Fast Object Segmentation by Growing Minimal Paths from a Single Point on 2D or 3D Images

Fethallah Benmansour · Laurent D. Cohen

Abstract In this paper, we present a new method for segmenting closed contours and surfaces. Our work builds on a variant of the minimal path approach. 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. This approach can also be used for finding an open curve giving extra information as stopping criteria. Detection of a variety of objects on real images is demonstrated. Using a similar 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.

F. Benmansour (✉) · L.D. Cohen
CEREMADE, UMR CNRS 7534, Université Paris Dauphine,
Place du Maréchal de Lattre de Tassigny, 75775 Paris Cedex 16,
France
e-mail: benmansour@ceremade.dauphine.fr

L.D. Cohen
e-mail: cohen@ceremade.dauphine.fr

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 [12], 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 smoothes the curve and an external part which guides the curve toward particular image features. For instance, the geodesic active contour model [5, 22] 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 [8] 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 interesting features 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 $(1/\tilde{\mathcal{P}})$. It satisfies the Eikonal equation and can be numerically solved efficiently using the Fast Marching Method. Moreover, as introduced in [9] we can compute simultaneously the Euclidean path length of the minimal path for each endpoint with the *Fast Marching Method*. This technic will be detailed in Sect. 2.2.

Fig. 1 (**a**) Original image.
(**b**) Build potential based on the
Laplacian. (**c**) Short cut.
(**d**) Minimal path with keypoints
detection can avoid short cut

Still, the minimal path technique may fail if the potential is too noisy, or not enough contrasted or if the objects of interest are long thin curvy curves. In this case a portion of the minimal path can be a short cut to the starting point, leading to wrong results, like in Fig. 1. We propose in this paper a new approach to avoid this drawback. A first version of this method was described in [2, 3]. The idea is to detect recursively new source points (called keypoints) along the curve of interest between the two given points. These points are automatically detected using a criterion based on the Euclidean length of minimal paths and are almost equi-distributed along the curve of interest. In [9] it was proposed to use the length of the minimal path in order to find the second extremity of the path. Here we iterate this idea many times in order to obtain intermediate points along the contour. Since the front propagates faster on the features, i.e. potential is smaller, 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. A good choice of parameter λ enables the front to propagates further in the direction of the thin curve without propagating in all directions and thus to reduce considerably the visited domain on the image.

In Sect. 2, we give some background on minimal path, Fast Marching, and finding the Euclidean length of the minimal path. In Sect. 3, we introduce a novel front propagation approach, based on the Fast Marching Method, to distribute a set of points on a codimension-1 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 as described before. By using the a-priori knowledge on the topology of the closed manifold, we devise a relevant criterion for stopping the keypoint detection and front propagation. For open curves, we propose in Sect. 3.2 two different stopping criteria. The first one needs more interaction by specifying an endpoint to reach and the second one is based on an approximation of the total contour length. These criteria are general for any dimension for closed manifolds or elongated structures. In Sect. 3.4, 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. In Sect. 3.5, we explain briefly the influence of the path length parameter λ. In Sect. 4, we explain how we can extend the previous ideas to 3D by using the right connectivity associated to the dimension. Segmentation results on a set of 2D and 3D images are presented. Finally, conclusions and perspectives follow in Sect. 5.

2 Background on Minimal Paths

2.1 Definitions

Given a 2D image $I : \Omega \to \mathbb{R}^+$ and two points $\mathbf{p}_1$ and $\mathbf{p}_2$, the underlying idea introduced by Cohen and Kimmel [8] is to build a potential $\mathcal{P} : \Omega \to \mathbb{R}^{*+}$ 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

Fig. 2 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**) Min- imal 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 $\mathcal{C}_{\mathbf{p}_1, \mathbf{p}_2}$

define $\mathcal{P}$ as a decreasing function of $\|\nabla I\|$ to extract image edges by finding a curve that globally minimizes the energy functional $E : \mathcal{A}_{\mathbf{p}_1, \mathbf{p}_2} \to \mathbb{R}^+$

$$E(\gamma) = \int_\gamma \{\mathcal{P}(\gamma(s)) + w\} \mathrm{d}s = \int_\gamma \tilde{\mathcal{P}}(\gamma(s)) \mathrm{d}s, \tag{1}$$

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 is a positive constant 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 min- imization problem is obtained through the computation of the *minimal action map* $\mathcal{U}_1 : \Omega \to \mathbb{R}^+$ 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:

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

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

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

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 func- tional (1) (see [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. 2), solving the following ordinary differential equation with standard numerical methods like Heun's or Runge-Kutta's:

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

Let us extend the definitions given so far to the case of mul- tiple 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

$\tilde{\mathcal{P}} : \Omega \to \mathbb{R}^{*+}$ and the set of n sources $\mathcal{S} = \{\mathbf{p}_1, \ldots, \mathbf{p}_n\}$ is the function $\mathcal{U} : \Omega \to \mathbb{R}^+$ defined by

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

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

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

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

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 the weighted distance) to $\mathbf{p}_j$ than to any other source of $\mathcal{S}$:

$$\mathcal{R}_j = \{\mathbf{x} \in \Omega; \mathcal{U}_j(\mathbf{x}) \le \mathcal{U}_i(\mathbf{x}), \ \forall i \in \{1, \ldots, n\}, \ i \ne j\}. \tag{7}$$

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 $\mathcal{V} : \Omega \to \{1, \ldots, n\}$ that assigns to any point of the domain Ω the index of its Voronoi region:

$$\forall \mathbf{x} \in \mathcal{R}_j, \quad \mathcal{V}(\mathbf{x}) = j. \tag{8}$$

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 the weighted distance.

