

This is a repository copy of *Correcting curvature-density effects in the Hamilton-Jacobi skeleton*.

White Rose Research Online URL for this paper:

<https://eprints.whiterose.ac.uk/1997/>

Article:

Hancock, E.R. orcid.org/0000-0003-4496-2028 and Torsello, A. (2006) Correcting curvature-density effects in the Hamilton-Jacobi skeleton. IEEE Transactions on Image Processing. pp. 877-891. ISSN 1057-7149

<https://doi.org/10.1109/TIP.2005.863951>

Reuse

Items deposited in White Rose Research Online are protected by copyright, with all rights reserved unless indicated otherwise. They may be downloaded and/or printed for private study, or other acts as permitted by national copyright laws. The publisher or other rights holders may allow further reproduction and re-use of the full text version. This is indicated by the licence information on the White Rose Research Online record for the item.

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.

Correcting Curvature-Density Effects in the Hamilton–Jacobi Skeleton

Andrea Torsello and Edwin R. Hancock

Abstract—The Hamilton–Jacobi approach has proven to be a powerful and elegant method for extracting the skeleton of two-dimensional (2-D) shapes. The approach is based on the observation that the normalized flux associated with the inward evolution of the object boundary at nonskeletal points tends to zero as the size of the integration area tends to zero, while the flux is negative at the locations of skeletal points. Nonetheless, the error in calculating the flux on the image lattice is both limited by the pixel resolution and also proportional to the curvature of the boundary evolution front and, hence, unbounded near endpoints. This makes the exact location of endpoints difficult and renders the performance of the skeleton extraction algorithm dependent on a threshold parameter. This problem can be overcome by using interpolation techniques to calculate the flux with subpixel precision. However, here, we develop a method for 2-D skeleton extraction that circumvents the problem by eliminating the curvature contribution to the error. This is done by taking into account variations of density due to boundary curvature. This yields a skeletonization algorithm that gives both better localization and less susceptibility to boundary noise and parameter choice than the Hamilton–Jacobi method.

Index Terms—Curvature, Hamilton–Jacobi equations, shape-description, two-dimensional (2-D) skeleton.

I. INTRODUCTION

THE skeletal abstraction of two-dimensional (2-D) and three-dimensional (3-D) objects has proven to be an alluring yet highly elusive goal for over 30 years in shape analysis. The topic is not only important in image analysis, where it has stimulated a number of important developments including the medial axis transform and iterative morphological thinning operators, but is also an important field of investigation in differential geometry and biometrics where it has led to the study of the so-called morphological skeleton [7].

The morphological skeleton of a shape is defined as the set of singularities in the inward evolution of the boundary with constant velocity. The dynamics of the boundary motion is described by the eikonal equation. This is a partial differential equation that governs the motion of a wave-front through a medium. In the case of a uniform medium, the equation is

$$\frac{\partial}{\partial t} \vec{C}(t) = \alpha \vec{N}(t) \quad (1)$$

where $\vec{C}(t) : [0, s] \rightarrow \mathbb{R}^2$ is the equation of the front at time t , $\vec{N}(t) : [0, s] \rightarrow \mathbb{R}^2$ is the equation of the normal to the wave

front in the direction of motion, and α is the propagation speed. As the wave front evolves, opposing segments collide, generating a singularity.

Broadly speaking, the representation and recognition of 2-D shapes based on the shock representation is a three stage process. First, the skeleton must be computed from the available shape boundary information [1], [2], [8], [9], [12], [18], [19], [28]. The second issue is how to use the extracted skeleton to represent the differential structure of the original boundary. Most of the approaches reported in the literature opt to use a structural characterization. For instance, Siddiqi and Kimia [29] labeled points on the skeleton using so-called shock-labels. According to this taxonomy of local differential structure, there are different classes associated with behavior of the radius of the bitangent circle inscribed in the shape. The so-called shocks distinguish between the cases where the bitangent circle has locally maximum radius, locally minimum radius, constant radius, or a radius which is strictly increasing or decreasing. In [30], Siddiqi *et al.* use this characterization as the basis for a shape matching approach. Kimia and Giblin opt for a simpler representation which is based just on the junctions and terminations of the skeleton [32]. With the skeletal representation to hand, the final step is the matching and indexation of the resulting shape representation [17], [21], [23], [26], [32]. There has also been a consolidated effort at extending these methods to volumetric imagery [6], [25], [34]. Skeletonization methods have found important applications medical image processing including angiography [33] and bronchoscopy [15].

Given the importance of skeletal representations, the quest for reliable and efficient ways of computing skeletal shape descriptors has been a topic of sustained activity. The problem is a complex and elusive one because it is based on the detection of singularities in the evolution of the boundary. The available methods for extracting the skeleton can be divided into three broad categories. The first class of methods are those that involve the use of *marching front techniques* which simulate the grassfire transform. These methods are concerned with iteratively propagating the boundary front over time. Singularities in the simulated evolution of the front indicate the locations of the skeleton. This class of algorithms can be further divided into a) thinning methods [1], [2], [5], where layers of pixels are sequentially peeled from the shape like the skin of an onion, and b) curve evolution methods [16], [31], where curve descriptors such as splines or snakes are transformed according to the eikonal equation. Thinning algorithms have a clear advantage in terms of simplicity. However, their performance is not invariant under Euclidean transformation unless weighted distance functions are used to approximate the Euclidean distance or the Eu-

Manuscript received August 2, 2004; revised March 29, 2005. The associate editor coordinating the review of this manuscript and approving it for publication was Dr. Nicolas Merlet.

A. Torsello is with the Dipartimento di Informatica, Ca' Foscari University of Venice, Venice Mestre 30172, Italy (e-mail: torsello@dsi.unive.it).

E. R. Hancock is with the Department of Computer Science, University of York, York YO10 5DD, U.K. (e-mail: erh@cs.york.ac.uk)

Digital Object Identifier 10.1109/TIP.2005.863951

clidean distance itself are used [3]. Curve evolution methods, on the other hand, are invariant under Euclidean transformation, but require a functional description of the boundary curve. Concrete examples include the use of second or higher order curves and splines [16], or local descriptors such as line segments or circular arc segments [31]. If, on the other hand, the shape is a binary silhouette on the image lattice, then curve evolution requires that a fit be performed to the shape-boundary, and this process both adds to the complexity of the method and it can also be adversely affected by noise. Furthermore, the quality of the extracted boundary curve depends strongly on the reliability of the fitted curve descriptors.

A second class of skeleton extraction algorithms is those that rely on the relationship between the Voronoi triangulation and the skeleton [18], [19]. This work is based on the property that as the number of control points on the object boundary increases, then so the locus of the centers of the triangles of the corresponding Voronoi triangulation of the shape converge to the skeleton. The consequence is that as the triangulation increasingly approximates the shape boundary, then, correspondingly, the centers of the triangles increasingly approximate the skeleton. The important advantages of this approach are that it offers invariance under Euclidean transformation, robustness to noise, that it is fast, and that it is simple to implement. However, its major drawback is the relatively slow convergence speed of the skeleton approximation with respect to the number of control points on the boundary. Hence, this class of algorithm is the natural choice either when the shape is already triangulated (as is often the case with 3-D models) or when it presents a natural triangulation, like a polygonal object. Otherwise, the necessity to tessellate the shape with a large number of triangles negates the advantage of speed.

The third, and final, class of algorithms rely on the analysis of the differential structure of the boundary. Kimmel *et al.* divide the shape-boundary into segments that are delineated by points of maximal curvature [14]. They show that the skeleton is a subset of the Voronoi diagram of these segments. Although this is a powerful approach, it requires direct estimates of the boundary curvature and this in turn relies on curves to be fitted to the boundary. Another important method that falls into this class results from the analysis of the boundary evolution dynamics using the Hamilton–Jacobi equations from classical mechanics [20]. This analysis leads to an eikonal equation which governs the boundary flow. Whenever this flow is nonsingular, the system is Hamiltonian, and, thus, conservative. However, when the system ceases to be conservative there are singularities in the flow of boundary evolution. When the boundary reaches the singularities a so-called shock forms. In the Hamilton–Jacobi setting, skeletal points are detected by searching for locations where the system ceases to be Hamiltonian. The resulting skeleton search method is algorithmically simple and fast. Furthermore, it works directly on the image lattice without the need to extract an intermediate curve description of the boundary. This also makes it relatively robust with respect to boundary noise.

In the first reported account of the Hamilton–Jacobi method, the analysis assumed that boundary evolution ceased to be Hamiltonian at locations where the divergence of the flow was nonzero [9], [27]. Unfortunately, this is not the case. Hence,

this initial work appears to overlook the fact that the linear density of the evolving boundary front is not constant where the front is curved. The result of changes in density is that the flux is not conservative, and, hence, the premise underpinning the skeletonization method does not hold. In a subsequent paper [28], the authors correct this oversight in the analysis by normalizing the flux by the perimeter of the integration area. The resulting normalized flux is still nonzero at nonskeletal locations. However, in the limit as the integration area shrinks to zero, the normalized flux does tend to zero at nonskeletal locations, and is negative on the skeleton itself. Unfortunately, when the integration is performed on the image lattice, the integration area is bounded from below by the pixel size and this introduces an error into the calculation of the normalized flux. Furthermore, there are locations where this error is unbounded.

