



## Open Archive Toulouse Archive Ouverte

OATAO is an open access repository that collects the work of Toulouse researchers and makes it freely available over the web where possible

This is a publisher's version published in:  
<http://oatao.univ-toulouse.fr/22611>

### Official URL

DOI : <https://doi.org/10.1016/j.gmod.2018.07.003>

**To cite this version:** Blanc Beyne, Thibault and Morin, Géraldine and Leonard, Kathryn and Hahmann, Stefanie and Carlier, Axel *A salience measure for 3D shape decomposition and sub-parts classification*. (2018) *Graphical Models*, 99. 22-30. ISSN 1524-0703

Any correspondence concerning this service should be sent to the repository administrator: [tech-oatao@listes-diff.inp-toulouse.fr](mailto:tech-oatao@listes-diff.inp-toulouse.fr)

# A salience measure for 3D shape decomposition and sub-parts classification

Thibault Blanc-Beyne<sup>a</sup>, Géraldine Morin<sup>a</sup>, Kathryn Leonard<sup>b</sup>, Stefanie Hahmann<sup>\*,c</sup>, Axel Carlier<sup>a</sup>

<sup>a</sup>IRIT - Toulouse Institute of Computer Science Research, France

<sup>b</sup>Occidental College, United States

<sup>c</sup>Univ. Grenoble Alpes, CNRS (LJK), INRIA, Grenoble INP, France

## A B S T R A C T

### Keywords:

Shape analysis  
Part hierarchy  
Medial axis  
WEDF  
Salience measure  
3D shape decomposition

This paper introduces a measure of significance on a curve skeleton of a 3D piecewise linear shape mesh, allowing the computation of both the shape's parts and their saliency. We begin by reformulating three existing pruning measures into a non-linear PCA along the skeleton. From this PCA, we then derive a volume-based salience measure, the 3D WEDF, that determines the relative importance to the global shape of the shape part associated to a point of the skeleton. First, we provide robust algorithms for computing the 3D WEDF on a curve skeleton, independent on the number of skeleton branches. Then, we cluster the WEDF values to partition the curve skeleton, and coherently map the decomposition to the associated surface mesh. Thus, we develop an unsupervised hierarchical decomposition of the mesh faces into visually meaningful shape regions that are ordered according to their degree of perceptual salience. The shape analysis tools introduced in this paper are important for many applications including shape comparison, editing, and compression.

## 1. Introduction

Many 3D applications make use of 3D shape decompositions in various forms, whether for shape comparison and compression, for editing purposes, or for animation. While much work has studied how to produce meaningful parts, there has been little focus on providing additional information on parts, such as their relative importance. In [1], the authors show that decomposition into a hierarchy of part and subparts is memorized by user, and such structure is relevant for shape recognition. Such a measure has been formalized by Hoffman and Singh [2] and called salience of parts. This measure reveals how important a part is to understand the shape as a whole. In this paper, we propose a shape decomposition method that naturally provides insights about parts salience.

The present work makes a first step towards automatic similarity-based shape decomposition by introducing a new salience measure for 3D shape decomposition and subparts classification. In particular, the proposed decomposition provides an automatic pre-processing for applications such as structuring a 3D shape (e.g. like [3] whose semantic descriptors relies on a decomposition of the shape), or similarity detection in a 3D shape, as required by Guy et al. [4].

A key observation is the fact that, despite a reputation for instability, a skeleton can provide a reliable basis for shape analysis when it is complemented with measures defined along it. Below, we introduce

such a measure, called *Weighted Extended Distance Function (WEDF)*, which is defined on a curve skeleton centered within a 3D shape. WEDF has been introduced recently for 2D shape contours [5] and successfully used for defining a salience measure and decomposition of 2D curves [6]. Generalizing to the 3D setting takes two steps: first defining WEDF on a 3D skeleton and deducing a partition, second, mapping the measure on the surface mesh. The second step assesses salience of surface mesh parts: it establishes an injective mapping between surface mesh vertices and the skeleton points while ensuring simultaneously that the inverse mapping produces coherent shape parts. In 2D, a straightforward association of boundary points to the WEDF values of the corresponding skeletal points generates a good partition of a shape [6], but the 3D setting is more challenging: a direct extension of the 2D approach to 3D generates a poor partition (see Fig. 10, where the legs run into the body of the ant). Accordingly, we propose a robust algorithm for appropriately transferring the salience measure on the skeleton to the surface mesh, enabling a decomposition into a consistent set of subparts that can be clustered according to their salience levels.

Given a WEDF measure coherently transferred to a surface mesh, we use unsupervised methods for decomposing a 3D shape into subparts and for classifying those parts according to their visual salience. This allows us to determine a *primary supporting shape* and *features* with decreasing saliency. In contrast to classical shape segmentation algorithms, our method provides a *graded salience-based parts decomposition*,

\* Corresponding author.

E-mail address: [stefanie.hahmann@inria.fr](mailto:stefanie.hahmann@inria.fr) (S. Hahmann).

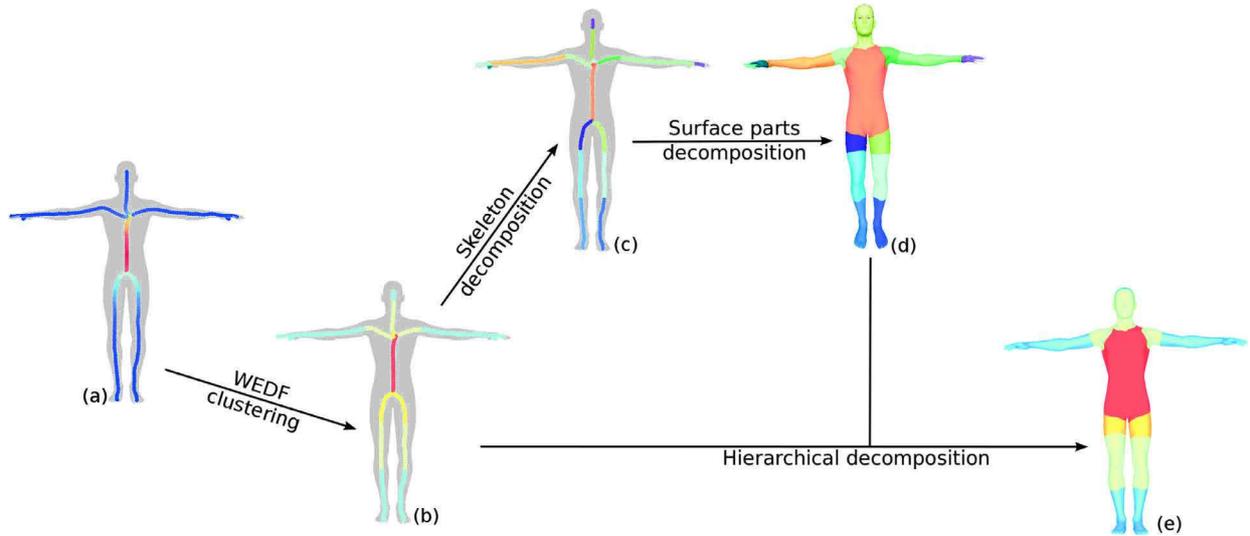


Fig. 1. Pipeline of our hierarchical decomposition. Starting from a 3D shape and its curve skeleton, we compute a new measure called WEDF on the curve skeleton (a) and, by clustering WEDF values, we decompose the skeleton into hierarchical parts (b). To each connected part on the skeleton –shown with a different color (c)– a connected region of the surface mesh is assigned (d). Then, a saliency value according to the hierarchy is assigned to each corresponding surface part (e) –parts of same importance get a similar color. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

where each part is labeled with a value characterizing its importance to the overall shape. Fig. 1 shows an example: our method not only segments the body parts from the body, but additionally labels each body part with an importance value so that parts with the same importance value belong to the same level in the importance hierarchy. All rose mesh parts belong to the main body, all green parts to the lowest hierarchy level, all yellow parts to a parent hierarchy level, and so on.

