Discrete deformable boundaries for the segmentation of multidimensional images

Jacques-Olivier Lachaud and Anne Vialard

LaBRI,* Université Bordeaux 1, 351 cours de la Libération, F-33405 Talence, France
(lachaud,vialard)@labri.u-bordeaux.fr

Abstract. Energy-minimizing techniques are an interesting approach to the segmentation problem. They extract image components by deforming a geometric model according to energy constraints. This paper proposes an extension to these works, which can segment arbitrarily complex image components in any dimension. The geometric model is a digital surface with which an energy is associated. The model grows inside the component to segment by following minimal energy paths. The segmentation result is obtained *a posteriori* by examining the energies of the successive model shapes. We validate our approach on several 2D images.

1 Introduction

A considerable amount of litterature is devoted to the problem of image segmentation (especially 2D image segmentation). Image components are determined either by examining image contours or by looking at homogeneous regions (and sometimes using both information). The segmentation problem cannot generally be tackled without adding to that information some *a priori* knowledge on image components, e.g., geometric models, smoothness constraints, reference shapes, training sets, user interaction. This paper deals with the segmentation problem for arbitrary dimensional images. We are interested in methods extracting an image component by deforming a geometric model. The following paragraphs present classical techniques addressing this issue.

Energy-minimizing techniques [4] have proven to be a powerful tool in this context. They are based on an iterative adaptation process, which locally deforms a parametric model. The model/image adequation is expressed as an energy, which is minimal when the model geometry corresponds to image contours. The continuity of the geometric model and tunable smoothness constraints provide a robust way to extract image components, even in noisy images. The adaptation process is sensitive to initialization since it makes the model converge on local minima within the image. The parametric definition of the model also restricts its topology to simple objects. Recent works now propose automated topology adaptation techniques to overcome this issue, both in 2D [10] and in 3D [6]. However, these techniques are difficult to extend to arbitrary dimensions.

* Laboratoire Bordelais de Recherche en Informatique - UMR 5800

Front propagation techniques have been proposed to avoid the topology restriction induced by the model parameterization. Instead of deforming a geometric model in the image, they assign a scalar value to each point of the image space. The evolution of the points value is governed by partial differential equations, similar to the heat diffusion equation. The model is then implicitly defined as a level-set of this space, which is called a front. The equations are designed to make the front slow down on strong contours and to minimize its perimeter (or area in 3D) [9]. The implicit definition of the front ensures natural topology changes. However, this technique is not designed to integrate *a priori* knowledge on the image component (e.g., other geometric criteria, reference shape).

In region growing methods [1], the extraction of an image component follows two steps: (i) seeds are put within the component of interest and (ii) these seeds grow by iteratively adding pixels to them according to a merging predicate (homogeneity, simple geometric criterion). These methods are interesting because on one hand they have a simple dimension independent formulation and on the other hand they can segment objects of arbitrary topology. However, they are not well adapted to the extraction of inhomogeneous components.

This paper proposes an original approach based on a discrete geometric model that follows an energy-minimizing process. The discrete geometric model is the *digital boundary* of an object growing within the image. The model energy is distributed over all the boundary elements (i.e., the *surfels*). The energy of each element depends on both the local shape of the boundary and the surrounding image values. The number of possible shapes within an image grows exponentially with its size. Therefore, the following heuristic is used to extract components in an acceptable time. The model is initialized as an object inside the component of interest. At each iteration, a set of connected elements (i.e., a *voxel patch*) is locally glued to the model shape. The size and position of this set are chosen so that the object boundary energy be minimized. This expansion strategy associated with proper energy definitions induce the following model behavior: strong image contours forms "wells" in the energy that hold the model, the growth is at the same time penalized in directions increasing the area and local curvature of the object boundary, the model grows without constraints elsewhere.

This model casts the energy-minimizing principle in a discrete framework. Significant advantages are thus obtained: reduced sensibility to initialization, modeling of arbitrary objects, arbitrary image dimension. The paper is organized as follows. Section 2 recalls some necessary digital topology definitions and properties. Section 3 defines the model geometry and its energy. Section 4 presents the segmentation algorithm. Segmentation results on 2D images are presented and discussed in Section 5.

