

Chapter 4

SKELETONS VIA SHOCKS OF BOUNDARY EVOLUTION

Kaleem Siddiqi

School of Computer Science & Centre for Intelligent Machines

McGill University, Canada

siddiqi@cim.mcgill.ca

Sylvain Bouix

Psychiatry Neuroimaging Laboratory

Harvard University Medical School, USA

sylvain@bwh.harvard.edu

Jayant Shah

Department of Mathematics

Northeastern University, USA

shah@neu.edu

Abstract In this chapter we develop the Hamiltonian formulation of the eikonal equation and then relate it to the specific case of Blum's grassfire flow, which gives the level sets of the Euclidean distance function to the boundary. This view provides an explicit association between medial loci and the singularities of this flow. In order to detect these singularities we consider the average outward flux of the gradient of the Euclidean distance function. This measure has very different limiting behaviors depending upon whether the region over which it is computed shrinks to a singular point or a non-singular one. At medial loci the limiting values are related to the object angle. We combine the flux measurement with a homotopy preserving thinning process applied in a discrete lattice. This leads to a robust algorithm for computing skeletons in 2D as well as 3D, which has low computational complexity. We also discuss a related approach to associate medial loci with locations where the gradient of the Euclidean distance function is multi-valued,

which uses the object angle as a measure of salience. Both approaches are illustrated with computational examples.

1. OVERVIEW

The grassfire flow is a special case of the eikonal equation which arises in geometric optics and has an associated Hamiltonian formulation. In this chapter we develop this formulation in Section 2. Whereas this presentation is somewhat detailed, the development provides a comparison with the more standard Lagrangian formulation and suggests an alternative method for simulating the grassfire flow. It also shows that in this general setting medial loci may be viewed as the locus of positions that correspond to shocks of this boundary flow. In order to detect these shocks we consider the average outward flux ideas introduced in Chapter 3, Sections 5 and 6 and provide an implementation based on digital distance functions in Section 2.3. In Section 3 we combine this implementation with a homotopy preserving thinning process and illustrate the overall algorithm with 2D and 3D examples. We then discuss a related approach which uses the object angle as a measure of salience in Section 4, where we provide comparative 2D and 3D results between the two methods.

We begin with a brief overview of skeletonization methods based on properties of the Euclidean distance function, since these are most closely related to the approaches discussed in this chapter. As explained in Chapter 1, the Euclidean distance transform of a binary image is obtained by labeling each point in the domain by its Euclidean distance to the boundary of the object. The locus of skeletal points coincides with the singularities of this function. Thus, approaches in this class attempt to detect local maxima of this function, or the corresponding discontinuities in its derivatives (Arcelli and di Baja, 1992; Leymarie and Levine, 1992; Kimmel et al., 1995; Gomes and Faugeras, 1999). A common approach is to find ridges in the digital distance map of the input image. Unfortunately, significant discretization problems are faced when working on a digital lattice, e.g., thresholding the Euclidean distance function is not sufficient to guarantee thin, homotopy preserving skeletons. Thus, more elaborate procedures have to be devised. For example, Malandain and Fernandez-Vidal (1998) obtain two sets based on thresholding a function of two heuristic measures to characterize the singularities of the Euclidean distance function. The two sets are combined using a topological reconstruction process. Pudney (1998) has also introduced a distance ordered homotopy preserving thinning procedure where points are removed in order of their distance from the boundary

while anchoring end points and centers of maximal balls identified from a chamfer distance function. In Chapter 5 some of the state of the art algorithms for obtaining skeletons from digital distance transforms will be discussed. These algorithms are different from methods based on shocks of the grassfire flow, which are the focus of the present chapter, in that the latter are motivated by properties of the Euclidean distance function in the continuum along with the discretization of appropriate measures to locate its singularities.

As discussed in Chapter 1, Blum's (1973) definition of the medial set in terms of a *grass fire* has also been used as a model for skeletonization. Imagine the inside of the object to be a field of grass, and let the background be a non-flammable material. If the boundary of the object is lit, the wavefront originating from the boundary will meet exactly where the medial set of the object lies, assuming that the field of grass has constant refractive index. Instead of simulating the boundary erosion as a digital thinning, the motion of the boundary \mathbf{W} under fire can be modeled as a continuous wavefront by the following partial differential equation:

$$\frac{\partial \mathbf{W}}{\partial t} = -\mathbf{n}, \quad (1)$$

where $-\mathbf{n}$ is the inward normal to the boundary. The medial axis of an object can be extracted by simulating the flow and detecting the points where the fronts meet (Xia, 1989; Leymarie and Levine, 1992; Kimmel et al., 1995). In these techniques, the contour has to be first segmented at curvature extrema, which is itself a challenging problem. Tek and Kimia (1999) have also proposed an interesting approach for calculating symmetry maps, which is based on the combination of a wavefront propagation technique with an exact (analytic) distance function. In their technique the representation must be pruned in order to distinguish salient branches from unwanted ones. These methods give good results and provide many of the properties the digital medial axis should have; however a generalization to the 3D case proves difficult.

2. OPTICS, MECHANICS AND HAMILTON-JACOBI SKELETONS

We now develop the connection between the grassfire flow and the formation of the skeleton. Let $\mathbf{W}(t)$ be the moving boundary (front) of the object (assumed to be a closed curve in 2D or a closed surface in 3D), $-\mathbf{n}$ the unit inward normal to \mathbf{W} and t a parameter to denote the family of evolved fronts. Let \mathbf{W}_0 be the initial boundary of the object.

The motion of the front is given by

$$\frac{\partial \mathbf{W}}{\partial t} = -f \mathbf{n}, \quad (2)$$

where the function $f = f(\mathbf{x})$ is the speed of the front and is related to the refractive index n of the medium. If n is a function solely of position \mathbf{x} , the medium is said to be *isotropic*. If it depends on both position and orientation, the medium is *anisotropic*. Furthermore, if n does not vary with position, the medium is said to be *homogeneous*. In the case of Blum's grass fire flow, f is the scalar constant $1/n$ and thus the medium is homogeneous and isotropic. Unfortunately, given an arbitrary boundary \mathbf{W}_0 no direct analytical method exists to detect the shocks of this equation.

2.1. MEDIAL LOCI AND THE EIKONAL EQUATION

A numerical approach to simulating flows of the type in Eq. 2 while handling topological changes is to use level set methods developed by Osher and Sethian (1988); Sethian (1996). Below we focus the discussion on the case of a *monotonically* advancing front where $f = f(\mathbf{x})$ has fixed sign for all points \mathbf{x} in the domain of \mathbf{W} . Let $\phi(\mathbf{x})$ be a graph of the solution, obtained by superimposing all the evolved fronts in time. In other words, $\phi(\mathbf{x})$ is the time at which the front crosses a point \mathbf{x} in the medium. We thus have

$$\phi(\mathbf{W}(t)) = t. \quad (3)$$

Taking the total derivative of ϕ with respect to time we get

$$\frac{d}{dt} \phi(\mathbf{W}(t)) = 1 \quad (4)$$

$$\frac{\partial \phi}{\partial \mathbf{x}} \cdot \frac{\partial \mathbf{W}}{\partial t} = 1 \quad (5)$$

$$\nabla \phi \cdot \frac{\partial \mathbf{W}}{\partial t} = 1. \quad (6)$$

Now substitute for $\frac{\partial \mathbf{W}}{\partial t}$ using equation (2),

$$\begin{aligned} \nabla \phi \cdot -f \mathbf{n} &= 1 \\ \|\nabla \phi\|^2 &= \frac{1}{f^2}. \end{aligned} \quad (7)$$

Equation (7) is a form of the well known *eikonal equation*, a key concept of geometrical optics (Luneburg, 1964; Stavroudis, 1972). It is also

at the core of a multitude of other scientific problems. A number of algorithms have been recently developed to numerically solve this equation, including Sethian's fast marching method (Sethian, 1996) which systematically constructs ϕ using only upwind values, Rouy and Tourin's (1992) viscosity solutions approach and Sussman et al.'s (1994) level set method for incompressible two-phase flows. However, none of these methods address the issue of shock detection explicitly, and more work has to be done to track shocks.

In the next section, we shall consider an alternate framework for solving the eikonal equation, which is based on the canonical equations of Hamilton. The technique is widely used in classical mechanics and rests on the use of a Legendre transformation (see Arnold (1989); Shankar (1994)), which takes a system of d second-order differential equations to a mathematically equivalent system of $2d$ first-order differential equations.

2.2. HAMILTONIAN DERIVATION OF THE EIKONAL EQUATION

2.2.1 Variational Principles. We begin by reviewing Lagrangian and Hamiltonian methods for solving minimization problems (Arnold, 1989; Shankar, 1994). Formally, let

$$\psi = \int_{t_0}^{t_1} \mathcal{L}(\mathbf{q}, \dot{\mathbf{q}}, t) dt, \quad (8)$$

be a functional over the space of curves $\{(\mathbf{q}, t) : \mathbf{q}(t) = \mathbf{q}, t_0 \leq t \leq t_1\}$, where $\dot{\mathbf{q}} = \mathbf{q}'(t)$. $\mathcal{L}(\mathbf{q}, \dot{\mathbf{q}}, t)$ is called the *Lagrangian* of the problem. We are interested in finding the curve $\gamma = \{(\mathbf{q}, t) : \mathbf{q}(t) = \mathbf{q}, t_0 \leq t \leq t_1\}$, such that $\psi(\gamma)$ is an extremum. The procedure is well known and leads to the *Euler-Lagrange* theorem.

Theorem 1 (Euler-Lagrange equation) *The curve γ is an extremal of the functional $\psi(\gamma) = \int_{t_0}^{t_1} \mathcal{L}(\mathbf{q}, \dot{\mathbf{q}}, t) dt$ on the space of curves joining (\mathbf{q}_0, t_0) and (\mathbf{q}_1, t_1) , if and only if the following Euler-Lagrange equations*

$$\frac{d}{dt} \frac{\partial \mathcal{L}}{\partial \dot{\mathbf{q}}} - \frac{\partial \mathcal{L}}{\partial \mathbf{q}} = 0 \quad (9)$$

are satisfied along the curve γ .

