International Journal of Computer Vision

DOI: 10.1007/s11263-006-6850-z

Fast Constrained Surface Extraction by Minimal Paths

ROBERTO ARDON
MEDYSIS, Philips France, 51, rue Carnot, 92120 Suresnes, France; CEREMADE, UMR 7534 Université Paris-Dauphine, Place du Marechal de Lattre de Tassigny, 75775 Paris cedex 16, France
roberto.ardon@centraliens.net

LAURENT D. COHEN
CEREMADE, UMR 7534 Université Paris-Dauphine, Place du Marechal de Lattre de Tassigny, 75775 Paris cedex 16, France
cohen@ceremade.dauphine.fr

Received April 7, 2004; Revised November 12, 2004; Accepted November 23, 2004

First online version published in xxx

Abstract. In this paper we consider a new approach for single object segmentation in 3D images. Our method improves the classical geodesic active surface model. It greatly simplifies the model initialization and naturally avoids local minima by incorporating user extra information into the segmentation process. The initialization procedure is reduced to introducing 3D curves into the image. These curves are supposed to belong to the surface to extract and thus, also constitute user given information. Hence, our model finds a surface that has these curves as boundary conditions and that minimizes the integral of a potential function that corresponds to the image features. Our goal is achieved by using globally minimal paths. We approximate the surface to extract by a discrete network of paths. Furthermore, an interpolation method is used to build a mesh or an implicit representation based on the information retrieved from the network of paths. Our paper describes a fast construction obtained by exploiting the Fast Marching algorithm and a fast analytical interpolation method. Moreover, a Level set method can be used to refine the segmentation when higher accuracy is required. The algorithm has been successfully applied to 3D medical images and synthetic images.

Keywords: active surfaces, active contours, minimal paths, level set method, object extraction

1. Introduction

Since their introduction by Kass et al. (1998) deformable models have been extensively used to find single and multiple objects in 2D and 3D images. The common use of these models consists in introducing an initial object in the image and transforming it until it reaches a wanted target. In most applications, the evolution of the object is done in order to minimize an energy attached to the image data, until a steady state is reached. One of the main drawbacks of this approach is that it suffers from local minima 'traps'. This happens when the steady state, reached by the active object, does not correspond to the target but to another local minimum of the energy. Thus, the active object initialization is a fundamental step, if it is too far from the target, local minima can block the active object evolution, and the target is never reached. On the other hand, when image quality is very low, the information contained in any energy derived from the image, may not lead to the desired segmentation. The model should then be able to take into account additional information given by the user.

P1: NVI
International Journal of Computer Vision SJNW694-04-NO00006850 February 10, 2006 9:36

Since the publication of Kass et al. (1998), much work has been done in order to free active models from the problem of local minima. A balloon force was early proposed in Cohen (1991) to make the model more active and to cope with the shrinking problem, but this force supposed a known direction in the evolution. The introduction of region dependent energies (Paragios, 2000; Cohen, 1997) and the use of shape priors approaches (Yuille et al., 1992; Cremers and Schnörr 2003; Tsai et al., 2003), contributed to create a more robust framework. Nonetheless, when looking for a precise object (like the left ventricle in 3D ultrasound images) if the initialization of the model is made by simple geometric objects (spheres, cylinders), too far from the targeted shape, most of the present models will fail. Tedious hand drawing initializations are thus often needed. In this work, we focus on a novel approach for 3D single object segmentation having a cylinder-like topology. Our contribution consists in exploiting two curves, introduced in the image by the user, in order to segment the object by a first approximation of a minimal energy surface that avoids unwanted local minima. The given curves are supposed to be drawn on the surface of the object to be segmented. They constitute the initialization of the 3D model, and the information they provide (for being drawn on the object to extract) is highly exploited, since the surface our algorithm generates is constrained to contain them. In order to avoid local minima 'traps', our algorithm builds a network of globally minimal paths, then a surface is interpolated by a novel analytical interpolation method we have developed. As an illustration of the situation we are working on, we give, in Fig. 1, an example of the user input to our algorithm for the segmentation of a 3D ultrasound volume of the left ventricle.

The outline of our paper is as follows: we begin in Section 2 by recalling the principles of geodesic active contours and surfaces as well as the global minimal paths framework. In Section 3 we explain how minimal paths can be used to build a network of paths that discretely approximates the surface to be segmented and that is not sensitive to the problem of local minima traps. In Section 4 we give the final step of our algorithm which is the generation of the surface from the network of paths. At last, in Section 5 we show some examples on synthetic data and real medical images.

2. Active Surfaces and Minimal Paths

2.1. Evolution Equations

Active surfaces as well as minimal paths resulted from deformable models introduced with the snakes model (kass et al., 1988). This model consisted in introducing a curve g into the image and making it evolve in order to minimize the energy,

$$E(g) = \int \alpha.\|g'(s)\|^2 + \beta.\|g''(s)\|^2 + \mathcal{P}(g(s))ds.$$

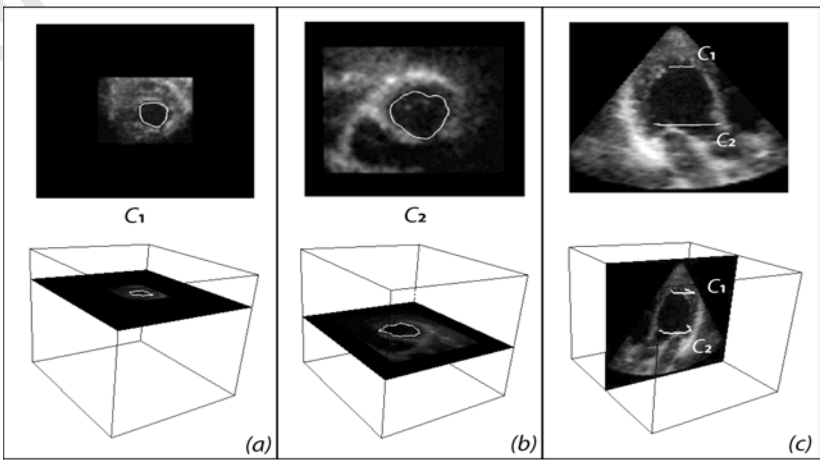

Figure 1. Three different slices of a 3D ultrasound volume of a left ventricle and the two user given curves $\mathcal{C}_1$ and $\mathcal{C}_2$. (a) and (b) show the two parallel slices where the curves are drawn. (c) shows a perpendicular slice to the curves in order to show their position with respect to the ventricle.

