

Deformable organisms for automatic medical image analysis

Tim McInerney^{a,b,*}, Ghassan Hamarneh^{b,c,d}, Martha Shenton^e, Demetri Terzopoulos^{b,f}

^a*School of Computer Science, Ryerson University, Toronto, ON M5B 2K3, Canada*

^b*Department of Computer Science, University of Toronto, Toronto, ON M5S 3H5, Canada*

^c*Mouse Imaging Centre, Hospital for Sick Children, Toronto, ON M5G 1X8, Canada*

^d*Department of Signals and Systems, Chalmers University of Technology, Göteborg 41296, Sweden*

^e*Laboratory of Neuroscience, Clinical Neuroscience Division, Department of Psychiatry, VA Boston Healthcare Systems, Brockton Campus, Harvard Medical School, Brockton, MA, USA*

^f*Computer Science and Mathematics Departments, Courant Institute, New York University, New York, NY 10003, USA*

Abstract

We introduce a new approach to medical image analysis that combines deformable model methodologies with concepts from the field of artificial life. In particular, we propose ‘deformable organisms’, autonomous agents whose task is the automatic segmentation, labeling, and quantitative analysis of anatomical structures in medical images. Analogous to natural organisms capable of voluntary movement, our artificial organisms possess deformable bodies with distributed sensors, as well as (rudimentary) brains with motor, perception, behavior, and cognition centers. Deformable organisms are perceptually aware of the image analysis process. Their behaviors, which manifest themselves in voluntary movement and alteration of body shape, are based upon sensed image features, pre-stored anatomical knowledge, and a deliberate cognitive plan. We demonstrate several prototype deformable organisms based on a multiscale axisymmetric body morphology, including a ‘corpus callosum worm’ that can overcome noise, incomplete edges, considerable anatomical variation, and interference from collateral structures to segment and label the corpus callosum in 2D mid-sagittal MR brain images.

© 2002 Elsevier Science B.V. All rights reserved.

Keywords: Medical image segmentation; Deformable models; Intelligent agents; Artificial life

1. Introduction

Medical imaging has become essential to the practice of medicine, but accurate, fully *automatic* medical image analysis (MIA) continues to be an elusive ideal. A substantial amount of knowledge is often available about anatomical structures of interest—characteristic shape, position, orientation, symmetry, relationship to neighboring structures, associated landmarks, etc.—and about plausible image intensity characteristics, subject to natural biological variability or the presence of pathology. Even so, MIA researchers have not yet succeeded in developing completely automatic segmentation techniques that can take full advantage of such prior knowledge to achieve segmentation accuracy and repeatability. Although it may be generally acknowledged that this will require the incorporation of context-based information within a robust

decision-making framework (Duncan and Ayache, 2000), we contend that current frameworks of this sort are inflexible and do not operate at an appropriate level of abstraction, which limits their potential to deal with the most difficult data sets.

Deformable models demonstrated early promise in image segmentation and they have become one of the most intensively researched segmentation techniques (McInerney and Terzopoulos, 1996). The classical deformable model methodology, which is epitomized by snakes (Kass et al., 1988), is based on the optimization of objective functions in conjunction with an interactive decision-making strategy that relies on human expert initialization and guidance. The difficult challenge in automating this technique is to develop intelligent initialization mechanisms, along with control mechanisms that can guide the optimization-driven segmentation process at an appropriately high level of abstraction. Researchers have tried in vain to obtain the right global behavior (i.e., on the scale of the entire image) by embedding nuggets of contextual knowl-

*Corresponding author.

E-mail address: tmcinern@scs.ryerson.ca (T. McInerney).

edge into the low-level optimization engine. As a result of such efforts, it has become painfully obvious that current deformable models have little to no explicit ‘awareness’ of where they are in the image, how their parts are arranged, or to what structures they or any neighboring deformable models are converging during the optimization process. To make progress towards full automation, we need to complement the powerful low-level feature detection and integration abilities of deformable models with flexible higher-level decision-making strategies.

With this in mind, we propose a new approach to automated MIA that augments deformable model methodologies with concepts from the field of Artificial Life (ALife) (see, e.g., Terzopoulos, 1999).¹ In particular, we develop *deformable organisms*, autonomous agents whose task is the automatic segmentation, labeling, and quantitative analysis of anatomical structures in medical images. Analogous to natural organisms capable of voluntary movement, our artificial organisms possess deformable bodies with distributed sensors, as well as (rudimentary) brains with motor, perception, behavior, and cognition centers. Deformable organisms are perceptually aware of the image analysis process. Their behaviors, which manifest themselves in voluntary movement and alteration of body shape, are based upon sensed image features, pre-stored anatomical knowledge, and a cognitive plan.

By synthesizing organisms in a bottom-up, layered fashion, we are able to separate the global, model-fitting control functionality from the local, feature detection and integration functionality, so that the deformable organism can make decisions about the segmentation process at an appropriately high level of abstraction. This layered architecture facilitates the incorporation of plans in the form of sequential search strategies; for example, plans that direct organisms to look initially for the most stable anatomical features in images before deforming or growing towards less stable features, and so on. The result is autonomous and, to some degree, intelligent segmentation algorithms that are aware of their progress and apply prior knowledge in a deliberate manner in different phases of the segmentation process. Our layered architecture and reusable behavior routines facilitate the rapid implementation of powerful, custom-tailored deformable organisms that can serve as new tools for automated segmentation, object-based registration, and the quantification of shape variation.

