Medial Residues of Piecewise Linear Manifolds

Erin W. Chambers* Tao Ju† David Letscher‡

Abstract

Skeleton structures of objects are used in a wide variety of applications such as shape analysis and path planning. One of the most widely used skeletons is the medial axis, which is a thin structure centered within and homotopy equivalent to the object. However, on piecewise linear surfaces, which are one of the most common outputs from surface reconstruction algorithms, natural generalizations of typical medial axis definitions may fail to have these desirable properties. In this paper, we propose a new extension of the medial axis, called the medial residue, and prove that it is a finite curve network homotopy equivalent to the original surface when the input is a piecewise linear surface with boundary. We also develop an efficient algorithm to compute the medial residue on a triangulated mesh, building on previously known work to compute geodesic distances.

1 Introduction

The medial axis of an object is a skeletal structure originally defined by Blum [1]. It is the set of points having more than one closest points (under the Euclidean distance metric) on the boundary of the object. The medial axis is centered within the object, homology equivalent to the object if it is an open bounded subset of $\mathbb{R}^n$ [11], and (at least) one dimension lower than that of the object. These properties make the medial axis ideal for many applications including shape analysis and robotic path planning.

We are interested in defining a similar skeletal structure on a surface S (with boundary) that inherits the properties of the medial axis. Such a structure could then be used for applications such as shape analysis of surface patches as well as path planning in non-planar domains. We are particularly interested in the case when S is piecewise smooth, which is more representative of typical outputs of discrete surface reconstruction algorithms (e.g., triangulated meshes) than globally smooth surfaces.

A natural approach would be to replace the Euclidean distances in the medial axis definition by geodesic distances over S [18]. Interestingly, as we will show in this paper (Section 3), several equivalent definitions of the medial axis may yield different structures when S is only piecewise smooth, and none of these definitions guarantees the two essential properties of the medial axis, namely being homotopy equivalent to the original surface and codimension one.

In this paper, we propose a new extension of the medial axis onto a piecewise linear surface S with boundary, which we call the *medial residue* (Section 4), and prove that the structure is a finite curve network that is always homotopy equivalent to S (Section 5). We also describe a

*Department of Mathematics and Computer Science, Saint Louis University, echambe5@slu.edu. Research supported in part by NSF grant (CCF 1054779).

†Department of Computer Science and Engineering, Washington University in St. Louis, taoju@cse.wustl.edu. Research supported in part by NSF grant (IIS 0846072).

‡Department of Mathematics and Computer Science, Saint Louis University, letscher@slu.edu.

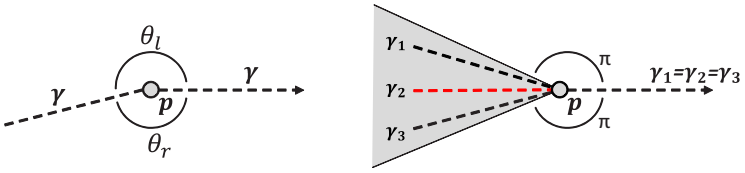

Figure 1: Left: Illustration of left and right curve angles. Right: at a concave vertex p, there may be infinitely many geodesic paths to the boundary (such as $\gamma_1, \gamma_2, \gamma_3$) sharing a common outgoing direction, but only one of them (γ_2) can be straight. Shaded region is the shadow rooted at p, made up of points whose shortest paths to the boundary go through p.

quadratic-time algorithm to compute this structure on a piecewise flat surface with boundary embedded in Euclidean space (Section 6).

2 Background and Definitions

We assume the reader is familiar with classical definitions of manifold topology, which can be found in books such as [8, 14, 3]. We shall only review definitions that are specifically relevant to our work.

A *piecewise linear surface* is a 2-manifold (with boundary) with a piecewise linear structure, whose presentation consists of a finite number of triangles glued together along with an intrinsic distance metric on each triangle that is a linear map. Our algorithmic results work in a more restricted class of *piecewise flat surface*, where the piecewise linear structure comes from an embedding of a triangulation of M into $\mathbb{R}^3$, so that each triangle will be isometric to a triangle in $\mathbb{R}^2$.

Given a vertex v of a piecewise linear surface which is contained in more than two triangles, let $\{f_1, f_2, \ldots, f_k\}$ be the faces to which v belongs, where $\theta_i(v)$ is the interior angle of f_i at vertex v. The *total angle* is the sum of all of these angles, $\theta(p) = \sum_i \theta_I(v)$. The *curvature* at v is the value $2\pi - \theta(p)$. A vertex is said to be *convex*, *flat* or *concave* if its curvature is positive, zero or negative.

A *curve* (or path) is (the image of) a map $p : [0, 1] \to M$; the length of the curve is generally the length of the image in M. A curve is a *geodesic* if it is locally shortest; in other words, no perturbation of the curve will result in a shorter curve. On a piecewise linear surface, geodesics and shortest paths are themselves piecewise linear maps. We say a curve γ *bisects* a piecewise differentiable curve X at time t if $\gamma(t) \in X$ and the two angles bounded by γ and the tangent of X at $\gamma(t)$ are equal. The *curve angles* θ_l and θ_r of a point p on a piecewise linear curve γ are the two angles to the left and right of the curve at p, where $\theta_l + \theta_r$ is the total vertex angle at that point p (see Figure 1 left).

A curve γ is considered *straight* if for each point $p \in \gamma$, the left and right curve angles are equal. This definition was introduced by Polthier and Schmies [15]. It is worth noting that Polthier and Schmies used the term "straight geodesic", and not simply straight. However, their straight geodesics might in fact not be geodesic (for example, it can go through a convex vertex). In this paper, the term straight geodesic will be used to denote a curve that is both straight and geodesic. Note that although there may be infinitely many geodesic paths to the boundary that go through a concave vertex p, only one of them is straight (see Figure 1 right). We call the region made up of points whose shortest paths to the boundary go through p the *shadow* rooted at p (shaded region in Figure 1 right).