The two first terms maintained the regularity of the curve and the last one was the data attachment term. The potential function $\mathcal{P}$, usually represented an edge detector that had lower values on edges. For example $\mathcal{P} = (1 + |\nabla I|^2)^{-1}$ if I is the image.

Caselles et al. improved the energy formulation in Caselles et al. (1997a,b) by introducing the geodesic active contour model and its surface extension. In their approach the evolution of an initial curve g_0 or surface $\mathcal{S}_0$ was driven by the minimization of the geodesic energies

$$E(g) = \int \mathcal{P}(g(s)l) \|g'(s)\| ds \quad \text{and}$$
$$E(\mathcal{S}) = \int\int \mathcal{P}(s(u,v)) \|\mathcal{S}_u \times \mathcal{S}_v\| du dv \qquad (1)$$

Hence, their model is geometrical, since it is no longer dependent on parameterization. Even though these models are only edge-driven, most of current approaches that integrate other information (region, texture, shape knowledge) are actually extensions. The most popular approach for solving the minimization problems (1) is to consider Euler-Lagrange equations (first variation of the energy) and derive from them the corresponding descent schemes:

$$\left\{ \frac{\partial g}{\partial t} = (\mathcal{P}\kappa - \nabla\mathcal{P}.\vec{n})\vec{n}, \quad g(\cdot, 0) = g_0 \right\} \quad \text{and}$$
$$\left\{ \frac{\partial \mathcal{S}}{\partial t} = (\mathcal{P}H - \nabla\mathcal{P}.\vec{N})\vec{N}, \quad \mathcal{S}(\cdot,\cdot,0) = \mathcal{S}_0 \right\} \qquad (2)$$

where H and κ are respectively the mean curvature of the surface and the curvature of the curve. $\vec{N}$ and $\vec{n}$ are their inward normals. This approach is limited by the fact that it can lead to local minima of the energy. This is of course true for their level set formulation as well (see for example Caselles et al. (1997a) Osher and Sethian (1988)). Therefore, in the next section we recall a method introduced in Cohen and Kimmel (1997) that allows to find the global minimum for the active contour energy (1) when imposing the two end points. This formulation does not use the curve evolution equation in (2).

2.2. *Global Minimal Paths Between Two Points*

Cohen and Kimmel give in Cohen and Kimmel (1997) a method to find the global minimal path, connecting two points p_1 and p_2, with respect to a given cost function $\mathcal{P}$. In other words, they find the global minimum of the geodesic active contour's energy (1) when imposing to the curve its two end points. They show that this globally minimal curve is obtained by following the opposite gradient direction on the minimal action map $\mathcal{U}_{p_1}$,

$$\mathcal{U}_{p_1}(q) = \inf_{g(0)=p_1, g(L)=q} \left\{ \int_0^L \mathcal{P}\big(g(s)\big) ds \right\},$$
$$\text{where } L \text{ is the length of } g. \qquad (3)$$

The minimal path between p_2 and p_1 is thus obtained by solving the problem:

$$\frac{dg}{ds}(s) = -\nabla\mathcal{U}_{p_1}(g(s)) \quad \text{with } g(0) = p_2. \qquad (4)$$

In order to compute $\mathcal{U}_{p_1}$, Cohen and Kimmel (1997) use the fact that this map is solution to the well known eikonal equation (a proof of this fact can be found in Bruckstein (1988)):

$$\|\nabla\mathcal{U}_{p_1}\| = \mathcal{P} \quad \text{and} \quad \mathcal{U}_{p_1}(p_1) = 0. \qquad (5)$$

Equation (4) can be numerically solved by simple ordinary differential equations techniques like Newton's or Runge-Kutta's. To numerically solve Eq. (5), classic finite differences schemes tend to be unstable. In Tsitsiklis (1995) Tsitsiklis introduced a new method that was independently reformulated by Sethian in Sethian (1996). It relies on a one-sided derivative looking in the direction of the information flow, and it gives a consistent approximation of the weak solution to Eq. (5). This algorithm is known as the Fast Marching algorithm and is now widely used and understood. It was used in Cohen and Kimmel (1997) to solve Eq. (5) and find globally minimizing contours in images. More details on its background and implementation can be found in Sethian (1999) and Cohen (2001). It is important to highlight a major advantage of this algorithm: it has an $O(N \log(N))$ complexity on a grid of N nodes, and only one grid pass is needed to give a first order approximation of the solution. An extension to 3D of Fast Marching and minimal paths is straightforward. The authors of Deschamps and Cohen (2001) used it to find centerlines in 3D tubular structures. The minimal path is obtained by gradient descent, solving Eq. (4), like in the 2D case.

To summarize, we are able, by imposing its two end points, to build a 3D global minimum path for the energy, without using an evolution equation subject to unwanted local minima traps. On the other hand, the goal of active surfaces is to locate a certain local minimum of energy (1) that agrees with the user's criteria. The problem is that during the evolution process the surface can be trapped by other local minima, or, additional information could be necessary in order to complete image information and achieve a new wanted minimum.

In what follows, we propose to use the global minimum property of the paths to generate a segmentation surface $\mathcal{S}_0$ from two curves drawn by the user. We reduce the 3D initialization to drawing these curves, instead of complicated volumes in the case of difficult images. We also use these curves as additional user information for avoiding unwanted local minima.

3. From Global Minimal Paths to 3D Surface

We propose to use a set of minimal paths, built between two constraining curves $\mathcal{C}_1$ and $\mathcal{C}_2$, to define a first approximation of an energy minimizing surface $\mathcal{S}_0$. The intuition behind this approach is that this set of global minimal paths is contained in a surface that would qualify for a good segmentation approach if, in the beginning, $\mathcal{C}_1$ and $\mathcal{C}_2$ are well located in the 3D image.