One way to reduce the effect of this error is to use interpolation techniques to compute the flux with subpixel precision [11]. In this paper, on the other hand, we circumvent this problem by eliminating the error. This is achieved by performing a Hamilton–Jacobi analysis of the boundary evolution under conditions where the flux-density varies due to curvature. Instead of using the gradient of the distance map, i.e., the velocity field of the eikonal equation, we use the momentum field. In other words, we multiply the velocity by the linear density of the boundary-front. The resulting field is conservative and, hence, zero at nonskeletal locations. Moreover, our analysis leads to a new skeleton extraction method. While the method is based on the analysis of a continuous momentum field, its discretization is straightforward since it does not rely on curve fitting to the object boundary and the computation is performed entirely on the image lattice. Hence, discretization requires only sampling on the pixel lattice. The resulting skeleton extraction algorithm is a thinning-based method, where the thinning process is controlled by the value of divergence of the momentum field integrated over the image lattice. We compare the resulting curvature corrected skeletonization method with the Hamilton–Jacobi method. The advantages of the new method are improved localization and stability, and a reduced sensitivity to the model parameters. On the other hand, basing the analysis on a continuous interpretation of the pre-image, that is of the continuous image function that is used to generate the discretely sampled image, allows for subpixel precision algorithms [11].

II. HAMILTON–JACOBI SKELETON

We commence by defining a distance-map that assigns to each point on the interior of an object the closest distance D from the point to the boundary (i.e., the distance to the closest point on the object boundary). The gradient of this distance-map is a field \vec{F} whose domain is the interior of the shape. The field is defined to be

$$\vec{F} = \nabla D \quad (2)$$

where $\nabla = (\partial/\partial x, \partial/\partial y)^T$ is the gradient operator. The trajectory followed by each boundary point under the eikonal equation is governed by the ordinary differential equation $\dot{\vec{x}} = \vec{F}(\vec{x})$,

where \vec{x} is the coordinate vector of the point. Siddiqi *et al.* assume that this dynamic system is Hamiltonian everywhere except on the skeleton [9], [27]. The original interpretation of this property was that at nonskeletal points the normalized flux field \vec{F} is conservative, i.e., $\nabla \cdot \vec{F} = 0$. However, the total inward flux through the boundary of the shape is nonzero. In fact, the flux is proportional to the length of the boundary.

The divergence theorem states that the integral of the divergence of a vector-field over an area is equal to the flux of the field over the enclosing boundary of that area. In our case, this implies that

$$\int_A \nabla \cdot \vec{F}(x) dx = \int_{\partial A} \vec{F} \cdot \vec{n} dl = \Phi_A(\vec{F}) \quad (3)$$

where A is an arbitrary area, \vec{F} is a vector field defined in A , dl is the length differential on the boundary ∂A of A , and $\Phi_A(\vec{F})$ is the outward flux of F through the boundary ∂A of the area A . This implies that, where the divergence is well defined, we have

$$\nabla \cdot \vec{F} = \lim_{|A| \rightarrow 0} \frac{\Phi_A(\vec{F})}{|A|}. \quad (4)$$

We extend the definition to points where the divergence is not well defined, by redefining the divergence at skeletal points using (4). Since the flux through the initial boundary is nonzero, by virtue of the divergence theorem within the interior of the shape, there are points where the system is not conservative. The nonconservative points are those where the boundary trajectory is not well defined, i.e., where there are singularities in the evolution of the boundary. These points are the so-called shocks or skeleton of the shape-boundary. Shocks are, thus, characterized by locations where

$$\lim_{|A| \rightarrow 0} \frac{\Phi_A(\vec{F})}{|A|} < 0$$

or, using the extended definition of the divergence

$$\nabla \cdot \vec{F} < 0.$$

A. Curvature in the Boundary Front

Unfortunately, in general, the flux of \vec{F} is not conservative. To illustrate this point, let us consider an instant in time t during the inward boundary evolution. The initial shape boundary has evolved under the eikonal equation to the front $\vec{C}(t)$ which is at every location orthogonal to \vec{F} . We would like to select a point $p \in \vec{C}(t)$ and compute the value of $\nabla \cdot \vec{F}(p)$ at this point. The value of this divergence is more easily computed in the Frenet frame of the front $\vec{C}(t)$ passing over point p . The Frenet frame of a plane curve $\gamma : [a, b] \leftarrow \mathbb{R}^2$ is the frame provided by the basis $\{v_{\parallel}, v_{\perp}\}$, where

$$v_{\parallel}(\gamma(s)) = \frac{\gamma'(s)}{\|\gamma'(s)\|}$$

$$v_{\perp}(\gamma(s)) = \pm \frac{\frac{\partial}{\partial s} v_{\parallel}(\gamma(s))}{\|\frac{\partial}{\partial s} v_{\parallel}(\gamma(s))\|}.$$

Here, the sign of v_{\perp} depends on the chosen orientation of the curvature. We chose to orient the curvature so that $\partial/\partial s v_{\parallel}(\vec{C}(t, s)) = \kappa(p)\vec{F}(p)$, that is, so that the curvature

$\kappa(p)$ of $\vec{C}(t)$ is positive when the curve bends toward the interior of the shape. Calculating the Frenet frame for the front $\vec{C}(t)$ at the point p , and selecting the inward orientation of the boundary curvature, we have $v_{\perp} = \vec{F}$. Furthermore, we have $v_{\parallel} = dl$, where dl is the arc length differential of $\vec{C}(t)$ at point p . Since the divergence operator is invariant under rotations, the divergence of the field calculated in the Frenet frame is

$$\nabla \cdot \vec{F} = \frac{\partial}{\partial v_{\parallel}} \vec{F} + \frac{\partial}{\partial v_{\perp}} \vec{F}.$$

Since $\|\vec{F}\| = 1$ everywhere, we have $\partial/\partial v_{\perp} \vec{F} = 0$. Moreover, since v_{\parallel} is an arc-length differential for the boundary front $C(t)$ at point p , we have $\partial/\partial v_{\parallel} \vec{F}(p) = -\kappa(p)$, where $\kappa(p)$ is the curvature at p of $\vec{C}(t)$. Hence, we have

$$\nabla \cdot \vec{F}(p) = -\kappa(p). \quad (5)$$

In other words, the divergence $\nabla \cdot \vec{F}$ is not always zero as predicted by the original Hamilton–Jacobi approach [27]. Rather, it is equal to the curvature of the front of the inward evolving boundary.

As a concrete example, consider a circle of unit radius centered in $(0, 0)^T$. The gradient of the distance map at point $(x, y)^T$ is $\nabla D = -1/\sqrt{x^2 + y^2}(x, y)^T$, and the divergence is $\nabla^2 D = -1/\sqrt{x^2 + y^2} \neq 0$.

B. Normalized Flux

This problem was recognized by Siddiqi *et al.* who corrected the analysis in a subsequent publication [28] by introducing the concept of normalized flux. With this modification to the analysis, the non-Hamiltonian points are detected by considering the flux through a circular area A of radius r normalized by the perimeter length $2\pi r$. According to the modification, nonskeletal points satisfy the condition

$$\lim_{r \rightarrow 0} \frac{\Phi_A(F)}{2\pi r} = 0.$$

This condition results from the fact that $\Phi_A(F) = \nabla \cdot \vec{F}(\xi)|A|$, where $\xi \in A$ and $|A| = \pi r^2$ is the area of the circle A . Hence, the limit of the normalized flux becomes

$$\lim_{r \rightarrow 0} \frac{\Phi_A(F)}{2\pi r} = \lim_{r \rightarrow 0} \frac{\nabla \cdot \vec{F}(\xi)}{2} r = 0.$$

Furthermore, in [28], the authors proved that the limit of the normalized flux at skeletal locations is less than a negative constant c , and that this constant depends only on the characteristics of the boundary of the original shape.

While this analysis is correct, it relies on the ability to calculate the limit of the normalized flux through a vanishingly small area. Unfortunately, on the image lattice, there is an obvious lower bound on the size of the integration area due to the pixel resolution. Hence, assuming a minimum integration radius of one pixel, the calculated normalized flux is

$$N\Phi_A(F)(p) = -\frac{\kappa(p)}{2}.$$

At most locations, the absolute value of the calculated normalized flux is much smaller than the constant c . However, near the endpoints of the skeleton, the curvature of the boundary front tends to infinity. Hence, at these locations the exact location of the skeletal points becomes somewhat elusive.




Fig. 1. Evolution of a boundary segment.

