Extreme Elevation on a 2-Manifold *

Pankaj K. Agarwal†, Herbert Edelsbrunner‡, John Harer§ and Yusu Wang†

Abstract

Given a smoothly embedded 2-manifold in $\mathbb{R}^3$, we define the elevation of a point as the height difference to a canonically defined second point on the same manifold. Our definition is invariant under rigid motions and can be used to define features such as lines of discontinuous or continuous but non-smooth elevation. We give an algorithm for finding points of locally maximum elevation, which we suggest mark cavities and protrusions and are useful in matching shapes as for example in protein docking.

1 Introduction

The starting point of our work is the desire to identify features that are useful in finding a fit between solid shapes in $\mathbb{R}^3$. We are looking for cavities and protrusions and a way to measure their size. The problem is made difficult by the interaction of these features, which typically exist on various scale levels. We therefore take an indirect approach, defining a real-valued function on the surface that is sensitive to the features of the shape. We call this the elevation function because it has similarities to the elevation measured on the surface of the Earth, but the problem for general surfaces is more involved and the analogy is not perfect.

Related work in protein docking. The primary motivation for work reported in this paper is protein docking, which is the computational approach to predicting protein interaction, a biophysical phenomenon at the very core of life. The phenomenon is clearly important and the interest in protein docking is correspondingly wide-spread. We refer to survey articles by Elcock et al. [12], Halperin et al. [15], and Janin and Wodak [17]. The idea of docking by matching cavities with protrusions goes back to Crick [8] and Connolly [7]. Connolly also introduced the idea of using the critical points of a real-valued function defined on the protein surface to identify cavities and protrusions. The particular function he used is the fraction of a fixed-size sphere that is buried inside the protein volume as we move the sphere center on the protein surface. In the limit, when the size of the sphere goes to zero, this function has the same critical points as the mean curvature function [3]. A similar but different function suggested for the same purpose is the atomic density [20]. Here we take the buried fraction of the ball bounded by the sphere but we also vary its radius from zero to about ten Angstrom. At every point of the protein surface, the function value is the fraction of buried volume averaged over the balls centered at that point.

Results. The main contribution of this paper is the description and computation of a new type of feature points that mark extreme cavities and protrusions on a surface embedded in $\mathbb{R}^3$. More specifically,

- we extend the concept of topological persistence [10] to form a pairing between all critical points of a function on a 2-manifold embedded in $\mathbb{R}^3$;
- we use the pairings obtained for a 2-parameter family of height functions to define the elevation function on the 2-manifold;
- we classify the generic local maxima of the elevation function into four types;
- we develop and implement an algorithm that computes all local maxima of the elevation function.

The elevation differs from Connolly's and the atomic density functions in two major ways: it is independent of scale and it provides, beyond location, estimates for the direction and size of shape features. Both additional pieces of information are useful in shape characterization and matching. The four generic types of local maxima are illustrated in Figure 1. In each but the first case, the maximum is obtained at an ambiguity in the pairing of critical points. In all cases,

*This research was partially supported by NSF under grants EIA-99-72879, EIA-01-31905, CCR-00-86013, CCR-02-04118, DMS-01-07621, and by the U.S.–Israel Binational Science Foundation.

†Department of Computer Science, Duke University, Durham, North Carolina.

‡Department of Computer Science, Duke University, Durham, and Raindrop Geomagic, Research Triangle Park, North Carolina.

§Department of Mathematics, Duke University, Durham, North Carolina.

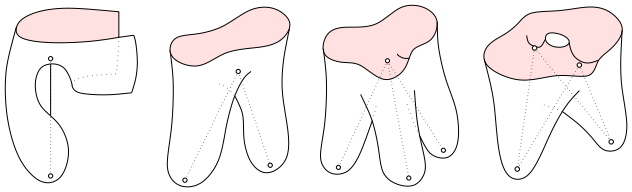

Figure 1: From left to right: a one-, two-, three-, and four-legged local maximum of the elevation function. In the examples shown, the outer normals at the endpoints of the legs are all parallel (the same). Each of the four types also exists with anti-parallel outer normals in various combinations.

the endpoints of the legs share the same normal line, and the legs have the same length if measured along that line. The case analysis is delicate and aided by a transformation of the original 2-manifold to its pedal surface, which maps tangent planes to points and thus expresses points with common tangent planes as self-intersections of the pedal surface. The algorithm we describe for enumerating all local maxima is inspired by our analysis of the smooth case but works on piecewise linear data.

Outline. Section 2 defines the pairing of the critical points. Section 3 introduces the height and elevation as functions on a 2-manifold. Section 4 describes a dual view of these concepts based on the pedal surface of the 2-manifold. Section 5 uses surgery to make elevation continuous and to define a stratified Morse function on the new 2-manifold. Section 6 characterizes the four types of generic local maxima. Section 7 sketches an algorithm for enumerating all local maxima. Section 8 presents preliminary experimental results for protein data. Section 9 concludes the paper.

2 Pairing

The elevation function is based on a canonical pairing of the critical points, which we describe in this section.

Traditional persistence. Let $\mathbb{M}$ be a connected and orientable 2-manifold and $f : \mathbb{M} \to \mathbb{R}$ a smooth function.[1] A point $x \in \mathbb{M}$ is *critical* if the derivative of f at x is identical 0, and it is *non-degenerate* if the Hessian at the point is invertible. It is convenient to assume that f is generic:

I. all critical points are non-degenerate;

II. the critical points have different function values.

A function that satisfies Conditions I and II is usually referred to as a *Morse function* [19]. It has three types of critical points: *minima*, *saddles* and *maxima* distinguished by the number of negative eigenvalues of the Hessian. Imagine we sweep $\mathbb{M}$ in the direction of increasing function value, proceeding along a level set of closed curves. We write $\mathbb{M}_a = \{x \in \mathbb{M} \mid f(x) \leq a\}$ for the swept portion of the 2-manifold. This portion changes the topology whenever the level set passes through a critical point. A component of $\mathbb{M}_a$ starts at a minimum and ends when it merges with another, older component at a saddle. A hole in the 2-manifold starts at a saddle and ends when it is closed off at a maximum. After observing that each saddle either merges two components or starts an new hole, but not both, it is natural to pair up the critical point that starts a component or a hole with the critical point that ends it. This is the main idea of topological persistence introduced in [10]. It is clear that a small perturbation of the function that preserves the sequence of critical events does not affect the pairing, other than by perturbing each pair locally. The method pairs all critical points except for the first minimum, the last maximum, and the $2g$ saddles starting the $2g$ cycles that remain when the sweep is complete. Here g is the *genus* of $\mathbb{M}$. These $2 + 2g$ unpaired critical points are the reason we need an extension to the method, which we describe next.

Extended persistence. It is natural to pair the remaining minimum with the remaining maximum. The remaining $2g$ saddles, S, contains g up-forking and g down-forking saddles. We wish to pair up-forking saddles with down-forking ones, and this can be achieved in a way that reflects how they introduce cycles during the sweep. This pairing is best described using the *Reeb graph* obtained by mapping each component of each level set to a point, as illustrated in Figure 2. As proved in [5], the Reeb graph has a basis of g

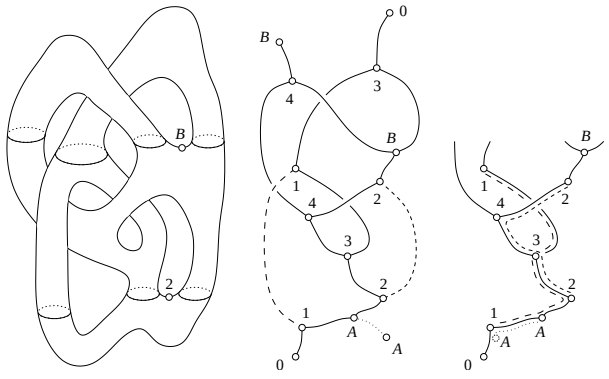

Figure 2: Left: a 2-manifold whose points are mapped to the distance above a horizontal plane. Middle: the Reeb graph in which the critical points of the function appear as degree-1 and degree-3 nodes. The labels indicate the pairing. Right: the tree representing the Reeb graph from slightly above B downwards.

cycles such that any cycle is the sum (modulo 2) of a subset of basis cycles. Each cycle has a unique lowest and a unique highest point, referred to as *lo-point* and *hi-point*. We

[1] We remark that the algorithms we describe below work for 2-manifolds with multiple components as well. We assume there is only one component for simplicity.

say the lo- and hi-point *span* this cycle but note that they may span more than one cycle. There is a one-to-one correspondence between lo- (hi-) points and up- (down-) forking saddles from S, thereby giving exactly g lo-points and g hi-points. We pair each lo-point x with the lowest hi-point y that spans a cycle with x. Note that x is also the highest lo-point that spans a cycle with y. Indeed, if it were not the case, then we could add the cycle spanned by y and x and the cycle spanned by y and the lo-point higher than x to get a cycle spanned by x and lower hi-point than y, a contradiction. This implies that each lo-point and each hi-point belongs to exactly one pair, giving a total of g pairs between up- and down-forking saddles from S as desired.