This is a system of d second order equations whose solutions depend on the $2d$ boundary conditions $\mathbf{q}(t_0) = \mathbf{q}_0$ and $\mathbf{q}(t_1) = \mathbf{q}_1$. This system of d second order equations can be transformed in a system of $2d$ first order

equations which is usually easier to solve. The transformation makes use of the theory of *Hamiltonian Canonical Variables*.

The key to the method is to exchange the roles of $\dot{\mathbf{q}}$ by the canonical variable $\mathbf{p} = \frac{\partial \mathcal{L}}{\partial \dot{\mathbf{q}}}$, commonly referred to as the *momentum*, and replace the Lagrangian $\mathcal{L}(\mathbf{q}, \dot{\mathbf{q}}, t)$ with the function $\mathcal{H}(\mathbf{q}, \mathbf{p}, t)$, called the Hamiltonian, such that the velocities now become the derived quantities

$$\dot{\mathbf{q}} = \frac{\partial \mathcal{H}}{\partial \mathbf{p}}.$$

This can be done by applying the following *Legendre transformation*:

$$\mathcal{H}(\mathbf{q}, \mathbf{p}, t) = \mathbf{p} \cdot \dot{\mathbf{q}} - \mathcal{L}(\mathbf{q}, \dot{\mathbf{q}}, t) \quad (10)$$

where the $\dot{\mathbf{q}}$'s are written as functions of \mathbf{p} 's. It is a simple exercise to verify that the above expression for the velocities $\dot{\mathbf{q}}$ then holds. This transformation is possible if the Lagrangian \mathcal{L} is non-degenerate, i.e., the determinant of its Hessian is non zero. One can also take partial derivatives of the Hamiltonian with respect to the \mathbf{q} 's and verify that

$$\frac{\partial \mathcal{H}}{\partial \mathbf{q}} = -\frac{\partial \mathcal{L}}{\partial \mathbf{q}}.$$

Using Eq. (9), $\frac{\partial \mathcal{L}}{\partial \mathbf{q}}$ can be replaced with $\dot{\mathbf{p}}$ to give *Hamilton's canonical equations*:

$$\dot{\mathbf{p}} = -\frac{\partial \mathcal{H}}{\partial \mathbf{q}}, \quad \dot{\mathbf{q}} = \frac{\partial \mathcal{H}}{\partial \mathbf{p}}. \quad (11)$$

Thus, in the Hamiltonian formalism one starts with the initial positions and momenta $(\mathbf{q}(t_0), \mathbf{p}(t_0))$ and integrates Eq. (11) to obtain the phase space $(\mathbf{q}(t), \mathbf{p}(t))$ of the system.

Using equation (8) and (10) it is straightforward to see that

$$\psi(\gamma) = \int_{t_0}^{t_1} \mathbf{p} \cdot \dot{\mathbf{q}} - \mathcal{H} dt. \quad (12)$$

These results are summarized in the following theorem:

Theorem 2 (Hamilton's equations) *Let $\mathbf{p} = \frac{\partial \mathcal{L}}{\partial \dot{\mathbf{q}}}$ and substitute in (9). The system of d second order Euler-Lagrange equations $\dot{\mathbf{p}} - \frac{\partial \mathcal{L}}{\partial \mathbf{q}} = 0$, is equivalent to the system of $2d$ first order equations (Hamilton's equations)*

$$\dot{\mathbf{p}} = -\frac{\partial \mathcal{H}}{\partial \mathbf{q}}, \quad \dot{\mathbf{q}} = \frac{\partial \mathcal{H}}{\partial \mathbf{p}},$$

where $\mathcal{H}(\mathbf{q}, \mathbf{p}, t) = \mathbf{p} \cdot \dot{\mathbf{q}} - \mathcal{L}(\mathbf{q}, \dot{\mathbf{q}}, t)$ is the Legendre transform of the Lagrangian viewed as a function of $\dot{\mathbf{q}}$.

THE LAGRANGIAN FORMALISM	THE HAMILTONIAN FORMALISM
The state of the system is described by $(\mathbf{q}, \dot{\mathbf{q}})$.	The state of the system is described by (\mathbf{q}, \mathbf{p}) .
The state may be represented by a point moving with a velocity in an d -dimensional configuration space.	The state may be represented by a point in a $2d$ -dimensional phase space.
The d coordinates evolve according to d second-order equations.	The $2n$ coordinates and momenta obey $2d$ first-order equations.
For a given \mathcal{L} several trajectories may pass through a given point in the configuration space.	For a given \mathcal{H} only one trajectory passes through a given point in the phase space.

Table 4.1 . A comparison of the Lagrangian and Hamiltonian formalisms, taken from (Shankar, 1994).

A comparison of the Lagrangian and Hamiltonian formalisms is presented in Table 4.1.

In the case that the extremals emanating from the point (\mathbf{q}_0, t_0) do not intersect elsewhere but instead form a so called “*central field of extremals*”, one can define the *action function* ϕ as the solution functional of our variational problem

$$\begin{aligned}
 \phi(\mathbf{q}, t) &= \min_{\gamma} (\psi(\gamma)) \\
 &= \int_{\gamma_*} \mathcal{L}(\mathbf{q}, \dot{\mathbf{q}}, t) dt \\
 &= \int_{\gamma_*} \mathbf{p} \cdot \dot{\mathbf{q}} - \mathcal{H} dt.
 \end{aligned} \tag{13}$$

where γ_* is an extremal curve. It can be shown that $\mathbf{p} = \frac{\partial \phi}{\partial \dot{\mathbf{q}}}$ and that the action function satisfies the *Hamilton-Jacobi equation* (Arnold, 1989)

$$\frac{\partial \phi}{\partial t} + \mathcal{H} \left(\mathbf{q}, \frac{\partial \phi}{\partial \dot{\mathbf{q}}}, t \right) = 0. \tag{14}$$

We now have all the necessary tools to formalize the connection between geometric optics and the monotonically advancing front(2).

2.2.2 Fermat’s Principle. Like most laws of classical physics, the equations of geometric optics can be derived from a *variational principle* (Stavroutis, 1972). In this context, the variational principle is called *Fermat’s principle* which states that a ray always chooses a tra-

jectory that minimizes the optical path length¹. Consider a (possibly inhomogeneous) isotropic medium with $n(\mathbf{x})$ its refractive index. The following calculations can then be carried out in arbitrary dimensions, although for simplicity we focus on the 2D case. The (2D) trajectory $\gamma(t) = (x(t), y(t))$ connecting two points $\gamma(t_0) = \mathbf{q}_0$ and $\gamma(t_1) = \mathbf{q}_1$ in the medium minimizes the following integral

$$\begin{aligned}\psi(\gamma) &= \int_{\mathbf{q}_0}^{\mathbf{q}_1} n \, ds \\ \psi(\gamma) &= \int_{t_0}^{t_1} n(x, y) \sqrt{\frac{dx^2}{dt} + \frac{dy^2}{dt}} \, dt \\ \psi(\gamma) &= \int_{t_0}^{t_1} L(x, y, \frac{dx}{dt}, \frac{dy}{dt}, t) \, dt,\end{aligned}\tag{15}$$

where ds is the line element along the ray. In order to proceed with the derivation, we must ensure that the Lagrangian L is non degenerate. Unfortunately, the determinant of the Hessian of L is equal to zero due to the fact that the variational problem is independent of the parameter t . Fortunately, the analysis can be done by choosing the projected coordinate y as the new variable of integration.

$$ds = \sqrt{dx^2 + dy^2} = \sqrt{1 + \frac{dx^2}{dy^2}} \, dy\tag{16}$$

Equation (15) then becomes

$$\begin{aligned}\psi(\gamma) &= \int_{y_0}^{y_1} n(x, y) \sqrt{1 + \frac{dx^2}{dy^2}} \, dy \\ &= \int_{y_0}^{y_1} \mathcal{L}(x, x') \, dy\end{aligned}\tag{17}$$

where $x' = \frac{dx}{dy}$. \mathcal{L} is sometimes referred to as the Fermat Lagrangian. We can apply the Euler-Lagrange equations (9) to the Fermat Lagrangian (17)

$$\frac{d}{dy} \frac{nx'}{\sqrt{1+x'^2}} - \frac{\partial n}{\partial x} \sqrt{1+x'^2} = 0\tag{18}$$

¹More precisely the path must be a local extremum and in rare cases may in fact be a maximum (Luneburg, 1964).

This second order partial differential equation is called the *ray equation* and is related to the eikonal equation. In order to develop this connection we turn to the theory of *Hamiltonian Canonical Variables*.

First, we write the Hamiltonian using (10) and (17)

$$\begin{aligned}\mathcal{H}(x, p_x) &= p_x x' - \mathcal{L}(x, x') \\ &= p_x x' - n\sqrt{1 + x'^2},\end{aligned}\tag{19}$$

with

$$p_x = \frac{\partial \mathcal{L}}{\partial x'} = \frac{nx'}{\sqrt{1 + x'^2}}.\tag{20}$$

We may solve for x' in terms of p_x ,

$$x' = \frac{\partial \mathcal{H}}{\partial p_x} = \frac{p_x}{\sqrt{n^2 - p_x^2}},\tag{21}$$

which leads to the Fermat Hamiltonian

$$\mathcal{H}(x, p_x) = -\sqrt{n^2 - p_x^2}.\tag{22}$$

Now, let us assume a central field of extremals, and define $\phi(x, y) = \min_{\gamma}(\psi(\gamma))$. Substituting (22) into (14) yields

$$\frac{\partial \phi}{\partial y} = \sqrt{n^2 - \frac{\partial \phi^2}{\partial x}}.\tag{23}$$

Squaring both sides one obtains

$$\|\nabla \phi\|^2 = n^2,\tag{24}$$

which is the eikonal equation (see Eq. 7) with speed $f(\mathbf{q}) = \frac{1}{n(\mathbf{q})}$. Observe that the speed of the front $f(\mathbf{q})$ is inversely proportional to the refractive index of the medium $n(\mathbf{q})$. We summarize the results in the following theorem.