1.1. Illustrative examples of deformable organisms

To date, we have developed several prototype deformable organisms based on an axisymmetric body morphology. This geometric representation in conjunction with a set

of multiscale deformation operators affords the motor center of the brain of the organism precise local growth and shape control at a variety of scales. We have also experimented with a variety of reusable behavior routines that support single organism behaviors and interacting, multiple organism behaviors. Interaction among organisms may be as simple as collision detection and the imposition of non-penetration constraints between two or more organisms in contact; or one or more organisms spawning a new organism and supplying it with appropriate initial conditions; or the sharing of statistical shape constraints and/or image appearance information between organisms. More complex rule-based interactions are also possible.

Fig. 1 illustrates deformable organisms with a nontrivial example involving the detection and detailed segmentation of the lateral ventricle, caudate nucleus, and putamen in the left and right halves of a transverse 2D MR image of the brain. Since the ventricles are the most discernible and stable structures, the segmentation process begins with the release of two ventricle organisms in the black background region outside the cranium, at the upper left and right edges of the image in subfigure (1). Performing a coordinated scanning behavior, the organisms proceed first to locate the tops of the ventricles, as shown in the zoomed in view of subfigure (2), and their inner and outer (with respect to the brain) boundaries (3)–(5). Next, both ends of each ventricle organism actively stretch to locate the upper and lower lobes of the ventricle (6), and then the organism fattens to finish segmenting the ventricle (7). Each organism employs the information that it has gleaned about the shape and location of the segmented ventricles to spawn and initialize a caudate nucleus organism in an appropriate location (8). Each caudate nucleus organism first stretches to locate the upper and lower limits of the caudate nucleus (9), then fattens until it has accurately segmented the caudate nucleus (10). From its bottom-most point in the image, each caudate nucleus organism then spawns and initializes a putamen organism (11), which then moves laterally outward towards the low-contrast putamen (12). Each putamen organism then rotates and bends to latch onto the nearer putamen boundary (13). Next, it stretches and grows along the boundary until it reaches the upper- and lower-most ends of the putamen (14), thus identifying the medial axis of the putamen (15). Since the edges of the putamen boundary near the gray matter are often weak, the organism activates an explicit search for an arc (parameterized by a single curvature parameter) that best fits the low-contrast intensity variation in that region, thus completing the segmentation (16).

As a second illustrative example, Fig. 2 shows a different type of axisymmetric deformable organism specialized to ribbon-like structures. The organism segments a vascular structure in a retinal angiogram. If it is afforded insufficient prior knowledge, the organism can latch onto the wrong overlapping vessel as shown in Fig.

¹The ALife approach to MIA was introduced in our MICCAI 2001 Conference paper (Hamameh et al., 2001).




Fig. 1. Automatic brain MR image segmentation by multiple deformable organisms. The sequence of images illustrates the temporal progression of the segmentation process. Deformable lateral ventricle (1–7), caudate nucleus (8–10), and putamen (11–16) organisms are spawned in succession and progress through a series of behaviors to detect, localize, and segment the corresponding structures in the MR image (see text).

2(b). However, given a suitable repertoire of behavior routines, the vessel organism can distinguish between overlapping vessels and deal with bifurcations (Fig. 2(c)). When a bifurcation is encountered, the organism spawns two new vessel organisms (Fig. 2(d)), each of which extends along a branch (Fig. 2(e)).

1.2. Overview

In the remainder of this article we focus on a different application of deformable organisms, the automatic segmentation of the corpus callosum in a variety of 2D mid-sagittal MR brain images. The corpus callosum is the largest white-matter tract in the human brain. It serves as the primary means of communication between the two cerebral hemispheres and mediates the integration of

cortical processes from opposite sides of the brain. The presence of morphologic differences in the corpus callosum in schizophrenics has been the subject of intense investigation (Rosenthal and Bigelow, 1972). The corpus callosum may also be involved in Alzheimer's dementia (Pantel et al., 1999), mental retardation (Marszal et al., 2000), and other disorders. MR imaging has allowed researchers to study corpora callosa in vivo in order to discover and quantify morphologic differences. The detailed, automatic segmentation of the corpus callosum is therefore considered an important, though difficult, MIA problem. To this end, we develop a prototype axisymmetric deformable organism, which we call a *corpus callosum worm*. We demonstrate that this organism can overcome poor image contrast, noise, diffuse or missing boundaries, considerable anatomical variation, and interference from




Fig. 2. Multiple deformable organisms segmenting vascular structures in an angiogram. (a) Automatic labeling of vessel overlap and bifurcation. (b) A simplistic vessel organism incorrectly bends into the more prominent overlapping vessel. (c) Appropriate high-level behaviors enable the vessel organism to identify the overlap and distinguish it from bifurcations. (d) Upon identifying a bifurcation, the organism spawns two new organisms, each of which proceeds along a branch. (e) The segmented vessel and branches.

collateral structures to segment and label the corpus callosum in a variety of MR brain images.

The remainder of the article is organized as follows. Section 2 motivates our artificial life approach to MIA and provides additional technical background. We describe in Section 3 the architectural characteristics of deformable organisms. Section 4 gives the details of the corpus callosum worm organism. Section 5 presents our automated segmentation results and validates them against manual segmentations. Section 6 discusses the implications of our approach and Section 7 draws conclusions.