3.1. Minimal Path Network

We wish to build a set of global minimal paths between the two constraining curves using the method outlined in the previous section. A naive numerical approach for this construction is to build minimal paths between all the points of the discretized versions of $\mathcal{C}_1$ and $\mathcal{C}_2$. Hence, each point of $\mathcal{C}1$ would be associated to all the points of $\mathcal{C}_2$. Clearly this would be computationally expensive (at least n actions maps to build and $n \times n$ gradient descents, if n is the number of points of the discretized versions of $\mathcal{C}_1$ and $\mathcal{C}_2$), and many of this numerous associations would not be relevant. Thus, we consider the following approach: We shall say that g is a path between a point p_1 and a curve $\mathcal{C}_1$ if $g(0) = p_1$ and $g(L) \in \mathcal{C}_1$. We then define surface $\mathcal{S}_0$ as the set of minimal energy paths $\{g^p_{\mathcal{C}_1}\}$ between curve $\mathcal{C}_1$ and all points p of the curve $\mathcal{C}_2$. More precisely, $\mathcal{S}_0 = \bigcup_{p\in\mathcal{C}_2}\{g^p_{\mathcal{C}_1}\}$.

As recalled in Ardon and Cohen (2003), the problem of computing $g^p_{\mathcal{C}_1}$, minimal path between $\mathcal{C}_1$ and p, can be addressed by performing a gradient descent on the action map $\mathcal{U}_{\mathcal{C}_1}$, defined by

$$\mathcal{U}_{\mathcal{C}_1}(p) = \inf_{\{g\ between\ p\ and\ \mathcal{C}_1\}} \left\{ \int_0^L \mathcal{P}(g(s))ds \right\}. \quad (6)$$

Furthermore it is easy to see that $\mathcal{U}_{\mathcal{C}_1}(p) = \inf_{q\in\mathcal{C}_1}\{\mathcal{U}_q(p)\}$, where $\mathcal{U}_q$ is the action map associated to point q defined in Section 2.2 by Eq. (3). This implies that the numerical estimation of $\mathcal{U}_{\mathcal{C}_1}$ can also be done using the Fast marching algorithm, initializing $\mathcal{U}_{\mathcal{C}_1}$ by $\mathcal{U}_{\mathcal{C}_1}(p) = 0$ if $p \in \mathcal{C}_1$ (a discretized version of it) and $\mathcal{U}_{\mathcal{C}_1}(p) = \infty$ otherwise. Indeed, this can be understood by recalling the fact that the value of $\mathcal{U}_{i,j,k}$ only depends on points among its six nearest neighbors whose values of $\mathcal{U}$ are inferior. Thus, when marching away from the points of $\mathcal{C}_1$, Fast marching will automatically compute $\inf_{q\in\mathcal{C}_1}\{\mathcal{U}_q(p)\}$.

Using $\mathcal{U}_{\mathcal{C}_1}$, we can now estimate $\mathcal{S}_0$. Consider a discretized version of $\mathcal{C}_2$ containing n_2 points $\{p_i\}_{i=1\ldots n_2}$. For each and every point p_i, by gradient descent on $\mathcal{U}_{\mathcal{C}_1}$, we build the minimal path between this point and $\mathcal{C}_1$, thus generating a finite set of paths from $\mathcal{C}_1$ to $\mathcal{C}_2$: $\{g^i_{\mathcal{C}_1}\}_{i=1\ldots n_2}$. The final numerical approximation of $\mathcal{S}_0$ will be the result of the interpolation of this network and concerns Section 4 of this paper. An illustration of $\mathcal{S}_0$ is given in Figs. 2(a) and (b) on a synthetic image. A potential adapted to finding the surface of the vase shown in Fig. 2(a) is used. The network, shown in Fig. 2(b), is built between two curves $\mathcal{C}_1$ and $\mathcal{C}_2$ drawn on the surface of the vase.

An important remark is that the definition of surface $\mathcal{S}_0$ is not symmetric. Indeed, in general $\bigcup_{p\in\mathcal{C}_2}\{g^p_{\mathcal{C}_1}\} \neq \bigcup_{p\in\mathcal{C}_1}\{g^p_{\mathcal{C}_2}\}$, and of course, the set of paths $\{g^i_{\mathcal{C}_1}\}_{i=1\ldots n_2}$ is different from its homologue set $\{g^i_{\mathcal{C}_2}\}_{i=1\ldots n_1}$. One could think of using this feature to generate a denser set of paths by looking for a surface that would be defined as the union of both networks. However, in practice, this symmetrical construction does not give satisfactory results.

An interesting particular case of the previous construction is obtained when curve $\mathcal{C}_1$ is reduced to a single point p_0. However, in this degenerated case, in order to obtain a coherent network $\mathcal{S}_0$, p_0 has to be situated in a specific location of the object to segment. This position corresponds to the maximum of the action map $\mathcal{U}_{\mathcal{C}_2}$ (solution to the eikonal equation taking zero values at curve $\mathcal{C}_2$) on the surface of the object to extract. This location is very difficult to find automatically, since the object is unknown; for the user, this

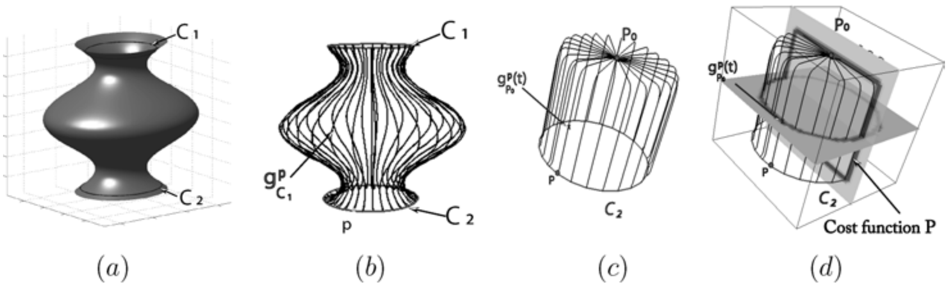

Figure 2. (a) is the original vase surface from which a 3D test image is generated. We also show the position of the constraining curves that are given by the user. (b) is the set of minimal paths ($\mathcal{S}_0$) generated between the two constraining curves. The paths are minimal with respect to a potential that takes small values on the vase's boundaries. Note that the paths of $\mathcal{S}_0$ lay on the vase's surface. (c) Set of Minimal Paths in the degenerated case: between a point and a curve lying on a closed cylinder. Point p_0 is located on the center of the upper face of the cylinder, it is the farthest point on the surface of the cylinder from curve $\mathcal{C}_2$. In (d) we superimposed to the network of paths the cost function $\mathcal{P}$ used for its construction.

point corresponds to the one being the farthest away from $\mathcal{C}_2$ on the surface. On Fig. 2(c) and (d) we give an illustration of this case on a synthetic image of a closed cylinder. As shown in Fig. 2(d) (where two slices of the 3D cost function are shown), point p_0 is the center of the upper part of the cylinder and curve $\mathcal{C}_2$ is drawn on the opposite side.