Theorem 3 (Geometrical Optics in Isotropic Media) *According to Fermat's principle, the trajectory of a ray in an inhomogeneous isotropic medium minimizes the following integral*

$$\psi(\gamma) = \int_{t_0}^{t_1} n(\mathbf{q}) \|\dot{\mathbf{q}}\| dt = \int_{y_0}^{y_1} n(x, y) \sqrt{1 + \frac{dx^2}{dy}} dy.$$

Its corresponding d second order Euler-Lagrange equations lead to the ray equation

$$\frac{d}{dy} \frac{nx'}{\sqrt{1+x'^2}} - \frac{\partial n}{\partial x} \sqrt{1+x'^2} = 0.$$

Transforming the problem into Hamiltonian form we obtain a system of $2d$ first order equations

$$\begin{aligned} \frac{dp_x}{dy} &= -\frac{\partial \mathcal{H}}{\partial x} \\ \frac{dx}{dy} &= \frac{\partial \mathcal{H}}{\partial p_x} \end{aligned}$$

where the Hamiltonian is given by $\mathcal{H}(x, p_x) = -\sqrt{n^2 - p_x^2}$. Assuming a central field of extremals, a graph of the solution surface

$$\phi(x, y) = \min_{\gamma} (\psi(\gamma)) = \int_{\gamma_*} n(x, y) \sqrt{1 + \frac{dx^2}{dy^2}} dy$$

exists and is called the action function. Furthermore

$$\begin{aligned} \frac{\partial \phi}{\partial y} &= -\mathcal{H}(x, \frac{\partial \phi}{\partial x}) \\ &= \sqrt{n^2 - \frac{\partial \phi^2}{\partial x}} \end{aligned}$$

which leads to the eikonal equation

$$\|\nabla \phi\|^2 = n^2.$$

Hence, there is an explicit connection between medial loci, a monotonically advancing front from an object's boundary, Fermat's principle and the eikonal equation. If one sets $n = 1/f = 1$, one gets the grass fire flow (2)

$$\frac{\partial \mathbf{W}}{\partial t} = -\mathbf{n},$$

whose shocks correspond to one of the Blum (1973) definitions of the medial set. Thus the medial set is the locus of positions where two or more front meets, i.e., where $-\mathbf{n}$ is not defined. If we turn back to the

variational formulation we obtain

$$\psi(\gamma) = \int_{t_0}^{t_1} n(\mathbf{q}) \|\dot{\mathbf{q}}\| dt \quad (25)$$

$$= \int_{y_0}^{y_1} \sqrt{1 + \frac{dx^2}{dy}} dy, \quad (26)$$

which is nothing other than the Euclidean distance $\|\mathbf{q}(t_1) - \mathbf{q}(t_0)\|$ between $\mathbf{q}(t_1)$ and $\mathbf{q}(t_0)$, with the extremal curve γ being a straight line. The associated action function is the *Euclidean distance function*

$$\phi(\mathbf{q}) = \min_{\mathbf{x} \in \mathbf{W}_0} \|\mathbf{q} - \mathbf{x}\|, \quad (27)$$

where $\mathbf{W}_0 = \mathbf{W}(t_0)$ is the initial boundary of the object and the inward normals to the evolving front \mathbf{W} are given by

$$-\mathbf{n} = \nabla\phi.$$

These constructs are illustrated for the outline of a panther shape in Fig. 4.1 (top and middle). Technically ϕ as an action function (13) is not defined at medial points, where the assumption of a central field of extrema is broken. Fortunately, ϕ in its form in Eq. (27) is defined and is continuous over \mathbb{R}^n . Thus we can conclude that the medial set corresponds to locations where the Euclidean distance function ϕ is singular and hence $\nabla\phi$ is multi-valued. This result is not novel; for example, Matheron (1988) noticed that ϕ is differentiable and that $\|\nabla\phi\| = 1$ for all points \mathbf{q} not lying on the medial set. However, by interpreting the grass fire flow using Fermat's principle we provide a view that is motivated by considerations in classical mechanics.

The view developed thus far suggests that if the Euclidean distance function ϕ is available, Hamilton's equations can be used to simulate the grassfire flow while associating its singularities with the Blum skeleton. There are implementation choices for obtaining ϕ including the fast marching technique (Sethian, 1996) and the use of digital distance transforms (Borgefors, 1984). The method that we develop in Section 2.3 uses the latter approach and considers the vector field $\nabla\phi$ for all points in the interior of the object. It exploits the fact that the limiting behavior of the average outward flux of this vector field through a shrinking circular neighborhood can be used to distinguish medial points from non-medial ones. A more recent implementation is designed to handle an object whose boundary is given by a mesh (Stolpner and Siddiqi, 2006); it uses accurate techniques for computing point-to-mesh distances developed in computational geometry.

2.3. DIVERGENCE, AVERAGE OUTWARD FLUX AND OBJECT ANGLE

As explained in Chapter 3, Section 4.5, it is possible to extend the standard divergence theorem to the case of a vector field defined on a region Ω in \mathbb{R}^n with smooth boundary \mathcal{B} defined by a skeletal structure (M, \mathbf{S}) . In particular, if $G = \nabla\phi$ denotes the unit vector field generating the grassfire flow for the region Ω with Blum medial axis M , then for a piecewise smooth region $\Gamma \subset \Omega$

$$\int_{\Gamma} \operatorname{div} G \, dV = \int_{\partial\Gamma} G \cdot \mathbf{n}_{\Gamma} \, dS + \int_{\tilde{\Gamma}} dM. \quad (28)$$

Here \mathbf{n}_{Γ} is the unit outward normal to Γ and $\tilde{\Gamma}$ consists of those $x \in \tilde{M}$ with smooth value \mathbf{S} such that the radial line determined by \mathbf{S} has non-empty intersection with Γ .

The first term on the right hand side of the above equation is the outward flux of the grassfire flow across $\partial\Gamma$. It differs from the divergence integral of G over Γ by the “medial volume of $\tilde{\Gamma}$ ”. As explained in Chapter 3, Section 4.6, it is this medial volume term that allows for the discrimination of medial points from non-medial ones (see also Siddiqi et al. (2002); Dimitrov et al. (2003)). The key idea is to consider the limiting behavior of the average outward flux (the outward flux normalized by the volume of $\partial\Gamma$) as Γ shrinks to a point \mathbf{x} . For points not on the medial locus the standard divergence theorem applies and it can be shown that the limiting value of the average outward flux is 0. However, for points on the medial axis the limiting average outward flux is non-vanishing, with the exception of edge points. Furthermore, the value that it attains can be bounded in terms of the quantities $\min(\rho)$ and M_{Γ} discussed in Chapter 3, Section 4.6. The quantity $\min(\rho)$ depends on both M and \mathbf{S} , and the quantity M_{Γ} is obtained from a medial density term that depends on the choice of Γ (see Chapter 3, Table 3.1 for the cases of Γ an n -disk or an n -cube, with $n = 2, 3$).

For the purposes of computation it is reasonable to choose Γ to be a n -disk, in which case a tighter bound can be obtained for the limiting values of the average outward flux. These calculations are presented in (Dimitrov et al., 2003) and (Dimitrov, 2003) for the case $n = 2$; very similar results hold for $n = 3$. The remarkable fact is that for Γ a disk the limiting values actually provide a function of the object angle θ (the angle between \mathbf{S} and M) at all generic points of M . These results are summarized in Table 4.2 for $n = 2$. For the case of the junction points the formula must be interpreted by considering the sum of the object angles for each of the 3 incoming branches at \mathbf{x} . Thus, an average outward

flux computation not only allows medial points to be distinguished from non-medial ones, but also yields the object angle θ (and hence \mathbf{S}) as a by-product for each skeletal point \mathbf{p} . As explained in Chapter 1, Section 2.3, this in turn allows for boundary reconstruction of the bi-tangent points $\mathbf{b}^{\pm 1}$ associated with each skeletal point, with the radius values obtained from the Euclidean distance function. In this sense the average outward flux of $\nabla\phi$ through a shrinking disk may be viewed as a type of flux invariant for detecting medial loci as well as characterizing the geometry of the implied boundary (Dimitrov et al., 2003; Dimitrov, 2003).

POINT TYPE	$\lim_{\epsilon \rightarrow 0} \frac{\mathcal{F}_\epsilon(\mathbf{x})}{2\pi\epsilon}$
Regular Points	$-\frac{2}{\pi} \sin \theta$
End-Points	$-\frac{1}{\pi} (\sin \theta - \theta)$
Junction Points	$-\frac{1}{\pi} \sum_{i=1}^n \sin \theta_i$
Non-Skeletal Points	0

Table 4.2. A summary of results from Dimitrov et al. (2003) relating the limiting values of the average outward flux of $\nabla\phi$ through a shrinking disk to the object angle θ for the case of $n = 2$. Here $\mathcal{F}_\epsilon(\mathbf{x})$ denotes the outward flux at a point \mathbf{x} of $\nabla\phi$ through $\partial\Gamma$, with Γ a disk of radius ϵ .

Algorithm 1 describes a discrete implementation of the average outward flux using a digital distance transform for ϕ . Fig. 4.1 illustrates its computation on the silhouette of a panther shape, where values close to zero are shown in medium grey. All computations are carried out on a rectangular lattice, although the bounding surface is shown in interpolated form. Strictly speaking, the average outward flux is desired only in the limit as the region shrinks to a point. However, the average outward flux over a very small neighborhood (a circle in 2D or a sphere in 3D) provides a sufficient approximation to the limiting values. A threshold on the average outward flux yields a close approximation to the medial set, as used in (Siddiqi et al., 1999). However, in general it is impossible to guarantee that the result obtained by simple thresholding is homo-

Algorithm 1: Average Outward Flux.

Data : Object Ω .

Result : Average Outward Flux Map.