3 The Medial Axis

Let X be a shape in Euclidean space. There are a variety of equivalent ways in which the medial axis of X could be defined. We will consider the following three:

1. Most commonly, the medial axis is defined as the set of points *without a unique closest point* on the boundary of the shape: $\mathcal{M}^{CP} = \{x \in X \mid \nexists \text{ unique } y \in \partial X \text{ with } d(x, \partial X) = d(x, y)\}$
2. Alternatively, the medial axis is the set of points *without a unique shortest path* to the boundary of the shape: $\mathcal{M}^{SP} = \{x \in X \mid \exists \text{ shortest paths } \gamma_1 \neq \gamma_2 \text{ from } x \text{ to } \partial X\}$
3. The medial axis is also the set of points *without a unique direction for shortest paths* to the boundary of the shape. We say two paths γ_1 and γ_2 with $\gamma_1(0) = \gamma_2(0)$ start in the same direction if there exists some $\epsilon > 0$ such that for all $t < \epsilon$, $\gamma_1(t) = \gamma_2(t)$ (or the curves can be reparameterized so that this holds): $\mathcal{M}^{SPD} = \{x \in X \mid \exists \text{ shortest paths } \gamma_1, \gamma_2 \text{ from } x \text{ to } \partial X$ that do not start in the same direction$\}$

The above definitions are all equivalent when X is a smooth manifold in any dimension:

Lemma 3.1. *If X is a smooth manifold then $M^{CP} = M^{SP} = M^{SPD}$.*

Proof. We first consider $\mathcal{M}^{SP}$ and $\mathcal{M}^{SPD}$. On a smooth manifold, given any point and any direction, there is a unique geodesic in that direction [3]. Therefore, no point can have two geodesics that begin in the same direction, so $\mathcal{M}^{SP} = \mathcal{M}^{SPD}$.

Now consider $\mathcal{M}^{CP}$. We know that $\mathcal{M}^{CP} \subseteq \mathcal{M}^{SP}$, since any point with two distinct closest points on the boundary must have two distinct shortest paths to the boundary. To show that $\mathcal{M}^{SP} \subseteq \mathcal{M}^{CP}$, we consider a point $x \in \mathcal{M}^{SP}$. If $x \notin \mathcal{M}^{CP}$, then it must have two shortest paths to the same point closest on the boundary. But since we are on a smooth manifold (so the boundary is also smooth), both curves must meet ∂X perpendicularly, which means they must be the same curve, giving a contradiction. □

However, when X is piecewise smooth, these three definitions yield different structures. More precisely, if X is any path metric space (where distances are realized by shortest paths), then we have the following relationship:

Lemma 3.2. *If X is any path metric space (where distances are realized by shortest paths), then $M^{CP} \subset M^{SP}$ and $M^{SPD} \subset M^{SP}$.*

Proof. The first inclusion follows from the fact that any point with at least two different closest points on the boundary will necessarily have at least two distinct paths to the boundary. The second inclusion follows because any point with at least two shortest paths that begin in different directions also have two different shortest paths to the boundary. □

More importantly, however, we will demonstrate that there are situations where none of the three definitions satisfy the desired properties of being one dimension lower than and homotopy equivalent to X.

First, consider the heart-shaped surface in Figure 2, which has an interior hole on top of a cylindrical protrusion. Note that $\mathcal{M}^{CP}$ excludes points like x in the picture, which has a single closest point q on the boundary (a C^0 corner point) but two shortest paths to q that go around the cylinder. As a result, $\mathcal{M}^{CP}$ consists of two disconnected components. On the other hand, x is included in $\mathcal{M}^{SP}$ and $\mathcal{M}^{SPD}$.

Next, consider the oval-shaped surface in Figure 3 (a). The surface has a concave vertex v with a large negative curvature that happens to have two shortest paths to two distinct boundary points (a non-generic situation). Since each point in the shadow rooted at v (shaded region in (b)) would have two distinct shortest paths to the boundary, both $\mathcal{M}^{CP}$ and $\mathcal{M}^{SP}$ include the

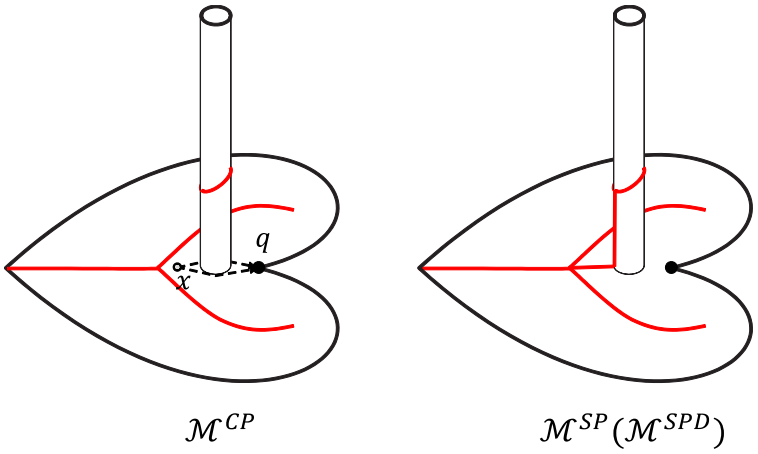

Figure 2: An example where $\mathcal{M}^{CP}$ (red) is not homotopy equivalent to the surface, but $\mathcal{M}^{SP}$ and $\mathcal{M}^{SPD}$ are.

2-dimensional shadow region. On the other hand, since any point in the shadow has a unique shortest path direction (that follows the geodesic to v), the entire shadow is excluded in $\mathcal{M}^{SPD}$, and $\mathcal{M}^{SPD}$ has an isolated vertex v that is disconnected from the rest of $\mathcal{M}^{SPD}$.

4 The Medial Residue

We now define our structure, called the medial residue, which is equivalent to existing definitions of the medial axis on a smooth manifold but possesses the desired properties of homotopy equivalence and co-dimension one on a piecewise linear surface. To make it clear that we are considering surfaces and not arbitrary manifolds from now on, we will use S instead of X to represent the shape.

We note that our medial residue is well defined on piecewise smooth manifolds, and that the majority of our results hold in these settings. However, our proof about homotopy and dimension holds only for piecewise linear surfaces, although we conjecture that the properties hold in more general settings as well.