The Reeb graph of a piecewise-linear function on a triangulation with n edges can be constructed in time $O(n \log n)$ using the algorithm in [5]. We now describe an algorithm that computes both the traditional persistence pairing and the extended persistence pairing as introduced above, given the Reeb graph $\mathcal{G}$ of $\mathbb{M}$.[2] It simulates the sweep of $\mathcal{G}$, maintaining a forest for $\mathbb{M}_a$ during the course. In particular, it takes the following steps at reaching a critical point x (i.e., a node in $\mathcal{G}$), merging two arcs across a degree-2 node whenever one is created.

`Case 1:` x is a minimum. We add a new tree, consisting of a single node, to the forest.

`Case 2:` x is an up-forking saddle. We turn the corresponding leaf into an internal node, adding two new leaves as its children.

`Case 3:` x is a down-forking saddle, connecting leaves u and v. We glue the two downward paths from u and v to their roots, and ends the gluing at y. In one case, y is the root of one tree (the higher one); y is a minimum, which we now pair with x (this corresponds to a traditional persistence pairing). In the other case, y is the lowest common ancestor of u and v; y is an up-forking saddle, and we pair it with x (this corresponds to an extended persistence pairing).

`Case 4:` x is a maximum. We pair it with its parent y and remove the joining edge together with the two nodes; y can be either an up-forking saddle, producing a traditional persistence pairing, or a minimum, producing an extended persistence pairing.

In order to perform these operations efficiently, we use the linking and cutting tree data structure proposed by Sleator and Tarjan [21]. It decomposes the forest into a family of vertex-disjoint paths, and each path is represented using a biased binary search tree. By maintaining a linking and cutting tree T, cases $1, 2$, and 4 can be handled in $O(n \log n)$ overall time. So we focus only on case 3. Given an instance of case 3, assume that the common ancestor of u and v, y, exists (the case where it does not exist can be processed similarly). We can find y in $O(\log n)$ time using the operations supported by the linking and cutting tree data structure. The only extra operation we need is to glue the path in T from u to y with that from v to y. Let m and k be the length of these two paths with $k \leq m$. We can perform the gluing operation by inserting each node from the shorter path into the longer one, which takes $O(k \log m)$ time as each path is represented using a weighted binary search tree.[3] Assume that there are a sequence of s gluing operations during the entire sweep, and the i'th operation glues a path of length k_i with one of length m_i for $k_i \leq m_i$. The overall time complexity is $\sum_{i=0}^{s} k_i \log m_i \log n \leq \sum_i k_i \log^2 n$.

If we regard the parent of each node as its successor, then T induces a partial order on the nodes of G. Let ϕ_i be the number of total orders that are extensions of the partial order induced by T after the i'th gluing operation. Since $T = \emptyset$ initially and a single path after all operations, $\phi_0 = n!$ and $\phi_s = 1$. The i'th gluing operation merges two paths of length m_i and k_i, $\phi_{i-1} = \frac{(m_i+k_i)!}{m_i!k_i!} \cdot \phi_i$. Therefore,

$$\ln \phi_{i-1} - \ln \phi_i \quad = \quad \ln \frac{(m_i + k_i)!}{m_i!k_i!} \quad \geq \quad k_i \ln(1 + \frac{m_i}{k_i}).$$

Hence,

$$\sum_{i=1}^{s} k_i \quad \leq \quad \sum_{i=0}^{s-1} (\ln \phi_{i-1} - \ln \phi_i) \quad \leq \quad n \ln n,$$

implying that the overall time for computing the persistence pairing is $O(n \log^2 n)$.

Symmetry. The negative function, $-f : \mathbb{M} \to \mathbb{R}$, has the same critical points as f. We claim that it also generates the same pairing.

SYMMETRY LEMMA. Two critical points x and y are paired for f iff they are paired for $-f$.

PROOF. The claim is true for the first minimum, x, and the last maximum, y. Every other pair of f contains at least one saddle. We assume without loss of generality that x is a saddle and that $f(x) < f(y)$. Consider again the sweep of the 2-manifold in the direction of increasing values of f. When we pass $a = f(x)$ we split a cycle in the level set into two. The two cycles belong to the boundary of $\bar{\mathbb{M}}_a$, the set of points with function value a or higher. If the two cycles belong to the same component of $\bar{\mathbb{M}}_a$, such as for the point labeled 2 in Figure 2, then x is a lo-point and y is the lowest hi-point that spans a cycle with x. The claim follows because x is also the highest lo-point that spans a cycle with y. If, on the other hand, the two cycles belong to two different components of $\bar{\mathbb{M}}_a$, such as for the point labeled B in Figure 2, then y is the lower of the two maxima that complete the

[2] We remark that the algorithm can in fact construct the Reeb graph and the pairing simultaneously in one sweeep. We assume the Reeb graph is given for simplicity.

[3] In fact, for our algorithm, to achieve an $O(n \log^2 n)$ bound, a balanced binary tree will suffice.

two components. In the backward sweep (the forward sweep for $-f$), y starts a component that merges into the other, older component at x. Again x and y are also paired for $-f$, which implies the claimed symmetry. □

3 Height and Elevation

In this section, we define the elevation as a real-valued function on a 2-manifold in $\mathbb{R}^3$.

Measuring height and elevation on Earth. Even on Earth, defining the elevation of a point x on the surface is a non-trivial task. Traditionally, it is defined relative to the mean sea level (MSL) in the direction of the measured point. In other words, the *MSL elevation* of a point x is the difference between the distance of x from the center of mass and the distance of the MSL from the center of mass in the direction of x. The difficulty of measuring height in the middle of a continent was overcome by introducing the *geoid*, which is a level surface of the Earth's gravitational potential and roughly approximates the MSL while extending it across land. The *orthometric height* above (or below) the geoid is thus more general and about the same as the MSL elevation. It is perhaps surprising that the geoid differs significantly from its best ellipsoidal approximation due to non-uniform density of the Earth's crust [13]. Standard global positioning systems (GPS) indeed return the *ellipsoidal height*, which is elevation relative to a standard ellipsoidal representation of the Earth's surface. They also include knowledge of the *geoid height* relative to the ellipsoid and compute the orthometric height of x as its ellipsoidal height minus the ellipsoidal height of the geoid in the direction of x.

A simplifying factor in the discussion of height and elevation on Earth is the existence of a canonical core point, the center of mass. For general surfaces, distance measurements from a fixed center make much less sense. We are interested in this general case, which includes surfaces with non-zero genus for which there is no simple notion of core. As on Earth, we define the elevation of a point x as the difference between two distances, except we no longer use a reference surface, such as the mean sea level or the geoid, but instead measure relative to a canonically associated other point on the surface. To explain how this works, we give different meanings to the 'height' of a point x, which we define for every direction, and the 'elevation' of the point, which is the difference between two heights. While height depends on an arbitrarily chosen origin, we will see that elevation is independent of that choice. Indeed, the technical concept of elevation, as introduced shortly, will be similar in spirit to the idea of orthometric height, with the exception that it substitutes the canonical associated point for a globally defined reference surface.