2 Preliminary definitions

A *voxel* is an element of the discrete n-dimensional space $\mathbb{Z}^n$, for $n \geq 2$. Some authors [11] use the term "spel" for a voxel in an n-dimensional space; since we

feel that no confusion should arise, we keep the term "voxel" for any dimension. Let M be a finite "digital parallelepiped" in $\mathbb{Z}^n$. An *image* I on $\mathbb{Z}^n$ is a tuple (M, f) where f is a mapping from the subset M of $\mathbb{Z}^n$, called the *domain* of I, toward a set of numbers, called the *range* of I. The *value* of a voxel $u \in M$ in the image I is the number $f(u)$. An *object* is any nonempty subset of the domain M. The *complement* of the object O in M is denoted by O^c.

Let ω_n be the adjacency relation on $\mathbb{Z}^n$ such that $\omega_n(u, v)$ is true when u and v differ of ± 1 on exactly one coordinate. Let α_n be the adjacency relation such that $\alpha_n(u, v)$ is true when $u \neq v$, and u and v may differ of either -1, 0, or 1 on any one of their coordinates. If ρ is any adjacency relation, a *ρ-path* from a voxel v to a voxel w on a voxel set A is a sequence $u_0 = v, \ldots, u_m = w$ of voxels of A such that, for any $0 \leq i < m$, u_i is ρ-adjacent to u_{i+1}. Its *length* is $m + 1$.

For any voxels u and v with $\omega_n(u, v)$, we call the ordered pair (u, v) a *surfel* (for "surface element" [11]). Any nonempty set of surfels is called a *digital surface*. For any given digital surface Σ, the set of voxels $\{v \mid (u, v) \in \Sigma\}$ is called the *immediate exterior* of Σ and is denoted by $\mathrm{IE}(\Sigma)$. The *boundary* ∂O of an object O is defined as the set $\{(u, v) \mid \omega_n(u, v)$ and $u \in O$ and $v \in O^c\}$.

Up to now, an object boundary is just viewed as a set. It is convenient to have a notion of surfel neighbors (i.e., a "topology") in order to define connected zones on an object boundary or to determine an object by tracking its boundary. In our case, this notion is compulsory to define a coherent model evolution. Besides, defining an object through its boundary is often faster.

Defining an adjacency between surfels is not as straightforward as defining an adjacency between voxels (especially for $n \geq 3$). The problem lies in the fact that object boundary components (through surfel adjacencies) must separate object components from background components (through voxel adjacencies). In this paper, we do not focus on building surfel adjacencies consistent with a given voxel adjacency. Consequently, given an object O considered with a voxel adjacency ρ, with either $\rho = \omega_n$ or $\rho = \alpha_n$, we will admit that it is possible to locally define a consistent *surfel adjacency* relation, denoted by β_O, for the elements of ∂O (3D case, see [3]; nD case, Theorem 34 of Ref. [7]). For the 2D case, Figure 2 shows how to locally define a surfel adjacency on a boundary.

The β_O-adjacency relation induces *β_O-components* on ∂O. *β_O-paths* on ∂O can be defined analogously to ρ-paths on O. The *length* of a β_O-path is similarly defined. The *β_O-distance* between two surfels of ∂O is defined as the length of the shortest β_O-path between these two surfels. The *β_O-ball* of size r around a surfel σ is the set of surfels of ∂O which are at a β_O-distance lesser or equal to r from the surfel σ. Let Σ be a subset of ∂O. We define the *border* $B(\Sigma)$ of Σ on ∂O as the set of surfels of Σ that have at least one β_O-neighbor in $\partial O \setminus \Sigma$. The *$k$-border* $B_k(\Sigma)$ of Σ on ∂O, $1 \leq k$, is the set of surfels of Σ which have a β_O-distance lesser to k from a surfel of $B(\Sigma)$.

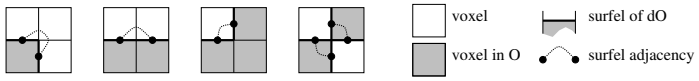

Fig. 1. Local computation of surfel adjacencies in the 2D case. The 4-adjacency (resp. 8-adjacency) has been chosen for the object elements (resp. background elements).

