

Nearly Convex Segmentation of Polyhedra Through Convex Ridge Separation

Guilin Liu
gliu2@gmu.edu

Zhonghua Xi
zxi@gmu.edu

Jyh-Ming Lien
jmlien@cs.gmu.edu

Technical Report GMU-CS-TR-2015-3

Abstract

Decomposing a 3D model into approximately convex components has gained more attention recently due to its ability to efficiently generate small decomposition with controllable concavity bound. However, current methods are computationally expensive and require many user parameters. These parameters are usually unintuitive and add unnecessary obstacles in processing a large number of meshes and meshes that are generated online in 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 requires only a single user parameter: concavity tolerance. We show that our method can generate 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 and model repairing.

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. A 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




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) Final decomposition. (d) CORISE used for automatic cage generation. (e) The convex hulls of CORISE components can be used to speedup computation in physically-based simulation.

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 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, 7, 6]. Finally, we demonstrate applications of CORISE, including physically-based simulation, cage generation and model repair 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, and Mumford-Shah model [3] that measures the variation within a segment. Methods that are not clustering based do exist. For example, the method proposed by Wang et al. [18] 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 [19] 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 [20] and convex objects [8].

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. [21] 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 [22] is based on a region growing approach controlled by convexity, compactness and part cost. To-down approaches are rare. For example, Ghosh et al. [23] 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. [24] 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 [24] 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 [25]. 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, Zimimer et al. [26] 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 $e = \{u, v\} \in E_H$. Since e is the shortest Euclidean path connecting u and v , let us call e a *bridge*. Because u and v must also be the vertices of P , we can always find a path connecting u and v on ∂P . We call the shortest geodesic path connecting u and v on ∂P a *valley* underneath the bridge e_β .




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.

A bridge e_β and its associated valley e_ν form a 3D polygon. We call this polygon the *concavity polygon*, which plays an important role in many stages of our algorithm. 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 $p \in e_\nu = (u, v)$ is projected to

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

where $\vec{u}\vec{p} \cdot \vec{u}\vec{b}$ is the inner product of $\vec{u}\vec{p}$ and $\vec{u}\vec{b}$, 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 .

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

$$\begin{aligned} \forall e \in E_R, \quad C(e) < \tau, \text{ and} \\ \exists e = (u \in V_R, v \in V_R), \quad C(e) > \tau, \end{aligned} \quad (2)$$

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 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. Figure 3 shows examples of convex ridges. We can also see that convex ridges are quite insensitive to surface noise in Figure 3. 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_ν after e_ν is split into two subpaths at a vertex of e_ν . Let $e_\nu = \{v_0, v_1, \dots, v_{n-1}\}$ be a deep valley, and let $e_\nu^k = \{v_0, v_1, \dots, v_{k-1}\}$ be a prefix subset of e_ν , where $k \leq n$. We further let $\hat{e}_\nu = \{v_{n-1}, v_{n-2}, \dots, v_0\}$ be the reverse of e_ν . The residual concavity of e_ν at the k -th vertex is then defined as:

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

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_ν is split, two or more new bridges must be formed to determine the concavity of e_ν '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 $RC(e_\nu, k)$ takes time linear to the size of e_ν , i.e. $O(n)$ for e_ν with n edges.*

Proof. See details in the supplementary material. \square

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_ν is *V-shaped* if there exists a vertex v_k such that the residual concavity of e_ν is less than the tolerable τ . Otherwise e_ν is *U-shaped*. An example of a *U-shaped* valley can be found in Fig. 4. If e_ν is *U-shaped*, we can always find the bottom of e_ν as the sub-polygon bounded by the prefix and suffix of e_ν that have residual concavities less than τ . In Fig. 4, 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 $C(e_\nu^i) < \tau$ and $C(\hat{e}_\nu^{(n-j)}) < \tau$, respectively; recall that \hat{e}_ν is the reverse of e_ν . 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 .




