

# Simultaneous Shape Decomposition and Skeletonization\*

Jyh-Ming Lien<sup>†</sup>

John Keyser<sup>‡</sup>

Nancy M. Amato<sup>§</sup>

Department of Computer Science, Texas A&M University

## Abstract

Shape decomposition and skeletonization share many common properties and applications. However, they are generally treated as independent computations. In this paper, we propose an iterative approach that simultaneously generates a hierarchical shape decomposition and a corresponding set of multi-resolution skeletons. In our method, a skeleton of a model is extracted from the components of its decomposition — that is, both processes and the qualities of their results are interdependent. In particular, if the quality of the extracted skeleton does not meet some user specified criteria, then the model is decomposed into finer components and a new skeleton is extracted from these components. The process of simultaneous shape decomposition and skeletonization iterates until the quality of the skeleton becomes satisfactory. We provide evidence that the proposed framework is efficient and robust under perturbation and deformation. We also demonstrate that our results can readily be used in problems including skeletal deformations and virtual reality navigation.

**CR Categories:** I.3.5 [COMPUTER GRAPHICS]: Computational Geometry and Object Modeling—Geometric algorithms, languages, and systems

**Keywords:** skeletonization, multi-resolution skeleton, convex decomposition

## 1 Introduction

Shape decomposition partitions a model into (visually) meaningful components. Recently shape decomposition has been applied to texture mapping [Sander et al. 2003], shape manipulation [Katz and Tal 2003], shape matching [Mangan and Whitaker 1999; Dey et al. 2003; Funkhouser et al. 2004], and collision detection [Li et al. 2001]. Early work on shape decomposition can be found in pattern recognition and computer vision; see surveys in [Rom and Medioni 1994; Wu and Levine 1997].

A *skeleton* is a lower dimensional object that essentially represents the shape of its target object. Because a skeleton is simpler than the original object, many operations, e.g., shape recognition and deformation, can be performed

\*This research supported in part by NSF Grants EIA-0103742, ACR-0081510, ACR-0113971, CCR-0113974, ACI-0326350, and by the DOE.

<sup>†</sup>e-mail:neilien@cs.tamu.edu

<sup>‡</sup>e-mail:keyser@cs.tamu.edu

<sup>§</sup>e-mail:amato@cs.tamu.edu

Copyright © 2006 by the Association for Computing Machinery, Inc.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from Permissions Dept, ACM Inc., fax +1 (212) 869-0481 or e-mail [permissions@acm.org](mailto:permissions@acm.org).

SPM 2006, Cardiff, Wales, United Kingdom, 06–08 June 2006.

© 2006 ACM 1-59593-358-1/06/0006 \$5.00




Figure 1: The skeleton (shown in the lower row) evolves with the shape decomposition (shown in the upper row).

more efficiently on the skeleton than on the full object. The process of generating such a skeleton is called *skeletonization*. Examples of *automatic* skeleton extraction include the Medial Axis Transform (MAT) [Blum 1967] and skeletonization into a *one dimensional poly-line skeleton* (or simply *1D skeleton*) [Capell et al. 2002; Lien and Amato 2006; Katz and Tal 2003].

Skeletons have been extracted from different sources, such as voxel (image) based data [Zhou and Toga 1999; Palenichka et al. 2002; Bitter et al. 2001], boundary represented models [Chuang et al. 2000; Amenta et al. 2001; Wu et al. 2003], and scattered points [Verroust and Lazarus 1999], and for different purposes, such as shape description [Sheehy et al. 1996; Shinagawa and Kunii 1991], shape approximation [Attali et al. 1994; Wyvill and Handley 2001], similarity estimation [Hilaga et al. 2001], collision detection [Bradshaw and O’Sullivan 2002; Hubbard 1996], biological applications [Amenta et al. 2002], navigation in virtual environments [Li et al. 1999], and animation [Teichmann and Teller 1998; Katz and Tal 2003].

Although it has been noted before that a good shape decomposition can be used to extract a high quality skeleton [Lien and Amato 2006; Katz and Tal 2003] and that a high quality skeleton can be used to produce a good decomposition [Li et al. 2001], this relationship between shape decomposition and skeleton extraction is a relatively unexplored concept, especially in 3D. Instead, when a relationship is noted, the skeletons are usually treated as an intermediate result or a by-product of the shape decomposition.

In this paper, we propose an integrated framework for simultaneous shape decomposition and skeleton extraction that not only acknowledges, but actually exploits the interdependence between these two operations. First, a simple skeleton is extracted from the components of the current decomposition. Then, this extracted skeleton is used to eval-




Figure 2: Simultaneous shape decomposition and skeleton extraction. The set  $\{C_i\}$  is a decomposition of the input model  $P$  and initially  $\{C_i\} = \{P\}$ .

uate the quality of the decomposition. If the skeleton is satisfactory under some user defined criteria, we report the skeleton and the decomposition as our final results. Otherwise, the components are further decomposed into finer parts using approximate convex decomposition (ACD) [Lien and Amato 2006; Lien and Amato 2004], which decomposes a given component by ‘cutting’ its most concave features. Figure 2 illustrates this proposed framework and Figure 1 shows an example of the co-evolution process of the shape decomposition and skeleton extraction.

As we will show, our proposed approach has several advantages and makes contributions as listed below.

- This recursive refinement strategy generates multi-resolution skeletons, from coarse to fine levels of detail, which are useful for some applications.
- *Divide-and-conquer* algorithms which operate on the decompositions or skeletons can be more efficient because refinement is applied to the more complex regions but not to areas with less variation.
- The extracted skeleton is invariant under translation, rotation, and uniform scale, and is not very sensitive to boundary noise and skeletal deformations.
- Our approach does not require any pre-processing, e.g., model simplification, or any post-processing, e.g., skeleton pruning, which are required by many of the existing methods, e.g., [Li et al. 2001; Katz and Tal 2003; Wu et al. 2003].
- Our framework is general enough to work for both 2D polygons and 3D polyhedra.

## 2 Related Work

Both shape decomposition and skeleton extraction have been studied for decades and there is a large amount of previous work. In this review, we concentrate on recent developments most relevant to our work.

**Shape decomposition.** Inspired by psychological studies, such as recognition by components [Biederman 1987] and the minima rule [Hoffman and Richards 1984; Hoffman and Singh 1997], methods have been proposed to partition models at salient features to produce visually meaningful components. In pattern recognition, Rom and Medioni [Rom and Medioni 1994] partition a model into a set of tubular (generalized cylinder) shapes according to their curvature properties. As a preprocessing step for mesh generation, Sonthi et al. [Lu et al. 1999] identify closed sets (loops) of edges with required convexity and use them to decompose a model into solid parts. However, these methods work best with simple models with sharp internal angles, such as mechanical parts.