3 The discrete deformable boundaries model

For the purpose of image segmentation, we introduce a geometric model to fit image components. The model *geometry* (i.e., its *shape*) is defined as an object in the image. It is equivalently defined as the boundary of this object. Note that the model is not necessarily connected. The model geometry is aimed to evolve from an initial shape toward an image component boundary. This evolution depends on an energy associated to each possible geometry of the model.

We identify the energy of an object to the energy of its boundary. The *energy* of an object boundary ∂O, denoted by $E(\partial O)$, is computed by summation of the energies of each of its surfels. To get finer estimates, the energy of a surfel may depend on a small neighborhood around it on the object boundary. That is why we use the notation $E(\sigma)_{\partial O}$ to designate the energy of the surfel σ with respect to ∂O (σ must be an element of ∂O). By definition, we set $E(\Sigma)_{\partial O} = \sum_{\sigma \in \Sigma} E(\sigma)_{\partial O}$, where Σ is any nonempty subset of an object boundary ∂O.

The surfel energy is the sum of several energy terms. Two types of surfel energies are distinguished: the surfel energies that only depend on the local geometry of the boundary around the surfel are called *internal energies*, the surfel energies that also depend on external parameters (e.g., local image values) are called *external energies*. The local geometric characteristics required for the surfel energy computation are based upon a neighborhood of the surfel on the object boundary: it is a β_O-ball on the object boundary, centered on the surfel. To simplify notations, we assume that this ball has the same size p for the computation of every surfel energy. The whole surfel energy is thus locally computed.

Internal energies allow a finer control on the model shape (e.g., smoothness). External energies express the image/model adequation or other external constraints (e.g., similarity to a reference shape). The following paragraphs present a set of internal and external energies pertinent to our segmentation purpose. This set is by no way restrictive. New energies can be specifically designed for a particular application to the extent that the summation property is satisfied.

Image features are generally not sufficient to clearly define objects: the boundary can be poorly contrasted or even incomplete. To tackle this problem, we use the fact that the shape which most likely matches an image component is generally "smooth". In our case, this is expressed by defining two internal energies. The *stretching energy* $E^s(\sigma)_{\partial O}$ of a surfel σ is defined as an increasing function of the area of the surfel σ. The *bending energy* $E^b(\sigma)_{\partial O}$ of a surfel σ is defined as an increasing function of the mean curvature of the surfel σ. Examples of area and mean curvature computations are given in Section 5. Note that these ener-

gies correspond to the internal energies of many deformable models [4], which regularize the segmentation problem.

In our context, we define a unique external energy, based on the image value information. Since the model evolution is guided by energy minimization, the *image energy* $E^I(\sigma)_{\partial O}$ of a surfel σ should be very low when σ is located on a strong contour. A simple way to define the image energy at a surfel $\sigma = (u, v)$ on an object boundary is to use the image gradient: $E^I(\sigma)_{\partial O} = -\|f(v) - f(u)\|^2$, if $I = (M, f)$. This definition is valid for arbitrary n.

The model grows by minimizing its energy at each step. Since the growing is incremental, an incremental computation of the energy would be pertinent. More precisely, the problem is to compute the energy of an object O' given the energy of an object O included in O'. In our case, the digital surfaces ∂O and $\partial O'$ generally have much more common surfels than uncommon surfels (as the model is growing, this assertion is more and more true). The set of the common surfels is denoted by Φ. The p-border of Φ on ∂O is identical to the p-border of Φ on $\partial O'$. We can thus denote it uniquely by $B_p(\Phi)$. The energy of ∂O and $\partial O'$ is expressed by the two following equations:

$$E(\partial O) = \sum_{\sigma \in \partial O} E(\sigma)_{\partial O} = E(\Phi \setminus B_p(\Phi))_{\partial O} + E(B_p(\Phi))_{\partial O} + E(\partial O \setminus \Phi)_{\partial O},$$

$$E(\partial O') = \sum_{\sigma \in \partial O'} E(\sigma)_{\partial O'} = E(\Phi \setminus B_p(\Phi))_{\partial O'} + E(B_p(\Phi))_{\partial O'} + E(\partial O' \setminus \Phi)_{\partial O'}.$$