Figure 3: Convex ridges found in these models. Each ellipsoid represents a convex ridge. The three figures on the right show the convex ridges on the ant model with random noise added. The convex ridges of last 3 pictures are all generated using $\tau=0.08$

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 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 $e_\beta = (a, b)$ be a bridge whose valley e_V is U -




Figure 4: 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 concavity 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 $RC(e_V^j, g)$.

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 e_V^j , a sub-valley between b and v_j , are shallow. See Fig. 4 for the illustration of e_V^i and e_V^j . Our goal is to define two valley separators in e_V^i and 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 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) :

$$(v_k, v_g), \text{ where } k = \arg \min_k RC(e_V^i, k), g = \arg \min_g RC(e_V^j, g).$$

From Fig. 4, 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.

5.2 Convex ridge separator

Graph cut is a powerful optimization tool that has been proven to be successful in segmentation. In this section, we formulate the mesh segmentation as a graph-cut problem. In the rest of this section, we will focus on how CORISE separates a convex ridge from others in a single iteration.

Data term To feed the graph cut, we define the data term using the so called *must-link regions*. Intuitively, each convex ridge has a must-link region which is a set of faces that must be segmented with the convex ridge. As we can see in Fig. 5, the boundary of must-link regions roughly define where the potential cut would be.

Potential Regions. In order to determine the must-link region, CORISE first finds a set of facets that can potentially belong to a give convex ridge R_k . Let e_V^i be the sub-valley incident to R_k , created from valley separator as described in Section 5.1. We further 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 . Now, from T_i , we




Figure 7: Shape approximation using the convex hulls of all components in the decompositions from CORiSE and HACD.

hierarchically merges facets and ensures that the concavities of all clusters are lower than the tolerable 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 smaller decomposition 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, and the bottleneck of WCSEg comes from the spectral clustering that involves solving the eigen-decomposition of a large pairwise matrix (whose size is square to the number of facets). 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 effective (use less components to provide tighter approximation). This is especially true in the human and octopus category. Fig. 8 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 we

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	CORiSE	WCSEg	HACD	CORiSE	WCSEg	HACD	CORiSE	WCSEg
human	12.10	9.85	18.85	84.80	2.33	677.92	1.23	1.31	1.24
cup	18.90	4.55	5.50	372.09	1.25	1362.30	2.66	3.15	3.02
glass	6.40	5.53	8.90	433.84	0.63	416.09	1.69	1.75	1.90
airplane	5.20	7.55	7.70	1103.85	1.07	431.91	1.30	1.37	1.25
ant	13.45	14.50	11.25	223.99	1.76	413.63	1.11	1.14	1.17
chair	9.50	10.64	16.35	297.72	1.86	1021.50	1.67	1.45	1.56
octopus	14.40	14.45	9.10	903.44	1.98	599.90	1.23	1.27	1.83
table	6.00	6.05	5.55	3385.97	1.25	2762.60	1.33	1.42	1.68
teddy	6.80	7.05	7.45	1548.05	2.16	1026.30	1.05	1.07	1.05
hand	9.10	9.00	9.20	205.92	1.67	670.00	1.14	1.27	1.16
plier	5.20	5.89	6.50	393.18	0.88	368.90	1.30	1.33	1.37
fish	5.05	5.20	6.25	2969.80	0.84	546.04	1.12	1.18	1.16
bird	5.55	6.39	11.10	547.89	0.73	515.03	1.36	1.35	1.26
armadillo	15.30	14.90	20.10	56.89	5.46	1914.80	1.12	1.19	1.12
bust	9.30	7.80	8.70	118.18	4.50	2285.80	1.05	1.11	1.08
mech	4.45	3.65	3.70	1286.43	1.04	2587.30	0.87	1.02	1.02
bearing	5.47	1.58	3.00	1095.92	0.25	1188.90	0.95	1.17	1.08
vase	9.80	5.15	5.65	1059.17	1.63	1247.60	1.06	1.12	1.10
four-leg	12.95	12.25	14.05	61.85	1.86	614.97	1.20	1.25	1.23
Average	9.21	8.00	9.42	849.95	1.75	1086.90	1.29	1.36	1.38

would cut 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. 13, 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. 9 shows the Rand Index (RI) scores obtained from seven methods on the same benchmark. Although 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