3.2. Projecting the Minimal Paths

Recall that functional $E(g) = \int_0^L \mathcal{P} l(g(s) l) ds$ is built by summing the cost function ($\mathcal{P}$) along the curve g. Hence, a minimal curve with respect to E establishes a balance between reducing its length, and following weak values of $\mathcal{P}$. In order to clarify the explanation that follows, we first consider a 2D situation, which corresponds to the illustration given on Fig. 3(a). The cost function is derived from a 2D image that contains an object we wish to extract. Suppose that this object presents a strong curvature on the neighborhood of a certain point p. Consider two other points, p_1 and p_2, also positioned on this object, and relatively far from p with respect to a characteristic size of the concavity. Then, as illustrated in Fig. 3(a), a minimal path between p_1 and p_2 will tend to avoid this concavity by 'cutting through' stronger values of $\mathcal{P}$, its length being toopenalizing otherwise. In the 3D case, minimal

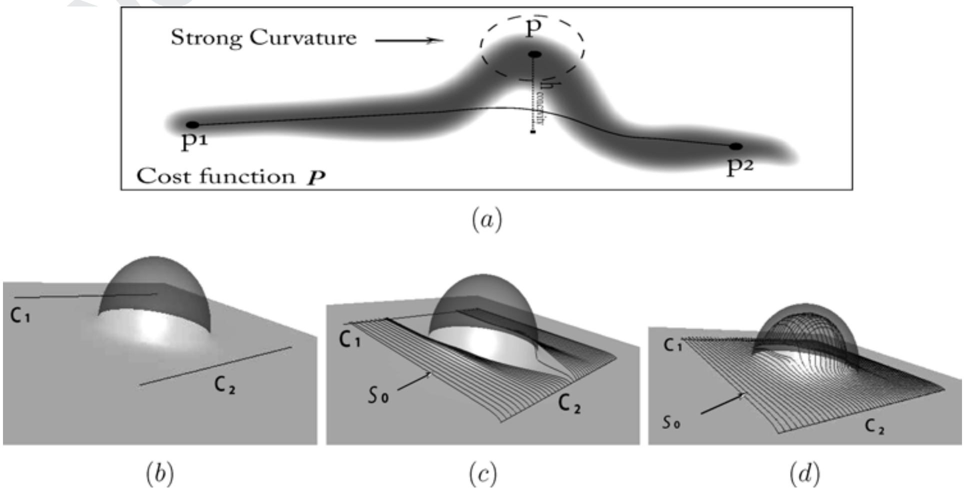

Figure 3. (a) Minimal path between points p_1 and p_2 that avoids a concavity of the object to segment. $h_{concavity}$ is the characteristic size of the concavity. (b) represents a half-sphere blended on a plane (transparent visualization) and $\mathcal{C}_1$ and $\mathcal{C}_2$ (black segments). (c) Result without constraints, set of paths $\{g^i_{\mathcal{C}_1}\}_{i=1\ldots n_2}$ taking a short cut around the sphere. (d) Result with constraints, sphere recovered.

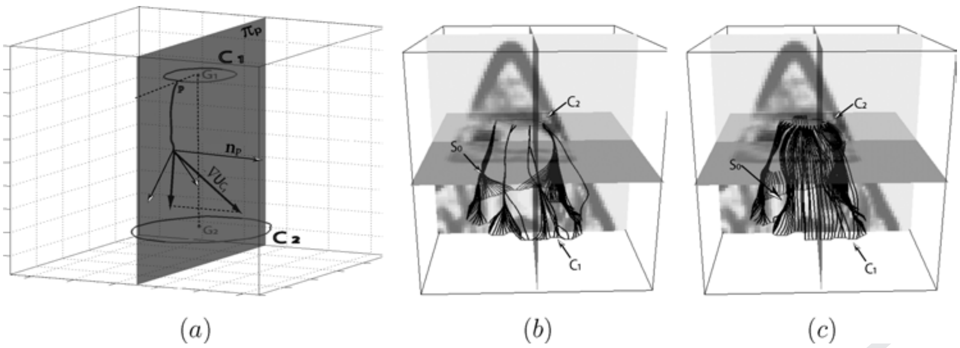

Figure 4. (a) Illustration of the construction of a projected path, it is done by projecting the vector field $\nabla\mathcal{U}_{\mathcal{C}_1}$ on plane π_p. (b) Shows the minimal path network obtained on an ultrasound image of the left ventricle without projecting. In transparency we gave three slices of the 3D volume. (c) is the projected network obtained in the same conditions.

paths of $\mathcal{S}_0$ have a similar behavior: if the surface of the object we wish to extract presents a strong localized mean curvature, the elements of $\mathcal{S}_0$ will tend to circumvent it. This constitutes a drawback in the use of $\mathcal{S}_0$ for a segmentation task: areas of the surface to extract presenting strong curvature, can be omitted by the minimal path network. Figure 3 illustrates a simple situation where the network $\{g^i_{\mathcal{C}_1}\}$ is unable to recover the expected surface. The cost function is constant on a surface which is the blending of a plane and half a sphere and has higher values on the background. Minimal paths tend to take a short cut around the sphere rather than 'climbing' on it, $\mathcal{P}$ has no influence (being constant on the surface) and the paths will minimize their length. In order to cope with this problem, we propose another approach for the construction of a segmenting surface $\mathcal{S}_0$, in the particular case where the user given curves, $\mathcal{C}_1$ and $\mathcal{C}_2$, do not intersect. Consider a family of planes $\Pi = \{\pi_p\}_{p\in\mathcal{C}_2}$, such that, for every p of $\mathcal{C}_2$, plane π_p contains this point and has a none empty intersection with $\mathcal{C}_1$. If $\vec{n_p}$ is the unit normal vector of plane π_p of Π, we call the projected minimal path $\tilde{g}^p_{\mathcal{C}_1}$, the solution of the following ordinary differential equation:

$$\frac{dg}{ds}(s) = -\nabla\mathcal{U}_{\mathcal{C}_1}(g) + (\nabla\mathcal{U}_{\mathcal{C}_1}(g).\,\vec{n_p}).\,\vec{n_p},$$

with $g(0) = p$. As it is shown in Fig. 4(a), this equation is obtained by replacing the vector field $\nabla\mathcal{U}_{\mathcal{C}_1}$ in Eq. (4) by its projection on plane π_p (whose normal is $\vec{n_p}$). $\mathcal{S}_0$ will be now defined as the union of the 'projected minimal paths': $\mathcal{S}_0 = \bigcup_{p\in\mathcal{C}_2}\{\tilde{g}^p_{\mathcal{C}_1}\}$. Figure 3(c) illustrates the network $\{\tilde{g}^i_{\mathcal{C}_1}\}$ of projected paths obtained with our half sphere. Π is the family of parallel planes which are orthogonal to $\mathcal{C}_1$ and $\mathcal{C}_2$ (n_p does not depend on p and π_p contains point p of $\mathcal{C}_2$). In practice, if $\mathcal{C}_1$ and $\mathcal{C}_2$ are two planar Jordan's curves, for each point p_i of $\mathcal{C}_2$, good choices for planes π_{p_i} are the planes passing through the following three points: G_1, belonging to the interior of $\mathcal{C}_1$, G_2 belonging to the interior of $\mathcal{C}_2$ and p_i. The normal vectors are then defined by, $\vec{n_{p_i}} = (\overrightarrow{G_1G_2} \wedge \overrightarrow{G_1p_i})(||\,\overrightarrow{G_1G_2} \wedge \overrightarrow{G_1p_i}\,||)^{-1}$.

In spite of the simplicity of this approach, the class of surfaces that can be segmented by evaluating their intersection with a plane, is quite large. This class contains at least those surfaces whose intersections with planes $\{\pi_{p_i}\}$ are connected.

In Fig. 4 we used this approach with a noisy ultrasound image of the left ventricle. Figure 4(b) shows the minimal path network obtained without the projected approach. Noise and the structure of the surface create strong curvature and many areas of the surface to extract are avoided by the network, the segmentation generated from this network will be of less precision. Here, the projection to planes is of great use due to the particular geometry of the ventricle: Fig. 4(c) shows how we manage to recover the areas that where missed by the unprojected network.

4. From the Network to the Surface

The final step for the generation of $\mathcal{S}_0$ is its construction through the interpolation of the network of paths. We consider two different approaches to generate $\mathcal{S}_0$. The first one is a novel analytical interpolation that uses the unprojected network of paths $\{g^i_{\mathcal{C}_1}\}_{i=1\ldots n_1}$; it exploits its particular structure which derives from the fact that minimal paths cannot cross without merging.

This method is fast and guaranties that the interpolated surface strictly contains all the paths of the network and the curves given by the user. The second, uses the variational approach proposed in Zhao et al. (2001). It can be applied to both, the unprojected and the projected network ($\{\tilde{g}^i_{\mathcal{C}_1}\}_{i=1\ldots n_1}$), but only ensures that the interpolated surface is close to the network but may not strictly contain all its paths.

4.1. Analytical Path Interpolation

In this section we present the construction of the interpolated surface from the unprojected network $\{g^i_{\mathcal{C}_1}\}_{i=1\ldots n_2}$ (henceforth noted $\{g^i\}_{i=1\ldots n}$ for simplicity). When the goal is to rapidly generate an approximation of the segmented surface (since we could miss areas of high curvature), this approach will be a good compromise between precision and efficiency. Being minimal paths, two paths belonging to $\{g^i\}$ may either have an empty intersection or merge (note that this is not the case for the elements of $\{\tilde{g}^i_{\mathcal{C}_1}\}_{i=1\ldots n_1}$). This particular configuration (see Fig. 5(a) for a schematic representation) of the network suggests to create sectors and interpolate the surface sector by sector (see definition below). Let s_1 and s_2 be parameterizations of $\mathcal{C}_1$ and $\mathcal{C}_2$ defined on the interval [0, 1]. Points $\{P^i_1\}$ and $\{P^i_2\}$ will be the intersection points of $\mathcal{C}_1$ and $\mathcal{C}_2$ with the network $\{g^i\}$ (see Fig. 5(a)). And $\{p^i_1\}$ and $\{p^i_2\}$ two families belonging to [0, 1] satisfying $C_1(s_1 = p^i_1) = P^i_1$ and $C_2(s_2 = p^i_2) = P^i_2$. For every $i \in \{1 \ldots n\}$ we define a sector as the following set of curves $\{g^i, g^{i+1}, \mathcal{C}^i_1, \mathcal{C}^i_2\}$ (as is shown on Fig. 5(b)). $\mathcal{C}^i_1$ and $\mathcal{C}^i_2$ are the restrictions of curves $\mathcal{C}_1$ and $\mathcal{C}_2$ to the intervals $[p^i_1, p^{i+1}_1]$ and $[p^i_2, p^{i+1}_2]$ respectively.

Our aim is to generate a parameterized surface $\mathcal{S}_0 : [0, 1]^2 \rightarrow \mathbb{R}^3$; $(u, v) \rightarrow \mathcal{S}_0(u, v)$, such that $\exists\{p^i\}_{1\le i\le n} \in [0, 1]^n$, verifying

$$(\mathcal{C}ond_1) : \forall i \in \{1....n\}\mathcal{S}_0(., v = p^i) \equiv g^i,$$
$$\mathcal{S}_0(u = 0, .) \equiv \mathcal{C}_1 \quad \text{and} \quad \mathcal{S}_0(u = 1, .) \equiv \mathcal{C}_2$$

meaning that the essential constraint on $\mathcal{S}_0$ is to contain curves $\mathcal{C}_1, \mathcal{C}_2$ and all paths $\{g^i\}$. Moreover, consider the restrictions $\mathcal{S}^i_0$ of $\mathcal{S}_0$ to the sets $[0, 1] \times [p^i_1, p^{i+1}_1]$. By imposing to $\mathcal{S}_0$ the following condition, $\forall u \in [0, 1]$ and $\forall i = 1 \ldots n - 1$.

$$(\mathcal{C}ond_2) : \ \partial_v \mathcal{S}^i_0(u, v = p^{i+1}) = \partial_v \mathcal{S}^{i+1}_0(u, v = p^{i+1})$$

we can build it locally continuously differentiable. In fact, it is easy to build $\mathcal{S}_0$ of class C^1 in the interior of each sector, difficulty arises only on the boundaries. The analytical construction that follows will guaranty that $\mathcal{S}_0$ will stay first order differentiable at the borders of each sector if paths do not merge, and continuous if they do.

