

Finding a Closed Boundary by Growing Minimal Paths from a Single Point on 2D or 3D Images

Fethallah Benmansour Stéphane Bonneau
Laurent D. Cohen

CEREMADE, Université Paris Dauphine

benmansour, bonneau, cohen @ceremade.dauphine.fr

Abstract

In this paper, we present a new method for segmenting closed contours and surfaces. Our work builds on a variant of the Fast Marching algorithm. First, an initial point on the desired contour is chosen by the user. Next, new keypoints are detected automatically using a front propagation approach. We assume that the desired object has a closed boundary. This a-priori knowledge on the topology is used to devise a relevant criterion for stopping the keypoint detection and front propagation. The final domain visited by the front will yield a band surrounding the object of interest. Linking pairs of neighboring keypoints with minimal paths allows us to extract a closed contour from a 2D image. Detection of a variety of objects on real images is demonstrated. Using a similar same idea, we can extract networks of minimal paths from a 3D image called Geodesic Meshing. The proposed method is applied to 3D data with promising results.

1. Introduction

Energy minimization techniques have been applied to a broad variety of problems in image processing and computer vision. Since the original work on snakes [9], they have notably been used for boundary detection. An active contour model, or snake, is a curve that deforms its shape in order to minimize an energy combining an internal part which smooths the curve and an external part which guides the curve toward particular image features. For instance, the geodesic active contour model [3, 17] relies on the minimization of a geometric energy functional that deforms an initial curve toward local geodesics in a Riemannian metric derived from the image. Whereas the geodesic active contour model presents significant improvements compared to the original snake model, the energy minimization process is still prone to local minima. Consequently, results strongly depend on the model initialization.

To overcome this issue, Cohen and Kimmel [6] introduced an approach to globally minimize the geodesic active contour energy, provided that two endpoints of the curve are initially supplied by the user. This energy is of the form $\int_{\gamma} \tilde{\mathcal{P}}$ where the incremental cost $\tilde{\mathcal{P}}$ is chosen to take lower values on the contour of the image, and γ is a path joining the two points. The solution of this minimization problem is obtained through the computation of the minimal action map associated to a source point. The minimal action map can be regarded as the arrival times of a front propagating from the source point with velocity $(1/\tilde{\mathcal{P}})$, and it satisfies the Eikonal equation. Therefore, we can compute simultaneously, and efficiently, the minimal action map and its Euclidean path length with the *Fast Marching Method* as will be detailed in section 2.2.

In section 3, we introduce a novel front propagation approach, based on the Fast Marching Method, to distribute a set of points on a codimension-1 closed manifold that is not known a priori, all starting from a single point (or more if desired) initialized on the desired object boundary. Each newly detected keypoint is immediately defined as a new source of propagation, and keypoints are detected with a criterion based on the Euclidean length of the minimal paths. Since the front propagates faster on the object boundary, the first point for which the length λ is reached, is located in this area (of small values of $\tilde{\mathcal{P}}$) and is a valuable choice as a keypoint. By using the a-priori knowledge on the topology of the manifold, we devise a relevant criterion for stopping the keypoint detection and front propagation. The criterion is general for any dimension. In section 4, we explain how to extract a codimension-1 closed manifold within the image using the previous results. The main idea is to link pairs of neighboring keypoints with minimal paths via gradient descent on the minimal action map. Segmentation results on a set of 2D and 3D images are presented. Finally conclusions and perspectives follow in section 5.

2. Background on minimal paths

2.1. Definitions

Given a 2D image $I : \Omega \rightarrow \mathbb{R}^+$ and two points \mathbf{p}_1 and \mathbf{p}_2 , the underlying idea introduced by Cohen and Kimmel [6] is to build a potential $\mathcal{P} : \Omega \rightarrow \mathbb{R}^{*+}$ which takes lower values near desired features of the image I . The choice of the potential \mathcal{P} depends on the application. For example, one can define \mathcal{P} as a decreasing function of $\|\nabla I\|$ to extract image edges by finding a curve that globally minimizes the energy functional $E : \mathcal{A}_{\mathbf{p}_1, \mathbf{p}_2} \rightarrow \mathbb{R}^+$

$$E(\gamma) = \int_{\gamma} \left\{ \mathcal{P}(\gamma(s) + w) \right\} ds = \int_{\gamma} \tilde{\mathcal{P}}(\gamma(s)) ds, \quad (1)$$

where $\mathcal{A}_{\mathbf{p}_1, \mathbf{p}_2}$ is the set of all paths connecting \mathbf{p}_1 to \mathbf{p}_2 , s is the arc-length parameter, $w > 0$ is a regularization term and $\tilde{\mathcal{P}} = (\mathcal{P} + w)$. A curve connecting \mathbf{p}_1 to \mathbf{p}_2 that globally minimizes the energy (1) is a *minimal path* between \mathbf{p}_1 and \mathbf{p}_2 , noted $\mathcal{C}_{\mathbf{p}_1, \mathbf{p}_2}$. The solution of this minimization problem is obtained through the computation of the *minimal action map* $\mathcal{U}_1 : \Omega \rightarrow \mathbb{R}^+$ associated to \mathbf{p}_1 . The minimal action is the minimal energy integrated along a path between \mathbf{p}_1 and any point \mathbf{x} of the domain Ω :

$$\forall \mathbf{x} \in \Omega, \mathcal{U}_1(\mathbf{x}) = \min_{\gamma \in \mathcal{A}_{\mathbf{p}_1, \mathbf{x}}} \left\{ \int_{\gamma} \tilde{\mathcal{P}}(\gamma(s)) ds \right\}. \quad (2)$$

The values of \mathcal{U}_1 may be regarded as the arrival times of a front propagating from the source \mathbf{p}_1 with velocity $(1/\tilde{\mathcal{P}})$. \mathcal{U}_1 satisfies the Eikonal equation

$$\begin{cases} \|\nabla \mathcal{U}_1(\mathbf{x})\| = \tilde{\mathcal{P}}(\mathbf{x}) & \text{for } \mathbf{x} \in \Omega, \\ \mathcal{U}_1(\mathbf{p}_1) = 0. \end{cases} \quad (3)$$

The map \mathcal{U}_1 has only one local minimum, the point \mathbf{p}_1 , and its flow lines satisfy the Euler-Lagrange equation of functional (1). Thus, the minimal path $\mathcal{C}_{\mathbf{p}_1, \mathbf{p}_2}$ can be retrieved with a simple gradient descent on \mathcal{U}_1 from \mathbf{p}_2 to \mathbf{p}_1 (see Fig. 1), solving the following ordinary differential equation with standard numerical methods like Heun's or Runge-Kutta's :

$$\begin{cases} \frac{d\mathcal{C}_{\mathbf{p}_1, \mathbf{p}_2}(s)}{ds} = -\nabla \mathcal{U}_1(\mathcal{C}_{\mathbf{p}_1, \mathbf{p}_2}(s)), \\ \mathcal{C}_{\mathbf{p}_1, \mathbf{p}_2}(0) = \mathbf{p}_2. \end{cases} \quad (4)$$