Methods that are applicable to models with general shapes also exist. Wu and Levine [Wu and Levine 1997] propose a partitioning method based on a simulated electrical charge distribution on the surface of a model. Mangan and Whitaker [Mangan and Whitaker 1999] and Page et al. [Page et al. 2003] decompose polygonal meshes by applying watershed segmentation with curvature computation. Li et al. [Li et al. 2001] decompose polygonal meshes at critical points along skeletons obtained via model simplification. Dey et al. [Dey et al. 2003] segment a model, in  $\mathbb{R}^2$  or  $\mathbb{R}^3$ , into *stable manifolds*, which are collections of Delaunay complexes of sampled points on the boundary. Katz and Tal [Katz and Tal 2003] cluster mesh facets into fuzzy regions, carefully partition facets in those regions, and successfully produce perceptually clean cuts between decomposed components. A similar approach using a different clustering technique can also be found in [Liu and Zhang 2004]. Interactive methods [Lee et al. 2004; Funkhouser et al. 2004] that identify features via human assistance have also been shown to produce high quality and clean decompositions.

**Skeletonization.** The Medial Axis (MA), Voronoi diagram, Shock graph and Reeb graph are common skeleton representations. Although the MA can represent a loss-less shape descriptor [Blum 1967], it is difficult and expensive to compute accurately in high ( $> 2$ ) dimensional space [Culver et al. 2004]. Several ideas for approximating the MA have been proposed, e.g., using Voronoi diagram, and its dual Delaunay triangulation [Amenta et al. 2001; Attali and Lachaud 2001; Dey and Zhao 2002], of densely sampled points from the object boundary. Shock graphs [Siddiqi et al. 1998; Cyr and Kimia 2001], another representation of the MA, encode the formation order and, therefore, the importance of each part of the MA. Reeb graphs, a type of 1D skeleton, extracted from various Morse functions, are a powerful tool for shape matching [Verroust and Lazarus 1999; Shinagawa et al. 1991; Attene et al. 2001; Hilaga et al. 2001]. Since Morse functions are defined on mesh vertices, re-meshing [Hilaga et al. 2001; Attene et al. 2001] is usually needed to generate a good (accurate) skeleton.

Several methods have been proposed to extract a skeleton from the components of a decomposition [Lien and Amato 2006; Katz and Tal 2003]. Skeletons can also be constructed by simplifying (contracting) a polygonal mesh to line segments [Li et al. 2001].

**Multi-scale and multi-resolution skeletons.** Multi-scale skeletons [Rom and Medioni 1993; Ogniewicz and Kubler 1995] consist of a set of skeletons,  $S_0, \dots, S_N$ , whose union represents a complete skeleton of the model.  $S_0$  is the most important part of the skeleton, representing global topology, while  $S_N$  encodes local features and is sensitive to local changes. Multi-resolution skeletons [Hilaga et al. 2001] consist of a set of skeletons,  $S_0, \dots, S_N$ , that encode topology at different levels of detail.  $S_0$  will have the coarsest skeleton and  $S_N$  will contain the most detailed information. This representation is desired for some applications. For instance, to extract similar items from a 3D database, a rough skeleton can be used to reject many unlikely models and incrementally refine the skeleton to get better matches. As previously mentioned, one of the features of our framework is that its recursive nature results in the construction of multi-resolution skeletons.

## 3 Preliminaries

Let  $P$  be a polyhedron represented by its boundary  $\partial P$  and let  $H_P$  be the convex hull of  $P$ .

**Approximate Convex Decomposition.** A set of components  $\{C_i\}$  is a *decomposition* of  $P$  if their union is  $P$  and all  $C_i$  are interior disjoint, i.e.,  $\{C_i\}$  must satisfy:

$$D(P) = \{C_i \mid \cup_i C_i = P \text{ and } \forall_{i \neq j} C_i^\circ \cap C_j^\circ = \emptyset\}, \quad (1)$$

where  $C^\circ$  is the open set of  $C$ .

A component  $C$  is  $\tau$ -*approximate* convex if  $C$  has *concavity* less than or equal to a tunable variable  $\tau$ . We use  $\text{concave}(C)$  to denote the concavity measurement of  $C$ . Therefore, the  $\tau$ -*approximate* convex decomposition of  $P$  is:

$$\text{ACD}_\tau(P) = \{C_i \in D(P) \mid \text{concave}(C_i) \leq \tau\}. \quad (2)$$

We define the concavity of a vertex  $x$  of  $C$  as the distance from  $x$  to the convex hull surface of  $C$  and the concavity of  $C$  as the maximum concavity of its vertices, i.e.,  $\text{concave}(C) = \max_{x \in C}(\text{concave}(x))$ . ACD iteratively identifies and resolves concave features with maximum concavity. Figure 3 shows an example of this process. Due to space limitations, we refer readers to [Lien and Amato 2004; Lien and Amato 2006] for details regarding ACD, including how ACD handles models with holes and handles.




Figure 3: (a) The input model. (b) The convex hull of the input model. The concavity of  $x$  is measured as the distance from  $x$  to the convex hull surface. (c) The shading of the model represents concavity, i.e., darker areas have higher concavity. (d) The model is decomposed by partitioning at the high concavity region (indicated by an arrow).

**The Principal Axis.** Let  $X$  be a set of points and  $\ell$  be a line. We define  $\text{dist}(X, \ell)$ , the distance from  $X$  to  $\ell$ , as  $\sum_X \text{dist}(x, \ell)$ , where  $x \in X$ . Then, the *principal axis* (PA) of a set of points  $X$  is a line  $\ell$  such that distance  $\sum_X \text{dist}(x, \ell)$  is minimized over all possible lines  $\kappa \neq \ell$ .

---

#### Algorithm 3.1 SSS( $P$ )

---

```

1:  $S = \text{Ext\_Skeleton}(P)$ 
2: if  $\text{Error}(P, S) \leq \tau$  then
3:   Report  $S$  as  $P$ 's skeleton and report  $P$  as a component
4: else
5:    $\{C_i\} = \text{Decompose}(P)$ 
6:   For each  $C \in \{C_i\}$  do return SSS( $C$ )

```

---

## 4 Framework

We propose a framework that simultaneously performs shape decomposition and skeleton extraction. For a given polyhedron  $P$ , **Simultaneous Shape decomposition and Skeleton extraction** (SSS) (see Algorithm 3.1) constructs a skeleton for the model from (local) skeletons extracted from each component of a decomposition, evaluates the extracted

skeleton components, and continues *refining* the decomposition and the associated skeleton components until the quality of the skeleton for each component is satisfactory, e.g., the error estimation of the skeleton for the respective component is smaller than a tunable threshold  $\tau$ .