The starting point of our definition is $\mathcal{M}^{SPD}$, which is more complete than $\mathcal{M}^{CP}$ in our first example (Figure 2) and remains low dimension in the second example (Figure 3). Our goal is to add low-dimensional components to $\mathcal{M}^{SPD}$ to restore the homotopy equivalence. Observe that, in our second example, the disconnection in $\mathcal{M}^{SPD}$ takes place in the shadow rooted at a concave vertex $v \in \mathcal{M}^{SPD}$, where the shortest paths from a point x in the shadow to the boundary would agree for some time and then diverge at v. Since we cannot include the entire shadow, which is 2-dimensional, we wish to keep one representative curve. A natural choice of such curve would be one that is "centered" with respect to the two diverging shortest paths at v. More precisely,

Definition 4.1. *The* medial residue*, $\mathcal{MR}$ consists of any point $x \in S$ such that either $x \in \mathcal{M}^{SPD}$ or where there are two distinct shortest paths from x to the boundary, γ_1 and γ_2, parameterized by arc length, which first intersect $\mathcal{M}^{SPD}$ at $v = \gamma_1(t) = \gamma_2(t)$ such that $\gamma = \gamma_1([0,t]) = \gamma_2([0,t])$ is straight and bisects the angle between the tangents of the two shortest paths from v to the boundary that are nearest to γ on its left and right side.*

The definition for a point $x \in \mathcal{MR} \setminus \mathcal{M}^{SPD}$ is illustrated in Figure 4 (a). Note that, by definition, every point on the common segment γ of the shortest paths from x to the boundary

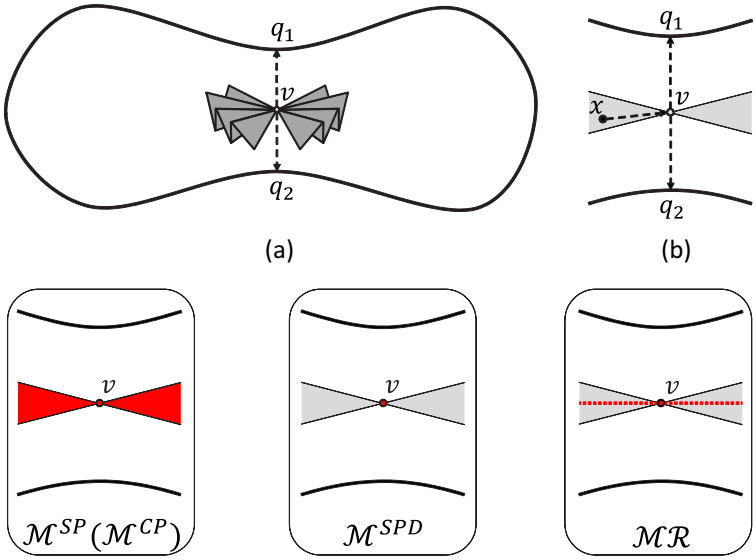

Figure 3: Top: a surface with a highly concave central vertex (a) and a zoom-in view (b). Bottom: different medial axis extensions (red): $\mathcal{M}^{CP}$ and $\mathcal{M}^{SP}$ are 2-dimensional, $\mathcal{M}^{SPD}$ has an isolated vertex, and $\mathcal{MR}$ is 1-dimensional and homotopy equivalent to the surface.

is also included in $\mathcal{MR} \setminus \mathcal{M}^{SPD}$. In fact, $\mathcal{MR} \setminus \mathcal{M}^{SPD}$ consists of straight geodesics that bisect shortest path directions at concave vertices of $\mathcal{M}^{SPD}$. Figure 4 (b) gives a generic picture of $\mathcal{MR}$ at a concave vertex of $\mathcal{M}^{SPD}$. The multiple shortest path directions divide the local neighborhood of the vertex radially into *sectors*. Each sector is bisected either by a curve in $\mathcal{M}^{SPD}$ (solid red line), if the sector's interior angle is less than 2π, or otherwise by a curve in $\mathcal{MR} \setminus \mathcal{M}^{SPD}$ (dotted red line).

Since any point in $\mathcal{MR} \setminus \mathcal{M}^{SPD}$ has two distinct shortest paths, we have $\mathcal{M}^{SPD} \subset \mathcal{MR} \subset \mathcal{M}^{SP}$. Since both $\mathcal{M}^{SPD}$ and $\mathcal{M}^{SP}$ are equivalent when S is a smooth manifold, this implies that our medial residue is also equivalent to the other definitions we mentioned earlier in a smooth manifold.

5 Medial Residue on Piecewise Linear Surfaces

In this section, we state and prove the central result in the paper, that the medial residue is a finite graph and is homotopy equivalent to the original surface. As previously mentioned, while we only prove this for piecewise linear surfaces (the main focus of this work), we conjecture that it also holds for piecewise smooth surfaces and higher dimensional manifolds as well.

Theorem 5.1. *If S is a piecewise linear surface with boundary then the medial residue of S is a finite graph that is a deformation retract of S.*

To prove this theorem, we will construct a deformation retract by incrementally "eroding" from the boundary, stopping at potentially interesting points along the way. To begin this process, we must understand what a neighborhood of the boundary of S looks like. Let $S_t = \{x \in S \mid d(x, \partial S) \geq t\}$; in other words, S_t is the set of points whose distance from the boundary of S is no less than t. The boundary of S_t is precisely the points with distance t to ∂S.

Our first step is to prove that shortest paths to the boundary are of finite complexity; in other words, they cannot cross the triangulation an unbounded number of times. If our PL

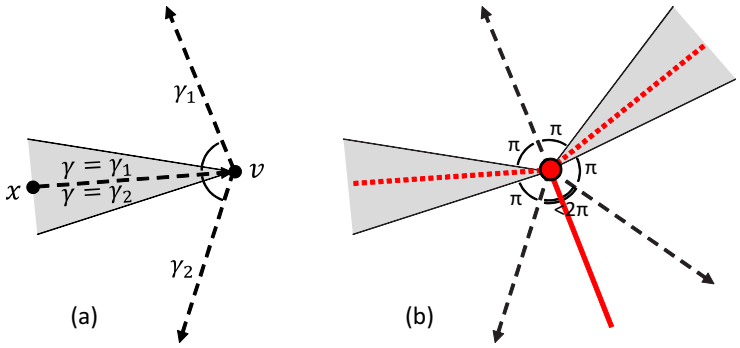

Figure 4: (a): illustration for the definition of a point $x \in \mathcal{MR} \setminus \mathcal{M}^{SPD}$. (b): a generic picture of $\mathcal{MR}$ at a concave vertex with multiple shortest path directions (solid line is $\mathcal{M}^{SPD}$ and dotted lines are $\mathcal{MR} \setminus \mathcal{M}^{SPD}$).

surface is a flat embedding in $\mathbb{R}^3$, this will follow easily since edges in the triangulation are shortest paths (and no two shortest paths can cross twice), but we must also have a bound for arbitrary PL surfaces. We note that variants of the following proof have been used in the normal surface community for at least 20 years; however, we are unaware of any published reference for such a bound on the number of possible intersections, so we have included it for completeness.

