University of Groningen

Part-type Segmentation of Articulated Voxel-Shapes using the Junction Rule

Reniers, Dennie; Telea, Alexandru

Published in:
EPRINTS-BOOK-TITLE

IMPORTANT NOTE: You are advised to consult the publisher's version (publisher's PDF) if you wish to cite from it. Please check the document version below.

Document Version
Publisher's PDF, also known as Version of record

Publication date:
2008

[Link to publication in University of Groningen/UMCG research database](#)

Citation for published version (APA):

Reniers, D., & Telea, A. (2008). Part-type Segmentation of Articulated Voxel-Shapes using the Junction Rule. In *EPRINTS-BOOK-TITLE* University of Groningen, Johann Bernoulli Institute for Mathematics and Computer Science.

Copyright

Other than for strictly personal use, it is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license (like Creative Commons).

The publication may also be distributed here under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license. More information can be found on the University of Groningen website: <https://www.rug.nl/library/open-access/self-archiving-pure/taverne-amendment>.

Take-down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Downloaded from the University of Groningen/UMCG research database (Pure): <http://www.rug.nl/research/portal>. For technical reasons the number of authors shown on this cover page is limited to 10 maximum.

Part-type Segmentation of Articulated Voxel-Shapes using the Junction Rule

Dennie Reniers¹ and Alexandru Telea²

¹Department of Mathematics and Computing Science, Eindhoven University of Technology, The Netherlands

²Institute of Mathematics and Computing Science, University of Groningen, The Netherlands

Abstract

We present a part-type segmentation method for articulated voxel-shapes based on curve skeletons. Shapes are considered to consist of several simpler, intersecting shapes. Our method is based on the junction rule: the observation that two intersecting shapes generate an additional junction in their joined curve-skeleton near the place of intersection. For each curve-skeleton point, we construct a piecewise-geodesic loop on the shape surface. Starting from the junctions, we search along the curve skeleton for points whose associated loops make for suitable part cuts. The segmentations are robust to noise and discretization artifacts, because the curve skeletonization incorporates a single user-parameter to filter spurious curve-skeleton branches. Furthermore, segment borders are smooth and minimally twisting by construction. We demonstrate our method on several real-world examples and compare it to existing part-type segmentation methods.

Categories and Subject Descriptors (according to ACM CCS): I.3.5 [Computer Graphics]: Geometric algorithms, languages, and systems

1. Introduction

Part-type segmentation is the task of decomposing a 3D shape into its meaningful components, which can be informally defined as those that one would intuitively perceive as the distinct, logical parts of the shape. Such segmentations are most suitable for shapes consisting of clearly articulated parts, such as animals, humans, and other natural shapes. Segmentation applications include shape analysis, shape matching, medical imaging, collision detection, and other geometric processing methods employing divide-and-conquer strategies.

In what identifies the logical parts of a shape, several directions can be taken. One is to consider a shape as the union of several simpler shapes, or *parts*. To detect these parts, some methods use the *minima rule* [HS97], which says that when two separate shapes interpenetrate, they do so in a concave contour of discontinuity of their surface tangent-planes. Methods can put the minima rule to use by placing *part cuts* at these curvature minima [LLS*05] for example, or by clustering surface elements based on curvature [KT03]. Another category of methods uses the curve skeleton to identify parts. The *curve skeleton* of a 3D shape is a 1D connected struc-

ture that is centered within the shape and efficiently captures the topology and articulation of the shape. It has been observed that the shape part-structure can be inferred from the curve-skeleton structure, that is, the curve skeleton allows for component-wise differentiation [CSYB05]. In this work, we refine this intuition: when two shapes interpenetrate, their respective curve skeletons join, resulting in a new junction in the curve skeleton located near the intersection. We call this the *junction rule*, in analogy with the minima rule.

In this paper, we present a new method for segmenting articulated voxel-shapes based on the junction rule. The particular curve-skeleton definition we use [DS06] associates with each curve-skeleton point a simple closed curve, i.e., a *Jordan curve*, on the shape surface. Each Jordan curve divides the surface into two parts and thus is a candidate part-cut. For each curve-skeleton junction, which signals the interpenetration of parts, we determine the number of part cuts that should be generated and their most suitable locations. We determine the suitability of a part cut using a new measure, called the “geodesicness” of the part cut. This ensures that the part-cuts are smooth, which is a desirable property. The saliency of the parts is guaranteed by

the curve-skeletonization stage, in which the curve skeleton is simplified using an intuitive scale-parameter, so that only significant curve-skeleton parts are retained. This makes our method robust to boundary noise. Our method is implemented for binary voxel-shapes, in which the shape is sampled on a regular grid, but in principle the same ideas are applicable to other shape representations (e.g. polyhedral).