The first step of the analytical interpolation is the introduction of a common parameterization on $\mathcal{C}_1$ and $\mathcal{C}_2$ (that will be noted v), and another (noted u) on all paths $\{g^i\}$. Parameter u is easy to find, it will be chosen as the normalized arc-length on each path $\{g^i\}$. In order to find v, let σ be an increasing one-to-one function on [0, 1], such that for every i, $\sigma(p^i_1) = p^i_2$. We perform a remapping of C_2 by σ and the new curve $\tilde{\mathcal{C}}_2 = \mathcal{C}_2 \circ \sigma$, satisfies for every $i \in \{1...n\}$ $\tilde{\mathcal{C}}_2(p^i_1) = \mathcal{C}_2 \circ \sigma(p^i_1) = \mathcal{C}_2(p^i_2) = P^i_2$. Which means that the same parameter values on $[0, 1]^n$ ($\{p^i_1\}$) correspond in each curve $\mathcal{C}_1$ and $\tilde{\mathcal{C}}_2$ to the intersection points with the set $\{g^i\}$. This leads us to choosing parameterization $v = s_1$ and henceforth working with $\mathcal{C}_1$ and $\tilde{\mathcal{C}}_2$. Finding an adequate σ function is a problem of a 1D constraint interpolation (since σ ought to be increasing). We use a piecewise cubic hermite interpolation (Fritsch and Carlson, 1980) to solve the problem.This function

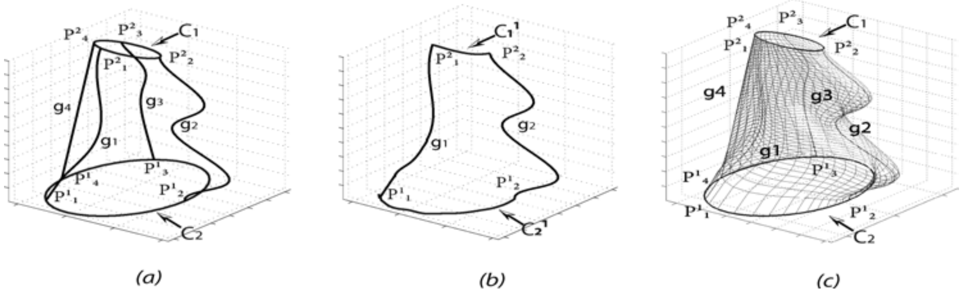

Figure 5. (a) Scheme illustrating a network that satisfies all the conditions for applying the analytical interpolation. (b) Illustrates our definition of a sector and (c) shows the interpolated surface, generated with our analytical method.

reflects the correspondence generated by the minimal paths between the two curves. We are now able to give an analytical expression of function $\mathcal{S}_0$ that satisfies conditions ($\mathcal{C}ond_1$) and ($\mathcal{C}ond_2$). For each sector i, we define the x-coordinate of the restriction of $\mathcal{S}_0$ by

$$\mathcal{S}^i_{0x}(u,v) = \mathcal{C}^i_x(u,v) + \left(1-\alpha^i_x(u,v)\right)\left(g^i_x(u) - \mathcal{C}^i_x(u,p^i)\right) + \alpha^i_x(u,v)\left(g^{i+1}_x(u) - \mathcal{C}^i_x(u,p^{i+1})\right)$$

where $\mathcal{C}^i_x(u,v) = (1-f(u))\mathcal{C}^i_{1x}(v) + f(u)\tilde{\mathcal{C}}^i_{2x}(v)$ (convex combination of the given curves). Function f can be chosen among all the differentiable functions on $[0,1]$ and must satisfy $f(0)=0$ and $f(1)=1$ (take for example $f(u)=u$). Each scalar α^i_x is the x-coordinate of a function α^i, which is tailored for satisfying ($\mathcal{P}_1$) and ($\mathcal{P}_2$); it is defined on the interval $[p^i_1, p^{i+1}_1]$ by

$$\alpha^i_x(u,v) = \frac{v-p^i}{p^{i+1}-p^i}\left(1 + \frac{p^{i+1}-v}{p^{i+1}-p^i}\left(\frac{v-p^i}{p^{i+1}-p^i}\cdot \left[2-\left(G^{i+1}_x(u)+G^{i-1}_x(u)\right)\right] + \left(G^{i-1}_x(u)-1\right)\right)\right).$$

with

$$\begin{aligned} G^i_x(u) &= g^{i+1}_x(u) - g^i_x(u) - \left(\mathcal{C}^i_x(u,p^{i+1}) - \mathcal{C}^i_x(u,p^i)\right), \\ &\quad \forall i \in \{1\ldots n-1\} \\ G^n_x(u) &= G^0_x(u) = g^1_x(u) - g^n_x(u) - \left(\mathcal{C}^n_x(u,p^1) - \mathcal{C}^n_x(u,p^n)\right) \\ G^{n+1}_x(u) &= G^1_x(u) \end{aligned}$$

The other two coordinates are obtained using the same formulas replacing x by y and then by z. Figure 5(c) shows the interpolated mesh generated from the set of curves in Fig. 5(a). A major advantage of this interpolation method is its calculation speed. Only elementary calculations are needed to generate the surface (there is no matrix inversion) and both information from the paths and from the initial curves are integrated in the process.

In Fig. 6 we show two interpolated surfaces generated by this method. Fig. 6(a) and (b) illustrate the fact that the interpolation combines both information coming from the network and from curves $\mathcal{C}_1$ and $\mathcal{C}_2$. Even when taking only four paths, the obtained surface is coherent with the shape of the user given curves. Figure 6(c) shows set $\{g^i\}$ obtained from a left ventricle image, Fig. 6(d) illustrates the interpolated surface.