Proposition 5.2. *The number of intersections between any shortest paths and the underlying triangulation of an arbitrary PL-surface is* $\leq |E| \cdot \frac{2\pi}{\delta} \cdot \max_{e \in E} \frac{l(e)}{c_e}$, *where* δ *is the minimum angle at any vertex of the triangulation,* $l(e)$ *is the length of the edge* e, *and* $c(e)$ *is the minimum distance between any pair of points on opposite edges of the quadrilateral formed by the two faces adjacent to an edge* e.

Proof. Consider any geodesic path γ. If we unfold the PL triangles that γ intersects into the plane, γ must unfold to a straight line since it is a geodesic. Therefore, γ cannot wind around a single vertex more than π radians, since such a path could not unfold to a straight line. Therefore, γ must cross some diagonal after at most $2\pi/\delta$ edge crossings, where δ is the minimum angle in any of the triangles of our PL surface.

Now consider a shortest path α. We will divide α up into diagonal crossings and crossings which wind around a vertex (similar to the streets, ports, and junctions in [5]). Since a shortest path is a geodesic, we already know that there are at most $2\pi/\delta$ crossings between each diagonal crossing; it remains only to bound the total number of diagonal crossings that α can have.

Fix an edge e. The path α can cross e diagonally many times. However, each such crossing adds a length $\leq c_e$ to the total length of α, since it must cross opposite edges of the rectangle formed by unfolding the triangles adjacent to e. Since α is a shortest path, this means that each such crossing must be at least c_e apart on e, or else we could form a shorter path by instead following along e. So the shortest path α can cross e at most $l(e)/c_e$ times.

Putting this together, we have at most $l(e)/c_e$ diagonal crossings of each edge, and at most $2\pi/\delta$ crossings between each diagonal crossing, yielding the desired bound. □

Next, we want to understand what the boundary of the surface looks like at each stage of the erosion process. Locally, ∂S_t consists of a union of straight edges and circular arcs. The straight edges correspond to points whose shortest paths to the boundary do not pass through a vertex of the triangulation, and the circular arcs to points whose shortest paths pass through a vertex. The previous proposition can be used to show that there are finitely many arcs and lines segments in ∂S_t:

Lemma 5.3. *Given a piecewise linear S, for all but finitely many values of t, ∂S_t is a curve. In the cases where ∂S_t is not a curve, ∂S_t is a graph.*

Proof. Consider a shortest path from a point x to the boundary of S where $d(x, \partial S) = t$. If the shortest path does not pass through a vertex, then there is some segment of ∂S_t around x that is a straight line running "parallel" to an edge of ∂S; to see this, simply unfold the triangles that the shortest paths crosses into a planar strip. If the shortest path passes through a vertex then, near x, there is a circular arc contained in ∂S_t. This gives a piecewise structure to ∂S_t since it is built from straight line segments and circular arcs. Each portion corresponds to a combinatorial type of shortest path to the boundary. The previous proposition implies that there are finitely many segments and arcs, directly implying that ∂S_t is a graph.

The straight lines correspond to solutions to equations of the form $d(x, L) = s$ for some line L, and the circular arcs correspond to equations of the form $d(x, v) + d(v, \partial S) = s$ for some vertex v of the triangulation. The vertices of the graph correspond to the intersection of three or more combinatorial types of shortest paths. The intersection of any such family consists of a unique point, which means that there are only finitely many possible vertices in the graph. Thus there are finitely many values of t where S_t is not a curve. Note these graph vertices have multiple directions for shortest paths to the boundary, so they lie in $\mathcal{M}^{SPD}$. □

Notice that $\mathcal{M}^{SPD}$ consists of points that have multiple shortest paths directions to the boundary. The above results allow us to bound the combinatorial types of these shortest paths. The points equidistant from the boundary in each shortest path direction are built locally from lines segments and circular arcs. So in a small neighborhood $\mathcal{M}^{SPD}$ consists of the intersection of two curves that are either lines or circles. Hence, $\mathcal{M}^{SPD}$ is built from lines, circles and parabolas. This leads to the following result:

Lemma 5.4. *If S is piecewise linear, then $\mathcal{M}^{SPD}$ is a finite graph.*

Proof. As observed in the proof of the previous lemma, there are finitely many combinatorial types of shortest paths to the boundary. $\mathcal{M}^{SPD}$ is built from intersections between multiple types of curves. The vertices come from the intersection of three or more as seen in the previous lemma. The rest come from the intersection of two curve types, which are at worst quadratic curves. □

Now we are ready to describe the deformation retract, which immediately implies that $\mathcal{MR}$ is homotopy equivalent to the original PL surface. We will build our deformation retract based on an erosion process which intuitively "pauses" at times $\{t_1, \ldots, t_k\}$, where each t_i corresponds to one or more of these three possibilities:

1. There is a vertex v of the triangulation of S with $d(v, \partial S) = t_i$.
2. ∂S_{t_i} is not a disjoint union of curves but instead forms a graph.
3. There is a vertex v of $\mathcal{M}^{SPD}$ with $d(v, \partial S) = t_i$.

The previous lemmas imply that the set of t_i's is finite. We will consider the sets S_{t_i} based on our level sets at times $\{t_1, \ldots, t_k\}$ described above, as well as the "slice" between two of our level sets, $C_i = \overline{(S_{t_i} \setminus S_{t_{i+1}})}$. The following lemma actually shows how we can construct the deformation retract.

Lemma 5.5. *For each t_i, $S_{t_{i+1}} \cup \mathcal{MR}$ is a deformation retract of $S_{t_i} \cup \mathcal{MR}$.*

Proof. Consider the slice region C_i between two level curves. One of several cases could occur depending on what happens on the boundaries of this region, as illustrated in Figure 5.

