape deformation, respectively. In a simple manner, the location, extent, type, and amplitude of deformation operators can be specified and combined to capture local and global intuitive shape deformations. Figure 15 demonstrates the capability of producing different types of spatially localized deformations by varying different operator parameters. Deformation parameters include deformation type (e.g., bending, stretching, bulging), location, extent, amplitude, and type (e.g., Gaussian, rectangular, pyramidal, spherical).

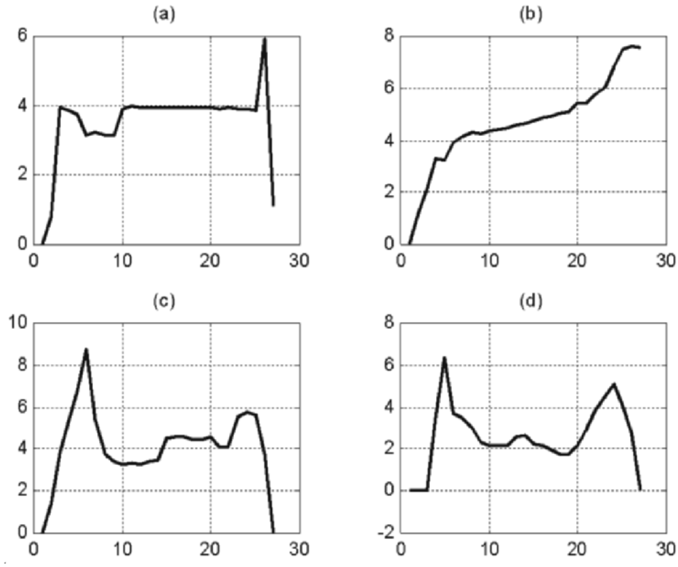

Figure 12. Resulting medial shape profiles after applying the fitting schedule: (a) length profile $L(m)$, (b) orientation profile $O(m)$, (c) left thickness profile $T^l(m)$, and (d) right thickness profile $T^r(m)$. Reprinted with permission from [63]. Copyright ©2004, World Scientific.

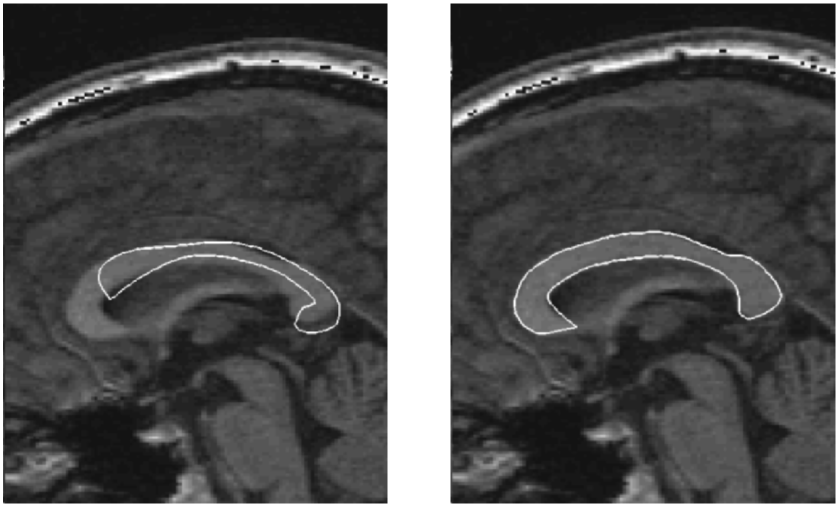

Figure 13. Close-up of the initial and final stages of the handcrafted fitting schedule. Reprinted with permission from [63]. Copyright ©2004, World Scientific.

Table 1. Deformation Schedule Used to Fit the Corpus Callosum Shape Model to MRI Data

Step	Deformation	Location	Scale	Variation mode/ Operator type	Variation mode weight/ Operator amplitude
1	Translation by ▼ 74, ▶ 24)				
2	Rotation counterclockwise by 10^o				
3	Scaling by 1.2				
4	Bend	1	8	2	$w = 0.5$
5	Bend	20	8	2	$w = -0.8$
6	Bend	22	6	2	$w = -0.75$
7	Bend	24	4	1	$w = 2.2$
8	Bend	1	4	2	$w = 1$
9	Stretch	6	4	1	$w = -1.5$
10	Stretch	26	1	1	$w = 2$
11	Left-bulge	15	7	1	$w = 3$
12	Left-bulge	18	3	1	$w = 2$
13	Left-bulge	6	12	1	$w = 3$
14	Left-bulge	5	3	1	$w = 3$
15	Right-squash	9	3	1	$w = -1$
16	Right-bulge	13	2	1	$w = 0.5$
17	Left-bulge	21	3	Gaussian	$\alpha = 0.3$
18	Left-bulge	21	7	Gaussian	$\alpha = 0.1$
19	Right-squash	24	2	Gaussian	$\alpha = -0.5$
20	Right-bulge	4	2	Bell	$\alpha = 1.7$
21	Right-bulge	6	3	Gaussian	$\alpha = 0.4$
22	Right-squash	1	3	Gaussian	$\alpha = -2.2$
23	Right-squash	25	1	Gaussian	$\alpha = -0.8$

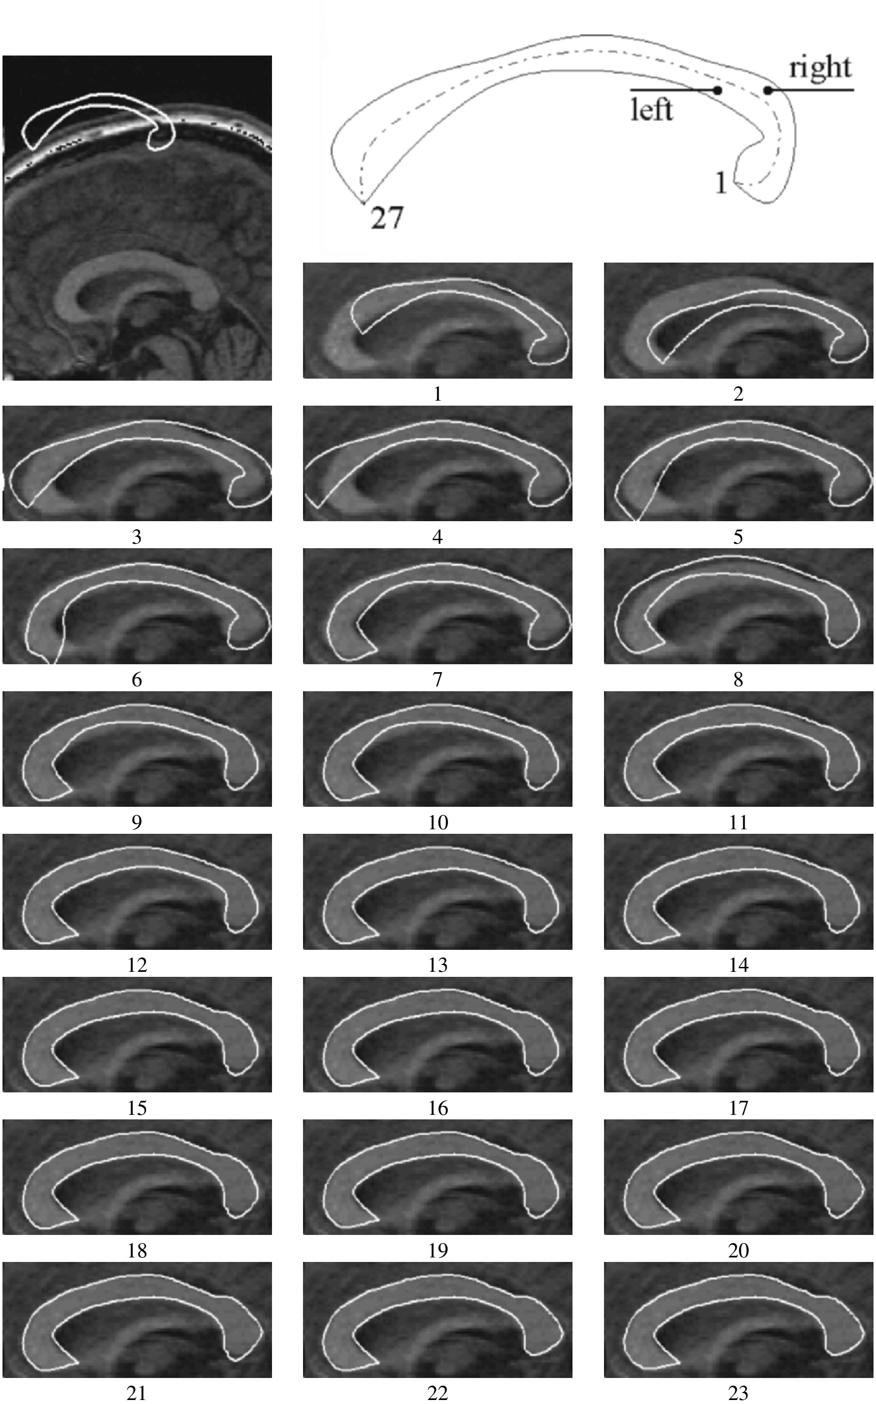

Figure 14. Progress of the handcrafted fitting schedule (fitting steps listed in Table 1). Reprinted with permission from [63]. Copyright ©2004, World Scientific.

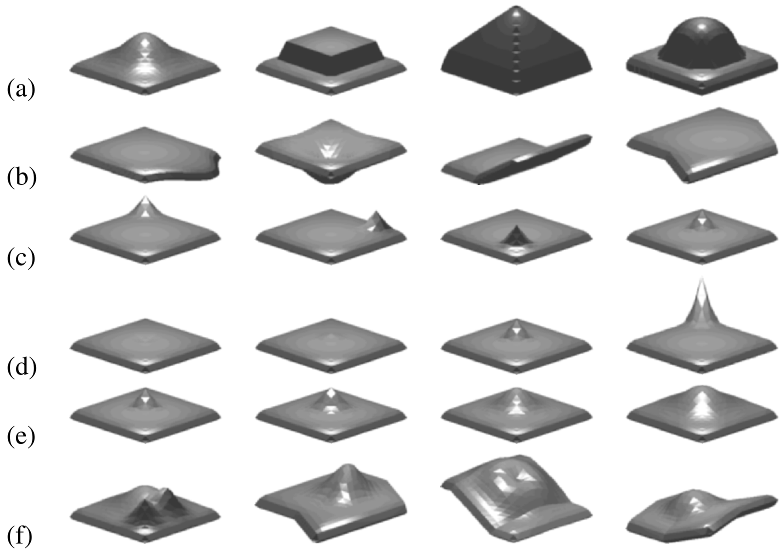

Figure 15. Deformations on a synthetic (slab) object. (a) Different operator types (left to right): Gaussian, rectangular, pyramidal, spherical. (b) Different deformation types: stretching, longitudinal bend, latitudinal bend. Different (c) locations, (d) amplitudes, and (e) extent of a bulge operator. (f) Combining stretching, longitudinal and latitudinal bend, and bulging deformations.

Because the effects of the deformation parameters are easily understood even to the user not familiar with the details of the shape representation, intuitive and accurate production of desired deformations is made possible. Results on real brain caudate nucleus and ventricle structures are presented in Section 4.5.

3.3.3. 2D Physics-based Shape Deformations

In this section we introduce the use of the physics-based shape representation and deformations method that is used for the deformable organism framework's geometrical and physical layers. This yields additional robustness by allowing intuitive real-time user guidance and interaction when necessary and addresses several of the requirements for an underlying shape representation and deformation layers stated previously. The deformations are realized through a physics-based framework implemented as meshes of connected nodes (mass-spring models). This inherently handles smoothness and continuity constraints, maintains the structural integrity of the body as it deforms, and facilitates intuitive user interaction. The mesh nodes and mesh connectivity are based on the medial axis of the object, yielding a shape representation and deformation method that naturally follows the geometry of the object. Various types of deformations at multiple locations and scale are controlled via operator- or statistics-based deformations that actuate springs. The method also provides statistics-based feasible deformations that are derived from a hierarchical (multiscale) regional (multi-location) principal component analysis.

In the following sections we present the construction of the dynamic mesh model, operator-based shape deformations, similarity transformations, and statistically based shape deformations, using internal spring actuation and external forces. We demonstrate results on both synthetic data and on a spring-mass model of the CC, obtained from 2D midsagittal brain MRI slices.

Dynamic Spring-Mass Mesh Model For our physics-based deformable organisms we use spring-mass models to represent object shapes (Figures 16 and 17). The model is made up of nodes (masses or particles) and springs (elastic links or

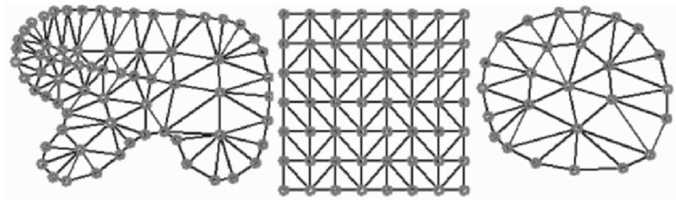

Figure 16. Examples of different synthetic spring-mass structures. Reprinted with permission from [80]. See attached CD for color version.