Let us extend the definitions given so far to the case of multiple sources and introduce other definitions which will be useful hereinafter. These definitions hold in dimension 2 and higher. The *minimal action map* associated to the potential $\tilde{\mathcal{P}} : \Omega \rightarrow \mathbb{R}^{*+}$ and the set of n sources $\mathcal{S} = \{\mathbf{p}_1, \dots, \mathbf{p}_n\}$ is the function $\mathcal{U} : \Omega \rightarrow \mathbb{R}^+$ defined by

$$\forall \mathbf{x} \in \Omega, \mathcal{U}(\mathbf{x}) = \min_{1 \leq j \leq n} \{ \mathcal{U}_j(\mathbf{x}) \},$$

where $\mathcal{U}_j(\mathbf{x}) = \min_{\gamma \in \mathcal{A}_{\mathbf{p}_j, \mathbf{x}}} \left\{ \int_{\gamma} \tilde{\mathcal{P}}(\gamma(s)) ds \right\}. \quad (5)$

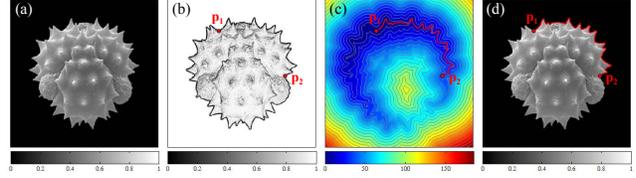


Figure 1. Extraction of an open contour from an electron microscopy image. (a) Original image I . (b) Potential $\mathcal{P} = (\|\nabla I\| + \varepsilon)^{-3}$, where ε is a small positive constant, and user-supplied points \mathbf{p}_1 and \mathbf{p}_2 . (c) Minimal action map \mathcal{U}_1 and minimal path $\mathcal{C}_{\mathbf{p}_1, \mathbf{p}_2}$ between \mathbf{p}_1 and \mathbf{p}_2 . (d) Image I and minimal path $\mathcal{C}_{\mathbf{p}_1, \mathbf{p}_2}$.

The map \mathcal{U} is a weighted distance map to the set of sources \mathcal{S} , and it satisfies the Eikonal equation

$$\begin{cases} \|\nabla \mathcal{U}(\mathbf{x})\| = \tilde{\mathcal{P}}(\mathbf{x}) & \text{for } \mathbf{x} \in \Omega, \\ \mathcal{U}(\mathbf{p}_j) = 0 & \text{for } \mathbf{p}_j \in \mathcal{S}. \end{cases} \quad (6)$$

The *Voronoi region* associated to the source $\mathbf{p}_j \in \mathcal{S}$, noted \mathcal{R}_j , is the locus of points of the domain Ω which are closer (in the sense of a weighted distance) to \mathbf{p}_j than to any other source of \mathcal{S} . The region \mathcal{R}_j is a connected subset of the domain Ω , and its boundary is noted $\partial \mathcal{R}_j$. The union of Voronoi regions and its complementary set, the *Voronoi diagram*, leads to a tessellation of the domain Ω , called the *Voronoi partition*.

The *Voronoi index map* is the function $\mathcal{V} : \Omega \rightarrow \{1, \dots, n\}$ that assigns to any point of the domain Ω the index of its Voronoi region :

$$\forall \mathbf{x} \in \mathcal{R}_j, \mathcal{V}(\mathbf{x}) = j. \quad (7)$$

If two Voronoi regions \mathcal{R}_i and \mathcal{R}_j are adjacent (i.e. if $\partial \mathcal{R}_i \cap \partial \mathcal{R}_j$ is a non-empty set), then the minimal path $\mathcal{C}_{\mathbf{p}_i, \mathbf{p}_j}$ passes through the point of $\partial \mathcal{R}_i \cap \partial \mathcal{R}_j$ which has the smallest \mathcal{U} value. This point, noted $\mathbf{m}_{i|j}$, is the *midpoint* of the minimal path $\mathcal{C}_{\mathbf{p}_i, \mathbf{p}_j}$ since it is equidistant to \mathbf{p}_i and \mathbf{p}_j in the sense of a weighted distance.

The *Euclidean path length map* is the function $\mathcal{L} : \Omega \rightarrow \mathbb{R}^+$ that assigns to any point \mathbf{x} of the domain Ω the Euclidean length of the minimal path between \mathbf{x} and the source which is the closest in the sense of a weighted distance :

$$\forall \mathbf{x} \in \mathcal{R}_j, \mathcal{L}(\mathbf{x}) = \int_{\mathcal{C}_{\mathbf{p}_j, \mathbf{x}}} ds. \quad (8)$$

Note that if $\tilde{\mathcal{P}}(\mathbf{x}) = 1$ for all $\mathbf{x} \in \Omega$, then the maps \mathcal{U} and \mathcal{L} are equal and both correspond to the Euclidean distance map to the set of sources \mathcal{S} .

Introduced first as a boundary detection method, minimal path techniques have been successfully applied to sundry problems (see [5] for a review), such as path planning [10], contour completion [4], tubular surface extraction [7], motion tracking [2], or remeshing of triangulated manifolds [11].

Table 1 : *Fast Marching Method for solving equation (6).*

- **Notation.**
 $\mathcal{N}_M(\mathbf{x})$ is the set of M neighbors of a grid point \mathbf{x} , where $M = 4$ in 2D and $M = 6$ in 3D.
- **Initialization.**
 For each grid point \mathbf{x} , do
 Set $\mathcal{U}(\mathbf{x}) := +\infty$, $\mathcal{V}(\mathbf{x}) := 0$ and $\mathcal{L}(\mathbf{x}) := +\infty$.
 Tag \mathbf{x} as *Far*.
 For each source $\mathbf{p}_j \in \mathcal{S}$, do
 Set $\mathcal{U}(\mathbf{p}_j) := 0$, $\mathcal{V}(\mathbf{p}_j) := j$ and $\mathcal{L}(\mathbf{p}_j) := 0$.
 Tag \mathbf{p}_j as *Trial*.
- **Marching loop.**
 While the set of *Trial* points is non-empty, do
 Find \mathbf{x}_{\min} , the *Trial* point with the smallest \mathcal{U} value.
 Tag \mathbf{x}_{\min} as *Alive*.
 For each grid point $\mathbf{x}_n \in \mathcal{N}_M(\mathbf{x}_{\min})$ which is not *Alive*, do
 $\{u, v, \ell\} := \text{UpdateSchemeFMM}(\mathbf{x}_n, \mathcal{N}_M(\mathbf{x}_n))$.
 Set $\mathcal{U}(\mathbf{x}_n) := u$, $\mathcal{V}(\mathbf{x}_n) := v$ and $\mathcal{L}(\mathbf{x}_n) := \ell$.
 If \mathbf{x}_n is *Far*, tag \mathbf{x}_n as *Trial*.