The first case is that portions of the boundary ∂S_{t_i} meet at a convex corner. This is shown in Figure 5(a), where the shortest paths are shown on the left and the deformation retract on

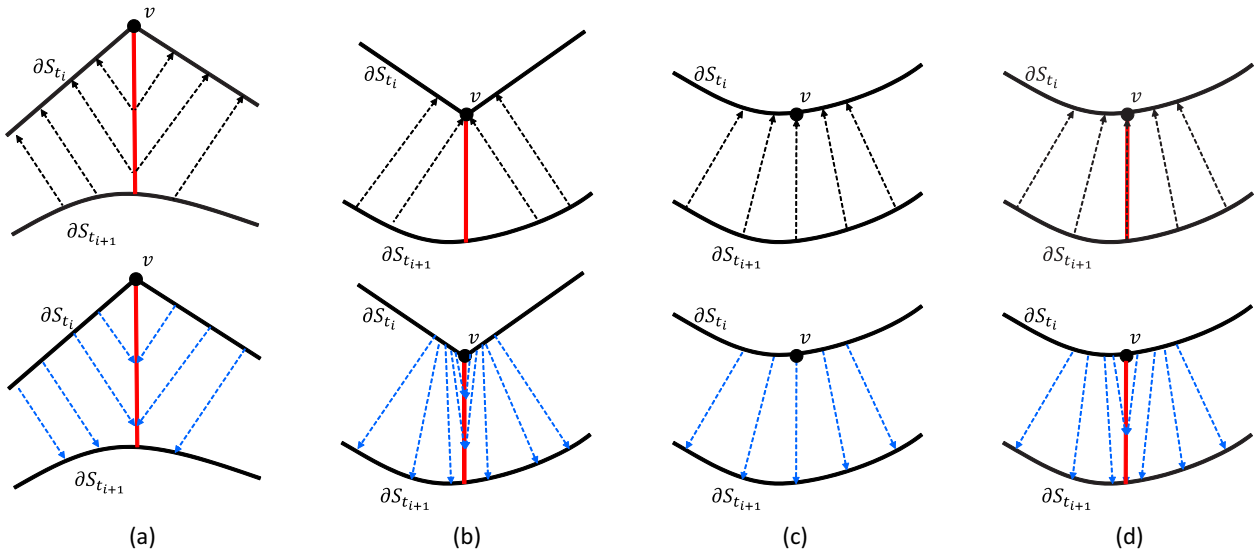

Figure 5: The shortest paths (black arrows), medial residue (red lines) and deformation retract (blue arrows) at the points in the slice region.

the right. At such a corner point v, there is a segment of $\mathcal{M}^{SPD}$ going from ∂S_{t_i} to $\partial S_{t_{i+1}}$, which bisects the convex corner at v. Shortest paths from points on this segment hit ∂S_{t_i} near v. The deformation retract follows these curves.

The second case is that portions of the boundary ∂S_{t_i} meet at a concave corner, see Figure 5(b). Note that the concave corner v must contain a shadow rooted at v where there is a cone of shortest paths going through v. By definition of $\mathcal{MR}$, if $v \in \mathcal{MR}$ then the bisector of the shadow will be in $\mathcal{MR} \setminus \mathcal{M}^{SPD}$. However, the deformation retract cannot simply follow the shortest paths exactly, as this would not be continuous at v; observe in the figure that points near v are taken to opposite sides of the bisector and do not move continuously. Instead, very near this point, the deformation retract will take points to either the bisector or the full shadow, as shown on the right in Figure 5(b). Note that the reparameterization continuously deforms points from ∂S_{t_i} onto the union of $\partial S_{t_{i+1}}$ and $\mathcal{MR}$.

In the third case, consider points $v \in \partial S_{t_i}$ where the ∂S_{t_i} is smooth. A single shortest path passes through v. If $v \notin \mathcal{MR}$, the deformation retract simply follows this path backwards, as shown in Figure 5(c). Otherwise, if $v \in \mathcal{MR}$, there is a segment of a bisector in $\mathcal{MR}$ that contains v and continues in a direction that is perpendicular to ∂S_{t_i}. In this situation (as in the second case above), the shortest paths cannot be used as a deformation retract as it would not be continuous at that v. However, a similar re-parameterization as in the second case can be used in a local neighborhood of this portion of $\mathcal{MR}$ to construct a deformation retract, as shown in Figure 5(d).

Note that it is possible that ∂S_{t_i} is a graph, in which case the deformation retract described in the three cases above can be applied to individual components of C_i that are incident to a point $v \in \partial S_{t_i}$.

□

Finally, to show that $\mathcal{MR}$ is a finite graph, we observe that the set of concave vertices in $\mathcal{M}^{SPD}$ and the set of sectors around each such vertex are both finite on a PL surface, which implies that $\mathcal{MR} \setminus \mathcal{M}^{SPD}$ consists of a finite number of straight geodesic paths that bisect these

