Intelligent Deformable Organisms: An Artificial Life Approach to Medical Image Analysis

Ghassan Hamarneh[1,2], Tim McInerney[3,2], and Demetri Terzopoulos[4,2]

[1] Department of Signals and Systems, Chalmers University of Technology, Göteborg, SE-41296, Sweden
ghassan@s2.chalmers.se, http://www.s2.chalmers.se
[2] Department of Computer Science, University of Toronto, Toronto, M5S 3H5, Canada
{ghassan, dt, tim}@cs.toronto.edu, http://www.cs.toronto.edu
[3] School of Computer Science, Ryerson University, Toronto, M5B 2K3, Canada
tmcinern@scs.ryerson.ca, http://www.scs.ryerson.ca
[4] Media Research Laboratory, New York University, New York, NY 10003, USA
dt@cs.nyu.edu, http://www.mrl.nyu.edu

Abstract. We propose and demonstrate a new paradigm for medical image analysis that borrows concepts from the field of artificial life. The paradigm prescribes intelligent agents in the form of artificial organisms whose primary objective is the proper segmentation and analysis of anatomical structures in medical images. An organism's 'awareness' of the segmentation process guides it to identify landmark anatomical features during its development. The organism possesses a 'brain' enabling it to make decisions and issue 'muscle' actuation commands, resulting in shape deformations of its body, comprising skeleton and limbs. Behavioral decisions are based upon perceived data, memorized information, and a pre-stored 'cognitive' plan. The underlying medial-based representation of the organism's bodily shape facilitates a variety of controlled deformations at multiple scales and locations. Our framework promises to lay the foundation for the construction of robust and automatic medical image analysis tools by combining deformable models and high-level a priori knowledge. We show how an intelligent corpus callosum agent, which takes the form of a worm, can deal with noise, incomplete edges, enormous anatomical variation, and occlusion in order to segment and label the corpus callosum in 2D mid-sagittal MRI slices of the brain.

1 Introduction

A key problem that impacts a host of medical data-driven applications and continues to defy solution is the automatic segmentation and labeling of anatomical structures from medical images. It is generally recognized that any solution to this problem requires the incorporation and effective utilization of context-based, constraining information within a robust decision-making framework [3]. Furthermore, the decision-making framework should operate at an appropriate level of abstraction. Currently, no single algorithm meets this requirement and consequently no algorithm can robustly segment a variety of structures in medical images over a range of data sets.

One of the most researched approaches in recent years - deformable model-based strategies [6] - have shown early promise, but they continue to be plagued by sensitivity to the huge variation in anatomy, the significant variability inherent to image data, and their own need for intelligent initialization conditions. These approaches optimize objective functions of one form or another. This move into the world of mathematical optimization and away from the expert system-like approaches of the late 1980's [3] has yielded a powerful variety of image analysis algorithms, but it has shifted the focus from what may be the key element in any viable highly automated solution: the decision-making framework itself. A basic underlying problem of optimization-based decision-making strategies is the lack of high-level control or guidance over the segmentation process. Intelligent, global (i.e. over the whole image) behavior of deformable models cannot easily emerge from the fundamentally local decisions based on energy-minimizing model deformations. Hierarchically organized or structure-oriented models, which shift their focus from structures with stable image features to structures with less stable features, are a step in the right direction [5,10]. However, high-level contextual knowledge remains underutilized because it is still too intertwined with the low-level optimization-based framework. Fundamentally, current deformable models have no explicit awareness of where they (or their parts) are or what they are looking for during the optimization process.

It is our contention that we must revisit ideas explored in earlier systems, and develop approaches that focus on higher-level reasoning strategies which may best leverage the feature detection and integration abilities of deformable models and other model-based techniques. We further contend that, within a layered architecture approach, the high-level reasoning layer must have knowledge about and control over the low-level model (or models) at all times. It should adhere to 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 knowledge to resolve regions where there is a deficiency of image feature information.

To achieve these goals, we propose a new paradigm for medical image analysis founded upon concepts from Artificial Life [12,13]. We create intelligent agents in the form of artificial organisms whose principal intention is the proper segmentation and analysis of anatomical objects. The organism's awareness of the segmentation process guides it to identify landmark anatomical features during its development. An organism possesses a brain enabling it to make high-level decisions and issue 'muscle' commands, resulting in shape deformations of its skeleton and limbs. Decisions are based upon the integration of perceived data, memorized information, and a pre-stored plan. We use an underlying medial-based representation of the organism's shape to facilitate a variety of controlled deformations at multiple locations and scales.

We believe that our novel framework lays the foundation for the construction of robust, automatic tools for medical image segmentation, object-based registration, and shape variation measurement, by harnessing the true potential of deformable models. We demonstrate our approach by constructing an intelligent "worm" to segment and label the corpus callosum in 2D mid-sagittal MRI slices of the brain.

1.1. ALife for Computer Graphics

Over the past several years, researchers have been exploring the modeling and simulation of living systems for computer graphics by applying concepts from an emerging field of scientific inquiry called Artificial Life, or ALife [12,13]. These new graphics models not only incorporate geometric and physics-based modeling techniques, but also attempt to simulate many of the natural processes that uniquely characterize living systems - including birth and death, growth and development, natural selection, evolution, perception, locomotion, manipulation, adaptive behavior, learning, and intelligence.