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

$$\forall \mathbf{x} \in \mathcal{R}_j, \quad \mathcal{L}(\mathbf{x}) = \int_{\mathcal{C}_{\mathbf{p}_j, \mathbf{x}}} \mathrm{d}s. \tag{9}$$

Table 1 Fast Marching Method for solving (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

          $\{u, v, \ell\} := \mathtt{UpdateSchemeFMM}\,(\mathbf{x}_n, \mathcal{N}_M(\mathbf{x}_n))$ (see text).

          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*.

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 [7] for a review), such as path planning [14], contour completion [6], tubular surface extraction [9], motion tracking [4], or remeshing of triangulated manifolds [15].

2.2 Fast Marching Method

The *Fast Marching Method* (FMM) is a numerical method introduced by Sethian in [17–19] and Tsitsiklis in [20] for efficiently solving the isotropic Eikonal equation on a Cartesian grid. In (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 (6) and computing the maps $\mathcal{U}$, $\mathcal{V}$ and $\mathcal{L}$ in a common computational framework (see Table 1).

2.2.1 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 [10]. 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 *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 $\mathcal{O}(N \log N)$, where N is the total number of grid points. Some authors [11, 13, 21] have proposed various versions of the fast marching in order to get an $\mathcal{O}(N)$ complexity. These approaches are less accurate on the result or assuming some hypotheses that are not met in our method.

2.2.2 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 output of routine $\mathtt{UpdateSchemeFMM}$ 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. Although with a different approach, Rouy and Tourin [16] showed that the correct viscosity solution for $\mathcal{U}_{i,j}$ is given by the following first order accurate scheme:

$$\left(\frac{\max\{(\mathcal{U}_{i,j} - \mathcal{U}_{i-1,j}), (\mathcal{U}_{i,j} - \mathcal{U}_{i+1,j}), 0\}}{h_x} \right)^2$$
$$+ \left(\frac{\max\{(\mathcal{U}_{i,j} - \mathcal{U}_{i,j-1}), (\mathcal{U}_{i,j} - \mathcal{U}_{i,j+1}), 0\}}{h_y} \right)^2$$
$$= (\tilde{\mathcal{P}}_{i,j})^2. \tag{10}$$

Fig. 3 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

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 (10), 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. 3). For each triangle, we attempt to solve two quadratic equations to get estimations of $\mathcal{U}_{i,j}$, $\mathcal{V}_{i,j}$ and $\mathcal{L}_{i,j}$. These estimations are respectively noted u, v and ℓ. As an example, consider the triangle with vertices $\{(i-1,j),(i,j),(i,j-1)\}$. Four cases may occur:

- If neither $(i-1,j)$ nor $(i,j-1)$ are *Alive*, information is not propagating through this triangle. Then, $u=+\infty$, $v=0$ and $\ell=+\infty$.

- If $(i-1,j)$ is *Alive* but $(i,j-1)$ is not, then $v=\mathcal{V}_{i-1,j}$. Furthermore, u and ℓ are the largest solutions of quadratic equations

$$\left(\frac{u-\mathcal{U}_{i-1,j}}{h_x}\right)^2=(\tilde{\mathcal{P}}_{i,j})^2 \quad \text{and} \quad \left(\frac{\ell-\mathcal{L}_{i-1,j}}{h_x}\right)^2=1.$$

- If $(i,j-1)$ is *Alive* but $(i-1,j)$ is not, then $v=\mathcal{V}_{i,j-1}$. Furthermore, u and ℓ are the largest solutions of quadratic equations

$$\left(\frac{u-\mathcal{U}_{i,j-1}}{h_y}\right)^2=(\tilde{\mathcal{P}}_{i,j})^2 \quad \text{and} \quad \left(\frac{\ell-\mathcal{L}_{i,j-1}}{h_y}\right)^2=1.$$

- If both $(i-1,j)$ and $(i,j-1)$ are *Alive*, then u is the largest real solution of quadratic equation

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

ℓ is the largest real solution of quadratic equation

$$\left(\frac{\ell-\mathcal{L}_{i-1,j}}{h_x}\right)^2+\left(\frac{\ell-\mathcal{L}_{i,j-1}}{h_y}\right)^2=1,$$

and v is chosen according to the smallest value of $\mathcal{U}$:

$$v=\begin{cases} \mathcal{V}_{i-1,j} & \text{if } \mathcal{U}_{i-1,j}\le\mathcal{U}_{i,j-1}, \\ \mathcal{V}_{i,j-1} & \text{otherwise.} \end{cases}$$

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. 3), but the 3D update scheme is very similar to the 2D one.

