

Divergence-Based Medial Surfaces

Sylvain Bouix and Kaleem Siddiqi

McGill University
School of Computer Science &
Center for Intelligent Machines
3480 University Street
Montreal, QC H3A 2A7, Canada
{sbouix,siddiqi}@cim.mcgill.ca

Abstract. The medial surface of a volumetric object is of significant interest for shape analysis. However, its numerical computation can be subtle. Methods based on Voronoi techniques preserve the object's topology, but heuristic pruning measures are introduced to remove unwanted faces. Approaches based on Euclidean distance functions can localize medial surface points accurately, but often at the cost of altering the object's topology. In this paper we introduce a new algorithm for computing medial surfaces which addresses these concerns. The method is robust and accurate, has low computational complexity, and preserves topology. The key idea is to measure the net outward flux of a vector field per unit volume, and to detect locations where a conservation of energy principle is violated. This is done in conjunction with a thinning process applied in a cubic lattice. We illustrate the approach with examples of medial surfaces of synthetic objects and complex anatomical structures obtained from medical images.

1 Introduction

Medial surface based representations are of significant interest for a number of applications in biomedicine, including object representation [20,30], registration [16] and segmentation [25]. Such descriptions are also popular for animating objects in graphics [31,23] and manipulating them in computer-aided design. They provide a compact representation while preserving the object's genus and retain sufficient local information to reconstruct (a close approximation to) it. This facilitates a number of important tasks including the quantification of the local width of a complex structure, e.g., the grey matter in the human brain, and the analysis of its topology, e.g., the branching pattern of blood vessels in angiography images. Graph-based abstractions of such data have also been proposed [9]. Despite their popularity, the stable numerical computation of medial surfaces remains a challenging problem. Unfortunately, the classical difficulties associated with computing their 2D analog, the Blum skeleton, are only exacerbated when a third dimension is added.

1.1 Background

The 2D skeleton of a closed set $A \subset \mathcal{R}^2$ is the locus of centers of maximal open discs contained within the complement of the set [5]. An open disc is maximal if there exists no other open disc contained in the complement of A that properly contains the disc. The medial surface of a closed set $A \subset \mathcal{R}^3$ is defined in an analogous fashion as the locus of centers of maximal open spheres contained in the complement of the set. It is often referred to as the 3D skeleton, though this term is misleading since it is in fact comprised of a collection of 3D points, curves and surfaces [3]. Whereas the above definition is quite general, in the current context we shall assume that the closed set A is the bounding surface of a volumetric object. Hence, this set will have two complementary medial surfaces, one inside the volume and the other outside it. In most cases we shall be referring to the former, though the development applies to both.

Interest in the medial surface as a representation for a volumetric object stems from a number of useful properties: i) it is a thin set, i.e., it contains no interior points, ii) it is homotopic to the volume, iii) it is invariant under Euclidean transformations of the volume (rotations and translations), and iv) given the radius of the maximal inscribed sphere associated with each medial surface point, the volumetric object can be reconstructed exactly. Hence, it provides a compact representation while preserving the object's genus and making certain properties explicit, such as its local width.

Approaches to computing skeletons and medial surfaces can be broadly organized into three classes. First, methods based on *thinning* attempt to realize Blum's grassfire formulation [5] by peeling away layers from an object, while retaining special points [2,14,19]. It is possible to define erosion rules in a lattice such that the topology of the object is preserved. However, these methods are quite sensitive to Euclidean transformations of the data and typically fail to localize skeletal or medial surface points accurately. As a consequence, only a coarse approximation to the object is usually reconstructed [19,4,14].

Second, it has been shown that under appropriate smoothness conditions, the vertices of the Voronoi diagram of a set of boundary points converges to the exact skeleton as the sampling rate increases [24]. This property has been exploited to develop skeletonization algorithms in 2D [21], as well as extensions to 3D [26, 27]. The dual of the Voronoi diagram, the Delaunay triangulation (or tetrahedralization in 3D) has also been used extensively. Here the skeleton is defined as the locus of centers of the circumscribed spheres of each tetrahedra [12,20]. Both types of methods preserve topology and accurately localize skeletal or medial surface points, provided that the boundary is sampled densely. Unfortunately, however, the techniques used to prune faces and edges which correspond to small perturbations of the boundary are typically based on heuristics. In practice, the results are not invariant under Euclidean transformations and the optimization step, particularly in 3D, can have a high computational complexity [20].