2. Motivation and background

Current model-based MIA frameworks utilize geometric and often physical modeling layers. The models are fitted to images by minimizing energy functions, simulating dynamical systems, or applying probabilistic inference methods, but they do not control this optimization process other than in primitive ways, such as monitoring convergence or equilibrium. Some deformable models incorporate prior information to constrain shape and image appearance and the observed statistical variation of these quantities (Cootes et al., 1995, 1999; Szekely et al., 1996). These models have no explicit awareness of where they or

their parts are, and therefore the effectiveness of such constraints is dependent upon appropriate model initialization. The lack of self-awareness may also prevent models from knowing when to trust the image feature information and ignore the constraint information or vice versa. The lack of optimization control can prevent these models from performing intelligent searches over their parameter spaces during the fitting process; i.e., the constraint information is applied more or less indiscriminately and, once set in motion, the optimization process continues ‘mechanically’ to completion.

Furthermore, because there typically is no active, deliberate search for stable image features, the models can latch onto nearby spurious features (Cootes et al., 1999). Their short-sighted decision-making abilities prevent these models from correcting missteps. Even if global optimization methods such as simulated annealing are employed to perform more extensive searches, the parameter space of the model is explored in a rather random fashion and there is no guarantee (other than an excruciatingly slow, asymptotic one) that the correct solution will be found. Moreover, it remains an open question whether suitable solution metrics can be defined for many MIA tasks using the ‘language’ of objective functions and low-level optimization, or even of probabilistic inference.

Alternatively, if a model is aware of itself and its

environment, it can potentially be programmed to perform a more intelligent search for correct solutions by exploiting global contextual knowledge more effectively. It may explore several alternative paths and choose the optimal result. For example, it is often possible to prioritize the strength or stability of different features of the target structure(s) and this knowledge may be significant in many segmentation scenarios. Indeed, hierarchical deformable model schemes that shift their focus from stable image features to less stable features have been explored (McInerney and Kikinis, 1998; Shen and Davatzikos, 2000). These features may occur at different locations and scales and may vary from low level landmark points to curves or surface patches to volumetric regions or to more complex features. However, without control over the fitting process and without a proper language with which to define the high level features, it may be difficult to exploit this information.

It is our contention that we must revisit ideas for incorporating knowledge that were explored usually with limited success in earlier MIA systems (e.g., the ALVEN cardiac left ventricular wall motion analyzer (Tsotsos et al., 1980)), and develop new algorithms that leverage top-down reasoning strategies against the powerful bottom-up feature detection and integration abilities of deformable models and other modern model-based MIA techniques.² To retain the core strengths of deformable models but afford them the ability to control themselves, it seems prudent to investigate analogies with living systems.

2.1. Artificial Life modeling

The modeling and simulation of living systems has defined an emerging scientific discipline known as Artificial Life (ALife).³ In recent years, the ALife paradigm has had substantial impact in computer graphics, giving impetus to several important avenues of research and development, including artificial plants and animals, behavioral modeling and animation, and evolutionary modeling (Terzopoulos, 1999). These graphical models typically employ geometric and physics-based techniques, as is characteristic of the deformable models used in MIA, but they also aspire to simulate many of the biological processes that characterize living systems—including birth and death, growth and development, natural selection, evolution, perception, locomotion, manipulation, adaptive behavior, learning, and cognition.

Most relevant to our MIA approach is the ALife modeling of animals. The key components of artificial

animals like the prototypical ‘artificial fishes’ (Terzopoulos et al., 1994) are synthetic bodies, including functional motor organs (contractile muscles), sensory organs (eyes, etc.) and, most importantly, brains with motor, perception, behavior, and learning centers. In the motor center of these brains, motor controllers coordinate muscle actions to carry out specific motor functions, such as locomotion and sensor control. The perception center incorporates perceptual attention mechanisms which support active perception that acquires information about the dynamic environment. The behavior center realizes an adaptive sensorimotor system through a repertoire of behavior routines that couple perception to action. The learning center in the brain enables the artificial animal to learn motor control and behavior through practice and sensory reinforcement.

To manage the complexity, artificial animals are best organized hierarchically, with each successive modeling layer adding to the more basic functionalities of underlying layers (Terzopoulos, 1999). At the base of the modeling hierarchy, a geometric modeling layer represents the morphology of the animal. Next, a physical modeling layer incorporates biomechanical principles to constrain geometry and emulate biological tissues. Further up the hierarchy, a motor control layer motivates internal muscle actuators to synthesize lifelike locomotion. Behavioral and perceptual modeling layers cooperate to support a reactive behavioral repertoire. The apex of the modeling pyramid, the domain of classical artificial intelligence, simulates the deliberative abilities of higher animals. Here, a cognitive modeling layer concerns how knowledge is represented and how automated reasoning and planning processes achieve high-level goals.

3. Deformable organism architecture

We create deformable organisms by adding high-level control layers (a ‘brain’) atop the standard geometric and physical layers of deformable models (Fig. 3). The deliberate activation of these lower layers allows the brain to control the fitting/optimization procedure. This high-level control is made possible by utilizing prior knowledge, memorized information, sensed image features, and even inter-organism interaction. The high-level control layers coupled with a multi-scale shape representation scheme provide the ability to define anatomical features in a very general fashion and to prioritize the search for these features, where the search priority is typically based on the stability of the features found in an image.

