

Sampled Medial Loci and Boundary Differential Geometry

Svetlana Stolpner
School of Computer Science
McGill University
Montréal, QC, Canada
sveta@cim.mcgill.ca

Sue Whitesides
Department of Computer Science
University of Victoria
Victoria, BC, Canada
sue@uvic.ca

Kaleem Siddiqi
School of Computer Science
McGill University
Montréal, QC, Canada
siddiqi@cim.mcgill.ca

Abstract

We introduce a novel algorithm to compute a dense sample of points on the medial locus of a polyhedral object, with a guarantee that each medial point is within a specified tolerance ϵ from the medial surface. Motivated by Damon’s work on the relationship between the differential geometry of the smooth boundary of an object and its medial surface [8], we then develop a computational method by which boundary differential geometry can be recovered directly from this dense medial point cloud. Experimental results on models of varying complexity demonstrate the validity of the approach, with principal curvature values that are consistent with those provided by an alternative method that works directly on the boundary. As such, we demonstrate the richness of a dense medial point cloud as a shape descriptor for 3D data processing.

1. Introduction

The ubiquity of polyhedral models in computer vision, computer graphics, medical imaging and solid modeling calls for algorithms for the analysis of the smooth object boundaries associated with them. Medial representations, introduced in [3], have found useful applications in this context, including object recognition and retrieval, object segmentation, object registration, statistical shape analysis, motion planning, object animation and path planning [19]. Strikingly, whereas the majority of these applications use a model comprised of medial points and their spoke vectors, few (if any) have capitalized on the fact that these quantities suffice to recover the differential geometry of the implied boundary.

In this article, we present:

1. A novel algorithm to obtain a dense set of points within a one-sided Hausdorff distance ϵ of the object’s medial surface, along with their spoke vectors in Section 3. Medial surface normals are then computed from the

spoke vectors, following which the medial point cloud may be segmented into its constituent smooth sheets.

2. A novel algorithm to obtain the principal curvatures and principal curvature directions of the two implied boundary patches on either side of each medial point in Section 5. These curvature estimates are then mapped onto the polyhedral object boundary B .

These results are the first computational investigations, to our knowledge, which relate medial differential geometry to boundary differential geometry for arbitrary polyhedral models. We demonstrate that a discrete approximation to the medial surface is sufficient to estimate local differential properties of the object boundary.

2. Background and Related Work

We begin with an overview of medial surfaces and their properties. Consider a solid object Ω with boundary B .

Definition 1 *The medial surface \mathcal{MS} of Ω is the locus of centres of all maximal spheres inscribed in Ω .*

As an illustrative example, Figure 1 (Right) depicts the medial surface of a box. This representation has a number of useful properties. First, the original object Ω is reconstructed by the union of its maximal inscribed spheres. Second, the medial surface is of the same homotopy type as Ω [15]. Third, the medial surface has a natural decomposition into parts. For a smooth object, the *generic* (i.e. stable under small perturbations of the object boundary) points of the \mathcal{MS} fall into a small number of classes [12]. The majority lie on sheets (manifolds with boundary) of type A_1^2 , i.e., they are points which have an order-1 contact with the boundary at exactly 2 distinct points. These sheets intersect at curves of A_1^3 points and are bounded by curves of A_3 points (see Figure 1 (Left)). For an A_1^2 point \mathbf{p} , the *object angle* θ is the angle between the vector from \mathbf{p} to either one of its two closest points on B and the tangent plane to \mathcal{MS}




Figure 1. Left: Different classes of points that compose the medial locus of a smooth 3D object [12]. Right: The medial surface of a box, with each sheet shown in a different colour. The object angle θ at the selected medial point is $\pi/4$. The spoke vectors are the black arrows.

at \mathbf{p} (see Figure 1 (Left)). We shall refer to the vectors from \mathbf{p} to its two closest points on B as the *spoke vectors*.

When Ω is a polyhedron, \mathcal{MS} is composed of bisectors of the faces, edges and vertices of B . The bisector of two such elements is a quadric surface and these surfaces intersect along curves of higher algebraic degree. Thus, the complete medial surface of a polyhedron can be complex and is impractical (and often unnecessary) to compute exactly for models with a large number of faces. A common strategy, therefore, is to only recover those parts of the medial surface which are “salient” with respect to an appropriate measure, such as the associated object angle [19].