From the surfel energy computation, it is easy to see that the surfels of $\Phi \setminus B_p(\Phi)$, common to both ∂O and $\partial O'$, hold the same energy on ∂O and on $\partial O'$. However, the energy of the surfels of $B_p(\Phi)$ may (slightly) differ whether they are considered on ∂O or on $\partial O'$. We deduce

$$\underbrace{E(\partial O') - E(\partial O)}_{\text{variation}} = \underbrace{E(\partial O' \setminus \Phi)_{\partial O'}}_{\text{created surfels}} - \underbrace{E(\partial O \setminus \Phi)_{\partial O}}_{\text{deleted surfels}} + \underbrace{E(B_p(\Phi))_{\partial O'} - E(B_p(\Phi))_{\partial O}}_{\text{surfels close to displacement}}.$$

To get efficient energy computations at each step, each surfel of the model stores its energy. When a model grows from a shape O to a shape O', the energy of only a limited amount of surfels will have to be computed: (i) the energy of created surfels and (ii) the energy of the surfels nearby those surfels.

4 Segmentation algorithm

In the energy-minimizing framework, the segmentation problem is translated into the minimization of a cost function in the space of all possible shapes. Finding the minimum of this function cannot be done directly in a reasonable time. Except for very specific problems (e.g., what is the best contour between two known endpoints), heuristics are proposed to extract "acceptable" solutions. For the snake, an "acceptable" solution is a local minimum. We propose a heuristic that builds a set of successive shapes likely to correspond to image components.

We first briefly outline the segmentation process. The model is initialized as a set of voxels located inside the object to be segmented. At each step, a voxel patch is added to the model. A *voxel patch* of radius k around a surfel σ on the boundary ∂O is the immediate exterior of the β_O-ball of size k around σ. To decide where a voxel patch is "glued" to O, its possible various locations are enumerated. Among the possible resulting shapes, the one with the smallest energy is chosen. Unlike most segmentation algorithms, this process does not converge on the expected image component. However, the state of the model at one step of its evolution is likely to correspond to the expected image component. The boundary of the object of interest is hence determined *a posteriori*. This technique is similar to the discrete bubble principle [2].

It is then possible to let the user choose the shape pertinent to his problem among the successive states of the model. A more automated approach can also be taken. Since the model growing follows a minimal energy path (among all possible shapes within the image), the image/model adequation can be estimated through the model energy at each step. Consequently, the shape of minimal energy often delineates an object which is a pertinent component of the image.

For now, k is a given strictly positive integer number. It corresponds to the radius of the voxel patch that is added on the model at each step. The following process governs the growing evolution of the model and assigns an energy to each successive model state (see Fig. 2):

1. Assign 0 to i. Let O_0 be equal to the initial shape. The initial shape is a subset of the image domain included in the component(s) to extract. Let E_0 be equal to $E(\partial O_0)$.
2. For all surfels $\sigma \in \partial O_i$, perform the following steps:
   (a) Extract the β_{O_i}-ball V_σ of radius k around σ. Define O_σ as $O_i \cup \mathrm{IE}(V_\sigma)$.
   (b) Incrementally compute $E(\partial O_\sigma)$ from $E(\partial O_i)$.
3. Select a surfel τ with minimal energy $E(\partial O_\tau)$.
4. Let O_{i+1} be equal to O_τ. Let E_{i+1} be equal to $E(\partial O_\tau)$. Increment i.
5. Go back to step 2 until an end condition is reached (e.g., $O = M$, user interaction, automated minimum detection).

In the experiments described in Section 5, the end condition is $O = M$, which corresponds to a complete model expansion.

5 2D experiments

In order to validate this segmentation approach, a 2D prototype has been implemented. The experiments emphasize the ability of our model to segment image components in various contexts: incomplete or weakly contrasted contours, inhomogeneous components. In 2D, voxels correspond to pixels and surfels are usually called pixel edges. We first consider how energies are computed and weighted. We then highlight the model capabilities on both synthetic and medical images.