ALife models for computer graphics are often organized hierarchically as a layered pyramid (Fig. 1b). At the base of the pyramid, the geometric modeling layer is used to represent the shape of the artificial life form. Climbing the pyramid, the physical modeling layer constrains the geometry by incorporating physical principles to simulate the dynamics of biological tissue and internal muscle actuators of animals to generate lifelike animation. Climbing still further, the behavioral modeling layer attempts to simulate the behaviors of animals such as locomotion, collision avoidance, foraging, etc., through sensory perception and active control of the physical modeling layer. At the apex of the pyramid, cognitive modeling has emerged as the use of artificial intelligence techniques, including knowledge representation, reasoning, and planning, to produce graphical characters with some level of deliberate intelligence. Cognitive modeling goes beyond behavioral modeling in that it governs what a character knows, how that knowledge is acquired, and how it can be used to direct autonomous characters to perform specific tasks. ALife modeling also investigates the possibility of applying evolutionary models to evolve graphical characters or aspects of their bodies and brains.

Artificial life for computer graphics has spawned several principal avenues of research and development, including artificial plants, animals, evolution, and behavioral models and animation.

The key ingredients of artificial animals for example, are the functional models of animal bodies and brains. The artificial fish [12], for example, is an autonomous agent with a deformable body actuated by internal muscles. The body of the fish harbors a brain with motor, perception, behavior, and learning centers. The motor system, comprising the actuators and a set of motor controllers (MCs), drives the dynamic model of the fish. The MCs are parameterized procedures dedicated to carrying out a specific motor function, such as 'swim-forward' or 'turn-left'. They translate natural control parameters such as forward speed or angle of the turn into detailed muscle actions.

The perception system relies on a set of virtual sensors to provide sensory information about the dynamic environment. The brain's perception center includes a perceptual attention mechanism that allows the artificial fish to train its sensors at the world in a task-specific way, hence filtering out sensory information superfluous to its current behavioral needs. The behavior center of the artificial fish's mind mediates between its perception system and its motor system. An intention generator (or action-selection mechanism), the fish's cognitive faculty, combines the fish's innate characteristics, the mental state, and the incoming stream of sensory information at every simulation time step. It uses this information to activate behavior routines. The behavior routines in turn compute the appropriate motor controller parameters to carry the fish one step closer to fulfilling its current intention.

The brains of artificial fishes are also able to learn. For example, the fish can learn to locomote through practice and sensory reinforcement. The learning center of the brain comprises a set of optimization-based motor learning algorithms that can discover and perfect motor controllers capable of producing efficient locomotion.

1.2. An ALife Modeling Paradigm for MIA: Motivation

Current deformable model-based approaches to medical image analysis utilize geometric and physical modeling layers only. The models are fitted by simulating dynamics or minimizing energy terms but do not monitor or control (except in a primitive sense) this optimization process. Sophisticated deformable models incorporate prior information to constrain shape and image appearance and statistical variation of these quantities [2,1,11]. However, these models have no explicit awareness of where they are and therefore the effectiveness of these constraints is dependent upon model starting conditions. This lack of awareness also prevents the models from knowing when to trust the image feature information and ignore the constraint information and vice versa. The constraint information is arbitrarily applied. Furthermore, because there is no active, explicit search for stable image features, the models are prone to latching onto incorrect features [1] simply due to their proximity and local decision-making. Once this latching occurs, the lack of control of the fitting procedure prevents the model from correcting this misstep. The result is that the local decisions that are made do not add up to intelligent global behavior. Finally, most deformable models do not have intuitive, multi-scale, multi-location deformation 'handles'. They are often boundary-based with no global shape descriptors making them inadequate for higher-level guidance. They are unable to perform high-level deformations, such as bending, sliding, or backing up, and therefore it becomes extremely difficult to develop reasoning or planning strategies for these models at the correct level of abstraction [6]. For example, when segmenting the corpus callosum (CC) in 2D mid-sagittal images, the 'vocabulary' that one uses should contain words that describe principal anatomical features of the CC, such as the genu, splenium, rostrum, fornix, 'ribbon' body (Fig. 2a), 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.

To overcome the numerous deficiencies in the deformable model approach yet retain its core strengths, we borrow the layered-pyramid architecture concepts from Artificial Life modeling. We add high-level controller layers (a 'brain') on top of geometric and physical (or deformation) layers to produce an intelligent deformable organism (Fig. 1). The planned activation of these lower layers allows us to control the fitting/optimization procedure. We use a natural, intuitive medial description of object shape plus medial-based statistics of (localized) shape variation as our shape knowledge representation scheme. The medial-based shape descriptors can be easily mapped onto anatomical features of an object. A primitive cognitive layer activates 'behavior' routines (e.g. find-splenium, find-genu, find-upper-boundary-of-CC) according to a plan or schedule[1]. The behavior routines subsequently activate 'motor' (i.e. deformation) controller routines, bringing the organism closer to its sole intention of object segmentation with each step. The plan (or plans) can be generated with the aid of a human expert since the behavior routines are defined using familiar anatomical terminology.