2.1. Medial Surface Computation

A variety of methods have been proposed for computing medial surfaces of polyhedral models. Amenta *et al.* [1] approximate the medial surface by a subset of the Voronoi vertices, the ‘poles’, of points sampled on the object’s surface. Leymarie and Kimia [14] compute a geometric directed hypergraph whose nodes are medial points, using exact bisector computations between clusters of surface points. Culver *et al.* [7] use exact arithmetic to compute the seams of the medial surface of polyhedra having a small number of faces. Etzion and Rappaport [10] perform spatial subdivision until each cell contains one vertex of the generalized Voronoi diagram of a polyhedron. Foskey *et al.* [11] approximate the medial surface of a mesh having a chosen object angle with axis-aligned facets, with computations accelerated by graphics hardware.

We now specialize to methods for computing the medial surface that are based on the average outward flux of the gradient of the Euclidean distance transform of the object Ω . This approach is well suited to our goal of recovering boundary differential geometry directly from medial geometry as, in this framework, it is possible to efficiently and ac-

curately compute a dense collection of medial points along with their spoke vectors.

Definition 2 The Euclidean distance transform of a solid Ω with boundary B is given by $D(\mathbf{p}) = \inf_{\mathbf{q} \in B} d(\mathbf{p}, \mathbf{q})$, where $p \in \Omega$ and $d(\mathbf{p}, \mathbf{q})$ denotes Euclidean distance.

$\nabla D : \mathbb{R}^3 \rightarrow \mathbb{R}^3$ is hence a unit vector field where each point in the interior of Ω is assigned the direction from the closest point on B , which is also an inward normal to B . The following useful relationships between the average outward flux \mathcal{AOF} of ∇D through a sphere S inside Ω with outward normal N_S and medial sheets in S are shown in [20]. If no medial sheet passes through S ,

$$\lim_{\text{area}(S) \rightarrow 0} \mathcal{AOF} = \lim_{\text{area}(S) \rightarrow 0} \frac{\iint_S \nabla D \cdot N_S dS}{\iint_S dS} = 0. \quad (1)$$

However, if a medial sheet with object angle θ passes through S ,

$$\lim_{\text{area}(S) \rightarrow 0} \frac{\iint_S \nabla D \cdot N_S dS}{\iint_S dS} = -\frac{1}{2} \sin(\theta). \quad (2)$$

We wish to apply the above continuous criterion to detect regions of space inside B that contain medial points. To do this, we voxelize the interior of B and circumscribe each voxel with a sphere S . We then sample m points \mathbf{p}_i on the surface of S . A discrete version of \mathcal{AOF} then is:

$$AOF = \frac{\sum_{i=1}^m \nabla D(\mathbf{p}_i) \cdot N_S(\mathbf{p}_i)}{m}. \quad (3)$$

Eqs. 1 and 2 require that the area of S shrinks to zero. When S is of non-zero size and B is the boundary of a polyhedron, [20] shows that if the contribution of certain sample points \mathbf{p}_i to the sum in Eq. 3 is ignored, then the quantity AOF is 0 if there are no medial points and proportional to the object angle of the medial points inside S otherwise.

2.2. Discrete Differential Geometry

The traditional approach to estimating curvatures on objects represented as triangle meshes is to perform computations directly on the boundary [17]. These methods may be divided into two classes: those that fit analytic functions to the data and those that work with the discrete data directly. In the former class, Cazals and Pouget [4] show that differential quantities evaluated for fitted polynomial surfaces of required degree converge to the true values given that a sampling condition on the boundary is met. In the latter class, Taubin [22] finds curvature tensors by considering estimates of normal curvature in a neighbourhood about each mesh vertex. The curvature values obtained from the curvature tensor in [6] approach the true values if a sampling condition on a particular kind of mesh is satisfied. Rusinkiewicz [18] performs this computation on a

per-triangle basis. The work of Taubin [22] is extended to 3D range data in [13].

