Nearly Convex Segmentation of Polyhedra Through Convex Ridge Separation

Guilin Liu, Zhonghua Xi, Jyh-Ming Lien*

George Mason University, Fairfax, VA, U.S.A.

Abstract

Decomposing a 3D model into approximately convex components has gained more attention recently due to its ability to efficiently generate small decompositions with controllable concavity bounds. However, current methods are computationally expensive and require many user parameters. These parameters are usually unintuitive thus adding unnecessary obstacles in processing a large number of meshes with various types and shapes or meshes that are generated online within applications such as video games. In this paper, we investigate an approach that decomposes a mesh P based on the identification of **convex ridges**. Intuitively, convex ridges are the protruding parts of the mesh P . Our method, called **CoRiSe**, extracts nearly convex components of P by separating each convex ridge from all the other convex ridges through the new concept of **residual concavity**. **CoRiSe** takes a single user parameter: concavity tolerance which controls the maximum concavity of the final components, as input, along with other two fixed parameters for encoding the smoothness term of graph cut. We show that our method can generate comparable (sometimes noticeably better) segmentation in significant shorter time than the current state-of-art methods. Finally, we demonstrate applications of **CoRiSe**, including physically-based simulation, cage generation, model repair and mesh unfolding.

Keywords: Segmentation, Convexity, Shape Approximation

1. Introduction

In interactive applications, it is necessary to create simplified representations of a mesh to support various computationally intensive operations. In this work, we are interested in obtaining such simplified representations via decomposition. Taking deformation as an example, a mesh that has been decomposed into visually meaningful parts (e.g. head, torso, limbs) eases the process of creating deformation at semantic level. On the other hand, if a mesh is caged and partitioned by a set of convex shells, artists can use these shells to perform physically-based deformation efficiently. In many situations, both of these higher-level (semantic) and lower-level (physical or geometric) deformations are required. While it is ideal to keep both representations, it is also desirable to have a unified representation. An unified representation not only reduces space requirement but also allows deformations created at various semantic levels to be applied to the original mesh through a single approximation. Therefore, we recognize the need to have decomposition methods, such as the one proposed in this paper, that provide users both visual saliency and bounded geometric properties.

The first type of decomposition above is called **shape decomposition**. In the past decade, significant progress has

been made in producing high quality results [1, 2, 3, 4, 5, 6, 7]. Shape decomposition provides implicit shape approximation and is useful for shape analysis and recognition and semantic level shape editing and deformation. The second type of decomposition can be accomplished through the Approximate Convex Decomposition (ACD) [8]. ACD decomposes a 3D mesh into nearly convex parts. Unlike shape decomposition, ACD provides explicit approximation with bounded approximation errors and is therefore suitable for lower level processing. For example, Bullet physics library uses hierarchical ACD (HACD) [9] to speed up the collision detection and response computation of the non-convex shapes, and Muller et al. [10] proposed an interactive tool to generate approximate convex parts for speeding up dynamic fracture simulation.

The most challenging task in developing such a decomposition method stems from the fact that current convex decomposition and shape segmentation methods are computationally expensive and require different parameter settings for different shapes which are usually unintuitive and make processing a massive number of meshes in applications difficult. Recently, learning based segmentation methods [4, 5] are proposed to learn common parameters in an unsupervised or supervised manner, but the computation time of these methods is often intolerably high (minutes to hours), in particularly, for interactive applications.

Contribution In this paper, we will describe an efficient decomposition method, called **CoRiSe**. The only

*Corresponding author

Email addresses: gliu2@gmu.edu (Guilin Liu),
zxi@masonlive.gmu.edu (Zhonghua Xi), jmlie@cs.gmu.edu
(Jyh-Ming Lien)

Figure 1: An example of **CoRiSe** with concavity tolerance $\tau = 0.05$. (a) Input mesh. (b) Ten convex ridges (shown as translucent ellipsoids) are connected by **deep valleys** (shown as the geodesic paths). Formal definition of convex ridge and valley can be found in Section 4. (c) Left is the decomposition result of the first iteration; right is final decomposition (the second iteration). Note that colors are just used to distinguish different parts. (d) **CoRiSe** used for automatic cage generation. (e) The convex hulls of **CoRiSe** components can be used to speedup the computation in physically-based simulation.

user input parameter is a concavity tolerance which directly controls the approximation error. The novelty of **CoRiSe** comes from the idea of **convex ridge** which can be efficiently and robustly determined. Essentially, **CoRiSe** extracts nearly convex components of a 3D mesh P by separating each convex ridge from all the other convex ridges using graph cut. An example of convex ridge and **CoRiSe** decomposition is shown in Fig. 1. Formal definition of convex ridge can be found in Section 4. In addition, through the new concept of **residual concavity**, **CoRiSe** is insensitive to the parameter in graph cut and requires only a single user parameter: concavity tolerance. We will discuss residual concavity and graph cut in Section 5.

Using the Princeton Segmentation Benchmark, we show that **CoRiSe** generates noticeably better segmentation in significant shorter time than the existing methods [9, 1]. Finally, we demonstrate applications of **CoRiSe**, including physically-based simulation, cage generation, model repair and mesh unfolding in Section 6. Fig. 1 shows two of these applications.

2. Related Work

Many methods have been developed to decompose 3D mesh models. Comprehensive surveys can be found in [11, 12, 13]. In this section, we will review more recent works on shape segmentation and convex segmentation. After this short review, we will point out that the needs for developing a more intuitive and efficient method, such as **CoRiSe**.

Shape segmentation Many of the existing single-shape segmentation methods are based on clustering mesh faces, such as k -means clustering [14], fuzzy clustering [15], mean-shift clustering [16], and spectral clustering [17]. Shape features, such as geodesic distance, local concavity, curvature, have significant impact on these clustering methods. More advanced features include shape diameter function [6] that is a measure of the diameter of the object's volume in neighborhood of a point, Mumford-Shah model [3] that measures the variation within a segment and

continuous visibility feature [18] that uses more restricted visibility to better capture shape than the traditional line of sight visibility. Methods that are not clustering based do exist. For example, the method proposed by Wang et al. [19] uses the training segmentation of 2D projection images. Random cuts method [7] uses other different segmentation algorithms with various parameters to generate a collection of segmentations to find consistent cut positions.

Recently, we see more techniques using data-driven approach. These methods usually provide more consistent segmentations over a set of models [20] and produce segmentations that are more natural via machine learning approaches [4, 5]. However, these methods are computationally intensive (require hours of computation) and are not suitable for interactive use.

Primitive segmentation While many algorithms focus on decomposing a model into visually meaningful parts, other algorithms focuses on partitioning models into geometric primitives such as ellipsoids [21] and convex objects [8]. In other words, multiple primitives are used to jointly approximate the original shape.