There are three important sub-routines that are required by Algorithm 3.1:  $\text{Ext\_Skeleton}(P)$ , which extracts a skeleton from a component  $P$ ,  $\text{Error}(P, S)$ , which estimates the quality of the extracted skeleton, and  $\text{Decompose}(P)$ , which separates  $P$  into sub-components when the extracted skeleton is not acceptable. We discuss methods for skeleton extraction  $\text{Ext\_Skeleton}(P)$  in Section 4.1, and methods for quality measurement  $\text{Error}(P, S)$  in Section 4.2. Recall that our choice for the  $\text{Decompose}(P)$  sub-routine is approximate convex decomposition.

### 4.1 Extracting Skeletons

In this section, we discuss two simple methods to extract a (*local*) skeleton from a component of a decomposition. These local skeletons can be connected to form a global skeleton of the input model. The centroid method is a simple approach that can result in skeletons that do not represent the shape of the object. The second method, based on the principal axis of a component, is slightly more expensive to compute, but leads to improved skeletons in some cases.

**Using Centroids.** One of the easiest ways to construct a skeleton for a component  $C$  (in a decomposition) is to connect the centroids of the openings, called *opening centroids*, on  $\partial C$  to the centroid of  $C$ . These openings are generated when a component is split into sub-components during the decomposition process,

Several similar methods for extracting skeletons have been proposed [Lien and Amato 2006; Katz and Tal 2003]. Although this approach is simple and generates fairly good results one of the major drawbacks of this type of skeleton is its inability to represent some types of shapes. For example, the skeleton of a cross-like model in Figure 4 extracted using its centroids is only a line segment instead of two crossing line segments. The method described next attempts to address this problem.




Figure 4: This example shows a problem that arises when skeletonization is based only on the centroids. Points  $b$  and  $d$  are the centers of the openings and  $a$ ,  $c$  and  $e$  are the centers of the components  $P_1$ ,  $P_2$  and  $P_3$ , respectively. This problem can be addressed using the principal axis.

**Using the Principal Axis.** In this method, we extract a skeleton from a component  $C$  (in a decomposition) using the principal axis of the convex hull  $H_C$  of  $C$ . Instead of connecting the centroids of  $C$ 's openings to the center of mass of  $C$ , we connect these centroids to the principal axis enclosed in  $H_C$ . Figure 5 shows an example of skeletons constructed in this manner.

Let  $\text{PA}(H_C)$  be a line through the center of mass of  $H_C$ , parallel to the principal axis of  $H_C$ . Our method connects an opening centroid to one of the  $k$  points on  $\text{PA}(H_C) \cap H_C$ . These  $k$  points, denoted by  $\mathcal{P}$ , evenly subdivide  $\text{PA}(H_C) \cap H_C$  into  $k+1$  line segments. The selection of the value of  $k$  is based on the desired minimum skeleton link length. Let  $\mathcal{P}' \subset \mathcal{P}$  be a set of points to which the opening centroids connect. Figure 5 illustrates  $\mathcal{P}$  and  $\mathcal{P}'$  with circles along  $\text{PA}(H_C)$ . Then, the final skeleton  $S$  of  $C$  contains line segments that connect the opening centroids to  $\mathcal{P}'$  and line segments that connect the  $\mathcal{P}'$ .

To minimize the chance of getting a long skeleton with many joints, we match the opening centroids to  $\mathcal{P}$  so that the cardinality of  $\mathcal{P}'$  and the distances from the opening centroids to  $\mathcal{P}'$  are minimized. We solve this optimization matching problem using dynamic programming. Details of how we find the optimal solution are discussed in Appendix A.

In cases where all the points in  $\mathcal{P}'$  lie only on one side of the center of mass  $c$  of  $H_C$ , e.g.,  $\mathcal{P}'$  in Figure 5(b), line segments that connect to the points in  $\mathcal{P}'$  are not enough to represent the entire component. In such cases, the skeleton will connect  $\mathcal{P}'$  with the end point of  $\mathcal{P}$  on the other side of the center of mass. Similarly, when  $\mathcal{P}'$  contains only  $c$ , the skeleton will connect  $c$  with the end points of  $\mathcal{P}$  on both sides of  $c$ , e.g., the skeleton of the component  $P_1$  in Figure 4 (using the principal axis).

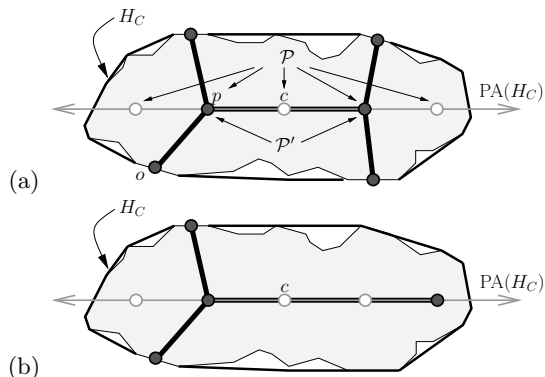


Figure 5: Using the principal axis of the convex hull  $H_C$  to extract a skeleton from a component. Skeletons are shown in dark thick lines and skeletal joints are shown in dark circles and  $c$  denotes the center of mass of  $H_C$ . (a) Opening centroids are connected to both sides of  $c$ . (b) Opening centroids are connected to only one side of  $c$ .

Figure 6 shows three skeletons: two extracted skeletons using the centroid and the principal axis methods, and one skeleton manually generated by a professional animator. One can see that the skeleton extracted using the principal axis is topologically more similar to the animator generated skeleton than the skeleton generated using the centroid method. In Section 5, we analyze the similarity of these skeletons using graph edit distance.

## 4.2 Measuring Skeleton Quality

Although several criteria exist for measuring the quality of a skeleton, the general principles we adopt are that the skeleton should reside in the interior of the model and it should encode the “topology” of the model’s shape. Thus, using these general criteria, our strategy to compute the quality of




Figure 6: Notice the differences of these skeletons at the torso, the head, and the fingers.

a skeleton  $S$  is to compare  $S$  with its associated component  $C$ . In this section, we propose three methods for measuring quality. This first method checks whether  $S$  intersects  $\partial C$  and the second method checks the topological representation of  $S$  w.r.t.  $C$ . In the third method, we propose an adaptive measurement based on the volume of the component.

An important property of these three methods is that the error of the skeleton becomes smaller as the decomposition becomes finer. This property is justified in Appendix C. Figure 8 shows extracted skeletons based on these three quality measurements.

**Checking penetration.** Our first method measures the quality of  $S$  by checking whether  $S$  intersects the component boundary  $\partial C$ . If so, the function  $\text{Error}(C, S)$  returns a large number (larger than the tolerable value  $\tau$ ). Otherwise, zero will be returned. The consequence is that  $C$  will be decomposed if  $\partial C \cap S \neq \emptyset$ .