A third class of methods exploits the fact that the locus of skeletal or medial surface points coincides with the singularities of a *Euclidean distance function* to the boundary. These approaches attempt to detect local maxima of the distance

function, or the corresponding discontinuities in its derivatives [1,15,13]. The numerical detection of these singularities is itself a non-trivial problem; whereas it may be possible to localize them, ensuring homotopy with the original object is difficult.

In recent work we observed that the grassfire flow leads to a hamilton-jacobi equation, which by nature is conservative in the smooth regime of its underlying phase space [29]. Hence, we suggested that a measurement of the net outward flux per unit volume of the gradient vector field of the Euclidean distance function could be used to associate locations where a conservation of energy principle was violated with medial surface points [28]. Unfortunately, in practice, the resulting medial surface was not guaranteed to preserve the topology of the object, since the flux computation was a purely local operation. The main contribution of the current paper is the combination of the flux measurement with a homotopy preserving thinning process applied in a cubic lattice. The method is robust and accurate, has low computational complexity and is now guaranteed to preserve topology. There are other promising recent approaches which combine aspects of thinning, Voronoi diagrams and distance functions [18,33,8,32]. In spirit, our method is closest to that of [18]¹ but is grounded in principles from physics. We illustrate the algorithm with a number of examples of medial surfaces of synthetic objects and complex anatomical structures obtained from medical images.

2 Hamiltonian Medial Surfaces

We shall first review the hamilton-jacobi formulation used to simulate the eikonal equation as well as detect singularities in [29,28]. Consider the grassfire flow

$$\frac{\partial S}{\partial t} = \mathcal{N} \quad (1)$$

acting on a 3D surface S , such that each point on its boundary is moving with unit speed in the direction of the inward normal \mathcal{N} . In physics, such equations are typically solved by looking at the evolution of the phase space of an equivalent Hamiltonian system. Since Hamiltonian systems are conservative, the locus of skeletal points (in 2D) or medial surface points (in 3D) coincides with locations where a conservation of energy principle is violated. This loss of energy can be used to formulate a natural criterion for detecting singularities of the distance function.

In more formal terms, let D be the Euclidean distance function to the initial surface S_0 . The magnitude of its gradient, $\|\nabla D\|$, is identical to 1 in its smooth regime. With $\mathbf{q} = (x, y, z)$, $\mathbf{p} = (D_x, D_y, D_z)$, associate to the surface $S \subset \mathbf{R}^3$, evolving according to Eq. 1, the surface $\tilde{C} \subset \mathbf{R}^6$ given by $\tilde{C} := \{(x, y, z, D_x, D_y, D_z) : (x, y, z) \in S, D_x^2 + D_y^2 + D_z^2 = 1, \mathbf{p} \cdot \mathbf{q} = 1\}$.

¹ Malandain and Fernandez-Vidal use a heuristic estimation of the singularities of a distance function to obtain an initial skeleton or medial surface, and then perform a topological reconstruction to ensure homotopy with the original shape.

The Hamiltonian function obtained by applying a Legendre transformation to the Lagrangian $L = \|\dot{\mathbf{q}}\|$ is given by

$$H = \mathbf{p} \cdot \dot{\mathbf{q}} - L = 1 - (D_x^2 + D_y^2 + D_z^2)^{\frac{1}{2}}.$$

The associated Hamiltonian system is:

$$\dot{\mathbf{p}} = -\frac{\partial H}{\partial \mathbf{q}} = (0, 0, 0), \quad \dot{\mathbf{q}} = \frac{\partial H}{\partial \mathbf{p}} = -(D_x, D_y, D_z). \quad (2)$$

\tilde{C} can be evolved under this system of equations, with $\tilde{C}(t) \subset \mathbf{R}^6$ denoting the resulting (contact) surface. The projection of $\tilde{C}(t)$ onto \mathbf{R}^3 will then give the parallel evolution of \mathcal{S} at time t , $\mathcal{S}(t)$. Note that the interpretation of Eq. 2 is quite intuitive: the gradient vector field \mathbf{p} does not change with time, and points on the boundary of the surface move in the direction of the inward normal with unit velocity.