In this paper, we are interested in producing nearly convex components. We found that many methods in the literature use again clustering (bottom-up) approach. Mamou and Ghorbel [9] proposed a method called Hierarchical Approximate Convex Decomposition (HACD) for 3D meshes. HACD iteratively clusters mesh facets by successively applying half-edge collapse decimation operations (see more detailed discussion in Section 6). Attene et al. [22] proposed to convert a model into tetrahedral mesh and then merge tetrahedra to form near convex components. However, this method requires human interaction, thus not suitable for segmenting a large number of models. The method proposed in [23] is based on a region growing approach controlled by convexity, compactness and part cost. Top-down approaches are rare. For example, Ghosh et al. [24] proposed a notion named **relative concavity** to model the concavity measure before and after every mesh cut. Nearly convex components are obtained by finding

the cuts that have largest relative concavities via dynamic programming. In many of these methods, several parameters are needed to be set to balance various features.

We would like to note that there are also methods that use convexity and concavity but do not produce segmentation with bounded convexity or concavity. For example, Au et al. [2] used concavity-aware field to form potential cuts and the final cuts are achieved by maximizing the cut score. Even though concavity is used, it doesn't have direct control of concavity for final components and several parameters and thresholds need to be set. Asafi et al. [25] defined the convexity using the line-of-sight and applied spectral clustering on the visibility matrix to achieve segmentation. van Kaick et al [1] applied merging on the initial result of [25] based on the Shape Diameter Function. However, since the pairwise visibility needs to be computed, the computation is time-consuming. Moreover, even though these two methods use convexity as clue for segmentation, they didn't directly control the convex/concave geometric property for the final components.

Concavity and its counterpart, convexity, have shown to be valuable for various decomposition methods. Intuitively, concavity of a polyhedron P measures how different P is from a convex object, which is typically the smallest convex object enclosing P , i.e. the convex hull $H(P)$ of P . Concavity can be measured globally from the difference between the volumes of P and $H(P)$ or the difference of their projections [26]. Concavity can also be measured locally for every point on the surface ∂P of P . Local concavity measures how far away a point on the surface of P is from the surface of $H(P)$, thus provides richer information to guide the decomposition process.

There have been several definitions of local concavity, such as curvature [2]. A more general definition of the concavity of a point p is the length of the shortest distances from $p \in \partial P$ to $H(P)$ without intersecting P . Due to the high computational complexity of 3D shortest-path problem, approximation is required. For example, Zimmer et al. [27] proposed to compute the shortest path by tessellating the space between ∂P and $\partial H(P)$ using constrained Delaunay triangulation (CDT). However, this solution required ∂P to be closed. Many other approaches resort to the Euclidean distance between the vertex and the convex hull in the outward normal direction of the vertex. It is clear that the association between ∂P and $\partial H(P)$ is usually complex and cannot be captured by the surface normals of P . This can result in inaccurate measurement of the concavity and lead to incorrect decomposition. Lien and Amato [8] proposed to determine the association by projecting the edges and then facets of $\partial H(P)$ to ∂P . Unfortunately, such an association between ∂P and $\partial H(P)$ is not always well defined as it usually depends on how P and $\partial H(P)$ are tessellated.

3. Overview

CoRiSe takes a single 3D mesh and a concavity tolerance parameter τ as input and decomposes this mesh into nearly convex components in a top-down fashion that includes three main steps:

Step 1. Build bridges and compute concavities Bridges are the convex hull edges, and the shortest geodesic path on the mesh connecting the end vertices of a bridge is called *valley*. Then, the concavity of a point in the valley is measured as the shortest-path distance to the bridge. Exact shortest geodesic path and shortest-path distance are computationally expensive thus we will discuss how we approximate them. We say that a valley is *deep* if the maximum concavity in the valley is greater than τ , otherwise the valley is *shallow*. It is important to remember that, throughout the entire segmentation process, concavity is only defined for points in the valleys and undefined elsewhere.

Step 2. Identify convex ridges In CoRiSe, a convex ridge consists of a group of convex-hull vertices connected *only* by shallow valleys. Intuitively, a convex ridge can be viewed as a *protruding* part of the mesh that is guaranteed to be convex enough. This implies that a continuous subset of the mesh containing two or more convex ridges must be too concave. Fig. 1(b) shows an example of 10 convex ridges.

Step 3. Partition using convex ridges CoRiSe segments the mesh by ensuring that each component in the segmentation contains only a single convex ridge. The segmentation process is bootstrapped by splitting each deep valley into shallow ones at the *valley separators*. Then CoRiSe applies graph cut using the region defined by valley separators as data term and edge angle and length as smoothness term to find optimal cut loops near these valley separators.

The above three steps are repetitively applied to the segmented part whose concavity is greater than τ .

4. Bridge, Valley and Convex Ridge

To approximate a mesh using convex shapes, we should establish relationships between the mesh and its convex hull. Let the input be a polyhedron $P = \{V_P, F_P, E_P\}$ composed of a list of vertices V_P and faces F_P connected by edges E_P . We assume that P consists of a single connected component. The convex hull of P is denoted as $H(P) = \{V_H, F_H, E_H\}$, where $V_H \subset V_P$. Let us consider a convex hull edge $e = \{a, b\} \in E_H$. Since e is the shortest Euclidean path connecting a and b , let us call e a *bridge*. Because a and b must also be the vertices of P , we can always find a path connecting a and b on ∂P . We call the shortest geodesic path connecting a and b on ∂P a *valley* underneath the bridge e_β .

A bridge e_β and its associated valley e_v form a 3D polygon. We call this polygon the *concavity polygon*, which plays an important role in many stages of our algorithm.

Figure 2: The left figure shows one bridge and its valley in 3D. The right part shows its projected concavity polygon in 2D. Concavities are measured in this 2D polygon using the shortest-path distance.

Figure 3: The concavity of the projected polygon with highly concave regions or self-intersections.

For example, we project the concavity polygon into a 2D space and only measure the concavity in the projected 2D space. More specifically, each point $p \in e_v = (a, b)$ is projected to

$$p' = (\vec{ap} \cdot \vec{ab}, \mathbf{d}(p, e_\beta)), \quad (1)$$

where $\vec{ap} \cdot \vec{ab}$ is the inner product of $\vec{ap}$ and $\vec{ab}$, and the function $\mathbf{d}$ defines the Euclidean distance between a point and a line segment. Fig. 2 shows an example of bridge, its valley and the projected 2D concavity polygon. Then, the concavity $C(p)$ of the point p is determined as the **shortest-path distance** between p' and the projected e_β . It is known that the shortest-path distance can be measured in $O(n)$ time for a polygon of size n .