2.2. Fast Marching Method

The *Fast Marching Method* (FMM) is a numerical method introduced by Sethian in [13, 14, 15] and Tsitsiklis in [16] for efficiently solving the isotropic Eikonal equation on a cartesian grid. In equation (6), the values of \mathcal{U} may be regarded as the arrival times of wavefronts propagating from each point of \mathcal{S} with velocity $(1/\tilde{\mathcal{P}})$. The central idea behind the FMM is to visit grid points in an order consistent with the way wavefronts propagate, i.e. with the Huygens principle. It leads to a single-pass algorithm for solving equation (6) and computing the maps \mathcal{U} , \mathcal{V} and \mathcal{L} in a common computational framework (see Table 1).

Ordered sweeping of grid points.

The FMM is a front propagation approach that computes the values of \mathcal{U} in increasing order, and the structure of the algorithm is almost identical to Dijkstra's algorithm for computing shortest paths on graphs [8]. In the course of the algorithm, each grid point is tagged as either *Alive* (point for which \mathcal{U} has been computed and frozen), *Trial* (point for which \mathcal{U} has been estimated but not frozen) or *Far* (point for which \mathcal{U} is unknown). The set of *Trial* points forms an interface between the set of grid points for which \mathcal{U} has been frozen (the *Alive* points) and the set of other grid points (the *Far* points). This interface may be regarded as a set of fronts expanding from each source until every grid point has been reached.

Let us denote by $\mathcal{N}_M(\mathbf{x})$ the set of M neighbors of a grid point \mathbf{x} , where $M = 4$ if Ω is a 2D domain and $M = 6$ if Ω is a 3D domain. Initially, all grid points are tagged as *Far*, except the points of \mathcal{S} that are tagged as *Trial*. At each iteration of the FMM one chooses the *Trial* point with the smallest \mathcal{U} value, denoted by \mathbf{x}_{\min} in Table 1. Then, \mathbf{x}_{\min} is tagged as *Alive* and the values of \mathcal{U} , \mathcal{V} and \mathcal{L} are updated for each point of the set $\mathcal{N}_M(\mathbf{x}_{\min})$ which is either *Trial* or

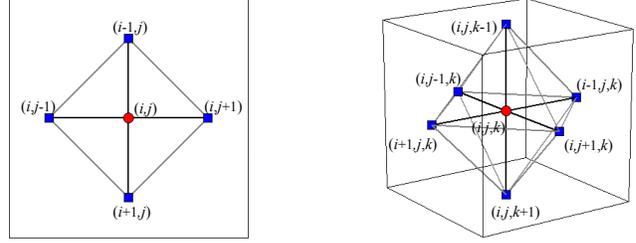


Figure 2. Connecting a grid point \mathbf{x}_n and the points of $\mathcal{N}_M(\mathbf{x}_n)$ with virtual edges forms four triangles on a 2D grid and eight tetrahedrons on a 3D grid.

Far. In order to satisfy a causality condition, the way \mathcal{U} , \mathcal{V} and \mathcal{L} are updated in the vicinity of \mathbf{x}_{\min} requires special care. The iteration ends by tagging every *Far* point of the set $\mathcal{N}_M(\mathbf{x}_{\min})$ as *Trial*. The algorithm automatically stops when all grid points are *Alive*.

The key to the speed of the FMM is the use of a priority queue to quickly find the *Trial* point with the smallest \mathcal{U} value. If *Trial* points are ordered in a min-heap data structure, the computational complexity of the FMM is $\mathcal{O}(N \log_2 N)$, where N is the total number of grid points.

Update scheme for the Fast Marching Method.

Here, we present a way to estimate \mathcal{U} , \mathcal{V} and \mathcal{L} for a grid point \mathbf{x}_n , i.e. a way to compute the outputs of routine `UpdateSchemeFMM` in Table 1. We limit ourselves to the 2D case, though extending the scheme to higher dimensions is straightforward. Adopting standard notation, we denote by $\mathcal{U}_{i,j}$ the value of \mathcal{U} at the grid vertex (i, j) associated to the point \mathbf{x}_n with coordinates $(i h_x, j h_y)$, where h_x and h_y are grid spacings in the x and y directions.

A discretized version of (6) is solved in order to compute $\mathcal{U}_{i,j}$. For the Eikonal equation, classic finite difference schemes tend to overshoot and are unstable. Rouy and Tourin [12] showed that the correct viscosity solution for $\mathcal{U}_{i,j}$ is given by the following first order accurate scheme :

$$\left(\frac{\max\{(\mathcal{U}_{i,j} - \mathcal{U}_{i-1,j}), (\mathcal{U}_{i,j} - \mathcal{U}_{i+1,j}), 0\}}{h_x} \right)^2 + \left(\frac{\max\{(\mathcal{U}_{i,j} - \mathcal{U}_{i,j-1}), (\mathcal{U}_{i,j} - \mathcal{U}_{i,j+1}), 0\}}{h_y} \right)^2 = (\tilde{\mathcal{P}}_{i,j})^2 \quad (9)$$

This is an upwind scheme : the forward and backward differences are chosen to follow the direction of the flow of information. Since the action can only grow due to the quadratic nature of equation (9), information is propagating *upwards* (from smaller to larger values of \mathcal{U}).

Connecting with virtual edges \mathbf{x}_n and the points of $\mathcal{N}_M(\mathbf{x}_n)$, i.e. \mathbf{x}_n and the four nearby grid points, forms four triangles (see Fig. 2). For each triangle, we attempt to solve two quadratic equations to get estimates of $\mathcal{U}_{i,j}$, $\mathcal{V}_{i,j}$ and

$\mathcal{L}_{i,j}$. These estimates are respectively denoted by u , v and ℓ . The Euclidean length ℓ is computed in the same manner as u by solving the equation $\|\mathcal{L}\| = 1$ in the same triangle. For each of four triangles, we get a triplet $\{u, v, \ell\}$. Finally, we choose the one with the smallest u (this is the triplet returned by the routine `UpdateSchemeFMM`). Note that if one needs to approximate $\nabla\mathcal{U}$, computing the derivatives of \mathcal{U} in the triangle used to estimate $\mathcal{U}_{i,j}$ gives a consistent approximation of $\nabla\mathcal{U}(\mathbf{x}_n)$. On a 3D grid, the neighborhood of a grid point is divided into eight tetrahedra (see Fig. 2), but the 3D update scheme is very similar to the 2D one.

3. Distribution of a set of points on a closed manifold

3.1. 2D case : distribution of a set of points on a closed curve

First, we consider the case where the domain Ω is a 2D domain. We assume that we are given an initial set $\mathcal{S} = \{\mathbf{p}_1, \dots, \mathbf{p}_n\}$ of points on a closed curve along which a potential $\tilde{\mathcal{P}} : \Omega \rightarrow \mathbb{R}^{*+}$ takes lower values. Note that the set \mathcal{S} may contain only one point.