The method that we develop for recovering boundary differential geometry in this paper is fundamentally different from the above techniques since it relies on medial geometry. In fact, to our knowledge there is very little computational work relating medial geometry to boundary geometry, with two exceptions. In [16] formulas are derived for the Gaussian and Mean curvatures for 3D boundaries based on derivatives along medial sheets, but this theory has not yet lead to implementations. For objects with non-branching medial topology, Yushkevich *et al.* [23] fit a single-sheet continuous medial representation (an *m-rep*) to medical image data and derive conditions to compute the implied boundary.

3. Dense and Precise Medial Points

As a starting point, we use the method of [20] to compute a set of voxels inside a polyhedral object Ω that are intersected by the medial surface. In this section, we show how these initial estimates can be refined to give a dense medial point cloud with a precision guarantee.

3.1. Precise Medial Point Location

We first describe a subroutine that uses information about the gradient of the Euclidean distance transform to locate points near the medial surface. Let (\mathbf{a}, \mathbf{b}) be a line segment connecting points \mathbf{a} and \mathbf{b} . The subroutine is based on the following lemma proved in [21]:

Lemma 1 *Let \mathbf{p} be a point in Ω . Let $\mathbf{q} = \mathbf{p} + \gamma \cdot \nabla D(\mathbf{p})$, such that \mathbf{q} is also in Ω and γ is a scalar. A medial point of Ω lies on (\mathbf{p}, \mathbf{q}) if and only if $\nabla D(\mathbf{p}) \neq \nabla D(\mathbf{q})$.*

Algorithm 1 RETRACT($\mathbf{p}, \mathbf{q}, \epsilon$)

Input: Points \mathbf{p}, \mathbf{q} inside Ω s.t. $\mathbf{q} = \mathbf{p} + \gamma \cdot \nabla D(\mathbf{p})$ and $\nabla D(\mathbf{p}) \neq \nabla D(\mathbf{q})$, tolerance ϵ .

Output: A point within ϵ of the medial surface of Ω and its two spoke vectors.

```

1: while  $d(\mathbf{p}, \mathbf{q}) > \epsilon$  do
2:    $\mathbf{m} = \frac{1}{2}(\mathbf{p} + \mathbf{q})$ 
3:   if  $\nabla D(\mathbf{m}) \neq \nabla D(\mathbf{p})$  then
4:      $\mathbf{q} = \mathbf{m}$ 
5:   else
6:      $\mathbf{p} = \mathbf{m}$ 
7:   end if
8: end while
9: Return  $\mathbf{p}, \nabla D(\mathbf{p}), \nabla D(\mathbf{q})$ .
```

By Lemma 1, Algorithm 1 finds a point on (\mathbf{p}, \mathbf{q}) that is within a user-specified tolerance ϵ of the medial surface




Figure 2. Left: The original polyhedra; Centre: the union of medial spheres; Right: medial points grouped into smooth sheets. The allowed tolerance between adjacent normals is 5.7 degrees.

MS. When \mathbf{p} and \mathbf{q} are points on a sphere S , these computations are fast because the search for closest boundary points is restricted to the set of closest polygons to S .

3.2. Dense, Salient Medial Points

In order to retain salient portions of the medial surface, we assign each medial point a measure of significance based on its object angle θ and also on the radius of its corresponding inscribed sphere. The object angle is a common simplification criterion used in Voronoi methods for medial surface construction, methods based on the grassfire flow, and also ones based on digital distance functions [19]. On

Model	# Triangles	Final Voxel Resolution	# Medial Points	Time (min)
Torus	16,000	3,033,136	64,195	5.93
Pear	86,016	5,177,526	18,147	22.24
Head	6,816	5,244,401	74,853	6.22
Venus	22,688	7,739,924	116,793	12.82

Table 1. Statistics for the computation of medial points.

the other hand, Chazal and Lieutier [5] prove that a simplification scheme based on radius preserves topology. Attali *et al.* [2] argue in favour of a simplification scheme based on both object angle and radius.