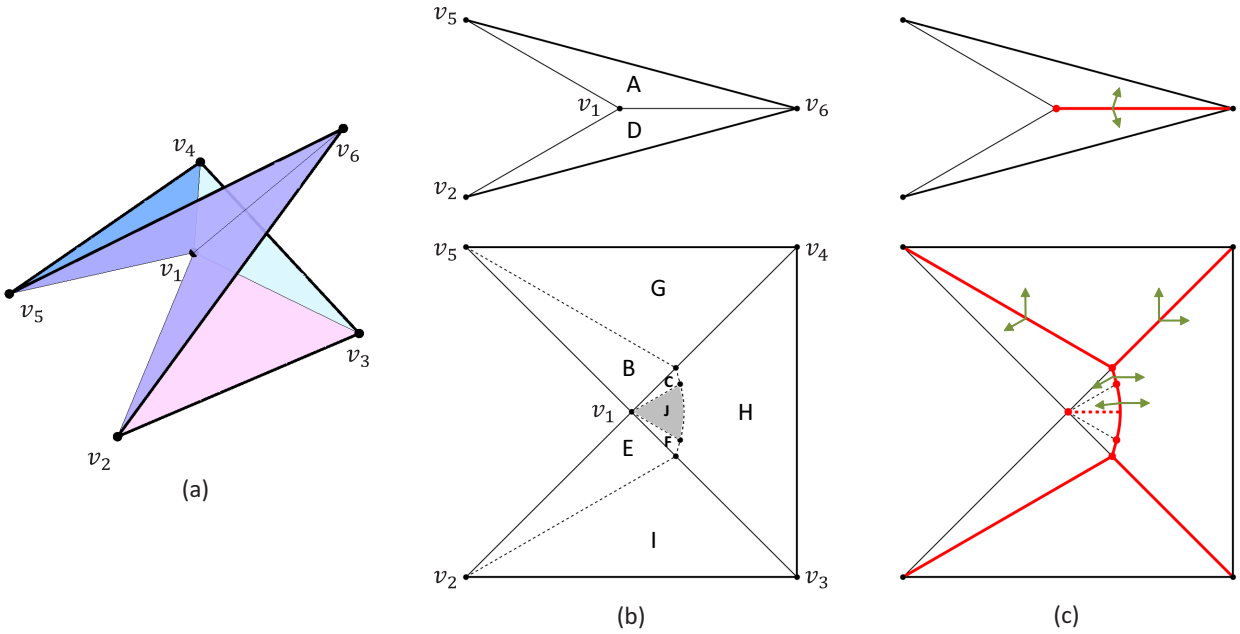

	A	B	C	D	E	F	G	H	I	J
Root	(v_5, v_6)	(v_5, v_6)	(v_5, v_6)	(v_2, v_6)	(v_2, v_6)	(v_2, v_6)	(v_4, v_5)	(v_3, v_4)	(v_2, v_3)	v_1
Last edge sequence		(v_1, v_5)	(v_1, v_5) (v_1, v_4)		(v_1, v_2)	(v_1, v_2) (v_1, v_3)				

Figure 6: (a): A simple surface consisting of 5 triangles incident to a concave vertex (v_1). (b): the surface is cut open along $\{v_5, v_1, v_2\}$ and the subdivision is shown on each part after flattening onto the plane. The root and last edge sequence of each labelled cell is listed below. (c): the medial residue consisting of $\mathcal{M}^{SPD}$ (red solid lines and red dots) and a bisector (red dotted line). Green arrows indicate shortest path directions for some segments in $\mathcal{M}^{SPD}$.

sectors. Together with our previous lemma that $\mathcal{M}^{SPD}$ is a finite graph, this completes the proof of Theorem 5.1.

6 Algorithm

We next give an overview of our algorithm to compute the medial residue on a piecewise flat surface with triangle faces in $\mathbb{R}^3$, a commonly used discretization in many applications.

We first recall some essential properties of shortest paths on a triangulated surface S [12, 7]. We assume the boundary ∂S consists of vertices and edges of some triangle faces. A shortest path p that connects any point $x \in S$ to the boundary ∂S originates either from a vertex or an interior point of an edge. In the latter case, p is orthogonal to that originating edge. The path p may go through some vertex of S, and if it does, both the left and right curve angles made by p at that vertex are greater than or equal to π. Away from the vertices, p is a straight line segment after unfolding the triangles that p goes through onto a plane. We call the last vertex visited by p before reaching x the *root* of p. If p does not go through any vertex, the root is the

originating vertex or edge on ∂S. The *last edge sequence* of p is the (possibly empty) sequence of edges that p goes through between the root and x.

The starting point of our algorithm is a subdivision of each triangle face into regions where the shortest paths have a common combinatorial structure. Given a face f, a root r (being either a vertex or edge), and an edge sequence E, a *cell* is the set of points $x \in f$ such that some shortest path from x to ∂S has root r and last edge sequence E. The curve segments that bound the cells (including both interior segments on f and the segments on the edges of f) form a graph, which is called the *subdivision graph*. An example subdivision graph is shown in Figure 6 (b) for the surface in (a). The subdivision can computed using an easy extension of existing methods [13, 12, 7] in $O(n^2 \log n)$ time and $O(n^2)$ space. This is described in Section 6.1.

Given the subdivision graph, our algorithm first identifies a subset of the graph as $\mathcal{M}^{SPD}$, then adds in the bisectors to form the complete $\mathcal{MR}$. Both steps can be done in $O(n^2)$ time and space, where n is the number of triangles of the surface. The overall process, taking into account the creation of the face subdivision, can be done in $O(n^2 \log n)$ time and $O(n^2)$ space. The two steps are described in Section 6.2 and 6.3. We assume exact arithmetic is used to precisely compute distances and angles.

6.1 Face subdivision

The subdivision graph can be obtained in two phases. In the first phase, we subdivide each triangle edge into intervals such that points within one interval share the same root and last edge sequence in their shortest paths to the boundary. This can be done in $O(n^2 \log n)$ time and $O(n^2)$ space using a "continuous Dijkstra" algorithm that was introduced by Mitchell et al. [12] for distance to point sources and later extended to both point and edge sources by Fort and Sellares [7].

In the second phase, the edge intervals are propagated into the interior of each triangle to form cells. Mount [13] gave an algorithm with $O(n)$ time and space within each triangle, which is a modification of the divide-and-conquer algorithm for computing generalized Voronoi diagram [16]. He also showed that the total number of cells, as well as the number of segments in the subdivision graph, on S is bounded by $O(n^2)$ ($O(n)$ for each face). Although the algorithm was presented for distances to point sources, it only needs to be slightly modified to work for edge sources as well. First, the "spokes" [16], which are used in the algorithm for efficient tracing of segments that bound a cell, are slightly different between a cell whose shortest paths are rooted at a vertex and a cell whose shortest paths are rooted at an edge. While the spokes emanate from the root vertex in the first case, they are parallel and orthogonal to the root edge in the second case. Second, the kinds of segments in the subdivision graph will include parabolas in addition to straight lines and hyperbolas. Neither adjustment increases the complexity of the algorithm or of the resulting subdivision graph.