Our hierarchy labels inherit the following two features from the relationship to the curve skeleton:

- the connectivity of sub-parts including a parent-child relationship starting from the main shape and decreasing monotonically along parts moving away from the main shape is provided, and
- the partition and saliency measure is stable under volume preserving shape deformation (with no topology change).

Both features allow for a straightforward classification of the extracted sub-shapes into visual saliency clusters, which then aids any subsequent computations such as correspondences or local symmetries (see Fig. 14).

After presenting related work and giving background in Section 2, we define WEDF on 1D skeletons embedded in 3D shapes in Section 4, describe its use for skeletal decomposition in Section 5, and introduce the mapping of the WEDF to the surrounding 3D mesh in Section 6. We then demonstrate the applications of the method for 3D shape decomposition and saliency-based evaluation of sub-parts.

## 2. Related work

### 2.1. 3D shape saliency

Saliency, or saliency, is a concept that has been widely studied in the literature. Borji and Itti, in their important state-of-the-art review [7] define saliency as a characterization of *parts of a scene – which could be objects or regions – that appear to an observer to stand out relative to their neighboring parts*. Most of the algorithms that characterize saliency in images use bottom-up approaches, determining what *stands out* using low-level cues such as color contrast [8].

A similar approach has been adopted in studies of 3D objects saliency. In their pioneer work, Lee et al. [9] define a mesh saliency measure based on local surface curvature at multiple scales. Many more

methods have then be devised following this local approach. Song et al. [10] exploit the apparent correlation between the log-spectrum of 3D mesh geometric Laplacian. They later extended their work, introducing a pooling scheme to introduce a more global saliency [11]. Nouri et al. compute local descriptors on each vertex and then a measure of saliency based on a weighted average at different scales [12]. These methods usually compare their results with outcome of user studies, that either captured eye fixations of users watching 3d objects [13], or matched manually designated interest points [14].

In our work, we do not aim at computing a local measure of saliency that would characterize the shape at the vertex level, but rather try to define the relative importance of parts of the shape relative to the whole, thus following Hoffman and Singh’s footsteps [2]. These authors emphasize the need for a measure of the saliency of visual parts, especially in the context of object recognition. A part saliency is, according to the authors, correlated to how much recognizing this part helps recognizing the object. The authors propose three features to compute parts saliency: relative size of the part, degree of protrusion and boundary strength. As a consequence, for the matter of shape understanding, it is important to be able to discriminate a shape into parts and get a sense of the relative importance of those parts. This is a natural outcome of the work we propose in this paper.

### 2.2. 3D shape decomposition

3D shape segmentation, or decomposing a shape in sub-parts, is an important research topic offering a wide range of methods emerging from its numerous applications such as compression, reverse engineering, editing, and comparison (see [15] for a recent survey). Supervised methods are based on user-segmented data and learning methods, offering a perceptually coherent segmentation as in [16,17]. Unsupervised methods make use of other tools, such as clustering methods or region-growing based on various shape characteristics (e.g. [18,19], or see [15] for a comparison), spectral analysis (e.g. [20]), spatial subdivision methods or explicit boundary extraction (e.g. [13,21]), or probabilistic models [22,23]. Other segmentation methods are based on topological information, like Reeb graphs [24] but depend on the function chosen for defining the graph [25]. Skeletons have been applied for segmentation on point clouds, where either the resulting skeletons are not connected [26,27] or do not provide a consistent segmentation [28], and on surface meshes [29], using the skeleton to

define a dense set of possible segmentation cuts from which a subset will be chosen based on boundary geometry [30]. Very few of these segmentation methods explicitly profit from the innate parent-child relation between sub-parts, or measure for the relative saliency over the parts they obtain.

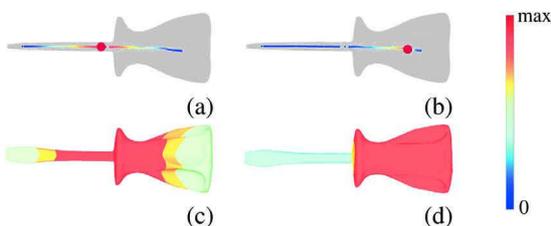
While most of the 3D methods described above do not make use of skeletons, in part because of their perceived instability and difficult to compute, recent methods for 2D shape analysis, decomposition and clustering [5,6,31] successfully exploit the geometry of the Blum medial axis for various shape tasks. These methods have shown that the use of the skeletons can be complemented with properly chosen functions that provide robustness and stability. One example of such a function is EDF [31] (*Extended Distance Function*) on 2D shapes which is a useful measure for pruning and has already been generalized to 3D, enabling improved 3D skeleton computations [32]. Another example is WEDF [5] (*Weighted Extended Distance Function*), which defines a measure for 2D sub-parts saliency. It has recently been validated through perceptual studies about cognitively meaningful shape importance [33] and used for 2D shape decomposition and similarity detection [6].

Our work builds on these 2D structures. We generalize WEDF to curve-skeletons on the medial axis of a 3D shape, develop an efficient algorithm and provide a perceptually consistent, graded saliency-based shape decomposition based on the variation of the WEDF along the curve-skeleton (see Fig. 2).

### 2.3. Depth functions in 3D

In 3D, the medial axis of a surface, which consists of the loci of centers of all enclosed maximal balls of the surface, is no longer simply a curve skeleton but generically includes two-dimensional components called medial surfaces that intersect transversally creating non-manifold regions. Each point on the medial axis is still associated to a radius function whose value is the shortest distance to the shape boundary, namely the radius of the corresponding medial ball. A direct extension to 3D of a depth function like EDF is not obvious as there is no single medial axis endpoint associated to a medial component. For this reason, it is more convenient to define functions on some kind of medial skeleton rather than on the full medial axis. Dey and Sun [34] are the first to introduce a mathematical definition and an algorithm to compute a medial curve skeleton based on singularities of a so-called medial geodesic function. Yan et al. [32] propose another definition and algorithm of a medial curve skeleton for 3D, based on ET and a burning process over the medial axis. They generalize their work in 2D to define the following functions for 3D on piecewise linear medial axes such as those generated by discrete surface meshes: *medial burn time* (MBT), *erosion thickness* (ET) and *curve burn time* (CBT).

In [32], these functions are meant to aid in pruning the medial axis



**Fig. 2. Shape center.** 3D EDF (a) and 3D WEDF (b) and their associated decomposition (c,d) on a 3D shape. The maximum WEDF value (b, red dot) is centered inside the handle, while the maximum EDF value lies in the tail (a). Our decomposition of the shape into hierarchical saliency levels based on each of the two measures given in second row, highlights the fact that considering length, like EDF (c), for measuring saliency does not align with perception, whereas considering saliency based on WEDF does (d). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

to obtain a simpler version that is more useful in applications, and to obtain a 1D skeletal shape representation. In contrast, our paper aims to provide a new function on the skeleton, similar to WEDF in 2D, able to measure part saliency. To this end, we build on MBT and CBT and derive an efficient algorithm not only for calculating a 3D WEDF along the skeleton but also for returning the function back to the surface mesh in order to determine a saliency-based shape decomposition and to structure the shape hierarchically.

### 3. Local analysis of the shape around the axis

Before presenting our algorithms in Section 4, let us briefly recall the existing depth functions relevant for our work and introduce an original interpretation of these three measures as a non-linear PCA.

The first two functions are defined on the full medial axis, and are used to extract a curve skeleton. The third is then defined on the extracted skeleton.