It is possible that the projected polygon can self-intersect. However, self-intersection seldom happens; all projected polygons of models studied in this work are free of self intersection. However, when the projected polygon is self-intersecting or has large concave regions, we use a more sophisticated concavity measure proposed in [28], which decomposes the polygon into several near convex parts and determines a hierarchy of bridges. The concavity will be measured by the accumulated distance of the point-to-adjacent bridge distance and bridge-to-bridge distances. For example, in Figure 3, the concavity of p will be the distance sum of the distance between p and β_1 and the distance between β_1 and ab .

The concavity of a valley e_v (and its associated bridge e_β) is defined as the maximum concavity of its vertices, i.e., $C(e_v) = C(e_\beta) = \max_{p \in e_v} (C(p))$. Let us call e_v **deep valley** if $C(e_v)$ is larger than the tolerance, otherwise e_v is called a **shallow valley**. For the convenience of our

future discussion, we also classify bridges into **high** and **low** bridges (as measured by the **deck height** of bridge) if their corresponding valleys are deep and shallow, respectively.

4.1. Convex Ridge

Now let us formally define the convex ridge. A convex ridge R of P is a graph representing a subset of convex hull $H(P)$. We say a subset of $R \subset H(P)$ is a convex ridge if all bridges of R are shallow and there are no high bridges connecting two vertices in R . More specifically, a convex ridge R must satisfy the following two conditions:

$$\begin{aligned} & \forall e \in E_R, C(e) < \tau, \text{ and} \\ & \nexists e = (a \in V_R, b \in V_R), C(e) > \tau, \end{aligned} \quad (2)$$

where V_R and E_R are the vertices and edges of R , respectively. In short, a convex ridge can only have vertices connected by bridges whose valleys are shallow.

The convex ridges of a given mesh is constructed by iteratively greedy clustering the convex hull vertices while ensuring that the properties in Eq. 2 is maintained for each cluster. The clustering proceeds by collapsing low bridges sorted in an ascending order based on the length of the associated valleys. Each final cluster would become a convex ridge. The detailed procedure can be found in Algorithm 1 in Appendix. Figure 4 shows examples of convex ridges. We can also see that convex ridges are quite insensitive to surface noise in Figure 4. In fact, as long as shallow valleys remain shallow, all convex ridges will be unaffected.

4.2. Residual concavity

Residual concavity measures the concavity of a valley e_v after e_v is split into two subpaths at a vertex of e_v . Let $e_v = \{v_0, v_1, \dots, v_{n-1}\}$ be a deep valley, and let $e_v' = \{v_0, v_1, \dots, v_{k-1}\}$ be a prefix subset of e_v , where $k \leq n$. We further let $\hat{e}_v = \{v_{n-1}, v_{n-2}, \dots, v_0\}$ be the reverse of e_v . The residual concavity of e_v at the k -th vertex is then defined as:

$$RC(e_v, k) = \max \left\{ C(e_v^k), C(\hat{e}_v^{(n-k+1)}) \right\}. \quad (3)$$

Recall that the concavity is always measured in the 2D concavity polygon projected using Eqn.1. It is important to note that once the valley e_v is split, two or more new bridges must be formed to determine the concavity of e_v' 's sub-valleys. Therefore, at the first glance, computing RC seems to be time consuming, but we have shown that RC can be computed in linear time.

Lemma 4.1. *The computation of residual concavity $RC(e_v, k)$ takes time linear to the size of e_v , i.e. $O(n)$ for e_v with n edges.*

Proof. See details in the Appendix Proof 8. □

Figure 4: Convex ridges found in these models. Each ellipsoid represents a convex ridge. Even though several convex ridges in these examples concentrate around small compact regions, the convex ridges in (c) manifest various ridge-like shapes (regions that are locally with higher “elevation”). Figures 4(e), 4(f) and 4(g) show the convex ridges on the ant model with random noise added. The convex ridges of figures 4(e), 4(f) and 4(g) are all generated using $\tau=0.08$.

4.3. Shape of a valley

The residual concavity gives us a way to estimate the shape of a valley. We say a valley e_v is **V-shaped** if there exists a vertex v_k such that the residual concavity of e_v is less than the tolerance τ . Otherwise e_v is **U-shaped**. An example of a U-shaped valley can be found in Fig. 5. If e_v is U-shaped, we can always find the bottom of e_v as the sub-polygon bounded by the prefix and suffix of e_v that have residual concavities less than τ . In Fig. 5, we show that the bottom of the U-shaped valley is defined as (v_i, v_j) , where i and j are the maximum indices such that $C(e_v^i) < \tau$ and $C(\hat{e}_v^{(n-j)}) < \tau$, respectively; recall that $\hat{e}_v$ is the reverse of e_v . As we will discuss in Section 5, the shape of valley can be used to determine the number of vertices needed to separate each convex ridge in the wired representation to ensure that its concavity is less than τ .

5. Convex Ridge Separation

CoRiSe segments a mesh by ensuring that each component in the segmentation contains only a single convex ridge. Essentially, **CoRiSe** finds such segmentation in two steps: first, **CoRiSe** only focuses on the wireframe representation of the mesh that consists of vertices and edges of all convex ridges and deep valleys. An example of such a representation can be found in Fig. 1(b). For each deep valley, **CoRiSe** determines one or two key vertices (depending on whether it is a V- or U-shaped valley) that can separate each deep valley into at least two shallow valleys. In Section 5.1, we will discuss how **CoRiSe** decomposes this wireframe representation. Then, once the wireframe is decomposed, **CoRiSe** switches back to the original model and applies a two-way graph cut to find optimal cut loops near the valley separators. This step will be discussed in Section 5.2.

5.1. Valley separators

To separate the convex ridges in the wireframe presentation, **CoRiSe** finds **valley separators** for each deep valley and ensures that each convex ridge only connects the subset of the valley that has concavity less than the tolerance. For a V-shaped valley e_v , the valley separator is a single vertex that has the largest concavity in e_v . By splitting

Figure 5: A deep U-shaped valley e_v and its valley separators v_k and v_g . The sub-valleys e_v^i and e_v^j are subsets of e_v whose residual concavity is less than the valley tolerance τ . The valley separators $v_k \in e_v^i$ and $v_g \in e_v^j$ minimize the residual concavities $RC(e_v^i, k)$ and $RC(e_v^j, g)$.

at the vertex with the largest concavity, we can guarantee that the two new sub-valleys of e_v are shallow.