The outline of this paper is as follows. In the next section we discuss related work. In Section 3 we briefly discuss previous work that our method uses, namely, the curve-skeleton definition [DS06] and its computation for voxel-shapes [RT07]. In Section 4, we introduce the *geodesicness measure* and based hereon, we place part-cuts that induce a segmentation of the shape into its logical parts. Section 5 presents results and a discussion. We evaluate the method according to criteria identified in literature and we compare our method to the relevant state-of-the-art methods. Section 6 concludes the paper.

2. Related Work

Segmentation methods can be classified by the type of output they produce, or by the type of information that they use. In the first classification, one can distinguish between patch- and part-type methods [Sha04]. Patch-type methods produce segments that satisfy certain well-defined geometric properties, such as flatness, whereas part-type methods try to produce meaningful components, often using cognition principles. Our method belongs to the part-type methods, to which we restrict our discussion.

Among the methods using *boundary* information, Katz and Tal [KT03] produce a fuzzy clustering of the surface-elements based on on geodesic and angular distances. After clustering, the fuzzy part-cuts are refined. In [KLT05], feature points are first extracted to improve pose-invariance, and each segment is made to represent at least one feature point. Lee et. al. [LLS*05] employ the minima rule to find partial feature-contours, which are then closed to form loops. The most salient loops are chosen, and a snake-based approach moves the loops to more suitable locations, forming part cuts. It is interesting to note that when closing the partial feature-contours, which were found using the minima rule, the authors use a “centricity” measure to ensure that the loops are perpendicular to the curve skeleton, something that users are found to be interested in.

Methods can also consider the shape’s *interior*. Mortara et al. [MPS*04] partition a shape into generic body and tubular parts, by sweeping spheres inside its volume. Attene et al. [AFS06] propose a segmentation based on fitting primitives, such as cylinders and spheres, inside the volume. Another way to consider the shape’s interior is by means of a skeletal structure, such as the curve skeleton. Our approach belongs to this category. To obtain meaningful part-cuts, this category of methods needs to signal events

Figure 1: The curve skeleton $\mathcal{C}$ with rainbow color map encoding $\rho_{\mathcal{C}}$, junctions $j, k, m \in \mathcal{J}_{\tau}$, a regular point $p \in \mathcal{C}$, feature sets $F(j)$ and $F(p)$, and shortest-path sets $\Gamma(j)$ and $\Gamma(p)$.

on the curve skeleton and map these events to the boundary in order to produce part cuts. Li et al. [LWTH01] use a planar space-sweep along the skeleton and signal both geometric and topological events of the cross sections. Cornea et al. [CSYB05] uses force-following of boundary particles inside the volume to compute the curve skeleton, and project skeleton segments back to the boundary. Reniers and Telea [RT07] detect curve-skeleton junctions, which naturally divide the boundary into connected components using Jordan curves that are associated with each curve-skeleton point. Lien et al. [LKA05] observe that segmentation and skeletonization share similar properties and propose a method that does both simultaneously. Another skeletal structure is the Reeb graph [HSKK01], produced from the iso-contours of a suitable mapping function on the surface. In [TVD07], the mapping function is defined as the curvature-constrained geodesic-distance to the closest feature points, detected as the extrema of another geodesic-based function. A subset of the contours is selected based on their topological and geometric evolution, and are used as part cuts. A comparative study of five recent part-type methods can be found in [AKM*06].

3. Preliminaries

We now briefly discuss the curve-skeleton computation, which is necessary for a good understanding of the rest of this paper.

The *surface skeleton* $\mathcal{S}$ of a 3D shape Ω with boundary $\partial\Omega$ is a 2D structure consisting of curved sheets, and is defined as those points in Ω having at least two boundary points at minimum distance [PSS*03]:

$$\mathcal{S}(\Omega) = \{p \in \Omega \mid |F(p)| \geq 2\}, \quad (1)$$

where F is the *feature transform*, which assigns to a point p the set of boundary points $F(p)$ at minimum distance of p : the *feature points* of p . Each sheet-point $p \in \mathcal{S}$ has exactly

two feature points. In limit cases, i.e., at sheet boundaries or sheet intersection-curves, $F(p)$ has a cardinality of three or more. In the following, we assume the general case of a surface-skeleton point having exactly two feature points, for notation simplicity.