- The *medial burn time* (MBT) is defined as the geodesic distance of a medial axis point to the shape boundary associated to the closest extremity of the widest medial axis sheet containing this point [32].
- The *erosion thickness* ET is defined as the difference between the width and the height: at a point  $x$  in a medial axis  $M$ ,  $ET(x) = MBT(x) - R(x)$ . ET is used in [32] to determine the scale at which to extract the curve skeleton.
- The *curve burn time* (CBT) is defined only for skeletal points, and measures the distance to the closest extremity of the shape on the longest skeleton curve containing the given medial skeleton curve point. CBT thus corresponds to a length computation, and models the direction in which the shape extends the most. Yan et al. [32] use CBT as a way to prune the medial skeleton to obtain a simplified version.

While these measures were intended for simplification of the medial axis, we show that they can be used as a powerful tool for shape understanding. In particular, the three functions  $R(x)$ ,  $MBT(x)$ ,  $CBT(x)$  along a curve skeleton (Fig. 5), when taken together with the paths through the medial axis whose lengths they represent, provide a non-linear version of a PCA decomposition of the shape: at a point  $x$  on the curve skeleton,  $CBT$  gives *length*,  $MBT$  gives *width* and  $R$  gives *height*, and the accompanying curves give the principal directions through the shape starting at  $x$ . Note that, unlike in PCA, the directions considered are not linear but follow the curvature of the medial axis along each important direction. Moreover, integrating these quantities (width and height) over the principal direction provides our volume-based saliency measure, the 3D WEDF as introduced in the next section.

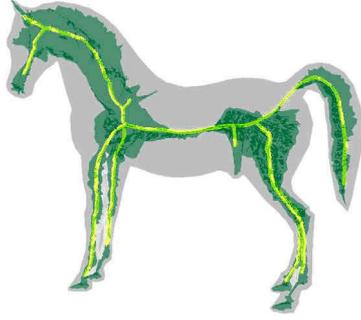
### 4. 3D weighted extended distance function

We now develop a definition for the weighted extended distance function (WEDF) in 3D for piecewise-linear medial surfaces.

The customary challenge in working with discrete medial surfaces is that, in practice, the connectivity of a discrete 3D medial axis is often not consistent with the theory, as noise in the surface mesh propagates to the medial axis sheet by generating regions where the connectivity is rarely manifold even when the underlying continuous model would be smooth. This prevents finding a direct, continuous path to the boundary on a medial axis sheet generated from a surface mesh. Such a path is however essential to define a partial order on the medial axis points. The next section introduces the tools necessary to overcome this difficulty.

#### 4.1. Preliminaries

The input of our algorithm is a closed, manifold triangle mesh representing the surface of a 3D object. From the mesh, we can compute a



**Fig. 3. Medial axis** (dark gray) and a **curve skeleton** (yellow) of a 3D shape. The medial axis is obtained using [35] and curve skeleton using [32]. This curve skeleton preserves the topology of the shape (light gray) and captures its features (ears, head, legs, tail). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

discrete version of the medial axis (e.g. [36,37]), and for each medial point  $x$  its burn time,  $MBT(x)$ , as computed in [32]. Note that each point  $x$  in the medial axis is also the circumcenter of an associated tetrahedron given by a Delaunay tetrahedralization. The set of all tetrahedra associated to all medial axis points provides a partition of the volume within the mesh.

We extract a curve skeleton that is a subset of the medial axis, a connected curve network capturing the structure of the shape. We assume the curve skeleton has the same topology as the shape, and is as centered as possible. In particular, the curve skeleton will be properly contained in the interior of the shape since it lies on the medial axis, see Fig. 3. Our method also works for non-centered skeletons, but the result may be less precise (a comparison is shown in Section 7.2). In our implementation, we use the method from [32] to compute the curve-skeleton, but we also consider mean curvature skeletons from [38,39] extracted via mean curvature flow, and curve-skeletons from [40,41]. Our methods are consistent for all types of curve-skeletons tested. Note, that we present results for surfaces of genus 0 generating curve skeletons with tree structures. In this case, there exists a single path on a curve skeleton between any of its two points. In other words, the path between an extremity of the curve skeleton and a skeleton point is unique.

#### 4.2. Defining 3D WEDF on a curve skeleton

We now define 3D WEDF values on a curve skeleton in order to provide a salience measure of subparts. In the discrete setting, where the shape is a polygonal curve and the medial axis a piecewise linear curve skeleton, WEDF is computed as the areas of the shape parts transversed along the WEDF path from the extremity  $y$  to the medial axis point  $x$ . An extension to a discrete 3D setting would be to sum up the volume traversing the medial axis from a part extremity until skeleton point. The direction of medial axis traversal is not well defined, however, because of the lack of correspondence between a curve skeleton point and points on a medial sheet (and also on the surface mesh). This prevents us from extending the 2D WEDF definition to 3D directly.

We pause to summarize the two main challenges we face and how we propose to solve them, before defining WEDF in more detail:

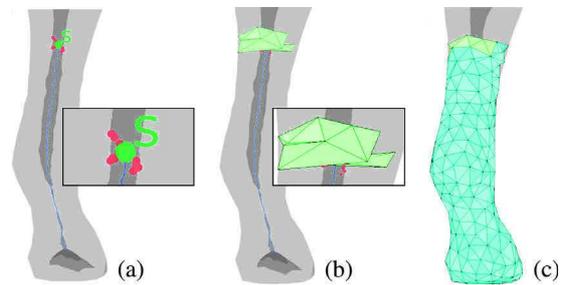
- 1. Skeleton traversal.** Since the discrete medial axis contains two-dimensional sheets so the medial axis cannot be used for an EDF-like path traversal of the 3D shape. Our solution is to base our computations on a one-dimensional skeleton inside the 3D shape. The CBT values can then be extended to CBT values for medial points on the manifesting path for the medial burn time (MBT) of the skeletal points.
- 2. Volumetric shape partition.** While in 3D, as in 2D, there is a correspondence between points on the medial axis and the points on

the surface mesh, namely those on the circumcircle of the Delaunay tetrahedralisation of the volume, smoothing the non-manifold noise on the medial axis often breaks the correspondence. As is true for the Delaunay triangulation and area in 2D, the theoretical Delaunay tetrahedralisation in 3D partitions the volume of the shape into regions associated to a medial point. In removing the noise from the medial axis, however, the correspondence between medial points and tetrahedra may be broken, and so we must establish a new correspondence between points on the medial axis and tetrahedra. Our solution depends on the skeleton traversal described above: we associate tetrahedra for multiple medial points to a single skeletal point so that summing volumes over all skeletal points results in the volume of the entire shape.

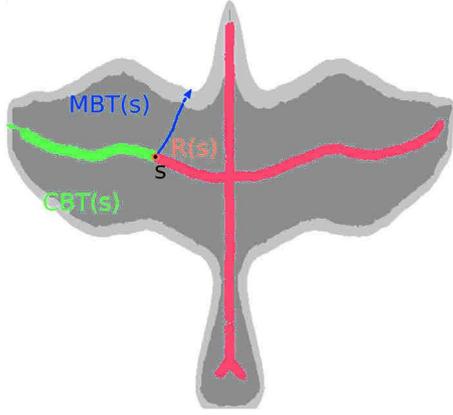
Let us explain our solution to these challenges in more detail now. In the following,  $m$  denotes a point on the medial axis  $M$  and  $s$  denotes a skeletal point, a point on the curve skeleton (and also on the medial axis, as the curve skeleton lies on the medial axis). As the curve skeleton is traversed, the WEDF must account for the volume of the entire shape part between a skeletal point and the surface boundary, not just the volume of the shape associated to the tetrahedra of the curve skeletal points. For each skeletal point  $s$ , we obtain the medial burn time ( $MBT$ ), which is the time required for a unit-speed grass-fire lit on the boundary of the medial axis to completely expose a medial point, and consider its burn path  $b_s$ . See [32] for a detailed description of the algorithm and rigorous definition in the discrete setting. The burn path is the path of length  $MBT(s)$  extending from the point  $s$  to a boundary point that represents the path of the fire front that exposed  $s$ , see Fig. 5. Let  $T(s)$  be the volume that is assigned to a skeleton point  $s$ : the sum of volumes of the tetrahedra of the medial points along the burn path of  $s$  (green tetrahedra in Fig. 4). Then the WEDF-volume assigned to  $s$ ,  $WEDF(s) = T(s) + \sum_{z \in \text{Skeleton\_path}} T(z)$ , is the volume  $T(s)$  plus the volumes  $T(z)$  associated to all preceding skeletal points  $z$  that are closer to the skeleton extremity than  $s$  along the CBT path to  $s$ . Then  $WEDF(s)$  can be thought of as the integral along the path along which CBT is computed within the curve skeleton, of the other two PCA-like values: radius and MBT. These volumes form the basis of the 3D WEDF computation.