III. MOMENTUM FIELD

The fact that the divergence of the field \vec{F} is nonzero can be easily understood by appealing to an analogy from physics. Let us assume that a fluid of uniform density flows from the boundary of the shape (which acts as a source) to the skeleton (which acts as a sink). If the fluid is incompressible, then the fluid density never changes and the flux of the velocity field \vec{F} is conservative everywhere except at points on the skeleton. If, on the other hand, the fluid is compressible, then as soon as a curved front compresses the fluid, the density changes and the velocity field is no longer conservative.

To develop this idea one step further, consider a segment $dl(t)$ of the boundary front $\vec{C}(t)$ at time t . We assume that this segment has average linear density $\hat{\rho}(t)$ (see Fig. 1). Under the eikonal equation, at time $t + \Delta t$ the boundary front segment $dl(t)$ has evolved to $dl(t + \Delta t)$. Since each of the points in $dl(t)$ are now contained in $dl(t + \Delta t)$, the total mass of the two segments is the same. However, if $dl(t)$ is curved then the lengths of the segments are different, i.e., $\|dl(t + \Delta t)\| \neq \|dl(t)\|$. Thus, the average density of $dl(t + \Delta t)$ is $\hat{\rho}(t + \Delta t) \neq \hat{\rho}(t)$. As a result, when the front is curved, then the density is not constant and we have to take into account mass effects. That is, we have to resort to the more general principle of conservation of mass.

Based on this physical intuition, we state that there is indeed a conservative field associated with the dynamics of the boundary, namely the momentum $\vec{M} = \rho\vec{F}$, where ρ is the scalar field that assigns to each point the linear density of the boundary front. As a result, we have

$$\nabla \cdot (\rho\vec{F}) = 0.$$

Applying the rules of product differentiation, we obtain the partial differential equation (PDE)

$$\nabla \rho \cdot \vec{F} = -\rho \nabla \cdot \vec{F}.$$

By setting $\sigma = \log(\rho)$, we can write the above PDE as a function of the log-density σ

$$\rho \nabla \sigma \cdot \vec{F} = -\rho \nabla \cdot \vec{F}.$$

Eliminating ρ from both sides, we obtain

$$\nabla \sigma \cdot \vec{F} = -\nabla \cdot \vec{F}. \quad (6)$$

This is a transport equation that can be reduced to the following set of ordinary differential equations (ODE) along the paths of the boundary points

$$\begin{cases} \frac{d}{dt} \sigma(s(t)) = -\nabla \cdot \vec{F}(s(t)) \\ \frac{d}{dt} s(t) = \vec{F}(s(t)) \end{cases} \quad (7)$$

where $s(t)$ is the trajectory of a boundary point under the eikonal equation.

These equations can be derived by analyzing the change of density of the segment dl in Fig. 1. To commence, we note that $\hat{\rho}(t)\|dl(t)\| = m$, where $dl(t)$ is the length of the boundary segment at time t , $\hat{\rho}(t)$ is its average linear density and $\kappa(t)$ is the curvature at time t . After a small interval of time Δt , the segment length will be

$$\|dl(t + \Delta t)\| = \|dl(t)\| \frac{\kappa(t)}{\kappa(t + \Delta t)} + O(\Delta t^2)$$

and the curvature

$$\kappa(t + \Delta t) = \frac{\kappa(t)}{1 - \kappa(t)\Delta t} + O(\Delta t^2).$$

From these equations, and the conservation of momentum, we have

$$\hat{\rho}(t + \Delta t) = \frac{m}{\|dl(t + \Delta t)\|} = \hat{\rho}(t) \frac{1}{1 - \kappa(t)\Delta t} + O(\Delta t^2).$$

Hence

$$\hat{\rho}(t + \Delta t) - \hat{\rho}(t) = \hat{\rho}(t) \frac{\kappa(t)\Delta t}{1 - \kappa(t)\Delta t} + O(\Delta t^2).$$

Taking the limit for $\Delta t \rightarrow 0$ and $\|dl\| \rightarrow 0$, we have

$$\frac{d}{dt} \rho(s(t)) = \kappa(s(t)) \quad (8)$$

where $s(t)$ is the trajectory to which the limit point of the segment dl tends as $\|dl\| \rightarrow 0$ under the eikonal equation. Integrating (8), we obtain

$$\log(\rho(s(t))) = \int_0^t \kappa(s(\tau)) d\tau.$$

From (5), we have $\kappa(p) = -\nabla \cdot \vec{F}$, yielding

$$\log(\rho(s(t))) = -\int_0^t \nabla \cdot \vec{F}(s(\tau)) d\tau.$$

Hence, integrating ρ , we obtain the vector field $\rho\vec{F}$ which satisfies the condition $\nabla \cdot (\rho\vec{F}) = 0$ at nonskeletal points. Hence, we have $\Phi_A(\rho\vec{F}) = 0$ for any region A not containing a skeletal point. The analysis of the flux of regions containing skeletal points is more complex. The problem we face is that both F and ρ are multivalued on the skeleton, and, hence, $\nabla \cdot (\rho\vec{F})$ is not defined. Although the divergence is not well defined, we can calculate the flux through any area containing a skeletal point.




Fig. 2. Flux through a circular region A containing a skeleton segment is equal to the length of the boundary generating the segment. Left: General case. Middle: Skeletal junction. Right: Skeletal endpoint.

To pursue this analysis, we turn our attention to Fig. 2. Consider a region A containing some skeletal points. Let us assume that the region A contains no endpoints or junctions. Further, without loss of generality, we can assume that A contains a single connected skeleton segment s . In fact, were it not the case, we could partition the region into several disjoint regions A_i , $i = 1, \dots, n$ each containing single connected skeleton segment s . The total flux through A will then be $\Phi_A = \sum_i^n \Phi_{A_i}$. Let e_1 and e_2 be the endpoints of s . There will be two points p_1^1 and p_2^1 on the boundary that converge toward e_1 under the eikonal equation. Similarly, there will be two boundary points p_1^2 and p_2^2 that converge to e_2 under the eikonal equation. Let l_1 be the boundary segment connecting p_1^1 and p_2^1 , and l_2 be the boundary segment connecting p_1^2 and p_2^2 . Further, let $b_1^1, b_2^1, b_1^2, b_2^2$ be the trajectories of $p_1^1, p_2^1, p_1^2, p_2^2$, respectively. Let B be the region contained within the segments $l_1, b_1^1, b_2^1, l_2, b_1^2, b_2^2$. Clearly, the fluxes through A and B are linked by the relation

$$\Phi_A(\rho F) = \Phi_B(\rho F) - \Phi_{B \setminus A}(\rho F) + \Phi_{A \setminus B}(\rho F) \quad (9)$$

where $X \setminus Y$ is the region obtained by subtracting from X the points that are in Y . Since the regions $B \setminus A$ and $A \setminus B$ do not contain any skeletal points, we have $\Phi_{B \setminus A}(\rho F) = \Phi_{A \setminus B}(\rho F) = 0$. Hence, we have

$$\Phi_A(\rho F) = \Phi_B(\rho F). \quad (10)$$

The flux through region B can easily be computed by summing the flux over all its defining boundary segments. Since $b_1^1, b_2^1, b_1^2, b_2^2$ are trajectories of the eikonal equation they are parallel to F at every point, and, hence, there is no flux passing through them, i.e., $\int_{b_i^j} (\rho F) \cdot \vec{n} dl = 0$. On the other hand, F is perpendicular to l_1 , and, hence, the flux through l_1 is $\int_{l_1} (\rho F) \cdot \vec{n} dl = -\|l_1\|$. Similarly, the flux through l_2 is $-\|l_2\|$. Hence, we have

$$\Phi_A(\rho F) = \Phi_B(\rho F) = -\|l_1\| - \|l_2\|. \quad (11)$$

We can now extend the proof to cases where A contains a junction or an endpoint. If the region A contains a junction from which n different branches originate, we can partition it into n

regions A_1, \dots, A_n in the following manner: Let j be the junction, there will be n boundary points p_1, \dots, p_n that converge to j under the eikonal equation, and let b_1, \dots, b_n be their trajectories. These branches split A into the required n regions. Since j is not in the interior of any of the regions, the flux through each of the regions can be computed using the method described above. On the other hand, if the region A contains an endpoint, then region B would consist of a single boundary segment l , which wraps around the endpoint together with two trajectories of the endpoints of l that collide at the single intersection of A with the skeleton. The computation of the flux would, otherwise, be identical yielding the result $\Phi_a(\rho F) = -\|l\|$.

The limit $\lim_{\|A\| \rightarrow 0} \Phi_A(\rho F) / \|A\|$ is not well-defined for an arbitrary sequence of regions of vanishingly small area. However, assuming that we can construct a sequence of circular regions A_r of radius r , containing a skeletal segment l_s , we have

$$\lim_{r \rightarrow 0} \frac{\Phi_{A_r}(\rho \vec{F})}{\|A_r\|} = \lim_{r \rightarrow 0} \frac{-4r \frac{dl}{ds}(\xi) + O(r^2)}{\pi r^2} = -\infty, \text{ if } \frac{dl}{ds} > 0.$$

Here, dl/dl_s is the ratio between boundary length and segment length, and ξ is a point in the skeletal segment l_s .