We currently use an axisymmetric representation of object shape plus medial-axis-based statistics of (localized) shape variation as our prior shape knowledge representation scheme. This scheme allows us to represent a class of 2D objects that are primarily ribbon or tube shaped. The axisymmetric shape descriptors may easily be mapped onto

²Note that the ‘visual routines’ theory of Ullman (1984), which was proposed as a model of intermediate visual perception, also includes a prominent top-down, sequential control mechanism.

³See (Levy, 1992) for an entertaining perspective of the artificial life field. Journals such as *Artificial Life* and *Adaptive Behavior* chronicle the state of the art.




Fig. 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.

anatomical features of an object. A primitive cognitive layer activates behavior routines (e.g., for a corpus callosum (CC) (Fig. 5(a)) organism: find-splenium, find-genu, find-upper-boundary-of-CC) according to a plan or schedule (Fig. 4). The behavior routines, in turn, activate

motor (i.e., deformation) controller routines or growth controller routines, enabling the organism to fulfill its goal of object segmentation. The plan (or plans) can be generated with the aid of a human expert, since the behavior routines are defined using familiar anatomical terminology.

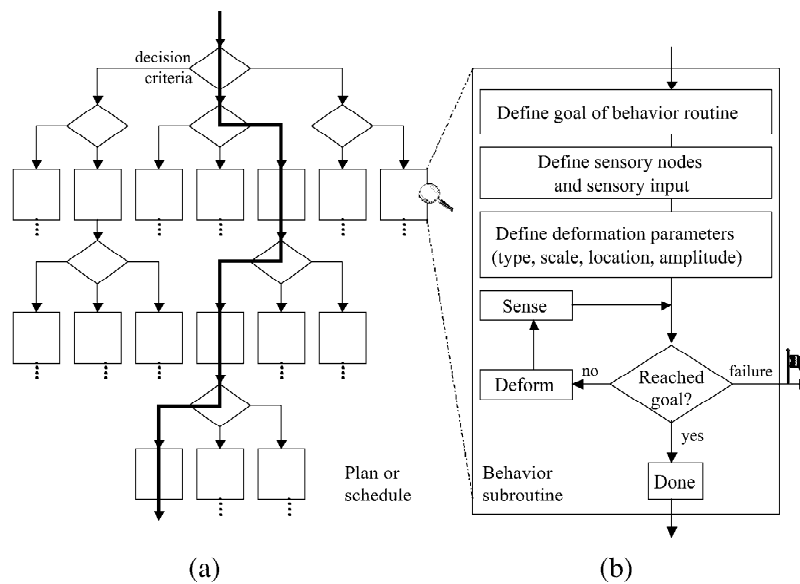


Fig. 4. (a) A procedural representation of a fragment of a deformable organism's plan or schedule. The organism goes through several behavior routines (bold path in (a)). (b) A simple example of a standard behavior routine.

3.1. Intelligent decision and control

The layered architecture allows an organism to make deformation decisions at the correct level of abstraction. An organism possesses a non-trivial ‘awareness’ (i.e. knows where it is and where its parts are and what it is seeking at every stage) and is therefore able to utilize global contextual knowledge effectively. An organism begins by searching for the most stable anatomical features in the image and then proceeds to the next best features and so on. Alternatively, an organism may interact with other organisms to determine optimal initial conditions or resolve conflicting views of data. Once stable features are found and labeled, an organism can selectively use prior knowledge in regions known to offer little or no feature information. That is, the organism intelligently ‘fills in’ the boundary in ways tailored to specific regions of interest in the target structure.

An organism carries out active, explicit searches for stable anatomical features. Its awareness allows it to perform these searches intelligently. 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 sub-optimal solutions. 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. If multiple plans exist, another plan can be selected and/or the search for the missing feature postponed until further information is available (from, for example, 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.

Explicit feature search requires powerful, flexible and intuitive model deformation control coupled with a flexible feature perception system. We currently achieve this with a set of ‘motor’ (i.e. deformation) controllers and medial-axis-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 $\langle \text{BEND-ANGLE}, \text{LOCATION}, \text{SCALE} \rangle$ or $\langle \text{STRETCH-LENGTH}, \text{LOCATION}, \text{SCALE} \rangle$ into detailed deformations. Medial-based profiles (Hamarneh and McInerney, 2001) 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 (multi-scale) and regional (multi-location) principal component analysis of the profiles. We describe these organism components in detail in Section 4.

Finally, 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 state and undergo controlled deformations as it carries out its model-fitting plan. Which 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. In summary, the ALife modeling paradigm provides a common framework and standard behavior subroutines upon which to build powerful and flexible ‘custom-tailored’ models with the potential for robustness and generality.

4. A corpus callosum (CC) worm organism

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

4.1. Shape representation (geometry)

We use axisymmetric shape profiles (Hamarneh and McInerney, 2001) to describe the body of the CC worm organism. 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 $m = 1, 2, \dots, N$, 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 (Figs. 5 and 6).

4.2. Motor system

4.2.1. Shape deformation (motor skills)

Aside from a few notable exceptions (e.g., Staib and Duncan, 1992; Terzopoulos and Metaxas, 1991), most deformable models do not have intuitive, multi-scale, multi-location deformation ‘handles’. The lack of global shape descriptors makes them cumbersome for higher-level guidance and they are unable to perform global deforma-




Fig. 5. (a) CC anatomical feature labels overlaying a reconstruction of the CC using the medial shape profiles shown in Fig. 6. (b) Diagram of shape representation.

tions, such as bending, and global motions such as sliding, or backing up. Therefore, it becomes extremely difficult to develop reasoning or planning strategies for such models.

