FITTING 3D DEFORMABLE MODELS TO MICROSCOPE IMAGES

Danillo Roberto Pereira and Jorge Stolfi

Institute of Computing – University of Campinas (UNICAMP)

dpereira@ic.unicamp.br   stolfi@ic.unicamp.br

Abstract –We describe a general method for fitting 3D models of deformable biological structures to microscope images. The method uses multiscale image matching techniques with gradual introduction of parameters, as well as specialized image synthesis methods and image distance metrics.

Keywords – 3D deformable models, fitting, allign, multiscale, biological models.

1. INTRODUCTION

1.1   MOTIVATION AND STATEMENT OF THE PROBLEM

The popularity of digital microscopes created an acute shortage of manpower to analyze and classify these images. In clinical microscopy examinations, for example, a single image that is produced in a matter of seconds may demand hours of visual inspection.

In this article we describe an algorithm to automate this task. We are mainly concerned with biological structures smaller than ≈ 1 mm imaged through a microscope of some sort. To simplify the discussion we will refer to such structures as *organisms*, even though our method can be applied also to isolated cells of multicellular organisms, organs, organelles, etc.

Each kind of organism, at a specific stage of life, usually has recognizable morphology. For example Figures 1 and 2 show the schematic morphology of erythrocytes (red blood cells) and a of *Paramecium* (an aquatic unicellular animal).
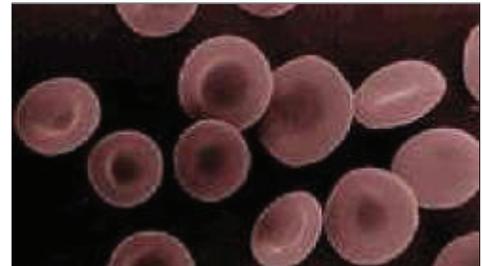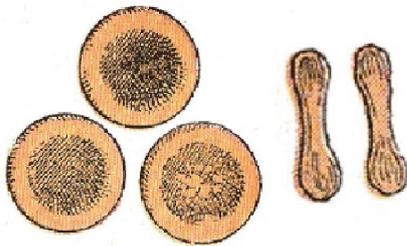

Figure 1: Schematic morphology of erythrocytes [1] (left) and a SEM image of the same (right).
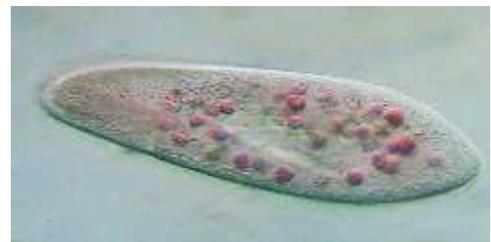

Figure 2: Schematic morphology of *Paramecium* sp. [1] (left) and an optical image of same [2] (right).

We consider here a common type of analysis which is the location and identification of certain organisms in a given image. A major difficulty in this problems is the fact that the same organism can produce very different images, depending on its orientation in space. Many organisms of interest have gelatinous and/or elastic consistency and therefore may suffer considerable deformation, as can be seen in Figure 3. They may also have visible internal structures which may be mobile and independently deformed; as well as ingested particles and externally attached debris. See Figure 4.

This task is made harder by the translucent nature of most microscopic biological structures, and by the complexities of the microscope imaging process — which may include very short depht-of-focus, interference and diffraction artfacts, and the common occurrence of dust and other debris over and around the organism.

Figure 3: Optical microscope images of *Caenorhabditis elegans*($\approx$ 1mm) [3, 4] showning deformation.

Figure 4: Microscopic image of a *Daphnia* [5] (left, $\approx 1mm$) and *Paramecium* sp (right) [6] (right, $\approx 0.3mm$), showing internal structures, ingested particles and attached debris.

Yet, with some experience, an human observer can identify microscopic images of the organisms of various kinds, and align them with three-dimensional diagrams like those of Figure 1 and 2. Our goal is to replicate this human ability in software, as far as possible.

Our algorithm can be adapted to solver other variants of this task, such as finding instances of a specified organism in a given image; identifing an imaged organism among a collection of known organisms; and locating all images from a given collection where a specified organism appears. Our method may also be useful in any application that requires matching a deformable 3D model to a 2D image, such as face recognition, industrial inspection of non-rigid parts, analysis of radiographs and angiographs, augmented reality and many others.

2. CHARACTERISTICS OF MICROSCOPIC IMAGES

There are several classes of microscopic images, defined by the type of instrument and the observation techniques used. There are two main microscope types commonly used in biology, *optical* and *electron*. Each could in principle be used *transmission mode* or *scattering mode*, and the image may be formed by *projection* onto a 2D sensor or by *scanning* a beam over the specimens.