As seen in Figure 8, skeletonization using penetration detection stops evolving after a few iterations and does not produce skeletons that represent the dragon or the bird.

**Measuring centeredness.** In the second method, we measure the offsets of  $S$  from the level sets of the geodesic distance map on  $\partial C$ . The value for each point in this map is the shortest distance to its closest opening of  $C$ . Ideally, a skeleton should pass through all connected components in all level sets. Therefore, this measurement method simply checks the number of times that  $S$  does not do so. An example of this measurement is shown in Figure 7.

Let  $L_C$  be all the connected components in the level sets of  $C$ . We define the error of a skeleton  $S$  as:

$$\text{Err}(C, S) = \frac{\sum_{l_c \in L_C} f(l_c, S)}{|L_C|}, \quad (3)$$

where  $f(l_c, S)$  returns 0 if  $S$  intersects component  $l_c$ , and 1 otherwise, and  $|L_C|$  is the total number of the connected components in  $C$ . Details of how we compute the level sets and  $f(l_c, S)$  are discussed in Appendix B.




Figure 7: The error measurement for this skeleton, which intersects level sets 4, 7 and 8, is  $\frac{3}{8}$ .



As seen in Figure 8, skeletonization using the centeredness measurement captures the shape of the dragon and the bird better than simply using penetration detection, but it over segments the tail of the bird and does not produce accurate skeletons in the feet of the dragon or the bird.

**Measuring convexity.** Our idea for the last quality measurement comes from the observation that in many cases the significance of a feature depends on its volumetric proportion to its “base”. For example, a 5 cm stick attached to a ball with 5 cm radius is a more significant feature than a 5 cm stick attached to a ball with 5 km radius. This intuition can be captured by the concept of the *convexity* of a component  $C$  defined as  $\text{convexity}(C) = \frac{\text{vol}(C)}{\text{vol}(H_C)}$ , where  $\text{vol}(X)$  is the volume of a set  $X$ . Thus, we can define the error measurement as:

$$\text{Err}(C, S) = 1 - \text{convexity}(C) . \quad (4)$$

Assume that the skeleton  $S$  is a good representation of the convex hull  $H_C$ . Then, a smaller difference between  $H_C$  and  $C$  means that  $S$  is a better representation of  $C$ . Thus, although the skeleton  $S$  is not included in Equation 4,  $S$  is implicitly considered in terms of  $H_C$ .

As seen in Figure 8, using convexity produces the most realistic skeleton that captures the overall shape of the dragon and the bird and also identifies the detailed features of their feet.




Figure 8: Final skeletons of a dragon polyhedron and a bird polygon extracted using different quality estimation functions: checking penetration, measuring centeredness, and measuring convexity. The maximum tolerable errors for centeredness and convexity are 0.2 and 0.3, respectively.

### 4.3 Putting it All Together

---

#### Algorithm 4.1 $SSS_{ACD}(P)$

---

- 1: Compute a skeleton  $S$  from  $P$  using the *Principal Axis* of  $H_P$ .
  - 2: Estimate the quality of  $S$  using *convexity*.
  - 3: **if**  $S$  is acceptable **then**
  - 4: Report  $S$  as  $P$ 's skeleton and report  $P$  as a component.
  - 5: **else**
  - 6:  $\{C_i\} = ACD(P)$ .
  - 7: **For** each  $C \in \{C_i\}$  **do** return  $SSS_{ACD}(C)$
- 

Algorithm 4.1 shows a fleshed-out version of the proposed simultaneous shape decomposition and skeletonization framework. Here we suggest using the principal axis,

convexity and approximate convex decomposition for local skeleton extraction, quality measurement and partitioning, respectively. Algorithm 4.1 is used for all the experiments in Section 5. We would like to emphasize that the choice of these methods is made based on our own experience. The framework is not restricted to these selected sub-routines, which can be replaced by other methods to fit particular needs of an application.

## 5 Implementation and Results

The experiments in this section are used to demonstrate the *efficiency*, the *robustness*, and several *applications* of the proposed method. The method was implemented in C++ and all these experiments are performed on a Pentium 2.0 GHz CPU with 512 Mb RAM. Seventeen decompositions and their associated skeletons are shown in Figures 8 to 13 and in Tables 1 and 2.

**Efficiency.** A summary of the studied models, which include several game characters, a high genus model, and two scanned models, and the skeletonization and decomposition time of these models is shown in Table 1. Table 1 shows that the processing time of SSS depends on both the size of the model and on the complexity of the shape. For example, even though the model in Figure 9 has the fewest triangles, its large genus (18) increases the processing time. In general, our proposed SSS method can handle models with thousands of triangles in less than a half a minute and scales well for models with tens or hundreds of thousands of triangles.

We further show that SSS is efficient by comparing our results to two recently proposed shape decomposition and skeletonization methods that have been shown to produce very promising results; see Figures 10 and 11, respectively. In both experiments, SSS generates results similar to those results reported previously but SSS can produce the shape decomposition and the skeletons about 30 times and 5 times faster than those methods reported in [Katz and Tal 2003] and [Wu et al. 2003], respectively. We note that there are no well-accepted criteria to compare the quality of these decompositions and skeletons quantitatively, and therefore we do not intend to claim that our results are necessarily better.










**Robustness.** In this set of experiments, we show that SSS is robust under perturbation and deformation, meaning that the shape decompositions and skeletons remain approximately the same after the input models are perturbed and deformed. The results are shown in Table 2.

Although there are no well accepted criteria to measure the differences among decompositions, we can measure the similarity of these skeletons, e.g., using graph edit distance [Bunke and Kandel 2000] which computes the cost of operations (i.e., inserting/removing vertices or edges) needed to convert one graph to another. In this paper, we simply associate one unit of cost with each operation.

We measure two types of distances, denoted as  $D_O$  and  $D_O^2$ .  $D_O$  is the graph edit distance from a skeleton to the skeleton extracted from the original mesh. Because removing or inserting a degree-two node does not change the topology of a graph, we are also interested in the distance, denoted as  $D_O^2$ , that does not count operations that create and remove degree-two nodes. Table 2 shows that  $D_O$  remains small for both perturbed and deformed models and  $D_O^2$  is zero in all cases.

**Applications.** The extracted skeleton can be readily used to create animations. We demonstrate this advantage by re-targeting motion captured data to the skeletons extracted using our method. In Figure 12, we show a sequence

Table 1: Experimental Results

Model									
	Figure 9	Figure 12	Table 2	Figure 6	Figure 11	Figure 12	Table 2	Figure 1	Table 2
Size	1,984	3,392	5,660	6,564	8,276	11,180	39,694	48,312	243,442
Time	15.6	2.6	1.7	1.5	8.8	3.4	19.4	30.1	73.3

Size is measured as the number of the triangles of each model and the processing time is measured in seconds.