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. 7). 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 (multi-scale) and regional principal component analysis to capture the shape variation statistics in a training set (Hamarneh and McInerney, 2001), we can keep the deformations consistent with prior knowledge of

possible shape variations. Whereas general statistically-derived shape models produce only global shape variation modes (Cootes et al., 1999; Szekely et al., 1996), we 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. 8(a)–(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. 8(e)–(h)). In general, operator- and statistics-based deformations can be combined (Fig. 8(i)) and expressed as

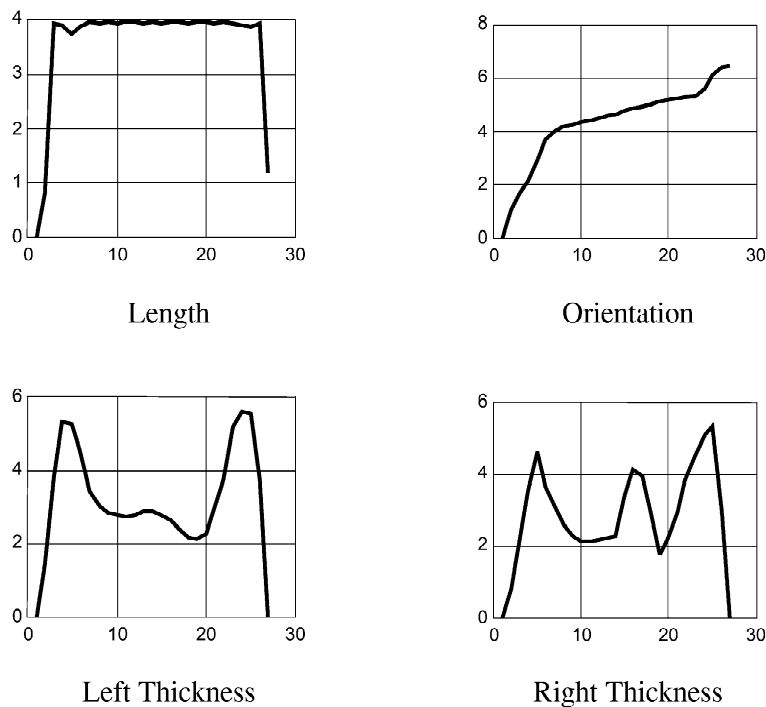


Fig. 6. Example medial shape profiles.




Fig. 7. Introducing a bulge on the CC upper boundary by applying a deformation operator on the upper thickness profile $T'(m)$. (a) $T'(m)$ before and (b) after applying operator. (c) Reconstructed shape before and (d) after applying operator.

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

where p is a shape profile, d is a deformation type (stretch, bend, left/right bulge), i.e. $p_d(m)$: $\{L(m), O(m), T'(m), T''(m)\}$, \bar{p} is the average shape profile, k is an operator profile (with unit 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 variation mode weights. Details can be found in (Hamarnah and McInerney, 2001).

4.2.2. 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, for example, smoothing the medial/left/right boundaries, interpolating a missing part of the thickness profile, moving the medial axis to a position midway between the left and right boundaries, and re-sampling the model by including more medial and boundary nodes.

4.3. Perception system

The perception system of our organism consists of a set of sensors that provide information. We can incorporate a variety of sensors from edge strength and edge direction detectors to snake 'feelers', etc. 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. 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 SOs) such as at its medial or boundary nodes, at curves or segments connecting different nodes. In our CC organism 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 the Canny edge detector) image, or even the result of a Hough transform. In general, a wide variety of image processing/analysis techniques can be applied to the input images.

4.4. Behavioral/cognitive system

An organism's cognitive center combines sensory information, memorized information, and instructions from a pre-stored segmentation plan to carry out active, explicit searches for stable 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 or growth controllers to complete a stage in the plan and bring an organism closer to fulfilling its object segmentation mission.

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




Fig. 8. Examples of controlled deformations. (a)–(c) Operator-based bulge deformation at varying locations/amplitudes/scales. (d) Operator-based stretching with varying amplitudes over entire CC. (e)–(g) Statistics-based bending of left end, right end, and left half of CC. (h) Statistics-based bulge of the left and right thickness over entire CC. (i) From left to right: (1) mean shape, (2) statistics-based bending of left half, followed by (3) locally increasing lower thickness using operator, followed by (4) applying operator-based stretch and (5) adding operator based bend to right side of CC.

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. The segmentation plan and its supporting behaviors give the organism an awareness of the segmentation process. This enables it to make very effective use of prior shape knowledge, which it deliberately applies as needed (for example, in anatomical regions of the target object where there is a high level of noise or known gaps in the object boundary edges).

We describe a detailed segmentation plan for the CC worm organism by example in the next section.

5. Experiments and results

The segmentation and labeling of the corpus callosum is an important first step for subsequent CC shape analysis and classification. In addition, it may also form an important stage in the automatic segmentation of the entire brain and all of its parts. Identification and labeling of key structures and shape features in the brain can ensure correct correspondence, a crucial part of subsequent brain comparison studies.

5.1. Corpus callosum data

Our data comprises a set of 2D mid-sagittal MR brain images. For each image in the dataset, the corpus callosum has been manually segmented by an expert. We use these expert segmentations to validate our automatic segmentation algorithms. Details of the data can be found in (Shenton et al., 1992).