Optical microscopes use light (usually visible or ultraviolet), almost always in transmission mode. The light that passes through the material is focused by a set of lenses, and projected onto a 2D sensor which generates the digital image. Specimens can be observed in their natural state, including live (see Figure 5 (left)) or after a process of preparation and staining. The preparation can result in flattening or cutting the organism into thin layers so that it becomes two-dimensional. Optical microscopes have a resolution limit of 200 nm, since optical lenses cannot focus light with smaller wavelength.

The electron microscope uses electrons instead of light. These particles have much shorter quantum wavelength than light photons, so they can form images with much greater resolution. In transmission mode the electrons that pass through the object are focused by magnetic "lenses" and projected onto a sensor. The fraction transmitted at each point of the object is represented in the image by a shade of gray. For this type of microscopy, the material must be cut into very thin layers. See Figure 5 (center).

Electron microscopes are often used in scanning mode. In the *scanning electron microscope* (SEM) the image is formed one pixel at time by a moving beam of electrons that sweeps over the specimen. Typically one measures the electrons scattered by the specimen instead of transmitted throught it. SEM can be used with specimens of arbitrary thickness, but generally requires the object to be coated with a thin layer of metal, since uncoated organic substances do not scatter enough electrons to form a usable image. SEM images typically have a characteristic three-dimensional shadowed appearance and show only the outer surface of the organism. See Figure 5 (right).

In this article, we consider only optical images of unprepared specimens, and images obtained with scanning electron microscopes — where the relationship between the 3D object and 2D image can be described by a geometric projection. We exclude images of prepared or sectioned specimens (such as most TEM images), since in them correspondence between the 3D schematic model of the organism and the 2D image is much more complex.
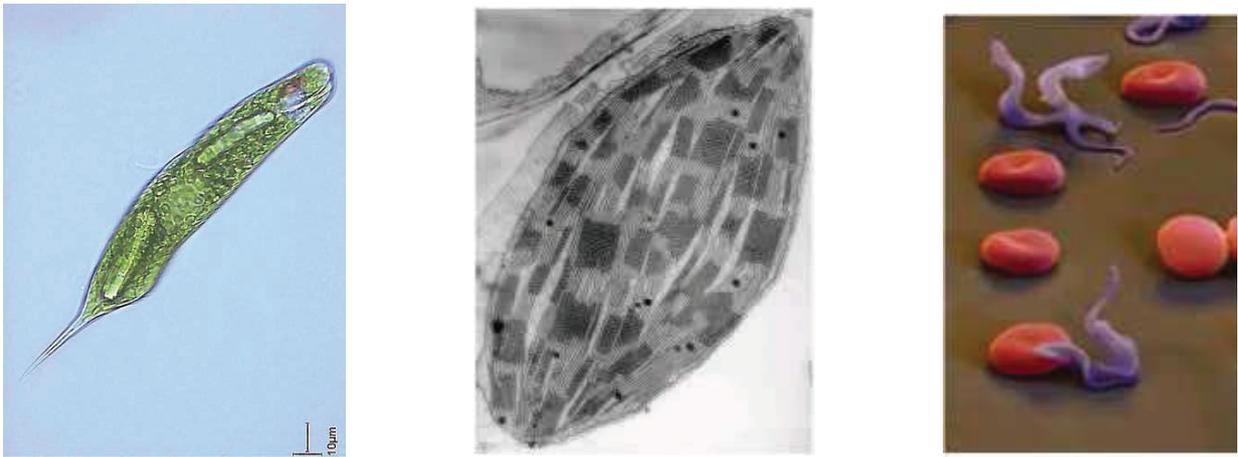
Figure 5: Images obtained by different types of microscopes [7]. Left: Optical microscope. Center: Transmission electron microscope. Right: Scanning electron microscope.

There is little published work on the rendering of transparent organisms taking into account lighting effects peculiar optical microscopy (such as diffraction and refraction) or scanning electron microscopy (such as electron scattering). Therefore, this problem need for a specialized methods of synthesis and/or processing and/or comparison of images that take into account these effects.

3. MODELING TECHNIQUES

The literature on the modeling of deformable biological structures is relatively modest [8,9] and that on modeling of microorganisms is virtually non-existent. Anyway, the adequate description of such structures and their deformations often requires ad hoc geometric modeling methods specific to each kind of organism. See Section 4.3.2.

3.1 MODELING BIOLOGICAL STRUCTURES

Our algorithm can be used with a wide variety of geometric modeling techniques. The models must be able to represent the approximate 'rest' shape of the organism and allow its normal deformations. For most recognition application the model need not be highly accurate or detailed. On the other hand the model must allow efficient rendering. Several techniques can be used for this purpose, such as triangular meshes, Bézier patches, constructive solid geometry (CSG), blobs, sweep solids, etc. See Figure 6.

Figure 6: Simple geometrical models of a red cell, using 8 Bézier patches (left), and a blob with three elements whose centers are indicated by the dots (right). Both models were rendered with POV-Ray [10].