Our algorithm for detecting salient medial points is as follows. For each voxel u that is known to contain medial points, we attempt to find a salient medial point inside u . For each of the sample points \mathbf{p} on the circumscribing sphere S of u , we check to see if $\nabla D(\mathbf{p}) \neq \nabla D(\mathbf{q})$, where $\mathbf{q} \neq \mathbf{p}$ is a point of intersection of the ray $\mathbf{p} - \nabla D(\mathbf{p})$ with S . If so, Algorithm 1 finds an approximate location of a medial point \mathbf{x} along the segment (\mathbf{p}, \mathbf{q}) , as well as its 2 spoke vectors. Note that the spoke vector $\nabla D(\mathbf{p})$ is exact, while $\nabla D(\mathbf{q})$ is approximate. The object angle is then half the angle between these spoke vectors and the radius is their magnitude. Point \mathbf{x} returned by Algorithm RETRACT is stored as the medial point associated with voxel u if \mathbf{x} lies inside voxel u and if either the object angle or radius threshold is satisfied. At most one medial point is found per voxel. Thus, the density of the medial points is controlled by the resolution of the voxels considered. Deciding if a voxel contains a salient medial point is done quickly by using efficient point to polygon mesh distance tests and taking advantage of temporal coherence between queries. Hence, we can choose a fine resolution that produces a large number of voxels.

Figure 2 (Centre) illustrates the union of the maximal spheres for the computed medial points with object angle threshold 0.3 radians and radius threshold 0.25 times the maximum dimension of the bounding box. These examples show that the boundary of the union of spheres associated with the salient medial point cloud provides a close visual approximation to the boundary B of the original object.

Statistics for the computation of medial points for the models shown in this paper are given in Table 1. The final voxel resolution is the number of finest-resolution voxels in the interior of the model. The number of medial points is the number of salient medial points found, at most one per finest-resolution voxel. Timings are shown for a single 3.4GHz Pentium IV processor with 3GB of RAM. Note that our code can easily be parallelized to speed up computations.

3.3. Medial Surface Normals

Given a medial point \mathbf{x}_0 and its two spoke vectors $\mathbf{U}_{\mathbf{x}_0}^a$ and $\mathbf{U}_{\mathbf{x}_0}^b$, these vectors make an equal angle with the tangent plane to \mathcal{MS} at \mathbf{x}_0 , $T_{\mathbf{x}_0}\mathcal{MS}$. We use this fact to derive the equation for the normal $\mathbf{N}(\mathbf{x}_0)$ to $T_{\mathbf{x}_0}\mathcal{MS}$ as follows. For $\mathbf{x} \in T_{\mathbf{x}_0}\mathcal{MS}$,

$$\mathbf{U}^a(\mathbf{x}_0) \cdot (\mathbf{x}_0 - \mathbf{x}) = \mathbf{U}^b(\mathbf{x}_0) \cdot (\mathbf{x}_0 - \mathbf{x}) \quad (4)$$

$$\Rightarrow (\mathbf{U}^a(\mathbf{x}_0) - \mathbf{U}^b(\mathbf{x}_0)) \cdot (\mathbf{x}_0 - \mathbf{x}) = 0. \quad (5)$$

Therefore, $\mathbf{U}^a(\mathbf{x}_0) - \mathbf{U}^b(\mathbf{x}_0)$ is the normal direction to $T_{\mathbf{x}_0}\mathcal{MS}$. Figure 2 shows the polyhedra (Left) and their medial surfaces (Right) computed by our algorithm. We illustrate the quality of the normal estimates by grouping medial points that have locally consistent normals. Specifically, for a medial point \mathbf{q} on a sheet that is inside voxel $u_{\mathbf{q}}$, we add to this sheet all medial points \mathbf{p} inside voxels in the 26-neighbourhood of $u_{\mathbf{q}}$ whose normals are within an allowed tolerance of the normal at \mathbf{p} . The result is a segmentation of the medial surface into smooth sheets of A_1^2 points.

4. Boundary Geometry from Medial Geometry

The work of Damon [8] demonstrates that surface differential geometry can be directly inferred from the rate of change of the spoke vectors with respect to the medial surface. In this section, we survey some of the results in [8].

If $B \subset \mathbb{R}^3$ is the boundary of an object $\Omega \subset \mathbb{R}^3$ with unit outward normals \mathbf{n} , then the rate of change of \mathbf{n} along B describes the curvature of B . Rather than studying the rate of change of \mathbf{n} as one moves along B , consider the rate of change of \mathbf{n} as one moves along medial surface \mathcal{MS} of Ω . Recall that smooth (A_1^2) medial points are equidistant from exactly 2 points on B . The vectors from smooth medial points to nearest points on B , the spoke vectors, are normal to B . Thus, by studying the rate of change of the spoke vectors we obtain information about the rate of change of the normals to the boundary.