When integrating the flux numerically on the pixel lattice, there will be a lower bound on the radius. Hence, the corresponding limiting value of the normalized flux will not be $-\infty$, but it will be negative since $dl/dl_s \geq 0$. In particular, the flux will be zero if and only if $dl/dl_s = 0$, that is, on pure ligatures. Ligatures are skeletal branches linked with high negative curvature on the boundary. They are not associated with boundary features, but serve the purpose of linking skeletal features and ensure the skeleton connectivity is preserved. In general, the flux through a 1-pixel circle around a skeletal point will be proportional to the length of the boundary segment generating the skeletal point. This measure has been used previously to distinguish between relevant skeletal branches and irrelevant branches caused by boundary noise [24].

It is worth noting that if we follow Siddiqi *et al.* [28] and normalize by the perimeter $\|\partial A_r\|$ of the region A_r , we obtain

$$\lim_{r \rightarrow 0} \frac{\Phi_{A_r}(\rho \vec{F})}{\|\partial A_r\|} = \lim_{r \rightarrow 0} \frac{-4r \frac{dl}{ds}(\xi) + O(r^2)}{2\pi r} = -\frac{2}{\pi} \frac{dl}{ds}.$$

Hence, factoring out the curvature effects, we obtain a clear geometrical interpretation of the value of the normalized flux at a skeletal point. It is simply proportional to the ratio between the boundary length and the skeletal length.

IV. BOUNDARY CURVE PARAMETERIZATION

There is another interpretation for the scalar field ρ derived from the analysis of the evolution of an arc-length parameterization of the boundary curve of the shape [13]. Let $\vec{C}(t)$ be a solution to the eikonal (1), where $\vec{C}(t) : [0, s] \rightarrow \mathbb{R}^2$ is the equation of the front at time t . Furthermore, let the differential ds of the parameterization s of the curve be, at time $t = 0$, an arc length differential, i.e., let $\|d/ds\vec{C}(0, s)\| = 1$. Clearly, ds will not remain a differential of arc length throughout the evolution of the curve. However, we can define a metric $g(t, s) = ds/dl$ that links the length at time t of the differential ds of the parameterization s , to the arc length differential dl . With this notation, we have

$$\rho(\vec{C}(t, s)) = \frac{1}{g(t, s)}.$$

Further, suppose that l is a segment on the initial boundary, and let l' be the corresponding segment on the front at time t . Since we assume unit density at time $t = 0$, the total mass of the segment l is equal to its length $|l|$. Since mass is conserved, the total mass of the segment l' will remain $|l|$. Hence, the average density is $\rho = |l|/|l'|$. Taking the limit as $|l'| \rightarrow 0$, we have $\rho = dl/ds$. With this definition of the front density, we can rewrite the momentum field as $\vec{M} = \vec{F}/g$. Following [13], the divergence of the momentum and velocity fields are given by

$$\begin{aligned} \nabla \cdot \left(\frac{\vec{F}}{g} \right) &= \frac{\partial \vec{F}}{\partial l} \frac{1}{g} + \frac{\partial \vec{F}}{\partial t} \frac{1}{g} = \frac{\partial \vec{F}}{g \partial s} \frac{1}{g} + \frac{\partial \vec{F}}{\partial t} \frac{1}{g} \\ &= \frac{1}{g^2} (-kg) + \frac{-\frac{\partial g}{\partial t}}{g^2} - \frac{k}{g} + \frac{kg}{g^2} = 0 \quad \text{and} \\ \nabla \cdot \vec{F} &= \frac{\partial \vec{F}}{\partial l} + \frac{\partial \vec{F}}{\partial t} = \frac{\partial \vec{F}}{g \partial s} + \frac{\partial \vec{F}}{\partial t} = \frac{1}{g} (-kg) + 0 = -k \end{aligned}$$

which is simply the result obtained via our analysis of the momentum field.

V. COMPUTING THE DENSITY

To obtain the momentum field, we need to integrate the density field over the interior of the shape. Since images have a finite resolution, we need to discretize the solution onto the image lattice.

One approach is to express the PDE in (6) as a system of difference equations. The difference equations form a linear system that can then be solved to obtain the log-density $\sigma = \log(\rho)$. The problem with this approach is that the skeleton is a set of singularities of the momentum field. Hence, the density can have very different values at opposite sides of a skeletal branch. The net effect is that the linear system will have no solution. In fact, even seeking an approximate solution using a residual descent method would result in oscillations near the skeleton.

A. Integration in Time

In order to overcome this problem, we need to ensure that the difference operators used in the equations never cross a skeletal branch. One way to guarantee this is to integrate the equation in the time domain. This must be done so that the formulae giving the value of ρ at points on the boundary front at time t reference values of ρ only at points in the fronts at previous times. We can realize this by integrating the ODE in (7) along the paths of the boundary points.

To do this, we opt to use the second-order Crank–Nicolson method [10]. For each point $(x, y) = \vec{C}(t, s)$ in the interior of the shape, we have the equation

$$\begin{aligned} \sigma(\vec{C}(t, s)) - \sigma(\vec{C}(t-1, s)) &= \\ &= -\frac{1}{2} [\nabla \cdot \vec{F}(\vec{C}(t, s)) + \nabla \cdot \vec{F}(\vec{C}(t-1, s))]. \end{aligned}$$

Solving for the log-density at time t , we obtain

$$\begin{aligned} \sigma(\vec{C}(t, s)) &= \sigma(\vec{C}(t-1, s)) \\ &= -\frac{1}{2} [\nabla \cdot \vec{F}(\vec{C}(t, s)) + \nabla \cdot \vec{F}(\vec{C}(t-1, s))]. \end{aligned} \quad (12)$$

Using this equation, we can calculate the log-density at a point on the evolving boundary at time t , referencing only values of the log-density at points that belong to the front at previous times. Since the evolution never crosses the skeleton, we are guaranteed not to cross skeletal branches during our calculations.

B. Integration in Space

Equation (12) allows us to integrate the log-density σ in the time domain along the evolution path followed by a boundary point. However, we have not shown how to calculate the integration path. Fortunately, we do not need to calculate every possible path. Let us assume that at time t the boundary front passes through the point $\vec{C}(t, s) = (x, y)^T$. The first-order approximation of the position of this point at time $t-1$ is

$$\vec{C}(t-1, s) = (x, y)^T - \vec{F}(x, y) = (x - F_x, y - F_y)^T.$$

Using this approximation, we can write (12) in the spatial domain instead of the time domain. As a result, the density is given by

$$\begin{aligned} \sigma(x, y) &= \sigma(x - F_x(x, y), y - F_y(x, y)) \\ &= -\frac{1}{2} [\nabla \cdot \vec{F}(x, y) + \nabla \cdot \vec{F}(x - F_x(x, y), y - F_y(x, y))]. \end{aligned} \quad (13)$$

Here, we have used a first-order approximation to compute the location of the point at time $t-1$. The reason for doing this is that it allows us to obtain a closed-form solution. The development of a second-order approximation, on the other hand, would require the solution of a system of two second-order equations. The solution does not exist in closed form and must be located numerically. This would have significantly compromised the speed of the integration step.

As shown in Fig. 3, the point $(x, y)^T - \vec{F}(x, y)$ does not belong to the image lattice. Hence, we need to interpolate it using the values at the four corners of the square containing the point. Note that the point $(x, y)^T$ is the last of the four points in the lattice that is visited by the evolving boundary. Hence, the interpolation is guaranteed to use points on the same side of a




Fig. 3. Integration along the boundary path.

skeleton. We opt to compute the quantity $f(x + a, y + b)$ with $a, b \in [0, 1)$ using the bilinear interpolation

$$(a - 1)(b - 1)f(x, y) + a(b - 1)f(x + 1, y) \\ + (a - 1)bf(x, y + 1) + abf(x + 1, y + 1).$$

With this interpolation, (13) becomes

$$[1 - (1 - |F_x|)(1 - |F_y|)]\sigma(x, y) \\ = |F_x|(1 - |F_y|)\sigma(x', y) \\ + (1 - |F_x|)|F_y|\sigma(x, y') + |F_x||F_y|\sigma(x', y') - \\ - \frac{1}{2}[\nabla \cdot \vec{F}(x, y) + \nabla \cdot \vec{F}(x - F_x, y - F_y)] \quad (14)$$

where $x' = x + \text{sgn}(F_x(x, y))$ and $y' = y + \text{sgn}(F_y(x, y))$.