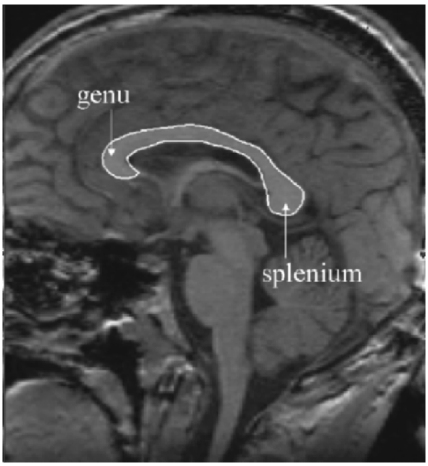

(a)

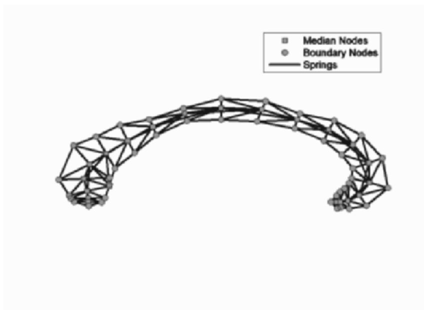

(b)

Figure 17. (a) Midsagittal MR brain image with the corpus callosum (CC) outlined in white. (b) CC mesh model showing medial and boundary masses. Reprinted with permission from [80]. Copyright ©2005, SPIE.

connecting segments), and can be deformed according to Eqs. (22)–(26) in Section 1.6.7 in Chapter 11.

Specifically, the deformations can be caused by internal or external forces. External forces are those external to the organism, including image forces and user-driven mouse forces (Figure 18), whereas internal forces result from spring actuation (in a manner analogous to muscle actuation in animals), causing the attached nodes to change position, thereby propagating the force throughout the model (Figure 19). Spring actuation is accomplished by changing the spring's rest length while continuously simulating the mesh dynamics.

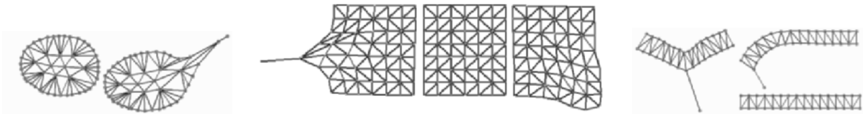

Figure 18. Examples of deformations via user interaction ("mouse" forces). Reprinted with permission from [62]. Copyright ©2003, SPIE. See attached CD for color version.

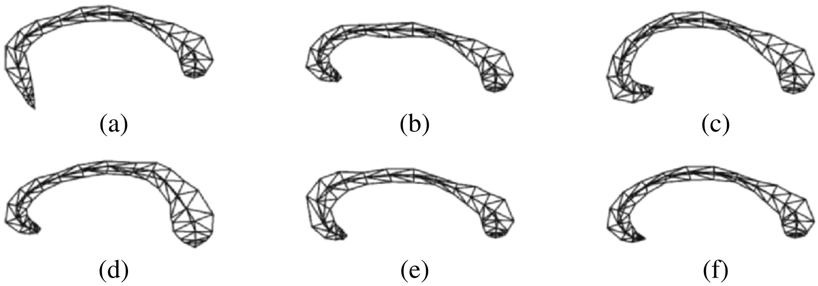

Figure 19. Examples of physics-based deformations of the CC organism using (a) user applied and (b) rotational external forces. Operator based (c) bending, (d) bulging, and (e) stretching deformations. (f) Statistics-based spring actuation. Reprinted with permission from [62]. Copyright ©2003, SPIE.

In order for deformable organisms to explore the image space, fit to specified regions, and take new shapes, they must be able to undertake a sequence of deformations. Some of these deformations take place independent of the topological design of the organism (general deformations), while others are designed specifically for medial-axis based worm-type organisms (medial-based deformations). Furthermore, in specific applications the general deformations can be modified to apply to specific anatomical regions of the model. By automatically fixing a group of nodes in place, the organism can perform deformations on specific regions of its geometrical model without affecting others. We use the terms *regional rotation/translation* and *boundary expansion* to refer to these types of deformations.

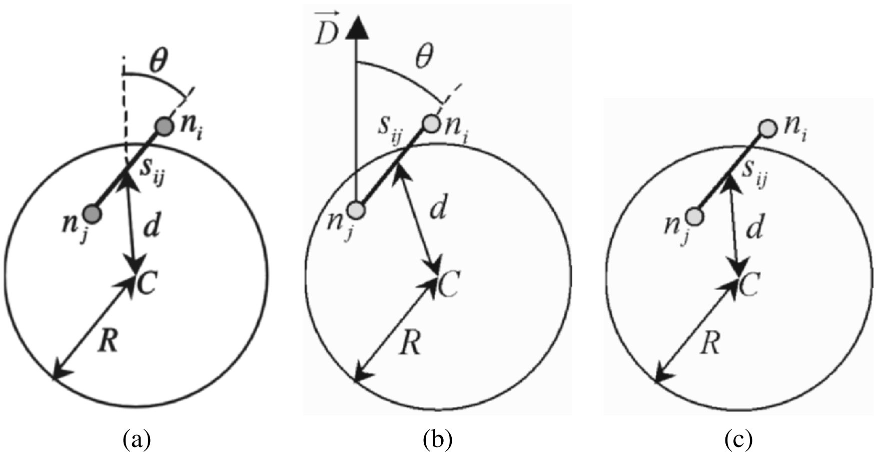

Figure 20. Definition of variables for (a) radial bulge, (b) directional bulge, and (c) localized scaling.

Boundary expansion refers to a sequence of stretching deformations whose direction is along thickness springs, while all nodes aside from the concerned boundary node are fixed. What follows is a description of the various types of deformations enabled.

Operator-Based Localized Deformations Using Internal Spring Actuation Localized bulging (radial bulge), stretching (directional bulge), scaling, bending, and tapering deformations are implemented using spring actuation. These operator-based deformations can be applied at different locations and scales with varying amplitudes. In the following paragraphs we describe how spring rest length must be changed to produce these different deformations.

To perform a (radial) bulge deformation we specify a center C and a radius R of a deformation region (Figure 20a) as well as a deformation amplitude K. We then update the rest length r_{ij} of each spring s_{ij} if at least one of its terminal nodes, n_i or n_j, lies within the deformation region, as follows:

$$r_{ij} = \left(\left(1 - \frac{d}{R}\right)\left(1 - \frac{2\theta}{\pi}\right)(K-1) + 1\right) r_{ij}^{\text{old}}, \tag{9}$$

where $\theta \in \left[0, \frac{\pi}{2}\right]$ is the angle between s_{ij}, and line L connecting the midpoint of the spring with C and d is the length of L (Figure 20a). The resulting effect of the above equation is that springs closer to C and with directions closer to the radial direction are affected more (Figure 21).

To perform a stretch (directional bulge) we again specify a deformation region and amplitude as well as a direction $\vec{D}$ (Figure 20b). We update the rest length of

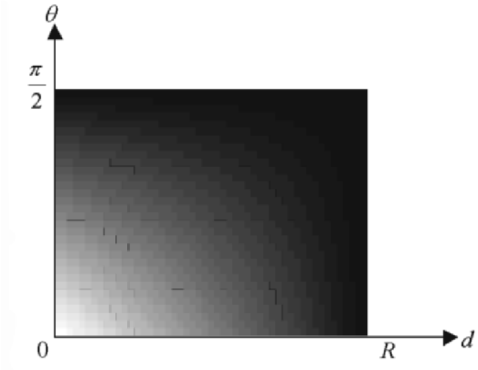

Figure 21. The coefficient (white = K, black = 1) by which r_{ij}^{old} is multiplied as a function of θ and d. Reprinted with permission from [62]. Copyright ©2003, SPIE.

each spring as in Eq. (9), where $\theta \in \left[0, \frac{\pi}{2}\right]$ is now defined as the angle between s_{ij} and $\vec{D}$ (Figure 20b). The resulting effect in this case is that springs closer to C and with directions closer to the stretch deformation direction are affected more (Figure 21).

A localized scaling deformation is independent of direction and requires only specification of a deformation region and amplitude (Figure 20c). The rest length update equation then becomes

$$r_{ij} = \left((1 - d/R)\,(K - 1) + 1\right) r_{ij}^{\text{old}}. \tag{10}$$

To perform localized bending, we specify bending amplitude K and two regions surrounding the medial axis (Figure 22). The rest lengths of the springs on one side of the medial are increased according to

$$r_{ij}^{1} = \left(\frac{d_1}{R_1}\left(1 - \frac{2\theta_1}{\pi}\right)(K - 1) + 1\right) r_{ij}^{1,\text{old}}, \tag{11}$$

while the rest lengths on the other side are decreased according to

$$r_{ij}^{2} = \left(\frac{d_2}{R_2}\left(1 - \frac{2\theta_2}{\pi}\right)\left(\frac{1}{K} - 1\right) + 1\right) r_{ij}^{2,\text{old}}. \tag{12}$$

To perform localized tapering, we specify tapering amplitude K and a region with a base (Figure 22). The rest lengths on one side of the base are increased according to

$$r_{ij}^{1} = \left(\frac{d_1}{R_1}\,(K - 1) + 1\right) r_{ij}^{1,\text{old}}, \tag{13}$$

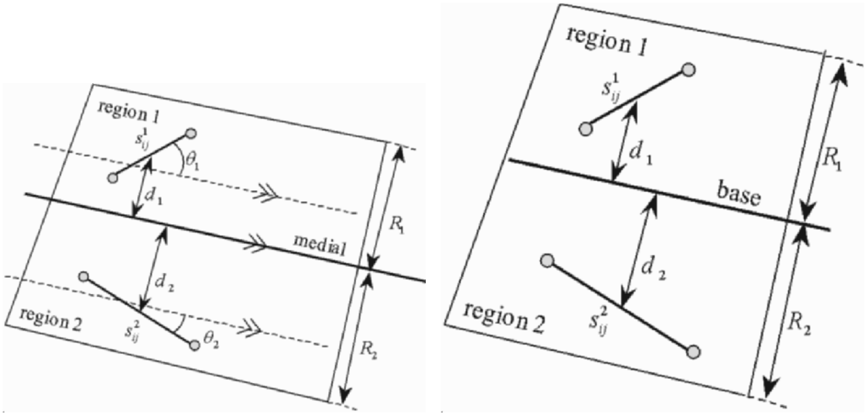

Figure 22. Definition of variables for deformation operators: bending (left), and tapering (middle). Reprinted with permission from [62]. Copyright ©2003, SPIE.

while those on the other side are decreased according to

$$r_{ij}^2 = \left(\frac{d_2}{R_2}\left(\frac{1}{K} - 1\right) + 1\right) r_{ij}^{2,\mathrm{old}}. \tag{14}$$

Different examples of localized operator-based deformations are shown in Figure 25.

Similarity Transformations Using Internal and External Forces Simple rotation and translation are implemented via the application of external forces, while scaling is implemented by muscle actuation. Rotation forces are applied on all nodes in a direction normal to the line connecting each node with the center of mass of the model, with a consistent clockwise/counterclockwise direction (Figure 23a). While translation forces are applied on all nodes in the direction of the desired translation (Figure 23b). Scaling by a factor of S is performed by changing the rest length of all the springs, i.e., $r_{ij} = S \cdot r_{ij}^{\mathrm{old}}$. Examples are shown in Figure 24.

Learned Deformations Using Internal Spring Actuation Learned or statistically based deformations are implemented via spring actuation. To facilitate intuitive deformations, springs are designed to be of different types: stretch springs, bend springs, or thickness springs. Stretch springs connect neighboring medial nodes, bending springs are hinge springs that connect nonconsecutive medial nodes, and thickness springs connect medial nodes with boundary nodes (Figure 26). Actuating the stretch springs causes stretch deformations, actuating hinge springs causes bend deformations, and actuating thickness springs causes bulging, squashing, or tapering deformations.

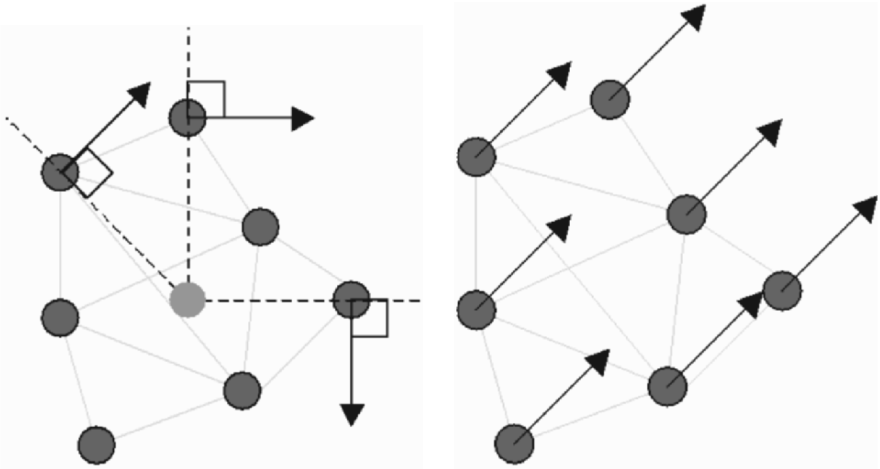

(a) (b)

Figure 23. External forces for performing a (a) rotation (light gray circle marks center of mass) and a (b) translation. Reprinted with permission from [62]. Copyright ©2003, SPIE. See attached CD for color version.