It is straightforward to show that all Hamiltonian systems are conservative [22, p. 172]:

Theorem 1. *The total energy $H(\mathbf{p}, \mathbf{q})$ of the Hamiltonian system (2) remains constant along trajectories of (2).*

Proof. The total derivative of $H(\mathbf{p}, \mathbf{q})$ along a trajectory $\mathbf{p}(t)$, $\mathbf{q}(t)$ of (2) is given by

$$\frac{dH}{dt} = \frac{\partial H}{\partial \mathbf{p}} \cdot \dot{\mathbf{p}} + \frac{\partial H}{\partial \mathbf{q}} \cdot \dot{\mathbf{q}} = \frac{\partial H}{\partial \mathbf{p}} \cdot \frac{\partial H}{\partial \mathbf{q}} - \frac{\partial H}{\partial \mathbf{p}} \cdot \frac{\partial H}{\partial \mathbf{q}} = 0.$$

Thus $H(\mathbf{p}, \mathbf{q})$ is constant along any trajectory of (2).

The analysis carried out thus far applies under the assumption of a *central field of extremals* such that trajectories of the Hamiltonian system do not intersect. Conversely, when trajectories intersect, the conservation of energy principle will be violated (energy will be absorbed). This loss of energy can be used to formulate a robust and efficient algorithm for detecting singularities of the distance function D , which correspond to medial surface points.

The key is to measure the flux of the vector field $\dot{\mathbf{q}}$, which is analogous to the flow of an incompressible fluid such as water. Note that for a volume with an enclosed surface, an excess of outward or inward flow through the surface indicates the presence of a *source*, or a *sink*, respectively, in the volume. The latter case is the one we are interested in, and the net outward flux is related to the divergence of the vector field. More specifically, the divergence of a vector field at a point, $\text{div}(\dot{\mathbf{q}})$, is defined as the net outward flux per unit volume, as the volume about the point shrinks to zero:

$$\text{div}(\dot{\mathbf{q}}) \equiv \lim_{\Delta v \rightarrow 0} \frac{\int_S \langle \dot{\mathbf{q}}, \mathcal{N} \rangle ds}{\Delta v} \quad (3)$$

Here Δv is the volume, S is its surface and \mathcal{N} is the outward normal at each point on its surface. This definition can be shown to be equivalent to the classical

definition of divergence as the sum of the partial derivatives with respect to each of the vector field's component directions:

$$\operatorname{div}(\dot{\mathbf{q}}) = \frac{\partial q_{x_1}}{\partial x_1} + \dots + \frac{\partial q_{x_n}}{\partial x_n} \quad (4)$$

However, Eq. 4 cannot be used at points where the vector field is singular, and hence is not differentiable. These are precisely the points we are interested in, and Eq. 3 offers significant advantages for detecting medial surface points. In particular, the numerator, which represents the net outward flux of the vector field through the surface which bounds the volume, is an index computation on the vector field. It is not surprising that this is numerically much more stable than the estimation of derivatives in the vicinity of singularities. Via the divergence theorem,

$$\int_v \operatorname{div}(\dot{\mathbf{q}}) dv \equiv \int_S \langle \dot{\mathbf{q}}, \mathcal{N} \rangle ds. \quad (5)$$

Hence, the net outward flux through the surface which bounds a finite volume is just the volume integral of the divergence of the vector field within that volume. Locations where the flux is negative, and hence energy is lost, correspond to sinks or medial surface points.

We now have a robust method for localizing medial surface points by discretizing Eq. 5, and thresholding to select points with negative total outward flux. However, since the computation is local, global properties such as the preservation of the object's topology, are not ensured. In our earlier work we have observed that the method gives accurate medial surfaces, but that as the threshold is varied new holes or cavities may be introduced and the medial surface may get disconnected [28]. The remedy, as we shall now show, is to introduce additional criteria along the lines of those incorporated in [18], to ensure that the medial surface is homotopic to the original object.

3 Homotopy Preserving Medial Surfaces

Our goal is to combine the divergence computation with a thinning process acting in the cubic lattice, such that as many points as possible are removed without altering the object's topology. A point is called a *simple* point if its removal does not change the topology of the object. Hence in 3D, its removal must not disconnect the object, create a hole, or create a cavity. We shall adopt a formal definition of a simple point introduced by Malandain et al. [17]. First we review a few basic concepts in digital topology.