The layered architecture approach allows the organism to make deformation decisions at the correct level of abstraction. The organism is 'self-aware' (i.e. knows where it and its parts are and what it is looking for at every stage) and therefore is able to utilize global contextual knowledge effectively. The organism begins by searching for the most stable anatomical features in the image and then proceeds to the next best feature and so on. Once stable features are found and labeled, the organism uses neighboring information and prior knowledge to determine the object boundary in regions where there is known to be little or no feature information. The organism carries out active, explicit searches for object features. It is not satisfied with the nearest matching feature but looks further within a region to find the best match and therefore avoids local minimum solutions. Furthermore, by carrying out explicit searches for features we ensure correct correspondence between the model and the data. If a feature cannot be found, the organism flags this situation. In the future, if multiple plans exist, another plan could potentially be selected and the search for the missing feature postponed until further information is available.

[1] Currently we use a single fixed plan but multiple plans with a plan selection scheme are possible.

Explicit feature search requires powerful, flexible and intuitive model deformation control. We achieve this with a set of 'motor' (i.e. deformation) controllers and medial-based deformation operators. Deformation controllers are parameterized 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 $<\mathrm{bend\_angle}, \mathrm{location}, \mathrm{scale}>$ or $<\mathrm{stretch\_length}, \mathrm{location}, \mathrm{scale}>$ into detailed deformations. Medial-based *profiles* [4] are used for shape representation, which follow the geometry of the structure and describe general and intuitive shape variation (stretch, bend, thickness). Shape deformations are obtained either as a result of applying deformation operators at certain locations and scales on the medial profiles, or by varying the weights of the main variation modes obtained from a hierarchical (multiscale) and regional (multi-location) principal component analysis of the medial profiles.

The perception system of our organism consists of a set of sensors that provide information. Sensors can be virtually anything - from edge strength and edge direction detectors to snake 'feelers'. Sensors can be focused or trained for specific image feature and image feature variation in a task-specific way and hence the organism is able to disregard sensory information superfluous to its current behavioral needs.

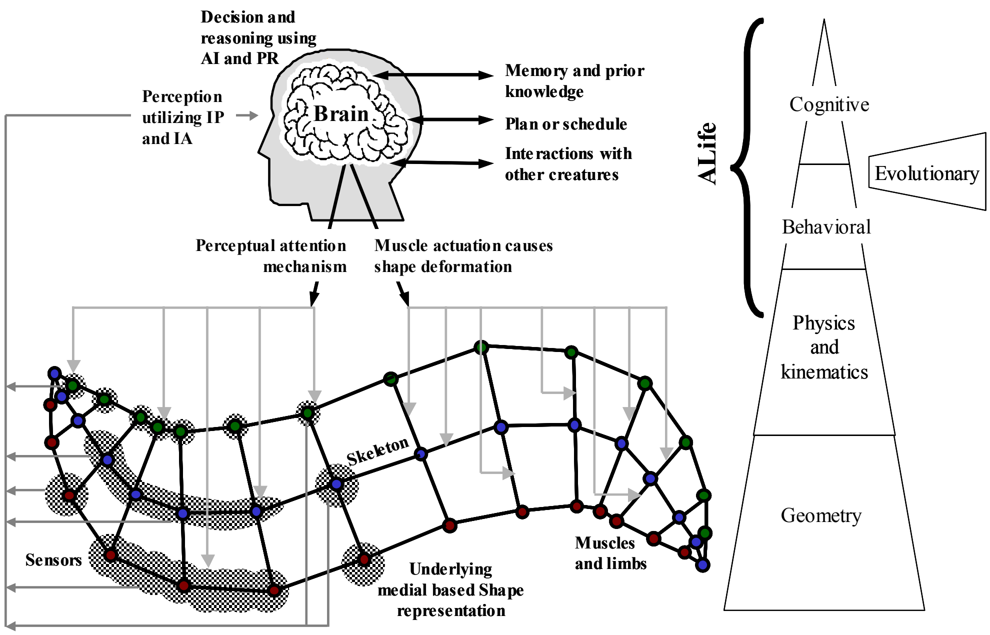

(a) **(b)**

Fig. 1. (a) Intelligent organism: Sensors are sensitive to specific stimuli. The perceived data is passed to the brain. The brain makes decisions based on perceived data, prior knowledge, and a plan. The brain issues 'muscle' actuation and sensor perception commands, the organism is then deformed and sensors perceive new data etc. (AI, PR, IP, IA: Artificial Intelligence, Pattern Recognition, Image Processing, Image Analysis). (b) The ALife modeling pyramid.

2 Intelligent Deformable Organisms for Medical Image Analysis

To demonstrate the potential of the intelligent organism approach to medical image analysis, we will describe the detailed construction of the layered-architecture for a corpus callosum organism, beginning with the lower layers and progressing upwards.