3.2 MODELING THE DEFORMATIONS

There is an extensive literature on modelling 2D deformations [11–14]. A specially on the problem of aligning two-dimensional images taken with the same pose and view point — for example, images of consecutive sections of an organism.

However these 2D techniques are of little use in our problem. The literature on 3D deformation modeling is more limited.

A straightforward solution is to vary the parameters of the geometric model, such as the position of the control points of Bézier patches, centroids and radii of blobs elements, etc. However this solution is very expensive since typical models (such as triangle or Bézier meshes) may have hundreds of independent parameters.

A better solution is to choose a small set of *deformation modes*, where each mode consist of a simultaneous displacement and/or deformation change in all the shape's parameters, in various directions and magnitudes. The overall deformation is then defined by a small number of coefficients that specify the amount or strength of each mode. This technique is usually applied to triangular meshes, but can be adapted to other modeling techniques such as Bézier meshes and blobs.

Another way to model the deformation is to immerse the geometric model within a coarse grid of deformable elements, such as tricubic Bézier hexahedra. Each hexahedron can be understood as cube of elastic material, whose position and shape in space are controlled by a grid of 64 control points. Any deformation applied to this mesh is then interpreted as a deformation of 3D space, and transferred to the model which is immersed in it. This technique is very general (since it does not depend on the kind or complexity of the geometric model), and can produce fairly smooth deformations; but is far more expensive to render.

4. OUR ALGORITHM

We describe here an algorithm to identify an occurence of a specified organism in a given microscopy image. Our algorithm requires a *tridimensional deformable geometric model* of the organism to be recognized, i. e, a computer description of its schematic morphology and its allowed variations and deformations. The model has a set of *parameters* that determine the *pose* (position and orientation in space) of the model relative the microscope's camera, and the *deformation* (change of length and width, folding, displacement of parts, etc.) of the organism relative to its "rest" shape. The algorithm *align* that abstract 3D model with the microscope image; that is, it rotates, translates and deforms the model so as to obtain the best match between the expected and the observed projection of the organism on the camera sensor. The output of the algorithm is a set of values for these parameters that match the organism's position and appearance in the image. This basic algorithm can be adapted to identify an imaged organism among several models, or to find all occurrences of a organism in one or more images.

4.1 THE EXPONENTIAL COST BARRIER

Most research on the recognition of deformable models considers a relatively small range of deformations and rotations of the model, characterized by a small number of parameters. For the organisms considered in this article, the deformations may involve displacements comparable to the size of the object, and may have large number of degrees of freedom. On the other hand, the deformations always have limits, so that the morphology is usually recognizable in the images in spite of them.

In real situations, the position and deformation of the model can vary over a wide region of the space of $\mathcal{D}$ parameters, whose radius R is much larger than the required accuracy δ. A "brute force" solution would enumerate a grid of points in the region $\mathcal{D}$ with step δ. However, the number of such points is the order of $(R/\delta)^n$, where n is the number of parameters. For typical values of R, δ and n, this number is astronomical.

4.2 MULTISCALE ALIGNMENT

The exponential cost barrier can be overcome by *multi-scale matching* techniques. Multi-scale techniques use two or more versions of the entity involved — in our case, the digital images — each with a different level of detail. The problem at hand is solved first in the coarsest scale. This solution is then used as an initial guess to solve the problem at the next scale, where there is a bit more detail [15]. This process is repeated until reaching the original scale. See Figure 7. Since only small adjustments are considered at each stage, the total running time is greatly reduced in comparison with the brute-force approach.
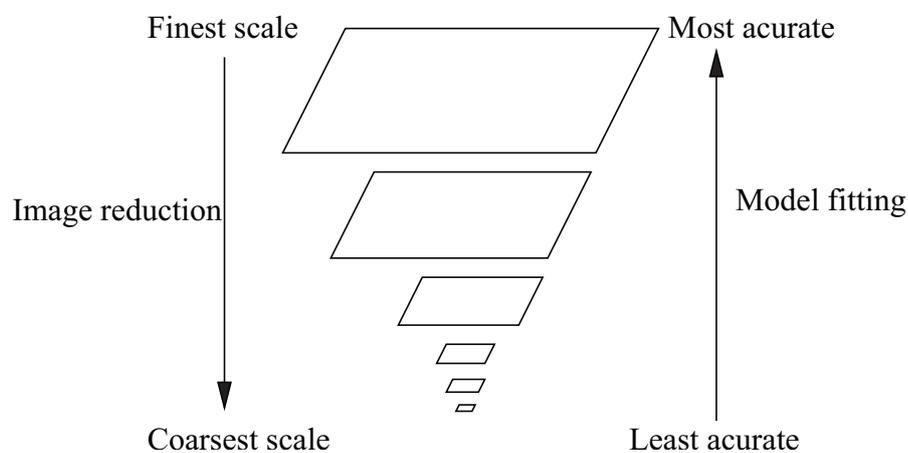
Figure 7: Schematic diagram of the multi-scale 3D-2D alignment.