Compute the digital Euclidean distance transform ϕ of the object (Borgefors, 1984);

Compute the gradient vector field $\nabla\phi$;

Compute the average outward flux of $\nabla\phi$:

for (*each point* \mathbf{x}) **do**

$$F(\mathbf{x}) = \frac{1}{n} \sum_{i=1}^n \langle \mathbf{n}_i, \nabla D(\mathbf{x}_i) \rangle;$$

(where \mathbf{x}_i is a n -neighbor of \mathbf{x} ($n = 8$ in 2D, $n = 26$ in 3D) and \mathbf{n}_i is the outward normal at \mathbf{x}_i of the unit sphere centered at \mathbf{x})

topic to the original shape. A high threshold may yield a connected set, but cannot guarantee that it is thin. A low threshold can yield a thin set, but it may be disconnected. The solution, as we shall show in the subsequent section, is to introduce additional constraints to ensure that the resulting medial set is homotopic to the shape. The essential idea is to incorporate a homotopy preserving thinning process, where the removal of points is guided by the average outward flux values. This leads to a robust and efficient algorithm for computing 2D and 3D medial loci.

3. HOMOTOPY PRESERVING MEDIAL LOCI

In this section we combine the average outward flux computation with a digital thinning process, where points are removed without altering the object's topology. In digital topology, a point is said to be *simple* if its removal does not change the topology of the object. In 2D, we shall consider rectangular lattices, where a point is a unit square with 8 neighbors, as shown in Fig. 4.2 (left). Hence, a 2D digital point is simple if its removal does not disconnect the object or create a hole. In 3D, we shall consider cubic lattices, where a point is a unit cube with 6 faces, 12 edges and 8 vertices. Hence, a 3D digital point is simple if its removal does not disconnect the object, create a hole, or create a cavity (Kong and Rosenfeld, 1989). We should note that the 2D version of the algorithm was first developed in (Dimitrov et al., 2000), and we review it here for completeness.

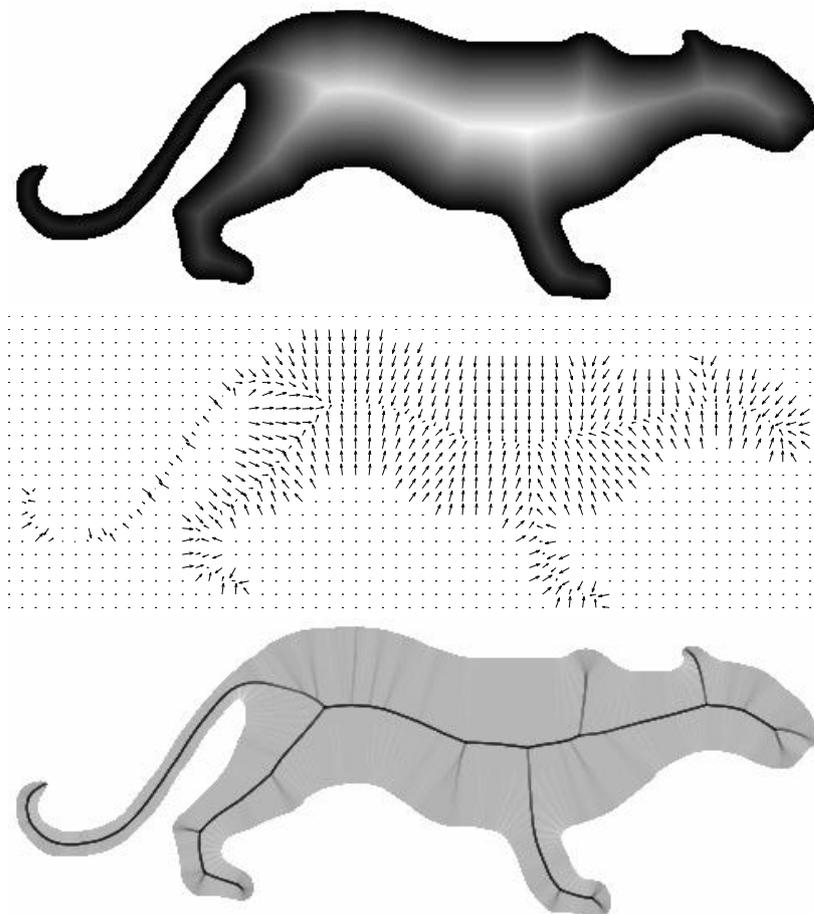


Figure 4.1. The Euclidean distance function ϕ to the boundary of a panther shape (top) with brightness proportional to increasing distance, its gradient vector field $\nabla\phi$ (center) and the associated average outward flux (bottom). Whereas the smooth regime of the vector field gives zero flux (medium grey), strong singularities give large negative values (dark grey) in the interior of the object. Adapted from (Dimitrov et al., 2000).

3.1. 2D SIMPLE POINTS

Consider the 3×3 neighborhood of a 2D digital point \mathbf{x} contained within an object and select those neighbors which are also contained within the object. Now construct a neighborhood graph by placing edges between all pairs of neighbors (not including \mathbf{x}) that are 4-adjacent or 8-

adjacent to one another. If any of the 3-tuples $\{2, 3, 4\}$, $\{4, 5, 6\}$, $\{6, 7, 8\}$, or $\{8, 1, 2\}$, are nodes of the graph, remove the corresponding diagonal edges $\{2, 4\}$, $\{4, 6\}$, $\{6, 8\}$, or $\{8, 2\}$, respectively. This ensures that there are no degenerate cycles in the neighborhood graph (cycles of length 3). Now, observe that if the removal of \mathbf{x} disconnects the object, or introduces a hole, the neighborhood graph will not be connected, or will have a cycle, respectively. Conversely, a connected graph that has no cycles is a tree. Hence, we have a criterion to decide whether or not \mathbf{x} is simple:

Proposition 4.1 *A 2D digital point \mathbf{x} is simple if and only if its 3×3 neighborhood graph, with cycles of length 3 removed, is a tree.*

A straightforward way of determining whether or not a connected graph is a tree is to check that its Euler characteristic $|V| - |E|$ (the number of vertices minus the number of edges) is identical to 1. This check only has to be performed locally, in the 3×3 neighborhood of a point P . Fig. 4.2 (right) shows an example neighborhood graph for which P is simple and hence can be removed.

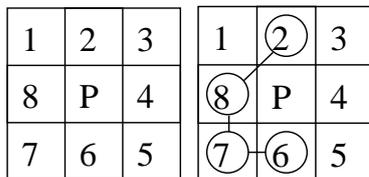


Figure 4.2. LEFT: A 3×3 neighborhood of a candidate point for removal P . RIGHT: An example neighborhood graph for which P is simple. There is no edge between neighbors 6 and 8 (see text).

3.2. 3D SIMPLE POINTS

In 3D a digital point can have three types of neighbors. Two points are *6-neighbors* if they share a face; two points are *18-neighbors* if they share a face or an edge; and two points are *26-neighbors* if they share a face, an edge or a vertex. This induces three *n-connectivities*, where $n \in \{6, 18, 26\}$, as well as three *n-neighborhoods* for \mathbf{x} , $N_n(\mathbf{x})$. An *n-neighborhood* without its central point is defined as $N_n^* = N_n(\mathbf{x}) \setminus \{\mathbf{x}\}$. An object A is *n-adjacent* to an object B , if there exist two points $\mathbf{x} \in A$ and $\mathbf{y} \in B$ such that \mathbf{x} is an *n-neighbor* of \mathbf{y} . A *n-path* from \mathbf{x}_1 to \mathbf{x}_k is a sequence of points $\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_k$, such that for all \mathbf{x}_i , $1 < i \leq k$,

\mathbf{x}_{i-1} is n -adjacent to \mathbf{x}_i . An object represented by a set of points O is *n-connected*, if for every pair of points $(\mathbf{x}_i, \mathbf{x}_j) \in O \times O$, there is a n -path from \mathbf{x}_i to \mathbf{x}_j .

Based on these definitions, Malandain et al. (1993) provide a topological classification of a point \mathbf{x} in a cubic lattice by computing two numbers:

- C^* : the number of 26-connected components 26-adjacent to \mathbf{x} in $O \cap N_{26}^*$
- \bar{C} : the number of 6-connected components 6-adjacent to \mathbf{x} in $\bar{O} \cap N_{18}$.

An important result with respect to our goal of thinning is the following:

Theorem 4 (Malandain et al. (1993)) \mathbf{x} is simple if $C^*(\mathbf{x}) = 1$ and $\bar{C}(\mathbf{x}) = 1$.

We can now determine whether or not the removal of a point will alter the topology of a digital object. When preserving homotopy is the only concern, simple points can be removed sequentially until no more simple points are left. The resulting set will be thin and homotopic to the object. However, without a further criterion the relationship to the medial set will be uncertain since the locus of surviving points depends entirely on the order in which the simple points are removed. In the current context, we have derived a natural criterion for ordering the thinning, based on the average outward flux of the gradient vector field of the Euclidean distance function.

3.3. AVERAGE OUTWARD FLUX ORDERED THINNING

Recall from Section 2.3, that the average outward flux of the gradient vector field of the Euclidean distance function can be used to distinguish non-medial points from medial ones. This quantity tends to zero for the former but approaches a negative number below a constant times $\langle \nabla\phi, \mathbf{N} \rangle$ for the latter, where \mathbf{N} is the one-sided normal to the medial axis or surface. Hence, the average outward flux provides a natural measure of the “strength” of a medial point for numerical computations. The essential idea is to order the thinning such that the weakest points are removed first and to stop the process when all surviving points are not simple, or have a total average outward flux below some chosen (negative) value, or both. This will accurately localize the medial set and also ensure homotopy with the original object. Unfortunately the result is not guaranteed to be a thin set, i.e., one without an interior.

One way of satisfying this last constraint is to define an appropriate notion of an end point. Such a point would correspond to the end point of a curve (in 2D or 3D), or a point on the rim of a surface, in 3D. The thinning process would proceed as before, but the threshold criterion for removal would be applied *only to end points*. Hence, all surviving points which were not end points would not be simple and the result would be a thin set.

In 2D, an end point will be viewed as any point that could be the end of a 4-connected or 8-connected digital curve. It is straightforward to see that such a point may be characterized as follows:

Proposition 4.2 (2D end point) *A 2D point \mathbf{x} could be an end point of a 1 pixel thick digital curve if, in a 3×3 neighborhood, it has a single neighbor, or it has two neighbors, both of which are 4-adjacent to one another.*

In 3D, the characterization of an end point is more difficult. An end point is either the end of a 26-connected curve, or a corner or point on the rim of a 26-connected surface.