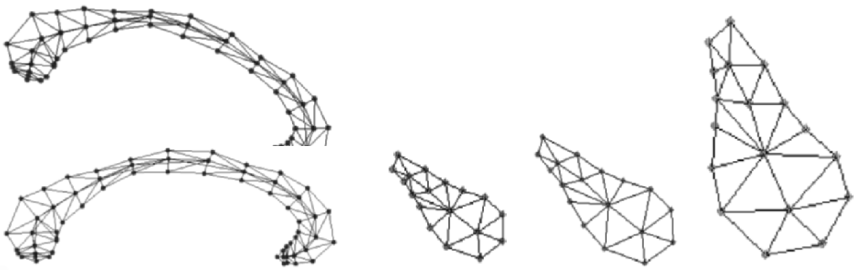

(a) (b)

Figure 24. Similarity transformation via external forces. (a) Rotating a model of the CC. (b) Scaling and rotating a synthetic model. Reprinted with permission from [62]. Copyright ©SPIE 2005.

Feasible mesh deformations (i.e., similar to what has been observed in a training set) at different locations and scales are obtained by actuating springs according to the outcome of a statistical analysis. Specifically, PCA is applied to a training set of spring rest lengths in the region corresponding to the desired localized deformations.

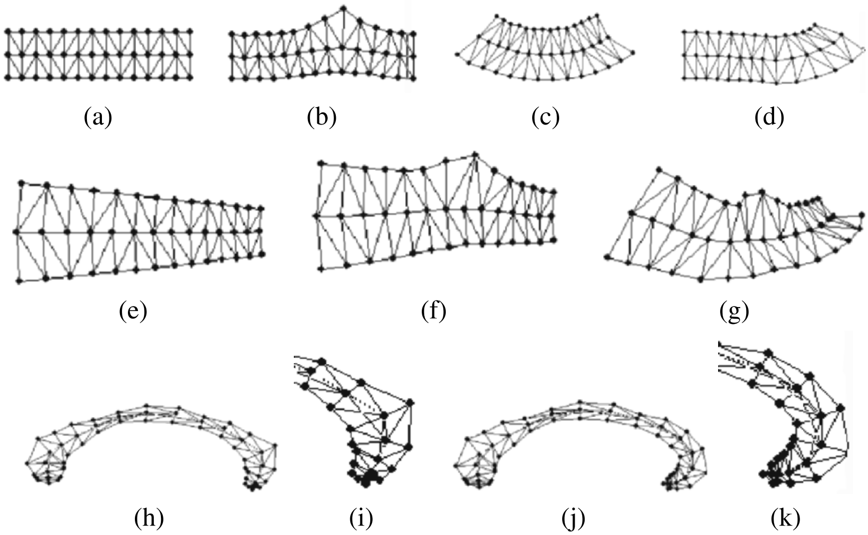

Figure 25. Examples of localized deformations: (a) initial synthetic object, (b) bulge, (c) bend, (d) bend at another location, (e) tapering, (f) tapering followed by a bulge, and (g) tapering followed by a bulge and a bend deformations. CC model (h) before and (j) after a localized bend. (i,k) Close-up versions of (h,j). Reprinted with permission from [62]. Copyright ©2003, SPIE.

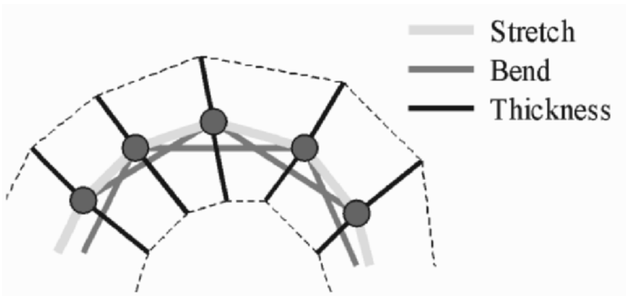

Figure 26. Spring types used for statistics-based deformations. Reprinted with permission from [62]. Copyright ©2003, SPIE. See attached CD for color version.

The set of rest lengths for the stretch springs (Figure 26) in a single example model are collected in a vector $\mathbf{r}_S$, i.e.,

$$\mathbf{r}_S = \{r_{ij} \forall i, j : s_{ij} \in \text{stretch springs}\}, \tag{15}$$

and similarly for the bending and left and right thickness springs: (Figure 26)

$$\begin{aligned} \mathbf{r}_B &= \{r_{ij} \forall i, j : s_{ij} \in \text{bend springs}\}, \\ \mathbf{r}_{TL} &= \{r_{ij} \forall i, j : s_{ij} \in \text{left thickness springs}\}, \\ \mathbf{r}_{TR} &= \{r_{ij} \forall i, j : s_{ij} \in \text{right thickness springs}\}. \end{aligned} \tag{16}$$

This gives

$$\begin{aligned} \mathbf{r}_S &= \left[\mathbf{r}_S^1, \mathbf{r}_S^2, \ldots, \mathbf{r}_S^{N_S}\right], \\ \mathbf{r}_B &= \left[\mathbf{r}_B^1, \mathbf{r}_B^2, \ldots, \mathbf{r}_B^{N_B}\right], \\ \mathbf{r}_{TL} &= \left[\mathbf{r}_{TL}^1, \mathbf{r}_{TL}^2, \ldots, \mathbf{r}_{TL}^{N_T}\right], \\ \mathbf{r}_{TR} &= \left[\mathbf{r}_{TR}^1, \mathbf{r}_{TR}^2, \ldots, \mathbf{r}_{TR}^{N_T}\right], \end{aligned} \tag{17}$$

where N_S, N_B, N_T are the numbers of stretch, bend, and left/right thickness springs, and the springs are ordered spatially (i.e., moving from one end of the medial to the other we encounter $\mathbf{r}_S^1, \mathbf{r}_S^2, \ldots, \mathbf{r}_S^{N_S}$). Performing global (traditional) PCA on corresponding variables in a training set gives

$$\begin{aligned} \mathbf{r}_S &= \bar{\mathbf{r}}_S + M_S \mathbf{w}_S, \\ \mathbf{r}_B &= \bar{\mathbf{r}}_B + M_B \mathbf{w}_B, \\ \mathbf{r}_{TR} &= \bar{\mathbf{r}}_{TR} + M_{TR} \mathbf{w}_{TR}, \\ \mathbf{r}_{TL} &= \bar{\mathbf{r}}_{TL} + M_{TL} \mathbf{w}_{TL}, \end{aligned} \tag{18}$$

where the columns of M_S, M_B, M_{TR}, M_{TL} are the main modes of spring length variation. Associated with each mode is the variance it explains.

For capturing the shape variations at different locations and scales, we study the variations in the rest lengths of the springs in the desired localized region. Furthermore, to decompose the variations into different types of general deformations, each statistical analysis of the spring length in a localized region is restricted to a specific type of deformation springs (Figure 26). Accordingly, the PCA becomes a function of the deformation type, location, and scale. For example, to analyze the local variation in object length (stretch), we perform a statistical analysis on the lengths of the stretch springs of that local region. In general, for a single deformation/location/scale- specific PCA we obtain

$$\mathbf{r}_{\text{def,loc,scl}} = \bar{\mathbf{r}}_{\text{def,loc,scl}} + M_{\text{def,loc,scl}} \mathbf{w}_{\text{def,loc,scl}}, \tag{19}$$

where def is the deformation type, being either S (for stretch), B (for bend), TL (for left thickness), or TR (for right thickness). The location and scale, determined by the choice of loc and scl, respectively, determine which springs are to be included in the analysis according to

$$\mathbf{r}_{\text{def,loc,scl}} = \left[\mathbf{r}_{\text{def}}^{\text{loc}}, \mathbf{r}_{\text{def}}^{\text{loc}+1}, \ldots, \mathbf{r}_{\text{def}}^{\text{loc}+\text{scl}-1}\right]. \tag{20}$$

For example, for the bending deformation at location "five" with scale "three" $(\text{def}, \text{loc}, \text{scl} = B, 5, 3)$ we have

$$\mathbf{r}_{\text{def,loc,scl}} = \mathbf{r}_{B,5,3} = \left[\mathbf{r}_B^5, \mathbf{r}_B^6, \mathbf{r}_B^7\right]. \tag{21}$$

The average values of the spring lengths are calculated according to

$$\bar{\mathbf{r}}_{\text{def,loc,scl}} = \frac{1}{N}\sum_{j=1}^{N} \mathbf{r}\,(j)_{\text{def,loc,scl}}, \tag{22}$$

where $\mathbf{r}\,(j)_{\text{def,loc,scl}}$ is $\mathbf{r}_{\text{def,loc,scl}}$ obtained from the jth training example, and N is the number of training examples. The columns of $M_{\text{def,loc,scl}}$ are the eigenvectors, $m_{\text{def,loc,scl}}$, of covariance matrix $C_{\text{def,loc,scl}}$. That is,

$$\{C\mathbf{m} = \lambda\mathbf{m}\}_{\text{def,loc,scl}}, \tag{23}$$

where

$$\left\{C = \frac{1}{N-1}\sum_{j=1}^{N} \left(r(j) - \bar{r}\right)\left(r(j) - \bar{r}\right)^T\right\}_{\text{def,loc,scl}}, \tag{24}$$

and where $\{\}_{\text{def,loc,scl}}$ denotes deformation type-, location-, and scale-specific PCA variables.

The data set needs to be aligned only with respect to scale. The statistical analysis of spring lengths is independent of orientation and translation. See the different examples in Figures 27–30.

3.3.4. 3D Physics-Based Shape Deformations

Increasing the dimensionality of the position, force, velocity, and acceleration vectors to 3D in Eqs. (22)–(26) (Section 1.6, Chapter 11) enables calculation of the 3D deformations, as opposed to 2D. External forces are provided by the volumetric image gradient and a drag force, while internal forces are supplied through Hooke's law and dampening spring forces.

Here we present a tubular geometric module parametrized by the distance between neighboring medial masses and the number of circumferential boundary masses (Figure 31). This layered medial shape representation enables intuitive deformations [62, 63, 76] wherein the medial axis governs the bending and

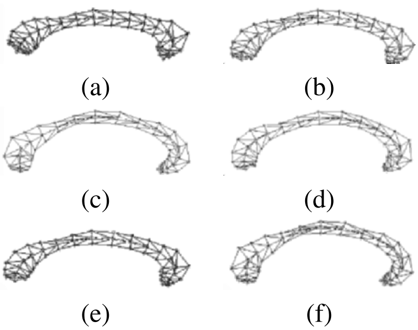

Figure 27. Sample corpus callosum mesh model deformations (1st PC for all deformation types over the entire CC) derived from the hierarchical regional PCA. Reprinted with permission from [62]. Copyright ©2003, SPIE.

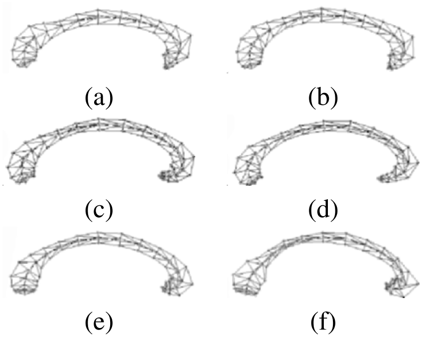

Figure 28. Sample CC mesh model deformations (2nd PC for all deformation types over the entire CC) derived from the hierarchical regional PCA. Reprinted with permission from [62]. Copyright ©2003, SPIE.

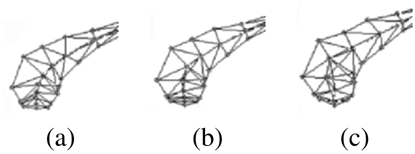

Figure 29. Statistical CC mesh model deformations: stretching the splenium. Reprinted with permission from [62]. Copyright ©2003, SPIE.

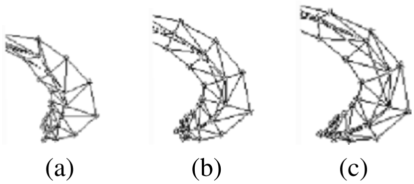

Figure 30. Statistical CC mesh model deformations: bending the genu. Reprinted with permission from [62]. Copyright ©2003, SPIE.

Figure 31. Topology of the vessel crawler showing: Masses (left), radial springs (middle), and stability springs (right) across sequential layers of the organism. See attached CD for color version.

stretching of the vessel crawler and links to boundary nodes to control thickness. As deformable organisms are typically modeled after their target structures, we employ this tubular topology to the task of vascular segmentation and analysis in volumetric medical images [86]. We provide results in Section 4.6.

3.4. Perception System

The perception system of the deformable organism consists of a set of sensors that provide image information to the organisms. Using virtual sensors an organism can collect data about the world they live in, such as image data or data from interaction with other organisms. For example, a single organism may have image intensity sensors, edge sensors, or texture sensors. A wide variety of image processing techniques can be used to enable the organism to perceive relevant features of their surroundings. Sensors can be focused or trained for a specific image feature in a task-specific way and hence the organism is able to disregard sensory information superfluous to its current behavioral needs. Different parts of the organism are dynamically assigned sensing capabilities and thus act as sensory organs or receptors. A sensor can either be on-board or off-board. On-board sensors are confined to the organism's body such as at its medial or boundary nodes, curves or segments connecting different nodes, or at internal geometric subregions, while off-board sensors are free floating. Once the sensory data are gathered, they are fed to the cognitive center of the brain for processing. In the following sections we highlight the main sensors used in a variety of deformable organisms.

- **Geometrical CC deformable organism:** In our implementation of the geometrical CC deformable organism, the sensors are made sensitive to different stimuli, including image intensity, image gradient magnitude and direction, multiscale edge-preserving diffusion filtered images [77, 78] a Canny-edge detected version of the image, and the result of a Hough transform applied to locate the top of the human skull in the brain MR image (results in Section 4.1).

- **Physics-based CC deformable organism:** This deformable organism was equipped with sensors for collecting statistical measures of image data such as mean and variance, in addition to intensity and edge strength and direction sensors (results in Section 4.4).