3.1 Digital Topology

In 3D digital topology, the input is a binary (foreground and background) image stored in a 3D array. We shall consider only cubic lattices, where a point is viewed as a unit cube with 6 faces, 12 edges and 8 vertices. For each point, three types of neighbors are defined:




Fig. 1. 6-neighborhoods, 18-neighborhoods and 26-neighborhoods in a cubic lattice.

- *6-neighbors*: two points are 6-neighbors if they share a face,
- *18-neighbors*: two points are 18-neighbors if they share a face or an edge, and
- *26-neighbors*: two points are 26-neighbors if they share a face, an edge or a vertex.

The above definitions induce three types of connectivity, denoted *n-connectivity*, where $n \in \{6, 18, 26\}$, as well as three different *n-neighborhoods* for x , called $N_n(x)$ (see Figure 1). A *n-neighborhood* without its central point is defined as $N_n^* = N_n(x) \setminus \{x\}$. A few more definitions are needed to characterize simple points:

- An object A is *n-adjacent* to an object B , if there exist two points $x \in A$ and $y \in B$ such that x is an *n-neighbor* of y .
- A *n-path* from x_1 to x_k is a sequence of points x_1, x_2, \dots, x_k , such that for all x_i , $1 < i \leq k$, x_{i-1} is *n-adjacent* to x_i .
- An object represented by a set of points O is *n-connected*, if for every pair of points $(x_i, x_j) \in O \times O$, there is a *n-path* from x_i to x_j .

Based on these definitions, Malandain et al. provide a topological classification of a point x in a cubic lattice by computing two numbers [17]: i) C^* : the number of 26-connected components 26-adjacent to x in $O \cap N_{26}^*$, and ii) \bar{C} : the number of 6-connected components 6-adjacent to x in $\bar{O} \cap N_{18}$. An important result with respect to our goal of thinning is that if $C^* = 1$ and $\bar{C} = 1$, the point is *simple*, and hence removable. When ensuring 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 shape. However, the relationship to the medial surface will be uncertain since the locus of surviving points will depend entirely on the order in which the simple points have been removed. In the current context, we have derived a natural criterion for ordering the thinning, based on the divergence of the gradient vector field of the Euclidean distance function.

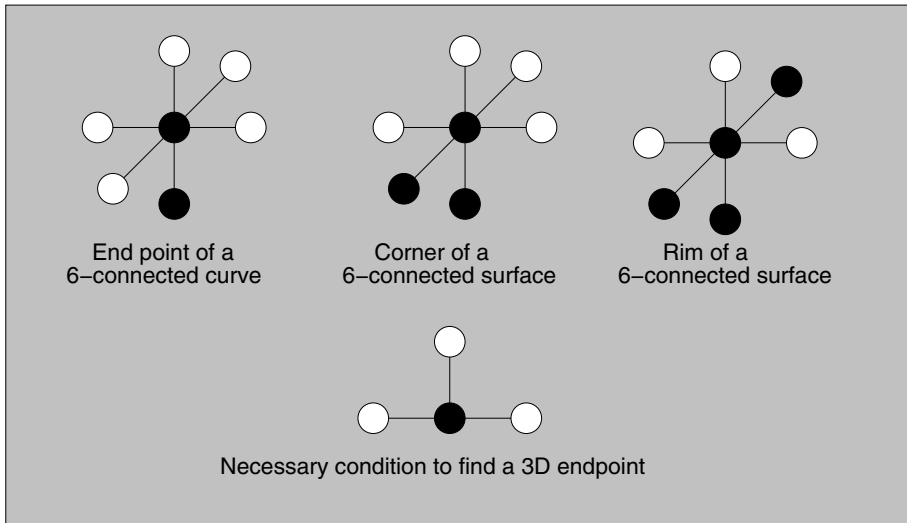


Fig. 2. An endpoint is defined as the end of a 6-connected curve or the corner or rim of a 6-connected surface in 3D. For each configuration, there exists at least one plane in which the point has at least three background 6-neighbors.