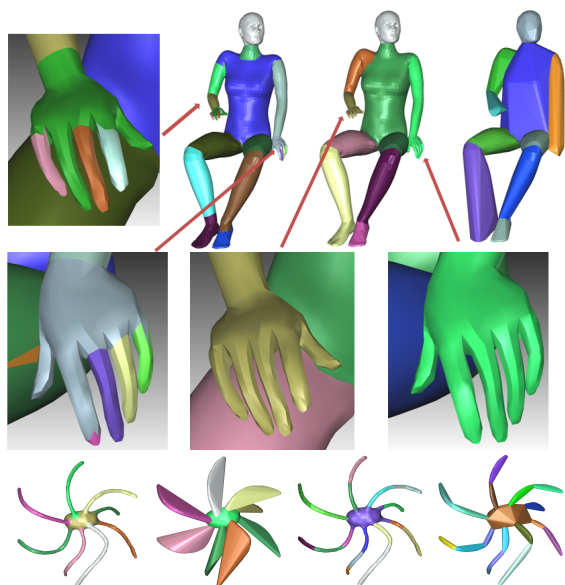


Figure 8: WCSeg vs. CORiSE Top row from left to right: WCSeg decomposition, CORiSE decomposition and convex hulls of CORiSE components. Bottom row from left to right: WCSeg decomposition, convex hulls of WCSeg components, CORiSE decomposition, and convex hulls of CORiSE components.

all single-model methods in Fig. 9.

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 parameters, 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. RC: Rand cut method; SDF: Shape Diameter Function method;

WCSeg	RC	SDF	HACD	CORiSE
≥ 4	≥ 2	≥ 2	7	2

7 Applications

CORiSE is designed to approximate a 3D shape with a




Figure 9: 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.

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 allows computations, such as collision response and penetration depth, more efficient. In these applications, we




Figure 10: Simulation created using the convex hulls of CORiSE.

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.




Figure 11: A cage created from CORISE decomposition.

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. [27] 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 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. Fig. 11 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. 12, 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.

8 Conclusion, Limitations and Future Work

In summary, in this paper we proposed a 3D decomposition method that produces components with bounded convexity. 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




Figure 12: (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.

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. 13(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. 13(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 13: (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.