For a U-shaped valley e_v , things are a bit more complex. Basically, **CoRiSe** splits e_v into three sub-valleys and guarantees that two sub-valleys incident to two convex ridges are shallow even though the bottom of e_v may not be. Let $e_b = (a, b)$ be a bridge whose valley e_v is U-shaped. Using the residual concavity defined in Eqn. 3 in Section 4.2, we can find two vertices v_i and v_j such that e_v^i , a sub-valley between a and v_i , and $\hat{e}_v^j$, a sub-valley between v_j and b , are shallow. See Fig. 5 for the illustration of e_v^i and $\hat{e}_v^j$. Our goal is to define two valley separators in e_v^i and $\hat{e}_v^j$, respectively. It is true that we can use v_i and v_j as the valley separators and still guarantee that e_v^i and $\hat{e}_v^j$ are shallow. However, the locations of v_i and v_j are quite arbitrary in most cases. Therefore, the valley separators for a U-shaped valley e_v are defined as a couple of vertices (v_k, v_g) :

$$(v_k, v_g), \text{ where } k = \arg \min_{k \in V} RC(e_v^i, k), g = \arg \min_{g \in V} RC(\hat{e}_v^j, g).$$

From Fig. 5, we can see that the valley separators v_k and v_g of e_v are identified at the locations of large concavity if the valley is split at v_i and v_j . Once the valley separators are identified for all deep valleys, all convex ridges are separated and only attached to shallow valleys. Next, **CoRiSe** switches back to the original representation and segments the mesh using graph cut.

Figure 6: Simplified graph cut. black nodes and blue nodes represent the must-link regions. They have infinite data term associated with source or sink. Red nodes represent the fuzzy region.

5.2. Convex ridge separator

Graph cut [29] is a powerful optimization tool that has been proven to be successful in segmentation. In the previous section, we have found separators for deep valleys. And the valley separators have roughly defined where the

cuts should be. However, since the valley separators are merely feature vertices capturing concavity constraints, separating convex ridges needs close cut loops. We formulate finding the cut loops as a simplified graph cut problem. We say simplified graph cut because the data term value is either infinity or zero, which corresponds must-link region and fuzzy region. The smoothness term is defined using dihedral angle and edge length. Fig. 6 shows the simplified graph cut formulation. In the rest of this section, we will discuss the data term and smoothness term in detail.

Data Term To formulate the graph cut problem, we introduce two terms: *must-link region* and *fuzzy region*. *Must-link region* would have infinite data term values associated with the *source* or *sink*. *Fuzzy region* would have zero data term values. For each iteration of graph cut, a facet belongs to either *must-link region* or *fuzzy region*. In other words, *must-link region* plus *fuzzy region* will cover the whole model. Intuitively, each convex ridge has a must-link region which is a set of faces that must be segmented with the convex ridge. To get the *must-link region* and *fuzzy region*, we will introduce another term named *potential region*. We call it *potential region* because a facet belonging to *potential region* has the chance to become either part of *must-link region* or part of *fuzzy region*. The ways to determine these regions are described as following: **Potential Region.** We will collect a *potential region* for each convex ridge R_k . This potential region is the union of the facets collected by R_k 's all the sub-valleys using geodesic distance. The way of how each sub-valley collects its facets and doing the union are described as following:

1. Collecting using Geodesic Distance: for each sub-valley e_v^i (e_v^i is created from valley separator as described in Section 5.1) incident to R_k , we let T_i be a set of facets whose geodesic distance to the end point of e_v^i in R_k is less than the path length of e_v^i .

2. Union: from T_i , we can determine a set U_{R_k} of facets that forms a superset containing R_k . The set U_{R_k} is the

Figure 7: Left: the potential region (including all the black and red regions). Right: must-link region and fuzzy region. In the right part, the convex ridges are differentiated by colors (4 in total). The colored circles (dots) are the end vertices of the deep valleys. The crossings are the detected valley separators. The red lines (region) are the fuzzy region. The colored lines attached to circles with same color are must-link region attached to corresponding convex ridge.

potential region of convex ridge R_k . More specially,

$$U_{R_k} = \bigcup_{1 \leq i \leq n} T_i^i, \quad (4)$$

where n is the number of sub-valleys incident to R_k .

Potential Region's Overlapping Issue. Two neighboring convex ridges R_k and R_j are likely to compete on certain facets. In this case, U_{R_k} and U_{R_j} overlap. The overlapping region of I_{R_k} with other convex ridges is then defined as:

$$I_{R_k} = \bigcap_{\forall j \neq k} U_{R_k} \setminus U_{R_j}. \quad (5)$$

Must-link and Fuzzy Regions. With U_{R_k} and I_{R_k} , we are ready to define the must-link and fuzzy regions. The must-link region C_{R_k} of R_k is simply U_{R_k} with facets in I_{R_k} removed,

$$C_{R_k} = U_{R_k} \setminus I_{R_k}. \quad (6)$$

Then the fuzzy region is the union of the facets that do not belong to any convex ridge's *must-link region*.

$$F_R = F \setminus \bigcup_k C_{R_k}, \quad (7)$$

where F is the facets of the mesh.

We give an example here. In the left part of Fig 7, the black segments show the potential region of a convex ridge. The black vertices are the vertices of a convex ridge. The crossing vertices are the valley separators. The right part of Fig 7 shows the must-link region and fuzzy region. The red segments represent the fuzzy region. Except the red, each of other colors stands for the must-link region for each convex ridge. There are three parts of fuzzy region (red regions). The left part corresponds to the area that is not covered by deep valley's potential region. The middle part is part of fuzzy region because it lies in the middle part of a U-shape valley and it is not within the potential region of either of end vertices from two convex ridges. The right part becomes fuzzy region due to the fact that it has been covered by two convex ridges.

Fig 8 shows the must-link regions on the 3D models. The filtered out regions are the fuzzy regions.

Figure 8: The must-link regions connected by the deep valleys.

Smoothness Term The smoothness term evaluates the similarity/compatibility between two adjacent triangle facets. In this paper, we use the same definition of smoothness as that in [15].

A parameter α is introduced to balance the importance of dihedral angle and edge length. The definition is as below: Let v_i and v_j be two adjacent vertices, namely are the two end vertices of an edge (denoted as e_{ij}) on the mesh. And let f_i and f_j to be the two triangle facets sharing this edge. Let θ_{ij} be the dihedral angle between f_i and f_j , $|e_{ij}|$ be the length of e_{ij} . Let μ_{θ} be a function of θ_{ij} such that μ_{θ} is a positive small number (the second parameter, fixed to be 0.3) for concave edges and 1 for convex edges. Then the smoothness term defined in [15] is:

$$w_{ij} = \alpha \left(\frac{\mu_{\theta}}{\mu_e} \right) + (1 - \alpha) \frac{|e_{ij}|}{\mu_e}, \quad (8)$$

where μ_{θ} and μ_e are the average values of θ and edge length, respectively, of the entire mesh. In this paper, we set $\alpha = 0.8$ (the third parameter, fixed) in all experiments.