Using (14), we can compute the value of the log-density $\sigma(x, y)$ using values of σ at the points spanned by the evolving boundary front before the point $(x, y)^T$. Hence, to calculate σ , all we need do is to iterate (14) through the interior points according to front arrival time. We commence from the points reached first by the boundary front and proceed to those reached last. Since the evolving boundary front is moving with constant unit velocity, the time taken by the front to reach the point with position $(x, y)^T$ is equal to its distance from the initial shape boundary.

Once we have the density to hand, we need to calculate the divergence of the momentum in every point on the image lattice. We opt to discretize (6) using the second-order approximation

$$\nabla \cdot (\rho \vec{F})(x, y) = [\sigma(x, y) \\ - \sigma(x - F_x, y - F_y)]e^{\sigma(x, y) - \frac{1}{2}\Delta\sigma} \\ + \frac{1}{2} \left[\nabla \cdot \vec{F}(x - F_x, y - F_y)e^{\sigma(x - F_x, y - F_y)} \right. \\ \left. + \nabla \cdot \vec{F}(x, y)e^{\sigma(x, y)} \right].$$

This corresponds to the second-order Crank–Nicolson method applied to the integration of the log-density σ .

VI. SKELETONIZATION

Once the divergence of the momentum field is to hand, we can extract the skeleton. The extraction process we adopt is similar to the one adopted by Siddiqi *et al.* [28]. The extraction is performed by thinning the shape by removing boundary points that

have energy absorption below a certain threshold, and whose removal would not cause the shape to be split into two disjoint parts. This is an important element of the thinning algorithm. In fact, while the method guarantees that the energy absorption is not negative at skeletal points, it can become arbitrarily small. In particular, as seen earlier, it becomes zero at pure ligature points. Without enforcing branch connectivity, the method would split the skeleton whenever the absorption falls below the threshold. On the other hand, with this check on the thinning process, we are guaranteed to obtain a connected skeleton. In particular, since the thinning is performed in distance order, the skeleton branch would, as expected, follow the gradient of the distance map. The remaining shape is further thinned to a 1-pixel wide skeleton, being careful to maintain the connectivity of the shape and to avoid shortening of the skeleton by eliminating endpoints. Expressed in terms of pseudocode the thinning process of the shape S is as follows:

```

For each point  $p$  in distance order
    if is_simple(S \setminus p) and  $-\Delta \cdot \rho \vec{F}(p) < \epsilon$ 
        then  $S = S \setminus p$ 
For each remaining point  $p$  in distance order
    if is_simple(S \setminus p) and not is_endpoint(S, p)
        then  $S = S \setminus p$ .
    
```

The predicate `is_simple` determines whether the shape is still connected after the removal of the point p . It does so by checking only the points in the neighborhood of p . The shape $S \setminus p$ is connected if the points in the neighborhood of p , excluding p , are connected. Similarly, `is_endpoint` determines whether p is an endpoint. It does so only by inspecting the neighborhood of p . The point is an endpoint if it has at most two neighboring points, and those points are 4-adjacent to each other.

It is worth noting that the only external input that this thinning algorithm requires is the detection of endpoints. In fact, if an algorithm were to return only the set of endpoints, the thinning process would reconstruct the same skeleton. On the other hand, the algorithm is highly dependent on the quality of the detection of the endpoints. Hence, an improvement on the detection and localization of the endpoints would result in an improvement on the extraction algorithm. Since the energy absorption is proportional to the length of the boundary-segment generating the skeletal point, the thinning algorithm is controlled by the branch relevance measure. By controlling the thinning process using the energy absorption as a skeleton-localization measure, we ensure that only endpoints associated with a boundary-segment longer than a threshold value are retained. As stated earlier, the length of the generating boundary is a well known branch-relevance measure often used in skeleton-pruning algorithms to eliminate branches produced by boundary noise. With our method we dispense of the need for a subsequent pruning process, by incorporating it into the skeleton-extraction step.

It is important to note that the addition of the momentum-integration step to the original Hamilton–Jacobi algorithm does not impact on the computational complexity of the method. In fact, the complexity of the step is linear in the number of interior points. This is equivalent to the calculation of the divergence and the thinning step which are used in the Hamilton–Jacobi

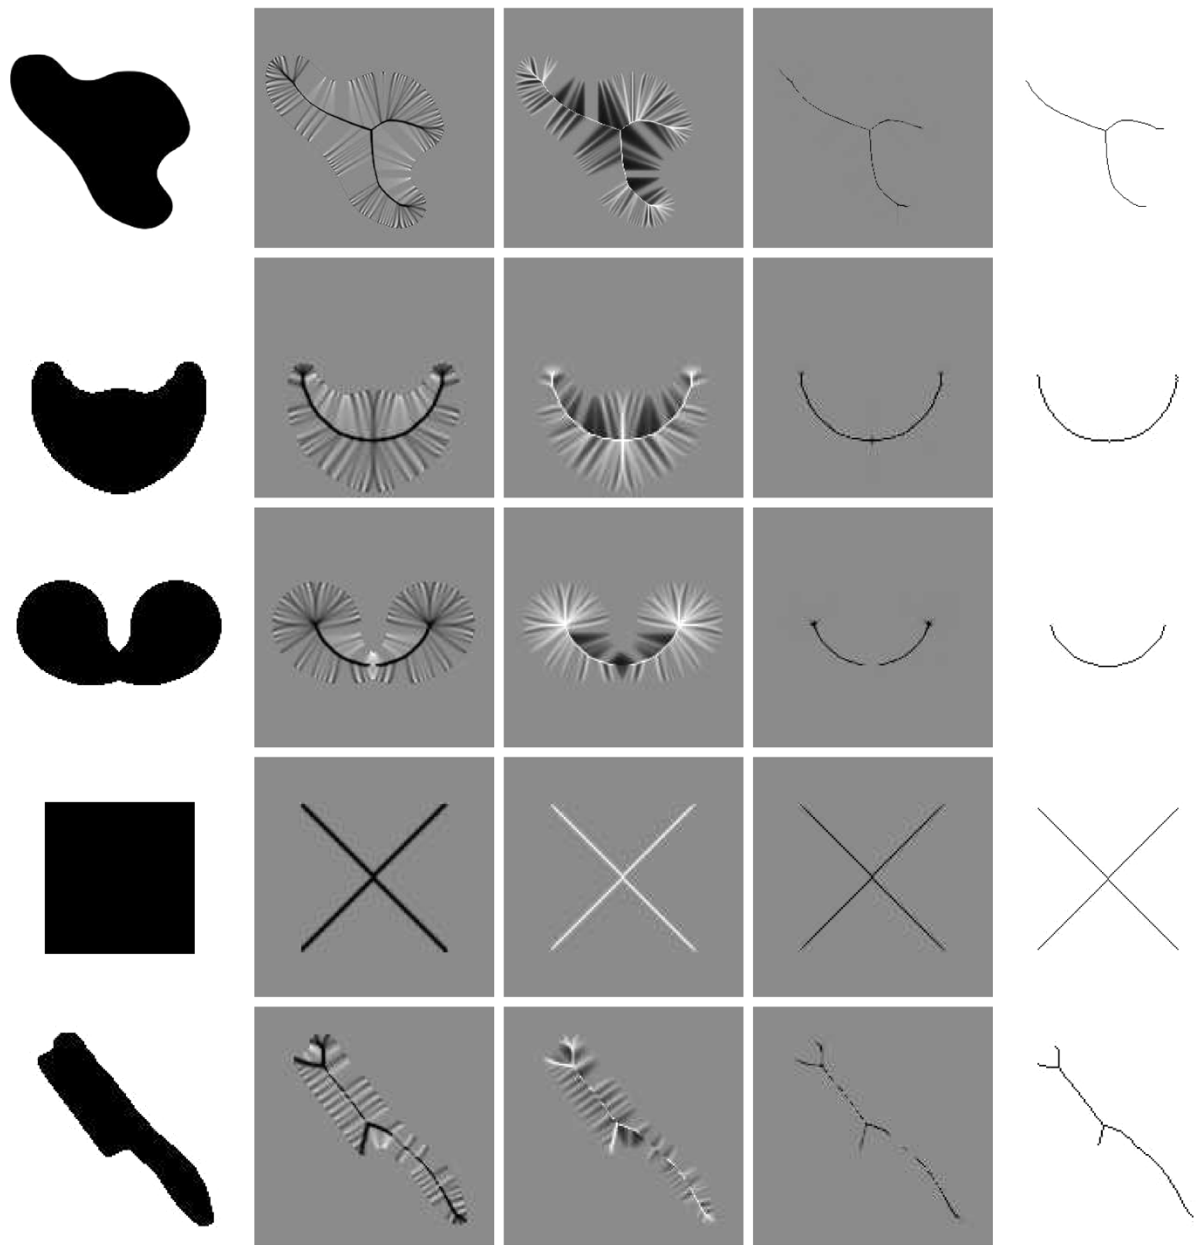


Fig. 4. Differences in the velocity and momentum fields. Left to right: Shape, (normalized) flux of \vec{F} , $\log(\rho)$, flux of $\rho\vec{F}$, and the skeleton extracted with the proposed approach.

approach too. These linear steps are dominated by the initial sort of the interior points which is $O(n \log n)$, where n is the number of interior points. Furthermore, in our implementation, the calculation of the distance map is done in a very naive way, giving it complexity $O(nm)$, where m is the number of boundary points.