4.2. *Variational Interpolation*

As was pointed out earlier, the analytical interpolation method can only be applied with the unprojected network $\{g^i_{\mathcal{C}_1}\}_{i=1\ldots n_1}$, since its particular structure (paths cannot cross without merging) is necessary. Nevertheless, considering the projected networks can improve results (see Fig. 4(b)). Unfortunately, a sector by sector approach can no longer be considered, for paths can cross without merging. In these situations one can hardly exploit the structure of the network, hence, a scattered data points interpolation has to be considered. We use the method proposed by Zhao et al. in Zhao, et al. (2001). We compute from the network $\{\tilde{g}^i_{\mathcal{C}_1}\}_{i=1\ldots n_1}$ a distance function d and we look for the surface $\mathcal{S}_0$ that minimizes energy $E(S) = \iint_S (d(x,y))^2 dx dy$. This is done by a gradient descent method similar to Eq. (2) and we have used a level set implementation. When using projected networks, this method gives satisfactory results since one can control the density of the paths by varying the number of points on $\mathcal{C}_1$ and $\mathcal{C}_2$.

5. Initializing Active Surface with $\mathcal{S}_0$, Applications

Having generated $\mathcal{S}_0$ by any of the previous methods, we may use it as the initial condition of the evolution Eq. (2). We have chosen a level set method for our implementations. If the analytical interpolation method is used, the construction of a higher dimensional function, $\phi_0 : \mathbb{R}^3 \to \mathbb{R}$ such that $\phi_0^{-1}(0) = \mathcal{S}_0$, is needed.

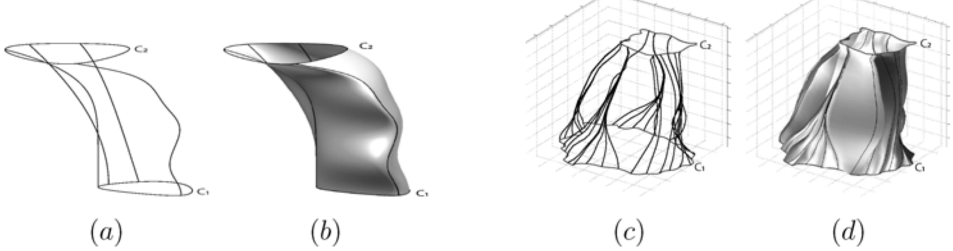

Figure 6. (a) Test network of four paths synthetically produced, $\mathcal{C}_1$ and $\mathcal{C}_2$ are the lower and upper curves. (b) Interpolated surface. (c) is the network of minimal paths obtained from an ultrasound image of the left ventricle. The user initialized the model by drawing the upper and lower closed curves. (d) is the analytically interpolated surface.

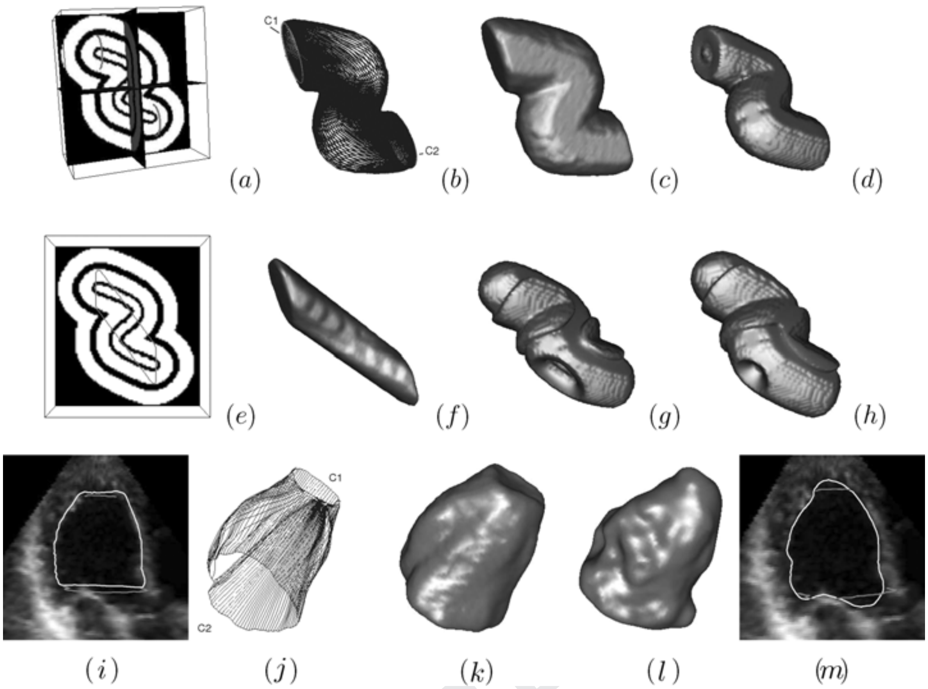

Figure 7. (a) View of different intersecting planes of a 3D volume with the two constraining curves drawn on it. (b) Network of paths obtained with our method. (c) Interpolated surface. (d) Surface after a few iterations of a level set. (e) and (f) Simple initialization of an active object. (g) surface after 150 iterations and (h) after 500 iterations. (i) A slice of the 3D ultrasound image, we also have drawn the projection of the user given curves and the intersection of our interpolated surface with this plane. (j) Set of paths. (k) Interpolated surface. (l) final segmentation after a few iterations of the level set, (m) Planar view of the same slice, intersection with the model evolved as a level set.

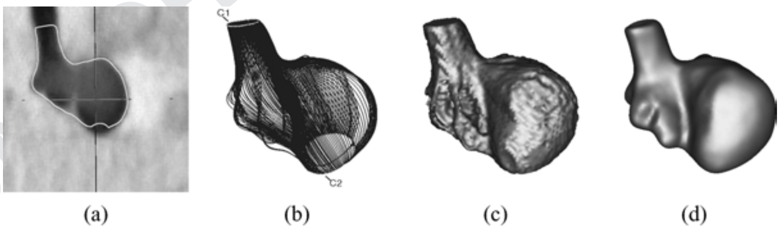

Figure 8. (a) Slice of a 3D MR image of an aneurysm. (b) Set of paths. (c) Interpolated surface. (d) final segmentation after a few iterations of a level set.