2.1. Shape Representation (Geometry)

We use medial-based shape profiles [4] to describe the body of the CC organism[2]. In this shape representation scheme, the CC anatomical structure is described with four shape profiles derived from the primary medial axis of the 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: 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 $m = 1,2,\ldots,N$ and N is the number of medial nodes. The length profile represents the distances between consecutive pairs of medial nodes, and the orientation profile represents the angles of the edges connecting the pairs of nodes. The thickness profiles represent the distances between medial nodes and their corresponding boundary points on both sides of the medial axis (Fig. 2, Fig. 3).

Currently we construct medial profiles only from the primary medial axis and have not considered secondary axes. This may prevent the CC organism from accurately representing highly asymmetrical (with respect to the primary axis) parts of some corpora callosa. We also realize that our medial shape representation needs improvement near the end caps. We are currently exploring these issues and issues related to the extension of our model to 3D and we intend to make full use of the considerable body of work of Pizer *et al* [7,8,9] on these topics.

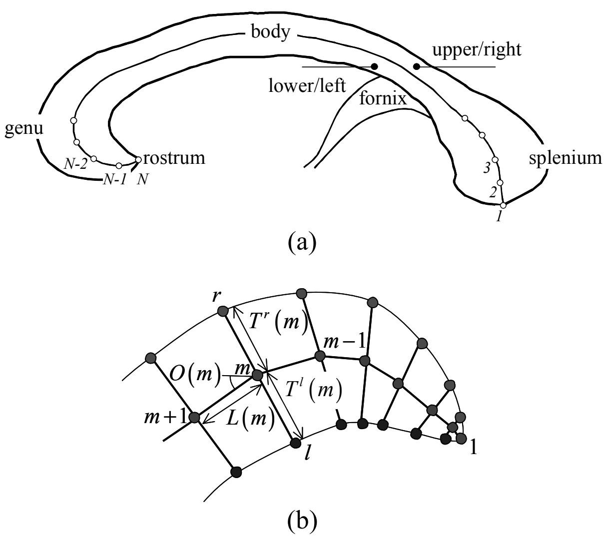

Fig. 2. (a) CC anatomical feature labels overlaying a reconstruction of the CC using the medial shape profiles shown in Fig. 3. (b) Diagram of shape representation.

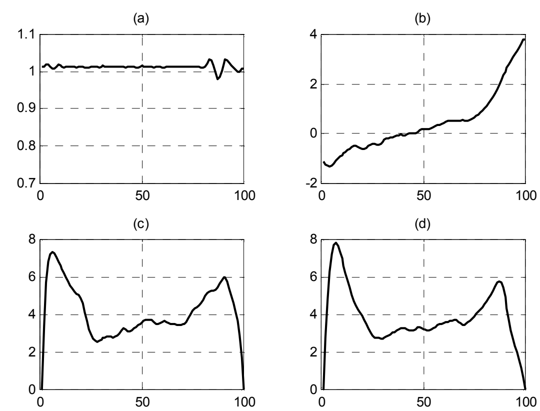

Fig. 3. Example medial shape profiles: (a) length profile $L(m)$, (b) orientation profile $O(m)$, (c) left thickness profile $T^l(m)$, and (d) right thickness profile $T^r(m)$.

2.2. Motor System

Shape Deformation (Motor Skills)

In addition to affine transformation abilities (translate, rotate, scale), we control organism deformation by defining deformation operators in terms of the medial-based shape profiles (Fig. 4). Controlled stretch (or compress), bend, and bulge (or squash) deformations are implemented as deformation operators acting on the length, orientation, or thickness profiles, respectively. Furthermore, by utilizing a hierarchical (multiscale) and regional principal component analysis to capture the shape variation statistics in a training set [4], we can keep the deformations consistent with prior knowledge of possible shape variations. Whereas general statistically-derived shape models produce global shape variation modes only [2,1,11], we

[2] Other shape representation and shape deformation schemes can be employed in the lower layers of the modeling pyramid as long as they provide sufficient shape description and deformation capabilities to the upper layers.

are able to produce spatially-localized feasible deformations at desired scales, thus supporting our goal of intelligent deformation planning.

Several operators of varying types, amplitudes, scales, and locations can be applied to any of the length, orientation, and thickness shape profiles (Fig. 5a-d). Similarly, multiple statistical shape variation modes can be activated, with each mode acting at a specified amplitude, location and scale of the shape profiles (Fig. 5e-h). In general, operator- and statistics-based deformations can be combined (Fig. 5i) and expressed as

$$p_d = \overline{p}_d + \sum_l \sum_s \left(M_{dls} w_{dls} + \sum_t \alpha_{dlst} k_{dlst} \right) \qquad (1)$$

where p is a shape profile, d is a deformation type (stretch, bend, left/right bulge), i.e. $p_d(m): \{L(m), O(m), T^l(m), T^r(m)\}$, $\overline{p}$ is the average (or original) shape profile, k is an operator profile (with unity amplitude), l and s are the location and scale of the deformation, t is the operator type (e.g. Gaussian, triangular, flat, bell, or cusp), α is the operator amplitude, the columns of M are the variation modes for a specific d, l, and s, and w contains the weights of the variation modes. Details can be found in [4].