Height, persistence and elevation. Let $\mathbb{M}$ be a smoothly embedded 2-manifold in $\mathbb{R}^3$. We assume that $\mathbb{M}$ is generic but it is too early to say what exactly that should mean. We define the height in a given direction as the signed distance from the plane normal to that direction and passing through the origin. Formally, for every unit vector $\mathbf{u} \in \mathbb{S}^2$, we call $f_{\mathbf{u}}(x) = \langle x, \mathbf{u} \rangle$ the *height* of x in the direction $\mathbf{u}$. This defines a 2-parameter family of height functions,

$$\text{Height} : \mathbb{M} \times \mathbb{S}^2 \to \mathbb{R},$$

where $\text{Height}(x, \mathbf{u}) = f_{\mathbf{u}}(x)$. The height is a Morse function on $\mathbb{M}$ for almost all directions. We pair the critical points of $f_{\mathbf{u}}$ as described in Section 2. Following [9], we define the *persistence* of a critical point as the absolute difference in height to the paired point: $\text{pers}(x) = \text{pers}(y) = |f_{\mathbf{u}}(y) - f_{\mathbf{u}}(x)|$.

Each point $x \in \mathbb{M}$ is critical for exactly two height functions, namely for the ones in the direction of its outer and inner normals: $\mathbf{u} = \pm \mathbf{n}_x$. We proved in Section 2 that the pairs we get for the two opposite directions are the same. Hence, each point $x \in \mathbb{M}$ has a unique persistence, which we use to introduce the *elevation function*,

$$\text{Elevation} : \mathbb{M} \to \mathbb{R},$$

defined by $\text{Elevation}(x) = \text{pers}(x)$. We note that the elevation is invariant under translation and rotation of $\mathbb{M}$ in $\mathbb{R}^3$.

Two-dimensional example. We illustrate the definitions of the height and elevation functions for a smoothly embedded 1-manifold $\mathbb{M}$ in $\mathbb{R}^2$. The critical points of $f_{\mathbf{u}} : \mathbb{M} \to \mathbb{R}$ are the points $x \in \mathbb{M}$ with normal vectors $\mathbf{n}_x = \pm \mathbf{u}$. Figure 3 illustrates a sweep in the vertical upward direction $\mathbf{u}$. Each critical point of $f_{\mathbf{u}}$ starts a component, ends a component by merging it into an older component, or closes the curve. The critical points that start components get paired with the other critical points. The elevation is zero at inflexion points and

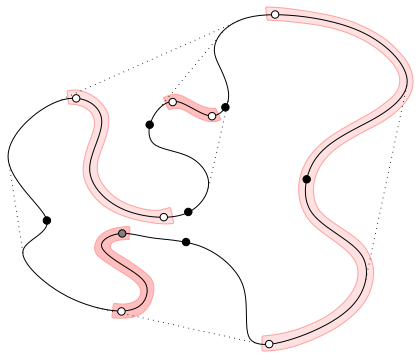

Figure 3: A 1-manifold with marked critical points of the vertical height function. The shaded strips along the curve connect paired critical points. The black and grey dots mark two- and one-legged elevation maxima.

increases as we move away in either direction. The function may experience a discontinuity at points that share tangent

lines with others, such as endpoints of segments that belong to the boundary of the convex hull. On the way towards a discontinuity, the elevation may go up and down, possibly several times. The elevation may reach a local maximum at points that either maximize the distance to a shared tangent line or the distance to another critical point in the normal direction. Examples of the first case are the black dots in Figure 3, where the elevation peaks in a non-differentiable manner. An example of the second case is the grey point, where the elevation forms a smooth maximum.

Singular tangencies. The elevation is continuous on $\mathbb{M}$, except possibly at points with singular tangencies. These points correspond to transitional violations of the two genericity conditions of Morse functions. Such violations are unavoidable as Height is a 2-parameter family within which we can transition from one Morse function to another:

- two critical points may converge and meet at a *birth-death point* where they cancel each other;
- two critical points may *interchange* their positions in the ordering by height, passing a direction at which they share the same height.

The first transition corresponds to an inflexion point of a geodesic on $\mathbb{M}$. Such points are referred to as *flat* or *parabolic*, indicating that their Gaussian curvature is zero. The second transition corresponds to two points $x \neq y$ that share the same tangent plane, $T_x = T_y$. Both types of singularities are forced by varying one degree of freedom and are turned into curves by varying the second degree of freedom. These curves pass through co-dimension two singularities formed by two simultaneous violations of the two genericity conditions. There can be two concurrent birth-death points, a birth-death point concurrent with an interchange, or two concurrent interchanges. In each case, the singularity is defined by two pairs of critical points and we get two types each because these pairs may be disjoint or share one of the points. See Table 1 for the features on $\mathbb{M}$ that correspond to the six types of co-dimension two singularities. We can now be more precise about what we mean by a generic 2-manifold.

GENERICITY ASSUMPTION A. The 2-parameter family of height functions on $\mathbb{M}$ has no violations of Conditions I and II for Morse functions other than the ones mentioned above (and enumerated in Table 1 below).

Some of these violations will be discussed in more detail later as they can be locations of maximum elevation. A second genericity assumption referring specifically to the elevation function will be stated in Section 5.

4 Pedal Surface

In this section, we take a dual view of the height and elevation functions based on a transformation of $\mathbb{M}$ to another surface in $\mathbb{R}^3$. We take this view to help our understanding of the singularities of Height, but it is of course also possible to study them directly using standard results in the field [1, 16].

Pedal function. Recall that T_x is the plane tangent to $\mathbb{M}$ that passes through $x \in \mathbb{M}$. The *pedal* p of x is the orthogonal projection of the origin on T_x. We write $p = \mathrm{Pedal}(x)$ and obtain a function

$$\mathrm{Pedal} : \mathbb{M} \to \mathbb{R}^3,$$

whose image $\mathbb{P} = \mathrm{Pedal}(\mathbb{M})$ is the *pedal surface* of $\mathbb{M}$ [2]. If the line $0x$ is normal to T_x then $p = x$. More generally, we can construct p by drawing the diameter sphere with center $x/2$ passing through 0 and x. This sphere intersects T_x in a circle with center $(x + p)/2$ that passes through x and $p = \mathrm{Pedal}(x)$. In fact, $\mathbb{P}$ is the evolute of the diameter spheres defined by the origin and the points $x \in \mathbb{M}$, as illustrated in Figure 4. The following three properties are

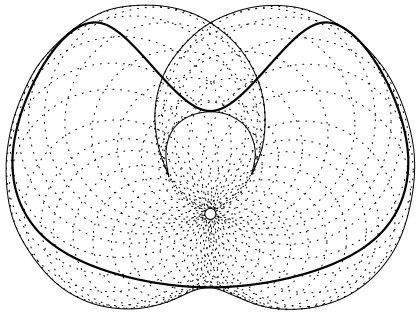

Figure 4: A smoothly embedded closed curve (boldface solid) and the image of the pedal function (solid) constructed as the evolute of the diameter circles (dotted) between the curve and the origin.

useful in understanding the correspondence between $\mathbb{M}$ and its pedal surface:

- points on $\mathbb{M}$ have parallel and anti-parallel normal vectors iff their images under the pedal function lie on a common line passing through the origin;
- the height of a point $x \in \mathbb{M}$ in the direction of its normal vector is equal to plus or minus the distance of $\mathrm{Pedal}(x)$ from the origin;
- from $p \in \mathbb{P}$ and the angle φ between the vector p and the normal $\mathbf{n}_p$ of $\mathbb{P}$ at p we can compute the radius ϱ of the corresponding diameter sphere and the preimage x at distance $2\varrho \sin\varphi$ from p in the direction normal to p and $p \times \mathbf{n}_p$.

The third property implies that the pedal surface determines the 2-manifold.

Tangents, heights, and pedals. We are interested in singularities of the pedal function as they correspond to directions

along which the height function is not generic. For example, a birth-death point of Height corresponds to a cusp point of $\mathbb{P}$. To see this recall that the birth-death point corresponds to a flat point $x \in \mathbb{M}$. A generic geodesic curve through this point has an inflexion at x, causing the tangent plane to reverse the direction of its rotating motion as we pass through x. Similarly, it causes a sudden reversal of the motion of the image point thus forming a cusp at $\mathrm{Pedal}(x)$. In contrast, an

Dictionary of Singularities		
$\mathbb{M}$	Height	$\mathbb{P}$
flat point double tangency	birth-death (bd) point interchange	cusp xing
Jacobi point triple tangency	2 bd-points 3 interchanges bd-pt. and interchange	dovetail point triple point cusp xing
	2 bd-points 2 interchanges bd-pt. and interchange	cusp-cusp overpass xing-xing overpass cusp-xing overpass

Table 1: Correspondence between singularities of tangents of the manifold, the 2-parameter family of height functions, and the pedal surface. There are two singularities of co-dimension one: curves of cusps and curves of self-intersections (xings). There are six singularities of co-dimension two.

interchange of Height, which corresponds to a plane tangent to $\mathbb{M}$ in two points, maps to a point of self-intersection (a xing) of $\mathbb{P}$. These two cases exhaust the co-dimension one singularities of Height, which are listed in the upper block of Table 1.