We propose here a variant of the FMM, called the *Fast Marching Method With keypoint Detection* (FMMWKD, see Table 2), to propagate fronts from each point of \mathcal{S} with velocity $(1/\tilde{\mathcal{P}})$ and sequentially detect, during the front propagation, a set of keypoints $\mathcal{S}^* = \{\mathbf{p}_{n+1}^*, \dots, \mathbf{p}_{n+m}^*\}$ on the closed curve along which $\tilde{\mathcal{P}}$ takes low values. Each newly detected keypoint is immediately defined as a new source of propagation, and keypoints are detected with a criterion based on the Euclidean length of minimal paths. This criterion depends on only one parameter, denoted λ . Front propagation and keypoint detection ceases as soon as the domain visited by the fronts contains the whole curve of interest.

The final domain visited by the fronts, denoted Ω_F , correspond to a band surrounding the curve of interest. Furthermore, the FMMWKD also enables the computation of the minimal action map $\mathcal{U} : \Omega_F \rightarrow \mathbb{R}^+$, the Voronoi index map $\mathcal{V} : \Omega_F \rightarrow \{1, \dots, n+m\}$ and the Euclidean path length map $\mathcal{L} : \Omega_F \rightarrow \mathbb{R}^+$ associated to the potential $\tilde{\mathcal{P}}$ and the set of sources $\mathcal{S} \cup \mathcal{S}^*$.

Keypoint detection and local correction of maps \mathcal{U} , \mathcal{V} and \mathcal{L} .

Initially, fronts are propagated from each point of \mathcal{S} with velocity $(1/\tilde{\mathcal{P}})$, until a grid point \mathbf{x} such that $\mathcal{L}(\mathbf{x}) \geq \lambda$ is tagged as *Alive*. This point is then defined as the first keypoint, denoted \mathbf{p}_{n+1}^* (see Fig. 3).

Such a criterion has already been used in [7] to find a minimal path given only one endpoint. Assuming that the point \mathbf{p}_{n+1}^* belongs to the Voronoi region \mathcal{R}_j when it is detected, this criterion ensures that the minimal path $\mathcal{C}_{\mathbf{p}_j, \mathbf{p}_{n+1}^*}$

minimizes the integral of $\tilde{\mathcal{P}}$ (along itself) over all open curves with Euclidean lengths greater than or equal to λ and with endpoints in \mathcal{S} . Therefore, \mathbf{p}_{n+1}^* is likely to belong to the curve along which the values of $\tilde{\mathcal{P}}$ are low.

Once the first keypoint has been detected, it is defined as a new source of propagation. It is unnecessary to restart the overall algorithm since values of \mathcal{U} , \mathcal{V} and \mathcal{L} which have already been estimated would not differ in the vicinity of initial sources (i.e. in the vicinity of points of \mathcal{S}). In order to limit the computational cost, one just needs to reinitialize \mathcal{U} , \mathcal{V} and \mathcal{L} in the following manner :

$$\mathcal{U}(\mathbf{p}_{n+1}^*) := 0, \quad \mathcal{V}(\mathbf{p}_{n+1}^*) := n+1, \quad \mathcal{L}(\mathbf{p}_{n+1}^*) := 0,$$

tag \mathbf{p}_{n+1}^* as *Trial* and continue front propagation. However, without any additional modification of the original FMM, final values of \mathcal{U} , \mathcal{V} and \mathcal{L} would be incorrect for grid points which are tagged as *Alive* when \mathbf{p}_{n+1}^* is detected and closer (in the sense of a weighted distance) to \mathbf{p}_{n+1}^* than to the initial sources. These errors would be solely due to the fact that, in the original FMM, values of \mathcal{U} , \mathcal{V} and \mathcal{L} are frozen for *Alive* points. An easy way to avoid this problem is just to let an *Alive* point be tagged as *Trial* again if it is closer to the new source of propagation than to initial sources. This algorithmic trick enables the local correction of \mathcal{U} , \mathcal{V} and \mathcal{L} in the neighborhood of \mathbf{p}_{n+1}^* .

Next, front propagation is continued until a grid point \mathbf{x} such that $\mathcal{L}(\mathbf{x}) \geq \lambda$ is tagged as *Alive*. This point is defined as the second keypoint, denoted \mathbf{p}_{n+2}^* , and is added to the set of sources. Afterward, front propagation is continued, and so on. Thus, during the front propagation, keypoints are sequentially detected on the curve along which $\tilde{\mathcal{P}}$ takes low values (see Fig. 3).

Stopping criterion for keypoint detection and front propagation.

In order to prevent the algorithm from distributing keypoints over the whole domain Ω , one needs to stop the keypoint detection as soon as the domain visited by the fronts contains the curve of interest. Note that even if this curve is unknown, we assume that it is closed. This topological assumption is used to devise a relevant criterion for stopping keypoint detection and front propagation.

One possible strategy is to take into account the Voronoi partition, and to stop keypoint detection as soon as each Voronoi region is adjacent to at least two other Voronoi regions (i.e. as soon as there exists a cycle of Voronoi regions). This strategy, although correct, is limited to the 2D case. To get a scheme which may be extended to higher dimensions, another strategy is employed in the FMMWKD. Let us denote by Ω_F the domain visited by the propagating fronts, defined as the set of grid points which are not *Far* (i.e. the set of grid points which are either *Alive* or *Trial*). In the FMMWKD, keypoint detection is stopped as soon as

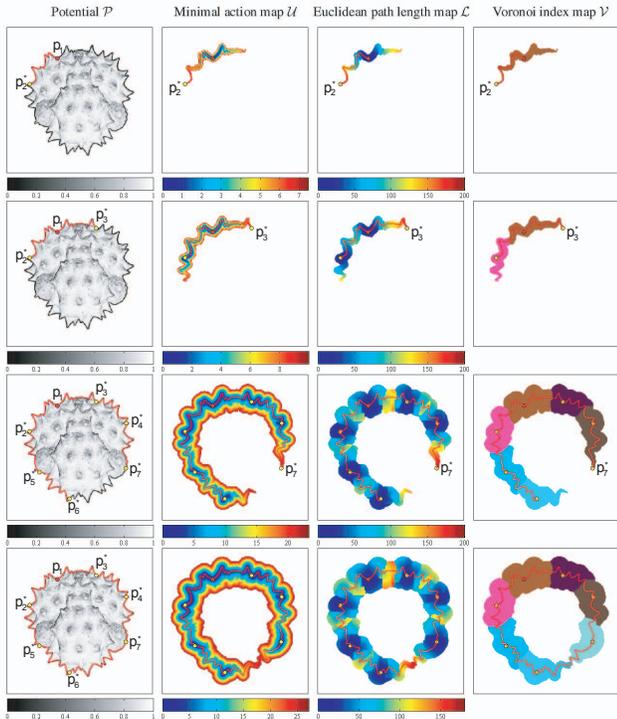


Figure 3. Intermediate and final results for the FMMWKD applied to the 2D potential of the Figure 3.b, with $S = \{p_1\}$ and $\lambda = 200$. The first, second and third rows show intermediate results obtained when are detected, respectively, p_2^* (the first keypoint), p_3^* (the second keypoint) and p_7^* (the last keypoint). The last row shows final results.