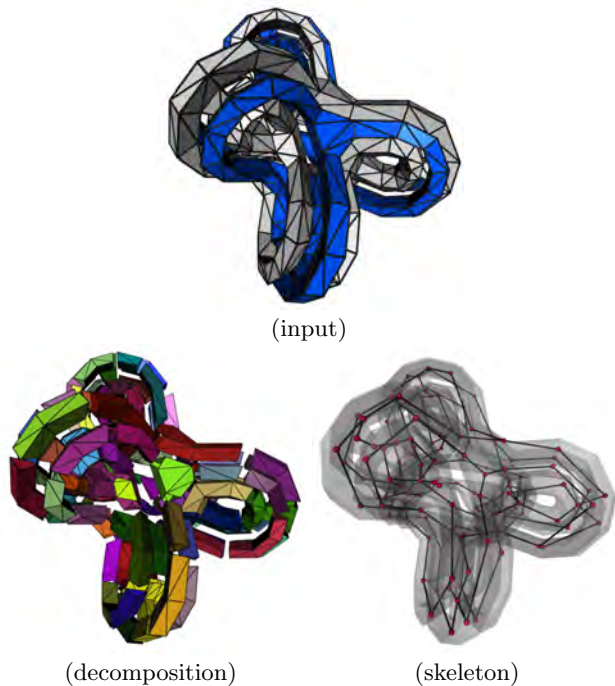


Figure 9: This figure shows the decomposition and the skeleton of a model with 18 handles.

of images obtained from a skeleton-based boxing animation of a baby and a robot using motion data captured from an adult male. Note that the baby and the robot models have different body proportions and rest poses. Due to the limited space in this paper, other animations, including walking and pushing a box, are provided on our webpages. We use motion captured data instead of a hand-made animation to show that the extracted skeletons are robust enough to be used by arbitrarily selected motions and not only carefully designed motion. The motions, i.e., joints angles, are manually copied from the captured data to the skeletal joints.

The extracted skeletons can also help to plan motion, e.g., for navigating in the human colon or removing a mechanical part from an airplane engine. Sampling-based motion planners have been shown to solve difficult motion planning problems; see a survey in [Barraquand et al. 1997]. These methods approximate the free configuration space (C-space) of a movable object by sampling and connecting random configurations to form a graph (or a tree). However, they also have several technical issues limiting their success on




Figure 10: The decomposition with 0.7 convexity and the associated skeleton of the dino-pet model (with 6,564 triangles) are computed in 1.5 seconds whereas Katz and Tal’s approach takes 57 seconds (on a P4 1.5 GHz CPU with 512 Mb RAM).

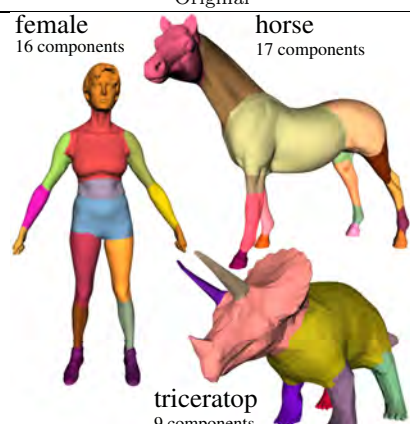
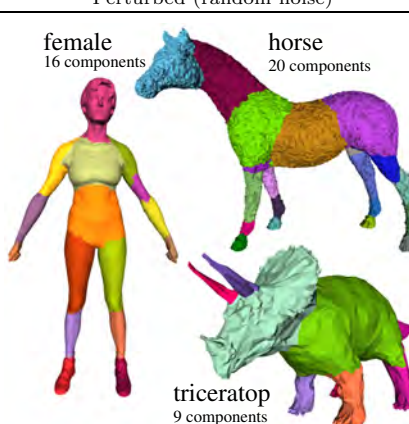
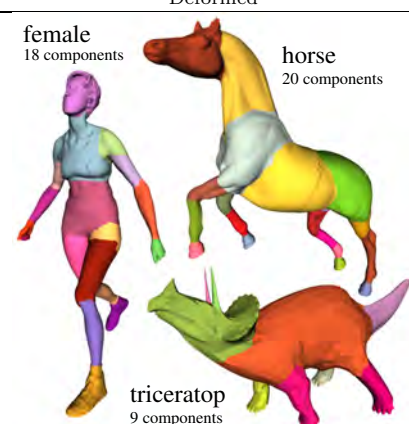
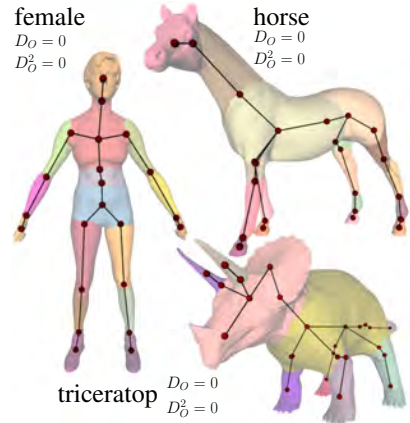
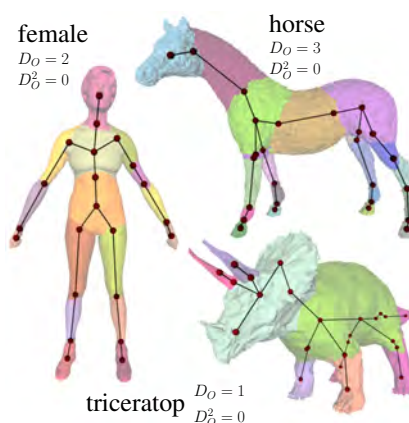
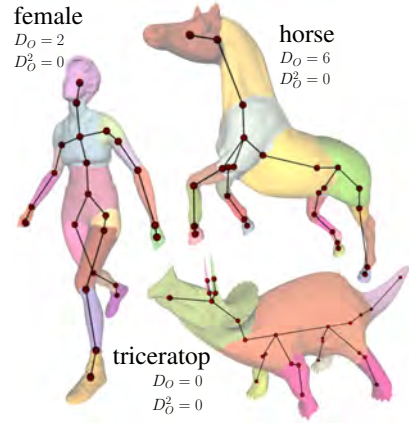
some important types of problems, such as the difficulty of finding paths that are required to pass through narrow passages [Wilmarth et al. 1999]. Using sampling biased toward the joints of the extracted skeleton, we can alleviate this so called “narrow passage” problem. Figure 13 shows that the graph constructed using the skeleton-based sampling can better represent the free C-space than using the uniform sampling [Kavraki et al. 1996] with the same number of samples. This is because more of the skeleton-based sampling samples are placed in the narrower regions. In addition, the connections between the samples in these narrow regions can be made easily because the components of the decomposition are nearly convex.