ϕ_0 can be computed as a signed distance map using Fast marching initialized with $\mathcal{S}_0$. The evolution of the level set will be done following $\frac{\partial\phi}{\partial t} = div(\mathcal{P}.\frac{\nabla\phi}{\|\nabla\phi\|})|\nabla\phi\|$, which is exactly the gradient descent of the geodesic active surface (1) in its level set formulation. For convergence, few iterations of ϕ will be needed, since $\mathcal{S}_0$ is already close to image features. Compared to using a level set approach from the beginning, our approach is much faster, needs no tedious 3D initializations, and avoids local minima by exploiting curves $\mathcal{C}_1$ and $\mathcal{C}_2$. Figure 7(a) presents a good example of a difficult to segment image because of the presence of many local minima. It is generated by three ‘S’ shaped tubes one inside the other. If one wishes to obtain the middle ‘S’ shaped tube, classical variational methods will fail (unless a very close initialization is given). Our method manages to extract the object when initialized by two curves given on the surface to extract. We compare it with the result of a geodesic active surface initialized with a cylinder (Fig. 7(e) and (f), and we observe in Fig. 7(g) and (h) that the model gets trapped by other local minima. Concerning ultrasound heart imaging, our method only needs two slices in order to build the entire volume of the left ventricle; this two curves can be, for example, two short axis segmentations as in Fig. 1(a) and (b). Figures 7(i) to (j) show the segmentation obtained. For this image of size 128^3, the generation of $\mathcal{S}_0$ took 25 seconds, the final segmentation 20 seconds more, on a 1.4 Ghz machine (512 MBy of RAM). In Fig. 8(a) to (d) , we show results on a MR image of an aneurysm. As for other previous examples, the user simply initialized the model by drawing two curves on two (non parallel) slices of the 3D image. On this image ($192 \times 168 \times 152$)), the total segmentation took 70 seconds on the same machine.

6. Conclusion

In this paper we have presented a method that generalizes globally minimal paths to surfaces. Our method allows to greatly simplify the initialization process of active surfaces. The model is initialized by two curves (eventually a curve and a well positioned point) instead of a volume. Our approach takes a maximum advantage of the information given by the user through the initialization curves, since the surface it generates is constrained to include those curves. Our method uses globally minimal paths to define and generate a surface which is a final segmentation or an initialization of an active surface model. Hence, in both cases, the final surface is not concerned by the problem of the local minima traps as all other active objects approach do. It is particularly well suited for medical image segmentation, in particular for ultrasound images segmentation. In cases where the image quality is very poor, our approach handles the introduction of additional information coming from the practitioner in a very natural manner. A few 2D segmentations can be enough to generate a coherent complete surface. We have also presented a novel interpolation method which is characterized by its simplicity and its efficiency.

References

Ardon, R. and Cohen, L. 2003. Fast constrained surface extraction by minimal paths. *2nd IEEE Workshop on Variational, Geometric and Level Set Methods in Computer Vision*, pp. 233–244.

Bruckstein, A. 1988. On shape from shading. *Computer Vision, Graphics, and Image Processing*, 44(2):139– 154.

Caselles, V., Kimmel, R., and Sapiro, G. 1997a. Geodesic active contours. *International Journal of Computer Vision*, 22(1):61–79.

Caselles, V., Kimmel, R., Sapiro, G., and Sbert, C. 1997b. Minimal-Surfaces based object segmentation. *IEEE Transactions On Pattern Analysis and Machine Intelligence*, 19(4):394–398.

Cohen, L. 1991. On active contour models and balloons. *Computer Vision, Graphics, and Image Processing: Image Understanding*, 53(2):211–218.

Cohen, L. 1997. Avoiding local minima for deformable curves in image analysis. In A. L. Mehaute, C. Rabut, and L. L. Schumaker (eds.), *Curves and Surfaces with Applications in CAGD*, Nashville.

Cohen, L. 2001. Multiple contour finding and perceptual grouping using minimal paths. *Journal of Mathematical Imaging and Vision*, 14(3).

Cohen, L. and Kimmel, R. 1997. Global minimum for active contour models: A minimal path approach. *International Journal of Computer Vision*, 24(1):57– 78.

Cremers, D. and Schnörr, C. 2003. Statistical Shape Knowledge in Variational Motion Segmentation. *Image and Vision Computing*, 21(1):77–86.

Deschamps, T. and Cohen, L. 2001. Fast extraction of minimal paths in 3D images and applications to virtual endoscopy. *Medical Image Analysis*, 5(4).

Fritsch, F.N. and Carlson, R.E. 1980. monotone piecewise cubic interpolation. *SIAM J. Numerical Analysis*, 17:238–246.

Kass, M., Witkin, A., and Terzopoulos, D. 1988. Snakes: Active contour models. *International Journal of Computer Vision*, 1(4):321–331.

Osher, S. and Sethian, J. 1988. Fronts propagating with curvature dependent speed: Algorithms based on the Hamilton-Jacobi formulation. *Journal of Computational Physics*, 79:12–49.

Paragios, N. 2000. Geodesic active regions and level set methods: Contributions and applications in artificial vision. Ph.D. thesis, Université de Nice Sophia-Antipolis, France.

Sethian, J. 1996. A fast marching level set method for monotonically advancing fronts. *Proceedings of the Natural Academy of Sciences*, 93(4):1591–1595.

Sethian, J. 1999. *Level set methods: Evolving Interfaces in Geometry, Fluid Mechanics, Computer Vision and Materials Sciences*, University of California, Berkeley: Cambridge University Press, 2nd edition.

Tsai, A., Jr., A. Y., Wells, W., Tempany, C., Tucker, D., Fan, A., Grimson, W. E., and Willsky, A. 2003. A shape-based approach to the segmentation of medical imagery using level sets. *IEEE Transactions on Medical Imaging*, 22(2):137–154.

Tsitsiklis, J.N. 1995. Efficient algorithms for globally optimal trajectories. *IEEE Transactions on Automatic Control*, 40(9):1528–1538.

Yuille, A., Hallinan, P., and Cohen, D. 1992. Feature extraction from faces using Deformable templates. *International Journal of Computer Vision*, 8(2):99–111.

Zhao, H., Osher, S. and Fedkiw, R. 2001. Fast surface reconstruction using the level set method. *Workshop on Variational and Level Set Methods In Computerj Vision*, pp. 194–201.