3 Distribution of a Set of Keypoints on a 2D Curve

3.1 Keypoints Detection and Local Update

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

We propose in this paper a new approach. The idea is to detect recursively new source points (called keypoints) along the curve of interest. These points are automatically detected using a criterion based on the Euclidean length of minimal paths and are almost equi-distributed along features of interest. The method is called Minimal Path method With Keypoint Detection (MPWKD, see Table 2). A front is propagated from each source point of $\mathcal{S}$ with velocity $(1/\tilde{\mathcal{P}})$. Keypoints are sequentially detected, on the curve of interest, during the front propagation and stored in a set $\mathcal{S}^*=\{\mathbf{p}_{n+1}^*,\ldots,\mathbf{p}_{n+m}^*\}$. 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. Initially, fronts are propagated from each point of $\mathcal{S}$ with velocity $(1/\tilde{\mathcal{P}})$, until a grid point $\mathbf{x}$ such that $\mathcal{L}(\mathbf{x})\ge\lambda$ is tagged as *Alive*. This point is then defined as the first keypoint, denoted $\mathbf{p}_{n+1}^*$ (see Fig. 6). Such a criterion has already been used in [9] 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

Table 2 Minimal path method with keypoint detection

- **Notation**

  $\mathcal{N}_{\text{M}}(\mathbf{x})$ is the set of M neighbors of a grid point $\mathbf{x}$, where M $= 4$ in 2D and M $= 6$ in 3D. $\mathcal{N}_{\text{M}^+}(\mathbf{x})$ 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 $\mathcal{U}(\mathbf{x}_{\min}) := 0$, $\mathcal{V}(\mathbf{x}_{\min}) := n + m$ and $\mathcal{L}(\mathbf{x}_{\min}) := 0$.

          $m := m + 1$.

      Else, do

          Tag $\mathbf{x}_{\min}$ as *Alive*.

          For each grid point $\mathbf{x}_{\text{n}} \in \mathcal{N}_{\text{M}}(\mathbf{x}_{\min})$, do

            If $\mathbf{x}_{\text{n}}$ is not *Alive*, do

              $\{u, v, \ell\} := \texttt{UpdateSchemeFMM}\,(\mathbf{x}_{\text{n}}, \mathcal{N}_{\text{M}}(\mathbf{x}_{\text{n}}))$.

              Set $\mathcal{U}(\mathbf{x}_{\text{n}}) := u$, $\mathcal{V}(\mathbf{x}_{\text{n}}) := v$ and $\mathcal{L}(\mathbf{x}_{\text{n}}) := \ell$.

              If (*StopDetection* = *FALSE*) and ($\mathbf{x}_{\text{n}}$ is *Far*), do

                  Tag $\mathbf{x}_{\text{n}}$ as *Trial* and as *Boundary*.

            Else if $\mathcal{V}(\mathbf{x}_{\text{n}}) \neq \mathcal{V}(\mathbf{x}_{\min})$, do (*local correction of the maps if needed*)

              $\{u, v, \ell\} := \texttt{UpdateSchemeFMM}(\mathbf{x}_{\text{n}}, \mathcal{N}_{\text{M}}(\mathbf{x}_{\text{n}}))$.

              If $u < \mathcal{U}(\mathbf{x}_{\text{n}})$, do

                  Set $\mathcal{U}(\mathbf{x}_{\text{n}}) := u$, $\mathcal{V}(\mathbf{x}_{\text{n}}) := v$ and $\mathcal{L}(\mathbf{x}_{\text{n}}) := \ell$.

                  Tag $\mathbf{x}_{\text{n}}$ as *Trial*.

      If $\mathbf{x}_{\min}$ is *Boundary*, do

          Tag $\mathbf{x}_{\min}$ as *Interior*.

          If *StopDetection* = *FALSE*, do

              StopDetection := $\texttt{StoppingCriterion}\,(\mathbf{x}_{\min})$.

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. Next, the process in iterated and assuming ℓ keypoints $\mathbf{p}^*_{n+1}, \ldots, \mathbf{p}^*_{n+\ell}$ have been added, front propagation from this set and $\mathcal{S}$ is made with velocity $1/\tilde{\mathcal{P}}$. 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 $\ell + 1$ keypoint, denoted $\mathbf{p}^*_{n+\ell+1}$, 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 Figs. 4, 5 and 6). At each iteration of the process, it is not necessary to start again the whole computation from scratch 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 update the value $\mathcal{U}$ and go on computation till criterion is reached. For this update, we reinitialize $\mathcal{U}$, $\mathcal{V}$ and $\mathcal{L}$ in the following manner:

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

tag $\mathbf{p}^*_{n+\ell}$ 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 $\mathbf{p}^*_{n+\ell}$ is detected and closer (in the sense of a weighted distance) to $\mathbf{p}^*_{n+\ell}$ than to the initial sources. These points cannot be updated 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 update of $\mathcal{U}$, $\mathcal{V}$ and $\mathcal{L}$ in the neighborhood of $\mathbf{p}^*_{n+\ell}$. Notice that in [6], in a different context, there was also an algorithmic trick to update the minimal action map in order to save computation time.