Segmentation using Graph Cut We solve $n-1$ times (n is the number of convex ridges) binary graph cut optimization using the max-flow algorithm. The order of $n-1$ graph cut is determined by the bounding box's diameter of convex ridges' vertices. We separate one convex ridge at one time using the graph cut. At each graph cut iteration i , we will assign infinite data term values associated with **source** for the facets in its **must-link region** C_{R_i} , and assign infinite data term values associated with **sink** for the facets in other convex ridges' **must-link region** $\forall j \neq i, C_{R_j}$. Facets in **fuzzy region** don't have data term with neither **source** nor **sink**. The smoothness term will be defined for any pair of neighboring facets using the Equation 8. We separate this convex ridge from the model and the rest of model will be used for applying region for next binary graph cut. Segmentation is performed iteratively until all components have smaller concavity than τ . Fig 1 (c) shows the first iteration decomposition result on the left and second (final) iteration decomposition result on the right.

Discussion One may argue that there are three issues in the above formulation of graph cut problem. 1. Why use must-link region? The must-link region is defined using valley separators, which roughly defines where the cut should be and satisfied the concavity constraints. The must-link region can fully utilize the rough cut positions proposed by valley separators. 2. Why use $n-1$

Figure 9: Approximate shapes using the convex hulls of decomposed parts. The results in this figure are generated by CoRiSe using models in the Princeton Segmentation Benchmark.

binary graph-cut? Another option is to k -way graph cut. However, choosing a right k is important for k -way graph cut. A proper k would be the number of final components. However, the problem is the number of final components is unknown and the number of convex ridges is not the real number corresponding to the number components. Some components may haven't been detected by the current convex hull. The real number should be $\geq n$. We apply $n-1$ times graph cut and postpone handling of the issue of undetected components in the further iteration of decomposition, because these components may be detected after we have already cut some components out of the model in previous iteration. 3. How to decide the order of cutting which convex ridge and how to deal with overlapping labeling issues? The order of applying graph cut is based on the heuristic that we handle the small components first (in terms of the bounding box's diameter of convex ridges' vertices). In for each time of graph cut, if we apply on the same region, there may be one issue that a facet has been labeled as "source" for several times. In this paper, after we apply one time of graph cut, we separate one convex ridges from the model and the rest of the model would be used as the applying region for further graph cut. In this way, we would be able to avoid the issue of overlapping labeling issues.

6. Experimental Results

We have implemented CoRiSe in C++. In order to evaluate CoRiSe, we compare CoRiSe extensively to existing approximate convex decomposition methods HACD [9] that is available in the Bullet Physics Library and WC-Seg [1]. We also compare CoRiSe with shape decomposition methods using the Princeton 3D Mesh Segmentation Benchmark [13], which includes 380 surface meshes of 19 different object categories. All the models are scaled to have a 1.0 radius bounding sphere. The quality of the decomposition is measured via the Rand Index that estimates the similarity by measuring facet pairwise label difference between the segmentations generated by CoRiSe and those generated by humans and the state-of-art methods [6, 7, 3, 1]. All experimental results are collected on a workstation with two Intel Xeon 2.30GHz CPUs and 32

GB memory. HACD and CoRise are implemented using C++ and WCSeg uses MATLAB.

6.1. Comparison of Approximation

Figure 10: Shape approximation using the convex hulls of all components in the decompositions from **CoRise** and HACD.

Given a concavity tolerance, we compare **CoRise** with HACD [9] and WCSeg [1] based on (1) the number of decomposition components, (2) decomposition time, and (3) approximation accuracy. HACD first simplifies the model using Quadric Error Metrics (QEM) and then hierarchically merges facets and ensures that the concavities of all clusters are lower than the tolerance concavity. WCSeg creates initial segmentation using line-of-sight concavity/convexity measurement and then uses Shape Diameter Function [6] to merge the initial segmentations. Table 1 shows the decomposition size, computation time, and the ratio between the volume of the convex hull approximation and the volume of the original model.

From Table 1, we can observe that, on average, **CoRise** generates fewer components than HACD and WCSeg. In addition, **CoRise** is significantly faster than HACD and WCSeg. This is because HACD requires many concavity evaluations in the bottom-up approach clustering process. Note that even though it is not totally fair to directly compare the computation time of **CoRise** and WCSeg shown in Table 1 due to WCSeg’s matlab implementation, doing spectral clustering that involves solving the eigen-decomposition of a large pairwise matrix (whose size is square to the number of facets) would still be a bottleneck using the faster C++ implementation. Without simplifying models, HACD can even take more than several hours to decompose a model with around 10,000 vertices.

Table 1 also shows the comparison of volume ratios. The volume ratio is defined as $\text{vol}(\cup_i(CH_i))/\text{vol}(P)$, where $\text{vol}(\cup_i(CH_i))$ is the volume of the union of all convex hulls and $\text{vol}(P)$ is the volume of the input mesh. From Table 1, HACD has smallest average volume ratios. However, the result is biased because HACD simplifies the input model before decomposition, thus these convex hulls created by HACD do not guarantee to enclose the original model. This may also produce noticeable penetration if these convex hulls are used in collision detection. We can also see that the approximation volume of **CoRise** is slightly smaller than WCSeg. However, because WCSeg produces more segments, we can say that **CoRise** is more

Figure 11: WCSeg vs. **CoRise**. Top row from left to right: WCSeg decomposition, **CoRise** decomposition and convex hulls of **CoRise** components. Bottom row from left to right: WCSeg decomposition, convex hulls of WCSeg components, **CoRise** decomposition, and convex hulls of **CoRise** components.

effective (use less components to provide tighter approximation). This is especially true in the human and octopus category. Fig. 11 shows an example of the segmentations of a human model and an octopus model from both **CoRise** and WCSeg. The main differences between **CoRise** and WCSeg stem from how concavity is used for segmentation. **CoRise** cares about the approximation error, thus ignores the fingers of the human model because the concavity between two fingers is smaller than the concavity tolerance. On the other hand, the legs of the octopus model are bended and **CoRise** cuts the legs into several components to satisfy the concavity tolerance while WCSeg merges them together.

6.2. Comparison using benchmark dataset

We should first note that **CoRise** aims at resolving the concavity of a model and is not designed for semantic segmentation. For example, in Fig. 17, the bird is decomposed into over 30 parts which is much more than that of other shape segmentation. This will certainly influence the evaluation if **CoRise** is compared to segmentations created by human. However, we feel it is important to compare **CoRise** to shape segmentation methods since, after all, convexity is one of the important properties used in shape segmentation, and we also believe that a visually meaningful decomposition can provide visually convincing simulation.

Fig. 12 shows the Rand Index (RI) scores obtained from seven methods on the same benchmark. Although

Table 1: Statistical results of the number of final components, computation time and the volume ratio which is approximations' volume compared to that of the original model on the Princeton 3D Mesh Segmentation Benchmark. All results are obtained using concavity tolerance = 0.1.