Co-dimension two singularities. There are six types of co-dimension two singularities listed in the lower block of Table 1. Perhaps the most interesting is formed by two concurrent birth-death points that share a critical point. As illustrated in Figure 5, left, the corresponding *dovetail point* in the pedal surface is endpoint of two cusps but also of a self-intersection curve. The second most interesting type is

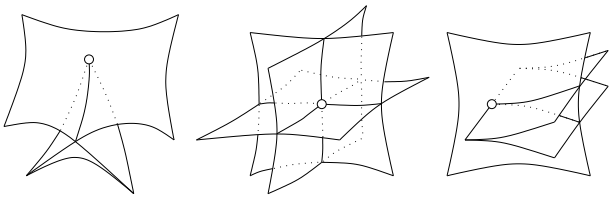

Figure 5: Left: a portion of the pedal surface in which a self-intersection and two cusps end at a dovetail point. Middle: three sheets of the pedal surface intersecting in a triple point. Right: a cusp intersecting another sheet of the pedal surface.

formed by two concurrent interchanges that share a critical point and therefore force a third concurrent interchange of the other two critical points. It corresponds to three self-intersection curves formed by three sheets of $\mathbb{P}$ that intersect in a *triple point*, as shown in Figure 5, middle. Third, we may have a concurrent birth-death point and interchange that share a critical point. As illustrated in Figure 5, right, this corresponds to a cusp curve that passes through another sheet of the pedal surface. There are three parallel types in which the concurrency happens in the same direction $\mathbf{u}$ but not in space. They correspond to two curves on the pedal surface that cross each other as seen from the origin but do not meet in $\mathbb{R}^3$. As before, a birth-death point corresponds to a cusp curve and an interchange to a curve of self-intersections.

5 Continuity

We are interested in the local maxima of the elevation function, which are the counterparts of mountain peaks and deepest points in the sea. But they are not well defined because the elevation can be discontinuous. We remedy this shortcoming through surgery.

Discontinuities at interchanges. As mentioned in Section 2, the pairs vary continuously as long as the height function varies without passing through interchanges and birth-death points (Conditions I and II). It follows that the elevation is continuous in regions where this is guaranteed. Around a birth-death point, the elevation is necessarily small and goes to zero as we approach the birth-death point. The only remaining possibility for discontinuous elevation is thus at interchanges, which happen when two points share the same tangent plane. As mentioned in Table 1, this corresponds to a point at which the pedal surface intersects itself. Figure 6 shows that discontinuities in the elevation can indeed arise at co-tangent points. We see four points with common vertical normal direction, of which y and z are co-tangent. Consider a small neighborhood of the vertical direction, $\mathbf{n}$, and observe that the critical points vary in neighborhoods of their locations for $f_{\mathbf{n}}$. The critical point near x changes its partner from the right side of y to the left side of z as it varies from left to right in the neighborhood of x. Similarly, the critical point near w changes its partner from the right side of z to the left side of y as it varies from left to right in the neighborhood of w. Since the height difference is the same at the time of the interchange, the elevation at x and w is still continuous. However, it is not continuous at y and at z, which both change their partners, either from x to w or the other way round. Not all interchanges cause discontinuities, only those that affect the pairing. These are the interchanges that affect a common topological feature arising during the sweep of $\mathbb{M}$ in the height direction.

Continuity through surgery. We apply surgery to $\mathbb{M}$ to obtain another 2-manifold $\mathbb{N}$ on which the elevation function is continuous. Specifically, we cut $\mathbb{M}$ along curves at which $\mathrm{Elevation} : \mathbb{M} \to \mathbb{R}$ is discontinuous, resulting in a 2-manifold with boundary, $\mathbb{B}$. Then we glue $\mathbb{B}$ along its

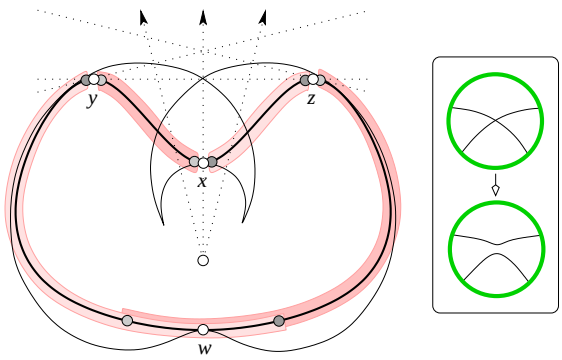

Figure 6: The four white points share the same normal direction, as do the four light shaded and the four dark shaded points. The strips indicate the pairing, which switches when the height function passes through the vertical direction. The insert on the right illustrates the effect of surgery at y and z on the pedal curve.

boundary, making sure that glued points have the same elevation. Formally, we cut by applying the inverse of a surjective map from $\mathbb{B}$ to $\mathbb{M}$, and glue by applying a surjective map from $\mathbb{B}$ to $\mathbb{N}$:

$$\mathbb{M} \xleftarrow{\text{Cut}^{-1}} \mathbb{B} \xrightarrow{\text{Glue}} \mathbb{N}.$$

As argued above, each boundary curve of $\mathbb{B}$ is defined by an interchange and corresponds to a self-intersection curve (a xing) of the pedal surface. The latter view is perhaps the most direct one in which surgery means cutting along xings and gluing the resulting four sheets in a pairing that resolves the self-intersection. This is illustrated in Figure 6 where on the right we see a self-intersection being resolved by cutting the two curves and gluing the upper and lower two ends. In the original boldface curve on the left, this operation corresponds to cutting at y and z and gluing the four ends to form two closed curves: one from y to x to $z = y$ and the other from y to w to $z = y$. As mentioned earlier, not all xings correspond to discontinuities and we perform surgery only on the subset that do. In general, a discontinuity follows a xing until it runs into a dovetail or a triple point. In the former case, the xing and the discontinuity both end. In the latter case, the xing continues through the triple point and the discontinuity may follow, turn, or even branch to other xings passing through the same triple point. Two possible configurations created by surgery in the neighborhood of a triple point p are illustrated in Figure 8 and 9 . Their particular significance in the recognition of local maxima will be discussed shortly. Whatever the situation, the subset of xings along which the elevation is discontinuous together with the gluing pattern across these xings provides a complete picture of how to use surgery to change $\mathbb{P}$ into a new surface, $\mathbb{Q}$. The 2-manifold $\mathbb{N}$ is the one for which this is the pedal surface: $\mathbb{Q} = \text{Pedal}(\mathbb{N})$. That $\mathbb{N}$ is indeed a manifold can be shown by (tediously) enumerating and examining all cases of cut-and-glue patterns that may occur. After surgery, we have a continuous function $\text{Elevation} : \mathbb{N} \to \mathbb{R}$. Furthermore, we have continuously varying pairs of critical points. To formalize this idea, we introduce a new map

$$\text{Antipode} : \mathbb{N} \to \mathbb{N}$$

that maps a point x to its paired point $y = \text{Antipode}(x)$. The function Antipode is a homeomorphism and its own inverse. We note in passing that we could construct yet another 2-manifold by identifying antipodal points. Each local maximum of the elevation function on this new manifold corresponds to a pair of equally high maxima in $\mathbb{N}$. This construction is the reason we will blur the difference between maxima and antipodal pairs of maxima in the next few sections.

Smoothness of Elevation**.** The elevation function on $\mathbb{N}$ is smooth almost everywhere. To describe the violations of smoothness, let $\text{Bd}\,\mathbb{B}$ denote the boundary of the intermediate manifold. Let $B = \text{Glue}(\text{Bd}\,\mathbb{B})$ and define $S = B \cup \text{Antipode}(B)$; S is the set of points at which the elevation function is not smooth. By Genericity Assumption A, S is a graph, consisting of *nodes* and *arcs*. We have degree-1 and degree-3 nodes that correspond to dovetail points and triple points in the pedal surface respectively, as well as degree-4 nodes that correspond to overpasses between xings. Each degree-4 node is the crossing of an arc in B and an arc in the antipodal image of B. We think of this construction as a *stratification* of $\mathbb{N}$. Its *strata* are

- the three kinds of nodes;
- the open and closed arcs;
- the open connected regions in $\mathbb{N} - S$.