More precisely, our algorithm begins by creating an initial list of *candidates*. Each candidate is a vector of parameters which specify a particular pose and deformation of the model. The initial candidate list $\mathcal{C}_m$ is prepared for the image at the coarset viable scale m, where the organism's image is expected to cover only a one or two pixels. At that scale, any spot on the image with that size is a possible instance of the target organism.

The algorithm then examines each successive scale in turn, from $m - 1$ down to 0. At each scale k, it scans the set of candidates $\mathcal{E}_{k+1}$ from the previous scale, and generates from them a new set of candidates $\mathcal{E}_k$. Each candidate c from $\mathcal{E}_{k+1}$ is examined to ascertain whether it is still compatible with this more detailed image at scale k. In this process, the parameters are re-adjusted to improve the accuracy of the match. The readjustment process often discards some candidates, and may generate new ones. At each step, the number of candidates is reduced by about one half; so that at the finest scale, there is only one surviving candidate — which is the result of the match.

To adjust the parameters, the algorithm uses a simple method reminiscent of the particle filter approach. Namely, the module replaces each candidate c by several perturbed candidates c' that differ only slightly from c. It then evaluates those perturbed candidates and selects the best ones as described in previous paragraph.

At each resolution scale, there may be some parameters which are not worth adjusting, because their effect is not observable at that scale. For instance, if the organism's projection is only one or two pixels wide, its position and size can be approximately determined; but its rotation can be estimated only at scales where the object is large enough for its overall shape to be visible. Parameters that affect only the position of small organelles or appendages are meaningful and relevant only at scales where these details are visible. Therefore, for each scale k our algorithm maintains a list of *active parameters* $\mathcal{P}_k$, which has only two or three parameters at scale m but grows to the full parameter set by the time it reaches scale 0.

4.3 MAIN LOOP

The core of our algorithm is the procedure that re-evaluates and refines the current candidates at each increased scale of resolution. It consists of four main modules: *candidate generation*, *image synthesis*, *image comparison* and *candidate pruning*.

4.3.1 IMAGE SYNTHESIS MODULE

The image synthesis module produces a synthetic image of the organism from its geometric model, given a particular set of pose and deformation parameters. The goal is to produce images that can be compared with real images.

For real optical microscopy images the organism must be modeled by transparent and possibly refractive materials. Ideally the rendering procedure should account for wave-related optical phenomena such as diffraction and interference. However, the physical simulation of light propagation as waves requires heavy calculations, and is probably impractical at present time.

In the case of images produced by scanning electron microscopes the organism can be considered opaque. The scattering of electrons obeys somewhat different laws from those used in computer graphics. However, since there is litle transmission, interference or diffraction, the image synthesis process is much simpler and largely compatible with the Newtonian particle model assumed in conventional ray-tracing.

For efficiency and flexibility reasons we have implemented a ray tracer that can handle several geometric primitives, such as spheres, cubes, cylinders and Bézier patches. This ray tracer acepts different types of surfaces finishes currently including diffuse, specular and transparent. By writing our tracer, rather than using POV-Ray or any other existing tracer, we were able to avoid all file I/O, and also to customize the shading model for our specific needs. Our tracer also generates a binary mask that defines the object's outline which is used by the comparison module.

4.3.2 IMAGE COMPARISON MODULE

The goal of the image comparison module is to assess how similar the synthetic microscopic image is to the actual image of the microorganism. A good match between the two images is evidence that the parameters of pose and deformation used in the synthesis are correct.

This module is closely linked to the image synthesis module, since the less sensitive the comparison metric is to the variations of illumination and diffraction effects, the less effort will be needed in the synthesis of microscopic images to simulate these effects. Thus, for example, we can eliminate many variations due to lighting and surface finish by smothing the microscope image and extracting its gradient. We use the normalized gradient operator

$$\blacktriangledown(I) = \frac{S_{3 \times 3}(I)}{\sqrt{V_{5 \times 5}(I) + \eta^2}} \tag{1}$$

where $S_{3 \times 3}$ is a Sobel filter with a 3×3 window, $V_{5 \times 5}$ is the local image variance computed using a 5×5 weight mask, and η is the assumed standard deviation of the imaging noise. This operator removes most of the image difference that is due to lighting and object colors, leaving mostly the object outlines. The Euclidean difference between the two normalized gradient images is therefore sensitive to the shape and position of the object. The comparison is limited to the region covered by the object's mask produced by image synthesis module. More precisely we calculate difference only inside the mask image and after that we divide by the mask's area. The difference between this normalized gradient and the ordinary gradient image can be seen in Figure 8.

Figure 8: Input image I (left) and output images of normalized gradient $|\blacktriangledown(I)|$ (center) and Sobel filter $|S_{3\times3}(I)|$ (right).