Let $\rho_S : \mathcal{S} \rightarrow \mathbb{R}_+$ be the length of the shortest geodesic on the boundary between the two feature points $F(p)$. Dey and Sun [DS06] observe that the function ρ_S is low near the periphery of the surface skeleton and increases toward the center and define the curve skeleton $\mathcal{C}$ as the local maximum ridge of ρ_S . They use an erosion step to erode ρ_S until a 1D connected structure remains. In [RVT08] a similar definition of $\mathcal{C}$ is presented that is more suitable for voxel shapes, and has the additional advantage that it does not require an erosion step. In this work we deal with voxel shapes, so we compute curve skeletons using the latter approach.

Let $\bar{F}(p)$ be the *extended* feature transform [RVT08], defined as $\bar{F} = \bigcup_{x,y,z \in \{0,1\}} F(p_x + x, p_y + y, p_z + z)$. Let $\Gamma(p)$ be the set of shortest paths on the shape boundary between each two feature voxels in $\bar{F}(p)$:

$$\Gamma(p) = \bigcup_{a,b \in \bar{F}(p)} \gamma(a,b), \quad (2)$$

where $\gamma(a,b)$ is the shortest path between voxels a,b , as discrete-space equivalent of the shortest geodesic. Shortest paths are computed as 26-connected weighted voxel-chains using Dijkstra's algorithm on the boundary graph. The definition of the curve skeleton $\mathcal{C}$ is then defined as:

$$p \in \mathcal{C} \Leftrightarrow \Gamma(p) \text{ contains a Jordan curve}, \quad (3)$$

where a *Jordan curve* is a simple closed curve on the shape boundary $\partial\Omega$. To detect whether Γ contains a Jordan curve we slightly dilate Γ on $\partial\Omega$ to obtain Γ' . If and only if Γ' has at least two boundaries, $\Gamma(p)$ contains a Jordan curve. For junctions we count $n \geq 3$ boundaries, where n indicates the number $\mathcal{C}$ -branches that meet in the junction. For more details on this curve-skeletonization algorithm see [RT07].

With each curve-skeleton voxel p inside the volume a set of components on the boundary is associated, as follows. The Jordan curve theorem states that a Jordan curve divides a genus 0 surface into exactly two connected components. Let $C : \mathcal{C} \rightarrow \mathcal{P}(\partial\Omega)$ be the component set associated with each point $p \in \mathcal{C}$, obtained by taking the connected boundary-components induced by $\Gamma(p)$. The component sets C can be used to define another importance measure $\rho_C : \mathcal{C} \rightarrow \mathbb{R}_+$ on the curve skeleton. Let the components in C be ordered by their areas, or voxel counts: $\forall_{1 \leq i < k} |C_i| \leq |C_{i+1}|$. The measure ρ_C for a point p is defined as the total area of the components excluding the largest-area component:

$$\rho_C(p) = |\partial\Omega \setminus C_k(p)| \quad (4)$$

In this manner, a *simplified curve-skeleton* $\mathcal{C}_\tau$ can be computed by discarding all points whose ρ_C is smaller than a user-parameter τ . Note that ρ_C is different from ρ_S , but similar in spirit. In [RVT08] the two measures ρ_S and ρ_C are

Figure 2: Schematic overview of the four junction types (a-d) and their desirable part-cuts (magenta line-segments). An n -junction ($n = \{0, 1, 2, 3\}$) generates $n + 1$ parts using n part cuts. Curve skeletons are blue.

combined into a single monotonic measure ρ by taking their maximum, whose thresholding produces simplified surface-skeletons that are always connected.

Curve-skeleton junctions are needed in the following stages of our pipeline, and are detected robustly as follows. Each boundary b of $\Gamma'(p)$ bounds a component $c \in C(p)$. Note that some b 's may bound the same component. We discard boundaries that bound components smaller than τ . If the number of remaining boundaries is at least three, voxel p is detected as a robust junction. The robust set of junctions is denoted J_τ . Figure 1 depicts F , $\mathcal{C}$, ρ_C , Γ , and J_τ for an example shape.

4. New Method

Section 3 explained how the simplified curve-skeleton $\mathcal{C}_\tau$ and its junctions J_τ can be robustly detected, and that with each point $p \in \mathcal{C}$ a shortest-path set $\Gamma(p)$ on the shape boundary is associated. Using the *junction rule*, namely that curve-skeleton junctions indicate the interpenetration of parts, it is straightforward to use the junctions directly as *critical points*, that is, to use their path sets as part cuts that divide the shape surface into disjoint segments.

In fact, this approach was taken in [RT07]. Although the resulting segmentations are satisfying for a large amount of shapes, it delivers non-intuitive results for some other ones (see e.g. Fig. 3b later in this section). We next explain why the junctions are suitable for indicating the interpenetration of parts, but why they should not be used as cut-generating points. For the sake of discussion, we assume the generic case of junctions having exactly three emanating branches. In non-generic cases, junctions may have more branches. Our algorithm, presented in Section 4.2, can also handle four or more branches.