Ω_F becomes a simply connected subset of Ω delimited by exactly two simply connected boundaries.

The set Ω_F may be divided into two subsets : the set of interior points, denoted $\text{int}(\Omega_F)$, and the set of boundary points, denoted $\partial\Omega_F$. In the original FMM, $\text{int}(\Omega_F)$ and $\partial\Omega_F$ respectively correspond to the set of *Alive* points and the set of *Trial* points. This is no longer true in the FMMWKD because of the local correction of \mathcal{U} , \mathcal{V} and \mathcal{L} in the neighborhood of a keypoint. That is why a second labelling is introduced in the FMMWKD : each grid point which is not *Far*, in addition to being tagged as *Alive* or *Trial*, is also tagged as *Interior* or *Boundary* depending on whether it belongs to $\text{int}(\Omega_F)$ or $\partial\Omega_F$. Noting that the iteration of the marching loop at which $\text{int}(\Omega_F)$ becomes a simply connected subset of Ω is also the iteration at which the number of simply connected components of $\partial\Omega_F$ increases for the first time, we just need to monitor the topological changes of $\partial\Omega_F$.

In the algorithm detailed in Table 2, the stopping criterion for keypoint detection is satisfied as soon as the routine `IsBoundarySplit` returns *TRUE*. This routine is called after the grid point \mathbf{x}_{\min} is moved from the set of *Trial* points to the set of *Alive* points, once some of the $M = 4$ neigh-

Table 2 : *Fast Marching Method With keypoint Detection.*

- **Notation.**
 $\mathcal{N}_M(\mathbf{x})$ is the set of M neighbors of a grid point \mathbf{x} , where $M = 4$ in 2D and $M = 6$ in 3D. $\mathcal{N}_{M^+}(\mathbf{x})$ is the set of M^+ neighbors of a grid point \mathbf{x} , where $M^+ = 8$ in 2D and $M^+ = 26$ in 3D.
- **Initialization.**
For each grid point \mathbf{x} , do
Set $\mathcal{U}(\mathbf{x}) := +\infty$, $\mathcal{V}(\mathbf{x}) := 0$ and $\mathcal{L}(\mathbf{x}) := +\infty$.
Tag \mathbf{x} as *Far*.
For each source $\mathbf{p}_j \in \mathcal{S}$, do
Set $\mathcal{U}(\mathbf{p}_j) := 0$, $\mathcal{V}(\mathbf{p}_j) := j$ and $\mathcal{L}(\mathbf{p}_j) := 0$.
Tag \mathbf{p}_j as *Trial* and as *Boundary*.
 $m := 1$, *StopDetection* := *FALSE*.
- **Marching loop.**
While the set of *Trial* points is non-empty, do
Find \mathbf{x}_{\min} , the *Trial* point with the smallest \mathcal{U} value.
If (*StopDetection* = *FALSE*) and $(\mathcal{L}(\mathbf{x}_{\min}) \geq \lambda)$, do
Here, \mathbf{x}_{\min} is defined as the keypoint \mathbf{p}_{n+m}^* .
Set $\mathcal{U}(\mathbf{x}_{\min}) := 0$, $\mathcal{V}(\mathbf{x}_{\min}) := n + m$, $\mathcal{L}(\mathbf{x}_{\min}) := 0$.
 $m := m + 1$.
Else, do
Tag \mathbf{x}_{\min} as *Alive*.
For each grid point $\mathbf{x}_n \in \mathcal{N}_M(\mathbf{x}_{\min})$, do
If \mathbf{x}_n is not *Alive*, do
 $\{u, v, \ell\} := \text{UpdateSchemeFMM}(\mathbf{x}_n, \mathcal{N}_M(\mathbf{x}_n))$.
Set $\mathcal{U}(\mathbf{x}_n) := u$, $\mathcal{V}(\mathbf{x}_n) := v$ and $\mathcal{L}(\mathbf{x}_n) := \ell$.
If (*StopDetection* = *FALSE*) and (\mathbf{x}_n is *Far*), do
Tag \mathbf{x}_n as *Trial* and as *Boundary*.
Else if $\mathcal{V}(\mathbf{x}_n) \neq \mathcal{V}(\mathbf{x}_{\min})$, do
 $\{u, v, \ell\} := \text{UpdateSchemeFMM}(\mathbf{x}_n, \mathcal{N}_M(\mathbf{x}_n))$.
If $u < \mathcal{U}(\mathbf{x}_n)$, do
Set $\mathcal{U}(\mathbf{x}_n) := u$, $\mathcal{V}(\mathbf{x}_n) := v$ and $\mathcal{L}(\mathbf{x}_n) := \ell$.
Tag \mathbf{x}_n as *Trial*.
If \mathbf{x}_{\min} is *Boundary*, do
Tag \mathbf{x}_{\min} as *Interior*.
If *StopDetection* = *FALSE*, do
StopDetection :=
`IsBoundarySplit` ($\mathbf{x}_{\min}, \mathcal{N}_{M^+}(\mathbf{x}_{\min})$).

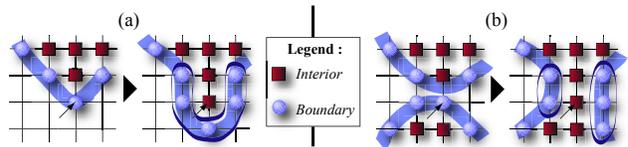


Figure 4. *Local test* applied in the vicinity of a grid point \mathbf{x}_{\min} (the point marked with an arrow) to detect a front collision. (a) \mathbf{x}_{\min} is a simple point of $\text{int}(\Omega_F)$ and $\#C(\partial\Omega_F \cap \mathcal{N}_{M^+}(\mathbf{x}_{\min})) = 1$. (b) Two fronts have collided in the neighborhood of \mathbf{x}_{\min} and $\#C(\partial\Omega_F \cap \mathcal{N}_{M^+}(\mathbf{x}_{\min})) = 2$.

bors of \mathbf{x}_{\min} have been tagged as *Boundary*. The routine `IsBoundarySplit` returns *TRUE* if both of the following tests are satisfied :

- **Local test for detecting a front collision.**

First, we check if some fronts collide in the vicinity of \mathbf{x}_{\min} . Let us denote by $\mathcal{N}_{M^+}(\mathbf{x}_{\min})$ the set of $M^+ = 8$ neighbors of \mathbf{x}_{\min} , and by $\partial\Omega_F \cap \mathcal{N}_{M^+}(\mathbf{x}_{\min})$ the set of points of $\mathcal{N}_{M^+}(\mathbf{x}_{\min})$ which are tagged as

Boundary. The local test simply relies on the computation of the number of 8-connected components of $\partial\Omega_F \cap \mathcal{N}_{M^+}(\mathbf{x}_{\min})$, denoted $\#C(\partial\Omega_F \cap \mathcal{N}_{M^+}(\mathbf{x}_{\min}))$. Most of the time, \mathbf{x}_{\min} is a simple point of $\text{int}(\Omega_F)$, and $\#C(\partial\Omega_F \cap \mathcal{N}_{M^+}(\mathbf{x}_{\min})) = 1$ (see Fig. 4.a). The local test is satisfied if $\#C(\partial\Omega_F \cap \mathcal{N}_{M^+}(\mathbf{x}_{\min})) > 1$, i.e. when there is a shock between some propagating fronts (see Fig. 4.b).