References

- [1] O. van Kaick, N. Fish, Y. Kleiman, S. Asafi, and D. Cohen-Or, "Shape segmentation by approximate convexity analysis," *ACM Trans. on Graphics*, vol. to appear, 2014.
- [2] O.-C. Au, Y. Zheng, M. Chen, P. Xu, and C.-L. Tai, "Mesh segmentation with concavity-aware fields," *Visualization and Computer Graphics, IEEE Transactions on*, vol. 18, no. 7, pp. 1125–1134, 2012.
- [3] J. Zhang, J. Zheng, C. Wu, and J. Cai, "Variational mesh decomposition," *ACM Transactions on Graphics (TOG)*, vol. 31, no. 3, p. 21, 2012.
- [4] E. Kalogerakis, A. Hertzmann, and K. Singh, "Learning 3d mesh segmentation and labeling," *ACM Transactions on Graphics (TOG)*, vol. 29, no. 4, p. 102, 2010.
- [5] Q. Huang, V. Koltun, and L. Guibas, "Joint shape segmentation with linear programming," in *ACM Transactions on Graphics (TOG)*, vol. 30, no. 6. ACM, 2011, p. 125.
- [6] L. Shapira, A. Shamir, and D. Cohen-Or, "Consistent mesh partitioning and skeletonisation using the shape diameter function," *The Visual Computer*, vol. 24, no. 4, pp. 249–259, 2008.
- [7] A. Golovinskiy and T. Funkhouser, "Randomized cuts for 3d mesh analysis," in *ACM Transactions on Graphics (TOG)*, vol. 27, no. 5. ACM, 2008, p. 145.
- [8] J.-M. Lien and N. M. Amato, "Approximate convex decomposition of polygons," *Comput. Geom. Theory Appl.*, vol. 35, no. 1, pp. 100–123, 2006.
- [9] K. Mamou and F. Ghorbel, "A simple and efficient approach for 3d mesh approximate convex decomposition," in *Image Processing (ICIP), 2009 16th IEEE International Conference on*, Nov. 2009, pp. 3501 – 3504.
- [10] M. Müller, N. Chentanez, and T.-Y. Kim, "Real time dynamic fracture with volumetric approximate convex decompositions," *ACM Transactions on Graphics (TOG)*, vol. 32, no. 4, p. 115, 2013.
- [11] M. Attene, S. Katz, M. Mortara, G. Patané, M. Spagnuolo, and 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, no. 6. Wiley Online Library, 2008, pp. 1539–1556.
- [13] X. Chen, A. Golovinskiy, and T. Funkhouser, "A benchmark for 3d mesh segmentation," *ACM Transactions on Graphics (TOG)*, vol. 28, no. 3, p. 73, 2009.
- [14] S. Shlafman, A. Tal, and S. Katz, "Metamorphosis of polyhedral surfaces using decomposition," in *Computer Graphics Forum*, vol. 21, no. 3. Wiley Online Library, 2003, pp. 219–228.
- [15] S. Katz and A. Tal, "Hierarchical mesh decomposition using fuzzy clustering and cuts," *ACM Trans. Graph.*, vol. 22, no. 3, pp. 954–961, 2003.
- [16] H. Yamauchi, S. Lee, Y. Lee, Y. Ohtake, A. Belyaev, and 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 and H. Zhang, "Mesh segmentation via spectral embedding and contour analysis," in *Computer Graphics Forum*, vol. 26, no. 3. Wiley Online Library, 2007, pp. 385–394.
- [18] Y. Wang, M. Gong, T. Wang, D. Cohen-Or, H. Zhang, and B. Chen, "Projective analysis for 3d shape segmentation," *ACM Transactions on Graphics (TOG)*, vol. 32, no. 6, p. 192, 2013.
- [19] A. Golovinskiy and T. Funkhouser, "Consistent segmentation of 3d models," *Computers & Graphics*, vol. 33, no. 3, pp. 262–269, 2009.
- [20] L. Lu, Y. Choi, W. Wang, and M. Kim, "Variational 3d shape segmentation for bounding volume computation," in *Computer Graphics Forum*, vol. 26, no. 3. Wiley Online Library, 2007, pp. 329–338.
- [21] M. Attene, M. Mortara, M. Spagnuolo, and B. Falcidieno, "Hierarchical convex approximation of 3d shapes for fast region selection," in *Computer Graphics Forum*, vol. 27, no. 5. Wiley Online Library, 2008, pp. 1323–1332.
- [22] V. Krevoy, D. Julius, and A. Sheffer, "Model composition from interchangeable components," in *Computer Graphics and Applications, 2007. PG'07. 15th Pacific Conference on*. IEEE, 2007, pp. 129–138.
- [23] M. Ghosh, N. M. Amato, Y. Lu, and J.-M. Lien, "Fast approximate convex decomposition using relative

concavity," *Computer-Aided Design*, vol. 45, no. 2, pp. 494–504, 2013.

- [24] S. Asafi, A. Goren, and D. Cohen-Or, "Weak convex decomposition by lines-of-sight," in *Computer Graphics Forum*, vol. 32, no. 5. Wiley Online Library, 2013, pp. 23–31.
- [25] Z. Lian, A. Godil, P. L. Rosin, and 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.
- [26] H. Zimmer, M. Campen, and 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.
- [27] C. Xian, H. Lin, and S. Gao, "Automatic cage generation by improved obbs for mesh deformation," *The Visual Computer*, vol. 28, no. 1, pp. 21–33, 2012.