Proposition 4.3 (3D end point) *In \mathbb{R}^3 , if there exists a plane that passes through a point \mathbf{x} such that the intersection of the plane with the object includes an open curve which ends at \mathbf{x} , then \mathbf{x} is an end point of a 3D curve, or is on the rim or corner of a 3D surface.*

This criterion can be discretized easily to 26-connected digital objects by examining 9 digital planes in the 26-neighborhood of \mathbf{x} as in (Pudney, 1998).

3.4. THE ALGORITHM AND ITS COMPLEXITY

The essential idea behind the average outward flux-ordered thinning process is to remove simple points sequentially, ordered by their average outward flux, until a threshold is reached. Subsequently, simple points are removed if they are not end points. The procedure converges when all remaining points are either not simple or are end points. The thinning process can be made very efficient by observing that a point which does not have at least one background point as an immediate neighbor cannot be removed, since this would create a hole or a cavity. Therefore, the only potentially removable points are on the border of the object. This suggests the implementation of the thinning process using a *heap* data structure. A full description of the procedure is given in Algorithm 2. The approach is computationally very efficient. With n the total

number of digital points within the original volume and k the number of points within the object, the worst case complexity can be shown to be $\mathcal{O}(n) + \mathcal{O}(k \log(k))$ (Siddiqi et al., 2002).

Algorithm 2: Topology Preserving Thinning.

Data : Object Ω , Average Outward Flux Map.

Result : (2D or 3D) Skeleton.

```

for (each point  $\mathbf{x}$  on the boundary of the object) do
  if ( $\mathbf{x}$  is simple) then
    insert( $\mathbf{x}$ , maxHeap) with AOF( $\mathbf{x}$ ) as the sorting key for insertion;
  while (maxHeap.size > 0) do
     $\mathbf{x}$  = HeapExtractMax(maxHeap);
    if ( $\mathbf{x}$  is simple) then
      if ( $\mathbf{x}$  is an end point) and (AOF( $\mathbf{x}$ ) < Thresh) then
        mark  $\mathbf{x}$  as a medial surface (end) point;
      else
        Remove  $\mathbf{x}$ ;
        for (all neighbors  $\mathbf{y}$  of  $\mathbf{x}$ ) do
          if ( $\mathbf{y}$  is simple) then
            insert( $\mathbf{y}$ , maxHeap) with AOF( $\mathbf{y}$ ) as the sorting key for insertion;

```

\bar{C}	C^*	TYPE
0	any	interior point
any	0	isolated point
1	1	border (simple) point
1	2	curve point
1	> 2	curves junction
2	1	surface point
2	> 2	surface-curve(s) junction
> 2	1	surfaces junction
> 2	≥ 2	surfaces-curves junction

Table 4.3 . The topological classification of Malandain et al. (1993).

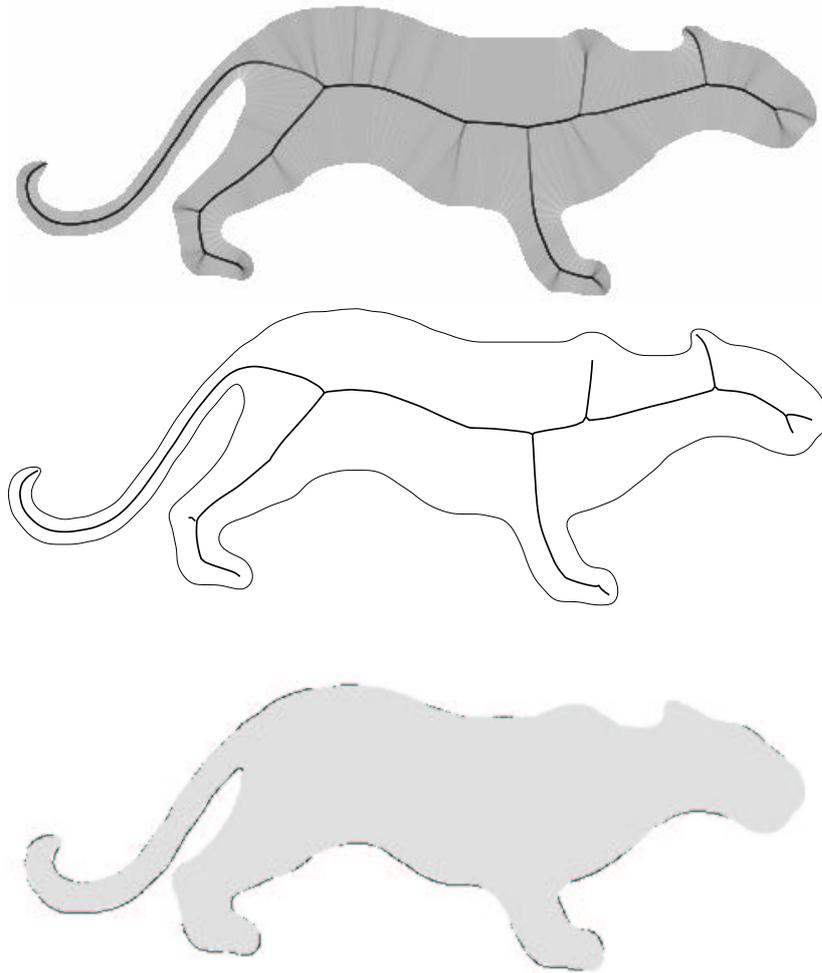


Figure 4.3. The topology preserving medial axis (top) is extracted from the average outward flux map of Fig. 4.1. The reconstruction as the envelope of the maximal inscribed disks of the medial axis is overlaid in grey on the original shape (bottom). Adapted from (Dimitrov et al., 2000).

3.5. LABELING THE MEDIAL SET

The classification of digital points on the 2D medial set is quite straightforward. Consider a circular path through the digital neighbors

of such a point P (see Fig. 4.2) and let n be the number of times this path intersects the medial set. The three generic possibilities are

- $n = 1$, in which case P is an endpoint.
- $n = 2$, in which case P is an interior (curve) point.
- $n = 3$, in which case P is a branch point.

The labeling of the 3D medial set is more subtle and relies on the classification of Malandain et al. (1993). Specifically, the numbers C^* and \bar{C} , described in Section 3.2, can be used to classify curve points, surface points, border points and junction points (see table 4.3).

However, certain junction points can be misclassified as surface points when special configurations of voxels occur. These points are relabeled in a second step using a new definition for simple surfaces described in (Malandain et al., 1993). The procedure is as follows.

Let \mathbf{x} be a surface point ($\bar{C} = 2$ and $C^* = 1$). Let $A_{\mathbf{x}}$ and $B_{\mathbf{x}}$ be the two connected components of $\bar{O} \cap N_{18}$ 6-adjacent to \mathbf{x} . Two surface points \mathbf{x} and \mathbf{y} are in an equivalence relation if there exists a 26-path $\mathbf{x}_0, \mathbf{x}_1, \dots, \mathbf{x}_i, \dots, \mathbf{x}_n$ with $\mathbf{x}_0 = \mathbf{x}$ and $\mathbf{x}_n = \mathbf{y}$ such that for $i \in [0, \dots, n-1]$, $(A_{\mathbf{x}_i} \cap A_{\mathbf{x}_{i+1}} \neq \emptyset$ and $B_{\mathbf{x}_i} \cap B_{\mathbf{x}_{i+1}} \neq \emptyset)$ or $(A_{\mathbf{x}_i} \cap B_{\mathbf{x}_{i+1}} \neq \emptyset$ and $B_{\mathbf{x}_i} \cap C_{\mathbf{x}_{i+1}} \neq \emptyset)$. A *simple surface* is then defined as any equivalence class of this equivalence relation. The 26-neighborhood of each previously classified surface point \mathbf{x} is then examined and when it is not a *simple surface*, \mathbf{x} is relabeled a junction point. Fig. 4.4 illustrates the labeling of the 3D medial set of a cylinder as two *simple* sheets connected by a 3D digital curve through two junction points.

The same definition can be used to extract the individual simple surfaces comprising the medial set of a 3D object. The idea is to find an unmarked surface point on a medial surface and use it as a “source” to build its associated simple surface using a depth first search strategy. The next simple surface is built from the next unmarked surface point and so on, until all surface points are marked.

3.6. EXAMPLES

We now illustrate the use of Algorithms 1 and 2 to produce average outward flux-based medial loci with several 2D and 3D examples.

2D Medial Loci. We assume that the input is a 2D binary image where the foreground and background are identified by distinct values. The implementation then uses an exact (signed) distance function to a piecewise circular arc interpolation of the boundary, which allows for sub-pixel computations (details are presented in Dimitrov et al. (2000,

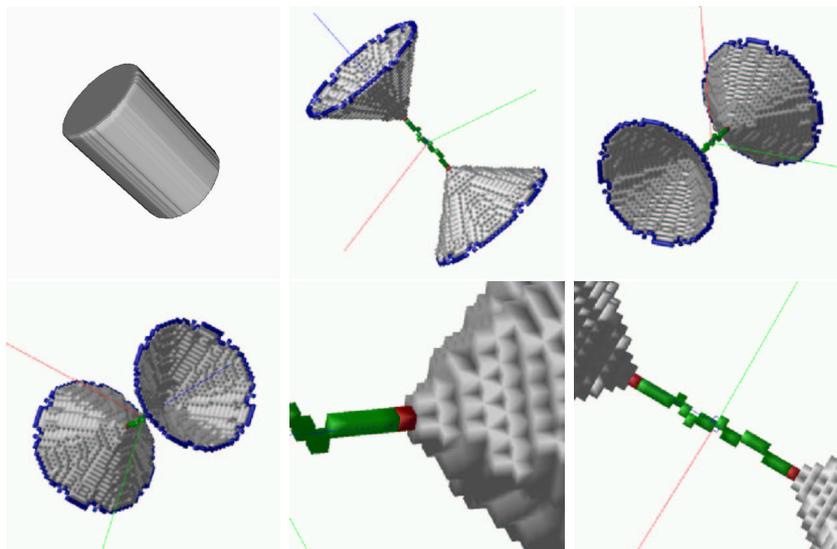


Figure 4.4. The 3D medial set of a cylinder is labeled as border points (blue), surface points (grey), curve points (green) and junction points (red).