## 6 Discussion and Conclusion

In this paper, we propose a framework that simultaneously generates shape decompositions and skeletons. This framework is inspired by the observation that both operations share many common properties and applications but are generally considered as independent processes. This framework extracts the skeleton from the components in a decomposition and evaluates the skeleton by comparing it to the components. The process of simultaneous shape decomposition and skeletonization iterates until the quality of the skeleton becomes satisfactory.

We studied two simple skeleton extraction methods, using the centroids and the principal axis, and three quality evaluation measurements, that compute penetration, centeredness and convexity, respectively. In the experiments, we demonstrate that the proposed framework is efficient, ro-

Table 2: Robustness tests using perturbed and skeletal deformed meshes.  $D_O$  is the graph edit distance between a skeleton extracted from a perturbed or deformed mesh and a skeleton extracted from the original mesh.  $D_O^2$  is  $D_O$  without counting operations on degree-2 nodes (which do not change the topology of the skeleton).

Shape Decomposition. 70% convexity		
Original	Perturbed (random noise)	Deformed
<p>female 16 components</p> <p>horse 17 components</p> <p>triceratop 9 components</p> 	<p>female 16 components</p> <p>horse 20 components</p> <p>triceratop 9 components</p> 	<p>female 18 components</p> <p>horse 20 components</p> <p>triceratop 9 components</p> 
Extracted Skeletons. 70% convexity		
Original	Perturbed	Deformed
<p>female <math>D_O = 0</math> <math>D_O^2 = 0</math></p> <p>horse <math>D_O = 0</math> <math>D_O^2 = 0</math></p> <p>triceratop <math>D_O = 0</math> <math>D_O^2 = 0</math></p> 	<p>female <math>D_O = 2</math> <math>D_O^2 = 0</math></p> <p>horse <math>D_O = 3</math> <math>D_O^2 = 0</math></p> <p>triceratop <math>D_O = 1</math> <math>D_O^2 = 0</math></p> 	<p>female <math>D_O = 2</math> <math>D_O^2 = 0</math></p> <p>horse <math>D_O = 6</math> <math>D_O^2 = 0</math></p> <p>triceratop <math>D_O = 0</math> <math>D_O^2 = 0</math></p> 

bust under perturbation and deformation, and can readily be used, e.g., to generate animations and plan motion.

There are several ways to extend the current work. First, there is a need to establish a systematic framework for comparing qualities of shape decomposition and skeletons using more quantitative measuring methods and benchmarks. Although the proposed quality measurements are based on a general idea of what a good skeleton should be, more studies are needed to investigate application-specific measurement criteria that should produce better and more “comparable” results. Second, not all models, such as a bowl, can have reasonable 1D skeletons. We are interested in using the same framework to extract the approximated medial axis from the components in a decomposition based on the idea that it is easier to extract the medial axis from a convex object than from a non-convex object. Finally, because the extracted skeletons and shape decompositions in our method co-evolve, we can provide more meaningful shape decompositions by using information from the extracted skeletons, e.g., merging components if the skeletons extracted from those components do not change the global skeleton made from the

entire decomposition.

## References

- AMENTA, N., CHOI, S., AND KOLLURI, R. K. 2001. The power crust, unions of balls, and the medial axis transform. *Computational Geometry* 19, 2-3, 127–153.
- AMENTA, N., CHOI, S., JUMP, M. E., KOLLURI, R. K., AND WAHL, T. 2002. Finding alpha-helices in skeletons. Tech. rep., Dept. of Computer Science, University of Texas at Austin.
- ATTALI, D., AND LACHAUD, J.-O. 2001. Delaunay conforming iso-surface; skeleton extraction and noise removal. *Computational Geometry: Theory and Applications* 19, 2-3, 175–189.
- ATTALI, D., BERTOLINO, P., AND MONTANVERT, A. 1994. Using polyballs to approximate shapes and skeletons. In *Proceedings of International Conference on Pattern Recognition (ICPR'94)*, 626–628.
- ATTENE, M., BIASOTTI, S., AND SPAGNUOLO, M. 2001. Remeshing techniques for topological analysis. In *Proc. of the Shape Modeling International (SMI'01)*, 142–151.






Figure 12: An animation sequence obtained from applying the boxing motion capture data to the extracted skeletons from a baby model and a robot model. The motion capture data (action number 13\_17) are downloaded from the CMU Graphics Lab motion capture database. The first two figures in the sequence are the shape decompositions and the skeletons of the baby and the robot. Note that not all joint motions from the data are used because the extracted skeletons have fewer joints.

BARRAQUAND, J., LATOMBE, L. K. J., LI, T., MOTWANI, R., AND RAGHAVAN, P. 1997. A random sampling scheme for path planning. *Int. J. of Rob. Res* 16, 6, 759–774.

BIEDERMAN, I. 1987. Recognition-by-components: a theory of human image understanding. *Psychol. Rev.* 94, 2, 115–147.

BITTER, I., KAUFMAN, A. E., AND SATO, M. 2001. Penalized-distance volumetric skeleton algorithm. *IEEE Transactions on Visualization and Computer Graphics* 7, 3, 195–206.

BLUM, H. 1967. A transformation for extracting new descriptors of shape. In *Models for the Perception of Speech and Visual Form*, W. Wathen-Dunn, Ed. MIT Press, 362–380.

BRADSHAW, G., AND O’SULLIVAN, C. 2002. Sphere-tree construction using dynamic medial axis approximation. In *Proceedings of the ACM SIGGRAPH symposium on Computer animation*, ACM Press, 33–40.

BUNKE, H., AND KANDEL, A. 2000. Mean and maximum common subgraph of two graphs. *Pattern Recogn. Lett.* 21, 2, 163–168.

CAPELL, S., GREEN, S., CURLESS, B., DUCHAMP, T., AND POPOVIC, Z. 2002. Interactive skeleton-driven dynamic deformations. *ACM Transactions on Graphics* 21, 3, 586–593.

CHUANG, J.-H., TSAI, C.-H., AND KO, M.-C. 2000. Skeletonization of three-dimensional object using generalized potential field. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 22, 11, 1241–1251.

CULVER, T., KEYSER, J., AND MANOCHA, D. 2004. Exact computation of the medial axis of a polyhedron. *Comput. Aided Geom. Des.* 21, 1, 65–98.

CYR, C. M., AND KIMIA, B. B. 2001. 3d object recognition using shape similarity-based aspect graph. In *ICCV’01*.

DEY, T. K., AND ZHAO, W. 2002. Approximate medial axis as a voronoi subcomplex. In *ACM Symposium on Solid Modeling and Applications*, 356–366.