In mid-sagittal MR imagery the middle of the CC is known to be located approximately halfway across the top of the head. A priori knowledge about the appearance of the CC is that it has a homogeneous intensity and is surrounded by the cerebrum, which appears darker and less homogeneous in intensity. Referring to the CC anatomy which is labeled in Fig. 5(a), the CC is a bright, stripe-like structure that curves downwards with respect to the top of the head and features a left and right ‘end-cap’. Its width is consistent (symmetric) in the middle portions. The left end cap (assuming the front of the head faces left) curves more rapidly towards the middle of the head and tapers to a point (the rostrum), making the left end-cap roughly triangular in shape. The right end-cap also curves inwards and is approximately circular (the splenium).

Although segmenting the CC may at first seem simple, it is a deceptively subtle task that can be quite challenging, especially when automatic labeling of the CC is a subsidiary requirement. Although the overall/global shape of the CC is relatively consistent, the local shape variation (over many scales and locations within the CC) is dramatic. The intensity of the CC also varies considerably from one MR image to another and there are often spurious imaging artifacts to contend with. There can be gaps of various sizes in the boundaries of the CC almost anywhere. Parts of the CC may narrow and bumps may appear almost anywhere.

The fornix is a thin structure that may or may not contact the CC in the mid-sagittal MR image on the underside of the CC. It is approximately the same brightness as the CC, and the size variation and position variation of the contact region varies considerably. Since the fornix is approximately the same brightness as the CC,

it cannot be distinguished by intensity alone. For these reasons, a standard deformable contour model or pure intensity-based techniques such as region growing will fail to extract automatically the correct boundary in many, if not most, cases—a single set of parameters does not exist that will guarantee a correct segmentation.

Even if a good solution is obtained, the subsequent task of labeling the CC parts remains difficult. Even more tightly constrained models will have problems dealing with the shape anomalies unless initialized very close to the target boundary—a task almost as difficult as the segmentation task itself. However, we do know that the upper middle boundary of the CC is consistently distinguishable. The position of the contact region between the fornix and CC is consistently right of middle. The left end-cap of the CC (the rostrum) is consistently triangular in shape (although a large variation in position, size, and shape of this triangle is exhibited) and the tip of the rostrum is a rather stable landmark.

We have constructed a CC organism that utilizes the above information to intelligently fit itself to the data. The information is used to assign measures of anatomical feature stability. We have also gathered multi-scale and

regional statistics of CC shape variation that can be used to keep model deformations consistent with prior knowledge of possible shape variation.

Fig. 9 illustrates the progression of the segmentation plan of the CC organism. Starting from an initial default position shown in subfigure (1), the CC organism goes through different behaviors as it progresses towards 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 organism activates behaviors to locate first the top of the head (subfigures (2–3)), then moves downwards through the gray and white matter in the image space to locate the upper boundary (4–7). The organism then bends to latch to the upper boundary (8) and activates a find-genu routine, causing the CC organism to stretch and grow along this boundary towards the genu (9–11). It then activates the find-rostrum routine causing the organism to back up, thicken (12), and track the lower boundary until reaching the distinctive rostrum (13–15). Once the rostrum is located, the find-splenium routine is activated and the organism stretches and grows in the other direction (15–16). The genu and splenium are easily detected by looking for a sudden




Fig. 9. Deformable corpus callosum organism progressing through a sequence of behaviors to segment the CC.




Fig. 10. Segmentation result (a) before and (b) after detecting and repairing the fornix dip. (c) The CC organism's self-awareness enables it to identify landmark parts.

change in direction of the upper boundary towards 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 (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 to the

potentially connected fornix structure (18–21). Nevertheless, the lower boundary might still dip towards the fornix so a successive step of locating where, if at all, the fornix does connect to 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




Fig. 11. Sample corpus callosum image segmentation results. (a)–(g) (top) CC organism overlaid on raw images; (bottom) manually segmented images (gray regions) with overlaid CC organism for comparison. The CC organism in segmentation sample (h) deviates in some places from the manually segmented result (i), but it accurately matches edges detected in the corpus callosum image (j).

applied only when and where required. If the fornix is indeed connected to the CC, any detected dip in the organism's boundary is repaired by interpolation using neighboring thickness values (Fig. 10). The thickness of the upper boundary is then adjusted to latch to the corresponding boundary in the image (22–26). At this point the boundary of the CC is located (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 (27) so it is re-parameterized until the medial nodes are halfway between the boundary nodes (28–30). Finally the upper and lower boundaries, which were reset in the previous step, are relocated (31–36) to obtain the final segmentation result shown in subfigure (36).

Fig. 11 shows several sample segmentation results. Table 1 presents quantitative error measurements for 26 corpora callosa relative to manual segmentations by an expert. Our validation study indicates the accuracy and of the corpus callosum organism in these test cases.