3.2 Divergence-Ordered Thinning

Recall from Section 2, that a conservation of energy principle is violated at medial surface points. The total outward flux of the gradient vector field of the Euclidean distance function is negative at such points, since they correspond to sinks.² More importantly, the magnitude of the total outward flux is proportional to the amount of energy absorbed, and hence provides a natural measure of the “strength” of a medial surface point, which we shall use 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 outward flux below some chosen (negative) value, or both. This will accurately localize the medial surface, 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 endpoint in a cubic lattice. Such a point would correspond to the endpoint of a curve, 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 endpoints*. Hence, all surviving points which are not endpoints would not be simple, and the result would be a thin set.

To facilitate this task, we shall restrict our definition of an endpoint to a 6-connected neighborhood. In other words, an endpoint is either the end of a

² Conversely, medial surface points of the background correspond to sources, with positive total outward flux.

6-connected curve, or a corner or point on the rim of a 6-connected surface. It is straightforward to enumerate the possible 6-connected neighborhoods of the endpoint, and to show that they fall into one of three configurations (see Figure 2). Notice that for each configuration, there exists at least one plane in which the point has at least three background 6-neighbors. This gives us a necessary condition to determine if a point is an endpoint according to our definition. Note that before performing this check, one must also verify that the point is simple.

3.3 The Algorithm

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. Once a border point is removed, only its neighbors may become removable. This suggests the implementation of the thinning process using a heap. A full description of the procedure can be found in Algorithm 1.

Algorithm 1 The divergence-ordered thinning algorithm.

Part I: Total Outward Flux

Compute the distance transform of the object D .

Compute the gradient vector field ∇D .

Compute the net outward flux of ∇D using Eq. 5

For each point P in the interior of the object

$$Flux(P) = \sum_{i=1}^n \langle N_i, \nabla D(P_i) \rangle,$$

where P_i is a 26-neighbor of P and N_i is the outward normal at P_i of the unit sphere in 3D, centered at P .

Part II: Homotopy Preserving Thinning

For each point P on the boundary of the object

if (P is simple)

insert(P , Heap) with $Flux(P)$

as the sorting key for insertion

While (Heap.size > 0)

$P = \text{HeapExtractMax}(\text{Heap})$

if (P is simple)

if (P is not an endpoint) or ($Flux(P) > Thresh$)

Remove P

for all neighbors Q of P

if (Q is simple)

insert(Q , Heap)

else mark P as a skeletal (end) point

end { if }

end { if }

end { while }

We now analyze the complexity of the algorithm. The computation of the distance transform, the gradient vector field and the total outward flux are all $\mathcal{O}(n)$ operations. Here n is the total number of points in the 3D array. The implementation of the thinning is more subtle. We claim an $\mathcal{O}(k \log(k))$ worst case complexity, where k is the number of points in the volumetric object. The explanation is as follows. At first store only the points that are on the outer layer of the object in a heap, using the total outward flux as the sorting key for insertion. The extraction of the maximum from the heap will provide the best candidate for removal. If this point is removable, then delete it from the object and add its *simple* (potentially removable) neighbors to the heap. A point can only be inserted a constant number of times (at most 26 times for a 26-neighborhood), and insertion in a heap, as well as the extraction of the maximum, are both $\mathcal{O}(\log(l))$ operations, where l is the number of elements in the heap. There cannot be more than k elements in the heap, because we only have a total of k points in the volume. The worst case complexity for thinning is therefore $\mathcal{O}(k \log(k))$. Hence, the complexity of the algorithm is $\mathcal{O}(n) + \mathcal{O}(k \log(k))$.

4 Examples




Fig. 3. FIRST COLUMN: Three views of a cube. SECOND COLUMN: The corresponding medial surfaces computed using the algorithm of [19]. THIRD COLUMN: The object reconstructed from the medial surfaces in the previous column. FOURTH COLUMN: The corresponding divergence-based medial surfaces. FIFTH COLUMN: The object reconstructed from the medial surfaces in the previous column.

We illustrate the algorithm with both synthetic data and volumes segmented from MR and MRA images. In these simulations we have used the D -Euclidean distance transform, which provides a close approximation to the true Euclidean distance function [6]. The only free parameter is the choice of the divergence value below which the removal of endpoints is blocked. For all examples, this was selected so that approximately 25% of the points within the volume had a lower divergence value.