Category	# of components			Time (sec)			Volume Ratio		
	HACD	WCseg	CoRiSe	HACD	WCseg	CoRiSe	HACD	WCseg	CoRiSe
human	12.10	18.85	9.85	84.80	677.92	2.33	1.23	1.24	1.31
cup	18.90	5.50	4.55	372.09	1,362.30	1.25	2.66	3.02	3.15
glass	6.40	8.90	5.53	433.84	416.09	0.63	1.69	1.90	1.75
airplane	5.20	7.70	7.55	1,103.85	431.91	1.07	1.30	1.25	1.37
ant	13.45	11.25	14.50	223.99	413.63	1.76	1.11	1.17	1.14
chair	9.50	16.35	10.64	297.72	1,021.50	1.86	1.67	1.56	1.45
octopus	14.40	9.10	14.45	903.44	599.90	1.98	1.23	1.83	1.27
table	6.00	5.55	6.05	3,385.97	2,762.60	1.25	1.33	1.68	1.42
teddy	6.80	7.45	7.05	1,548.05	1,026.30	2.16	1.05	1.05	1.07
hand	9.10	9.20	9.00	205.92	670.00	1.67	1.14	1.16	1.27
plier	5.20	6.50	5.89	393.18	368.90	0.88	1.30	1.37	1.33
fish	5.05	6.25	5.20	2,969.80	546.04	0.84	1.12	1.16	1.18
bird	5.55	11.10	6.39	547.89	515.03	0.73	1.36	1.26	1.35
armadillo	15.30	20.10	14.90	56.89	1,914.80	5.46	1.12	1.12	1.19
bust	9.30	8.70	7.80	118.18	2,285.80	4.50	1.05	1.08	1.11
mech	4.45	3.70	3.65	1,286.43	2,587.30	1.04	0.87	1.02	1.02
bearing	5.47	3.00	1.58	1,095.92	1,188.90	0.25	0.95	1.08	1.17
vase	9.80	5.65	5.15	1,059.17	1,247.60	1.63	1.06	1.10	1.12
four-leg	12.95	14.05	12.25	61.85	614.97	1.86	1.20	1.23	1.25
Average	9.21	9.42	8.00	849.95	1,086.90	1.75	1.29	1.38	1.36

Figure 12: Rand Index (RI) comparison on the Princeton Benchmark data. **human**: human cut; **PT**: CoRiSe with a τ for each model; **CT**: CoRiSe with a τ for each category; **CR9**: CoRiSe with $\tau = 0.09$; **CR10**: CoRiSe with $\tau = 0.1$; Additional comparisons evaluated using other metrics such as consistency error, cut discrepancy, Hamming distance all show similar trend as RI and can be found in the supplementary materials.

CoRiSe performs worse than the learning-based approach, CoRiSe outperforms some single-model based methods [6, 7] that require more user parameters, such as number of clusters. If CoRiSe can choose a concavity tolerance for each category, its performance is comparable to [1] which also has several parameters and thresholds. Moreover, if CoRiSe chooses concavity tolerance for each model individually, just as other methods selecting component size for each model, its RI score is lower than all single-model methods in Fig. 12.

Compared with other shape-segmentation methods, CoRiSe has three major advantages: First, CoRiSe provides the bounded convexity that can be used by broader applications. Second, CoRiSe has less number of param-

eters, while others require several user parameters for each model to specify component number, soft-clustering number, or smooth factors. Table 2 summarizes the number of parameters of each method. Finally, CoRiSe is efficient without simplification as pre-processing.

Although it is true that only a small subset of concavity tolerances produce semantic decomposition of a given shape, those that do create semantic decompositions reflect the geometric property of the shape. We encourage readers to look at our supplementary materials which show the RI, HD, CE and CD scores for a range of concavity tolerances.

Table 2: Number of parameters of segmentation methods (including fixed parameters). **WCseg**: the parameters will be used for encoding pairwise matrix, computing Shape Diameter Function and post-merging; **RC**: rand cut uses the statistics of other approaches' decomposition results. The parameter number will be more the sum of those approaches; **SDF**: Shape Diameter Function method ≥ 5 parameters (number of rays per facet; cone angle; component number of GMM; two fixed parameters in encoding smoothness term same as CoRiSe; other parameters); **HACD**: many parameters in mesh simplification and thresholds for angles and edges during decomposition; **CORISE**: 3 parameters (the first is the concavity tolerance; the other two are the fixed parameters in encoding smoothness term).

WCseg	RC	SDF	HACD	CoRiSe
> 5	> 5	≥ 5	7	3

7. Applications

CoRiSe is designed to approximate a 3D shape with

Figure 13: Simulation created using the convex hulls of **CoRiSe**.

a set of convex objects. In this section, we demonstrate several applications using this type of approximation.

7.1. Physically-based Simulation

Physics simulation libraries, such as Bullet and Box2D, use multiple convex shapes to approximate the original non-convex shape. Approximating a shape with bounded concavity error allows computations, such as collision response and penetration depth, become much more efficient. In these applications, we care about the number of approximation components and the approximation error. Exact convex decomposition has no approximation error, however, the number of components is usually prohibitively huge, thus the computation, e.g., penetration depth, is inherently expensive. It is therefore desirable to have smaller number of components while the approximation error is bounded. From the comparisons between HACD and **CoRiSe** shown in Table 1, **CoRiSe** provides following advantages: 1). **CoRiSe** guarantees to enclose the original model while HACD does not; 2). **CoRiSe** is much faster than HACD; 3). **CoRiSe** produces smaller segmentation given the same concavity tolerance.

We have successfully integrated **CoRiSe** with the Bullet library. This allows us to generate convincing physically-based simulations using only the convex hulls of **CoRiSe** decompositions. We encourage the readers to view the supplementary video.

7.2. Cage-Based Deformation

Cage based deformation technique aims to be an easy-to-use tool for graphics modelling, texturing and animation. The advantage of cage-based space deformation is its simplicity, flexibility and efficiency. Most cages in cage-based deformation are created manually, and the process can be tedious. Xian et al. [30] proposed an automatic cage generation technique based on iteratively splitting the bounding boxes. However, their method requires mesh voxelization that is both time consuming and, more importantly, ignores important structural features of the input shape. For example, their cage usually misses vertices near

Figure 14: A cage created from **CoRiSe** decomposition.