DEY, T. K., GIESEN, J., AND GOSWAMI, S. 2003. Shape segmentation and matching with flow discretization. In *Proc. Workshop on Algorithms and Data Structures*, 25–36.

FUNKHOUSER, T., KAZHDAN, M., SHILANE, P., MIN, P., KIEFER, W., TAL, A., RUSINKIEWICZ, S., AND DOBKIN, D. 2004. Modeling by example. *ACM Trans. Graph.* 23, 3, 652–663.

HILAGA, M., SHINAGAWA, Y., KOHMURA, T., AND KUNII, T. L. 2001. Topology matching for fully automatic similarity estimation of 3d shapes. In *Proceedings of the 28th annual conference on Computer graphics and interactive techniques*, 203–212.

HOFFMAN, D., AND RICHARDS, W. 1984. Parts of recognition. *Cognition* 18, 65–96.

HOFFMAN, D., AND SINGH, M. 1997. Saliency of visual parts. *Cognition* 63, 29–78.

HUBBARD, P. M. 1996. Approximating polyhedra with spheres for time-critical collision detection. *ACM Transactions on Graphics (TOG)* 15, 3, 179–210.

KATZ, S., AND TAL, A. 2003. Hierarchical mesh decomposition using fuzzy clustering and cuts. *ACM Trans. Graph.* 22, 3, 954–961.

KAVRAKI, L. E., SVETSKA, P., LATOMBE, J. C., AND OVERMARS, M. H. 1996. Probabilistic roadmaps for path planning in high-dimensional configuration spaces. *IEEE Trans. Robot. Automat.* 12, 4 (August), 566–580.

LEE, Y., LEE, S., SHAMIR, A., COHEN-OR, D., AND SEIDEL, H.-P. 2004. Intelligent mesh scissoring using 3d snakes. In *Proceed-*




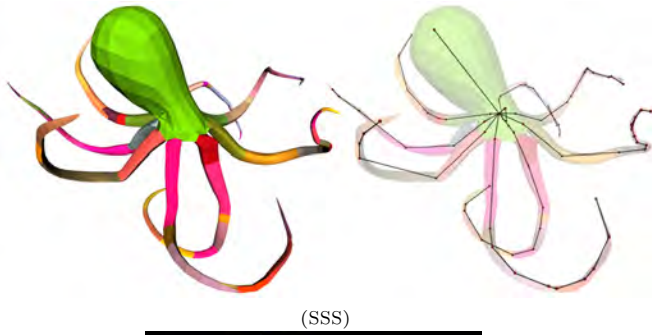



Figure 13: Narrow passage motion planning. (a) A difficult motion planning problem in which the robot is required to pass through four narrow windows to move from the start configuration (bottom) to the goal configuration (top). (b) A shape decomposition and a skeleton of the environment using the proposed method. In (c), uniform sampling results in a poor representation of the robot’s free C-space. Uniform sampling samples 350 uniformly distributed collision-free configurations and connects these samples into a graph. In this example, there are five connected components in the graph. In (d), the skeleton-based sampling results in a better representation of the robot’s free C-space. The skeleton-based sampling samples 350 collision-free configurations around the joints of the skeleton and connects these samples into a graph. In this example, the graph has one single connected component. (e) A collision-free path found by connecting the start and the goal configurations to the graph generated using the skeleton-based sampling.

- ings of the 12th Pacific Conference on Computer Graphics and Applications (PG’04), 279–287.
- LI, T.-Y., LIEN, J.-M., CHIU, S.-Y., AND YU, T.-H. 1999. Automatically generating virtual guided tours. In *Proc. IEEE Computer Animation (CA)*, 99–106.
- LI, X., TOON, T. W., AND HUANG, Z. 2001. Decomposing polygon meshes for interactive applications. In *Proceedings of the 2001 symposium on Interactive 3D graphics*, 35–42.
- LIEN, J.-M., AND AMATO, N. M. 2004. Approximate convex decomposition of polygons. In *Proc. 20th Annual ACM Symp. Computat. Geom. (SoCG)*, 17–26.
- LIEN, J.-M., AND AMATO, N. M. 2006. Approximate convex decomposition of polyhedra. Tech. Rep. TR06-002, Parasol Lab, Dept. of Computer Science, Texas A&M University, Jan.
- LIU, R., AND ZHANG, H. 2004. Segmentation of 3d meshes through spectral clustering. In *Proceedings of the 12th Pacific Conference on Computer Graphics and Applications (PG’04)*, 298–305.
- LU, Y., GADH, R., AND TAUTGES, T. J. 1999. Volume decomposition and feature recognition for hexahedral mesh generation. In *Proc. 8th International Meshing Roundtable*, 269–280.
- MANGAN, A. P., AND WHITAKER, R. T. 1999. Partitioning 3d surface meshes using watershed segmentation. *IEEE Transactions on Visualization and Computer Graphics* 5, 4, 308–321.
- OGNIEWICZ, R., AND KUBLER, O. 1995. Hierarchic voronoi skeletons. *Pattern Recognition* 28, 3, 343–359.
- PAGE, D. L., KOSCHAN, A. F., AND ABIDI, M. A. 2003. Perception-based 3d triangle mesh segmentation using fast marching watersheds. In *Proceedings of the 2003 Conference on Computer Vision and Pattern Recognition (CVPR ’03)*, 27–32.
- PALENICHKA, R. M., ZAREMBA, M. B., AND DU QUEBEC, U. 2002. Multi-scale model-based skeletonization of object shapes using self-organizing maps. In *16th International Conference on Pattern Recognition (ICPR’02)*, 10143–10147.
- ROM, H., AND MEDIONI, G. 1993. Hierarchical decomposition and axial shape description. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 15, 10, 973–981.
- ROM, H., AND MEDIONI, G. 1994. Part decomposition and description of 3d shapes. In *Proc. International Conference of Pattern Recognition*, 629–632.
- SANDER, P. V., WOOD, Z. J., GORTLER, S. J., SNYDER, J., AND HOPPE, H. 2003. Multi-chart geometry images. In *Proceedings of the Eurographics/ACM SIGGRAPH symposium on Geometry processing*, 146–155.
- SHEEHY, D. J., ARMSTRONG, C. G., AND ROBINSON, D. J. 1996. Shape description by medial surface construction. *IEEE Trans. Visualizat. Comput. Graph.* 2, 1 (Mar.), 62–72.
- SHINAGAWA, Y., AND KUNII, T. L. 1991. Constructing a reeb graph automatically from cross sections. *IEEE Computer Graphics and Applications* 11, 6, 44–51.
- SHINAGAWA, Y., KUNII, T. L., AND KERGOSIEN, Y. L. 1991. Surface coding based on morse theory. *IEEE Comput. Graph. Appl.* 11 (Sept.), 66–78.
- SIDDIQI, K., SHOKOUFANDEH, A., DICKINSON, S. J., AND ZUCKER, S. W. 1998. Shock graphs and shape matching. In *ICCV*, 222–229.
- TEICHMANN, M., AND TELLER, S. 1998. Assisted articulation of closed polygonal models. In *Proceedings of the Eurographics Workshop*, 254–254.
- VERROUST, A., AND LAZARUS, F. 1999. Extracting skeletal curves from 3d scattered data. In *International Conference on Shape Modeling and Applications*, IEEE Computer Society, 194–201.
- WILMARTH, S. A., AMATO, N. M., AND STILLER, P. F. 1999. Motion planning for a rigid body using random networks on the medial axis of the free space. In *Proc. ACM Symp. on Computational Geometry (SoCG)*, 173–180.
- WU, K., AND LEVINE, M. D. 1997. 3d part segmentation using simulated electrical charge distributions. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 19, 11, 1223–1235.
- WU, F.-C., MA, W.-C., LIU, P.-C., LAING, R.-H., AND OUHYOUNG, M. 2003. Skeleton extraction of 3d objects with visible repulsive force. In *Computer Graphics Workshop 2003, Hualien, Taiwan*.
- WYVILL, G., AND HANDLEY, C. 2001. The “thermodynamics” of shape. In *Proc. of the Shape Modeling International (SMI’01)*, 2–8.
- ZHOU, Y., AND TOGA, A. W. 1999. Efficient skeletonization of volumetric objects. *IEEE Transactions on Visualization and Computer Graphics* 5, 3, 196–209.