3.2 Stopping Criterion on an Open Curve

In this section we propose two different stopping criteria for keypoint detection and front propagation on an open curve. The first simple case is when the user could provide more than a single source point, for example a source point and a destination. One can stop the algorithm as soon as the destination is reached (as was shown in Fig. 1). The user has to provide at least two points, a starting point and a destination (two simple clicks), and the path length parameter λ to avoid short cut. This approach needs minimal user interaction and might be used for many different applications on 2D images.

We also propose another stopping criterion based on the *total path length*. Here we suppose that only one single point is provided by the user, i.e. $\mathcal{S} = \{\mathbf{p}_1\}$ and the idea is to detect automatically the keypoints until a total minimal path length is reached. Such a criterion has already been used in [9] to find a minimal path given only one endpoint. In Sect. 3.1, we explained how one can compute, on a common framework, the minimal action map $\mathcal{U}$ and the minimal path length map $\mathcal{L}$ with respect to the set $\mathcal{S} \cup \mathcal{S}^*$. A good way to compute an approximation of the total minimal path length $\mathcal{L}_T$ consists on solving the same Euclidean Eikonal equation: $\|\nabla \mathcal{L}_T\| = 1$ but without the local update as done is section 3.1, since this local update gives the minimal path length

Fig. 4 Approximation of the total minimal path length. (**a**) A line potential with contrast equal to 2: the velocity on the line is twice the velocity on the background. (**b**) The minimal action map $\mathcal{U}$ associated to some distributed keypoints. (**c**) The Voronoi map $\mathcal{V}$. (**d**) The minimal path length map $\mathcal{L}$. (**e**) The real total minimal path length map to the top source point. (**f**) The approximated total minimal path length map to the top source point as explained on Sect. 3.2. (**g**) Approximation error of the total minimal path length

Fig. 5 Stopping criteria based on the total length. (**a**) An image potential build from the image of Fig. 1. (**b**) The minimal action map $\mathcal{U}$ associated to some distributed keypoints with $\lambda = 30$. (**c**) The Voronoi map $\mathcal{V}$. (**d**) The minimal path length map $\mathcal{L}$. (**e**) The approximated total minimal path length map to the red source. (**f**) Detected keypoints and minimal paths

map to the set of points $\mathcal{S} \cup \mathcal{S}^*$, and one wants the minimal path length with respect to the source point $\mathbf{p}_1$. This computation gives us a good approximation of the total path length on elongated structures (see Fig. 4). The error is maximal on the perpendicular direction of the elongated structure and almost nil along the structure. On Fig. 5, the user provides a single source point, the path length parameter λ and a total length parameter to reach. Thus one can find features of interest by a minimal interaction consisting in a single click and two parameters.

3.3 Stopping Criterion for a Closed Curve

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.

A first 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, a second strategy is proposed. 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 MPWKP, keypoint detection is stopped as soon as the border of Ω_F is delimited by exactly two connected subsets. 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 MPWKP because of the local correction of $\mathcal{U}$,