6.2 Computing $\mathcal{M}^{SPD}$

First, we observe the following relation between $\mathcal{M}^{SPD}$ and a subdivision graph:

Lemma 6.1. *$\mathcal{M}^{SPD}$ is a subset of the subdivision graph.*

Proof. We first show that cells have disjoint interiors. Consider a cell c in a triangle f with root r and last edge sequence E. The distance to the boundary over the interior of c is either a linear or quadratic function, since it is defined by a distance function in an unfolded sequence of planar triangles where r (the root of the unfolded sequence of triangles) is either an edge or a point in the sequence of triangles. This function is unique for each pair of $\{r, E\}$. Since two distinct polynomial distance functions cannot have a region of points with non-zero area where

the two functions are equal, c cannot share a non-trivial region with another cell that has a different root or last edge sequence.

Since cells have disjoint interiors, the shortest paths from any interior point x in c to the boundary all have root r and last edge sequence E. These paths have to follow a common direction at x, as the distance function over c has no critical points. Hence x is not in $\mathcal{M}^{SPD}$. □

For a point x on the subdivision graph, the set of possible directions of shortest paths to the boundary starting at x can be obtained by examining the cells that x is incident to. For each incident cell on a face f whose root is r and last edge sequence is E such that x is away from r, the straight line direction from x to r (or the orthogonal direction to r if r is an edge) in the plane of f after unfolding the triangles along E is a shortest path direction of x. It follows that x is in $\mathcal{M}^{SPD}$ if there are two incident cells of x that give different shortest path directions at x. However, there are infinitely many points on the subdivision graph, and point-by-point checking for shortest path directions would be impractical. Fortunately, we only have to do the checking once for each segment in the subdivision graph, as ensured by the following lemma:

Lemma 6.2. *Let h be a segment in the subdivision graph, then either all interior points of h lie in $\mathcal{M}^{SPD}$ or none of them does.*

Proof. Suppose $x \in h$ has a unique shortest path direction to the boundary. We will show that all points on h have a unique shortest path direction as well.

Let the two cells c_1, c_2 on the two sides of h have roots r_1, r_2 and last edge sequences E_1, E_2. Unfold triangles respectively along E_1 and E_2 onto a common plane L that contains h, and let the locations of the two roots on the respectively unfolded triangles be r_1^* and r_2^*. Denote the distance function defined by these two unfolded roots over L as f_1, f_2, such that $f_i(z) = d_i + \|z - r_i^*\|$ if r_i is a vertex with distance d_i to the boundary and $f_i(z) = n_i^* \cdot (z - q_i^*)$ if r_i is an edge on the boundary whose normal and midpoint location after unfolding are respectively n_i^* and q_i^*. Note that the shortest path direction at x given by cell c_i after unfolding is opposite to the gradient direction of f_i at x. We consider two cases, and in each case we show that both f_1, f_2 have the same gradient direction at any point on h:

1. If $r_1^* = r_2^*$, we have $f_1 = f_2$ (in case of vertex roots, $d_1 = d_2$ because f_1, f_2 are identical on x which is away from r_1^* and r_2^*).
2. If $r_1^* \neq r_2^*$, $f_1 \neq f_2$ but they have the same value and gradient direction at x. The only possible scenario would be either 1) both r_1, r_2 are vertex roots and $\{r_1^*, r_2^*, x\}$ are co-linear, or 2) (without loss of generality) r_1 is a vertex root and r_2 is an edge root and the line connecting x and r_1^* is orthogonal to r_2^*. In both cases, the loci where $f_1 = f_2$, which contains h, is a straight line and the gradient direction of both functions at any point on that line follows the line.

□

The algorithm simply goes through each element (a vertex or a segment open at its ends) of the subdivision graph. For each element l, it picks an arbitrary point $x \in l$ and gathers shortest path directions at x by examining each incident cell of l. l is included in $\mathcal{M}^{SPD}$ as soon as two distinct shortest path directions are found. An example output of the algorithm is shown in Figure 6 (c), which also shows the shortest path directions for several segments in $\mathcal{M}^{SPD}$.

Since computing the shortest path direction given a cell takes constant time, the complexity of the algorithm is proportional to the number of pairs of an element and an incident cell, which is linear to the number of elements in the subdivision graph. The algorithm uses a data structure that maintains adjacency between cells and subdivision graph elements, which is again linear to the complexity of the graph. Hence computing $\mathcal{M}^{SPD}$ can be done in $O(n^2)$ time and $O(n^2)$ space.

6.3 Computing $\mathcal{MR} \setminus \mathcal{M}^{SPD}$

We use a tracing algorithm to compute bisectors that make up $\mathcal{MR} \setminus \mathcal{M}^{SPD}$. For each sector bounded by shortest path directions at some concave vertex $v \in \mathcal{M}^{SPD}$, we start tracing a straight and shortest path from v in the bisecting direction of the sector. Note that each such bisector also bisects a shadow rooted at v. Tracing proceeds in a cell-by-cell manner, creating straight line segments within each cell and maintaining straightness while marching to the next cell. Tracing ends when the path hits a segment or vertex of the subdivision graph that belongs to $\mathcal{M}^{SPD}$. An example bisector is shown in Figure 6 (c). Note that in this example, the bisector connects the two disconnected components of $\mathcal{M}^{SPD}$.

To efficiently identify the shadows, we first observe that the set of shadows rooted at v are precisely covered by those cells incident at v whose shortest paths to the boundary have v as their root (e.g., the shaded cell J in Figure 6 (b)). In general, each shadow is covered by a sequence of such cells $\{c_1, c_2, \ldots, c_k\}$ such that c_i and c_{i+1} share a common segment in the subdivision graph. Given the adjacency structure between cells in the subdivision graph, detecting all shadows at v takes time linear to the number of cells incident to v, and hence detecting all shadows at all vertices can be done in $O(n^2)$ time. Since the number of shadows is no more than the number of cells, the storage needed is bounded by $O(n^2)$.

Tracing within a cell involves intersecting a line with several low-degree algebraic curves. Since the intersection of a cell with a shortest path to the boundary is a single line segment [13], tracing in a cell can be done in time linear to the number of segments of the cell. Marching from one cell to the next can be done in constant time using an adjacency structure. To bound the complexity of tracing all bisectors, the key is to observe that each cell can contain a non-trivial portion of at most one bisector. This is because only a cell whose shortest paths to the boundary are rooted at some vertex may contain a bisector rooted at the vertex, and the angle made by any two bisectors rooted at a vertex is at least 2π. So the total tracing time for all bisectors is bounded by the sum of number of segments over all cells, which is $O(n^2)$. Tracing uses $O(n^2)$ space since it adds only a constant amount of additional data per element of the subdivision graph.