We now formally define WEDF in the piecewise linear case. Let  $s$  be a point on a curve skeleton  $C$  of the medial axis  $M$  associated to a surface mesh  $S$ . Then the value of  $WEDF(s)$  is given by one of the following:

- If  $s$  is a skeleton extremity,  $WEDF(s)$  is the volume of the cap



**Fig. 4. Volume contribution of a skeleton point.** The point  $s$  is a point on the curve skeleton. In red, the medial axis points (or tetrahedra circumcenters) associated to the point  $s$  thanks to the distance  $D$  (Eq. (1)) (a). In green, the tetrahedra whose circumcenter is one of the medial axis points previously marked in red (b).  $T(s)$  is the sum of the volume of these green tetrahedra.  $WEDF(s)$  is the volume of the subshape whose curve skeleton starts at  $s$  and goes towards closest extremities of the curve skeleton, here the volume of the ankle and foot below  $s$  (c). The tetrahedra whose volume contributes to WEDF of previous points ( $WEDF_{prev}$  in Eq. (2)) are in light blue. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)



**Fig. 5. Depth functions at 3D curve skeleton point  $s$ :** curve burn time  $CBT(s)$ , medial burn time  $MBT(s)$  and radius  $R(s)$ . Together they provide a global analysis of the shape around  $s$ , respectively *depth*, *width* and *height* ( $R(s)$  towards us).

between  $s$  and the corresponding points on  $S$ .

- If  $s$  is a junction point where multiple branches of  $C$  come together, then  $WEDF(s)$  corresponds to the sum of the volumes of all incoming branches, namely those branches that are closer to an skeletal extremity than  $s$ , together with the volumes of the tetrahedra assigned to  $s$ .
- If  $s$  is a generic branch point of a branch  $b$  (neither an extremity nor a junction), then  $WEDF(s)$  is the sum of the volumes assigned to the neighboring point in  $b$  that is closer to an extremity than  $s$ , together with the volumes of the tetrahedra assigned to  $s$ .

Note that because MBT is not defined in the continuous setting [32], we cannot provide a continuous definition for WEDF.

#### 4.3. Computing WEDF on the 3D curve skeleton

In practice, 3D WEDF is computed for an input surface mesh, on a curve-skeleton consisting of points on the medial axis of the surface mesh. The sampling density of the medial axis or skeleton points may affect the smoothness of the WEDF, but not its actual value. Above, we discuss the volume  $T(s)$  assigned to a skeletal point  $s$ . More precisely,  $T(s)$  is the sum of the volumes of the Delaunay tetrahedra whose circumcenters  $m \in M$  are closest to  $s$  according to the following distance function:

$$D(m, s) = | |MBT(m) - MBT(s)| - d(m, s) | \quad (1)$$

where  $d$  is the Euclidean 3D distance and  $MBT$  the medial burn time. Note that because MBT is defined only for discrete medial axes, we can define WEDF only in the discrete setting as well. Fig. 4 shows the tetrahedra and their circumcenters associated to a skeleton point  $s$ . Note that the proposed distance function  $D$  ensures that a medial point  $m$  is necessarily contributing to a skeleton point  $s$  that lies on the same sheet of the medial axis, as illustrated in Fig. 6.

The definition of  $T(s)$  ensures that points  $m$  associated with  $s$  are lying along the burn path  $b_s \subset M$  reaching the closest border of the medial axis according to the medial burning time.  $T(s)$  therefore measures the volume associated to all points on the burn path of  $s$  to the shape boundary, see Fig. 4. By construction, each medial point  $m$  is associated to a unique skeleton point  $s$ . Taking Delaunay tetrahedra across all  $m \in M$  gives the volume of the entire shape, therefore taking  $T(s)$  across all  $s \in C$  partitions the shape volume into tetrahedra sets associated to each  $s \in C$ .

To compute WEDF values along the skeleton, we begin at a skeletal extremity and progress inward. During the skeleton traversal, the WEDF value at a point is then computed from the value at the previous point:

$$WEDF(s) = T(s) + WEDF_{prev} \quad (2)$$

$$WEDF_{prev} = \begin{cases} 0 & \text{at extremities;} \\ WEDF(y) & \text{if } s \text{ is a regular point, and } y \\ & \text{its neighbor with computed WEDF;} \\ \sum_{i \geq 1} WEDF(y_i) & \text{if } s \text{ is a junction, and all neighbors } y_i \\ & \text{but one } y_0 \text{ have a WEDF value.} \end{cases}$$

Here in Eq. (2), we define a *summed WEDF*, since at a junction  $s$  of the curve-skeleton the volume  $WEDF(s)$  is here the *sum* of values from all incoming skeleton branches. This differs from the 2D WEDF definition [5], where *max* values at branch junctions are taken. We use the *summed WEDF* because it is necessary to insure its robustness to surface noise. Indeed, with a 3D max-version of WEDF, small branches of the skeleton originating from noise would capture the volume erroneously around them (by catching tetrahedra volume previously assigned to a main branch) and thus diminish the WEDF on the global shape. An illustration of the change in WEDF due to small branches is given by Fig. 7. Conversely, in the summed version the volume captured by the small branch is still contributing to the volume of the main branch, and therefore, WEDF values on main branches are stable.

Fig. 8 shows a color-coding of the WEDF function on curve skeletons for 3D objects. When the object and its curve skeleton have some symmetries, the WEDF remains symmetric. The WEDF value is maximum at a point whose value is the total volume within the surface of the object, and then decreases on any path towards curve skeleton extremities. As in 2D, WEDF along a skeleton path between two extremities is unimodal.

## 5. Clustering the skeleton into parts according to salience

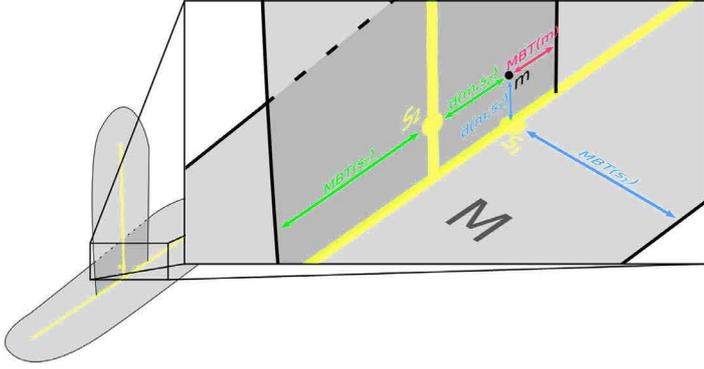
To demonstrate one avenue for application of our WEDF computation, we use the WEDF values defined on the skeleton to develop a salience measure for each point on the surface mesh. Examples of part salience on some well-known 3D shapes are shown in Section 7.