Table 1

Mean, maximum, and standard deviation of the shortest distances between automatically extracted and expert segmented CC boundaries. The last column lists pixel distances between the automatically and manually labeled rostrum tip

Case no.	Mean error (pixels)	Max. error (pixels)	S.D. (pixels)	Rostrum tip (pixels)
1	0.50	1.51	0.35	2.27
2	0.77	2.81	0.69	1.99
3	0.48	1.32	0.30	0.66
4	0.42	1.46	0.31	1.72
5	0.53	1.99	0.38	1.26
6	0.53	2.22	0.41	0.20
7	0.54	2.31	0.41	1.22
8	0.53	1.60	0.34	1.16
9	0.46	1.50	0.32	0.97
10	0.44	2.14	0.36	2.17
11	0.43	1.71	0.33	0.53
12	0.59	2.19	0.51	2.14
13	0.57	2.46	0.51	3.66
14	0.66	3.39	0.65	1.29
15	0.43	1.54	0.28	1.68
16	0.46	1.34	0.29	1.48
17	0.44	2.19	0.40	1.27
18	0.67	2.24	0.46	2.95
19	0.53	1.47	0.35	0.74
20	0.70	4.67	0.85	1.46
21	0.48	1.87	0.37	2.40
22	0.58	3.35	0.63	3.05
23	0.46	1.38	0.32	0.64
24	0.47	1.47	0.30	0.64
25	0.83	4.52	1.05	2.38
26	0.66	3.26	0.68	0.43
Overall avg. errors (pixels)	0.54	2.23	0.46	1.55

6. Discussion

It is generally acknowledged that the success of automated MIA depends on the effective use of all available global contextual knowledge. Current deformable model frameworks are making more sophisticated use of prior shape and appearance knowledge, but they are attempting to do so within the confines of standard local and/or global optimization methods. This hampers the application of prior knowledge and does not facilitate its incorporation into deliberately customized global searches for good solutions. Consequently, while current frameworks may work well on data close to the norm, data with abnormal or spurious features may cause these methods to fail in many segmentation scenarios. We believe that higher-level control and guidance of the model optimization is necessary in order to use global contextual knowledge effectively and completely. Unfortunately, current optimization-based frameworks are not amenable to the addition of such higher-level controllers.

The deformable organism approach, with its layered architecture, is an attempt to construct a framework that has the necessary properties. The cost, however, is increased complexity. Our framework is admittedly more complex than standard deformable model frameworks. We have attempted to structure this complexity in a manageable manner by adopting ALife modeling concepts and terminology and using a layered architecture. Nevertheless, we currently have no automatic scheme for designing suitable brains for deformable organisms. Identifying the correct set of behaviors, implementing the behaviors and setting behavior parameters, choosing appropriate sensors, and planning behaviors are all issues with which we are currently experimenting. For example, in Fig. 12 the CC organism has slightly mislabeled the rostrum and was only able to partially repair the minor fornix dip. We are continuing to fine-tune the image feature sensors of the CC organism. However, we believe these issues are primarily implementation issues and the underlying principle of using higher-level guidance is sound. We also intend to explore ways to aid in the construction of brains, such as supervised learning methods, and we are attempting to identify common behaviors that can be used by many organisms. We have already determined that once the CC brain had been constructed, the other brains (lateral ventricles, caudate nucleus, etc.) were much easier to construct by reusing or slightly modifying existing behavior routines (although admittedly these organisms have not been as well tested).

Another motivation for frameworks with explicit guidance mechanisms is that the potential for human intervention during the segmentation process is maintained. Completely automatic analysis of all data sets may be an unrealistic goal, even in the long term. It may be more achievable to design for highly-automated processing with the ability to flag abnormal situations, allowing for varying




Fig. 12. Our current fornix-detector was not able to detect this subtle fornix dip, and the rostrum was slightly mislabeled.

degrees of human expert intervention, and providing the ability to continue with the processing if the situation can be easily repaired. As a simple example, the CC organism may not be able to find the rostrum tip (Fig. 13(a,b)). This failure is flagged and the human expert intervenes by deforming the model (Fig. 13(c)) to fit the rostrum tip. The expert reactivates the CC organism and the organism proceeds to complete the segmentation (Fig. 13(d)).

With regard to bottom-up processing, a deformable organism's layered architecture permits the replacement of the lower levels with alternative shape representation and

deformation schemes, so long as they provide sufficient support for intuitive, multi-scale and multi-location deformation of an organism's body. For the axisymmetric deformable organisms presented in this paper, we associate medial profiles only with a 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 certain 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




Fig. 13. Example of human intervention for accurately locating the rostrum. (a) The organism latches onto the lower boundary of the CC and then activates the find-rostrum behavior routine. (b) Utilizing knowledge of the rostrum shape and orientation, the organism was not able to detect a plausible rostrum tip and flags the user. (c) The user intervenes and positions the model correctly. (d) The organism is reactivated to complete the segmentation.

the extension of our model to 3D and we intend to make use of the considerable body of work of Pizer and his co-workers on these topics (e.g., Pizer and Fritsch, 1999).

7. Conclusion

We have introduced deformable organisms, a promising new paradigm for automatic medical image analysis that applies artificial life modeling concepts to deformable models, for the purposes of incorporating and utilizing all available prior knowledge and global contextual information. The organism framework enables us to separate global top-down, model-fitting control functionality from the local, bottom-up, feature integration functionality. This separation allows us to define model-fitting controllers or ‘brains’ in terms of the high-level anatomical features of anatomical objects of interest, rather than low-level image features. The result is an ‘intelligent’ deformable organism that is continuously ‘aware’ of the progress of the segmentation, allowing it to apply prior knowledge about target anatomical objects in a deliberate fashion.

We demonstrated the potential of our novel approach by constructing a ‘worm’ organism and releasing it into mid-sagittal MR brain images in order to segment and label the corpus callosum. The axisymmetric shape representation of the organism affords its brain with precise control over the lower-level model deformation layer, which is crucial for a detailed CC analysis. Our validation study of the corpus callosum organism indicates its accuracy in several test cases. As a demonstration of the potential generality of our approach, additional multi-organism segmentation results of a more elaborate, albeit preliminary, nature were presented, including combined ventricle/caudate nucleus/putamen segmentation in a brain MR image (Fig. 1) and overlapping/bifurcating vessels in an angiogram (Fig. 2).