(SSS)



(Wu et al. [2003])

Figure 11: The decomposition with 0.7 convexity and the associated skeleton of the octopus model (with 8,276 triangles) are computed in 8.8 seconds whereas Wu et al.’s approach takes 53 seconds (on a P4 1.5 GHz CPU with 512 Mb RAM) using a simplified version of this model (with 2,000 triangles).

## Appendix

### A From a Principal Axis to a Skeleton

Here, we show how a local skeleton can be computed using the principle axis. Our goal is to find a match  $M : O \rightarrow \mathcal{P}$  from the opening centroids  $O$  to the points  $\mathcal{P}$  on the principle axis so that the total length of the match and the number of the matched points (joints) in  $\mathcal{P}$  is minimized. We let the score function  $F$  of a match  $M$  be defined as

$$F(M) = s_1 \cdot |M| + s_2 \cdot J(M), \quad (5)$$

where  $|M|$  and  $J(M)$  are the length and joint size of match  $M$ , and  $s_1$  and  $s_2$  are user specified scalars. In this paper,  $s_1$  and  $s_2$  are constants set to ten and one, resp. A brute force approach to find an optimal solution will take  $O(|\mathcal{P}|^{|O|})$  time, where  $|\mathcal{P}|$  and  $|O|$  are the number of vertices in  $\mathcal{P}$  and  $O$ , respectively. This exponential time complexity is in general impractical for most applications.

The main idea of finding the optimal match is to group opening centroids  $O$  and each group will connect to a point in  $\mathcal{P}$ . After knowing how  $O$  is grouped, it takes  $O(|\mathcal{P}||O|)$  time to find solution.

Grouping  $O$  can be done using dynamic programming. An observation that enables us to group  $O$  is that two centroids are likely to be grouped when their closest points in  $\mathcal{P}$  are close. Thus, we first sort  $O$  according to their closest points of  $\mathcal{P}$  and then group sorted  $O$ . A dynamic programming is shown in Algorithm A.1 to group  $O$ . In Algorithm A.1, we use  $G[i, j]$  to denote the optimal solution for the sub-problem  $\{O_i, \dots, O_j\}$  and use  $\langle G_i G_j \rangle$  and  $G_i G_j$  to denote the joint of two groups  $G_i$  and  $G_j$  with and without merging  $G_i$  and  $G_j$  to one group.

---

### Algorithm A.1 Optimal Matching( $O, \mathcal{P}$ )

---

```

1: for  $i \in \{1, \dots, |O|\}$  do
2:    $G[i, i] = O_i$ 
3: for  $l \in \{2, \dots, |O|\}$  do
4:   for  $i \in \{1, \dots, |O| - l + 1\}$  do
5:      $j = i + l - 1$ 
6:      $G[i, j] = \langle O_i \dots O_j \rangle$ 
7:      $score = F(G[i, j], \mathcal{P})$  { $F$  is defined in Eqn. 5}
8:     for  $k \in \{i \dots, j - 1\}$  do
9:        $s = F(G[i, k]G[k + 1, j], \mathcal{P})$ 
10:      if  $s_1 < score$  then
11:         $G[i, j] = G[i, k]G[k + 1, j]$ 
12:       $score = s_1$ 

```

---

### B Compute level sets and centeredness

A level set of a component  $C$  in a decomposition is a set of points on the surface  $\partial C$  of the component with the same geodesic distance to the closest opening of  $C$ . A connected component in a level set is a list of connected edges, which usually forms a loop on  $\partial C$ . A level set can have one or multiple connected component(s). These level sets can be computed, similar to the construction process of a Reeb graph [Shinagawa and Kunii 1991], by flooding the entire  $\partial C$  from the boundaries of the openings of  $C$ . In each iteration of this flooding process, the wavefronts will propagate from the visited vertices to unvisited vertices via incident edges.

To compute centeredness, we need to know how a skeleton  $S$  intersects the level sets of  $C$ , i.e., we need the function  $f(l_c, S)$  use in Eqn 3, which returns zero if  $S$  intersects the level set  $l_c$ . The function  $f(l_c, S)$  can be implemented by simply checking the intersection between each line segment of  $S$  and the triangulation of  $l_c$ .

### C ACD increases skeleton quality

In this section, we show that the error measurements of a skeleton described in Section 4.2, i.e., penetration, centeredness, and convexity, decrease as the input model is decomposed. This is a critical property, which allows the SSS framework to terminate.

**Lemma C.1.** *Let  $S$  be the skeleton of a polyhedron  $P$  and let  $S'$  be the skeleton of the components of the ACD of  $P$ . The error estimation of  $S'$  must be smaller than the error estimation of  $S$  measured using penetration, centeredness, and convexity defined in Section 4.2.*

*Proof.* We show that all error measurements become zero if the input model is convex. For penetration, because the segments connecting any two points inside the convex object must not intersect its boundary, a skeleton will never penetrate the object. For the same reason, the skeleton must not be ‘outside’ of any level set of a convex component. Finally, because the convexity of a convex object is one, its error must be zero.

Since ACD decomposes  $P$  into more convex components after each iteration, the error measurements of  $S'$  will be closer to zero than  $S$  for all three types of measurements.  $\square$