Figures 3 and 4 compare our approach with the parallel thinning method introduced by Manzanera et al. [19]. The results reveal that both frameworks are robust, and yield structures that are homotopic to the underlying object. However, note that the latter method yields only a subset of the “true” medial surface for these data sets, and hence only a coarse approximation to the object is possible. In contrast, a near perfect reconstruction is possible from the divergence-based medial surfaces.




Fig. 4. FIRST COLUMN: Three views of a cylinder. SECOND COLUMN: The corresponding medial surfaces computed using the algorithm of [19]. THIRD COLUMN: The object reconstructed from the medial surfaces in the previous column. FOURTH COLUMN: The corresponding divergence-based medial surfaces. FIFTH COLUMN: The object reconstructed from the medial surfaces in the previous column.

Recall from Section 3 that the computation of the distance transform, its gradient vector field and the total outward flux are all $\mathcal{O}(n)$ operations. Hence, one may be tempted to simply threshold the divergence map below a certain (negative) value to obtain a medial surface extremely efficiently, as in [28]. Unfortunately, due to discretization, the result will not always be satisfactory. On the one hand, if the threshold is too high, slight perturbations of the boundary will be represented, and the resulting structure will not be a thin set. On the




Fig. 5. FIRST COLUMN: Four views of the ventricles of a brain, segmented from volumetric MR data using an active surface. SECOND COLUMN: The corresponding medial surfaces obtained by thresholding the divergence map. THIRD COLUMN: The divergence-based medial surfaces obtained using the same threshold, but with the incorporation of homotopy preserving thinning. FOURTH COLUMN: The ventricles reconstructed from the divergence-based medial surfaces in the previous column.

other hand, lowering the threshold can provide a thin set, but at the cost of altering the object's topology. This is illustrated in Figure 5 (second column) where the medial surfaces corresponding to the views in the first column are accurate and thin, but have holes. This motivates the need for the topological constraints along with the characterization of endpoints discussed in Section 3. Observe that with the same threshold as before, the divergence-based thinning algorithm now yields a thin structure which preserves topology, Figure 5 (third column). The ventricles reconstructed from the medial surfaces in the fourth column are shown in the fifth column.

Next, we illustrate the robustness of the approach on a (partial) data set of blood vessels obtained from an MRA image of the brain, in Figure 6. The blood vessels have complex topology with loops (due to pathologies), and are already quite thin in several places. The bottom row illustrates the accuracy of the method, where the medial surfaces are shown embedded within the original data set. Generically these structures are thin sheets which approach 3D curves when the blood vessels become perfectly cylindrical. In a number of medical applications where the objects are tubular structures, an explicit reduction of the medial surface to a set of 3D curves is of interest [11,7,8,33]. There is a straightforward modification of our framework which allows this, provided that certain special points on the medial surface have been identified. The essential




Fig. 6. TOP ROW: Blood vessels segmented from volumetric MRA data, with magnified parts shown in the middle and right columns. MIDDLE ROW: The corresponding divergence-based medial surfaces. THIRD ROW: The divergence-based medial surfaces (solid) are shown within the vessel surfaces (transparent).

idea is to preserve such special points, but to remove all other simple points in a sequence ordered by their divergence values. This is illustrated for a portion of the vessel data in Figure 7, where the endpoints of 3 branches were selected as special points. Observe that the result is now composed of three 1 voxel wide 26-connected 3D digital curves.

As a final example, Figure 8 illustrates the medial surface of the sulcii of a brain, where we have shown an X, Y and Z slice through the volume. Observe that the medial surface is well localized, and captures the complex topology of the object's shape.




Fig. 7. LEFT COLUMN: Blood vessels segmented from volumetric MRA data, with a magnified portion shown in the second row. MIDDLE COLUMN: The divergence-based 3D curves. RIGHT COLUMN: The divergence-based 3D curves are shown embedded within the vessel data.

5 Conclusions

We have introduced a novel algorithm for computing medial surfaces which is robust and accurate, computationally efficient, invariant to Euclidean transformations and homotopy preserving. The essential idea is to combine a divergence computation on the gradient vector field of the Euclidean distance function to the object's boundary with a thinning process that preserves topology. The characterization of simple (or removable) points is adopted from [17], but we have also introduced a notion of an endpoint of a 6-connected structure, in order that the algorithm may converge to a thin set. We have illustrated the advantages of the approach on synthetic and real binary volumes of varying complexity.