4.3.3 CANDIDATE GENERATION MODULE

The candidate generation module is given a list of candidates $\mathcal{E}_{k+1}$ that were obtained at the next coarser resolution level $k+1$, and generates from them a list $\mathcal{E}_k$ of *raw candidates* for the current level k. First, the module expands the active parameter list $\mathcal{P}_{k+1}$ by adding any parameter which was irrelevant at scale $k+1$ but is relevant at scale k; thus obtaining the current list $\mathcal{P}_k$. Then, for each candidate c in the list $\mathcal{E}_{k+1}$, and each parameter i in $\mathcal{P}_k$, the module determines the *current value* $c.p_i$, a *perturbation interval* $c.R_i$, and an *increment step* $c.\delta_i$. The current value $c.p_i$ is the value assigned to it in the previous scale in c if $i \in \mathcal{P}_{k-1}$, or a default value otherwise. The variables $c.R_i$ and $c.\delta_i$ may be specified by the user, or may be computed from the current parameter values stored in c in a model-specific way. Finally, the module adds to the list $\mathcal{E}_k$ one or more clones of c, by varying each parameter i over the range $c.p_i + c.R_i$ with step $c.\delta_i$, in all possible combinations. The process is started by assuming that at scale $m+1$ there are no active parameters ($\mathcal{P}_{m-1} = \{\}$), and there is only one empty candidate in the set $\mathcal{E}_{m+1} = \{()\}$.

4.3.4 CANDIDATE PRUNING MODULE

The candidate pruning method receives the list $\mathcal{E}_k$ of raw candidates and eliminates the worst ones, leaving at most a specified number N_k of candidates. For each candidate c in $\mathcal{E}_k$, the module renders the organism's model with the parameter values specified by c, and compares the resulting synthetic image with the given image at scale k. Let $c.d$ be the distance between the two images. The module then sorts the candidates by increasing $c.d$ field, and puts in $\mathcal{E}_k$ the first N_k candidates of their list, discarding the rest.

5. TESTS

We report here a few tests of our algorithm using one geometric object (a cube) and two simple organisms (*Volvox*, and *Caenorhabditis elegans*). Each organism was modeled in an abstract *object reference system* whose Z axis was assumed to be parallel to the optical axis of the microscope. In all tests, the goal was to determine the pose and deformation of that model in a set of given 64×64 pixels images.

All models have two *position parameters* X, Y which are the coordinates of the organism's reference point (centroid) on the image. These parameters are always measured in pixels of the scale 0 image. The parameters X, Y are always introduced at the coarserst level m, with initial range $[0, 64]$ and increments steps 2^m. At any other scale k, the increment step δ is set to 2^k (the current pixel size) and R is set to $[-\delta, +\delta]$. Thus, at each sucessive scale k the candidate generation module will pertub the position of each candidate by $+2^k$, 0 or -2^k in each coordinate, that is $+1$, 0 or -1 at that scale. The pruning at scale m leaves only those candidates which correspond to single-pixel spots on the scale m image.

All three models have also a *size parameters* S, which is defined as the logarithm base 2 of the organism's diameter. In all tests we are seeking ocurences of the organism with diameter between 4 and 8 pixels. At scale m the organism is assumed to be at most one pixel wide, so $m = 3$ and the parameter S is set to 1. At each sucessive scale k the step δ is set to 2^{k-m}, and the range is set to $[-\delta, +\delta]$.

Rotation parameters are expressed as arc lengths on a circle whose radius is the organism's approximate radius in directions perpendincular to the rotation axis. Thus, for instance, if the organism is approximately spherical the unit of angle is approximate $1/2^{S-1}$. The rotation range is determined by the model parameters.

5.1 CUBE

The cube model (Figure 9) is determined by the coordinates $(X, Y, 0)$ of the center, the size parameter S, and a 3D rotation matrix M. The parameter S is the logarithm in base 2 of the cube's diagonal. The matrix M is a the composition of three sucessive rotations about the X,Y, and Z axes, by specified angles θ_X, θ_Y and θ_Z respectively.
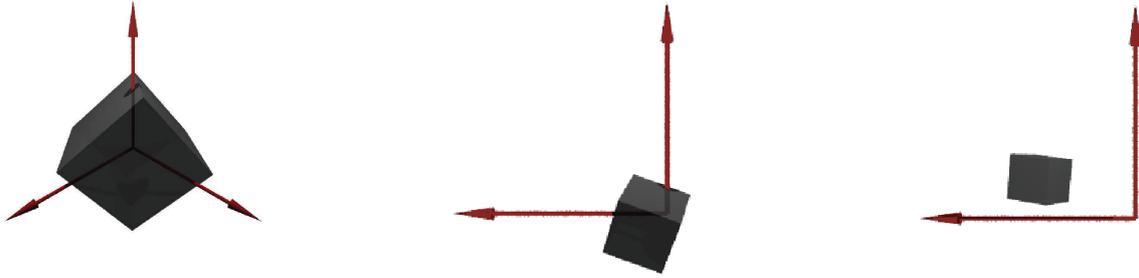
Figure 9: The cube model (left) and two images of it with diferents pose paramaters (center and right).