(a)

(b)

(c)

(d)

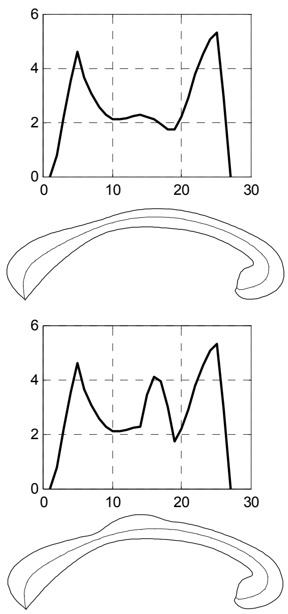

Fig. 4. Introducing a bulge on the upper boundary by applying a deformation operator on the upper thickness profile: (a) $T^r(m)$ before and (c) after applying the operator. (b) The reconstructed shape before and (d) after the operator.

a

b

c

d

e

f

g

h

i

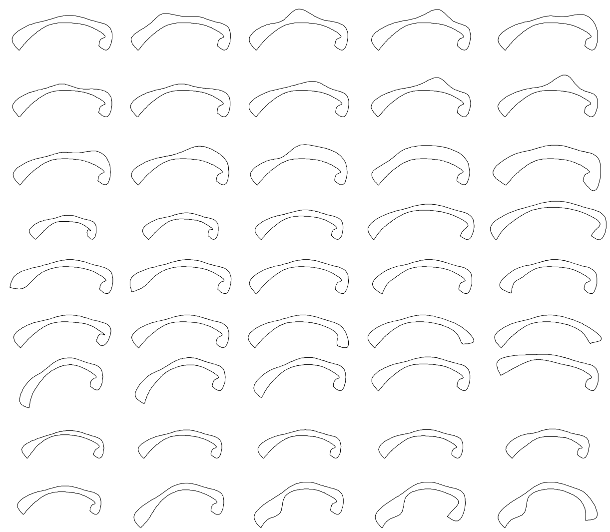

Fig. 5. Examples of controlled deformations: (a)-(c) Operator-based bulge deformation at varying locations, amplitudes, and scales. (d) Operator-based stretching with varying amplitudes over the entire structure. (e)-(g) Statistics-based bending of the left end, the right end, and the left half of the structure. (h) Statistics-based bulge of the left and right thickness over the entire structure. (i) From left to right: (1) mean shape, (2) statistics-based bending of the left half, followed by (3) locally increasing the lower thickness using an operator, followed by (4) applying an operator-based stretch and (5) adding an operator based bend to the right side of the structure.

Deformation (Motor) Controllers

We build upon the organism's low-level motor skills to construct high-level motor controllers. These parameterized procedures carry out complex deformation functions such as sweeping over a range of rigid transformation parameters, sweeping over a range of stretch/bend/thickness amplitudes at a certain location and scale, bending at increasing scales, moving a bulge on the boundary etc. Other high-level deformation capabilities include smoothing the medial or left or right boundaries, interpolating a missing part of the thickness profile, moving the medial axis (the organism's 'spine') to a position midway between the left and right boundaries, re-sampling the model by including more medial and boundary nodes, and shortening or elongating the organism while maintaining the original thickness and orientations at appropriate arc lengths.

2.3. Perception System

Different parts of the organism are dynamically assigned sensing capabilities and thus act as sensory organs (SOs) or receptors. The locations of the SOs are typically confined to the organism's body (on-board SO's) such as at its medial or boundary nodes, at curves or segments connecting different nodes, or at other inner regions. In our implementation, the SOs are made sensitive to different stimuli such as image intensity, image gradient magnitude and direction, a non-linearly diffused version of the image, an edge detected (using Canny's edge detector) image, or even the result of a Hough transform (applied to find the top of the human skull in the MR image). In general, a wide variety of image processing (IP) and image analysis (IA) techniques can be applied to the original image source and thus act as focusers or filters of the 'outside world' signals (Fig. 6). The sensed data are fed to the cognitive center of the brain for processing.

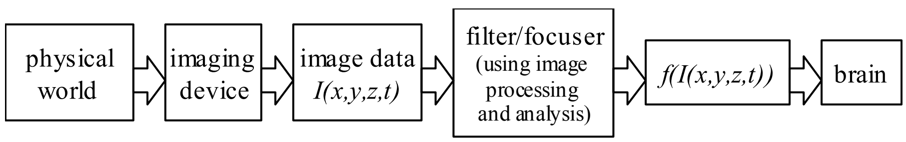

Fig. 6. Perception System.

2.4. Behavioral/Cognitive System

The organism's cognitive center combines sensory information, memorized information, and instructions from a pre-stored segmentation plan to carry out active, explicit searches for object features by activating 'behavior' routines. Behavior routines are designed based on available organism motor skills, perception capabilities, and available anatomical landmarks. For example, the routines implemented for the CC 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. The behavior routines subsequently activate the deformation controllers to complete a stage in the plan and bring the organism closer to its intention of object segmentation.

The segmentation plan provides a means for human experts to 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 corner-shaped rostrum of the CC is always very clearly defined in an MRI image, then the find-rostrum behavior should be given a very high priority. Adhering to the segmentation plan and defining it at a behavioral level imbues the organism with awareness of the segmentation process. This enables it to make very effective use of prior shape knowledge – it is applied only in anatomical regions of the target object where there is a high level of noise or known gaps in the object boundary edges etc.