Consider a smooth medial point $\mathbf{x}_0 \in \mathcal{MS}$ and let $\mathbf{U}^a, \mathbf{U}^b$ be the 2 spoke vectors at \mathbf{x}_0 . Let $\mathbf{U}_1^a = \mathbf{U}^a / \|\mathbf{U}^a\|$, $\mathbf{U}_1^b = \mathbf{U}^b / \|\mathbf{U}^b\|$. Damon defines

$$S_{rad}^a(\mathbf{v}) = -\text{proj}_{\mathbf{U}^a}(\frac{\partial \mathbf{U}_1^a}{\partial \mathbf{v}}) \text{ and} \quad (6)$$

$$S_{rad}^b(\mathbf{v}) = -\text{proj}_{\mathbf{U}^b}(\frac{\partial \mathbf{U}_1^b}{\partial \mathbf{v}}), \quad (7)$$

where $\text{proj}_{\mathbf{U}}$ denotes projection into $T_{\mathbf{x}_0}\mathcal{MS}$ along \mathbf{U} [8]. Refer to Figure 3.

Choosing an orthonormal basis $\{\mathbf{v}_1, \mathbf{v}_2\}$ for $T_{\mathbf{x}_0}\mathcal{MS}$, let S_{rad}^a, S_{rad}^b denote the matrix representation of the two radial shape operators. Then, as in the case of the standard shape operator, the eigenvectors are principal radial




Figure 3. An illustration of $S_{rad}^a(\mathbf{v})$ for a given \mathbf{U}^a (adapted from [8]).

curvature directions and the eigenvalues are principal radial curvatures. Consider the correspondence between medial points and their closest boundary points given by $\psi^a : \mathcal{MS} \rightarrow B$, $\psi^b : \mathcal{MS} \rightarrow B$. Denote $\psi^a(\mathbf{x}_0)$ by \mathbf{x}_0^a and $\psi^b(\mathbf{x}_0)$ by \mathbf{x}_0^b . Then it can be shown ([8]):

Theorem 1 *There is a bijection between the principal curvatures κ_i^a of B at \mathbf{x}_0^a and the principal radial curvatures κ_{ri} of M at \mathbf{x}_0 given by*

$$\kappa_i^a = \frac{\kappa_{ri}}{1 - r\kappa_{ri}},$$

where $r = \|\mathbf{U}^a\|$. Furthermore, the principal radial curvature directions corresponding to κ_{ri} are mapped by $d\psi^a$ to the principal curvature directions corresponding to κ_i^a . The case of S_{rad}^b is symmetric.

Thus, given the radial shape operator and the spoke vectors at a smooth medial point, one can find principal curvatures and principal curvature directions on the object boundary.

5. Boundary Geometry from Sampled Medial Geometry

Working only with the set of medial points and spoke vectors sampled on salient portions of the medial surface, we now describe a numerical method to obtain boundary curvature estimates.

5.1. Imposing Smoothness on the Boundary

In our case, B is the boundary of a polyhedron and is piecewise planar. According to [8], if $r < \min_i \frac{1}{\kappa_{ri}}$ for all medial points, B is C^1 or better. Hence, if B is not smooth, for some medial points, $r \geq \min_i \frac{1}{\kappa_{ri}}$. The case $r > \min_i \frac{1}{\kappa_{ri}}$ happens when B locally self-intersects. The case $r = \min_i \frac{1}{\kappa_{ri}}$ occurs when B is not differentiable, *i.e.* along edges and on vertices. Thus, the computation of S_{rad} using the spoke vectors \mathbf{U}^a , \mathbf{U}^b which are the normals to a polyhedral boundary is not immediately valid. Figure 4 illustrates a scenario where two nearby medial points share the same nearest boundary point. This happens often for

concave regions of the polyhedron. It can be easily verified that in this case, indeed, $r = \frac{1}{\kappa_{ri}}$. We impose smoothness on the \mathbf{U}^a and \mathbf{U}^b vectors by performing a local Gaussian weighted average with neighbouring vectors. Specifically, the weights of the contribution of vectors at \mathbf{x}_1 to vectors at \mathbf{x}_0 are given by the Gaussian function $\frac{1}{(2\pi)^{3/2}\sigma} e^{-\frac{\|\mathbf{x}_0 - \mathbf{x}_1\|_2^2}{2\sigma^2}}$, with $\sigma = 0.25$. Note that this corresponds to smoothing the boundary reconstructed from the medial points.