The input images were photos of cubical macroscopic objects against relatively simple backgrounds. The results are shown in Figure 10 and the Table 1. The parameters θ_X, θ_Y and θ_Z are introduced only in scale 1. The increment δ of θ_X, θ_Y, θ_Z is set to 22.5^o and range is $[-90^o, +90^o]$ for scale 1 and δ is set to 11.25^o and range $[-11.5^o, +11.5]$ in the finest scale.
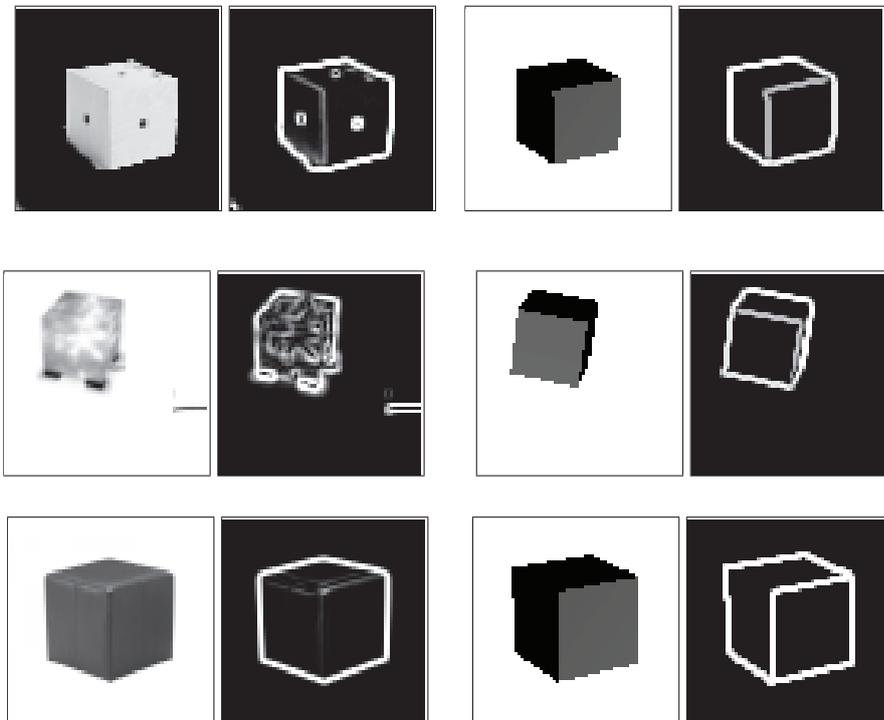

Figure 10: Left: Input images of three cubical objects and their normalized gradient images. Right: Synthetic images of the cube model fitted by our program and their normalized gradient images.

Table 1: Candidates for the cube test at each scale.

Scale	3	2	1	0
Raw candidates	64	540	39366	729
Adjusted candidates		4	2	1

5.2 *VOLVOX*

Our *Volvox* model was a hollow translucent sphere containing a variable number (zero or more) of smaller spheres. See Figure 11. Since the depth cordinate of the spheres cannot be easily determined from the images, we assumed that all spheres have their

centers on the plane $Z = 0$. The degrees of freedom are the radius r_0 and position $(X_0, Y_0, 0)$ of the main sphere, the number n of spheres, and their radii r_i and positions $(X_i, Y_i, 0)$, $i = 1, 2, 3, ..., n$.
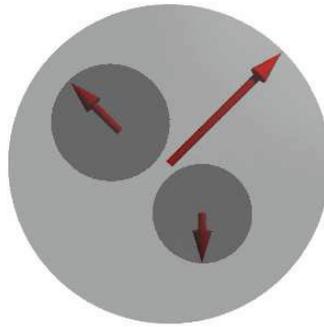

Figure 11: Image of a simple model of *Volvox*.