- **Global test for detecting a topological change of $\partial\Omega_F$.**

When the local test is satisfied, we need to check if the different components of $\partial\Omega_F \cap \mathcal{N}_{M^+}(\mathbf{x}_{\min})$ are also disconnected at a global scale. The global test is satisfied if the front collision has split an 8-connected component of $\partial\Omega_F$ into several 8-connected components.

Such a test is easy to implement. For instance, consider the case where $\#C(\partial\Omega_F \cap \mathcal{N}_{M^+}(\mathbf{x}_{\min})) = 2$. Let \mathbf{x}_1 and \mathbf{x}_2 be two grid points such that \mathbf{x}_1 belongs to the first component of $\partial\Omega_F \cap \mathcal{N}_{M^+}(\mathbf{x}_{\min})$ and \mathbf{x}_2 to the second. We just have to visit all grid points which belong to the same 8-connected component of $\partial\Omega_F$ as \mathbf{x}_1 , and assign to each visited point a temporary label. Then, the global test is satisfied if \mathbf{x}_2 has not been labeled.

Since the scheme used to detect the iteration at which the keypoint detection has to be stopped mainly requires tests at a local scale, it is considerably less computationally expensive than globally counting the number of connected components of $\text{int}(\Omega_F)$ and $\partial\Omega_F$ at each iteration of the marching loop. Moreover, note that special care is required to deal with the fact that a propagating front may reach the border of the domain Ω . We suggest adding virtual points along each border of the discrete grid and tagging as *Boundary* every virtual point in the neighborhood of an *Interior* point lying on the border of the grid. This ensures that any connected component of $\text{int}(\Omega_F)$ is completely delimited by a connected set of *Boundary* points.

Once the keypoint stopping criterion is satisfied, no more grid points are moved from the set of *Far* points to the set Ω_F , and computations are continued until correct values of \mathcal{U} , \mathcal{V} and \mathcal{L} have been assigned to each point of Ω_F . The front propagation is thus limited to a band surrounding the curve of interest.

3.2. 3D case : distribution of a set of points on a closed surface

Now, we consider the case where the domain Ω is a 3D domain, and we assume that we are given a few initial $\mathcal{S} = \{\mathbf{p}_1, \dots, \mathbf{p}_n\}$ points distributed on a closed surface along which a potential $\tilde{\mathcal{P}} : \Omega \rightarrow \mathbb{R}^{*+}$ takes lower values.

The FMMWKD, as it has been introduced in section 3.1, may be straightforwardly extended from a 2D to a 3D

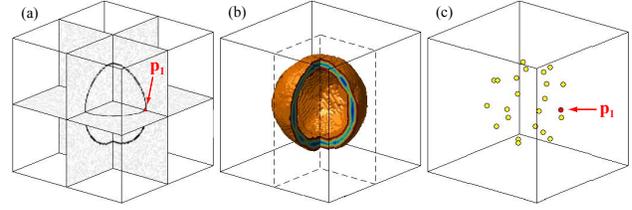


Figure 5. Distribution of a set of points on a sphere, by applying the FMMWKD to a 3D synthetic potential with $\mathcal{S} = \{\mathbf{p}_1\}$. (a) Cut view of Ω showing the initial point \mathbf{p}_1 and values of the potential $\tilde{\mathcal{P}}$. (b) Cut view showing values of the minimal action map \mathcal{U} inside Ω_F . (c) Set of points $\mathcal{S} \cup \mathcal{S}^*$.

framework, in order to distribute a set of points on a closed surface. The overall algorithm (see Table 2) is similar in 2D and 3D, with the difference that 4-connectivity and 8-connectivity on a 2D grid, becomes 6-connectivity and 26-connectivity on a 3D grid.

Therefore, fronts are propagated from each point of \mathcal{S} with velocity $(1/\tilde{\mathcal{P}})$ and, during the front propagation, keypoints are sequentially detected on the closed surface along which $\tilde{\mathcal{P}}$ takes low values. Keypoint detection and front propagation are led until a front collision splits a 26-connected component of $\partial\Omega_F$ into several 26-connected components. When the algorithm ends, we finally get a set of points $\mathcal{S} \cup \mathcal{S}^*$ distributed on the closed surface (see Fig. 5), but also the maps \mathcal{U} , \mathcal{V} and \mathcal{L} associated to the set of sources $\mathcal{S} \cup \mathcal{S}^*$. These maps are defined on a domain Ω_F , such that Ω_F is a simply connected subset of Ω delimited by two simply connected surfaces (an inner boundary and an outer boundary).

4. Geodesic meshing of a closed manifold

4.1. Extraction of a closed contour from a 2D image

The FMMWKD may be used to extract a closed contour from a 2D image I given a single contour point \mathbf{p}_1 in an easy and fast manner. Once a potential $\tilde{\mathcal{P}}$ has been derived from the image I such that $\tilde{\mathcal{P}}$ takes low values along the contour of interest, applying the FMMWKD with $\mathcal{S} = \{\mathbf{p}_1\}$ gives a set of points $\mathcal{S} \cup \mathcal{S}^*$, but also the maps \mathcal{U} and \mathcal{V} associated to $\mathcal{S} \cup \mathcal{S}^*$ and defined on a domain $\Omega_F \subset \Omega$. Linking pairs of points in the set $\mathcal{S} \cup \mathcal{S}^*$ by minimal paths finally enables the extraction of the desired contour.

Here, we explain first how to exploit the maps \mathcal{U} and \mathcal{V} to link two neighboring sources of $\mathcal{S} \cup \mathcal{S}^*$ with a minimal path. Then, we describe how to correctly choose the pairs of points of the set $\mathcal{S} \cup \mathcal{S}^*$ that should be linked, in order to build a cyclic sequence of minimal paths that follows the contour of interest.

Linking a pair of neighboring sources with a minimal path.