Fig. 6 Intermediate and final results for the MPWKP applied to the 2D potential of the Fig. 2(**b**) with $\mathcal{S} = \{\mathbf{p}_1\}$ and $\lambda = 200$. The first, second and third rows show intermediate results obtained when are detected, respectively, $\mathbf{p}_2^*$ (the first keypoint), $\mathbf{p}_3^*$ (the second keypoint) and $\mathbf{p}_7^*$ (the last keypoint). The last row shows final results with 7 keypoints
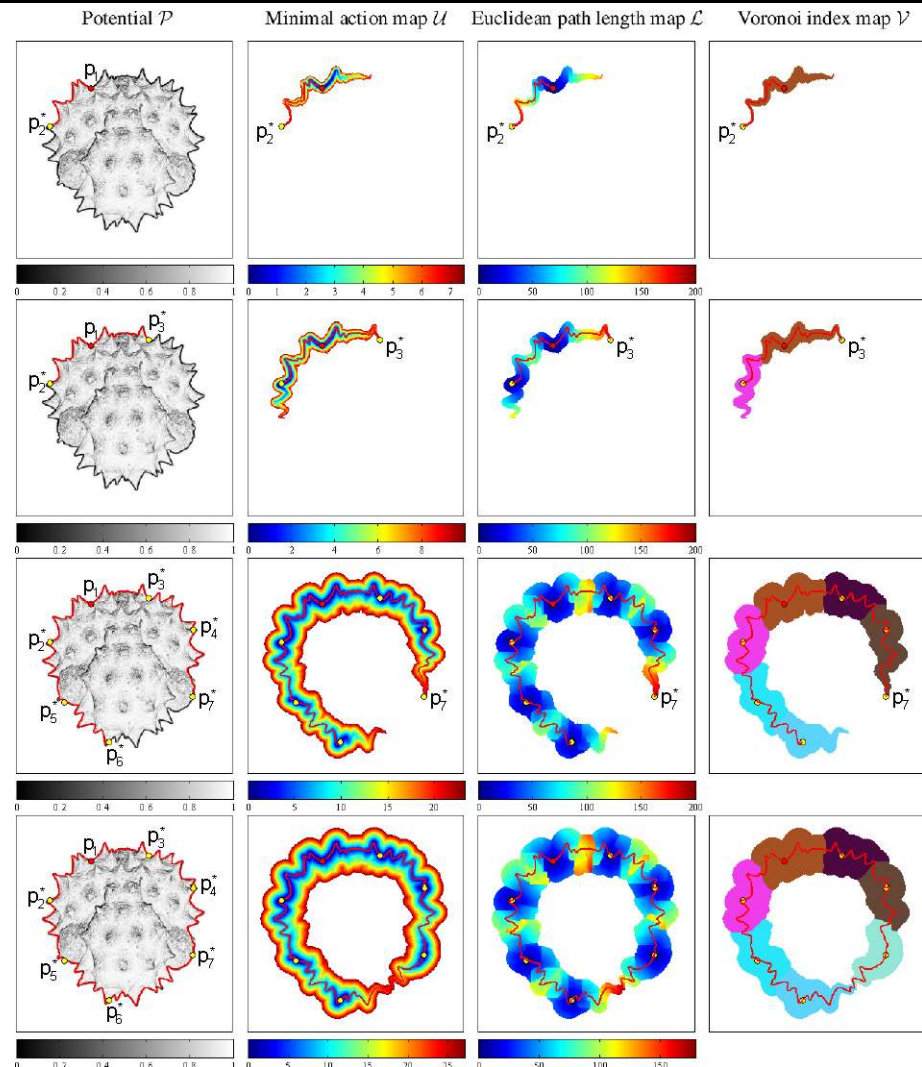
$\mathcal{V}$ and $\mathcal{L}$ in the neighborhood of a keypoint. That is why a second labelling is introduced in the MPWKP: each grid point which is not *Far*, in addition to being tagged as *Alive* or *Trial*, is also tagged as *Interior* or *Boundary* depending on whether it belongs to $\text{int}(\Omega_F)$ or $\partial\Omega_F$. In the case of one source point, it is clear that the visited domain Ω_F surrounds the object of interest when its boundary $\partial\Omega_F$ splits into exactly two connected subsets. In the case of multiple sources, the number of connected boundaries of $\partial\Omega_F$ first decreases then increases when the band surrounding the object of interest is obtained. Thus, 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 `StoppingCriterion` 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 $\text{M} = 4$ neighbors of $\mathbf{x}_{\min}$ have been tagged as *Boundary*. The rou-

tine `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 $\mathcal{N}_{\text{M}^+}(\mathbf{x}_{\min})$ the set of $\text{M}^+ = 8$ neighbors of $\mathbf{x}_{\min}$, and by $\partial\Omega_F \cap \mathcal{N}_{\text{M}^+}(\mathbf{x}_{\min})$ the set of points of $\mathcal{N}_{\text{M}^+}(\mathbf{x}_{\min})$ which are tagged as *Boundary*. The local test simply relies on the computation of the number of 8-connected components of $\partial\Omega_F \cap \mathcal{N}_{\text{M}^+}(\mathbf{x}_{\min})$, denoted $\#C(\partial\Omega_F \cap \mathcal{N}_{\text{M}^+}(\mathbf{x}_{\min}))$. Most of the time, $\mathbf{x}_{\min}$ is a simple point of $\text{int}(\Omega_F)$, and $\#C(\partial\Omega_F \cap \mathcal{N}_{\text{M}^+}(\mathbf{x}_{\min})) = 1$ (see Fig. 7(a)). The local test is satisfied if $\#C(\partial\Omega_F \cap \mathcal{N}_{\text{M}^+}(\mathbf{x}_{\min})) > 1$, i.e. when there is a shock between some propagating fronts (see Fig. 7(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 $\partial\Omega_F \cap \mathcal{N}_{\text{M}^+}(\mathbf{x}_{\min})$ are also disconnected at a global scale. The global test is satisfied

Fig. 7 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 $int(\Omega_F)$ 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 $\#C(\partial\Omega_F \cap \mathcal{N}_{M^+}(\mathbf{x}_{min})) = 2$

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 $int(\Omega_F)$ 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 $int(\Omega_F)$ 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.

3.4 Extraction of a Contour from a 2D Image

In the previous sections we gave a way to obtain a set of keypoints that describe a contour curve, given an initial set of points. Our initial goal is to obtain a contour as a set of minimal paths that link pairs of points. We explain in this section how this process takes place. The MPWKD may be used to extract a closed or an open 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 MPWKD 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.