### 5.1. Partitioning skeleton points into salience levels

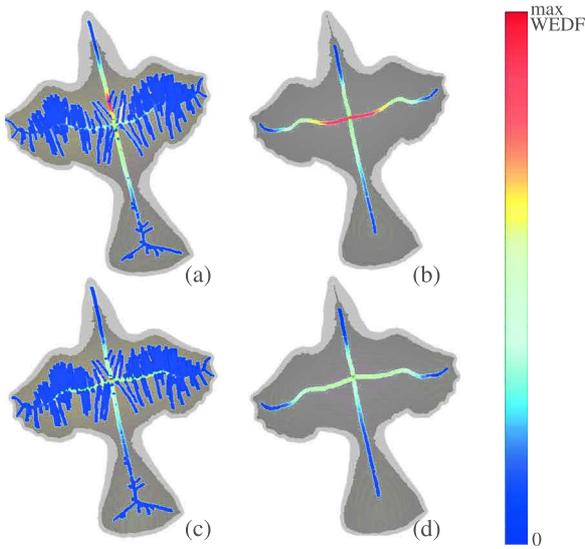
We perform unsupervised clustering on the WEDF values in order to identify parts at similar salience levels. The clustering we propose is a two-step process, where the first step determines the *main shape* of the skeleton and the second step identifies *features* with decreasing saliency values. For a genus 0 surface mesh, the skeleton has a tree-like connectivity, where the root corresponds to the highest 3D WEDF value. Moving away from the root, the salience values monotonically decreases along the branches of the skeleton.

*Identifying the main shape.* We select main skeletal points as a subset of the skeletal points in two steps. On the 3D WEDF values, we perform a seeded  $k$ -means clustering with two clusters where the seeds are the minimum and the maximum WEDF values of the junction skeletal points being clustered. We then associate all remaining skeleton points to these two clusters based on their WEDF values. All skeletal points belonging to the cluster with the centroid of largest WEDF value are selected as the main shape points.

*Identifying features.* The other 3D WEDF clusters are determined among those WEDF values of junction points and neighbors which are not in the main shape cluster: we again apply a  $k$ -means clustering with  $k$  in a given interval (in practice, we take  $k \in \{2, 3, \dots, 20\}$ ). An estimation of the best number of classes  $k$  is given by gap analysis [42]. Each cluster represents a salience level relevant for the shape. The level value  $L$  is chosen as the median WEDF value of the cluster: larger 3D WEDF values correspond to higher salience levels, i.e. closest to main shape, while



**Fig. 6. The distance function  $D$ .** In gray, two intersecting medial axis sheets belonging to the medial axis  $M$ . We distinguish the point  $m \in M$  and  $s_1$  and  $s_2$  as two points on the skeleton  $C$ . The skeleton lies on the medial axis. Despite  $s_1$  being closer to  $m$  than  $s_2$  in Euclidean distance, (i.e.,  $d(m, s_1) < d(m, s_2)$ ),  $m$  will be associated to  $s_2$  because  $D(m, s_2) < D(m, s_1)$  (since  $MBT(m) - MBT(s_2) \approx d(m, s_2)$  and  $MBT(m) - MBT(s_1) \neq d(m, s_1)$ ).



**Fig. 7. Robustness of the summed WEDF.** Two different curve skeletons for the same 3D shape, one with many branches (a,c) and a more filtered version with only main branches (b,d). Top row (a,b): WEDF is computed by taking the max value of incoming branches at intersections [6]. we observe a lack of robustness; the distribution of WEDF is different on the main branches of the two skeletons. In particular, the branch containing the greatest WEDF value is not the same. Bottom row (c,d): our proposed *summed* WEDF computed by summing WEDF values of incoming branches at junctions. The distribution of WEDF remains the same along the main branches of the skeletons since the volume counted on small branches also contributes to the main branch value.

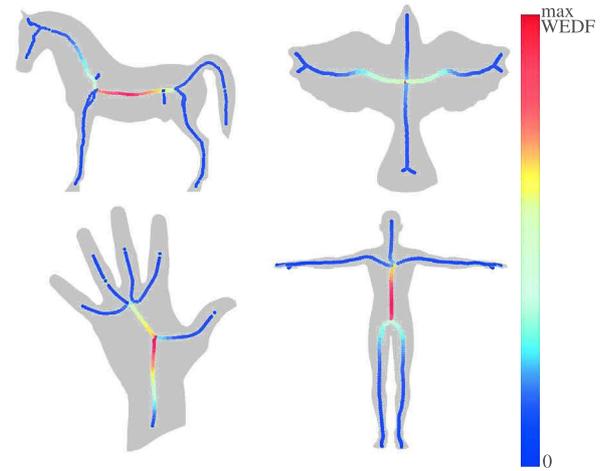
lower values are less significant, corresponding to lower salience levels.

**Clustering skeleton points.** The curve skeleton is partitioned into clusters according to the clustering of its 3D WEDF values. We thus obtain a partition of the curve-skeleton into salience levels. Note that any path on the curve skeleton going from a point belonging to the main shape to an extremity point will traverse levels of salience in decreasing order, see Fig. 9. This partial ordering of points of the curve skeleton, coherent with the tree topology of the medial axis, is the basis for the graded salience-based parts decomposition described in the next section.

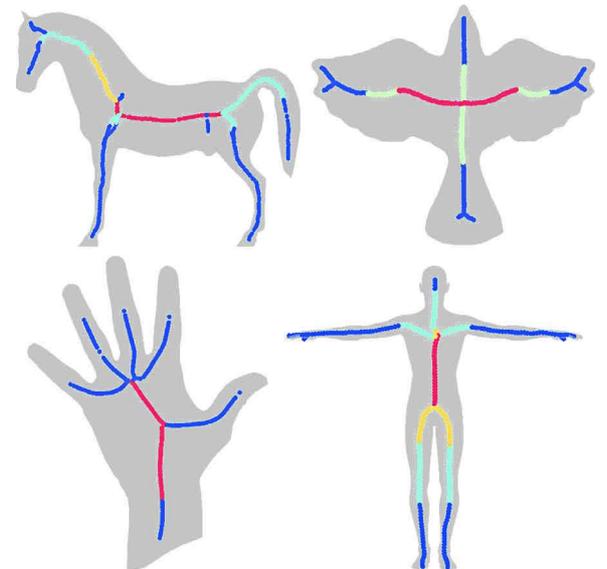
This process allows cross-parts comparison of salience for all parts, and is therefore not restricted to those with junctions: for example, the thumb in the hand of Fig. 9 is identified as a feature even if there is no junction in the curve skeleton, because its salience level matches the salience level of the other fingers.

## 6. Saliency measure for 3D shape surface decomposition and hierarchy

After having decomposed the skeleton into graded skeletal sub-parts



**Fig. 8. WEDF of 3D shapes.** Low values are colored in dark blue high values in dark red. We can see that the WEDF grows slowly along branches that are surrounded by a limited volume (such as legs, arms, tail or finger) and grows faster on branches surrounded by a higher volume (belly or palm for instance). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)



**Fig. 9. Decomposition of the curve-skeleton into graded salience levels,** from main shape (dark red) to less significant parts (dark blue). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

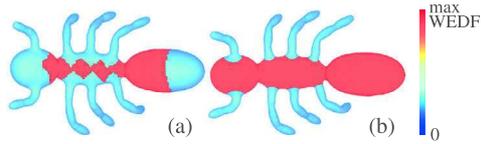


Fig. 10. The naive tetrahedra clustering (a) and our proposed clustering (b). We can see that the naive clustering gives a decomposition where sub-parts of lower saliency (like antennae or legs) are indeed entering into the parent shape of higher saliency (here head or body). The naive tetrahedra clustering is therefore not sufficient. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

(Fig. 1(d)), the second step of our method is transfer the parts decomposition to the shape’s surface, in order to partition the mesh into graded parts labeled by their saliency (Fig. 1(f)).