- **2D Vessel Crawler:** Once vessel-crawling deformable organisms are initialized in a target vascular system, they use a variety of sensory input modules to make decisions about which direction to grow in, and where bifurcations are located. For example, the 2D angiogram vessel crawler is equipped with an off-board arc sensor (Figure 32) in order to differentiate between overlapping vessels and authentic branch points. This crawler is also equipped with simple sensors for measuring image intensity and edge magnitude and direction (results in Section 4.3).

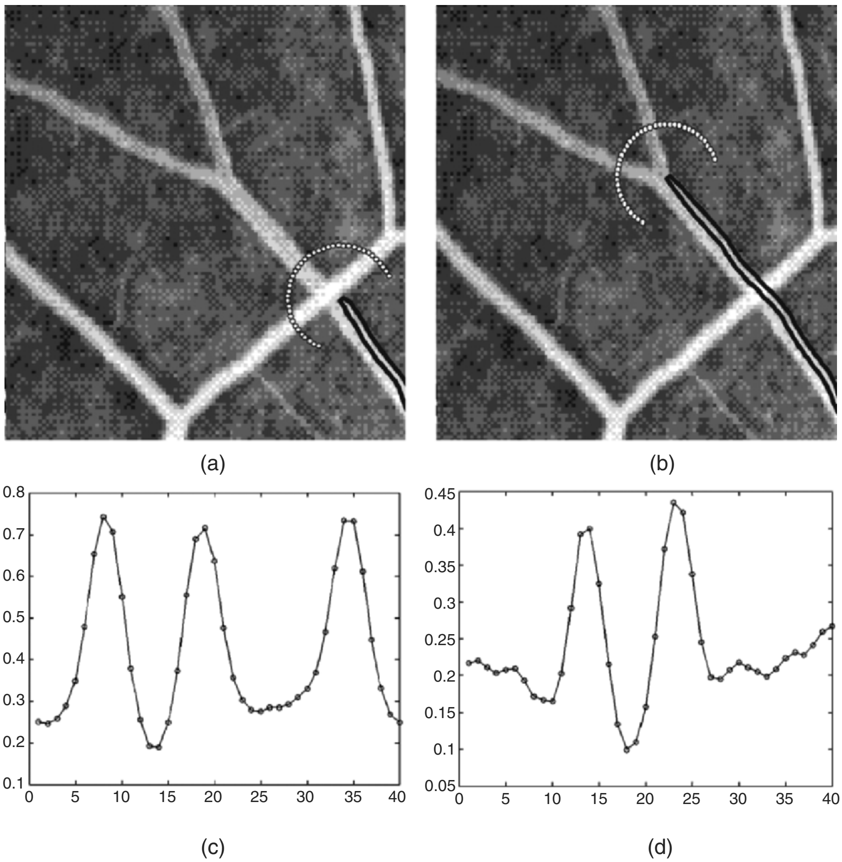

Figure 32. Off-board sensors (arc of white nodes in (a) and (b)) measure image intensity (along the arc). This results in an intensity profile exhibiting three distinct peaks when an overlapping vessel is ahead (c) and only two peaks in the case of a bifurcation (d).

- **3D vessel crawler:** The 3D vessel crawler utilizes two primary sensory modules, a vesselness sensor and a Hessian sensor, which both aide in guiding the crawler to the direction in which it should grow. The vesselness sensor is an extension of the 2D arc intensity sensor (Figure 32) to a 3D hemispherical sensor collecting vesselness values [79] (Figure 33). The Hessian sensor uses the smallest eigenvalued eigenvector $\vec{v}_1$ of Hessian matrix H.

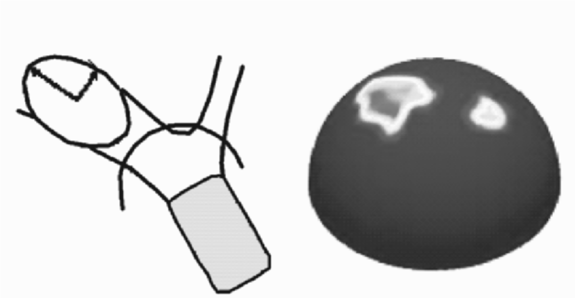

Figure 33. A vessel crawler (left, in gray) utilizing an off-board hemispherical sensor (shown as an arc in the left-hand image). The sensor (shown in 3D on the right) collects vesselness measures, guiding it as it crawls along the vessel and detects branching points. See attached CD for color version.

3.5. Behavioral Layer

A behavior is classically defined as the action or reaction of a life form in response to external or internal stimuli. The deformable organism's behaviors are a set of actions performable in response to perceived input or user control. Specifically, behavioral routines are designed based on available organism motor skills, perception capabilities, and available anatomical landmarks. For example, the routines implemented for the geometrically based CC deformable organism include: find-top-of-head, find-upper-boundary-of-CC, find-genu, find-rostrum, find-splenium, latch-to-upper-boundary, latch-to-lower-boundary, find-fornix, thicken-right-side, thicken-left-side, back-up (more details in Section 4.1). Each behavior routine subsequently activates the appropriate deformation or growth controllers to complete a stage in the plan and bring an organism closer to its target shape. For example, the latch-to-upper-boundary behavior causes the CC organism to latch the upper-half of itself to the top of the CC (Figure 34). In what follows we describe the primary behaviors of two example deformable organisms.

3.5.1. Example Behaviors of the Physics-Based CC Deformable Organism

We describe below the different behavioral routines necessary to achieve successful segmentation of the CC using the physics-based deformable organism (Figure 5b). Its anatomically tailored behaviors include: model initialization, global and regional alignment, expansion and contraction, medial alignment, fitting to boundary, and detecting and repairing segmentation inaccuracies. We describe the key behaviors below; for complete details refer to [80].

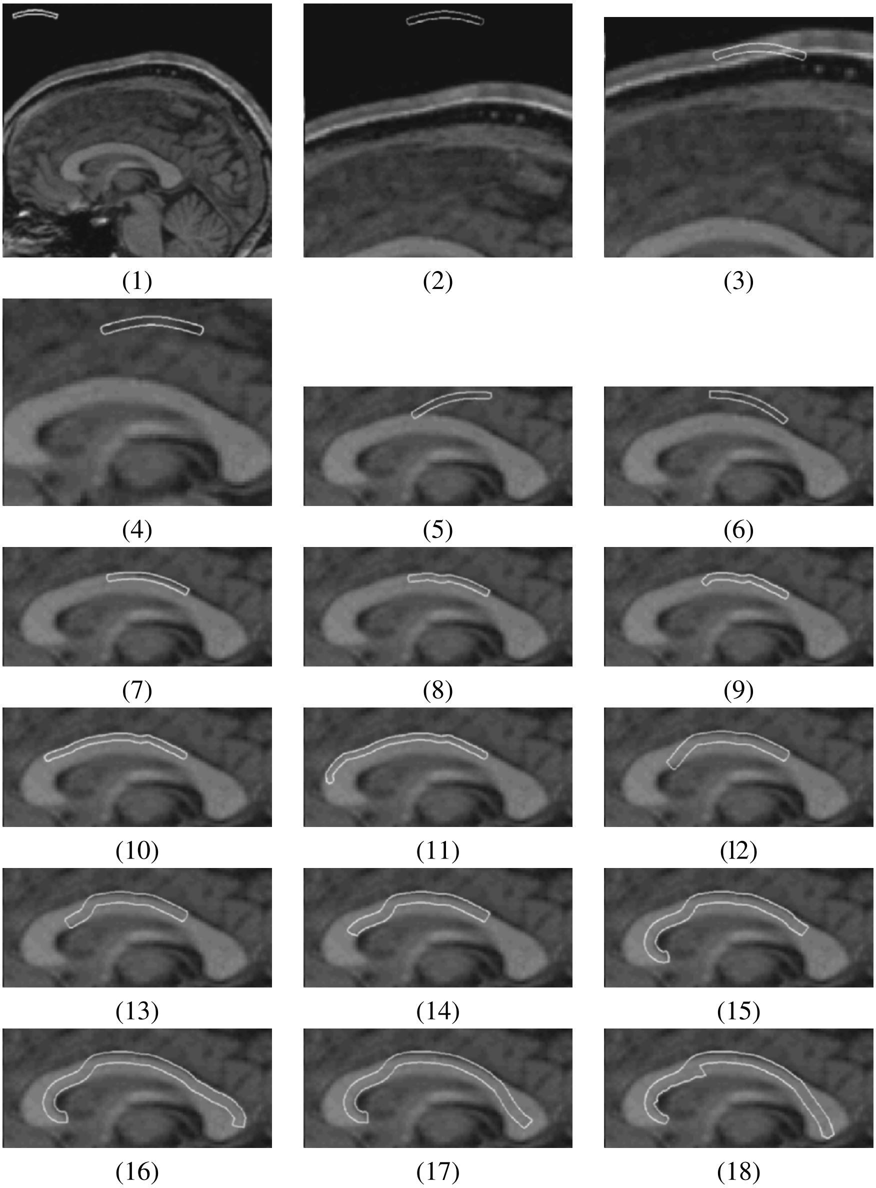

Figure 34. Deformable corpus callosum organism progressing through a sequence of behaviors to segment the CC (results continued on next page). Reprinted with permission from [85]. Copyright ©2002, Elsevier.

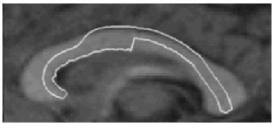
(19)
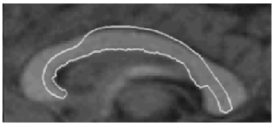
(20)
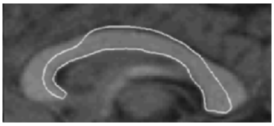
(21)
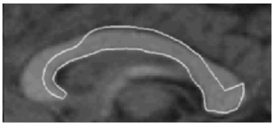
(22)
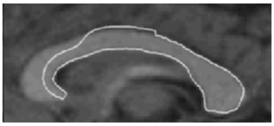
(23)
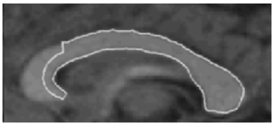
(24)
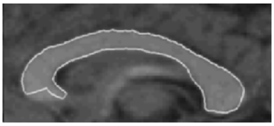
(25)
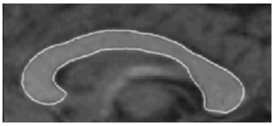
(26)
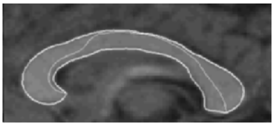
(27)
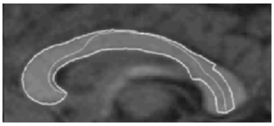
(28)
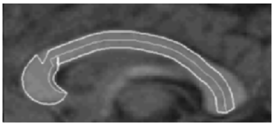
(29)
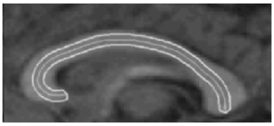
(30)
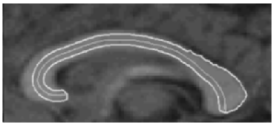
(31)
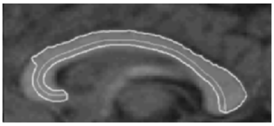
(32)
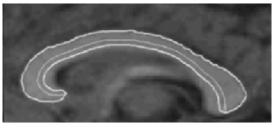
(33)
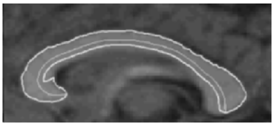
(34)
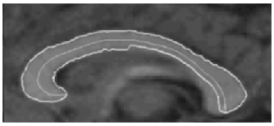
(35)
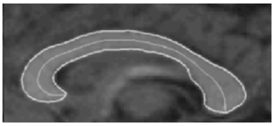
(36)

Figure 34. (continued). Results continued from previous page.

- **Model Initialization:** To begin its search for the CC, the deformable organism must be initialized to an appropriate location in the image using robust anatomical features. The organism's first step is to locate the top of the head using a modified Hough transform [81] to search for the largest ellipse (skull). Then, in response to the output of the skull sensor, a CC template is built consisting of a rectangular window connecting two squares to approximate the main body, genu, and splenium, respectively (see Figure 1). Maximizing average image intensity and minimizing intensity variance over the shape model yields several candidate locations for the three CC parts. Finally, a decision is made and the set of candidate parts that exhibits the strongest edge connectivity and exhibits maximal distance between the parts is selected.

- **Global and Regional Alignment:** The global alignment phase is designed to find the best overall location for the model using external forces applied to the entire organism to rotate and position it as necessary. The organism first locates the optimal horizontal and vertical location using translational forces, and then finds the best alignment using rotational forces. After searching its local space, optimal translations are chosen using a position decision based upon a sensory module configured to monitor local image mean, and variance across the entire model. As the CC is generally a bright, homogenous region, the decision function attempts to maximize the image mean, while minimizing the variance (more details presented in Section 3.6.1). During the regional alignment behavior, the organism aligns its anatomical parts to corresponding sections of the CC through successive phases. During each phase, rotational and translational forces are applied to only one part of the CC model. The phases are ordered as splenium-, genu-, and finally rostrum-alignment, as this particular ordering favors segmenting regions with stable features before proceeding to other regions. Optimal positions are decided upon using a position decision function utilizing the same sensory module as used in the global alignment phase.