The segmentation plan for the CC organism serves to illustrate the ability to harness global contextual knowledge. A CC organism is released into a 2D sagittal MRI brain image. It then goes though different 'behaviors' as it progresses towards its goal. Since the upper boundary (Fig. 2a) 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 then move downwards (through the gray and white matter) in the image space to locate the upper boundary. The organism then bends to latch to the upper boundary and activates a find-genu routine, causing the CC organism to stretch and grow along this boundary towards the genu. Once the genu is located, the find-splenium routine is activated and the organism stretches and grows in the other direction. The genu and splenium are easily detected by looking for a sudden change in direction of the upper boundary towards the middle of the head. Once the genu is found, the organism knows that the lower boundary opposite to the genu is well defined so it backs up and latches to the lower boundary. It then activates the find-rostrum behavior that tracks the lower boundary until reaching the distinctive rostrum. At the splenium end of the CC, the organism backs up and finds the center of a circle that approximates the splenium end cap. The thickness of the upper boundary is then adjusted to latch to the corresponding boundary in the image. 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 to the potentially occluding fornix. Nevertheless, the lower boundary might still dip towards the fornix so a successive step of locating where, if any, 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. At this point the CC organism has almost reached its goal. However, at this stage the medial axis is not in the middle of the CC organism so the medial axis is re-parameterized until the medial nodes are halfway between the boundary nodes.

3 Results

We have released our CC organism into several 2D mid-sagittal MRI slices of the brain. Starting from an initial default position (Fig. 7.1), the organism moves towards the top of the head (Fig. 7.2-3). Moving downwards, it then searches for the CC's upper boundary (Fig. 7.4-7) and latches to it (Fig. 7.8). Then it progresses (to the left within the figure) to find the genu (Fig. 7.9-11). It then backs-up and thickens (Fig. 7.12) preparing for tracking the lower boundary (Fig. 7.13-15) until it reaches the rostrum. The opposite side (right side) of the organism then starts tracking the upper boundary until reaching the splenium (Fig. 7.15-16). The right end is then moved to the center of a splenium-approximating circle (Fig. 7.17). The organism then thickens its lower boundary (Fig. 7.18-21) and upper boundary (Fig. 7.22-26). Although the boundary of the CC is now located (Fig. 7.26) the position of the medial (Fig. 7.27) needs to be fixed. Thus, the locations of the inner and outer boundaries are sacrificed to reparameterize the medial (Fig. 7.28-30). Finally the lower and upper boundaries are re-located again (Fig. 7.31-36) to obtain the final segmentation result (Fig. 7.36). Other segmentation results (Fig. 8) and several validated examples (Fig. 9) are also shown. In addition, we demonstrate the detection and repairing of the fornix (Fig. 10) and the organism's self-awareness (Fig. 11).