3.4.1 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 $(\tilde{\mathbf{x}}_i, \tilde{\mathbf{x}}_j) \in \mathcal{R}_i \times \mathcal{R}_j$. 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, $(\tilde{\mathbf{x}}_i, \tilde{\mathbf{x}}_j)$ 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 $\Sigma\mathcal{U}(\tilde{\mathbf{x}}_i, \tilde{\mathbf{x}}_j)$ 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 $\tilde{\mathbf{x}}_i$ and $\tilde{\mathbf{x}}_j$ have been found, the minimal paths $\mathcal{C}_{\tilde{\mathbf{x}}_i, \mathbf{p}_i}$ and $\mathcal{C}_{\tilde{\mathbf{x}}_j, \mathbf{p}_j}$ may be retrieved by performing two gradient descents on $\mathcal{U}$, respectively from $\tilde{\mathbf{x}}_i$ to $\mathbf{p}_i$ and from $\tilde{\mathbf{x}}_j$ to $\mathbf{p}_j$. Since the two paths $\mathcal{C}_{\tilde{\mathbf{x}}_j, \mathbf{p}_i}$ and $\mathcal{C}_{\tilde{\mathbf{x}}_j, \mathbf{p}_j}$ 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 $\mathbf{p}_j$.

3.4.2 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 neighbor-

ing 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. 8). 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 $\Sigma\mathcal{U}(\tilde{\mathbf{x}}_i, \tilde{\mathbf{x}}_j)$ defined above.

3.4.3 Results on 2D Data

In Fig. 8, 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. For these images, we used a 1.6 Ghz PC with 512 MB of RAM to obtain this segmentation in under a second.

3.5 Influence of the Parameter λ on the Results of the MPWKD

The influence of the parameter λ on the final results of the MPWKD is twofold (see Fig. 9). On the one hand, tuning

λ is a direct way to control the density of keypoints. Since λ corresponds to the spacing between two successive keypoints along the curve of interest, the final number of keypoints is inversely proportional to the value given to λ. On the other hand, tuning λ is an indirect way to control the width of the band in which the front propagation is limited, i.e. an indirect way to control the area of the domain Ω_F. The way the MPWKD is built ensures that λ is an upper bound of the Euclidean path length map $\mathcal{L}$ whenever a new keypoint is detected. Nevertheless, if λ is too large, short cuts may occur, and if λ is smaller than the width of the structure one wants to extract, many keypoints will be detected on it. Thus, the smaller the value given to λ is, the smaller the number of grid points visited during the front propagation is and the faster the algorithm is. In a sense, the MPWKD may be regarded as a way to limit the front propagation to a small neighborhood around the manifold of interest.
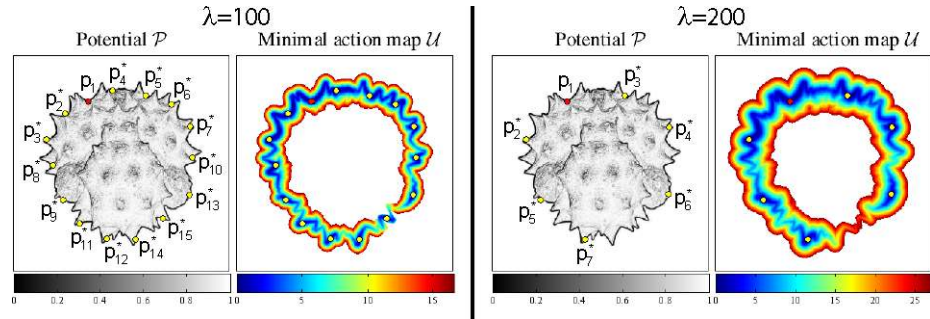
4 Distribution of a Set of Point on a Surface or on a 3D Elongated Structure

4.1 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 points $\mathcal{S} = \{\mathbf{p}_1, \ldots, \mathbf{p}_n\}$ distributed on a closed surface along which a potential $\tilde{\mathcal{P}} : \Omega \to \mathbb{R}^{*+}$ takes lower values.

The MPWKD, as it has been introduced in Sect. 3.1, may be straightforwardly extended from a 2D to a 3D 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. 10), 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 connected subset of Ω delimited by two connected surfaces (an inner boundary and an outer boundary). We will see in Sect. 4.3 how to use these points.

4.2 Distribution of a Set of Keypoints on a 3D Elongated Structure

Ideas presented on Sect. 3.2 can be extended to 3D straightforwardly. Then, one can use this algorithm to detect 3D elongated structure by using the proposed criteria. The first criterion needs more interaction, so the user must provide a source point a destination point and path length parameter λ. For the second criterion, the user provides at least one source point, the parameter λ and a total path length to reach, which corresponds to the maximum length of arteries for example (see Fig. 11). We conclude that using this method, with minimal interaction, one can have a first good segmentation on elongated structures like arteries.