Figure 7 illustrates the construction by showing how such a stratification may look like.

When restricted to every stratum, the elevation function is smooth, but still not a Morse function. For example, all points from lines of inflexion have elevation identical to 0, forming lines of local minima for Elevation. We now complete our description of what we mean by a generic 2-manifold.

GENERICITY ASSUMPTION B. The local maxima of Elevation on $\mathbb{N}$ are isolated.

The implication of this assumption becomes more clear after we enumerate the generic types of local maxima of the elevation function in next section. In particular, this means that surfaces such as spheres and cylinders are not generic under this assumption.

6 Elevation Maxima

In this section, we enumerate the generic types of local maxima of the elevation function. They come in pairs in $\mathbb{N}$ which, by inverse surgery, form multi-legged creatures in $\mathbb{M}$.

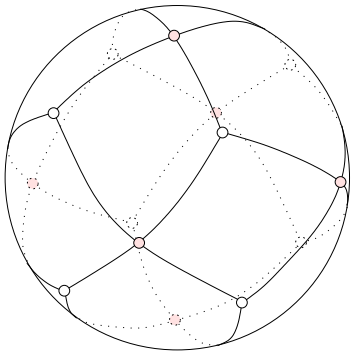

Figure 7: Stratification of the 2-sphere obtained by overlaying a spherical tetrahedron with its antipodal image. The (shaded) degree-4 nodes are crossings between B and its antipodal image.

Classification of local maxima. Depending on its location, a point $x \in \mathbb{N}$ can have one, two, or three preimages under surgery. We call this number its *multiplicity*, $\mu(x)$. Specifically, x has multiplicity three if it is a node of the graph B, it has multiplicity two if it lies on an arc of B, and it has multiplicity one otherwise. Degree-4 nodes in the stratification correspond to antipodal pairs of points with multiplicity two each. Let now $x \in \mathbb{N}$ be a local maximum of the elevation function. We know that x is not a flat point of $\mathbb{M}$, else its elevation would be zero. This simple observation eliminates five of the eight singularities in Table 1. Furthermore, the assumption of a generic 2-manifold $\mathbb{M}$ implies that a multiplicity three point can only be paired with a multiplicity one point. This leaves the following four possible types of local maxima x:

$$\left.\begin{array}{l} \textit{one-legged} \\ \textit{two-legged} \\ \textit{three-legged} \\ \textit{four-legged} \end{array}\right\} \text{if} \left\{\begin{array}{l} \mu(x) = \mu(y) = 1, \\ \mu(x) = 1 \text{ and } \mu(y) = 2, \\ \mu(x) = 1 \text{ and } \mu(y) = 3, \\ \mu(x) = \mu(y) = 2, \end{array}\right.$$

where $y = \mathrm{Antipode}(x)$; see Figure 1. We sometimes call the preimages of x the *heads* and those of y the *feet* of the maximum. The most exotic of the four types is perhaps the four-legged maximum, which corresponds to an overpass of two xings in the pedal surface or, equivalently, a degree-4 node in the stratification. The image under Pedal of x lies on one xing and the image of y lies on the other. Both maxima have two preimages under surgery, which makes for a complete bipartite graph with two heads, two feet, and four legs.

Neighborhood patterns. Given a point $x \in \mathbb{M}$, take an open neighborhood of x on $\mathbb{M}$. Denote by $\mathrm{B}(\mathbf{n}_x) \subseteq \mathbb{S}^2$ the image of this neighborhood under Gauss map[4], and refer to it as the neighborhood of $\mathbf{n}_x$. If x is not a flat point (i.e., the Gaussian curvature at x is not zero), then $\mathrm{B}(\mathbf{n}_x)$ is homeomorphic to an open disk, and there is a one-to-one map from the neighborhood of x to that of $\mathbf{n}_x$ under Gauss map. In the following discussion, we study only non-flat points since flat points will not possibly be maxima of the elevation function.

It is instructing to look at the local neighborhood of a maximum x in $\mathbb{M}$. Most interesting is the three-legged type, with feet y_1, y_2, y_3. A small perturbation of the normal direction can change the ambiguous pairing of x with all three to an unambiguous pairing of a point in the neighborhood of x with a point in the neighborhood of one of the feet. We indicate this by labeling the points in the neighborhood of $\mathbf{n}_x$ (i.e., $\mathrm{B}(\mathbf{n}_x)$) with the indices of the feet, as shown in Figure 8. The three curves passing through n_x correspond to the

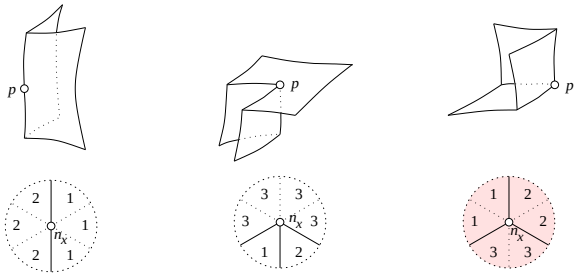

Figure 8: The three sheets of $\mathbb{Q}$ after cutting and gluing the neighborhood of a triple point p in $\mathbb{P}$ at the top, and the corresponding pairing patterns in the neighborhood of $\mathbf{n}_x$, at the bottom. The (shaded) Mercedes star is necessary for a three-legged maximum.

three xings passing through the triple point $p \in \mathbb{P}$. Note that in generic cases, such curves should pass through each other at $\mathbf{n}_x$ in a transversal manner as long as x is not a flat point. Hence, they decompose the neighborhood into six slices corresponding to the six permutations of the three feet. The labeling indicates the pairing and reflects the surgery at these feet and, equivalently, at the corresponding triple point in the pedal surface. Only the rightmost pattern in Figure 8 corresponds to a maximum, the reason of which will become clear later after we introduce and prove the necessary projection conditions for elevation maxima. We call this pattern the *Mercedes star* property of three-legged maxima.

There are in fact two ways to apply surgery at a three-legged maximum, one of which is already shown in Figure 8. We illustrate the neighborhood patterns of the other in Figure 9. The neighborhood pictures for the remaining three types of maxima are simpler. For a one-legged maximum we have an undivided disk, which requires no surgery. For a two- or four-legged maximum we have a disk divided into two halves and there is only one way to do the surgery.

Necessary projection conditions. Given a maximum x of $\mathrm{Elevation} : \mathbb{N} \to \mathbb{R}$ with $y = \mathrm{Antipode}(x) \in \mathbb{N}$, recall that the corresponding heads and feet $x_i, y_i \in \mathbb{M}$ have parallel tangent planes, or more precisely, all x_i's (resp. y_i's) are

[4]The Gauss map takes each point on $\mathbb{M}$ to its normal vector on $\mathbb{S}^2$.

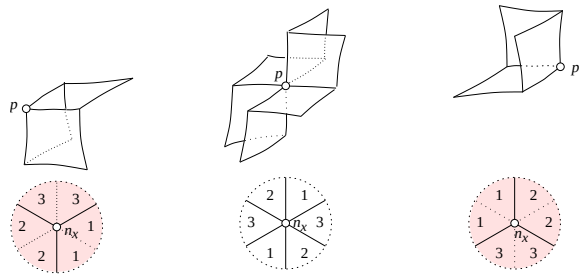

Figure 9: The second type of surgery pattern at a triple point: The three sheets of $\mathbb{Q}$ after cutting and gluing the neighborhood of a triple point p in $\mathbb{P}$, at the top, and the corresponding pairing patterns in the neighborhood of $\mathbf{n}_x$, at the bottom.

co-tangent. Furthermore, they should satisfy the following necessary properties, stated as *Projection Conditions*:

PROJECTION CONDITIONS. The point x is a maximum of the elevation function only if

- $\#\text{legs} = 1$: $\mathbf{n}_x$ is parallel or anti-parallel to $y - x$;
- $\#\text{legs} = 2$: $\mathbf{n}_x$, $y_1 - x$, and $y_2 - x$ are linearly dependent and the orthogonal projection of x onto the line of the two feet lies between y_1 and y_2;
- $\#\text{legs} = 3$: the orthogonal projection of x onto the plane of the three feet lies inside the triangle spanned by y_1, y_2 and y_3;
- $\#\text{legs} = 4$: the orthogonal projections of the segments x_1x_2 and y_1y_2 onto a plane parallel to both have a non-empty intersection.