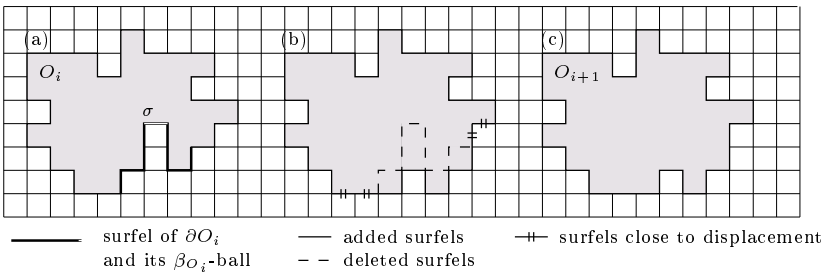

Fig. 2. A step in the model growing. (a) σ is a surfel of the model boundary at step i. (b) The voxel patch of radius 4 around σ is added to the object. (c) Model shape at step $i+1$ if σ is chosen as the best place to make the model grow.

5.1 Energy computation

For our experiments, the energies for a surfel $\sigma = (u, v)$ are computed as follows:

- The stretching energy $E^s(\sigma)_{\partial O}$ is the contribution of the surfel σ to the model perimeter. The boundary ∂O can be viewed as one or several 4-connected paths. We adapt the Rosen-Profitt estimator to measure the contribution of σ to the boundary length. On the boundary path, σ has two endpoints p_1 and p_2. These points can either be "corner" points or "non-corner" points. We set $E^s(\sigma)_{\partial O} = (\psi(p_1), \psi(p_2))/2$, where $\psi(p) = \psi_c = 0.670$ if p is a corner point and $\psi(p) = \psi_{nc} = 0.948$ otherwise. Note that other perimeter estimators could be used (e.g., see [5]).
- The bending energy $E^b(\sigma)_{\partial O}$ is computed as the ratio of the two distances l and D, defined as follows. The β_O-ball of size p around σ forms a subpath of ∂O, which has two endpoints e_1 and e_2. Its length l is computed by the same Rosen-Profitt estimator as above. The distance D is the Euclidean distance of e_1 and e_2. Many curvature estimation methods could be implemented (angular measurement, planar deviation, tangent plane variation, etc. [12]).
- The image energy $E^I(\sigma)_{\partial O}$ is defined as in Section 3 as $-\|f(v) - f(u)\|^2$. It is the only external energy used in the presented experiments.

The energy of a surfel σ on the boundary ∂O is the weighted summation of the above-defined energies:

$$E(\sigma)_{\partial O} = \alpha_s E^s(\sigma)_{\partial O} + \alpha_b E^b(\sigma)_{\partial O} + \alpha_I E^I(\sigma)_{\partial O},$$

where α_s, α_b and α_I are positive real numbers whose sum is one. To handle comparable terms, internal energies are normalized to $[0, 1]$. The image energy is normalized to $[-0.5, 0.5]$ to keep in balance two opposite behaviors: (i) should the image energy be positive, the shape of minimal energy would be empty, (ii) should the image energy be negative, long and sinuous shape would be favored. Note that most classical deformable models choose a negative image energy function. These coefficients allow us to tune more precisely the model behavior

on various kinds of images. A set of coefficients pertinent to an image will be well adapted to similar images.

5.2 Results

In all presented experiments, we consider the model with the 4-connectedness. The surfel adjacency β_O is therefore defined as shown on Fig. 2. The parameter p is set to 4.

Since our model searches for contours in images, inhomogeneous components can efficiently be segmented. We illustrate this ability on the test image of Fig. 3a. This test image raises another segmentation issue: contours are weak or even inexistant on some locations both between the disk and the ring and between the disk and the background. Fig. 3b emphasizes three steps in the model evolution: the initial shape, the iteration when the model lies on the disk–ring boundary (a local minimum in the energy function), the iteration when the model lies on the ring–background boundary (the minimum of the energy function). Fig. 3c displays the energy function. At each iteration, several patch radius sizes (i.e., k) have been tested for each surfel: k was between 0 and 6. The model successfully delineates the two boundaries during its evolution, although the first boundary induces a less significant minimum in the energy function than the second one: the first boundary is indeed not as well defined as the second one.