Figure 15: (a) Noisy mesh with missing facets. The bottom figure provides a close-up view. (b) **CoRiSe** results. (c) Top: Convex hulls of 16 segmented parts in (b). Bottom: Convex hulls of HACD segmentation (138 components). All results are obtained using concavity tolerance 0.1.

the concave regions of the input mesh. Moreover, their method is not suitable for interactive applications as for a shape with only 10k vertices, their method can take several minutes. **CoRiSe** generates a cage in almost real-time. We first use **CoRiSe** to generate a nearly convex decomposition of the shape. Then convex hulls are computed for each part. These convex hulls are slightly expanded and simplified. The final cage is the union of the convex hulls. Fig. 14 shows a cage created for the Armadillo model and its deformation using Green coordinates. More examples can be found in the supplementary materials.

7.3. Model Repair

All the models used in Table 1 are watertight. However, many digitized models have degenerated features and may require mesh repair. The convex hulls of **CoRiSe** decomposition provide a convenient way to generate a watertight representation. In Fig. 15, we show a mesh that has 45% of the facets removed, and **CoRiSe** successfully decomposes the mesh and produces better approximation than HACD does. In fact, **CoRiSe** will be able to produce similar results as long as the mesh edge connectivity has not been significantly changed. For models with many holes, **CoRiSe** can get help from methods such as [31] that can be used to ensure that the projected geodesic path of the bridges is still well defined.

Figure 16: Left: Segmentation result using **CoRiSe**. Right: Paper craft created by unfolding/folding each nearly convex component.

7.4. Origami

Origami has many applications in material engineering, manufacturing and packaging, such as solar panel packing and self-folding origami. Edge unfolding a convex polyhedron to a non-overlapping planer polygon (called a net) by cutting along its edges remains an open problem [32], in particular for non-convex shapes, which usually need to be segmented then a net can be found for each component [34, 35]. Using **CoRiSe** to decompose a non-convex shape into nearly convex components can significantly reduce the computational time of finding the nets and make folding easier [36]. Fig. 16 shows an example of this approach.

8. Conclusion, Limitations and Future Work

In summary, in this paper we proposed a 3D decomposition method that produces components with bounded concavity. Our method called **CoRiSe** meets the needs of real-time simulations and computer games better than the existing methods. We also showed that the decomposition results can produce semantically meaningful decomposition with low computation cost when proper concavity tolerances are given. The decomposition is achieved by identifying *convex ridges* and then solving a graph cut optimization problem formulated with *valley separators* and *must-link regions*. Comparing **CoRiSe** with other methods on the public benchmark dataset, we show that **CoRiSe** can generate results close to manual decomposition. Comparing **CoRiSe** with other approximate convex decomposition methods, we show that **CoRiSe** is significantly more efficient.

Known Limitations. Each of the major step of **CoRiSe** is quite robust. For example, the convex ridges are always identified at expected locations for all the models that we tested (see the supplementary material for more examples). A main limitation of **CoRiSe** is that its semantic meaning of **CoRiSe** components is significantly influenced by the concavity tolerance τ . When τ is too small, small components are produced and the semantic meaning of these components diminishes. For example in Fig. 17(b) and (c), even though **CoRiSe** still provides a tight approximation of the initial model using a relatively small concavity tolerance (0.05), many small components do not have significant meanings.

Another limitation of **CoRiSe** is that **CoRiSe** may not decompose a cup-like shape if none of the valleys reach the bottom of the cup's concavity (see Fig. 17(d)). This is caused by the fact that a valley is the (approximated) shortest geodesic path. We believe that this can be addressed by ensuring that the vertex p with largest concavity of a valley has the locally maximum concavity. If this is not the case, then we iteratively descend p until a point p' with locally maximum concavity is reached. A new valley is then computed by enforcing it to pass through p' .

Figure 17: (a) **CoRiSe** result of a bird. (b) **CoRiSe** result using $\tau = 0.05$. (c) The corresponding nearly convex approximation of (b). (d) **CoRiSe** result of a cup.

Acknowledgement

This work is in part supported by US NSF IIS-096053, CNS-1205260, EFRI-1240459, AFOSR FA9550-12-1-0238 and in part by the Korean Federation of Science and Technology Societies (KOFST) grant funded by the Korean government.

References

- [1] O. van Kaick, N. Fish, Y. Kleiman, S. Asafi, D. Cohen-Or, Shape segmentation by approximate convexity analysis, ACM Trans. on Graphics to appear.
- [2] O.-C. Au, Y. Zheng, M. Chen, P. Xu, C.-L. Tai, Mesh segmentation with concavity-aware fields, Visualization and Computer Graphics, IEEE Transactions on 18 (7) (2012) 1125 – 1134.
- [3] J. Zhang, J. Zheng, C. Wu, J. Cai, Variational mesh decomposition, ACM Transactions on Graphics (TOG) 31 (3) (2012) 21.
- [4] E. Kalogerakis, A. Hertzmann, K. Singh, Learning 3d mesh segmentation and labeling, ACM Transactions on Graphics (TOG) 29 (4) (2010) 102.
- [5] Q. Huang, V. Koltun, L. Guibas, Joint shape segmentation with linear programming, in: ACM Transactions on Graphics (TOG), Vol. 30, ACM, 2011, p. 125.
- [6] L. Shapira, A. Shamir, D. Cohen-Or, Consistent mesh partitioning and skeletonisation using the shape diameter function, The Visual Computer 24 (4) (2008) 249 – 259.
- [7] A. Golovinskiy, T. Funkhouser, Randomized cuts for 3d mesh analysis, in: ACM Transactions on Graphics (TOG), Vol. 27, ACM, 2008, p. 145.
- [8] J.-M. Lien, N. M. Amato, Approximate convex decomposition of polygons, Comput. Geom. Theory Appl. 35 (1) (2006) 100 – 123. doi:http://dx.doi.org/10.1016/j.comgeo.2005.10.005.
- [9] K. Mamou, F. Ghorbel, A simple and efficient approach for 3d mesh approximate convex decomposition, in: Image Processing (ICIP), 2009 16th IEEE International Conference on, 2009, pp. 3501 – 3504.