- **Detecting and Repairing Segmentation Inaccuracies:** An example of the organism's ability to incorporate anatomical knowledge is its ability to detect and repair the fornix dip (Figure 1). Typically, the anatomical structure of a corpus callosum exhibits a symmetrical property about its medial axis. Consequently, if the organism is nonsymmetrical in the region where the fornix dip is located, then that area needs repair. In order to repair itself, the organism simply mirrors its top half about the medial axis throughout the affected area. To ensure validity, the organism then checks if the deformation has been beneficial using a position decision function designed to maximize local image mean, and minimize local image variance (Section 3.6.1).

3.5.2. Example Behaviors of the 3D Vessel Crawler

As the behaviors vary from organism to organism in both design and complexity, we also provide examples of those used in our 3D vessel crawler. Each of the vessel crawler's key decisions results in the execution of the appropriate behavior, using the concluded locally optimal parameters such as scale, estimated vessel mean and variance, etc. The primary behaviors available to the organism are to grow, to fit the vessel wall, and to spawn new organisms.

- **Growing:** Once initialized at a seed point, the vessel crawler must grow outward along the vasculature by sequentially adding new layers centered around the final sensor position using the 3D hemispherical and Hessian-based sensors (described in Section 3.4). As it grows, each new layer must be created and subsequently connected to the current end most layer (Figure 31). The newest layer is aligned to the provided direction vector, and then connected via a consistent clockwise ordering to prevent mesh twisting. Once connected the model can be fit to the image data.

- **Fitting**: The organism fits itself to the vessel boundary using 3D image gradient driven deformations simulated by the physics layer (Section 3.3.4). Connections to the previous layer provide smoothness, while stiffer circumferential springs provide local stability to noise, and flexible radial springs allow deformation to the vessel boundary (Figure 31).

- **Spawning new organisms**: In order to branch off and explore bifurcations, the organism must be able to spawn new vessel crawlers. Each spawned vessel crawler is initialized based on the optimal parameters detected by the parent. For example, the initial radius is based on the parent's detected radius at the spawn point using the spherical sensory module (Section 3.4).

3.6. Cognitive System

The cognitive layer of the architecture combines memorized information, prior anatomical knowledge, a segmentation plan, and an organism's sensory data (Section 3.4) in order to initiate behaviors, carry out shape deformations, change sensory input, and make decisions toward segmenting the target structure (Figure 4). A single fixed path or multiple paths with a plan selection scheme can be implemented. In the first case the organism follows a sequential flow of control, proceeding directly from one behavior to the next. In the latter case, the organism decides between different options within the plan, thus taking a different path than it might have given a different image or different initialization conditions.

Most often the segmentation plan is subdivided into different behavioral stages with subgoals that are easy to define and attain (e.g., locating the upper boundary of an anatomical structure). Consequently, the segmentation plan provides a means for human experts to intuitively incorporate global contextual knowledge. It contains instructions on how best to achieve a correct segmentation by optimally prioritizing behaviors. If we know, for example, that the superior boundary of the CC is consistently and clearly defined in an MRI image, then the find-upper-boundary behavior should be given a very high priority as segmenting stable features will provide good initialization for less-stable ones. Adhering to the segmentation plan and defining it at a behavioral level infuses the organism with awareness of the segmentation process. This enables it to make effective use of prior shape knowledge — it is applied only in anatomical regions of the target object where a high level of noise or gaps in the object's boundary edge are known to exist.

In the following two subsections we first describe a sequential deformable organism using a pre-designed segmentation plan with a known and fixed sequence of behaviors. As previously noted, a general and more flexible approach is to have different options available for which behaviors the organisms can carry out next. We then present and describe a deformable organism with an adaptive segmentation plane, under which the organism's sequence of behaviors is adaptively selected during its life span.

3.6.1. Sequential Behavior Selection: Physics-Based CC Deformable Organism

Sequential control takes form in the way of event or schedule driven actions, and sensory or knowledge-based decisions. The actions are initiated by the behavioral layer and simulated by the physics layer (Section 3.3.3). Decisions are made at scheduled times according to user-provided decision functions. Using specialized decision functions, the deformable organism can decide on the best position across a sequence of performed deformations, and whether or not to continue its expansion/contraction behavior (Figure 42). We provide one such example below; for complete details see [80].

- **Shape Configuration Decision.** The organism decides on the best new shape configuration (for the whole organism or for a part of its body) by first capturing sensory input as it is deforming to different shape configurations, and then choosing the configuration at which the sensory input maximizes a certain objective function. For example, the following strategy causes the deformable organism to prefer bright homogenous regions (like the CC) over dim nonuniform regions. The organism first collects n values for $d_i = \alpha m_i + (1 - \alpha)v_i^{-1}$, where m_i and v_i are image intensity and variance, respectively, at n different shape configurations; it then decides to take on the shape configuration that gives the maximum d.

3.6.2. Adaptive Behavior Selection: 3D Vessel Crawler

The vessel crawler can make a number of key decisions at any point in its life span, where each decision can be based on sensory input, anatomical knowledge, or user interaction. The outcome of each decision can affect the vessel crawler's internal settings, or how it should navigate the segmentation plan. It should be noted that the user is able to override any of these key decision functions at any time during the organism's life cycle, and hence can illicit intuitive real-time control over the segmentation process. A particular decision that affects the choice of subsequent behaviors of the organism is whether a bifurcation has been detected or not.

- **Bifurcation Verification:** The analysis of the vesselness values returned by the hemispherical sensor (Figure 33) is used to detect bifurcations. Basically, if two disjoint high-vesselness magnitude regions appeared on the hemispherical slice, then this would indicate a bifurcation. If a bifurcation is verified, then two new organisms are spawned, one for each branch, while a single high-vesselness region indicates a single branch and therefore a "grow behavior" is initiated.

4. RESULTS AND APPLICATIONS

4.1. Geometrical CC Deformable Organism

We first present a detailed segmentation plan for the CC organism that serves to illustrate an ability to harness global contextual knowledge. A CC organism is released into a 2D midsagittal MRI brain image from an initial default position (Figure 34.1). It then goes though different "behaviors" as it progresses toward its goal. As the upper boundary of the CC is very well defined and can be easily located with respect to the top of the head, the cognitive center of the CC activates behaviors to first locate the top of the head (Figure 34.2–3); it then moves downward through the gray and white matter in the image space to locate the upper boundary (Figure 34.4–7). The organism then bends to latch onto the upper boundary (Figure 34.8) and activates a find-genu routine (refer to Figure 1 for CC anatomy), causing the CC organism to stretch and grow along this boundary toward the genu (Figure 34.9–11). It then activates the find-rostrum routine causing the organism to back up, thicken (Figure 34.12), and track the lower boundary until reaching the distinctive rostrum (Figure 34.13–15). Once the rostrum is located, the find-splenium routine is activated and the organism stretches and grows in the other direction (Figure 34.15–16). The genu and splenium are easily detected by looking for a sudden change in direction of the upper boundary toward the middle of the head. At the splenium end of the CC, the organism backs up and finds the center of a circle that approximates the splenium end cap (Figure 34.17). The lower boundary is then progressively tracked from the rostrum to the splenium while maintaining parallelism with the organism's medial axis in order to avoid latching onto the potentially occluding fornix (Figure 34.18–21). Nevertheless, the lower boundary might still dip toward the fornix, so a successive step of locating where, if anywhere, the fornix does occlude the CC is performed by activating the find-fornix routine (making use of edge strength along the lower boundary, its parallelism to the medial axis, and statistical thickness values). Thus, prior knowledge is applied only when and where required. If the fornix does indeed occlude the CC, any detected dip in the organism's boundary is repaired by interpolation using neighboring thickness values (Figure 37). The thickness of the upper boundary is then adjusted to latch onto the corresponding boundary in the image (Figure 34.22–26). At this point the boundary of the CC is located (Figure 34.26), and the CC organism has almost reached its goal. However, at this stage the medial axis is not in the middle of the CC organism (Figure 34.27) so it is re-parametrized until the medial nodes are halfway between the boundary nodes (Figure 34.28–30). Finally, the upper and lower boundaries, which were reset in the previous step, are relocated (Figure 34.31–36) to obtain the final segmentation result (Figure 34.36). Other CC segmentation (Figure 35), validation results (Figure 36), and a demonstration of the organism's self-awareness (Figure 38) are presented.

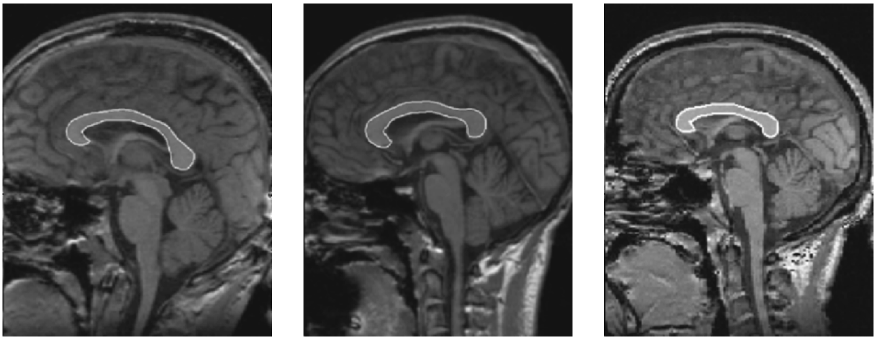

Figure 35. Segmentation results. Reprinted with permission from [38]. Copyright ©2002, Springer.

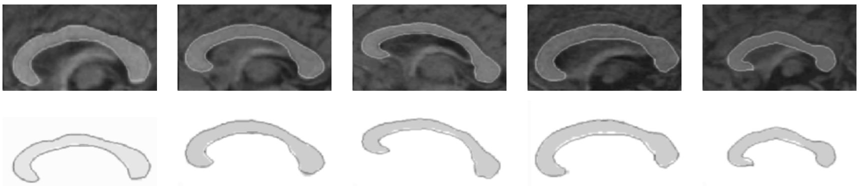

Figure 36. Segmentation results (top), also shown (in black) over manually segmented (gray) corpora callosa (bottom). Reprinted with permission from [85]. Copyright ©2002, Elsevier.

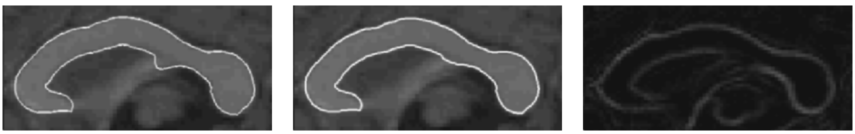

(a) (b) (c)

Figure 37. Segmentation result (a) before and (b) after detecting and repairing the fornix dip. (c) Note the weak gradient magnitude where the fornix overlaps the CC. Reprinted with permission from [85]. Copyright ©2002, Elsevier.

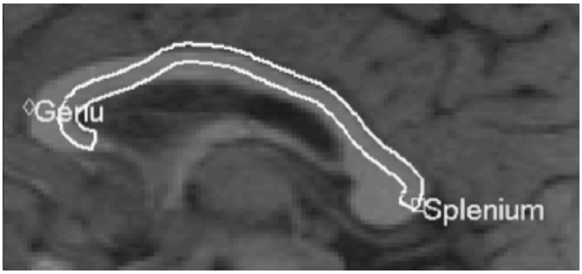

Figure 38. The CC organism's self-awareness enables it to identify landmark parts. Reprinted with permission from [85]. Copyright ©2002, Elsevier.

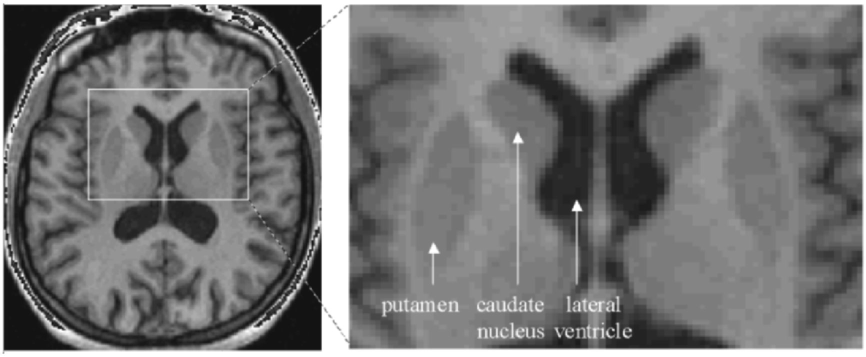

Figure 39. The lateral ventricle, caudate nucleus, and putamen shown in transversal brain MRI slice.

4.2. Simple Interacting Organisms

Simple interacting organisms were created to locate the lateral ventricles, caudate nuclei, and putamina in the left and right halves of transversal MR brain images (Figure 39). Since the ventricles are the most stable of the above structures, a ventricle organism is first released (Figure 40.1). It proceeds to locate the top of the ventricle (Figure 40.2) and its inner and outer (with respect to the brain) boundaries (Figure 40.3–5). Both ends of the ventricle organism are actively stretched to locate both the upper and the lower lobes of the ventricle (Figure 40.6). The ventricle organism then passes information about the shape and location of the segmented ventricle (Figure 40.7) to the caudate nucleus (CN) organism, which is initialized accordingly in a suitable position (Figure 40.8).