VII. EXPERIMENTAL COMPARISON

In this section, we attempt to characterize the differences between the Hamilton–Jacobi skeletonization method and our density-corrected approach. We commence by providing a qualitative analysis of the difference in the divergence of the velocity and momentum fields. Second, we provide an analysis of the noise and thresholding sensitivity of the two methods. Finally, we provide a more quantitative analysis of the localization properties of the two skeletonization methods.

Figs. 4 and 5 show, for a few selected shapes from our database, the values of the flux through a unit circle of the velocity field $\Phi_1(\vec{F})$, the computed log-density $\log(\rho)$, and the flux through a unit circle of $\Phi_1(\rho\vec{F})$. Note that, having fixed the radius for the calculation of the flux, the flux $\Phi_1(\vec{F})$ and the normalized flux $N\Phi_1(\vec{F})$ differ only by a multiplicative constant. In these diagrams white (grayscale 255) corresponds to a large positive value, black (grayscale value 0) to a large negative value and zero is represented by the grayscale value 128. To better show the differences, the contrast of the images is strongly enhanced. This is done by applying to each point intensity a sigmoidal function with slope on 0 equal to 10.

It is clear from the diagrams that the divergence of the velocity field at nonskeletal points is not zero where the boundary evolution front is curved. That is, in correspondence with a curved




Fig. 5. Differences in the velocity and momentum fields. Left to right: Shape, (normalized) flux of \vec{F} , $\log(\rho)$, flux of $\rho\vec{F}$, and the skeleton extracted with the proposed approach.

boundary. The value of the flux through an area that does not contain any section of the skeleton is, in general, an order of magnitude smaller than the value calculated over an area that contains a skeletal branch. However, near the endpoints of the skeletal branches the values become comparable. This can be observed as a blurred dark region around the endpoints. Furthermore, quantization in the localization of the shape causes the initial boundary to be very jagged. This high-frequency, low-amplitude noise is transported and amplified throughout the velocity field creating stripes with high local curvature in the evolving front. This, in turn, yields a noisy and poorly localized skeleton. By contrast, the density correction in the momentum field dampens this noise. As for computation time, the extraction of the skeleton using our method took, on average, 2.18 s, versus 1.72 s using the Hamilton–Jacobi method. Here, the di-




Fig. 6. Discretization error on boundary localization.

mension of images was 256×256 pixels and the computations were performed on a PC with a 1 GHz CPU.




Fig. 7. Effect of smoothing on skeleton extraction. Left to right: The divergence of the velocity field, the uncorrected Hamilton–Jacobi skeleton, the divergence of the momentum field, and the skeleton extracted using the density-corrected method.

A. Noise Sensitivity

Our skeletonization method depends on our ability to calculate the distance map D and its gradient. This, in turn, depends on the correct localization of boundary points. Unfortunately, due to the truncation effects caused by the finite precision of the image lattice, the extracted boundary presents discretization errors. This problem is most pronounced for jagged boundary edges. Fig. 6 illustrates the problem. The dashed line represents the original boundary of the shape and the gray squares represent the boundary pixels in the image lattice. Due to this discretization, the observed boundary is equal to the solid line. Clearly, commencing from this observed boundary, the distance map will diverge considerably from its correct value. The effects of this discretization error will be even more dramatic on the gradient $\vec{\nabla}D$. To overcome quantization noise from the object boundary, we need to smooth the observed shape boundary and select an appropriate skeletonization threshold. To smooth the shape boundary, we approximate shape diffusion [13]. In order to approximate the skeletonization of the diffused shape, it is not necessary to explicitly calculate the diffused shape and distance map. Instead, it is sufficient to approximate the gradient ∇D of the distance map. Let D be the distance map of the original shape and D_τ be the distance map of the shape after a diffusion with parameter τ . Furthermore, let $G_\tau * f$ be a Gaussian smoothing of a function f with standard deviation τ . For small τ , we have

$$\nabla D_\tau \approx \frac{\nabla G_\tau * D}{\|\nabla G_\tau * D\|}.$$

Hence, we can approximate the diffusion of the image by smoothing the distance map by convolving it with a Gaussian, and normalizing the resulting gradient. This approach allows us to approximate the gradient of the distance map of an ideal diffused boundary calculated with subpixel precision, without actually calculating the diffused boundary with subpixel precision.

If either the smoothing radius or the threshold is too large, then some of the branches of the skeleton will be thinned away. If, on the other hand, the selected values are too small, then the detected skeleton will have large numbers of spurious branches (see Fig. 7). In this section, we characterize the effects of

the smoothing radius and the skeletonization threshold on the quality of the detected skeleton.

Fig. 7 displays the effects of (top) very low and (bottom) very high values of the smoothing radius and the skeletonization threshold on a test shape. The diagram shows, left to right, the divergence of the velocity field, the uncorrected Hamilton–Jacobi skeleton, the divergence of the momentum field, and the skeleton extracted using the density-corrected method. The images in the top row were smoothed with a kernel of width 0.75 pixels and the skeleton was extracted using a threshold of 0.2. For the images in the bottom row the width of the smoothing kernel was 1.25 pixels and the threshold was 0.7. In general, a smoothing radius of the order of one pixel is sufficient to counter boundary noise caused by discretization error. These results demonstrate that the density corrected method is much less sensitive to the amount of smoothing and to the value of the threshold.

Fig. 8 shows the effect of smoothing on the direction of the velocity field ∇D obtained from two test images. Here, the smoothing radius is varied from 0 pixels to 8 pixels. It is clear that without smoothing, then the extracted vector field is very noisy, especially near the boundaries of the objects. Conversely, as the smoothing radius increases, then the vector field becomes smoother. With very high smoothing kernel radii, there is distortion near the skeletal branches. This is due to the interaction between the Gaussian kernel and the watershed of the distance map. However, the spatial extent of the distortion is limited by the smoothing radius. Moreover, the distortions are symmetric about the skeletal branches. Hence, they have little effect on the location of the branch. Furthermore, the interactions with the watershed, and, hence, the distortions, are reduced as we approach the endpoints of the skeletal branches.

The improved skeleton extraction and reduced parameter sensitivity have a direct effect on the usability of the method for later shape recognition based on the skeletal representation. The reduced parameter sensitivity limits the amount of tuning required for shape extraction prior to recognition. This allows the method to be require less user intervention. The improved accuracy of the extraction process has obvious implications for the




Fig. 8. Effect of smoothing on the direction of the velocity field.

quality of the skeletal representation and, hence, in the performance of the recognition process.

Sensitivity to boundary noise is a problem with all skeleton extraction methods, and is a consequence of the high sensitivity of the skeletal representation to boundary deformation. In practice, this sensitivity is compounded with the accumulation of discretization errors due to high frequency boundary noise. This means that an algorithm that is correct in the continuous domain, can be adversely affected by errors in the discrete domain. Our solution to this problem is to filter to remove the high frequency

components from the shape boundary in order to reduce the effects of discretization error.

A complementary method described in the literature is to extract a noisy skeleton and then prune the extraneous branches [8], [18], [19], [31]. A problem with this approach is that the determination of the endpoints of the unpruned skeletal branches is still affected by the high-frequency boundary noise. Hence, a branch that should terminate within the interior of the shape, and which lies at the center of a large radius (i.e., low curvature) bi-tangent circle, will be prolonged. In fact it can extend from close




Fig. 9. Smoothing improves localization of endpoints.

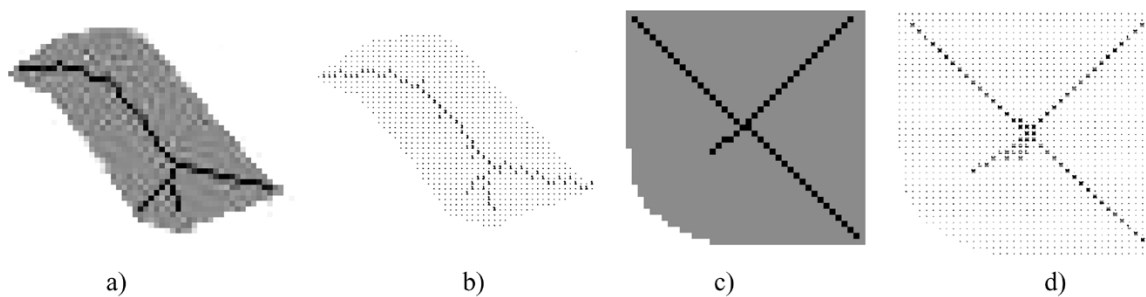


Fig. 10. Controlling the thinning process using the flux of the momentum field eliminates the need for postprocessing to eliminate spurious branches or loops: (a) momentum-based; (b) path-based; (c) momentum-based; (d) path-based.

to the boundary, to the center of a low radius (i.e., high curvature) bitangent circle created by a high frequency boundary feature. This effect is illustrated in Fig. 9 and draws on the ideas presented [31]. Like our method, the algorithm described by Tek and Kimia in [31] is correct in the continuous domain. However, they chose to adopt the pruning approach. The result is that the extracted skeletal branches are, for the most part, the same. However, their method extends the branches to the center of circles that are bitangent to boundary features that are potentially the artefacts of discretization noise.