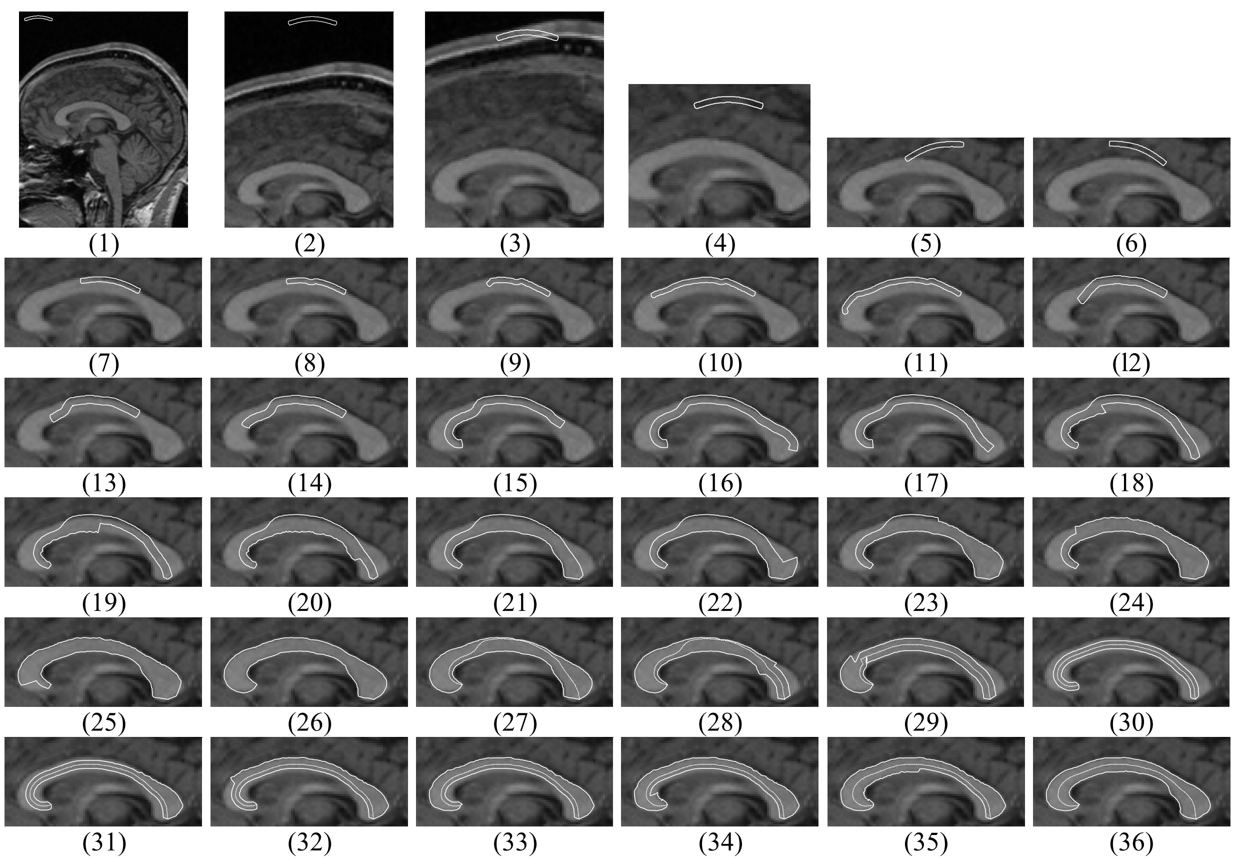

Fig. 7. Intelligent CC organism progressing through a sequence of behaviors to segment the CC.

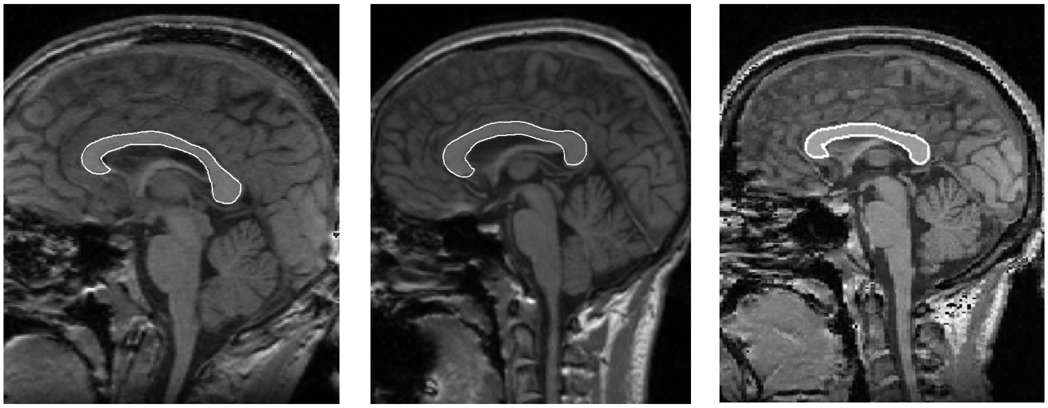

Fig. 8. Example segmentation results.

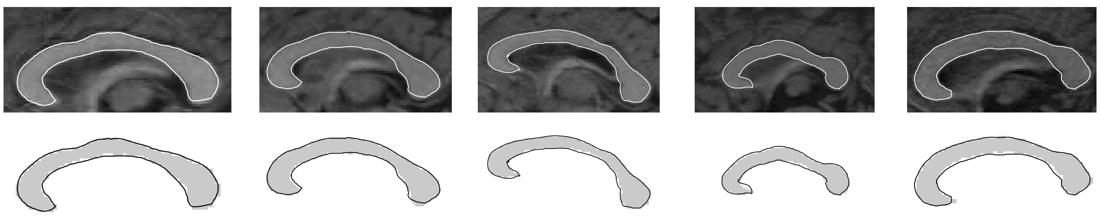

Fig. 9. Example segmentation results (top), also shown (in black) over manually segmented (gray) CC (bottom).

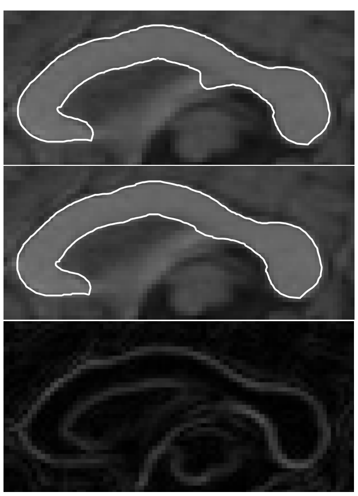

Fig. 10. Detecting and repairing the fornix dip. Before (left), after (middle), and the gradient magnitude (right).

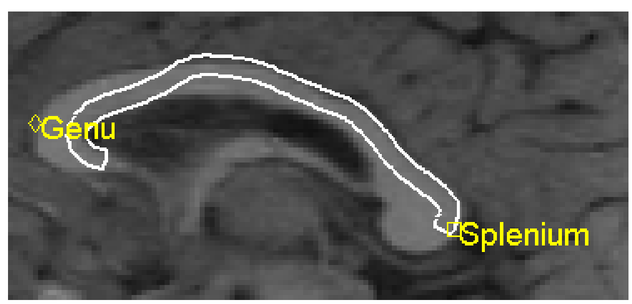

Fig. 11. The CC organism's self-awareness makes it capable of identifying landmark parts.