- [10]M. Müller, N. Chentanez, T.-Y. Kim, Real time dynamic fracture with volumetric approximate convex decompositions, ACM Transactions on Graphics (TOG) 32 (4) (2013) 115.
- [11]M. Attene, S. Katz, M. Mortara, G. Patané, M. Spagnuolo, A. Tal, Mesh segmentation-a comparative study, in: Shape Modeling and Applications, 2006. SMI 2006. IEEE International Conference on, IEEE, 2006, pp. 7 – 7.
- [12]A. Shamir, A survey on mesh segmentation techniques, in: Computer graphics forum, Vol. 27, Wiley Online Library, 2008, pp. 1539 – 1556.
- [13]X. Chen, A. Golovinskiy, T. Funkhouser, A benchmark for 3d mesh segmentation, ACM Transactions on Graphics (TOG) 28 (3) (2009) 73 – 73.
- [14]S. Shlafman, A. Tal, S. Katz, Metamorphosis of polyhedral surfaces using decomposition, in: Computer Graphics Forum, Vol. 21, Wiley Online Library, 2003, pp. 219 – 228.
- [15]S. Katz, A. Tal, Hierarchical mesh decomposition using fuzzy clustering and cuts, ACM Trans. Graph. 22 (3) (2003) 954 – 961. doi:http://doi.acm.org/10.1145/882262.882369.
- [16]H. Yamauchi, S. Lee, Y. Lee, Y. Ohtake, A. Belyaev, H. Seidel, Feature sensitive mesh segmentation with mean shift, in: Shape Modeling and Applications, 2005 International Conference, IEEE, 2005, pp. 236 – 243.
- [17]R. Liu, H. Zhang, Mesh segmentation via spectral embedding and contour analysis, in: Computer Graphics Forum, Vol. 26, Wiley Online Library, 2007, pp. 385 – 394.
- [18]G. Liu, Y. Gingold, J.-M. Lien, Continuous visibility feature, in: 28th IEEE Conference on Computer Vision and Pattern Recognition (CVPR), IEEE, Boston, MA, 2015, pp. 1182 – 1190.
- [19]Y. Wang, M. Gong, T. Wang, D. Cohen-Or, H. Zhang, B. Chen, Projective analysis for 3d shape segmentation, ACM Transactions on Graphics (TOG) 32 (6) (2013) 192.
- [20]A. Golovinskiy, T. Funkhouser, Consistent segmentation of 3d models, Computers & Graphics 33 (3) (2009) 262 – 269.
- [21]L. Lu, Y. Choi, W. Wang, M. Kim, Variational 3d shape segmentation for bounding volume computation, in: Computer Graphics Forum, Vol. 26, Wiley Online Library, 2007, pp. 329 – 338.
- [22]M. Attene, M. Mortara, M. Spagnuolo, B. Falcidieno, Hierarchical convex approximation of 3d shapes for fast region selection, in: Computer Graphics Forum, Vol. 27, Wiley Online Library, 2008, pp. 1323 – 1332.
- [23]V. Kreavoy, D. Julius, A. Sheffer, Model composition from interchangeable components, in: Computer Graphics and Applications, 2007. PG' 07. 15th Pacific Conference on, IEEE, 2007, pp. 129 – 138.
- [24]M. Ghosh, N. M. Amato, Y. Lu, J.-M. Lien, Fast approximate convex decomposition using relative concavity, Computer-Aided Design 45 (2) (2013) 494 – 504.
- [25]S. Asafi, A. Goren, D. Cohen-Or, Weak convex decomposition by lines-of-sight, in: Computer Graphics Forum, Vol. 32, Wiley Online Library, 2013, pp. 23 – 31.
- [26]Z. Lian, A. Godil, P. L. Rosin, X. Sun, A new convexity measurement for 3d meshes, in: Computer Vision and Pattern Recognition (CVPR), 2012 IEEE Conference on, IEEE, 2012, pp. 119 – 126.
- [27]H. Zimmer, M. Campen, L. Kobbelt, Efficient computation of shortest path-concavity for 3d meshes, in: Computer Vision and Pattern Recognition (CVPR), 2013 IEEE Conference on, IEEE, 2013, pp. 2155 – 2162.
- [28]G. Liu, Z. Xi, J.-M. Lien, Dual-space decomposition of 2d complex shapes, in: 27th IEEE Conference on Computer Vision and Pattern Recognition (CVPR), IEEE, Columbus, OH, 2014.
- [29]Y. Boykov, O. Veksler, R. Zabih, Fast approximate energy minimization via graph cuts, Pattern Analysis and Machine Intelligence, IEEE Transactions on 23 (11) (2001) 1222 – 1239.
- [30]C. Xian, H. Lin, S. Gao, Automatic cage generation by improved obbs for mesh deformation, The Visual Computer 28 (1) (2012) 21 – 33.
- [31]K. Crane, C. Weischedel, M. Wardetzky, Geodesics in Heat: A New Approach to Computing Distance Based on Heat Flow, ACM Trans. Graph. 32.
- [32]J. O'Rourke, Unfolding polyhedra (2008).
- [33]W. Schlickerrieder, Nets of polyhedra, Master's thesis, Technische Universität Berlin (1997).
- [34]R. Straub, H. Prautzsch, Creating optimized cut-out sheets for paper models from meshes, Tech. rep., Karlsruhe Institute of Technology (2011).
- [35]S. Takahashi, H.-Y. Wu, S. H. Saw, C.-C. Lin, H.-C. Yen, Optimized topological surgery for unfolding 3d meshes, in: Computer Graphics Forum, Vol. 30, Wiley Online Library, 2011, pp. 2077 – 2086.
- [36]Z. Xi, J.-M. Lien, Continuous unfolding of polyhedra - a motion planning approach, in: 2015 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), Hamburg, Germany, 2015, pp. 3249 – 3254.

Appendix

Appendix A. Algorithm for building the convex ridges

The following algorithm shows the details of how we compute convex ridge by iteratively collapsing the convex hull edges (i.e. bridges) whose valleys are shallow.

Algorithm 1 Build Convex Ridges

- 1: Sort bridges B in an ascending order by the valleys length
 - 2: $V \leftarrow$ the end vertices of all bridges B
 - 3: Create $||V||$ clusters each of which contains a vertex in V
 - 4: **for each** shallow bridge $b_i = (a_1, a_2)$ in B **do**
 - 5: $g_1 \leftarrow$ the cluster contains a_1
 - 6: $g_2 \leftarrow$ the cluster contains a_2
 - 7: **if** $\$e = (b_1 \in g_1, b_2 \in g_2)$ such that $C(e) \geq \tau$ **then**
 - 8: merge g_1 and g_2
 - 9: Report each cluster as a convex ridge
-

Appendix B. Proof of Linear Time Complexity of Residual Concavity

Lemma 8.1. *The computation of residual concavity $RC(e_v, k)$ takes time linear to the size of e_v , i.e. $O(n)$ for e_v with n edges.*

Proof. Similar to our previous definition, the concavity of a prefix e_v^k is defined as the shortest-path distance between e_v^k and edges of $H(e_v^k)$ inside the 2D concavity polygon (made of edges of e_β and e_v). Due to the special property of the concavity polygon P_v , we can determine if a convex hull edge $e = (v_i, v_j)$ of $H(e_v^k)$ is inside P_v by the order of v_i and v_j . That is, if $i < j$, then $e = (v_i, v_j) \in P_v$. Since e_v^k is a polyline, $H(e_v^k)$ can also be constructed in linear time $O(k)$, which makes the computation of residual concavity $RC(e_v, k)$ linear to the size of e_v , i.e. $O(n)$. $\square$