Figure 4. The case when $\kappa_{ri} = -\text{proj}_{\mathbf{U}} \frac{\partial \mathbf{U}_1}{\partial \mathbf{v}_1} = \frac{1}{r}$.

5.2. Derivatives on Medial Sheets

Given a normal \mathbf{N} to a medial point \mathbf{x} , we now need to measure the rate of change of the spoke vectors \mathbf{U}^a and \mathbf{U}^b in the directions given by the basis $\{\mathbf{v}_1, \mathbf{v}_2\}$ to the tangent plane at \mathbf{x} , $T_{\mathbf{x}}\mathcal{MS}$. We disambiguate between \mathbf{U}^a and \mathbf{U}^b directions using the consistent normal estimates of [9]. Given a particular step size Δ for the derivatives, a medial point \mathbf{x}_1 that is approximately distance Δ from \mathbf{x} gives the direction \mathbf{v}_1 . In order to ensure that \mathbf{x}_1 lies on the same medial sheet as \mathbf{x} , we check that the normals to \mathbf{x}_1 and \mathbf{x} are similar. For the $\mathbf{v}_2 = \mathbf{N} \times \mathbf{v}_1$ direction (recall that $\mathbf{v}_1, \mathbf{v}_2$ must give an orthonormal basis for the tangent plane at \mathbf{x}) we cannot hope to find a sampled point which lies exactly at $\mathbf{x}_2 = \mathbf{x} + \Delta \mathbf{v}_2$, because we are working with a discrete set of medial points. We therefore estimate \mathbf{U}^a and \mathbf{U}^b values at \mathbf{x}_2 as a Gaussian weighted average of the spoke vectors of the medial point samples in the vicinity of \mathbf{x}_2 . Once again, to restrict computations to the same medial sheet, only those points in the vicinity that have similar normals to \mathbf{N} are considered.

5.3. Computing Surface Curvature

Given $\frac{\partial \mathbf{U}_1^a}{\partial \mathbf{v}}$, $\frac{\partial \mathbf{U}_1^b}{\partial \mathbf{v}}$, where $\mathbf{v} \in \{\mathbf{v}_1, \mathbf{v}_2\}$, we now compute the S_{rad} operator at medial points and use its eigenvalues to find the principal curvatures $\kappa_i^a, \kappa_i^b, i = 1, 2$ at the 2 nearest points on the boundary B to the medial points.

In recovering the curvature of the original boundary B , we are presented with the obstacle that the set of nearest boundary points to the set of medial points we have computed does not include points on certain regions of B . This happens because we retain salient medial points only. Hence, the geometry described by the medial points is that of the boundary of the union of our medial spheres, rather

than that of the original boundary B . For medial points that contribute a large spherical patch to the reconstruction of the shape, we only have curvature information about its 2 nearest boundary points (refer to Figure 5 for a 2D illustration). These points correspond to A_3 points in Giblin and Kimia’s classification[12], or ‘edge points’ (see Figure 1 (Left)). Surface curvature values at missing locations are computed by propagating known surface curvature values of neighbouring surface points along the boundary B .




Figure 5. The medial surface of this object consists of low-object angle segments (dashed lines) and a high-object angle segment (bold line). When approximating the medial surface with a set of points, we only retain points on high-object angle segments of the medial surface. The boundary of the union of the associated medial circles approximates the original object. Medial point m is equidistant from points A and B on the boundary. Surface curvatures at points A and B are found using the radial shape operator, while for point C a different strategy is used.