To be fair, though, the smoothing method used in our algorithm has the disadvantage that for certain degenerate cases, boundary diffusion could generate new skeletal features [4]. In the airplane in the bottom right-hand panel of Fig. 9, it is interesting to note that the skeleton is not symmetrically placed along the shape axis. The reason for this is that the skeleton has a 4-junction which is an unstable configuration. Hence, any small boundary error would cause a break in the symmetry.

Compared to alternative thinning methods [1], [2], [5], the proposed approach is certainly more complex, requiring the integration of the momentum field before the thinning process is performed. However, by using the flux of the momentum field to control the thinning process, there is no need for postprocessing steps for either pruning spurious branches or eliminating spurious loops from the skeleton. Fig. 10 shows a comparison of the skeleton extracted with our method and that extracted using the path-based method described in [5]. Although both algorithms extract similar skeletons, from the results shown in Fig. 10(c) and 10(d), it is clear that the thinning process described in [5] creates a spurious loop that has to be removed by postprocessing.

Fig. 11 plots the number of detected points on the skeleton of the test shape as a joint function of the smoothing radius and the skeletonization threshold. Ideally, as the smoothing radius increases, then the number of detected points should reach a plateau very rapidly, and then drop abruptly to a lower plateau as a new feature of the shape is smoothed away. The amount of smoothing required to reach a new plateau should be independent of the value of the threshold. Fig. 11(a) shows the number of points extracted by the two methods. The results are superimposed as separate surfaces. It is clear from the plot that of the two methods, the density-corrected method reaches the plateau faster as we increase the threshold or the smoothing radius. Moreover, it maintains the plateau for longer. Fig. 11(b) and 11(c) show the results separately for the Hamilton–Jacobi and the density-corrected methods. Here, the location closest to the viewer direction is the plateau side of the plots. The ridges in the forefront show the drop in the number of skeletal points due to the smoothing away of an image feature. In both cases, the drop is sudden but the ridge in the Hamilton–Jacobi plot shows a higher dependence on the threshold.

B. Skeleton Localization

In this section, we characterize the localization properties of the skeleton extracted using the Hamilton–Jacobi method and the new density corrected method on a wide variety of shapes. To this end, we investigate how the values of the divergence of the velocity and of the momentum field are distributed over the distance to the extracted skeleton. Fig. 12 plots a histogram of the distribution of nonskeletal points as a function of distance and divergence value for the test shapes. The figure shows

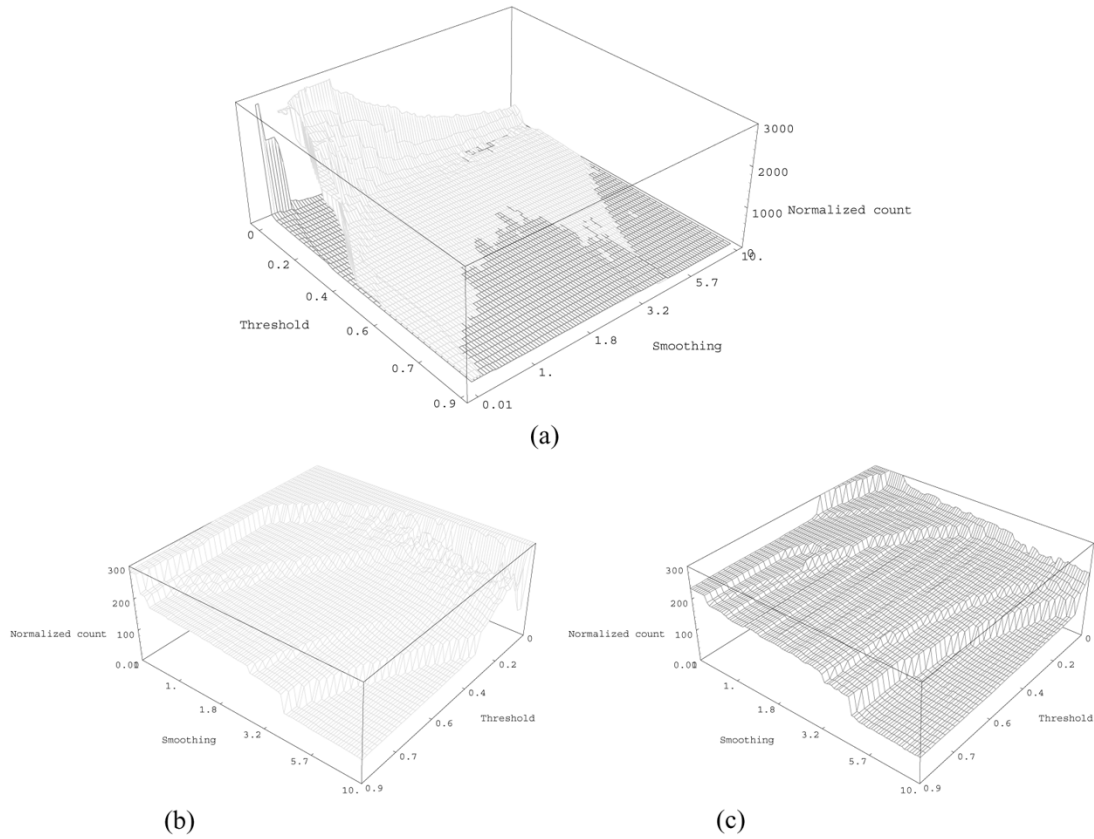


Fig. 11. Effect of smoothing and threshold on skeleton extraction: (a) both; (b) Hamilton–Jacobi; (c) density corrected.




Fig. 12. Histogram over value of (negative) divergence of the field and distance to skeleton: (a) Hamilton–Jacobi; (b) density corrected.

that the Hamilton–Jacobi skeleton has a nonnegligible tail for high divergence values, even at large distance from the extracted skeleton.

We have also performed an experiment aimed at quantifying the localization of the skeleton on a database of shapes. We have used a database of 50 shapes and have histogrammed the distribution of field divergence as a function of the distance to the skeleton. We have repeated this procedure for both the velocity field and the momentum field. For each shape, we take the mean of the relevant divergence-distribution as a measure of divergence-localization.

Fig. 13 shows histograms of this divergence-localization measure accumulated over all the shapes in our database. We

have divided the histogram contents into eight bins of average divergence-distance. In Fig. 13(a), we show the localization histogram for the velocity field. The mean of this distribution is 2.52, while the variance is 0.34. Fig. 13(b) is the corresponding histogram for the momentum field. The mean of this distribution is 1.46, while the variance is 0.28. The density correction clearly leads to a better localization of the skeleton.

VIII. CONCLUSION

In this paper, we present an improvement of the Hamilton–Jacobi method for skeleton extraction that reduces its sensitivity to high curvature. Our analysis takes into account variations of




Fig. 13. Histogram of divergence-localization on a database of 50 shapes: (a) Hamilton-Jacobi; (b) density corrected.

density due to boundary curvature. It is based on the momentum field, rather than velocity field, where the local density is integrated along the path followed by the boundary points. This yields a skeletonization algorithm that is both better localized and less susceptible to boundary noise than the Hamilton-Jacobi method. Our analysis of the effects of boundary noise show that high noise still affect the extraction algorithm. This intrinsic sensitivity is compounded with a higher incidence of discretization error in correspondence with high frequencies in the boundary features. To counter this, we smooth the boundary by approximating a diffusion operator. While this approach reduces the effect of discretization errors, it might smooth out important high-frequencies features and, in some degenerate cases, create artifacts in the form of new skeletal branches. One way to counter this problem could be to adopt other techniques to reduce the effects of discretization noise. For instance, subpixel precision methods such as the one developed by Dimitrov *et al.* in [11] suggest themselves.