2003)). Following this, Algorithms 1 and 2 are used with same average outward flux threshold for each example. Fig. 4.3 (middle) shows the sub-pixel medial axis for the panther silhouette. The accuracy of the representation is illustrated in Fig. 4.3 (bottom), where the shape is reconstructed as the envelope of the maximal inscribed discs associated with each medial axis point. Fig. 4.5 depicts sub-pixel medial axes for a number of other shapes.

3D Medial Loci. Next we illustrate the algorithm with both synthetic data and graphical models of 3D objects. In both cases we assume that the input is a 3D binary array. We then use the D -Euclidean distance function (Borgefors, 1984) which provides a good approximation to the true distance function and apply Algorithms 1 and 2. Once again, the only free parameter is the choice of the average outward flux threshold below which the removal of end points is blocked. For these examples, the value was selected so that approximately 25% of the points within the volume had a lower average outward flux.

Fig. 4.6 depicts the 3D medial loci of a rectangular parallelepiped and a cylinder, as well as the reconstructions of the original objects by superimposing the maximal inscribed spheres at the locus of all 3D

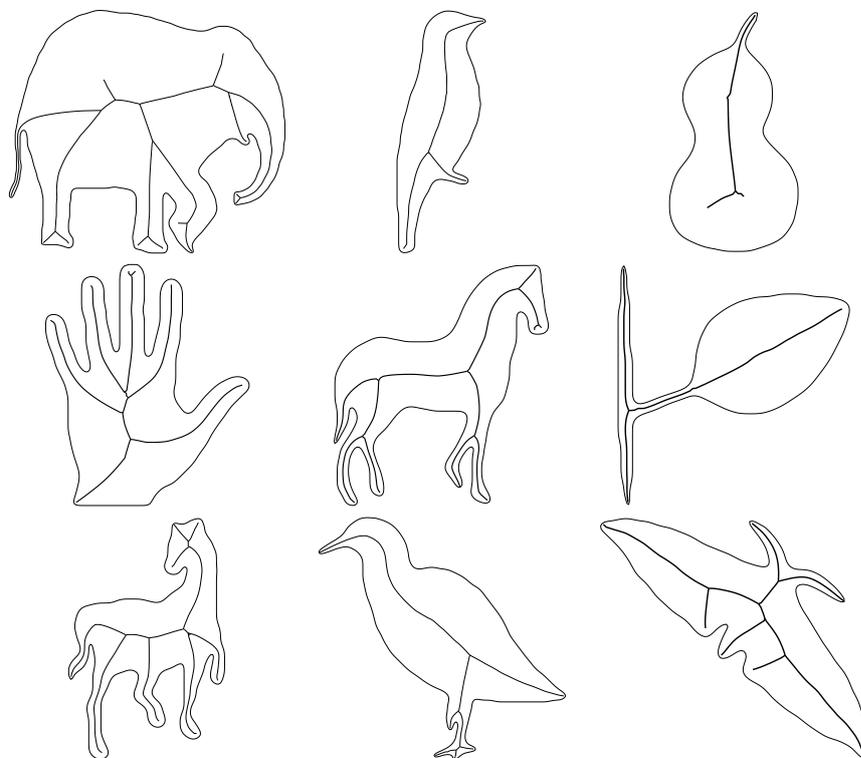


Figure 4.5. Sub-pixel medial sets for a range of 2D shapes, obtained by average outward flux-ordered thinning. We thank Pavel Dimitrov for providing his implementation of the algorithm presented in (Dimitrov et al., 2003).

medial points. As one would expect, the medial set of the parallelepiped is comprised of planes bisecting the adjacent faces and the medial set of the cylinder consists of two cone-like structures which intersect and share a central sheet. Fig. 4.7 depicts the medial loci of various graphical models. These models were originally described in the Virtual Reality Modeling Language (VRML) format. Each model was then rescaled to fit in a $128 \times 128 \times 128$ cubic lattice and was then voxelized using a level set based implementation of a surface extraction method on the cloud of 3D (discrete) surface points (Zhao et al., 2001; Savadjiev et al., 2003). The color map on the skeleton of each binary volume shown in Fig. 4.7 represents the value of the radius function at each medial point, as indicated by the associated color bars. Additional results are presented in Fig. 4.10 (middle column).

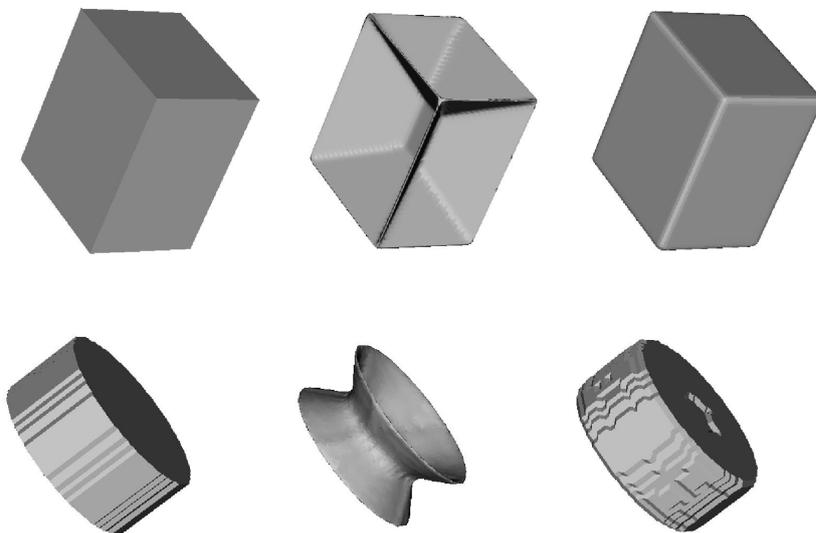


Figure 4.6. FIRST COLUMN: The original 3D objects. SECOND COLUMN: The corresponding average outward flux-based medial loci. THIRD COLUMN: The objects reconstructed from the medial loci in the second column.

4. AN OBJECT ANGLE APPROACH

We now discuss a related approach to detecting medial loci as locations where the gradient of the Euclidean distance function $\nabla\phi$ is multi-valued (Shah, 2005). This approach exploits the relationship between the gradient vector field and the object angle, using the latter as measure of salience by which to prune an initial coarse representation. It is therefore similar in spirit to the average outward flux-based algorithms discussed in Section 3 above, although in its implementation on a discrete lattice it does not incorporate homotopy preserving thinning.

In the following an object is assumed to be a connected open bounded set $\Omega \subset \mathbb{R}^n$ with boundary $\partial\Omega$. Its skeleton is denoted by M and its local geometry is described using the notation in Chapter 1 (Fig. 1.7, Section 2.3). In particular, at a regular medial point \mathbf{p} , the object angle is θ and the associated boundary pre-image points are $\mathbf{b}^{\pm 1}$. The approach exploits the property that in the complement of $\partial\Omega \cup M$, the gradient lines of ϕ are straight lines and $\|\nabla\phi\| = 1$. If \mathbf{x} is a point not on $\partial\Omega \cup M$, the gradient line of ϕ passing through \mathbf{x} connects it to its unique closest point on $\partial\Omega$. However, if \mathbf{x} is a regular point \mathbf{p} of M there are exactly two closest boundary points, \mathbf{b}^+ and \mathbf{b}^- , and the maximal inscribed

sphere centered at \mathbf{p} touches $\partial\Omega$ only at these points. We recall from Chapter 1 the relationship between the object angle and the arc-length derivative of Euclidean distance function (the radius) along the skeleton

$$\cos \theta = \frac{d\phi}{ds}. \quad (29)$$

We also have

$$\sin \theta = \frac{1}{2} \|D^+\phi - D^-\phi\| \quad (30)$$

where $D^\pm\phi$ denotes the directional derivatives in the directions $\overrightarrow{\mathbf{b}^\pm\mathbf{p}}$.

The above formula can be used to extend the notion of $\sin \theta$ to all points in Ω . In particular, for points in the complement of $\partial\Omega \cup M$, because $\nabla\phi = 0$ is continuously differentiable, $D^+\phi = D^-\phi$, so $\sin \theta = 0$.

The result is that the value of $\sin \theta$ can be used to detect the skeleton. This ‘‘object angle’’ extension to all points in Ω/M is referred to as the *grey skeleton* in Shah (2005). It is defined everywhere except on the set J of points of M where the maximal inscribed disk has more than 2 points of contact with $\partial\Omega$. Under the assumption that J has co-dimension ≥ 2 , the strategy is to determine M/J from $\sin \theta$ in the complement of $\partial\Omega \cup J$ and then extend it to portions of J that lie in the closure of M/J . This strategy still leaves out the non-generic cases, such as discs, circular cylinders and balls, where the closure of M/J does not contain all of J . However, since $\sin \theta$ is used only for the purpose of extracting M , it is sufficient to use a simple approximation to define it at points in J . At a point \mathbf{p} in J , pick two points, \mathbf{b}^+ and \mathbf{b}^- , among the set of boundary points nearest to \mathbf{p} , such that the object angle θ determined by the vectors $\overrightarrow{\mathbf{b}^+\mathbf{p}}$ and $\overrightarrow{\mathbf{b}^-\mathbf{p}}$ is the largest possible. $\sin \theta$ is then chosen as the object angle definition at \mathbf{p} . We note that the association between the object angle and the limiting behavior of the average outward flux of $\nabla\phi$ through a shrinking disc developed in Section 2.3 is a more general construction since it applies to all generic cases of medial loci. A possible alternative is to integrate the Laplacian of ϕ in the sense of distributions.