We note that in related work, Malandain and Fernandez-Vidal obtain two sets based on thresholding a function of two measures, ϕ and d , to characterize the singularities of the Euclidean distance function [18]. The first set preserves topology but captures many unwanted details and is not thin, while the second set provides a better approximation to the skeleton or medial surface, but often alters the object's topology. The two sets are combined using a topological reconstruction process. Whereas empirical results have been good, the choice of appropriate thresholds are context dependent. More seriously, ϕ and d , as well as their combination, are all based on heuristics.




Fig. 8. TOP ROW: Medial surfaces of the sulci of a brain, segmented from an MR image. The three columns represent X, Y and Z slices through the volume. The cross section through the medial surface in each slice is shown in black, and the object is shown in grey. BOTTOM ROW: A zoom-in on a selected region of the corresponding slice in the top row, to show detail.

In contrast, our method is rooted in a physics-based analysis of the gradient vector field of the Euclidean distance function, which shows that a conservation of energy principle is violated at medial surface points. This justifies the use of the divergence theorem to compute the total outward flux of the vector field, and to locate points where energy is absorbed. It should be clear that whereas we have focussed on the interior of an object, the medial surface of the background can be similarly obtained by locating points that act as sources, and have positive total outward flux. Furthermore, both medial surfaces can be located with sub-voxel accuracy by using the local gradient vector field to shift the final set of digital points. In related work we have demonstrated this idea for 2D shapes, where a similar framework was used to compute sub-pixel 2D skeletons and skeletal graphs [10].

In future work we plan to incorporate the topological classification of [17] to parse the medial surface and obtain a more abstract representation of it, e.g., as a graph. We shall also explore the possibility of finding necessary as well as sufficient conditions for defining endpoints in a cubic lattice.

Acknowledgements We are grateful to Allen Tannenbaum and Steve Zucker for collaborations on the hamilton-jacobi formulation. Louis Collins, Georges Le

Goualher, Belinda Lee and Terry Peters kindly supplied the medical data. This research was supported by CFI, FCAR and NSERC.

References

1. C. Arcelli and G. S. di Baja. Ridge points in euclidean distance maps. *Pattern Recognition Letters*, 13(4):237–243, 1992.
2. C. Arcelli and G. Sanniti di Baja. A width-independent fast thinning algorithm. *IEEE PAMI*, 7(4):463–474, July 1985.
3. V. I. Arnold. *Mathematical Methods of Classical Mechanics*. Springer-Verlag, 1989.
4. G. Bertrand. A parallel thinning algorithm for medial surfaces. *Pattern Recognition Letters*, 16:979–986, 1995.
5. H. Blum. Biological shape and visual science. *Journal of Theoretical Biology*, 38:205–287, 1973.
6. G. Borgefors. Distance transformations in arbitrary dimensions. *CVGIP*, 27:321–345, 1984.
7. G. Borgefors, I. Nystrom, and G. S. D. Baja. Skeletonizing volume objects part ii: From surface to curve skeleton. In *SSPR'98*, pages 220–229, 1998.
8. G. Borgefors, I. Nystrom, and G. S. D. Baja. Computing skeletons in three dimensions. *Pattern Recognition*, 32:1225–1236, 1999.
9. E. Bullitt, S. Aylward, A. Liu, J. Stone, S. K. Mukherjee, C. Coffey, G. Gerig, and S. M. Pizer. 3d graph description of the intracerebral vasculature from segmented mra and tests of accuracy by comparison with x-ray angiograms. In *IPMI'99*, pages 308–321, 1999.
10. P. Dimitrov, C. Phillips, and K. Siddiqi. Robust and efficient skeletal graphs. In *CVPR'2000*, Hilton Head, South Carolina, June 2000.
11. Y. Ge, D. R. Stelts, X. Zha, D. Vining, and J. Wang. Computing the central path of the colon from ct images. In *SPIE International Symposium on Medical Imaging*, volume 3338(1), pages 702–713, San Francisco, 1998.
12. J. A. Goldak, X. Yu, and A. K. abd Lingxian Dong. Constructing discrete medial axis of 3-d objects. *Int. Journal of Computational Geometry and Applications*, 1(3):327–339, 1991.
13. J. Gomez and O. Faugeras. Reconciling distance functions and level sets. Technical Report TR3666, INRIA, April 1999.
14. T.-C. Lee and R. L. Kashyap. Building skeleton models via 3-d medial surface/axis thinning algorithm. *CVGIP: Graphical Models and Image Processing*, 56(6):462–478, November 1994.
15. F. Leymarie and M. D. Levine. Simulating the grassfire transform using an active contour model. *IEEE PAMI*, 14(1):56–75, jan 1992.
16. A. Liu, E. Bullitt, and S. M. Pizer. 3d/2d registration via skeletal near projective invariance in tubular objects. In *MICCAI'98*, pages 952–963, 1998.
17. G. Malandain, G. Bertrand, and N. Ayache. Topological segmentation of discrete surfaces. *Int. Journal of Computer Vision*, 10(2):183–197, 1993.
18. G. Malandain and S. Fernandez-Vidal. Euclidean skeletons. *Image and Vision Computing*, 16:317–327, 1998.
19. A. Manzanera, T. M. Bernard, F. Preteux, and B. Longuet. Medial faces from a concise 3d thinning algorithm. In *ICCV'99*, pages 337–343, Kerkyra, Greece, September 1999.