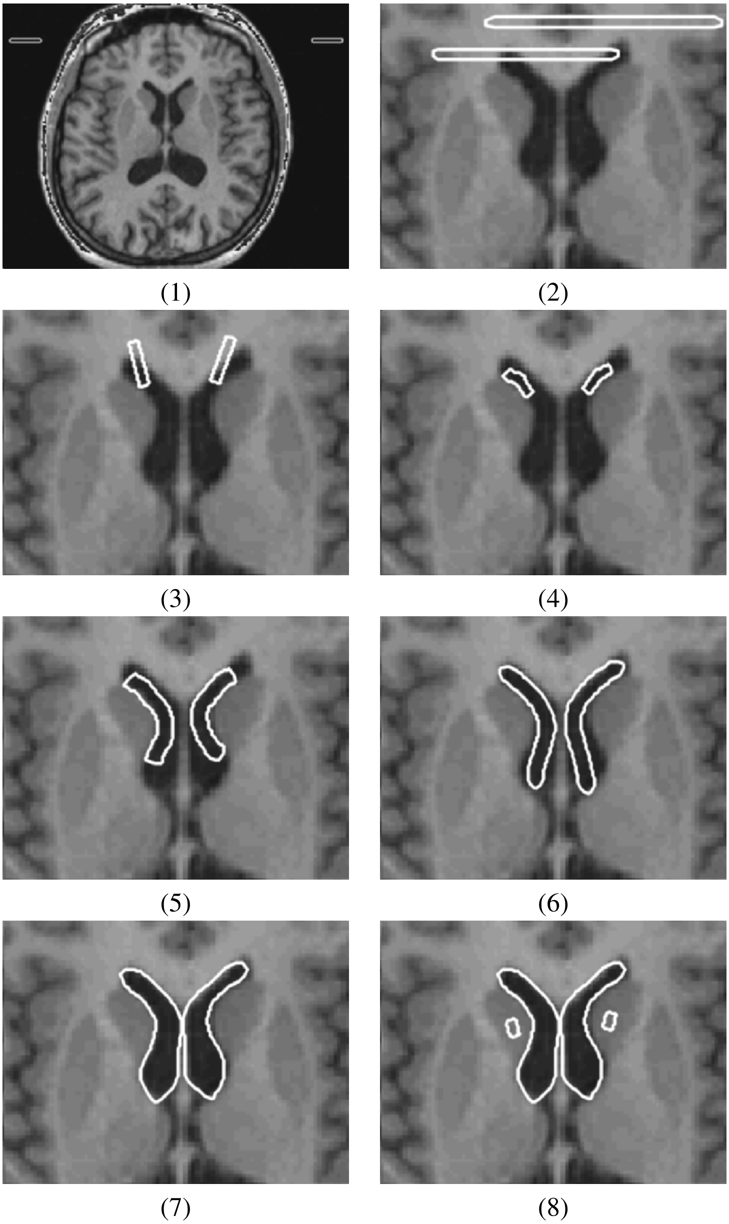

Figure 40. Deformable lateral ventricles (1–16), caudate nuclei (CN) (8–16), and putamina (11-16) organisms progressing through a sequence of behaviors to locate the corresponding structures in an MR brain image. Reprinted with permission from [85]. Copyright ©2002, Elsevier. (results continued on next page)

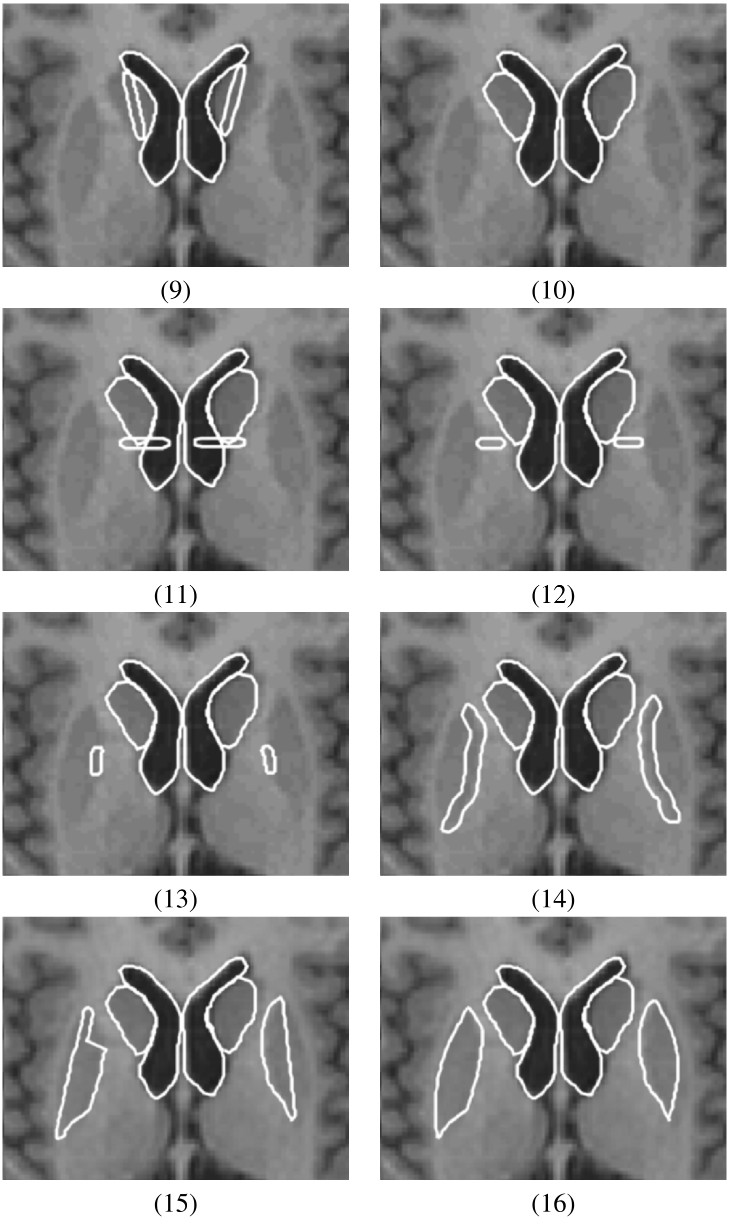

Figure 40. (continued). Results continued from previous page.

The CN organism segments the CN by stretching to locate its upper and lower limits (Figure 40.9) and thickening to latch onto its inner and outer boundaries (Figure 40.10). The CN organism passes information about the location of its lowest point (in the image) to the putamen organism, which is initialized accordingly (Figure 40.11). The putamen organism moves toward the putamen in the brain image (Figure 40.12) and then rotates and bends to latch onto the nearer putamen boundary (Figure 40.13). It then stretches and grows along the boundary until reaching the upper- and lower-most ends of the putamen (Figure 40.14), which identifies the medial axis of the putamen (Figure 40.15). Since the edges of the putamen boundary near the gray matter are usually weak, the organism activates an explicit search for an arc (parametrized only by one parameter controlling its curvature) that best fits the weak, sparse edge data in that region (Figure 40.16).

4.3. 2D Vessel Crawler

We also present the result of segmenting vessels in an angiogram (Figure 41) using a 2D artery crawler. Without proper constraints the vessel organism latches onto the wrong overlapping vessel (Figure 41a). However, adding additional sensors and high-level constraints enables the organism to distinguish between overlapping vessels and bifurcations (Figure 41b). When the latter is encountered, two new organisms are born from the original main branch organism, one for each branch (Figure 41c). Figure 32 demonstrates how this (overlap vs. bifurcate) decision strategy is implemented.

4.4. Physically Based Deformable CC Organism

In this section we exemplify the use of physics-based shape deformations (Section 3.3.3) within the deformable organisms framework yielding additional robustness by allowing intuitive real-time user guidance and interaction when necessary. The organism's behaviors and awareness of the different stages of the plan (Figure 1) enables it to not only segment the CC but also label anatomical regions (Figure 43a). Here we demonstrate the physics-based deformable organisms, with an underlying dynamic spring-mass mesh model, through the fully automatic segmentation and labeling of the CC in 2D midsagittal MRI slices. We also present further improvement of the segmentation results through minor, intuitive user interaction (Figure 43).

4.5. Caudate Nucleus and Lateral Ventricle Shape Modeling Using Medial Patches

We show two examples of medial patch-based deformations for fitting a 3D shape model to target brain structures. In the first example, we demonstrate the ability of performing intuitive and controlled manual 3D shape deformations to fit the medial patch representation to a lateral ventricle in a head MRI (Figure 44). We

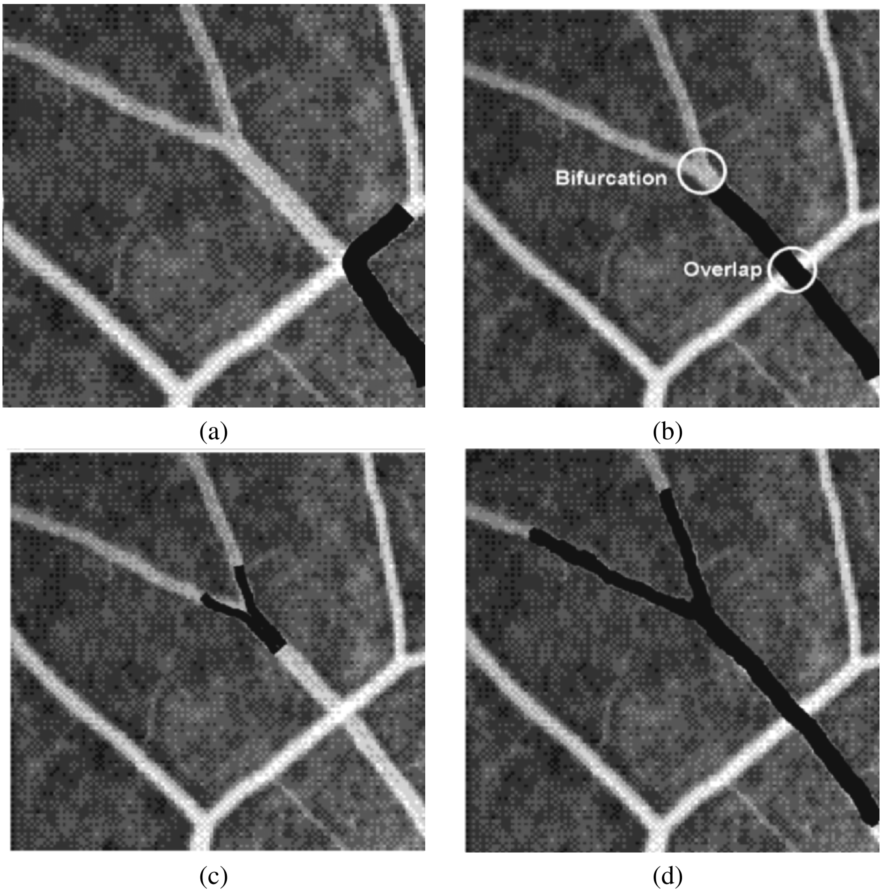

Figure 41. Segmenting vessels in an angiogram. (a) A deformable organism turning right and latching onto the wrong overlapping vessel. (b) High-level constraints enable the organism to differentiate between overlapping vessels and bifurcations. (c) Two new organisms are born upon identifying a bifurcation. (d) The segmented main vessel and its two branches. Reprinted with permission from [85]. Copyright ©2002, Elsevier.

then show an example of automatically fitting the 3D medial patch shape model to the binary image of a caudate nucleus from a 3D brain MRI (Figure 45). An initial medial sheet estimate is positioned at the plane spanned by the two main principal components of the locations of object points in the binary image. For each point in the medial sheet, rays are cast in both directions perpendicular to the medial sheet (along the third eigenvector) until they encounter an object boundary. The two boundary locations above and below the medial sheet are recorded. The nodes of the medial sheet are repositioned to be halfway between the top and

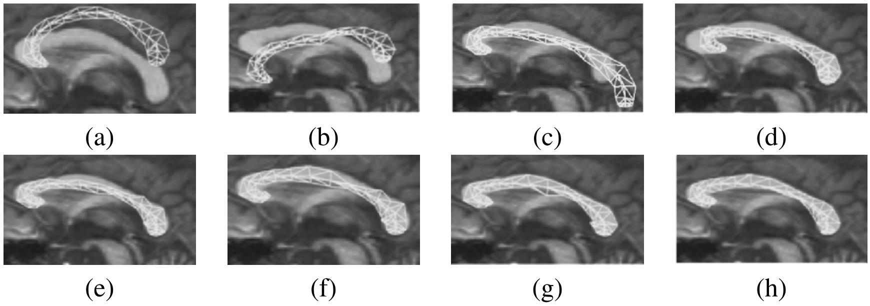

Figure 42. Progress of segmentation through its primary phases: (a) global model alignment, (b) model part alignment through (c) expansion and (d) contraction, (e) medial-axis alignment, (f) fitting to boundary, (g) detecting and (h) repairing fornix dip. Reprinted with permission from [80]. Copyright ©2005, SPIE. See attached CD for color version.

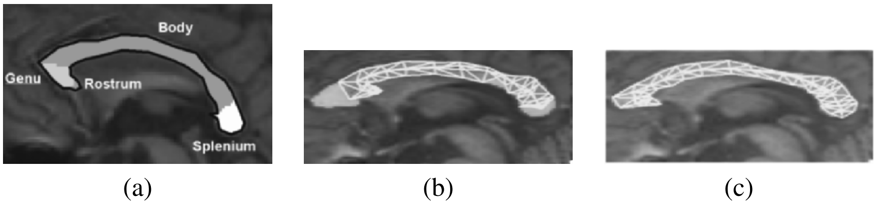

Figure 43. (a) Automatic labeling of important anatomical regions of the CC. (b) Before and (c) after intuitive manual intervention to improve the segmentation (red color corresponds to areas of erroneous segmentation). Reprinted with permission from [80]. Copyright ©2005, SPIE. See attached CD for color version.