In summary, x is a local maximum only if $\mathbf{n}_x$ is either a positive or a negative linear combination of the vectors $y_j - x_i$. We now prove the above necessary conditions for one-legged maxima, the correctness of these conditions for all other types of maxima can be shown using a similar but more involved argument.

Given a one-legged maximum x of $\text{Elevation} : \mathbb{N} \to \mathbb{R}$ with $y = \text{Antipode}(x)$, let x and y also denote their preimage in $\mathbb{M}$ before the surgery. Assume without loss of generality that x lies at the origin, $\mathbf{n} = \mathbf{n}_x = -\mathbf{n}_y = (0, 0, 1)$, $b = 0$ and $c \geq 0$ for $y = (a, b, c)$. By definition, $\text{Elevation}(x) = \text{Elevation}(y) = c$, which is the height difference between x and y in the vertical direction.

We parameterize points in $\mathrm{B}(\mathbf{n}) \subseteq \mathbb{S}^2$ by $\mathbf{n}(\theta, \alpha)$ as illustrated in Figure 10, where $\theta \in [0, 2\pi)$ and $\alpha \in (0, \epsilon)$ for an arbitrarily small $\epsilon > 0$. Next, let $x(\theta, \alpha) \in \mathbb{M}$ (resp. $y(\theta, \alpha) \in \mathbb{M}$) be the point in the neighborhood of x (resp. y) with normal $\mathbf{n}(\theta, \alpha)$. Denote by $h(\theta, \alpha)$ the height difference between $x(\theta, \alpha)$ and $y(\theta, \alpha)$ in direction $\mathbf{n}(\theta, \alpha)$, i.e.,

$$\begin{aligned} h(\theta, \alpha) &= \text{Elevation}(x(\theta, \alpha)) = \text{Elevation}(y(\theta, \alpha)) \\ &= \langle\, y(\theta, \alpha) - x(\theta, \alpha)\ ,\ \mathbf{n}(\theta, \alpha)\,\rangle. \end{aligned}$$

Obviously, x is a maximum if and only if $h(\theta, \alpha) \leq c$ for all $\theta \in [0, 2\pi)$, with a small enough $\epsilon > 0$.

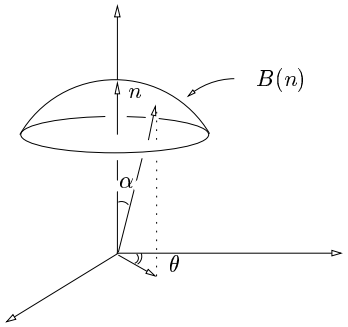

Figure 10: Normal n and the parameterization of its neighborhood $\mathrm{B}(n)$, which is a spherical patch from $\mathbb{S}^2$.

On the other hand, simple geometry shows that

$$\begin{aligned} \mathbf{n}(\theta, \alpha) &= (\sin\alpha\cos\theta, \sin\alpha\sin\theta, \cos\alpha), \quad \text{and} \\ y(\theta, \alpha) &= (a + \rho_y(\theta)\sin\alpha\cos\theta, b + \rho_y(\theta)\sin\alpha\sin\theta, \\ &\quad c - \rho_y(\theta)(1 - \cos\alpha)), \end{aligned}$$

where $\rho_y(\theta)$ denote the radius of curvature at position y in direction parameterized by θ. Similarly, we can compute $x(\theta, \alpha)$. It then follows that

$$\begin{aligned} h(\theta, \alpha) &= \langle\, y(\theta, \alpha) - x(\theta, \alpha)\ ,\ \mathbf{n}(\theta, \alpha)\,\rangle \\ &= a\sin\alpha\cos\theta + b\sin\alpha\sin\theta + c\cos\alpha \\ &\quad +(\rho_x(\theta) + \rho_y(\theta))(1 - \cos\alpha). \end{aligned}$$

As $b = 0$, we then have

$$\begin{aligned} h(\theta, \alpha) - c &= (1 - \cos\alpha)[a\cot\frac{\alpha}{2}\cos\theta \\ &\quad - c + (\rho_x(\theta) - \rho_y(\theta))]. \end{aligned} \tag{1}$$

The first term in the bracket dominates the above value for a small enough α, as $\cot\frac{\alpha}{2}$ tends to infinity in that case. This implies that x is a maximum if and only if (i) $a = 0$; and (ii) for any $\theta \in [0, 2\pi)$, $c \geq \rho_x(\theta) + \rho_y(\theta)$. Note that the Projection Condition for one-legged maxima is the same as (i) and is thus indeed necessary. Furthermore, if we add (ii), then the conditions are also sufficient. In fact, the necessary Projection Conditions for all other types of maxima can be made sufficient by adding appropriate constraints on the curvature of $\mathbb{M}$ at x_i's and at y_i's, and, if x is three-legged, it also needs the Mercedes star property. These curvature conditions are not used in our algorithm for piecewise-linear manifolds. So we omit their description.

It is now easy to see that the Mercedes star property is necessary for a three-legged maximum. This is because that for all the remaining neighborhood patterns in Figure 8 and 9, there exist pairs of antipodal normals in $\mathrm{B}(n_x)$ that are marked by the same index. For example, in the middle pattern in Figure 8, there are such normals both marked by index 3. By Eqn(1), $(h(\theta, \alpha) - c)$ and $(h(-\theta, \alpha) - c)$ have opposite signs, where $c = \text{Elevation}(x)$, that is, they cannot both be negative. Therefore x cannot be a maximum in this case.

7 Algorithm

In this section, we describe an algorithm for constructing all points with locally maximum elevation. The input is a piecewise linear 2-manifold embedded in $\mathbb{R}^3$. The running time of the algorithm is polynomial in the size of the 2-manifold.

Smooth vs. piecewise linear. We consider the case in which the input is a two-dimensional simplicial complex in $\mathbb{R}^3$. This data violates some of the assumptions we used in our mathematical considerations above. This causes difficulties which, with some effort, can be overcome. For example, it makes sense to require that K be a 2-manifold but not that it be smoothly embedded. The 2-parameter family of height functions is well-defined and continuous but not smooth. The definition of the elevation function is more delicate as it makes reference to point pairs in all possible directions. For any given direction, we get a well-defined collection of pairs, but how can we be sure that the pairs for different directions are consistent? The difficulty is rooted in the fact that a vertex in K can be critical for more than one direction and it may be paired with different vertices in different directions. To rationalize this phenomenon, we follow [9] and think of K as the limit of an infinite series of smoothly embedded 2-manifolds. A vertex of K gets resolved into a small patch with a two-dimensional variety of normal directions. Even as the patch shrinks toward the vertex, the variety of normal directions may remain fixed or at least not contract. For different directions in this variety, the corresponding points on the patch may be paired with points from different other patches. It thus seems natural that in the limit a vertex would be paired to more than one other point.

To make this idea concrete, we introduce a combinatorial notion of the variety of normal directions. Let σ be a simplex in K (a vertex, edge, or triangle), let x be a point in the interior of σ, and let $\mathbf{n} \in \mathbb{S}^2$ be a direction. We say x is *critical* for the height function in the direction $\mathbf{n}$ if

(i) $\langle \mathbf{n}, z - x \rangle = 0$ for all points z of σ;

(ii) the lower link of σ is not contractible to a point.

For example, the empty lower link of a minimum and the complete circle of a maximum are both not contractible. Let $\mathbf{N}(x) \subseteq \mathbb{S}^2$ be the set of directions along which x is critical. Generically, the set $\mathbf{N}$ for a point inside a triangle of K is an antipodal pair of points, that for a point on an edge is an antipodal pair of open great-circle arcs, and that for a vertex is an antipodal pair of open spherical polygons. Here, the word 'generic' applies to a simplicial complex in $\mathbb{R}^3$, where it simply means that the vertices are in general position. Computationally, this assumption can be simulated by a symbolic perturbation [11]. We write $\mathbf{N}(x, y, \ldots)$ for the common intersection of the sets $\mathbf{N}$ of x, y and so on.