The numerical procedure for estimating θ at each point $\mathbf{x} \in \Omega$ relies on the determination of the gradient directions $\overrightarrow{\mathbf{b}^+\mathbf{x}}$ and $\overrightarrow{\mathbf{b}^-\mathbf{x}}$. These are the directions in which the directional derivative of ϕ is maximum and equals 1. We compute the directional derivatives of ϕ in all *inward* radial directions at every point \mathbf{x} and determine the directions in which it is nearly equal to 1. On a discrete grid there are only finitely many radial directions available depending upon the size of the neighborhood N , and this determines the numerical accuracy of estimating θ . The key point is that since the gradient lines are straight lines and $\|\nabla\phi\| = 1$,

fairly large neighborhoods can be used. If there are exactly two rays from \mathbf{x} along which the directional derivative of ϕ approximately equals 1, then that point must be on the smooth part of the skeleton and the angle between the two rays may be used to calculate the value of $\sin \theta$ at that point. The presence of more than two such gradient directions at a point \mathbf{x} indicates that it is in J and the value of $\sin \theta$ there is determined as indicated above. Finally, if there is a single gradient direction at a point \mathbf{x} , $\sin \theta$ is set to 0 and \mathbf{x} is not in M . This leads to Algorithm 3 for obtaining the object angle extension $\sin \theta$.

Algorithm 3: Object Angle Extension.

Data : Object Ω .

Result : (2D or 3D) Object Angle Extension.

1. Compute the Euclidean distance transform ϕ of the object Ω .

2. **foreach** $\mathbf{x} \in \Omega$ **do**

 Scan the boundary $\partial\Gamma$ of a neighborhood Γ of \mathbf{x} .

foreach $\mathbf{y} \in \partial\Gamma$ **do**

 Calculate the derivative of ϕ in the direction $\overrightarrow{\mathbf{y}\mathbf{x}}$

$$D_{\mathbf{y}}(\phi) = \frac{\phi(\mathbf{x}) - \phi(\mathbf{y})}{|\overrightarrow{\mathbf{y}\mathbf{x}}|}. \quad (31)$$

 Determine the local maxima of $D_{\mathbf{y}}(\phi)$ along $\partial\Gamma$. Select those which are approximately equal to 1, within the tolerance determined by the size of Γ . Label them as $D_{\mathbf{y}_i}$, where $i = 1, \dots, N$ and N is the number of maxima.

if $N = 1$ **then**

 Set $\sin \theta = 0$ at \mathbf{x} .

if $N = 2$ **then**

 Set $\sin \theta = \frac{1}{2} \|D_{\mathbf{y}_1} - D_{\mathbf{y}_2}\|$ at \mathbf{x} .

else

 Set $\sin \theta = \max_{i,j,i \neq j} \frac{1}{2} \|D_{\mathbf{y}_i} - D_{\mathbf{y}_j}\|$ at \mathbf{x} .

The use of the above object angle extension algorithm is illustrated in Fig. 4.8 (left) for the outline of a cat shape, with darker shades corresponding to higher values for $\sin \theta$, for both the object Ω and its comple-

ment. These results were obtained using an adaptive square neighborhood Γ on a discrete lattice. It was chosen to be as large as possible, but no greater than 15×15 , provided that all points in Γ were also in the object Ω (or in its complement in the case of the external medial locus). These results can be compared qualitatively with the average outward flux computation of Fig. 4.1 (bottom), although in the latter example the results are shown only within the object. Fig. 4.8 (right) depicts those medial loci obtained by thresholding and retaining only points \mathbf{x} with $\theta > 75^\circ$. As with the average outward flux measure a procedure has to be devised to process the object angle extension image to obtain medial loci. The procedure suggested in Shah (2005) is to carry out a type of reconstruction based on two object angle thresholds, as follows.

- 1 Choose two thresholds $\bar{\theta}$ and $\underline{\theta}$ for angle θ with $\bar{\theta} > \underline{\theta}$.
- 2 Threshold the object angle image at $\bar{\theta}$ to retain only those loci with large enough object angle.
- 3 Extend the branches of the thresholded skeleton in the direction of increasing ϕ provided that ϕ remains greater than $\underline{\theta}$.

These ideas are quite similar in spirit to the topological reconstruction process developed by Malandain and Fernandez-Vidal (1998), where an initial set of fragmented loci are glued together to produce a skeleton that is topologically correct, in that it preserves homotopy type. Both the above approaches face some numerical challenges due to the difficulty of following directions on a discrete lattice. These approaches are distinct in spirit to the homotopy preserving thinning approach of Algorithm 2, where digital points are removed in a manner inversely proportional to object angle, while taking care to preserve homotopy type. An alternative class of methods for obtaining salient medial loci from an object angle image is based on the construction and minimization of appropriate energy functionals in 2D and 3D (Shah, 2005). This latter approach has also been applied to the segmentation of features such as protrusions and indentations.

4.1. EXAMPLES

We now present both 2D and 3D examples of the object angle approach developed above, and we compare them to those from the average outward flux method. As long as the threshold $\underline{\theta}$ is sufficiently low, the resulting skeleton is a connected set. The procedure is relatively insensitive to the choice of the thresholds $\bar{\theta}$ and $\underline{\theta}$ except around certain critical values.

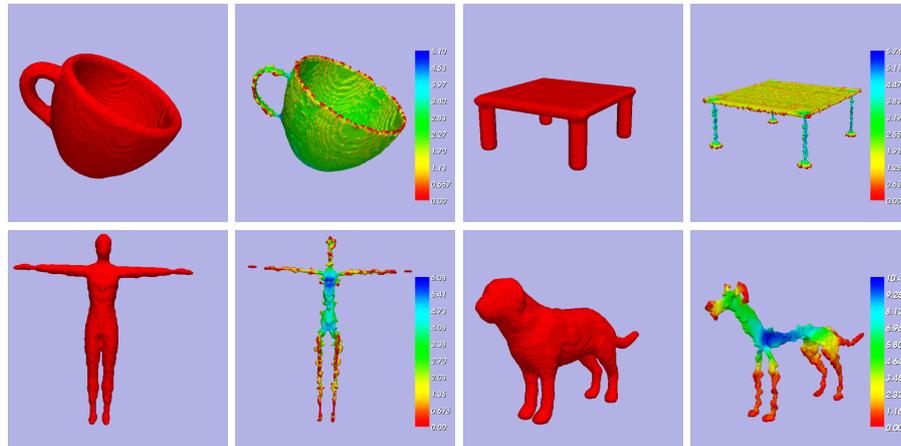


Figure 4.7. Medial sets for a variety of 3D shapes, obtained by average outward flux-ordered thinning. The color map on the skeletons represents the radius function, with values increasing from red to blue. Results on additional models are presented in Fig. 4.10 (middle column).

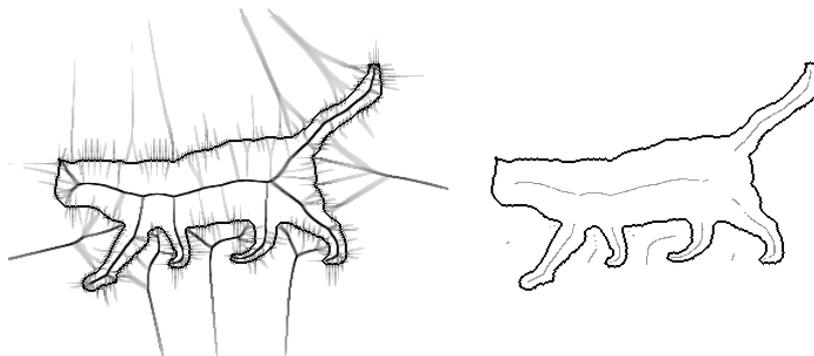


Figure 4.8. LEFT: The object angle extension $\sin \theta$ with darker shades depicting larger θ . RIGHT: The locus of positions with $\theta > 75^\circ$ obtained by thresholding the object angle image.

2D Medial Loci. Fig. 4.9 shows examples of 2D medial loci obtained using the object angle extension approach of Algorithm 3 followed by the reconstruction process from thresholds, for 3 distinct objects. In these examples the reconstruction thresholds were fixed to be $\underline{\theta} = 10^\circ$ and $\bar{\theta} = 75^\circ, 60^\circ$. Qualitatively these results are comparable to those obtained using the average outward flux based method (Fig. 4.5).

3D Medial Loci. The robust numerical implementation of the reconstruction process from thresholds is a significant challenge in 3D, in part due to the difficulty of following the direction of increasing ϕ on a discrete lattice. The object angle extension measure however, as provided by Algorithm 3, can be used in much the same fashion as the average outward flux measure to guide a homotopy preserving thinning process. In the right column of Fig. 4.10 we show examples of 3D medial loci obtained using this idea, with those obtained using average outward flux-ordered thinning shown in the middle column for comparison. The results are in qualitative agreement.

5. DISCUSSION AND CONCLUSION

In this chapter we have shown how considerations from classical mechanics and geometric optics lead to fresh insights into the computation of medial loci. We have considered the Hamiltonian system associated with the vector field underlying Blum's grassfire flow, i.e., the gradient, $\nabla\phi$ of the Euclidean distance function ϕ to the object boundary $\partial\Omega$. We have shown that an average outward flux measure computed over shrinking neighborhoods provides an effective way for detecting medial loci, which is a special case of the medial integrals developed in Chapter 3. For the case of shrinking disks this measure also reveals the object angle, so an explicit mapping from medial points back to the object boundary is provided. We have also discussed a related approach by which medial loci are detected via an extension of the object angle to Ω , along with a procedure to determine locations where $\nabla\phi$ is multi-valued.

The approaches to detecting singularities of the grassfire flow developed in this chapter have the advantage that their development is rooted in principles from classical mechanics. Furthermore, the numerical methods based on average outward flux and the object angle are well supported by an extended divergence theorem that applies to vector fields having a discontinuity along a locus of points and also have the property that the algorithms are essentially the same in 2D and 3D. In their numerical implementation the techniques are further refined to obtain topologically correct medial loci on a discrete lattice.