4 Conclusions

Robust, automatic medical image analysis requires the incorporation and intelligent utilization of global contextual knowledge. We have introduced a new paradigm for medical image analysis that applies concepts from artificial life modeling to meet this requirement. By constructing a deformable model-based framework in a layered fashion, we are able to separate the 'global' model-fitting control functionality from the local feature integration functionality. This separation allows us to define a model-fitting controller or 'brain' in terms of the high-level anatomical features of an object rather than low-level image features. The layered-architecture approach also provides the brain layer with precise control over the lower-level model deformation layer. The result is an intelligent organism that is continuously aware of the progress of the segmentation allowing it to effectively apply prior knowledge of the target object. We have demonstrated the potential of this approach by constructing a Corpus Callosum organism and releasing it into MRI brain images in order to segment and label the corpus callosum.

A number of interesting aspects of our approach are currently being considered for further exploration. These include extending our model to 3D (which would involve generating medial surfaces), designing a motion tracking plan and releasing an organism into 4D dynamic ‘environments’ (i.e. 4D images), and exploring the use of multiple plans and plan selection schemes. An interesting application we are currently planning is the construction of an ‘artery crawler’ organism that is released into the central artery in angiographic images and allowed to ‘crawl’ along arteries looking for branches, stenoses and aneurysms. Another important direction of research is the use of multiple organisms that communicate contextual image information to each other and avoid collisions (i.e. are ‘aware’ of each other) and a global brain that coordinates organism activation. This area of research has direct application to the segmentation of images with multiple anatomical structures such as brain MR images.

We argue that although the control and decision-making ability provided by the layered architecture approach raises issues of how to organize these abilities (i.e. construct the organism’s brain), the tremendous flexibility and control over the segmentation process is essential to guarantee robust automatic segmentation. In addition, human experts can help with plan generation and we also intend to explore the application of learning algorithms, such as genetic algorithms, for automatic generation of optimal plans. Finally, the layered architecture also provides the option of adding higher levels of cognitive modeling, knowledge representation, reasoning, and planning.

Acknowledgements

GH was partially funded by the Visual Information Technology (VISIT) program, Swedish Foundation for Strategic Research (SSF). The MRI data was provided by Martha Shenton, Harvard Medical School.

References

1. Cootes, T., Beeston, C., Edwards, G., Taylor, C.: A Unified Framework for Atlas Matching using Active Appearance Models. International Conference on Image Processing in Medical Imaging (1999) 322-333.
2. Cootes, T., Cooper, D., Taylor, C., Graham, J.: Active Shape Models-Their Training and Application. Computer Vision and Image Understanding 61(1) (1995) 38-59.
3. Duncan, J., Ayache, N.: Medical Image Analysis: Progress Over Two Decades and the Challenges Ahead. IEEE Transactions on Pattern Analysis and Machine Intelligence 22(1) (2000) 85 -106.
4. Hamarneh, G., McInerney, T.: Controlled Shape Deformations via Medial Profiles. Vision Interface (2001).
5. McInerney, T., Kikinis, R.: An Object-based Volumetric Deformable Atlas for the Improved Localization of Neuroanatomy in MR Images. MICCAI (1998) 861-869.
6. McInerney, T., Terzopoulos, D.: Deformable Models in Medical Image Analysis: A Survey. Medical Image Analysis 1(2) (1996) 91-108.
7. Pizer, S., Fletcher, P., Fridman, Y., Fritsch D., Gash, A., Glotzer, J., Joshi, S., Thall A., Tracton, G., Yushkevich, P., Chaney, E.: Deformable M-Reps for 3D Medical Image Segmentation. Submitted to Medical Image Analysis (2000).
8. Pizer, S., Fritsch, D., Low, K., Furst, J.: 2D & 3D Figural Models of Anatomic Objects from Medical Images. Mathematical Morphology and Its Applications to Image Processing, Kluwer Computational Imaging and Vision Series, H.J.A.M. Heijmans, J.B.T.M. Roerdink, Eds. Amsterdam, (1998) 139-150.
9. Pizer, S., Fritsch, D.: Segmentation, Registration, and Measurement of Shape Variation via Image Object Shape. IEEE Transactions on Medical Imaging 18(10) (1999) 851-865.
10. Shen, D., Davatzikos, C.: An Adaptive-Focus Deformable Model Using Statistical and Geometric Information. IEEE Transactions Pattern Analysis and Machine Intelligence 22(8) (2000) 906 -913.
11. Szekely, G., Kelemen, A., Brechbuehler C., Gerig, G.: Segmentation of 3D Objects From MRI Volume Data Using Constrained Elastic Deformations of Flexible Fourier Surface Models. Medical Image Analysis 1(1) (1996) 19-34.
12. Terzopoulos, D., Tu, X., Grzeszczuk, R.: Artificial Fishes: Autonomous Locomotion, Perception, Behavior, and Learning in a Simulated Physical World. Artificial Life 1(4) (1994) 327-351.
13. Terzopoulos, D.: Artificial Life for Computer Graphics. Communications of the ACM 42(8) (1999) 32-42.