4.3 Geodesic Meshing of a Closed Surface from a 3D Image

Once we obtain a set of keypoints that describe the desired surface, using minimal paths that link pairs of points as in Sect. 3.4 does not lead to a surface but to a mesh of surface. 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 MPWKD 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 MPWKD 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 straightforward 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

Fig. 11 (Color online)
Distribution of a set of points on
a tubular structure. (**a**) Original
image used as potential.
(**b**) Values of the minimal action
map $\mathcal{U}$ on the visited domain.
(**c**) The front of the visited
domain, with detected keypoints
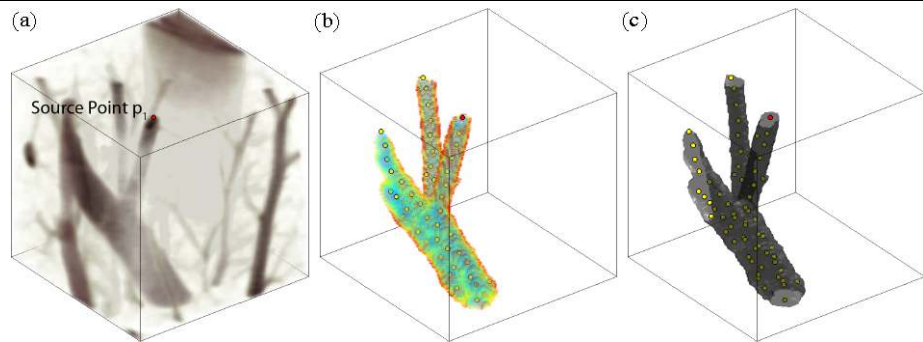(*yellow*) and the source point
(*red*)

Fig. 12 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
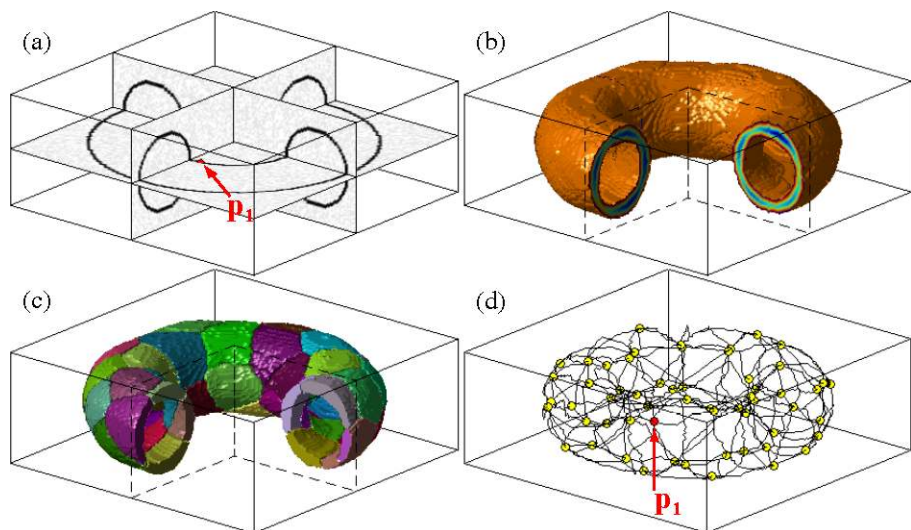

Fig. 13 Geodesic meshing of a
torus from a 3D synthetic
potential and a single surface
point $\mathbf{p}_1$. (**a**) Cut view of Ω
showing the initial point $\mathbf{p}_1$ and
values of the potential $\mathcal{P}$.
(**b**) Cut view showing values of
the minimal action map $\mathcal{U}$
inside Ω_F. (**c**) 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

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 Figs. 12 and 13).

As seen in Sect. 3.4, 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. [1], to obtain a complete closed surface.

5 Conclusion

In this paper we have proposed a new method to segment 2D object boundaries or to segment partially 3D object boundaries. Our method needs minimal interaction: a single source point and one or two real parameters are required, depending on the object topology. From a set of provided source points, which might be reduced to a single source point, a front is propagated and new keypoints are iteratively detected on the boundary. The detection of these keypoints is based on the Euclidean length of minimal path to the set of source points with respect to a given metric depending on

the image. A tricky algorithm has been introduced allowing local correction of the maps. Thus, computation time is considerably reduced. We proposed different stopping criteria. The first one is based on the total Euclidean path length from the sources. This criterion might be applied for 2D open object boundaries or 3D elongated structures and requires a total path length parameter to reach provided by the user. The second criterion is very simple. An endpoint is required, and the front propagation stops when this endpoint is reached. Finally, a specific stopping criterion for 2D or 3D closed object boundaries is proposed. This criterion is based on the front topology. The keypoints detection and front propagation is stopped when the visited domain surrounds the object of interest. Then, a few parameters as a source point and one or two real values can be enough to generate a coherent object boundary segmentation.