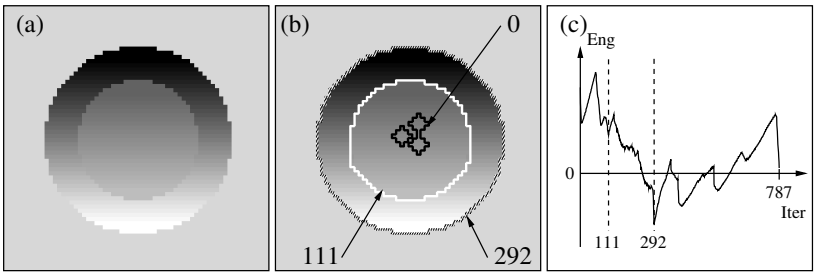

Fig. 3. Inhomogeneous components segmentation. (a) Test image: a circle filled with a shading from gray value 98 (top) to gray value 189 (bottom), a ring encircling it filled with a sharper shading from 0 (top) to 255 (bottom), a homogeneous background of gray value 215. (b) Three significant steps in the model evolution: initial shape, disk–ring boundary, ring–background boundary. (c) Energy curve for this evolution: the two extracted boundaries correspond to local minima of the energy function.

The second experiment illustrates the robustness of the segmentation process compared to the initial shape. The test image is a MR image[1] of a human heart at diastole (Fig. 4a). Our objective is to segment the right ventricle. This image

[1] Acknowledgements to Pr. Ducassou and Pr. Barat, Service de Médecine Nucléaire Hôpital du Haut Levêque, Bordeaux, France.

component is inhomogeneous in its lower part and presents weak contours on its bottom side. The other contours are more distinct but are somewhat fuzzy. All these defects can be apprehended on Fig. 4b-c, which show the image after edge detection. The middle row (Fig. 4d-f) presents three evolutions, one per column, with three different initial shapes. Each image depicts three or four different boundaries corresponding to significant steps in the model evolution. The bottom row (Fig. 4g-i) displays the corresponding energy curve. For this experiment, only the patch radius size 3 is tested for each surfel (i.e., $k = 3$). Whichever is the initial shape, the model succeeds in delineating the right ventricle. The left ventricle may also be delineated in a second stage (near the end of the expansion), but it is more hazardous: the proposed initial shapes are indeed extremely bad for a left ventricle segmentation.