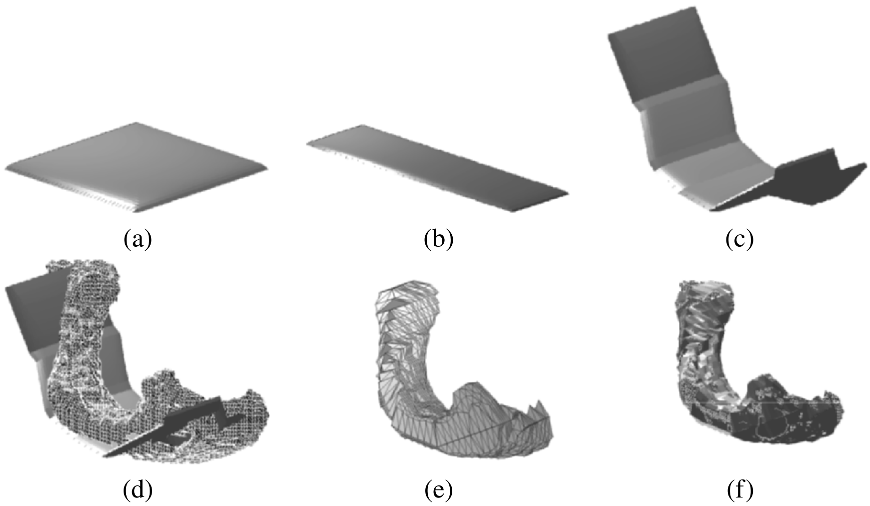

Figure 44. Fitting a 3D shape model to a lateral ventricle: (a) initial 3D model, (b) uniform shrinking along the x-axis, (c) bending deformations, (d) medial sheet bent into approximate bisecting position of target structure, (e) final 3D shape reconstructed from the medial sheets, and (f) overlay of 3D model on target structure.

bottom boundaries. Rays are now cast along vectors normal to the medial sheet until they encounter the boundaries. The procedure is iterated until the change in the locations of the medial nodes is too small.

4.6. 3D Vessel Crawler

In modern internal medicine, noninvasive imaging procedures are often crucial to the diagnosis of cardiovascular, pulmonary, renal, aortic, neurological, and abdominal diseases [82]. Common amongst the diagnosis of these diseases is the use of volumetric angiography to highlight branching vasculature. We present segmentation and analysis results of vasculature from volumetric magnetic resonance angiographic (MRA) data using the 3D vessel crawler. We ran the vessel crawler on an MRA image (Figure 46). The vessel crawler was initialized using a single seed point for each of the three root-most vessels. As shown, it was able to detect and track the vast majority of connected vessel segments. The crawler grows along vessels, latches onto their boundaries, detects bifurcations, spawns new child organisms, etc. All of these actions are described under the behavioral layer of the organism, and they rely upon both the physical and geometrical layers for implementation. For example, for crawling along the vessel the cognitive center gathers sensory input using the "vesselness" and "Hessian" sensory modules (Section 3.4), and elicits the acts of "growing" and then "fitting" to conform to the vascular walls (Section 3.5.2). In turn, each of these methods relies upon the physical and geometrical layers to carry out tasks, such as maintaining model stability through the application of stabilization springs.

In addition to their robust segmentation, deformable organisms are able to perform intuitive analysis, and labeling of the target structure. The 3D vessel crawler extracts clinically relevant features as it crawls though the vasculature, including the distance metric (DM), sum of angles metric (SOAM), and inflection count metric (ICM) [82]. In addition to those features the vessel crawlers are able to label branch points and vessel segments, determine branching angles, cross-sectional radius, vessel segment length and vessel volume, as well as highlight potential problem areas, the vascular regions they affect, and the shortest path to those locations from any target point (Figure 47).

5. SUMMARY

In this chapter we presented deformable organisms — an artificial life framework for medical image analysis. Deformable organisms extend the classical bottom–up, data-driven, deformable models to include additional top–down, knowledge-driven capabilities. This extension is realized by complementing the geometrical and physical layers of deformable models with behavioral and cognitive layers, as well as sensory modules.

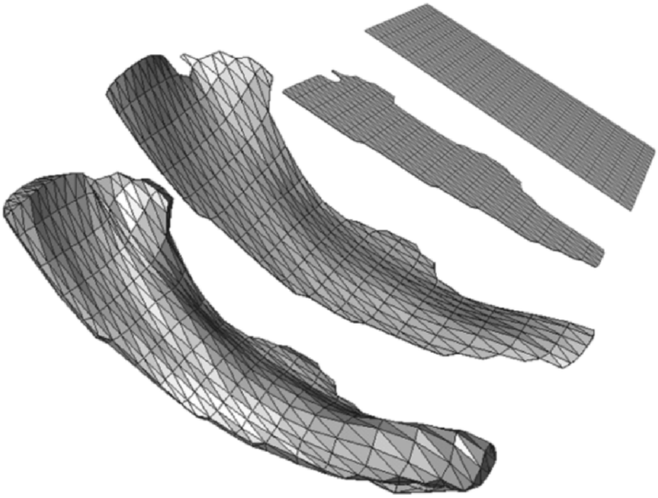

Figure 45. Caudate nucleus (CN) represented using a medial patch. Top to bottom: initial rectangular planar medial sheet, planar sheet cropped to match CN projection, curved medial sheet placed equidistant from the upper and lower CN boundaries. Thickness values are associated with each node in the medial sheet yielding a 3D surface of the CN.

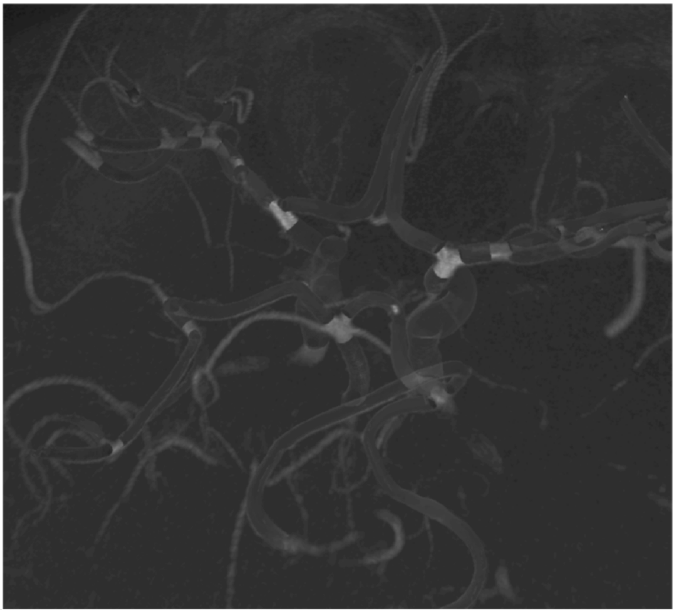

Figure 46. Maximum intensity projection rendering of an MRA showing the vessel crawler in orange. See attached CD for color version.

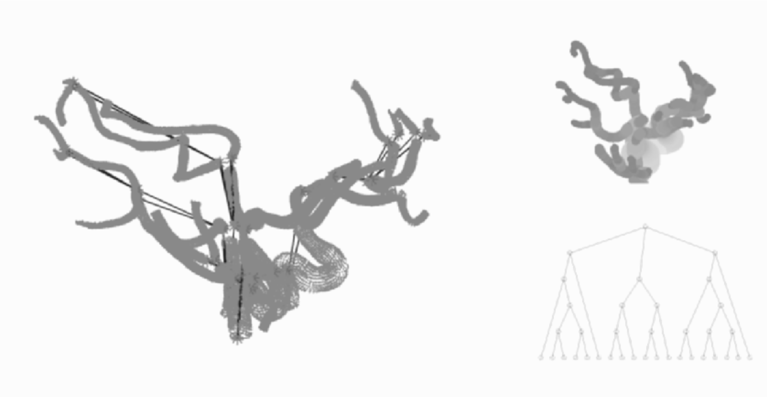

Figure 47. Three example representations of a vascular system: directed acyclic graph (left), a plot of the vessels with color corresponding to radial thickness (top), and a tree representing to the structure of a segmented vessel (bottom). See attached CD for color version.

The cognitive layer is the organism's brain, which is mainly responsible for decision-making based on sensory input, prior anatomical knowledge, a pre-stored segmentation plan, and interaction with other organisms. The cognitive layer controls the organism's perceptual capabilities by dynamically assigning on-board and off-board sensors. It also controls the sequence of behaviors that the organisms should adopt. The different behaviors in turn activate bodily deformations carried out through the lower geometrical and physical layers of the organism.

The need for top–down control of the shape deformations implies a requirement to develop shape representation techniques that respond to intuitive and controlled deformation commands (e.g., "move forward," "bend left"). To this end, we presented examples of pure geometrical and physically based shape representations that provide the desired deformation control capabilities. Specifically, we described the use of medial profiles/patches and 2D/3D deformable spring mass mesh models as examples of the geometric and physical layers of the deformable organism framework.

We presented examples of deformable organisms designed for different medical image analysis problems (e.g., segmentation of brain structures, analysis of vasculature) and highlighted their main behavioral routines, decision-making strategies, and sensory capabilities.

We believe the layered architecture of artificial life is a promising paradigm for medical image analysis, capable of incorporating state-of-the-art low-level image processing algorithms, high-level anatomical expert knowledge, and advanced planning and optimization strategies in a single modular design.

6. NOTES

1. See [28] as an example of previous work on, and motivation for, segmenting the corpus callosum.
2. Earlier mention of the use of AL in conjunction with segmentation appeared in [35, 36] However, as these methods rely on local, not global, decision-making, they strongly resemble traditional region-growing methods. For a survey of other applications of AL, see [37].

7. REFERENCES

1. Yoo TS. 2004. *Insight into images: principles and practice for segmentation*. Wellesley, MA: AK Peters Ltd.
2. Robb RA. 2000. *Biomedical imaging*. New York: Wiley-Liss.
3. Dhawan A. 2003. *Medical image analysis*. Wiley-IEEEE Press.
4. Bankman I. 2000. *Handbook of medical imaging: processing and analysis*. New York: Academic Press.
5. Sonka M, Fitzpatrick J. 2000. *Handbook of medical imaging*. Bellingham, WA: SPIE.
6. McInerney T, Kikinis R. 1998. An object-based volumetric deformable atlas for the improved localization of neuroanatomy in MR images. In *Proceedings of the first international conference on medical image computing and computer-assisted intervention (MICCAI'98)*. Lecture Notes in Computer Science, Vol. 1496, pp. 861–869. Berlin: Springer.
7. Shen D, Davatzikos C. 2000. An adaptive-focus deformable model using statistical and geometric information. *IEEE Trans Pattern Anal Machine Intell* **22**(8):906–913.
8. Tsotsos J, Mylopoulos J, Covvey H, Zucker S. 1980. A framework for visual motion understanding. *IEEE Trans Pattern Anal Machine Intell* **2**(6):563–573.
9. Draper BA, Hanson AR, Riseman EM. 1993. Learning blackboard-based scheduling algorithms for computer vision. *Int J Pattern Recognit Artif Intell* **7**(2):309–328.
10. Draper BA, Hanson AR, Riseman EM. 1996. Knowledge-directed vision: control, learning, and integration. *Proc IEEE Signals Symbols* **84**(11):1625–1637.
11. Crevier D, Lepage R. 1997. Knowledge-based image understanding systems: a survey. *Comput Vision Image Understand* **67**(2):161–185.
12. Strat T, Fischler M. 1991. Context-based vision: recognizing objects using information from both 2D and 3D imagery. *IEEE Trans Pattern Anal Machine Intell* **13**(10):1050–1065.
13. Poli R. 1996. Genetic programming for image analysis. In *Proceedings of the first international conference on genetic programming*, pp. 363–368. Ed JR Koza, DE Goldberg, DB Fogel, RL. Cambridge: MIT Press.
14. Martin MC. 2002. Genetic programming for real world robot vision. In *Proceedings of the international conference on intelligent robots and systems (IEEE/RSJ)*, pp. 67–72. Washington, DC: IEEE. Available at http://martincmartin.com/papers/GeneticProgramming ForRealWorldRobotVisionIROS2002Martin.pdf
15. Liu J, Tang Y. 1999. Adaptive image segmentation with distributed behavior-based agents. *IEEE Trans Pattern Anal Machine Intell* **21**(6):544–550.
16. Germond L, Dojat M, Taylor C, Garbay C. 2000. A cooperative framework for segmentation of MRI brain scans. *Artif Intell Med* **20**(1):77–93.
17. Boucher A, Doisy A, Ronot X, Garbay C. 1998. A society of goal-oriented agents for the analysis of living cells. *Artif Intell Med* **14**:183–199.
18. Rodin V, Harrouet F, Ballet P, Tisseau J. 1998. oRis: multiagents approach for image processing. In *Proceedings of the SPIE conference on parallel and distributed methods for image processing II*, Vol. 3452, 57–68. Ed S Hongchi, PC Coffield. Bellingham, WA: SPIE.