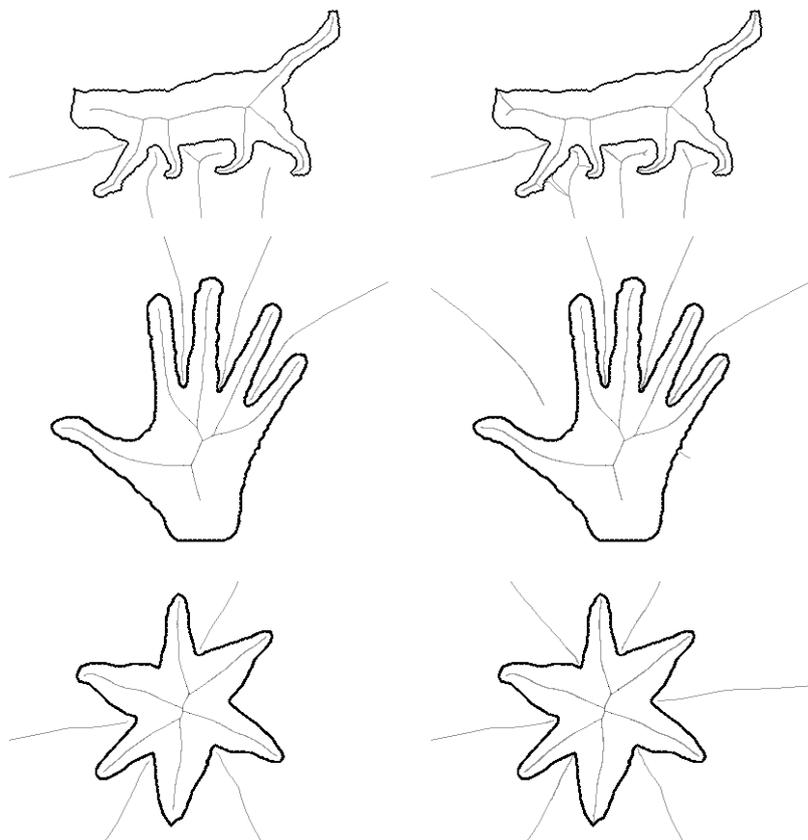


Figure 4.9. External and internal 2D medial loci obtaining using the object angle extension approach, followed by reconstruction from thresholds. LEFT: $\bar{\theta} = 75^\circ, \underline{\theta} = 10^\circ$ RIGHT: $\bar{\theta} = 60^\circ, \underline{\theta} = 10^\circ$.

A potential weakness faced by the approaches in this chapter that numerical errors due to approximations on a discrete lattice do occur. For example, in their current implementation in 3D, average outward flux based skeletons use an approximate digital distance transform, as well as a discrete lattice for thinning. Thus ϕ is imprecise, and the detection of medial loci is only as accurate as the grid sampling. In 2D some of these limitations have been overcome by developing a notion of subpixel medial loci, where the boundary is represented in a continuous fashion and an exact Euclidean distance function is used (Dimitrov et al., 2003; Dimitrov, 2003). In 3D more recent work has extended the average outward flux computation by using computational geometry techniques to work with a polyhedral mesh as $\partial\Omega$ and to use “exact” Euclidean



Figure 4.10. A comparison of medial sets for various 3D shapes (left column), obtained by average outward flux-ordered thinning (middle column), and object angle extension order-ordered thinning (right column). The results are in qualitative agreement.

distance functions to sample points on a shrinking sphere densely and continuously. This in turn leads to a coarse to fine algorithm for computation where the grid is iteratively refined at regions (voxels) through which the medial locus passes (Stolpner and Siddiqi, 2006). Preliminary

evidence suggests that the medial loci thus obtained have the regularity properties to be expected of the Blum skeleton, which are discussed in Chapter 2. Thus, these refined algorithms have the potential to lead to a computational instantiation of the medial to boundary theory developed in chapter 3.

Acknowledgments We are grateful to James Damon, Pavel Dimitrov, Carlos Phillips, Allen Tannenbaum and Steven Zucker for collaborations towards the development of the skeletonization techniques reviewed in this article. We thank Peter Savadjiev and Juan Zhang for their help with the numerical examples. This work was supported by grants from the Natural Sciences and Engineering Research Council of Canada, the Canadian Foundation for Innovation and FQRNT Québec.

Bibliography

- Arcelli, C. and G. S. di Baja: 1992, ‘Ridge Points in Euclidean Distance Maps’. *Pattern Recognition Letters* **13**(4), 237–243. [158]
- Arnold, V. I.: 1989, *Mathematical Methods of Classical Mechanics*. Springer-Verlag. [161, 163]
- Blum, H.: 1973, ‘Biological Shape and Visual Science’. *Journal of Theoretical Biology* **38**, 205–287. [166]
- Borgefors, G.: 1984, ‘Distance Transformations in Arbitrary Dimensions’. *Computer Vision, Graphics, and Image Processing* **27**, 321–345. [167, 170, 178]
- Dimitrov, P.: 2003, ‘Flux Invariants For Shape’. *M.Sc. thesis, School of Computer Science, McGill University*. [168, 169, 186]
- Dimitrov, P., J. N. Damon, and K. Siddiqi: 2003, ‘Flux Invariants for Shape’. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, Vol. 1. Madison, WI, pp. 835–841. [168, 169, 177, 179, 186]
- Dimitrov, P., C. Phillips, and K. Siddiqi: 2000, ‘Robust and Efficient Skeletal Graphs’. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. Hilton Head, South Carolina, pp. 417–423. [170, 171, 176, 177]
- Gomes, J. and O. Faugeras: 1999, ‘Reconciling Distance Functions and Level Sets’. Technical Report TR3666, INRIA. [158]
- Kimmel, R., , D. Shaked, and N. Kiryati: 1995, ‘Skeletonization Via Distance Maps and Level Sets’. *Computer Vision and Image Understanding* **62**(3), 382–391. [158, 159]
- Kong, T. Y. and A. Rosenfeld: 1989, ‘Digital Topology: Introduction and Survey’. *Computer Vision, Graphics, and Image Processing* **48**(3), 357–393. [170]

- Leymarie, F. and M. D. Levine: 1992, 'Simulating the Grassfire Transform Using an Active Contour Model'. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **14**(1), 56–75. [158, 159]
- Lunenburg, R. Y.: 1964, *Mathematical Theory of Optics*. Berkeley and Los Angeles: University of California Press. [160, 164]
- Malandain, G., G. Bertrand, and N. Ayache: 1993, 'Topological Segmentation of Discrete Surfaces'. *International Journal of Computer Vision* **10**(2), 183–197. [173, 175, 177]
- Malandain, G. and S. Fernandez-Vidal: 1998, 'Euclidean Skeletons'. *Image and Vision Computing* **16**, 317–327. [158, 183]
- Matheron, G.: 1988, 'Examples of Topological Properties of Skeletons'. In: J. Serra (ed.): *Image analysis and Mathematical Morphology*, Vol. 2. Academic Press, pp. 217–238. [167]
- Osher, S. and J. Sethian: 1988, 'Fronts Propagating with Curvature Dependent Speed: Algorithms Based on Hamilton-Jacobi Formulation'. *Journal of Computational Physics* **79**, 12–49. [160]
- Pudney, C.: 1998, 'Distance-Ordered Homotopic Thinning: A Skeletonization Algorithm for 3D Digital Images'. *Computer Vision and Image Understanding* **72**(3), 404–413. [158, 174]
- Rouy, E. and A. Tourin: 1992, 'A Viscosity Solutions Approach to Shape-from-Shading'. *SIAM Journal of Numerical Analysis* **29**(3), 867–884. [161]
- Savadjiev, P., F. P. Ferrie, and K. Siddiqi: 2003, 'Surface Recovery from 3D Point Data Using a Combined Parametric and Geometric Flow Approach'. In: *International Workshop on Energy Minimization Methods in Computer Vision and Pattern Recognition*, Vol. LNCS 2683. Lisbon, Portugal, pp. 325–340, Springer-Verlag. [179]
- Sethian, J. A.: 1996, 'A Fast Marching Level Set Method for Monotonically Advancing Fronts'. *Proceedings of the National Academy of Sciences, USA* **93**, 1591–1595. [160, 161, 167]
- Shah, J.: 2005, 'Grayscale Skeletons and Segmentation of Shapes'. *Computer Vision and Image Understanding* **99**(1), 96–109. [180, 181, 183]
- Shankar, R.: 1994, *Principles of Quantum Mechanics*. Plenum Press. [161, 163]

- Siddiqi, K., S. Bouix, A. Tannenbaum, and S. W. Zucker: 1999, ‘The Hamilton-Jacobi Skeleton’. In: *Proceedings of the IEEE International Conference on Computer Vision*. Kerkyra, Greece, pp. 828–834. [169]
- Siddiqi, K., S. Bouix, A. Tannenbaum, and S. W. Zucker: 2002, ‘Hamilton-Jacobi Skeletons’. *International Journal of Computer Vision* **48**(3), 215–231. [168, 175]
- Stavroudis, O. N.: 1972, *The Optics of Rays, Wavefronts and Caustics*. Academic Press. [160, 163]
- Stolpner, S. and K. Siddiqi: 2006, ‘Revealing Significant Medial Structure in Polyhedral Meshes’. In: *Proceedings of the Third International Symposium on 3D Data Processing, Visualization and Transmission*. [167, 187]
- Sussman, M., P. Smereka, and S. Osher: 1994, ‘A Level Set Approach for Computing Solutions to Incompressible Two-Phase Flow’. *Journal of Computational Physics* **114**, 146–154. [161]
- Tek, H. and B. B. Kimia: 1999, ‘Symmetry Maps of Free-Form Curve Segments Via Wave Propagation’. In: *Proceedings of the IEEE International Conference on Computer Vision*. Kerkyra, Greece, pp. 362–369. [159]
- Xia, Y.: 1989, ‘Skeletonization Via the Realization of the Fire Front’s Propagation and Extinction in Digital Binary Shapes’. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **11**(10), 1076–1086. [159]
- Zhao, H. K., S. Osher, and R. Fedkiw: 2001, ‘Fast Surface Reconstruction Using the Level Set Method’. In: *IEEE Workshop on Variational and Level Set Methods*. Vancouver, BC, pp. 104–111. [179]