This transfer is a challenging process. A naive method would map the cluster level of curve skeleton points onto medial axis points using the distance function introduced for the computation of WEDF (see Eq. (1) in Section 4.3). And, since each medial axis point is the circumcenter of a tetrahedra, assigning the cluster level value to its tetrahedra would directly lead to a partition of the volume of the shape into cluster levels (Fig. 4). However, it can be observed that lower level clusters, corresponding to features of lower saliency, carve holes along the branches inside the parent shape: Delaunay tetrahedra whose circumcenters are medial points close to the center of the main shape may be associated to a curve skeleton point belonging to a lower level branch (we see this effect on the left image of Fig. 10 where green triangles emerge from inside the main shape near attachment area of legs).

We therefore propose a two-step process to extend transfer saliency levels from the skeletal points to the surface, using vertices instead of tetrahedra. We first assign each vertex of the surface a saliency level by considering all medial points whose tetrahedra contain the vertex, and then determine a level value for the triangular faces from the levels of the associated vertices.

### 6.1. Assigning saliency levels to mesh vertices

We start with a surface mesh  $S$  and a curve skeleton  $C$ , where each skeletal point  $s$  is attributed a saliency level  $L(s)$ . Each  $s$  belongs to a cluster  $cluster(s)$ , and its associated level  $L = L(s)$ , the median of the WEDF values within the cluster  $cluster(s)$ . For each mesh vertex  $v \in S$ , we compute the skeleton point  $s$ , which is the closest to  $v$  using the following distance function:

$$F(v, s) = \frac{d(v, s)}{L(s)} \quad (3)$$

with  $d(v, s)$  the Euclidean distance between  $v$  and  $s$ .

This function warps the space to favor clusters of higher value: skeleton points belonging to high-level clusters have an  $F$ -distance smaller than the Euclidean distance to the mesh vertices, while skeleton points belonging to clusters with lower value have an  $F$ -distance distance greater than the Euclidean distance to the mesh vertices. This way, curve-skeleton points belonging to higher saliency-level clusters attract boundary points to the more significant part in terms of saliency (body) rather than attaching to a less significant feature. See Fig. 10, where mesh vertices are wrongly attached to the leg using the Euclidean distance (left), but correctly labeled with a saliency attached to the body (right).

### 6.2. Assigning saliency levels to mesh faces

After this first step of assigning mesh vertices to saliency levels of skeleton points, we label the mesh faces using the levels assigned to the three vertices belonging to a face. If all vertices are assigned to the same

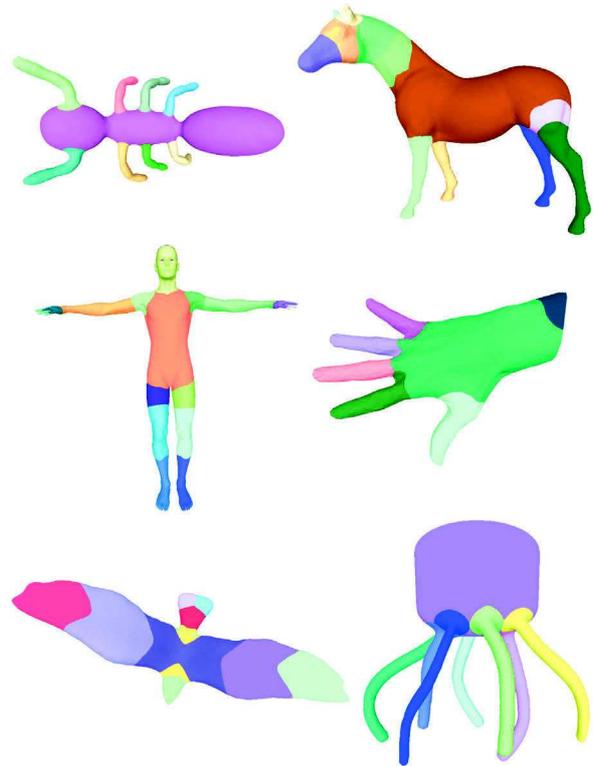


Fig. 11. Shape segmentations based on our hierarchical decomposition; here each connected cluster on the curve skeleton generates a different part of the mesh partition.

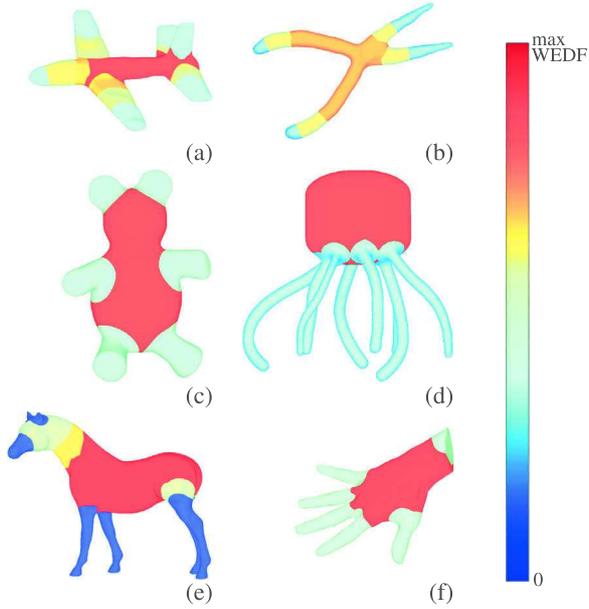
level, we assign the face to this level, that is, the face saliency value is the cluster saliency. In saliency transition zones, where vertices of a face are assigned to different levels, we subdivide the face (1–4 subdivision) using points linearly interpolated on the three edges according to saliency levels of the vertices, and assign levels to subregions accordingly. This generates a smooth transition between parts of the shape belonging to different levels, as shown in the following experiments.

## 7. Experiments

### 7.1. Validation

First, Fig. 11 shows the partition of different meshes created by considering each cluster on the skeleton as a different part, *without* grouping by saliency value. Note that the decomposition is coherent, and when the shape is close to symmetric, the decomposition preserves this symmetry. Also observe that different parts, like the fingers of the hand, or the arms of the octopus, are well separated (thanks to the topology of the curve medial axis).

In Fig. 12, we have now grouped parts according to saliency, so each part with the same color has been determined to be at the same saliency level. Moreover, because the skeleton has a tree structure, the saliency levels are monotonically decreasing from the main shape to the extremities of the shape, leading to a graded partition into subparts. We can see in all these results that the core of the shape is well-centered, while shape details are localized at extremities. Also, the shape parts that are similar (legs, fins or antennae) are assigned to the same hierarchy level. The proposed segmentation then gives not only a partition of the shape, but also assigns a level of saliency to each feature that is global on the shape, and known to be consistent with perception in 2D [5,33].



**Fig. 12.** Shape decompositions on several shapes, with saliency-based partial ordering of the extracted sub-parts. The color of a part represents its centroid value. The main shape is often centered in the shape (a,b,c,e,f) or located at the part with greatest volume (d). Details and part of lowest saliency are at the extremities of the shapes (wings, ears, legs, fingers). While (a) and (e) are being decomposed in four levels, (b) is cut in three levels and (c), (d) and (f) only need two levels. Note the capacity of our approach to model shape with a medial axis close to the curve skeleton (e) but also shapes with no tubular parts (f) having a two dimensional medial sheet.