References

1. Ardon, R., Cohen, L.D., Yezzi, A.: A new implicit method for surface segmentation by minimal paths in 3d images. Appl. Math. Optim. **55**(18), 127–144 (2007)
2. Benmansour, F., Bonneau, S., Cohen, L.D.: Finding a closed boundary by growing minimal paths from a single point on 2d or 3d images. In: MMBIA07, pp. 1–8 (2007)
3. Bonneau, S.: Chemins minimaux en Analyse d'images: Nouvelles contributions et applications à l'imagerie biologique. PhD thesis, Université Paris-Dauphine, France (2006)
4. Bonneau, S., Dahan, M., Cohen, L.D.: Single quantum dot tracking based on perceptual grouping using minimal paths in a spatiotemporal volume. IEEE Trans. Image Process. **14**, 1384–1395 (2005)
5. Caselles, V., Kimmel, R., Sapiro, G.: Geodesic active contours. Int. J. Comput. Vis. **22**, 61–79 (1997)
6. Cohen, L.D.: Multiple contour finding and perceptual grouping using minimal paths. J. Math. Imaging Vis. **14**, 225–236 (2001)
7. Cohen, L.D.: Minimal paths and fast marching methods for image analysis. In: Paragios, N., Chen, Y., Faugeras, O. (eds.) Mathematical Models in Computer Vision: The Handbook. Springer, Berlin (2005)
8. Cohen, L.D., Kimmel, R.: Global minimum for active contour models: a minimal path approach. Int. J. Comput. Vis. **24**, 57–78 (1997)
9. Deschamps, T., Cohen, L.D.: Fast extraction of minimal paths in 3D images and applications to virtual endoscopy. Med. Image Anal. **5**, 281–299 (2001)
10. Dijkstra, E.W.: A note on two problems in connection with graphs. Numer. Math. **1**, 269–271 (1959)
11. Falcão, A.X., Udupa, J.K., Miyazawa, F.K.: An ultra-fast user-steered image segmentation paradigm: Live-wire-on-the-fly. IEEE Trans. Med. Imaging **19**(1), 55–62 (2000)
12. Kass, M., Witkin, A., Terzopoulos, D.: Snakes: active contour models. Int. J. Comput. Vis. **1**, 321–331 (1988)
13. Kim, S.: An $\mathcal{O}(N)$ level set method for Eikonal equations. SIAM J. Sci. Comput. **22**, 2178–2193 (2001)
14. Kimmel, R., Sethian, J.A.: Optimal algorithm for shape from shading and path planning. J. Math. Imaging Vis. **14**, 237–244 (2001)
15. Peyré, G., Cohen, L.D.: Geodesic remeshing using front propagation. Int. J. Comput. Vis. **69**, 145–156 (2006)
16. Rouy, E., Tourin, A.: A viscosity solution approach to shape from shading. SIAM J. Numer. Anal. **29**, 867–884 (1992)
17. Sethian, J.A.: A fast marching level set for monotonically advancing fronts. Proc. Natl. Acad. Sci. **93**, 1591–1595 (1996)
18. Sethian, J.A.: Fast marching methods. SIAM Rev. **41**, 199–235 (1999)
19. Sethian, J.A.: Level Set Methods and Fast Marching Methods. Cambridge University Press, Cambridge (1999)
20. Tsitsiklis, J.N.: Efficient algorithms for globally optimal trajectories. IEEE Trans. Autom. Control **40**, 1528–1538 (1995)
21. Yatziv, L., Bartesaghi, A., Sapiro, G.: $\mathcal{O}(N)$ implementation of the fast marching algorithm. J. Comput. Phys. **212**, 393–399 (2006)
22. Yezzi, A., Kichenassamy, S., Kumar, A., Olver, P., Tannenbaum, A.: A geometric snake model for segmentation of medical imagery. IEEE Trans. Med. Imaging **16**, 199–209 (1997)

Fethallah Benmansour received the M.S. degree in applied mathematics and computer science in 2003 from Ecole Polytechnique, Palaiseau, France, and the M.S degree in applied mathematics and computer vision from Ecole Normale Supérieure de Cachan, Cachan, France in 2004. He is currently pursuing the Ph.D. in applied mathematics at CEREMADE, Université Paris-Dauphine, Paris, France. His research interests are in variational methods and their application to image analysis, for medical and biomedical image segmentation. He is also working on optimization problems, and their applications like traffic congestion and territory protection.

Laurent D. Cohen was at Ecole Normale Superieure Ulm in Paris from 1981 to 1985. He received Master's and Ph.D. degrees in Applied Mathematics from Paris 6 in 1983 and 1986. From 1985 to 1987, he was member at the Computer Graphics and Image Processing group at Schlumberger Palo Alto Research, California and Schlumberger Montrouge Research, and remained consultant there for a few years afterwards. He began working with INRIA, France in 1988, mainly with the medical image understanding group Epidaure. Since 1990, he is Research Scholar (Charge then Directeur de Recherche) with CNRS in the Applied Mathematics and Image Analysis group at CEREMADE, University Paris-Dauphine. His research interests and teaching at the university are applications of variational methods and Partial Differential Equations to Image Processing and Computer Vision, like deformable models, minimal paths, surface reconstruction, Image registration, Image segmentation and restoration. A large part of his work was motivated by medical image applications in collaboration with hospitals or industry. He obtained CS 2002 Prize for Image and Signal Processing. He has been member in program committees for boards for about 30 international conferences.