20. M. Näf, O. Kübler, R. Kikinis, M. E. Shenton, and G. Székely. Characterization and recognition of 3d organ shape in medical image analysis using skeletonization. In *IEEE Workshop on Mathematical Methods in Biomedical Image Analysis*, 1996.
21. R. Ogniewicz. *Discrete Voronoi Skeletons*. Hartung-Gorre, 1993.
22. L. Perko. *Differential Equations and Dynamical Systems*. Springer-Verlag, 1986.
23. S. M. Pizer, A. Thall, and D. T. Chen. M-reps: A new object representation for graphics. *ACM Transactions on Graphics (submitted)*, 1999.
24. M. Schmitt. Some examples of algorithms analysis in computational geometry by means of mathematical morphology techniques. In J. Boissonnat and J. Laumond, editors, *Lecture Notes in Computer Science, Geometry and Robotics*, volume 391, pages 225–246. Springer-Verlag, 1989.
25. T. B. Sebastian, H. Tek, J. J. Crisco, S. W. Wolfe, and B. B. Kimia. Segmentation of carpal bones from 3d ct images using skeletally coupled deformable models. In *MICCAI'98*, pages 1184–1194, 1998.
26. D. J. Sheehy, C. G. Armstrong, and D. J. Robinson. Shape description by medial surface construction. *IEEE Transactions on Visualization and Computer Graphics*, 2(1):62–72, 1996.
27. E. C. Sherbrooke, N. Patrikalakis, and E. Brisson. An algorithm for the medial axis transform of 3d polyhedral solids. *IEEE Transactions on Visualization and Computer Graphics*, 2(1):44–61, 1996.
28. K. Siddiqi, S. Bouix, A. Tannenbaum, and S. W. Zucker. The hamilton-jacobi skeleton. In *ICCV'99*, pages 828–834, Kerkyra, Greece, September 1999.
29. K. Siddiqi, A. Tannenbaum, and S. W. Zucker. A hamiltonian approach to the eikonal equation. In *EMMCVPR'99*, pages 1–13, York, UK, July 1999.
30. G. D. Stetten and S. M. Pizer. Automated identification and measurement of objects via populations of medial primitives, with application to real time 3d echocardiography. In *IPMI'99*, pages 84–97, 1999.
31. M. Teichmann and S. Teller. Assisted articulation of closed polygonal models. In *9th Eurographics Workshop on Animation and Simulation*, 1998.
32. H. Tek and B. B. Kimia. Symmetry maps of free-form curve segments via wave propagation. In *ICCV'99*, pages 362–369, Kerkyra, Greece, September 1999.
33. Y. Zhou, A. Kaufman, and A. W. Toga. 3d skeleton and centerline generation based on an approximate minimum distance field. *Int. Journal of the Visual Computer*, 14(7):303–314, 1998.