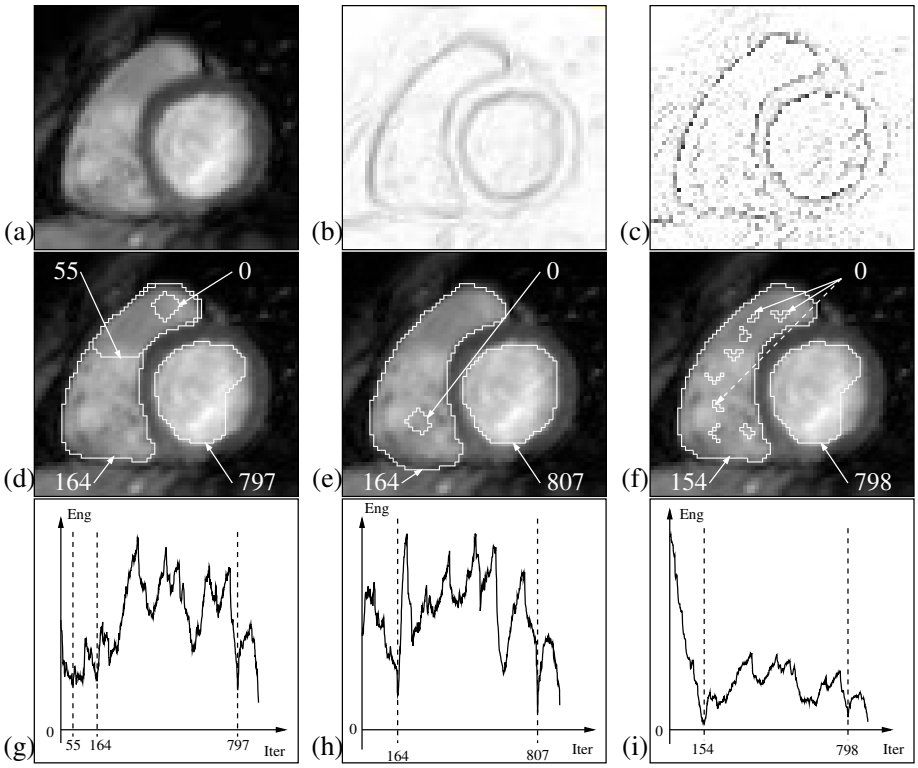

Fig. 4. Robustness of the segmentation to initialization. (a) Test image: a MR image of a human heart at diastole. (b) Image after Sobel filtering. (c) Image after Laplace edge detection. The bottom two rows depict the model behavior for three different initial shapes (energy parameters are set to $\alpha_s = 0.25$, $\alpha_b = 0.25$, $\alpha_I = 1$). The middle row (d-f) shows significant steps in the model evolution (initialization, important local minima). The corresponding energy curves are drawn on the bottom row figures (g-i).

Our prototype does not include all the optimizations which could be implemented for the 2D case (e.g., efficient traversal of surfel adjacency graphs, various precomputations). For the heart image, whose size is 68×63, the complete model evolution takes 349s. The right ventricle is detected after 40s.

6 Conclusion

We have presented a discrete deformable model for segmenting image components. The segmenting process is carried out by expanding a digital surface within the image under internal and external constraints. The external constraint (i.e., the image energy) stops the model expansion on strong contours. At the same time, the internal constraints regularize the model shape. The robust framework of energy-minimizing techniques is thus preserved. The significant advantages of this model are its dimension independence and its ability to naturally change topology. The first results on both synthetic and real-life data are very promising. They underline the model abilities to process images with poor contours and inhomogeneous components. Moreover, our segmentation technique is less sensitive to initialization than classical energy-minimizing techniques. Further information can be found in [8].

References

1. J.-P. Cocquerez and S. Philipp. *Analyse d'images : filtrage et segmentation*. Masson, Paris, Milan, Barcelone, 1995.
2. Y. Elomary. *Modèles déformables et multirésolution pour la détection de contours en traitement d'images*. PhD thesis, Univ. Joseph Fourier, Grenoble, France, 1994.
3. G.T. Herman and D. Webster. A topological proof of a surface tracking algorithm. *Computer Vision, Graphics, and Image Processing*, 23:162–177, 1983.
4. M. Kass, A. Witkin, and D. Terzopoulos. Snakes: Active contour models. *International Journal of Computer Vision*, 1(4):321–331, 1987.
5. J. Koplowitz and A. M. Bruckstein. Design of perimeter estimators for digitized planar shapes. *IEEE PAMI*, 11(6):611–622, 1989.
6. J.-O. Lachaud and A. Montanvert. Deformable meshes with automated topology changes for coarse-to-fine 3D surface extraction. *Medical Image Analysis*, 3(2):187–207, 1999.
7. J.-O. Lachaud and A. Montanvert. Continuous analogs of digital boundaries: A topological approach to iso-surfaces. *Graphical Models and Image Processing*, 62:129–164, 2000.
8. J.-O. Lachaud and A. Vialard. Discrete deformable boundaries for image segmentation. Research Report 1244-00, LaBRI, Talence, France, 2000.
9. R. Malladi, J. A. Sethian, and B. C. Vemuri. Shape Modelling with Front Propagation: A Level Set Approach. *IEEE PAMI*, 17(2):158–174, 1995.
10. T. McInerney and D. Terzopoulos. Medical Image Segmentation Using Topologically Adaptable Surfaces. In *Proc. of CVRMed-MRCAS, Grenoble, France*, volume 1205 of *LNCS*, pages 23–32. Springer-Verlag, 1997.
11. J. K. Udupa. Multidimensional Digital Boundaries. *CVGIP: Graphical Models and Image Processing*, 56(4):311–323, July 1994.

12. M. Worring and A. Smeulders. Digital curvature estimation. *Computer Vision, Graphics and Image Processing: Image Understanding*, 58(3):366–382, 1993.