The input images were microscope photos of isolated *Volvox* specimens. The results are shown in Figure 12 and Table 2. The parameters X_i, Y_i and r_i are introduced in the scale 1. The increment δ of X_i and Y_i assume same values of δ of X and Y; but the range is determined by diference betwen radii of main sphere and internal sphere as set to $[-(r_0 - r_i), +(r_0 - r_i)]$. The δ of r_i have the same value of the size parameter r_0.
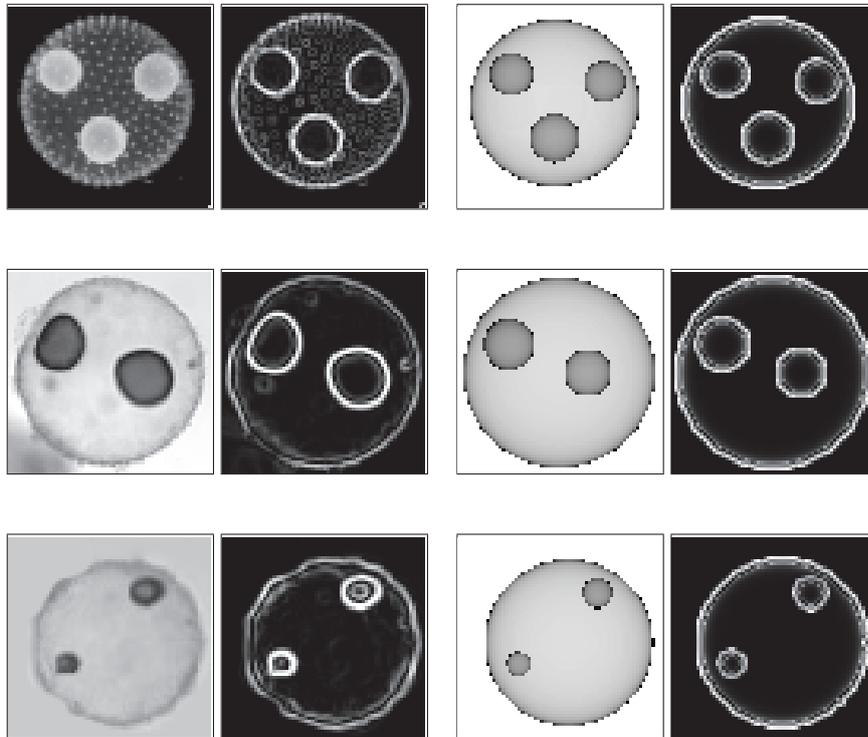

Figure 12: Left: Three input images of *Volvox* sp. and their normalized gradient images. Right: Synthetic images of the model as adjusted by our program, and their normalized gradient images.

Table 2: Candidates for the *Volvox* test at each scale.

Scale	3	2	1	0
Raw candidates	64	540	54000	2187
Adjusted candidates		4	2	1

5.3 *CAENORHABDITIS ELEGANS*

For *Caenorhabditis elegans*, we built a surface model with four bi-cubic Bézier patches with the control grid shown in Figure 13. Ten of those points are arbitrary, while the remainin 54 points are computed from theses 10 so that the surface is C_1 except at the ends.

Figure 13: Our *Caenorhabditis elegans* model. The Bézier control points are shown at left; the white points are free, and the black ones are computed from them.

Deformation of the *Caenorhabditis elegans* was modeled as a displacement $(X, Y, 0)$ of the barycenter, the rotation by some angle θ_Z, around the Z axis, and a global fattening, stretching and bending modification of the 10 free control points, controled by three parameters: the nominal diameter D of the body, its total length L, and the bending angle α. The efect of the three parameters is ilustrated in Figure 14.

$L = 2.0$
$D = 0.5$
$\alpha = 0^o$

$L = 2.2$
$D = 0.6$
$\alpha = -20^o$

$L = 2.0$
$D = 0.5$
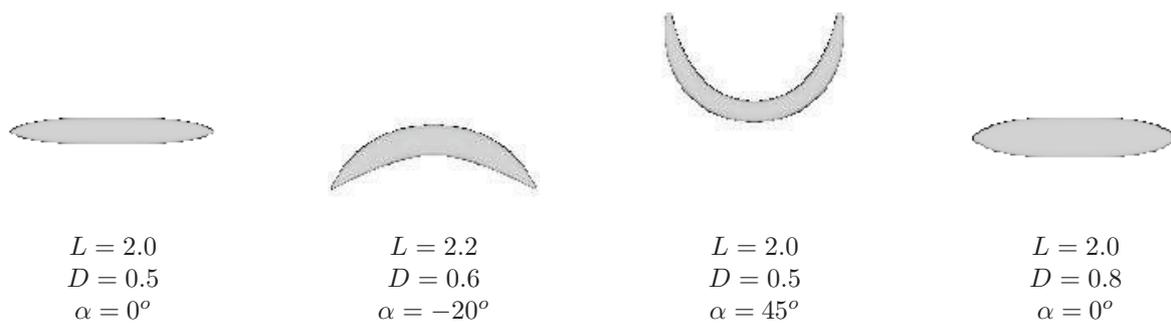$\alpha = 45^o$

$L = 2.0$
$D = 0.8$
$\alpha = 0^o$

Figure 14: Deformations of the *Caenorhabditis elegans* model.

The images were microscope photos of *Caenorhabditis elegans* sp. The results are shown in Figure 15 and Table 3. The parameters θ_Z, α are introduced only in scale 1. The increment δ of θ_Z and α is set to 22.5^o and range is $[0^o, +90^o]$ for scale 1 and δ is set to 11.25^o and range $[-11.5^o, +11.5]$ in the finest scale.

Table 3: Candidates for the *Caenorhabditis elegans* test at each scale.