## 7.2. Robustness of the graded partition

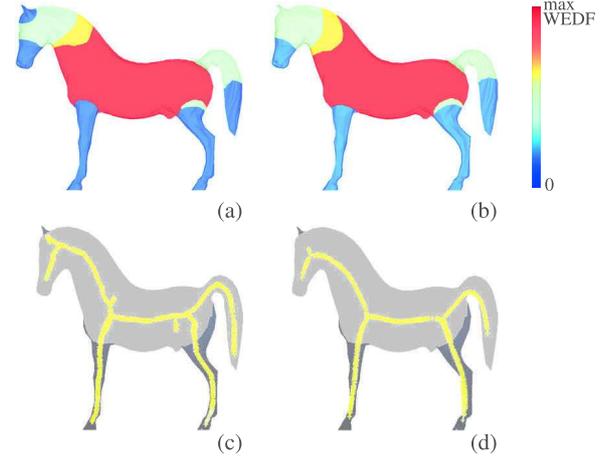
We also recall that the proposed graded segmentation is robust. First, computing the saliency level on the curve skeleton reduces the instability of the medial axis to boundary noise. A curve skeleton may have a lot of small branches, or be simplified. We have seen in Fig. 7 that the proposed computation of WEDF (summing the incoming values at junctions) provides stable WEDF values on the main branches, and thus insures a consistent segmentation.

We also experiment with different curve skeletons as input. Curve skeletons are not uniquely defined for 3D objects, and the computation of a curve skeleton is not part of our work. Even though our method works successfully for different curve skeleton types, the quality of the curve skeleton plays an important role, see Fig. 13. A good curve skeleton is well-centered and captures all small protrusions with branches associated to each feature (left) whereas a coarser skeleton leads to a coarser graded decomposition (right).

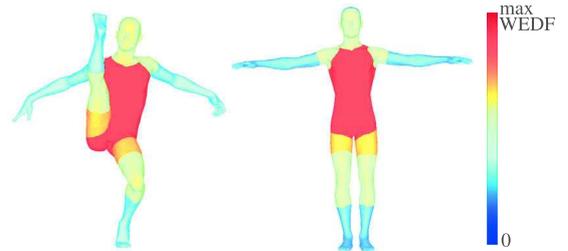
Fig. 14 shows that our segmentation is naturally stable under constant volume shape deformation, since the WEDF measure generating the saliency levels is volumetric. In other words, the saliency levels stay consistent as moving parts articulate. The coloring between the two human figures remains consistent for the legs and arms, with only slight changes on the shoulders. Similarly, in Fig. 15 the decomposition is very consistent across the different poses (a) to (d).

## 8. Discussion and conclusion

We have produced a graded, saliency-based parts partition from a importance measure on a curve skeleton of the shape integrating local PCA-like representation of a shape on the skeleton. Our method determines a level of relative importance for each sub-part, and labels features of the same saliency level as belonging to a group with the same perceptual importance. Moreover, because the decomposition respects the topology of the skeleton, it produces a parent-child



**Fig. 13.** Comparison between two segmentations (a,b) using different curve skeletons (c,d). The (a) decomposition uses the skeleton from [32], which is better centered and which correctly captures small protrusions (such as ears), while the first decomposition (b), uses the mean curvature skeleton (d) from [38], and is neither well-centered (see the tail) nor capturing correctly the small protrusions (such as ears). The (a) segmentation gives better results as small protrusions create smaller parts and as parts (legs, tail) are more finely cut, confirming our preference for the skeleton (c).



**Fig. 14.** Stability of the segmentation relative to pose variation without change of topology.

relationship between parts in a given shape. At the highest level, the main shape contains the barycenter of the shape. Moving away from the center, parts decrease in saliency as a parent level gives way to a child. For genus 0 shapes, having a skeleton with tree structure, the parent of each part will be unique. As a result, the connectivity produced by the skeleton allows for ordered parts relationships in a way that connectivity induced by cluster adjacency computed directly on the surface does not: for example, the ears of the horse on Fig. 12 have two adjacent regions on the head, themselves neighbors, which would lead to a triangle connection (loop) between the three adjacent regions when using surface connectivity: skeletal connectivity results in two ears each connected to a head. Thus our decomposition is enriched by both a topological relationship between subparts, and an importance value attached to each subpart.

While our decomposition has consistent results on different curve skeletons, it still depends on the quality of the curve skeleton used to capture the shape. Indeed, as shown in Figs. 3 and 13, the results may vary depending on the skeleton's geometry and topology. However, for volume preserving deformation, the proposed decomposition is robust and gives a consistent labeling of the shape parts as long as the topology is preserved. As expected, the generalization of WEDF to 3D leads to an analysis of the shape that seems consistent with perception, as shown in 2D in [5,33]. Future work will include a user study to confirm the perceptual coherence.

In future work, we also plan to extend the present work in order to analyze the similarity between segmented parts belonging to the same saliency level, by using for example a combination of PCA-like measures

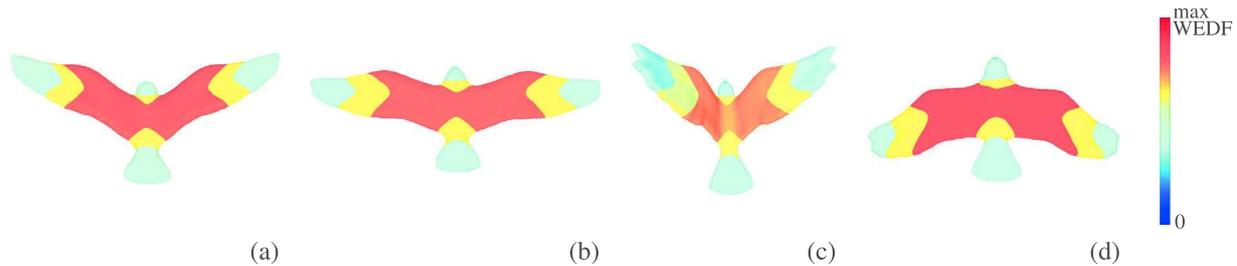


Fig. 15. Several shape decompositions of a bird shape in different positions. The decomposition is stable under deformations. This part decomposition into saliency levels may therefore be used in animation.

on the skeleton. This would enable to produce inputs to recent methods, such as SimSelect [4] or Pattern-Driven Colorization [43], where the user is required to annotate candidate similarities by hand.

Like in the 2D setting [33], our decomposition and saliency could be validated by a user study asking to differentiate the different parts of a 3D object and tag the main shape, and the importance of each different part.