Consider a source $\mathbf{p}_i \in \mathcal{S} \cup \mathcal{S}^*$. The Voronoi region \mathcal{R}_i associated to \mathbf{p}_i may be deduced from the map $\mathcal{V} : \mathcal{R}_i = \{\mathbf{x} \in \Omega_F; \mathcal{V}(\mathbf{x}) = i\}$. If there exists a Voronoi region \mathcal{R}_j and a couple of grid points $(\mathbf{x}_i, \mathbf{x}_j) \in \mathcal{R}_i \times \mathcal{R}_j$ such that $(\mathbf{x}_i, \mathbf{x}_j)$ are 8-connected neighbors, then the Voronoi regions \mathcal{R}_i and \mathcal{R}_j are adjacent and the sources \mathbf{p}_i and \mathbf{p}_j are neighboring. In this case, the midpoint of the minimal path $\mathcal{C}_{\mathbf{p}_i, \mathbf{p}_j}$ may be approximated by a couple of grid points $(\mathbf{m}_{i|j}, \mathbf{m}_{j|i}) \in \mathcal{R}_i \times \mathcal{R}_j$. Among all pairs of grid points $(\mathbf{x}_i, \mathbf{x}_j) \in \mathcal{R}_i \times \mathcal{R}_j$ which are 8-connected neighbors, $(\mathbf{m}_{i|j}, \mathbf{m}_{j|i})$ is the one that minimizes the accumulated energy $\Sigma\mathcal{U}$ defined by

$$\Sigma\mathcal{U}(\mathbf{x}_i, \mathbf{x}_j) = \mathcal{U}(\mathbf{x}_i) + \mathcal{U}(\mathbf{x}_j) + \frac{h(\mathbf{x}_i, \mathbf{x}_j)}{2} (\tilde{\mathcal{P}}(\mathbf{x}_i) + \tilde{\mathcal{P}}(\mathbf{x}_j)),$$

where $h(\mathbf{x}_i, \mathbf{x}_j)$ denotes the spacing between the grid points \mathbf{x}_i and \mathbf{x}_j . Note that $\Sigma\mathcal{U}(\mathbf{m}_{i|j}, \mathbf{m}_{j|i})$ is the energy integrated along $\mathcal{C}_{\mathbf{p}_i, \mathbf{p}_j}$, i.e. the minimal weighted distance between the sources \mathbf{p}_i and \mathbf{p}_j . Once the grid points $\mathbf{m}_{i|j}$ and $\mathbf{m}_{j|i}$ have been found, the minimal paths $\mathcal{C}_{\mathbf{m}_{i|j}, \mathbf{p}_i}$ and $\mathcal{C}_{\mathbf{m}_{j|i}, \mathbf{p}_j}$ may be retrieved by performing two gradient descents on \mathcal{U} , respectively from $\mathbf{m}_{i|j}$ to \mathbf{p}_i and from $\mathbf{m}_{j|i}$ to \mathbf{p}_j . Since the two paths $\mathcal{C}_{\mathbf{m}_{i|j}, \mathbf{p}_i}$ and $\mathcal{C}_{\mathbf{m}_{j|i}, \mathbf{p}_j}$ are the two halves of $\mathcal{C}_{\mathbf{p}_i, \mathbf{p}_j}$, connecting them to each other finally gives the minimal path $\mathcal{C}_{\mathbf{p}_i, \mathbf{p}_j}$ that links the neighboring sources \mathbf{p}_i and \mathbf{p}_j .

Building a cyclic sequence of minimal paths to extract a closed contour.

Assuming that the set $\mathcal{S} \cup \mathcal{S}^*$ contains at least three points, linking each source of $\mathcal{S} \cup \mathcal{S}^*$ to the two closest neighboring sources (in the sense of a weighted distance) via minimal paths gives a cyclic sequence of minimal paths that follow the desired closed contour (see Fig. 6). Note that finding the two closest neighboring sources of a given source only relies on the minimal weighted distance between a pair of neighboring sources $(\mathbf{p}_i, \mathbf{p}_j)$, which is given by the accumulated energy $\Sigma\mathcal{U}(\mathbf{m}_{i|j}, \mathbf{m}_{j|i})$ defined above.

Results on 2D data.

In figure 6, we show the segmentation results on microscopy images, found on the ANSP Algae Image Database from the Phycology Section, Patrick Center for Environmental Research, The Academy of Natural Sciences at <http://diatom.acnatsci.org/AlgaeImage/>. For these images, we used a 1.6Ghz PC with 512 MB of RAM to obtain this segmentation in under a second.

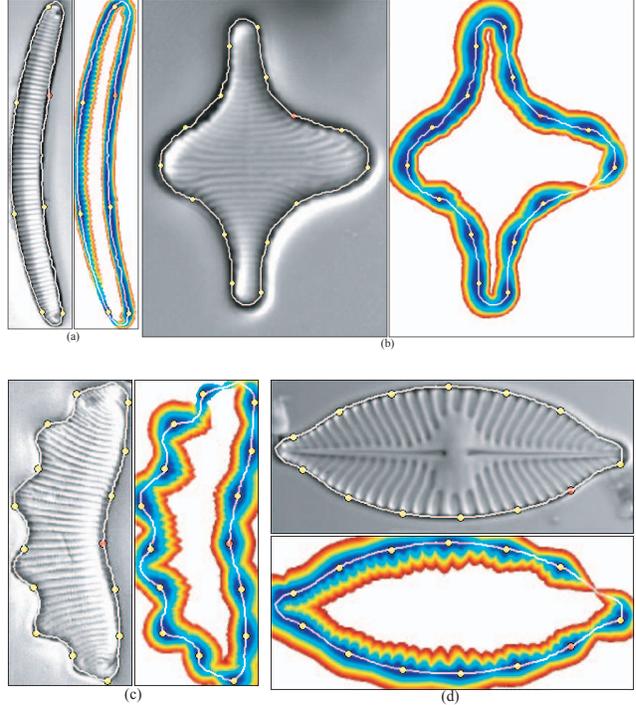


Figure 6. Extraction of a closed contour from a 2D microscopy image. Potential \mathcal{P} , set of sources $\mathcal{S} \cup \mathcal{S}^*$, Minimal action map and cyclic sequence of minimal paths. (a) Image size 101×521 , $\lambda = 180$. (b) 385×532 , $\lambda = 80$. (c) 153×380 , $\lambda = 60$. (d) 1032×435 , $\lambda = 160$.

4.2. Geodesic meshing of a closed surface from a 3D image

The FMMWKD may also be used to mesh a closed surface from a 3D image I and a single surface point. Consider a potential $\tilde{\mathcal{P}}$, derived from the image I which takes lower values along a closed surface, and a single surface point \mathbf{p}_1 . Applying the FMMWKD with $\mathcal{S} = \{\mathbf{p}_1\}$ generates a cloud of points $\mathcal{S} \cup \mathcal{S}^*$ distributed on the whole surface of interest. Moreover, the FMMWKD enables the computation of the maps \mathcal{U} and \mathcal{V} associated to the set $\mathcal{S} \cup \mathcal{S}^*$ and defined in a domain Ω_F that surrounds the underlying surface.