19. Bazen AM, van Otterlo M. 2001. A reinforcement learning agent for minutiae extraction from fingerprints. In *Proceedings of the Belgium–Netherlands Artificial Intelligence Conference (BNAIC'01)*, pp. 329–336. Ed B. Kröose, M de Rijke, G Schreiber, M van Someren. Washington, DC: IEEE Computer Society.
20. McInerney T, Terzopoulos D. 1996. Deformable models in medical image analysis: a survey. *Med Image Anal* **1**(2):91–108.
21. Montagnat J, Delingette H, Ayache N. 2001. A review of deformable surfaces: topology. *Image Vision Comput* **19**(14):1023–1040.
22. Osher S, Paragios N. 2003. *Geometric level set methods in imaging vision and graphics*. Berlin: Springer.
23. Caselles V, Kimmel R, Sapiro G. 1997. Geodesic active contours. *Int J Comput Vision* **22**(1):61–79.
24. Sethian JA. 1996. *Level set methods: evolving interfaces in geometry, computer vision and material sciences*. Cambridge: Cambridge UP.
25. Cootes TF, Cooper D, Taylor CJ, Graham J. 1995. Active shape models: their training and application. *Comput Vision Image Understand* **61**(1):38–59.
26. Cootes T, Beeston C, Edwards G, Taylor C. 1999. A unified framework for atlas matching using active appearance models. In *Proceedings of the 17th international conference on image processing in medical imaging (IPMI'99)*, pp. 322–333. Berlin: Springer.
27. Szekely G, Kelemen A, Brechbuehler Ch, Gerig G. 1996. Segmentation of 3D objects from MRI volume data using constrained elastic deformations of flexible Fourier surface models. *Med Image Anal* **1**(1):19–34.
28. Lundervold A, Duta N, Taxt T, Jain A. 1999. Model-guided segmentation of corpus callosum in MR images. In *Proceedings of the IEEE computer society conference on computer vision and pattern recognition*, Vol. 1, pp. 231–237. Washington, DC: IEEE Computer Society.
29. Cohen LD. 1991. On active contour models and balloons. *Comput Vision Graphics Image Process: Image understand* **53**(2):211–218.
30. Xu C, Prince J. 1998. Snakes, shapes, and gradient vector flow. *IEEE Trans Image Process* **7**(3):359–369.
31. Cootes TF, Edwards GJ, Taylor CJ. 2001. Active appearance models. *IEEE Trans Pattern Anal Machine Intell* **23**(1):681–685.
32. Leventon M, Grimson W, Faugeras O. 2000. Statistical shape influence in geodesic active contours. In *Proceedings of the IEEE computer society conference on computer vision and pattern recognition*, Vol. 1. pp. 316–323. Washington, DC: IEEE Computer Society.
33. Warfield SK, Kaus M, Jolesz FA, Kikinis R. 2000. Adaptive, template moderated, spatially varying statistical classification. *Med Image Anal* **4**(1):43–55.
34. Hamarneh G, Gustavsson T. 2000. Statistically constrained snake deformations. In *Proceedings of the international conference on systems, man, and cybernetics*, Vol. 3, pp. 1610–1615. Washington, DC: IEEE Computer Society.
35. Choi C, Wirth M, Jennings A. 1997. Segmentation: artificial life and watershed transform approach. In *Sixth international conference on image processing and its applications*, Vol. 1, pp. 371–375. Washington, DC: IEEE Computer Society.
36. Kagawa H, Kinouchi M, Hagiwara M. 1999. Image segmentation by artificial life approach using autonomous agents. In *Proceedings of the international joint conference on neural networks*, Vol. 6, pp. 4413–4418. Washington, DC: IEEE Computer Society.
37. Kim K-J, Cho S-B. 2006. A comprehensive overview of the applications of artificial life. *Artif Life* **12**(1):153–182.
38. Hamarneh G, McInerney T, Terzopoulos D. 2001. Deformable organisms for automatic medical image analysis. In *Proceedings of the fourth international conference on medical image computing and computer-assisted intervention (MICCAI'01)*. Lecture Notes in Computer Science, Vol. 2208, pp. 66–75. Berlin: Springer.
39. Terzopoulos D. 1999. Artificial life for computer graphics. *Commun ACM* **42**(8):32–42

40. Bookstein F. 1997. *Morphometric tools for landmark data: geometry and biology*. Cambridge: Cambridge UP.
41. Costa L, Cesar Jr R. 2000. *Shape analysis and classification: theory and practice*. Boca Raton, FL: CRC Press.
42. Dryden I, Mardia K. 1998. *Statistical shape analysis*. New York: John Wiley and Sons.
43. Lachaud J, Montanvert A. 1999. Deformable meshes with automated topology changes for coarse-to-fine three-dimensional surface extraction. *Med Image Anal* **3**(1):1–21.
44. Mandal C, Vemuri B, Qin H. 1998. A new dynamic FEM-based subdivision surface model for shape recovery and tracking in medical images. In *Proceedings of the first international conference on medical image computing and computer-assisted intervention (MICCAI'98)*. Lecture Notes in Computer Science, Vol. 1496, pp. 753–760. Berlin: Springer.
45. Miller J, Breen D, Lorensen W, O'Bara R, Wozny MJ. 1991. Geometrically deformed models: a method for extracting closed geometric models from volume data. In *Proceedings of the 18th annual conference on computer graphics and interactive techniques (SIGGRAPH'91)*, Vol. 25, pp. 217–226. New York: ACM Press.
46. Montagnat J, Delingette H. 1997. Volumetric medical image segmentation using shape constrained deformable models. In *Proceedings of the second international conference on computer vision, virtual reality and robotics in medicine (CVRMed-MRCAS'97)*. Lecture Notes in Computer Science, Vol. 1205, pp. 13–22. Berlin: Springer.
47. Barr A. 1984. Global and local deformations of solid primitives. In *Proceedings of the 11th annual conference on computer graphics and interactive techniques (SIGGRAPH'84)*, Vol. 18, pp. 21–30. New York: ACM Press.
48. Coquillart S. 1990. Extended free form deformations: a sculpting tool for 3d geometric modeling. In *Proceedings of the 17th annual conference on computer graphics and interactive techniques (SIGGRAPH'90)*, Vol. 24, pp. 187–196. New York: ACM Press.
49. Sederberg TW, Parry SR. 1986. Free-form deformation of solid geometric models. In *Proceedings of the 13th annual conference on computer graphics and interactive techniques (SIGGRAPH'86)*, Vol. 4, pp. 151–160. New York: ACM Press.
50. Singh K, Fiume E. 1998. Wires: a geometric deformation technique. In *Proceedings of the 25th annual conference on computer graphics and interactive techniques (SIGGRAPH'98)*, Vol. 99, No. 1, pp. 405–414. New York: ACM Press.
51. Terzopoulos D, Metaxas D. 1991. Dynamic 3D models with local and global deformations: deformable superquadrics. *IEEE Trans Pattern Anal Machine Intell* **13**(7):703–714.
52. Burel G, Henocq H. 1995. Three-dimensional invariants and their application to object recognition. *Signal Process* **45**(1):1–22.
53. Davatzikos C, Tao X, Shen D. 2003. Hierarchical active shape models: using the wavelet transform. *IEEE Trans Med Imaging* **22**(3)414–423.
54. McInerney T, Terzopoulos D. 1995. A dynamic finite-element surface model for segmentation and tracking in multidimensional medical images with application to cardiac 4D image analysis. *Comput Med Imaging Graphics* **19**(1):69–83.
55. Mortenson M. 1997. *Geometric modeling*. New York: Wiley.
56. Blum H. 1973. Biological shape and visual science. *Theor Biol* **38**:205–287.
57. Attali D, Montanvert A. 1997. Computing and simplifying 2D and 3D continuous skeletons. *Comput Vision Image Understand* **67**(3):261–273.
58. Borgefors G, Nystrom I, Baja GSD. 1999. Computing skeletons in three dimensions. *Pattern Recognit* **32**:1225–1236.
59. Bouix S, Dimitrov P, Phillips C, Siddiqi K. 2000. Physics-based skeletons. In *Proceedings of the vision interface conference*, pp. 23–30.
60. Dimitrov P, Damon J, Siddiqi K. 2003. Flux invariants for shape. In *Proceedings of the IEEE computer society conference on computer vision and pattern recognition (CVPR'03)*, pp. 835–841. Washington, DC: IEEE Computer Society.

61. Fritsch D, Pizer S, Yu L, Johnson V, Chaney E. 1997. Segmentation of medical image objects using deformable shape loci. In *Proceedings of the 15th international conference on information processing in medical imaging (IPMI'97). Lecture notes in computer science*, Vol. 1230, pp. 127–140. Berlin: Springer.
62. Hamarneh G, McInerney T. 2003. Physics-based shape deformations for medical image analysis. In *Proceedings of the 2nd SPIE conference on image processing: algorithms and systems (SPIE-IST'03)*, Vol. 5014, pp. 354–362. Ed ER Dougherty, JT Astola, KO Egiazarian.
63. Hamarneh G, Abu-Gharbieh R, McInerney T. 2004. Medial profiles for modeling deformation and statistical analysis of shape and their use in medical image segmentation. *Int J Shape Model* **10**(2)187–209.
64. Leymarie F, Levine MD. 1992. Simulating the grassfire transform using an active contour model. *IEEE Trans Pattern Anal Machine Intell* **14**(1):56–75.
65. Pizer S, Fritsch D. 1999. Segmentation, registration, and measurement of shape variation via image object shape. *IEEE Trans Med Imaging* **18**(10):851–865.
66. Pizer S, Gerig G, Joshi S, Aylward SR. 2003. Multiscale medial shape-based analysis of image objects. *Proc IEEE* **91**(10):1670–1679.
67. Sebastian T, Klein P, Kimia B. 2001. Recognition of shapes by editing shock graphs. In *Proceedings of the international conference on computer vision (ICCV'01)*, pp. 755–762. Washington, DC: IEEE Computer Society.
68. Siddiqi K, Bouix S, Tannenbaum A, Zucker S. 2002. Hamilton-Jacobi skeletons. *Int J Comput Vision* **48**(3):215–231.
69. Duta N, Sonka M, Jain A. 1999. Learning shape models from examples using automatic shape clustering and Procrustean analysis. In *Proceedings of the 17th international conference on information processing in medical imaging (IPMI'99)*. Lecture Notes in Computer Science, Vol. 1613, pp. 370–375.
70. Yushkevich P, Joshi S, Pizer SM, Csernansky J, Wang L. 2003. Feature selection for shape-based classification of biological objects. In *Proceedings of the 17th international conference on information processing in medical imaging (IPMI'03)*. Lecture Notes in Computer Science, Vol. 2732, pp. 114–125.
71. Styner M, Gerig G, Lieberman J, Jones D, Weinberger D. 2003. Statistical shape analysis of neuroanatomical structures based on medial models. *Med Image Anal* **7**(3):207–220.
72. Lu C, Pizer S, Joshi S. 2003. A Markov random field approach to multiscale shape analysis. In *Proceedings of the conference on scale space methods in computer vision.* Lecture Notes in Computer Science, Vol. 2695, pp. 416–431. Available at http://midag.cs.unc.edu/pubs/papers/ScaleSpace03_Lu_shape.pdf.
73. Grenander U. 1963. *Probabilities on algebraic structures*. New York: Wiley.
74. Fletcher P, Conglin L, Pizer S, Joshi S. 2004. Principal geodesic analysis for the study of nonlinear statistics of shape. *IEEE Trans Med Imaging* **23**(8):995–1005.
75. Bill JR, Lodha SK. 1995. Sculpting polygonal models using virtual tools. In *Proceedings of the conference on the graphics interface*, pp. 272–279. New York: ACM Press.
76. O'Donnell T, Boult T, Fang X, Gupta A. 1994. The extruded generalized cylinder: a deformable model for object recovery. In *Proceedings of the IEEE computer society conference on computer vision and pattern recognition (CVPR'94)*, pp. 174–181. Washington, DC: IEEE Computer Society.
77. Perona P, Malik J. 1990. Scale-space and edge detection using anisotropic diffusion. *IEEE Trans Pattern Anal Machine Intell* **12**(7)629–639.
78. Weickert J. 1998. *Anisotropic diffusion in image processing*. ECMI Series. Stuttgart: Teubner.
79. Frangi AF, Niessen WJ, Vincken KL, Viergever MA. 1998. Multiscale vessel enhancement filtering. In *Proceedings of the first international conference on medical image computing and computer-assisted intervention (MICCAI'98)*. Lecture Notes in Computer Science, Vol. 1496, pp. 130–137. Berlin: Springer.

80. Hamarneh G, McIntosh C. 2005. Physics-based deformable organisms for medical image analysis. In *Proceedings of the SPIE conference on medical imaging: image processing*, Vol. 5747, pp. 326–335. Bellingham, WA: SPIE.
81. Kimme C, Ballard DH, Sklansky J. 1975. Finding circles by an array of accumulators. *Commun Assoc Comput Machinery* **18**:120–122.
82. Bullitt E, Gerig G, Aylward SR, Joshi SC, Smith K, Ewend M, Lin W. 2003. Vascular attributes and malignant brain tumors. In *Proceedings of the sixth international conference on medical image computing and computer-assisted intervention (MICCAI'03)*. Lecture Notes in Computer Science, Vol. 2878, pp. 671–679. Berlin: Springer.
83. Terzopoulos D. 1986. *On matching deformable models to images*. Technical Report 60, Schlumberger Palo Alto Research. Reprinted in *Topical Meeting MachineVision, Technical Digest Series* **12**:160–167 (1987).
84. Kass M, Witkin A, Terzopoulos D. 1987. Snakes: active contour models. *Int J Comput Vision* **1**(4):321–331.
85. McInerney T, Hamarneh G, Shenton M, Terzopoulos D. Deformable organisms for automatic medical image analysis. *J Med Image Anal* **6**(3):251–266.
86. C. McIntosh and G. Hamarneh, Vessel Crawlers: 3D Physically-based Deformable Organisms for Segmentation and Analysis of Tubular Structures in Medical Images. *IEEE Conference on Computer Vision and Pattern Recognition*, 2006, pp. 1084–1091.