6. Experimental Results

We present numerical results on polyhedral objects of varying topology and surface complexity, and with varying medial surface branching topology: a three-hole torus, a pear, a head, and a Venus model.¹ To show surface curvatures, we use the colormap in Figure 6 (Bottom Right). Here, red corresponds to a convexity, blue to a concavity, green to a saddle-shaped region, yellow to a cylindrical patch curving toward the object (the non-zero principal curvature is positive), cyan to a cylindrical patch curving away from the object (the non-zero principal curvature is negative) and white to a flat region. Figure 6 (Left) shows the surface curvature estimates on the medial surface and Figure 6 (Middle) shows these estimates projected onto the boundary B . For comparison, Figure 6 (Right) shows the result of applying Rusinkiewicz’s method [18] directly on B .

Although there are subtle numerical differences, the results obtained by the two methods are qualitatively very consistent. We emphasize that the implementation of our




Figure 7. The principal curvature directions recovered on the surface of a cup using only the medial geometry.

method works at a fixed spatial resolution on the set of medial points, whereas the method in [18] is aided by the quality of the polygonal mesh. The majority of surface regions recovered by our method are correctly coloured. Examples include the holes and the sides of the tori; the stem, neck and base of the pear; the eye sockets, neck, chin and nose of the head; and the neck, shoulders, breasts, chest, thighs, and navel of the Venus model.

In order to demonstrate the validity of the principal curvature direction estimates from medial geometry, Figure 7 illustrates these on a model of a cup. As explained in Section 4, these estimates are obtained from the eigenvectors of S_{rad} . Note how the directions are orthogonal and correspond to the directions of maximal bending (red) and minimal bending (blue), as one would expect.

There are certain challenges faced in the implementation of our method, including the difficulty of performing computations near edges of medial sheets (A_3 curves), where the tangent plane cannot be accurately estimated, and on narrow sheets. Numerical errors can be mitigated, at least in part, by increasing the number of medial points (the voxel resolution) at the cost of increased computation time. We also note that the union of medial spheres, whose geometry the medial points describe, is a slightly different object than the original polyhedral object, as object angle simplification removes some small-scale boundary details. This fact should be taken into account when comparing the two different curvature estimates in Figure 6.

Since ground truth curvatures are not available for such complicated objects as those shown in Figure 6, it is difficult to numerically evaluate the quality of our curvature estimates. However, for the simple example of a portion of a cylindrical cup wall, such that one side of the wall has positive and zero principal curvatures, while the other side has negative and zero principal curvatures, such a numerical comparison is feasible. Based on estimates of exterior and interior radii of the cup wall, we obtain an estimates for the positive and negative principal curvatures, respectively. The distribution of the curvatures estimated with our method and that of [18] is shown in Table 2. We observe

¹Models from the Princeton Shape Benchmark, <http://vcg.isti.cnr.it/polycubemaps/models/>, and <http://www.cs.princeton.edu/gfx/proj/sugcon/models/>




Figure 6. Left: The medial surface coloured according to principal curvature on the object boundary. Centre: Projection of curvature values on the medial surface to the boundary of the object. Right: Curvature values obtained using the method in [18]. The colourmap used is shown in the bottom right corner. See the associated text for a discussion of these results.

	Positive	Negative	Zero
μ	0.069834	-0.087302	1.205×10^{-6}
σ	0.000101	0.000148	8.707×10^{-6}
μ	0.068380	-0.087144	-8.973×10^{-7}
σ	0.003784	0.002784	1.428×10^{-5}

Table 2. Mean (μ) and standard deviation (σ) of the 3 different types of curvature on a section of a cylindrical cup wall (shown in Figure 7) consisting of 551 vertices obtained using the method of [18] (Top) and our method (Bottom). We estimate the true positive curvature to be approximately 0.069898 and the true negative curvature to be approximately -0.0875045 .

that our relative error for the positive curvature is 2.2% and 0.4% for the negative curvature.

7. Conclusions

We have described a method to compute a dense sample of salient medial surface points, along with their spoke vectors. This discrete point-based approximation to the medial surface is shown to be sufficient to recover qualitatively consistent surface curvatures on the object boundary. As such, this discrete representation is a rich shape descriptor for applications in computer vision, computer graphics and related domains.