Then, it is quite straight forward to devise an algorithm for meshing the surface of interest, interpreting the cloud of points $\mathcal{S} \cup \mathcal{S}^*$ as a set of mesh vertices and the Voronoi partition of Ω_F as a way to derive the mesh connectivity. As in the 2D case, two points of the set $\mathcal{S} \cup \mathcal{S}^*$ may be considered as neighboring sources if their respective Voronoi regions are adjacent, and two neighboring sources may be linked with a minimal path by performing two gradient descents on \mathcal{U} from the path midpoint. Linking each pair of neighboring sources with a minimal path finally gives a *geodesic mesh* that describes the underlying surface (see Fig. 7 and Fig. 8).

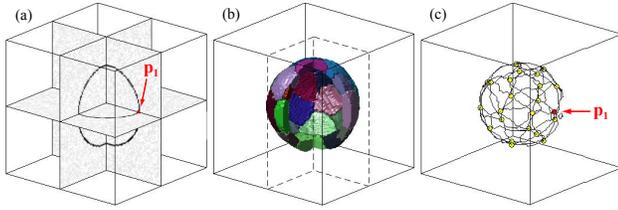


Figure 7. Geodesic meshing of a sphere from a 3D synthetic potential and a single surface point \mathbf{p}_1 . (a) Cut view of Ω showing the initial point \mathbf{p}_1 and values of the potential \mathcal{P} . (b) Cut view showing values of the Voronoi index map \mathcal{V} inside Ω_F . (c) Set of points $\mathcal{S} \cup \mathcal{S}^*$ and geodesic mesh.

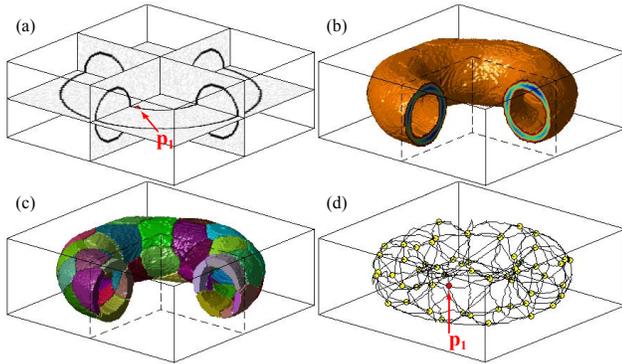


Figure 8. Geodesic meshing of a torus from a 3D synthetic potential and a single surface point \mathbf{p}_1 . (a) Cut view of Ω showing the initial point \mathbf{p}_1 and values of the potential \mathcal{P} . (b) Cut view showing values of the minimal action map \mathcal{U} inside Ω_F . (c) Cut view showing values of the Voronoi index map \mathcal{V} inside Ω_F . (d) Set of points $\mathcal{S} \cup \mathcal{S}^*$ and geodesic mesh.

5. Conclusion and Perspectives

We have presented a new fast front propagation approach for closed contour segmentation. Our method is interactive. At least one keypoint and the Euclidean length parameter λ have to be given by the user. The way the FMMWKD is built ensures that λ is an upper bound of the Euclidean path length map \mathcal{L} whenever a new keypoint is detected. Thus, the smaller the value given to λ is, the smaller the number of grid points visited during the front propagation is. In a sense, the FMMWKD may be regarded as a way to limit the front propagation to a small neighborhood around the manifold of interest. Better still, FMMWKD speeds up the segmentation process.

As seen in Section 4.1, it is straightforward to extract a closed contour from a 2D image. One may only get a mesh of minimal paths on a closed surface. Future work includes a new step based on a recent implicit method by Ardon et al. *et al.* [1], to obtain a complete closed surface.

References

- [1] R. Ardon and L. D. Cohen. Fast constrained surface extraction by minimal paths. *International Journal of Computer Vision*, **69**:127–136, 2006. 8
- [2] S. Bonneau, M. Dahan, and L. D. Cohen. Single quantum dot tracking based on perceptual grouping using minimal paths in a spatiotemporal volume. *IEEE Transactions on Image Processing*, **14**:1384–1395, 2005. 2
- [3] V. Caselles, R. Kimmel, and G. Sapiro. Geodesic active contours. *International Journal of Computer Vision*, **22**:61–79, 1997. 1
- [4] L. D. Cohen. Multiple contour finding and perceptual grouping using minimal paths. *Journal of Mathematical Imaging and Vision*, **14**:225–236, 2001. 2
- [5] L. D. Cohen. Minimal paths and fast marching methods for image analysis. In N. Paragios, Y. Chen, and O. Faugeras, editors, *Mathematical Models in Computer Vision: The Handbook*. Springer, 2005. 2
- [6] L. D. Cohen and R. Kimmel. Global minimum for active contour models: a minimal path approach. *International Journal of Computer Vision*, **24**:57–78, 1997. 1, 2
- [7] T. Deschamps and L. D. Cohen. Fast extraction of minimal paths in 3D images and applications to virtual endoscopy. *Medical Image Analysis*, **5**:281–299, 2001. 2, 4
- [8] E. W. Dijkstra. A note on two problems in connection with graphs. *Numerische Mathematic*, **1**:269–271, 1959. 3
- [9] M. Kass, A. Witkin, and D. Terzopoulos. Snakes: active contour models. *International Journal of Computer Vision*, **1**:321–331, 1988. 1
- [10] R. Kimmel and J. A. Sethian. Optimal algorithm for shape from shading and path planning. *Journal of Mathematical Imaging and Vision*, **14**:237–244, 2001. 2
- [11] G. Peyré and L. D. Cohen. Geodesic remeshing using front propagation. *International Journal of Computer Vision*, **69**:145–156, 2006. 2
- [12] E. Rouy and A. Tourin. A viscosity solution approach to shape from shading. *SIAM Journal on Numerical Analysis*, **29**:867–884, 1992. 3
- [13] J. A. Sethian. A fast marching level set for monotonically advancing fronts. *Proceedings of the National Academy of Sciences*, **93**:1591–1595, 1996. 3
- [14] J. A. Sethian. Fast marching methods. *SIAM Review*, **41**:199–235, 1999. 3
- [15] J. A. Sethian. *Level Set Methods and Fast Marching Methods*. Cambridge University Press, 1999. 3
- [16] J. N. Tsitsiklis. Efficient algorithms for globally optimal trajectories. *IEEE Transactions on Automatic Control*, **40**:1528–1538, 1995. 3
- [17] A. Yezzi, S. Kichenassamy, A. Kumar, P. Olver, and A. Tannenbaum. A geometric snake model for segmentation of medical imagery. *IEEE Transactions on Medical Imaging*, **16**:199–209, 1997. 1