Scale	3	2	1	0
Raw candidates	64	8100	12690	729
Adjusted candidates		4	2	1

6. CONCLUSIONS

The tests show that our algorithm can find a deformable organism model in a given image. Moreover, the method can be generalized to other models and other applications. In the future, we wish refine (increase the number of Bézier patches) the organisms models, thus the degress of freedom will be increased. Thereby we plan to use non-linear programation to fitting models, because can reduce number of candidates and improve the results.
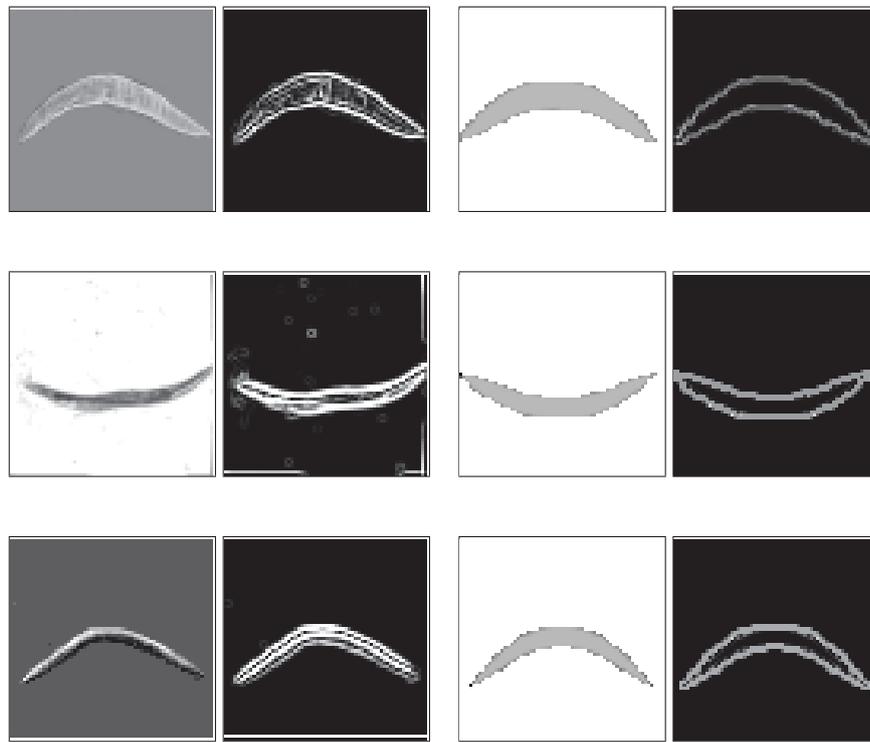
Figure 15: Left: Three input images of *Caenorhabditis elegans* and their normalized gradient images. Right: Synthetic images of the model as adjusted by our program and their normalized gradient images.

ACKNOWLEDGMENT

REFERENCES

[1] Image copied from website http://www.wikipedia.org.

[2] Image copied from website http://www.marksimmons.org/microscope/microimages/pages/paramecium2.htm.

[3] Image copied from website http://www.ncbi.nlm.nih.gov/projects/genome/guide/nematode/index.html.

[4] Image copied from website http://www.kiwicrossing.com/drrussell/projects.html.

[5] Image copied from website http://www.appslabs.com.au/daphnia.htm.

[6] Image copied from website http://botit.botany.wisc.edu/images/130/symbiosis_images/Paramecium_w_ Chlorella_MC.html.

[7] Image copied from website http://www.ufmt.br/bionet/conteudos/01.09.04/tip_mic.htm.

[8] T. Chen, X. Wang, D. N. Metaxas and L. Axel. "3D cardiac motion tracking using Robust Point Matching and meshless deformable models". In *ISBI*, pp. 280–283, 2008.

[9] T. Mcinerney, T. Mcinerney and D. Terzopoulos. "Deformable Models in Medical Image Analysis: A Survey", 1996.

[10] Image rendered by the software POV-Ray. Available at http://www.povray.org/.

[11] T. W. Sederberg and S. R. Parry. "Free-form deformation of solid geometric models". *SIGGRAPH Comput. Graph.*, vol. 20, no. 4, pp. 151–160, 1986.

[12] M. F. Karl, K. Rohr and H. S. Stiehl. "Elastic Registration of Medical Images Using Radial Basis Functions with Compact Support". In *In Proc. of CVPR'99*, pp. 402–407, 1999.

[13] S. Belongie, J. Malik and J. Puzicha. "Matching Shapes", 2001.

[14] H. Chui and A. Rangarajan. "A New Algorithm for Non-Rigid Point Matching". *Computer Vision and Pattern Recognition, IEEE Computer Society Conference on*, vol. 2, pp. 2044, 2000.

[15] H. C. da Gama Leitão and J. Stolfi. "A Multiscale Method for the Reassembly of Two-Dimensional Fragmented Objects". *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 24, no. 9, pp. 1239–1251, 2002.