Finite candidate sets. Given a candidate for a maximum, we can use the extended persistence algorithm to decide whether or not it really is a maximum. More specifically, we need a point x and a direction $\mathbf{n}$ along which the sweep defining the pairing proceeds. The details of this decision algorithm will be discussed shortly. We use the Projection Conditions, which are necessary for local maxima, to get four kinds of candidates:

$\#\text{legs} = 1$**:** pairs of points x and y on K with the direction $(y - x)/\|y - x\|$ contained in $\mathbf{N}(x, y)$;

$\#\text{legs} = 2$**:** triplets of points x, y_1, y_2 such that the orthogonal projection z of x onto the line of y_1 and y_2 lies between the two points and the direction $(z - x)/\|z - x\|$ is contained in $\mathbf{N}(x, y_1, y_2)$;

$\#\text{legs} = 3$**:** quadruplets of points x, y_1, y_2, y_3 such that the orthogonal projection z of x onto the plane of y_1, y_2, y_3 lies inside the triangle and the direction $(z-x)/\|z - x\|$ is contained in $\mathbf{N}(x, y_1, y_2, y_3)$;

$\#\text{legs} = 4$**:** quadruplets of points x_1, x_2, y_1, y_2 such that the shortest line segment zw connecting the lines of x_1, x_2 and y_1, y_2 also connects the two line segments and the direction $(z - w)/\|z - w\|$ is contained in $\mathbf{N}(x_1, x_2, y_1, y_2)$.

With the assumption of a generic simplicial complex K, we get a finite set of candidates of each kind. Since this might not be entirely obvious, we discuss the one-legged case in some detail. Let σ and τ be two simplices and x and y points in their interiors. For a generic K, the intersection of normal directions, $\mathbf{N}(x, y)$, is non-empty only if one of the two simplices is a vertex or both are edges. If $x = \sigma$ is a vertex then y is necessarily the orthogonal projection of x onto τ, which may or may not exist. If σ and τ are both edges then xy is necessarily the line segment connecting σ and τ and forming a right angle with both, which again may or may not exist. In the end, we get a set of $O(n^2)$ candidate pairs x and y, where n is the number of edges in K. For the two-legged case, we get $O(n^3)$ candidates, each a triplet of vertices or a pair of vertices together with a point on an edge. For the three- and four-legged cases, we get $O(n^4)$ candidates, each a quadruplet of vertices, giving a total of $O(n^4)$ candidates.

Verifying candidates. Let $x, y \in \mathbb{N}$ be a pair of points whose heads and feet all have parallel or anti-parallel normal directions. In the smooth case, the necessary and sufficient conditions for x and y to define an elevation maximum consists of three parts:

(a) the Projection Conditions of Section 5;

(b) the requirement that $y = \text{Antipode}(x)$;

(c) the curvature constraint alluded to in Section 5.

We subsume the Mercedes star property in (b) since it depends on the antipodality map or, equivalently, on the pairing by extended persistence. In the piecewise linear case, we only have (a) and (b) because the concentration of the curvature at the edges and vertices renders (c) redundant. We

have seen above how to translate (a) to the piecewise linear case. It remains to test (b), which reduces to answering a constant number of *antipodality queries*: given a direction $\mathbf{n}$ and a critical point x of $f_{\mathbf{n}}$, find the paired critical point y. This is part of what the algorithm described in Section 2 computes if applied to a sweep of K in the direction $\mathbf{n}$. More precisely, the algorithm computes one of the possible pairs, if applied in non-generic directions in which two or more vertices share the same height. Most of our candidates generate non-generic directions, and we cope with this situation by running the algorithm several times, namely once for each combination of permutations of the heads and of the feet. Each combination corresponds to a generic direction that is infinitesimally close to the non-generic direction. The largest number of combinations is six, which we get for three-legged maxima. This is also how we decide the Mercedes star property: each of the three feet is the answer to exactly two of the six antipodality queries. Letting n be the number of edges, the algorithm takes time $O(n \log^2 n)$ to answer the antipodality query. Since we have $O(n^4)$ candidates to test, this amounts to a total running time of $O(n^5 \log^2 n)$.

8 Experiments

We implemented the algorithm described in Section 7 and used it on surface representations of a few protein structures. We describe the findings to illustrate how the concepts introduced in the earlier sections might be applied.

Elevation on surface. We discuss the experimental findings for chain A of the protein complex with pdbid *1brs*, which we downloaded from the protein data bank. It contains 864 atoms, not counting the hydrogens which are too small to be resolved in the x-ray experiment and are not part of the structure. The particular surface representation we use is the molecular skin [4], which is similar to the better known molecular surface [**?**]. The reason for our choice is the availability of triangulating software and the guaranteed smooth embedding. The computed triangulation displayed in Figure 11 has slightly more than 50 thousand vertices after some simplification. Given the triangulation of the skin surface, we compute the elevation for each vertex of it and visualize it in Figure 11: Recall that each vertex has a range of directions associated with it that makes it critical, from which we choose an arbitrary one to compute its elevation value.

Number of maxima. Table 2 gives the number of maxima of each type computed for the skin triangulation for protein 1brs. We show in a separate row the number of additional maxima paired by extended persistence as introduced in Section 2. Since the genus of this particular surface is zero, all these maxima lie on the convex hull of the surface.

We notice that there are significantly more two-legged than other types of maxima. The reasons is perhaps the particular shape of molecules in which covalently bonded atoms form small dumbbells which invite two-legged maxima with one foot on each atom. These dumbbells are rotationally symmetric and form surface patches with non-generic elevation function, which further contributes to the abundance of two-legged maxima. The configurations required to form one-legged and three-legged maxima are considerably more demanding, but when they occur the maxima tend to have higher elevation. This observation is quantified in Figure 12, which sorts the maxima in the order of decreasing elevation. We see that for each threshold, the fraction of three-legged maxima higher than that elevation is significantly larger than the fractions of two- and four-legged maxima. The difference is even more pronounced for one-legged maxima of which four of the five have elevation exceeding 5 Angstrom. The statistics for other proteins are similar.

#legs	one	two	three	four
#max (trad.)	5	3,617	728	1,103
#max (addl.)	15	0	6	0

Table 2: The number of maxima for the molecular skin of the *1brs* protein structure obtained via traditional persistence (second row) and the additional maxima obtained by its extension (third row).

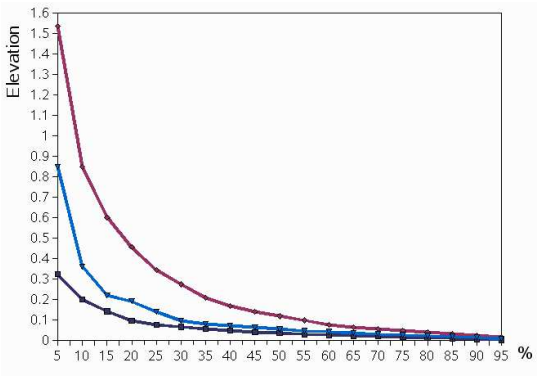

Figure 12: The percentage of maxima with elevation exceeding the threshold marked on the vertical axis. From top to bottom: the curves for the three-legged, four-legged, and two-legged maxima.

High elevation maxima. We are indeed mostly interested in high elevation maxima as the others are likely consequences of insignificant surface fluctuations or artifacts of the piecewise linear nature of the data. Figure 13 shows the top one hundred maxima on the skin surface of *1brs*. Each antipodal pair of maxima is represented by its one or two heads and one, two, or three feet.

One might expect that the binding site of a protein would perhaps have more or higher maxima. We did not observe any such trend in the few cases we studied. It seems that maxima are more or less uniformly distributed over the surface. This should be contrasted to the finding that in many

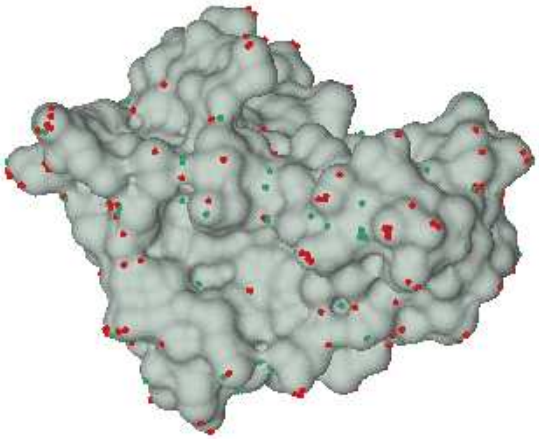

Figure 13: The one hundred pairs of maxima with highest elevation. The heads are marked by light and the feet by dark dots.

cases the pocket with the largest volume identifies the location of the binding site [18]. The elevation is indeed a less specific measurement with respect to a single surface, and we expect its primary use to be in the study of interactions between two or more shapes.