REFERENCES

- [1] C. Arcelli and G. Sanniti di Baja, "A width-independent fast thinning algorithm," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. PAMI-7, no. 4, pp. 463-474, Apr. 1985.
- [2] —, "Ridge points in euclidean distance maps," *Pattern Recognit. Lett.*, vol. 13, pp. 237-243, 1992.
- [3] C. Arcelli and G. Sanniti di Baja, "Euclidean skeleton via center-of-maximal-disc extraction," *Image Vis. Comput.*, vol. 11, pp. 163-173, 1993.
- [4] J. August, A. Tannenbaum, and S. W. Zucker, "On the evolution of the skeleton," in *Proc. Int. Conf. Computer Vision*, 1999, pp. 315-522.
- [5] C. Sanniti di Baja and E. Thiel, "Skeletonization algorithm running on path-based distance maps," *Image Vis. Comput.*, vol. 14, pp. 47-57, 1996.
- [6] I. Bitter, A. E. Kaufman, and M. Sato, "Penalized-distance volumetric skeleton algorithm," *IEEE Trans. Vis. Comput. Graph.*, vol. 7, no. 3, pp. 195-206, Jul/Sep. 2001.
- [7] H. Blum, "Biological shape and visual science (part i)," *J. Theoret. Biol.*, vol. 38, pp. 205-287, 1973.
- [8] G. Borgefors, G. Ramella, and G. Sanniti di Baja, "Multi-scale skeletons via permanence ranking," in *Advances in Visual Form Analysis*. Singapore: World Scientific, 1997, pp. 31-42.
- [9] S. Bouix and K. Siddiqi, "Divergence-based medial surfaces," in *Proc. Computer Vision ECCV*, vol. 1, ch. , 2000, pp. 603-618.
- [10] J. Crank and P. Nicolson, "A practical method for numerical evaluation of solutions of partial differential equations of the heat conduction type," in *Proc. Cambridge Philos. Soc.*, vol. 43, 1947, pp. 50-67.
- [11] P. Dimitrov, J. N. Damon, and K. Siddiqi, "Flux invariants for shape," in *Proc. IEEE Conf. Computer Vision Pattern Recognition*, vol. I, 2003, pp. 835-841.
- [12] D. Ivanov, E. Kuzmin, and S. Burtsev, "An efficient integer-based skeletonization algorithm," *Comput. Graph.*, vol. 24, pp. 41-51, 2000.
- [13] B. B. Kimia, A. R. Tannenbaum, and S. W. Zucker, "On the evolution of curves via a function of curvature, I: the classical case," *J. MAA*, vol. 163, no. 2, pp. 438-458, 1992.
- [14] R. Kimmel, D. Shaked, N. Kiryati, and A. M. Bruckstein, "Skeletonization via distance maps and level sets," *Comput. Vis. Image Understand.*, vol. 62, no. 3, pp. 382-391, 1995.
- [15] A. P. Kiraly, J. P. Helferty, E. A. Hoffman, G. McLennan, and W. E. Higgins, "Three-dimensional path planning for virtual bronchoscopy," *IEEE Trans. Med. Imag.*, vol. 23, no. 11, pp. 1365-1379, Nov. 2004.
- [16] F. Leymarie and M. D. Levine, "Simulating the grassfire transform using an active contour model," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 14, no. 1, pp. 56-75, Jan. 1992.
- [17] D. Macrini *et al.*, "View-based 3-D object recognition using shock graphs," in *Proc. Int. Conf. Pattern Recognition*, 2002, pp. 24-28.
- [18] R. L. Ogniewicz, "A multiscale mat from voronoi diagrams: The skeleton-space and its application to shape description and decomposition," in *Aspects of Visual Form Processing*. Singapore: World Scientific, 1994, pp. 430-439.
- [19] R. L. Ogniewicz and O. Kübler, "Hierarchic voronoi skeletons," *Pattern Recognit.*, vol. 28, no. 3, pp. 343-359, 1995.
- [20] S. J. Osher and J. A. Sethian, "Fronts propagating with curvature dependent speed: Algorithms based on Hamilton-Jacobi formulations," *J. Comput. Phys.*, vol. 79, pp. 12-49, 1988.
- [21] M. Pelillo, K. Siddiqi, and S. W. Zucker, "Matching hierarchical structures using association graphs," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 21, no. 11, pp. 1105-1120, Nov. 1999.
- [22] M. Schmitt, "Some examples of algorithms analysis in computational geometry by means of mathematical morphological techniques," presented at the Geometry Robotics, 1989. LNCS 391.
- [23] T. S. Sebastian, P. N. Klein, and B. B. Kimia, "Recognition of shapes by editing shock graphs," in *Proc. Int. Conf. Computer Vision*, vol. I, 2001, pp. 755-762.
- [24] D. Shaked and A. M. Bruckstein, "Pruning medial axes," *Comput. Vis. Image Understand.*, vol. 69, no. 2, pp. 156-169, 1998.
- [25] E. C. Sherbrooke, N. M. Patrikalakis, and E. Brisson, "An algorithm for the medial axis transform of 3D polyhedral solids," *IEEE Trans. Vis. Comput. Graph.*, vol. 2, no. 1, pp. 44-61, Mar. 1996.
- [26] A. Shokoufandeh, S. J. Dickinson, K. Siddiqi, and S. W. Zucker, "Indexing using a spectral encoding of topological structure," in *Proc. IEEE Conf. Computer Vision Pattern Recognition*, Jun. 1999, pp. 491-497.
- [27] K. Siddiqi, S. Bouix, A. Tannenbaum, and S. W. Zucker, "The Hamilton-Jacobi skeleton," in *Proc. Int. Conf. Computer Vision*, Sep. 1999, pp. 828-864.
- [28] —, "Hamilton-Jacobi skeletons," *Int. J. Computer Vision*, vol. 3, pp. 215-231, 2002.
- [29] K. Siddiqi and B. B. Kimia, "A shock grammar for recognition," in *IEEE Conf. Computer Vision Pattern Recognition*, Jun. 1996, pp. 507-513.
- [30] K. Siddiqi, A. Shokoufandeh, S. J. Dickinson, and S. W. Zucker, "Shock graphs and shape matching," *Int. J. Comput. Vis.*, vol. 35, no. 1, pp. 13-32, 1999.
- [31] H. Tek and B. B. Kimia, "Symmetry maps of free-form curve segments via wave propagation," *Int. J. Comput. Vis.*, vol. 54, no. 1-3, pp. 35-81, 2003.

- [32] S. Tirthapura, D. Sharvit, P. Klein, and B. B. Kimia, "Indexing based on edit-distance matching of shape graphs," in *Proc. SPIE Int. Symp. Voice, Video, Data Communications*, 1998, pp. 25–36.
- [33] P. J. Yim, P. L. Choyke, and R. M. Summers, "Gray-scale skeletonization of small vessels in magnetic resonance angiography," *IEEE Trans. Med. Imag.*, vol. 19, no. 4, pp. 568–576, Apr. 2000.
- [34] Y. Zhou and A. W. Toga, "Efficient skeletonization of volumetric objects," *IEEE Trans. Vis. Comput. Graph.*, vol. 5, no. 3, pp. 196–209, Jul./Sep. 1999.



Andrea Torsello received the "Laurea" degree (with honors) in computer science from the Ca' Foscari University of Venice, Venice, Italy, in 1997, and the Ph.D. degree in computer science at the University of York, York, U.K., in 2004.

He currently holds a research position at the Ca' Foscari University of Venice. His research interests are in the areas of computer vision and pattern recognition, particularly the application of stochastic and structural approaches to shape analysis. He is currently co-editing a special issue of *Pattern Recognition* on "similarity-based pattern recognition."



Edwin R. Hancock received the degree in physics and the Ph.D. degree in high-energy physics from the University of Durham, Durham, U.K., in 1977 and 1981, respectively.

For ten years, he was a Researcher in the fields of high-energy nuclear physics and pattern recognition at the Rutherford-Appleton Laboratory (now the Central Research Laboratory of the Research Councils), Cheshire, U.K. During this period, he also held adjunct teaching posts at the University of Surrey, Surrey, U.K., and the Open University,

Milton Keynes, U.K. In 1991, he joined the University of York, York, U.K., as a Lecturer in the Department of Computer Science. He was promoted to Senior Lecturer in 1997 and to Reader in 1998. In 1998, he was appointed to a Chair in Computer Vision. He now leads a group of some 15 faculty, research staff, and Ph.D. students working in the areas of computer vision and pattern recognition. He has been a Guest Editor for special issues of the journals *Image and Vision Computing* and *Pattern Recognition*. He has been a member of the Editorial Board for *Pattern Recognition*. He has published over 80 journal papers and 300 refereed conference publications. He has been on the program committees for numerous national and international meetings. In 1997, with Marcello Pelillo, he established a new series of international meetings on energy minimization methods in computer vision and pattern recognition. His main research interests are in the use of optimization and probabilistic methods for high- and intermediate-level vision. He is also interested in the methodology of structural and statistical pattern recognition. He is currently working on graph matching, shape-from-X, image databases, and statistical learning theory. His work has found applications in areas such as radar terrain analysis, seismic section analysis, remote sensing, and medical imaging.

Prof. Hancock became a Fellow of the International Association for Pattern Recognition in 1998. He has been a member of the Editorial Board of the IEEE TRANSACTIONS ON PATTERN ANALYSIS AND MACHINE INTELLIGENCE. He was awarded the Pattern Recognition Society medal in 1991 and an Outstanding Paper Award in 1997 by the journal *Pattern Recognition*.