7 Future Directions

As an extension, we are exploring a similar "residue" definition for the *cut locus*, which is closely related to the medial axis. The cut locus of a point x is formally defined in terms of the tangent space, but it can be thought of as the set of points where minimizing geodesics starting at x are no longer minimizing. On piecewise linear surfaces even defining the cut locus is a challenge, since the tangent space is not well defined in a non-smooth setting. In practice, the most commonly used definition of the cut locus of a point x in a non-smooth setting is the closure of the set of points which have two distinct geodesics to x. This simple definition has the same difficulties as the medial axis on piecewise linear surfaces, since it can contain 2-dimensional sheets. As a result, algorithms for computing cut locus on a triangulated mesh either use approximation [2], remain limited to convex surfaces [10], or compute a variation (often called the cut graph) where only edges of the triangulation may be used [4, 6].

There is an obvious connection between the cut locus of a point and the medial axis, since the cut locus can be thought of as the medial axis of the surface with a puncture around the point; see for example Wolter's work connecting the two [19]. The cut locus is homotopy equivalent to the space minus the source point, and so it is useful for getting information about the topology of a space. However, just like the medial axis, the cut locus may fail to be a low-dimensional structure on a piecewise smooth surface; in fact, even on smooth Riemannian manifolds, the cut locus can form any infinite graph [9], so computing it in this setting is impossible. However, by replacing ∂S in Definition 4.1 with a point source x, we can similarly define a *cut residue* from

x. The homotopy and curve network properties in Theorem 5.1 should hold for cut residue when S is a piecewise linear surface. Since the face subdivision on a piecewise flat surface in $\mathbb{R}^3$ can be computed from any set of vertex or edge sources, our algorithm for computing the medial residue should be directly applicable (with the same complexity) to compute the cut residue on such surfaces.

Another future direction for work is extending our medial residue definition to higher dimensional piecewise linear manifolds. In fact, the definitions of straight curves and bisectors can be generalized to higher dimensions. Essentially, a straight curve is one where at every point the angle between the incoming and outgoing curve is maximized. Bisectors generalize to curves whose angle to the boundary is as large as possible. With these ideas, the medial residue can be defined for any piecewise-linear manifold in any dimension. We in fact conjecture that in every dimension the medial residue has the correct dimension and remains homotopy equivalent the the manifold.

Last but not least, we are also working on an implementation of the medial residue algorithm described here. The implementation will allow us to evaluate the practical performance of the algorithm on real inputs. Since the practical performance of the "continuous Dijkstra" algorithm by Mitchell et al. [12] has been shown to be much better than the theoretical bounds [17], we are hopeful that the same observation can be made with our algorithm.

References

[1] H. Blum. A transformation for extracting new descriptors of form. *Models for the Perception of Speech and Visual Form*, pages 362–80, 1967.

[2] Tamal K. Dey and Kuiyu Li. Cut locus and topology from surface point data. In *Proceedings of the 25th Annual Symposium on Computational Geometry*, SCG '09, pages 125–134, New York, NY, USA, 2009. ACM.

[3] Manfredo P. Do Carmo. *Differential Geometry of Curves and Surfaces*. Prentice Hall, 1976.

[4] Jeff Erickson and Sariel Har-Peled. Optimally cutting a surface into a disk. In *Proceedings of the 18th Annual Symposium on Computational Geometry*, SCG '02, pages 244–253, New York, NY, USA, 2002. ACM.

[5] Jeff Erickson and Amir Nayyeri. Tracing compressed curves in triangulated surfaces. In *Proceedings of the 2012 symposuim on Computational Geometry*, SoCG '12, pages 131–140, New York, NY, USA, 2012. ACM.

[6] Jeff Erickson and Kim Whittlesey. Greedy optimal homotopy and homology generators. In *Proceedings of the 16th Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '05, pages 1038–1046, Philadelphia, PA, USA, 2005. Society for Industrial and Applied Mathematics.

[7] M. Fort and J.A. Sellares. Generalized source shortest paths on polyhedral surfaces. In *Proceedings of the 23rd European Workshop on Computational Geometry*, pages 186–189, March 2007.

[8] Allen Hatcher. *Algebraic Topology*. Cambridge University Press, 2002.

[9] Jin ichi Itoh and Costin Vîlcu. Every graph is a cut locus. Preprint on arxiv:1103.1759, 2011.

[10] Jin-ichi Itoh and Robert Sinclair. Thaw: A tool for approximating cut loci on a triangulation of a surface. *Experimental Mathematics*, 13(3):309–325, 2004.

[11] André 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. Solid Modeling Theory and Applications.

[12] Joseph S. B. Mitchell, David M. Mount, and Christos H. Papadimitriou. The discrete geodesic problem. *SIAM J. Comput.*, 16(4):647–668, August 1987.

[13] David Mount. Voronoi diagrams on the surface of a polyhedron. Technical report, Dept. of Computer Science, Univ. of Maryland, Baltimore, MD, 1985.

[14] James Munkres. *Topology.* Pearson, 2000.

[15] Konrad Polthier and Markus Schmies. Straightest geodesics on polyhedral surfaces. In *ACM SIGGRAPH 2006 Courses*, SIGGRAPH '06, pages 30–38, New York, NY, USA, 2006. ACM.

[16] Micha Sharir. Intersection and closest-pair problems for a set of planar discs. *SIAM J. Comput.*, 14(2):448–468, 1985.

[17] Vitaly Surazhsky, Tatiana Surazhsky, Danil Kirsanov, Steven J. Gortler, and Hugues Hoppe. Fast exact and approximate geodesics on meshes. *ACM Trans. Graph.*, 24(3):553–560, July 2005.

[18] F.-E. Wolter and K.-I. Friese. Local and global geometric methods for analysis interrogation, reconstruction, modification and design of shape. In *Proceedings of the International Conference on Computer Graphics*, CGI '00, pages 137–151, Washington, DC, USA, 2000. IEEE Computer Society.

[19] Franzerich Wolter. Cut locus and medial axis in global shape interrogation and representation. In *MIT Design Laboratory Memorandum 92-2 and MIT Sea Grant Report*, 1992.