Several interesting aspects of our approach are currently being considered for further exploration. These include developing 3D deformable organisms, designing a motion tracking plan and releasing an organism into 4D (dynamic 3D) images, exploring the use of multiple plans and plan selection schemes, developing more sophisticated organism interactions, and exploring the use of human experts and learning algorithms to automate the process of generating optimal cognitive plans.

Acknowledgements

We would like to thank Rafeef Abugharbieh for her assistance with the figures. The research reported herein was supported by Communications and Information Technology Ontario, an Ontario Centre of Excellence, and by the Natural Sciences and Engineering Research Council of

Canada. GH was funded in part by the Visual Information Technology (VISIT) program, Swedish Foundation for Strategic Research (SSF). This work was also supported, in part, from NIH grants MH 50740 and NIH MH 01110 awarded by Dr. Shenton.

References

- Cootes, T., Cooper, D., Taylor, C., Graham, J., 1995. Active shape models—their training and application. *Computer Vision and Image Understanding* 61 (1), 38–59.
- Cootes, T., Beeston, C., Edwards, G., Taylor, C., 1999. A unified framework for atlas matching using active appearance models. In: *Proc. Image Processing in Medical Imaging Conf.*, Visegrad, Hungary, pp. 322–333.
- Duncan, J., Ayache, N., 2000. Medical image analysis: Progress over two decades and the challenges ahead. *IEEE Trans. Pattern Analysis and Machine Intelligence* 22 (1), 85–106.
- Hamareh, G., McInerney, T., 2001. Controlled shape deformations via medial profiles. In: *Proc. Vision Interface*, Ottawa, Canada, pp. 252–258.
- Hamareh, G., McInerney, T., Terzopoulos, D., 2001. Deformable organisms for automatic medical image analysis. In: *Proc. Medical Image Computing and Computer-Assisted Intervention*, Utrecht, The Netherlands, pp. 66–75.
- Kass, M., Witkin, A., Terzopoulos, D., 1988. Snakes: Active contour models. *International J. of Computer Vision* 1 (4), 321–331.
- Levy, S., 1992. *Artificial Life*. Pantheon.
- Marszal, E., Jamrow, E., Pilch, J. et al., 2000. Agenesis of corpus callosum: Clinical description and etiology. *Journal of Child Neurology* 15, 401–405.
- McInerney, T., Kikinis, R., 1998. An object-based volumetric deformable atlas for the improved localization of neuroanatomy in MR images. In: *Proc. Medical Image Computing and Computer-Assisted Intervention*, Cambridge, MA, pp. 861–869.
- McInerney, T., Terzopoulos, D., 1996. Deformable models in medical image analysis: A survey. *Medical Image Analysis* 1 (2), 91–108.
- Pantel, J., Schroder, J., Jauss, M. et al., 1999. Topography of callosal atrophy reflects distribution of regional cerebral volume reduction in Alzheimer’s disease. *Psychiatry Research* 90, 180–192.
- Pizer, S., Fritsch, D., 1999. Segmentation, registration, and measurement of shape variation via image object shape. *IEEE Trans. Medical Imaging* 18 (10), 851–865.
- Rosenthal, R., Bigelow, L.B., 1972. Quantitative brain measurements in chronic schizophrenia. *British Journal of Psychiatry* 121, 259–264.
- Shen, D., Davatzikos, C., 2000. An adaptive-focus deformable model using statistical and geometric information. *IEEE Trans. Pattern Analysis and Machine Intelligence* 22 (8), 906–913.
- Shenton, M., Kikinis, R., Jolesz, F., Pollak, S., LeMay, M., Wible, C., Hokama, H., Martin, J., Metcalf, D., Coleman, M. et al., 1992. Abnormalities in the left temporal lobe and thought disorder in schizophrenia: A quantitative magnetic resonance imaging study. *New England J. Medicine* 327, 604–612.
- Staib, L., Duncan, J., 1992. Boundary finding with parametrically deformable models. *IEEE Trans. Pattern Analysis and Machine Intelligence* 14 (11), 1061–1075.
- Szekely, G., Kelemen, A., Brechbuehler, C., Gerig, G., 1996. Segmentation of 3D objects from MRI volume data using constrained elastic deformations of flexible Fourier surface models. *Medical Image Analysis* 1 (1), 19–34.
- Terzopoulos, D., 1999. Artificial life for computer graphics. *Communications of the ACM* 42 (8), 32–42.
- Terzopoulos, D., Metaxas, D., 1991. Dynamic 3D models with local and global deformations: Deformable superquadrics. *IEEE Trans. Pattern*

- Analysis and Machine Intelligence 13 (7), 703–714.
- Terzopoulos, D., Tu, X., Grzeszczuk, R., 1994. Artificial fishes: Autonomous locomotion, perception, behavior, and learning in a simulated physical world. *Artificial Life* 1 (4), 327–351.
- Tsotsos, J., Mylopoulos, J., Covvey, H.D., Zucker, S.W., 1980. A framework for visual motion understanding. *IEEE Trans. Pattern Analysis and Machine Intelligence* 2 (6), 563–573.
- Ullman, S., 1984. Visual routines. *Cognition* 18, 97–159.