Acknowledgements

This research was funded by NSERC Canada and FQRNT Québec. We thank the reviewers for their helpful comments. We are grateful to Szymon Rusinkiewicz for making the software package *trimesh2* publically available. We thank Tamal Dey for providing us with the *NormFet* software. We thank Emmanuel Piuzé-Phaneuf for assistance in generating Figure 7.

References

- [1] N. Amenta, S. Choi, and R. Kolluri. The Power Crust, Unions of Balls, and the Medial Axis Transform. *Computational Geometry: Theory and Applications*, 19(2-3):127–153, 2001.
- [2] D. Attali, J.-D. Boissonnat, and H. Edelsbrunner. Stability and computation of the medial axis — a state-of-the-art report. In *Mathematical Foundations of Scientific Visualization, Computer Graphics, and Massive Data Exploration*. Springer-Verlag, 2007.
- [3] H. Blum. Biological Shape and Visual Science. *Journal of Theoretical Biology*, 38:205–287, 1973.
- [4] F. Cazals and M. Pouget. Estimating differential quantities using polynomial fitting of osculating jets. In *Symposium on Geometry processing*, pages 177–187, 2003.
- [5] F. Chazal and A. Lieutier. The λ -medial axis. *Graphical Models*, 67(4):304–331, 2005.
- [6] D. Cohen-Steiner and J.-M. Morvan. Restricted delaunay triangulations and normal cycle. In *Symposium on Computational geometry*, pages 312–321, 2003.
- [7] T. Culver, J. Keyser, and D. Manocha. Exact computation of the medial axis of a polyhedron. *Computer Aided Geometric Design*, 21(1):65–98, 2004.
- [8] J. Damon. Determining the geometry of boundaries of objects from medial data. *International Journal of Computer Vision*, 63(1):45–64, 2005.
- [9] T. K. Dey and J. Sun. Normal and feature approximations from noisy point clouds. In *Foundations of Software Technology and Theoretical Computer Science*, pages 21–32, 2006.
- [10] M. Etzion and A. Rappoport. Computing voronoi skeletons of a 3-d polyhedron by space subdivision. *Computational Geometry: Theory and Applications*, 21:87–120, 2002.
- [11] M. Foskey, M. C. Lin, and D. Manocha. Efficient computation of a simplified medial axis. In *Solid modeling and applications*, pages 96–107, 2003.
- [12] P. J. Giblin and B. B. Kimia. A formal classification of 3D medial axis points and their local geometry. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 26(2):238–251, February 2004.
- [13] E. Hameiri and I. Shimshoni. Estimating the principal curvatures and the darboux frame from real 3d range data. In *3D Data Processing, Visualization, and Transmission*, pages 258–267. IEEE Computer Society, 2002.
- [14] F. F. Leymarie and B. B. Kimia. The medial scaffold of 3d unorganized point clouds. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 29(2):313–330, 2007.
- [15] A. Lieutier. Any open bounded subset of \mathbb{R}^n has the same homotopy type as its medial axis. *Computer-Aided Design*, 36(11):1029–1046, 2004.
- [16] L. R. Nackman and S. Pizer. Three dimensional shape description using the symmetric axis transform I: Theory. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 7(2):187–202, 1985.
- [17] S. Petitjean. A survey of methods for recovering quadrics in triangle meshes. *ACM Computing Surveys*, 34(2):211–262, 2002.
- [18] S. Rusinkiewicz. Estimating curvatures and their derivatives on triangle meshes. In *3D Data Processing, Visualization, and Transmission*, pages 486–493, 2004.
- [19] K. Siddiqi and S. Pizer, editors. *Medial representations: mathematics, algorithms and applications*. Springer, 2008.
- [20] S. Stolpner and K. Siddiqi. Revealing significant medial structure in polyhedral meshes. In *3D Data Processing, Visualization, and Transmission*, pages 365–372, 2006.
- [21] S. Stolpner and S. Whitesides. Medial axis approximation with bounded error. In *International Symposium on Voronoi Diagrams*, pages 171–180, 2009.
- [22] G. Taubin. Estimating the tensor of curvature of a surface from a polyhedral approximation. In *International Conference on Computer Vision*, page 902, 1995.
- [23] P. A. Yushkevich, H. Zhang, and J. C. Gee. Continuous medial representation for anatomical structures. *Transactions on Medical Imaging*, 25(12):1547–64, 2006.