## References

- [1] G.H. Bower, A.L. Glass, Structural units and the redintegrative power of picture fragments. *J. Exp. Psychol.: Hum. Learn.Memory* 2 (4) (1976) 456.
- [2] D.D. Hoffman, M. Singh, Saliency of visual parts, *Cognition* 63 (1) (1997) 29–78.
- [3] V. Léon, N. Bonneel, G. Lavoué, J.-P. Vandeborre, Continuous semantic description of 3d meshes, *Comput. Graph.* 54 (2016) 47–56.
- [4] E. Guy, J.-M. Thiery, T. Boubekeur, Simselect: similarity-based selection for 3d surfaces, *Comput. Graph. Forum* 33 (2) (2014) 165–173.
- [5] L.J. Larsson, G. Morin, A. Begault, R. Chaine, J. Abiva, E. Hubert, M. Li, B. Paniagua, G. Tran, et al., Identifying Perceptually Salient Features on 2D Shapes, *Research in Shape Modeling*, Springer, 2015, pp. 129–153.
- [6] K. Leonard, G. Morin, S. Hahmann, A. Carlier, A 2d shape structure for decomposition and part similarity, 2016 23rd International Conference on Pattern Recognition (ICPR), (2016), pp. 3216–3221.
- [7] A. Borji, L. Itti, State-of-the-art in visual attention modeling, *IEEE Trans. Pattern Anal. Mach. Intell.* 35 (1) (2013) 185–207.
- [8] L. Itti, C. Koch, E. Niebur, A model of saliency-based visual attention for rapid scene analysis, *IEEE Trans. Pattern Anal. Mach. Intell.* 20 (11) (1998) 1254–1259.
- [9] C.H. Lee, A. Varshney, D.W. Jacobs, Mesh saliency, *ACM Transactions on Graphics (TOG)*, 24 ACM, 2005, pp. 659–666.
- [10] R. Song, Y. Liu, R.R. Martin, P.L. Rosin, Mesh saliency via spectral processing, *ACM Trans. Graph.* 33 (1) (2014) 6.
- [11] R. Song, Y. Liu, R.R. Martin, K.R. Echavarria, Local-to-global mesh saliency, *Vis. Comput.* (2016) 1–14.
- [12] A. Nouri, C. Charrier, O. Lézoray, Multi-scale mesh saliency with local adaptive patches for viewpoint selection, *Signal Process. Image Commun.* 38 (2015) 151–166.
- [13] D.H. Kim, I.D. Yun, S.U. Lee, A new shape decomposition scheme for graph-based representation, *Pattern Recognit.* 38 (5) (2005) 673–689.
- [14] X. Chen, A. Saporov, B. Pang, T. Funkhouser, Schelling points on 3d surface meshes, *ACM Trans. Graph.* 31 (4) (2012) 29.
- [15] P. Theologou, I. Pratikakis, T. Theoharis, A comprehensive overview of methodologies and performance evaluation frameworks in 3d mesh segmentation, *Comput. Vision Image Understanding* 135 (2015) 49–82.
- [16] Z. Xie, K. Xu, L. Liu, Y. Xiong, 3d shape segmentation and labeling via extreme learning machine, *Comput. Graph. Forum* 33 (5) (2014) 85–95.
- [17] H. Benhabiles, G. Lavoué, J.-P. Vandeborre, M. Daoudi, Learning boundary edges for 3d-mesh segmentation, *Comput. Graph. Forum* 30 (8) (2011) 2170–2182.
- [18] Y.-K. Lai, S.-M. Hu, R.R. Martin, P.L. Rosin, Rapid and effective segmentation of 3d models using random walks, *CAGD* 26 (6) (2009) 665–679.
- [19] F. Bergamasco, A. Albarelli, A. Torsello, A graph-based technique for semi-supervised segmentation of 3d surfaces, *Pattern Recognit. Lett.* 33 (15) (2012) 2057–2064.
- [20] Y. Fang, M. Sun, M. Kim, K. Ramani, Heat-mapping: a robust approach toward perceptually consistent mesh segmentation, *IEEE Computer Vision and Pattern Recognition (CVPR)*, (2011), pp. 2145–2152.
- [21] O.K.-C. Au, Y. Zheng, M. Chen, P. Xu, C.-L. Tai, Mesh segmentation with concavity-aware fields, *IEEE Trans. Vis. Comput. Graph.* 18 (7) (2012) 1125–1134.
- [22] S. Pulla, A. Razdan, G. Farin, Improved curvature estimation for watershed segmentation of 3-dimensional meshes, *IEEE Trans. Vis. Comput. Graph.* 5 (4) (2001) 308–321.
- [23] A. Koschan, Perception-based 3d triangle mesh segmentation using fast marching watersheds, *Computer Vision and Pattern Recognition, Proceedings, 2 IEEE*, 2003, pp. II–27.
- [24] S. Berretti, A. Del Bimbo, P. Pala, 3d mesh decomposition using reeb graphs, *Image Vis. Comput.* 27 (10) (2009) 1540–1554.
- [25] S. Biasotti, D. Giorgi, M. Spagnuolo, B. Falcidieno, Reeb graphs for shape analysis and applications, *Theor. Comput. Sci.* 392 (2008) 5–22.
- [26] T. El-Gaaly, V. Froyen, A. Elgammal, J. Feldman, M. Singh, A bayesian approach to perceptual 3d object-part decomposition using skeleton-based representations, *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence*, AAAI Press, 2015, pp. 3762–3768.
- [27] X. Li, T.W. Woon, T.S. Tan, Z. Huang, Decomposing polygon meshes for interactive applications, *Proceedings of the 2001 Symposium on Interactive 3D Graphics*, ACM, 2001, pp. 35–42.
- [28] J. Kustra, A. Jalba, A. Telea, Shape segmentation using medial point clouds with applications to dental cast analysis, *Proc. of the 9th Int. Conf. on Computer Vision Theory and Applications (VISAPP 2014)*, (2014), pp. 162–170.
- [29] C. Feng, A.C. Jalba, A.C. Telea, Part-based segmentation by skeleton cut space analysis, *Int. Symposium on Mathematical Morphology and Its Applications to Signal and Image Processing*, Springer, 2015, pp. 607–618.
- [30] J. Tierny, J.-P. Vandeborre, M. Daoudi, Topology driven 3d mesh hierarchical segmentation, *Shape Modeling and Applications, 2007. SMI'07. IEEE International Conference on*, IEEE, 2007, pp. 215–220.
- [31] L. Liu, E.W. Chambers, D. Letscher, T. Ju, Extended grassfire transform on medial axes of 2d shapes, *CAD* 43 (11) (2011) 1496–1505.
- [32] Y. Yan, K. Sykes, E. Chambers, D. Letscher, T. Ju, Erosion thickness on medial axes of 3d shapes, *ACM Trans. Graph.* 35 (4) (2016) 38:1–38:12.
- [33] A. Carlier, K. Leonard, S. Hahmann, G. Morin, M. Collins, The 2d shape structure dataset: a user annotated open access database, *Comput. Graph.* 58 (2016) 23–30.
- [34] T.K. Dey, J. Sun, Defining and computing curve-skeletons with medial geodesic function, *Symposium on Geometry Processing*, 6 (2006), pp. 143–152.
- [35] J. Giesen, B. Miklos, M. Pauly, C. Wormser, The scale axis transform, *Proceedings of the Twenty-fifth Annual Symposium on Computational Geometry*, ACM, 2009, pp. 106–115.
- [36] N. Amenta, M. Bern, M. Kamvysselis, A new voronoi-based surface reconstruction algorithm, *Proceedings SIGGRAPH*, ACM, 1998, pp. 415–421.
- [37] A. Tagliasacchi, T. Delame, M. Spagnuolo, N. Amenta, A. Telea, 3D skeletons: a state-of-the-art report, *Comput. Graph. Forum* 35 (2) (2016) 573–597.
- [38] A. Tagliasacchi, I. Alhashim, M. Olson, H. Zhang, Mean curvature skeletons, *Comput. Graph. Forum* 31 (5) (2012) 1735–1744.
- [39] T.-C. Lee, R.L. Kashyap, C.-N. Chu, Building skeleton models via 3-d medial surface axis thinning algorithms, *CVGIP: Graph. Models Image Process.* 56 (6) (1994) 462–478.
- [40] M. Livesu, F. Guggeri, R. Scateni, Reconstructing the curve-skeletons of 3d shapes using the visual hull, *IEEE Trans. Vis. Comput. Graph.* 18 (11) (2012) 1891–1901.
- [41] M. Livesu, R. Scateni, Extracting curve-skeletons from digital shapes using occluding contours, *Vis. Comput.* 29 (9) (2013) 907–916.
- [42] R. Tibshirani, G. Walther, T. Hastie, Estimating the number of clusters in a data set via the gap statistic, *J. R. Stat. Soc.: Ser. B* 63 (2) (2001) 411–423.
- [43] G. Leifman, A. Tal, Pattern-driven colorization of 3d surfaces, 2013 IEEE Conference on Computer Vision and Pattern Recognition, (2013), pp. 241–248.