9 Discussion

The main contribution of this paper is the definition of elevation as a real-valued function on a 2-manifold embedded in $\mathbb{R}^3$ and the computation of all local maxima. The logical next step in this research is the exploitation of the maxima in protein docking and other shape matching problems. It would be worth exploring extensions of our results to manifolds with boundary and to manifolds of dimension three or higher. A crucial first step will have to be the generalization of the concept of extended persistence to these more general topological spaces. The algorithm presented in Section 5 enumerates all local maxima of the elevation function, without computing the elevation function itself, other than at a collection of candidate points. This approach is suggested by the ambiguities that arise in the definition of the elevation function for piecewise linear data. Unfortunately, it implies the fairly high running time of $O(n^5 \log^2 n)$ in the worst case. Can the maxima be enumerated more efficiently than that? Is there an algorithm that enumerates all maxima above some elevation threshold without computing the maxima below the threshold?

References

[1] V. I. ARNOLD. *Catastrophe Theory.* Springer-Verlag, Berlin, Germany, 1984.

[2] J. W. BRUCE AND P. J. GIBLIN. *Curves and Singularities.* Second edition, Cambridge Univ. Press, England, 1992.

[3] F. CAZALS, F. CHAZAL AND T. LEWINER. Molecular shape analysis based upon the Morse-Smale complex and the Connolly function. *In* "Proc. 19th Ann. Sympos. Comput. Geom., 2003", 351–360.

[4] H.L. CHENG, T. K. DEY, H. EDELSBRUNNER AND J. SULLIVAN. Dynamic skin triangulation. *Discrete Compt. Geom.* **25** (2001), 525–568.

[5] K. COLE-MCLAUGHLIN, H. EDELSBRUNNER, J. HARER, V. NATARAJAN AND V. PASCUCCI. Loops in Reeb graphs of 2-manifolds. *Discrete Comput. Geom.*, to appear.

[6] M. L. CONNOLLY. Analytic molecular surface calculation. *J. Appl. Crystallogr.* **6** (1983), 548–558.

[7] M. L. CONNOLLY. Shape complementarity at the hemoglobin albl subunit interface. *Biopolymers* **25** (1986), 1229–1247.

[8] F. H. C. CRICK. The packing of alpha-helices: simple coiled coils. *Acta Crystallogr.* **6** (1953), 689–697.

[9] H. EDELSBRUNNER, J. HARER AND A. ZOMORODIAN. Hierarchical Morse-Smale complexes for piecewise linear 2-manifolds. *Discrete Comput. Geom.* **30** (2003), 87–107.

[10] H. EDELSBRUNNER, D. LETSCHER AND A. ZOMORODIAN. Topological persistence and simplifi cation. *Discrete Comput. Geom.* **28** (2002), 511–533.

[11] H. EDELSBRUNNER AND E. P. MÜCKE. Simulation of Simplicity: a technique to cope with degenerate cases in geometric algorithms. *ACM Trans. Graphics* **9** (1990), 66–104.

[12] A. H. ELCOCK, D. SEPT AND J. A. MCCAMMON. Computer simulation of protein-protein interactions. *J. Phys. Chem.* **105** (2001), 1504–1518.

[13] W. FRACZEK. Mean sea level, GPS, and the geoid. *ArcUsers Online*, ERSI Web Sites: www.esri.com/news/arcuser/0703/summer-2003.html, 2003.

[14] M. GORESKY AND R. MACPHEARSON. *Stratified Morse Theory.* Springer-Verlag, Heidelberg, Germany, 1988.

[15] I. HALPERIN, B. MA, H. WOLFSON AND R. NUSSINOV. Principles of docking: an overview of search algorithms and a guide to scoring functions. *Proteins* **47** (2002), 409–443.

[16] A. HATCHER AND J. WAGONER. *Pseudo-Isotopies of Compact Manifolds.* Société Mathématique de France, 1973.

[17] J. JANIN AND S. J. WODAK. The structural basis of macromolecular recognition. *Adv. Protein Chem.* **61** (2002), 9–73.

[18] J. LIANG, H. EDELSBRUNNER AND C. WOODWARD. Anatomy of protein pockets and cavities: measurement of binding site geometry and implications for ligand design. *Protein Sci.* **7** (1998), 1884–1897.

[19] J. MILNOR. *Morse Theory*. Princeton Univ. Press, New Jersey, 1963.

[20] J. C. MITCHELL, R. KERR AND L. F. TEN EYCK. Rapid atomic density measures for molecular shape characterization. *J. Mol. Graph. Model.* **19** (2001), 324–329.

[21] D. D. SLEATOR AND R. E. TARJAN. A data structure for dynamic trees. *J. Comput. Sys. Sci.* **26** (1983), 362–391.

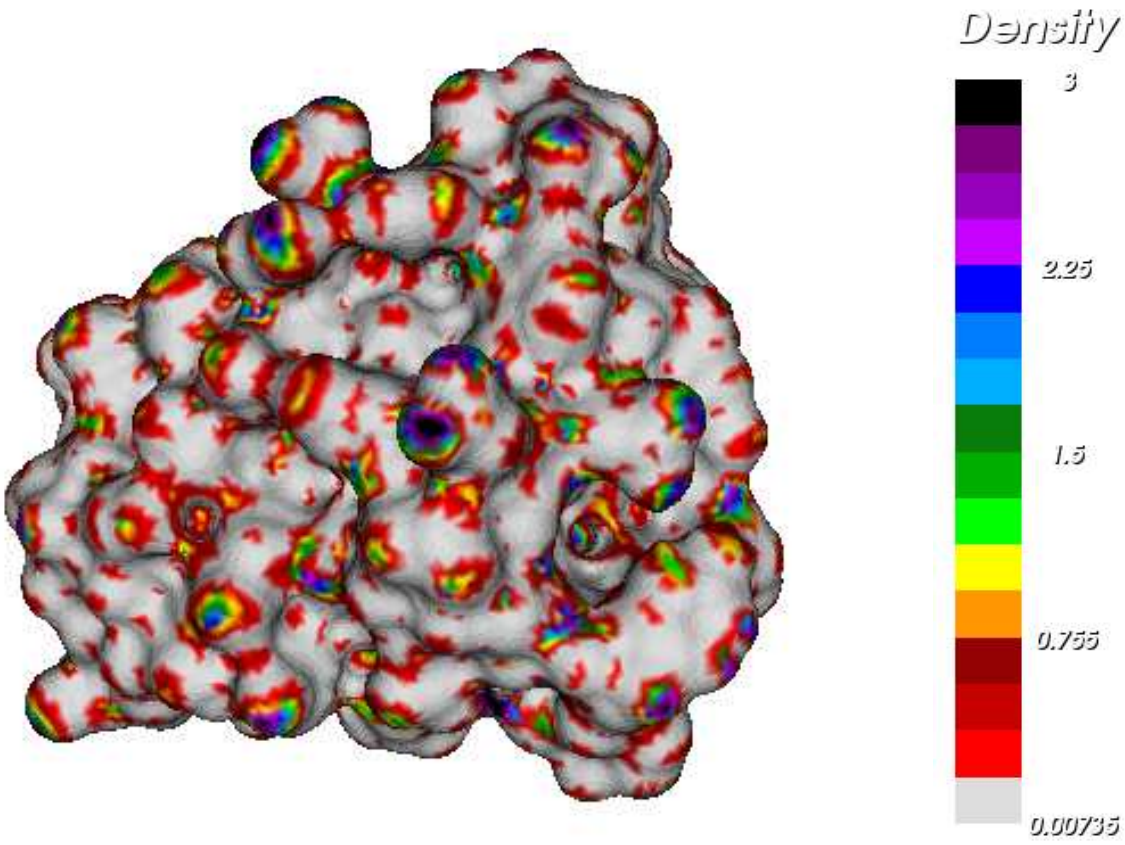

Figure 11: Visualization of elevation on skin surface for protein 1brs. Roughly, the higher the elevation is, the darker the color is.