First, by using junctions as critical points, three part cuts per junction are generated, because three branches come together in a junction. However, on actual shapes we may perceive junctions either indicate the intersection of 4, 3, or 2 parts, or not indicate a part cut at all. Let an n -junction be a junction for which n part cuts should be generated, that is, it represents the intersection of $n + 1$ parts. Figure 2

Figure 4: The geodesicness measure computed for a curve-skeleton point p with feature points a, b and $\Gamma(p) = \{\gamma^1(a, b), \gamma^2(a, b)\}$. High geodesicness ($\sigma \approx 1$) (a). Medium geodesicness ($\sigma \approx 0.5$) (b).

schematically shows how one would intuitively place part cuts for four different shapes. Each of the four shapes gives rise to a different number of part cuts. We have observed that 1-junctions (Fig. 2c) occur most frequently in real-world shapes, whereas 3-junctions (Fig. 2a) are rare. Fig. 2d shows a 0-junction: it does not represent any part-cut at all. The junction results from the fact that the curve skeleton extends into the two corners of the rectangle. Although the shape could be interpreted as consisting of three interpenetrating parts, namely as the two corners and the rectangle's rump, this is non-intuitive and no part cuts should be generated. Considering that a continuous deformation exists between each two of these shapes, it is clear that junction type is inherently subjective. For example, if the disc of Fig. 2a becomes smaller and smaller, the situation transforms into that of Fig. 2b. Indeed, part-type shape segmentation can be considered a subjective task [AKM*06]. To summarize, to segment using the junction rule, we should be able to differentiate between the different junction types.

The second issue with using junctions as critical points is that the path sets of the junctions are ill-suited as part cuts as they do not tightly split off the parts. Figure 3b illustrates this using a chamfered box with three attached tubes. The path set of junction j (Fig. 3a) does not split off the two tubes accurately: each tube contains a large part of the box. The desirable segmentation is shown in Fig. 3e.

We address the two issues raised above, as follows. For each junction j , we search along each branch emanating from j for the first suitable critical point, where suitability is determined by our new geodesicness measure, introduced in Section 4.1. We find at most three candidate critical-points: the number of branches at j . Next, we select among the candidate critical-points the final critical-points that split off valid parts. This is detailed in Section 4.2.

4.1. Geodesicness measure

Recall from Section 3 that the Jordan curve $\Gamma(p)$ that is associated with each regular (non-junction) curve-skeleton point p consists of two shortest geodesics $\Gamma(p) = \{\gamma^1(a, b), \gamma^2(a, b)\}$ between p 's feature points a, b . Because

the piecewise geodesic Jordan curves are used as part cuts, the borders are smooth for the larger part, a desirable property for any segmentation method [AKM*06]. However, this property does not need to hold at the two feature points $F(p)$ where the two geodesics come together under an angle. We next introduce the *geodesicness measure* $\sigma : \mathcal{C} \rightarrow [0..1]$ on the curve skeleton which measures the degree to which the total Jordan curve $\Gamma(p)$ is geodesic.

A geodesic is defined as a curve on a surface whose geodesic curvature is everywhere vanishing. Thus, our measure σ should be 1.0 iff the two geodesics meet at an angle of 180 degrees in both feature points. In that the case the whole Jordan curve is a geodesic. Using differential geometry, we can project $\Gamma(p)$ in a small neighborhood of a onto the tangent plane at a , to determine the angle that the two geodesics make. However, as we are working in voxel-space and shapes can be noisy, a relatively large neighborhood must be taken, but then the projection is not accurate.

We take a more robust approach instead. Let c be the midpoint of one geodesic $\gamma^1(a, b)$, and d of the other $\gamma^2(a, b)$ (see Figure. 4a). We compute the shortest geodesic $\gamma(c, d)$. If $\gamma(c, d)$ is a subset of Γ , it means that Γ is a shortest geodesic between c, d , because a shortest geodesic between any two points is unique. Furthermore, it can be seen that the more $\gamma(c, d)$ deviates from Γ , the more Γ deviates from a geodesic at the feature points. This deviation can be measured by the surface area spanned by the geodesic triangle acd , or, in terms of difference in length $\|\cdot\|$, which we choose because it is more easily computed:

$$\sigma(p \in \mathcal{C}) = \frac{\|\gamma(c, d)\|}{\|\gamma(c, a)\| + \|\gamma(a, d)\|}. \quad (5)$$

This is a dimensionless measure attaining values between 0 and 1, the latter indicating a geodesic Jordan curve. The case of junctions, in which Γ consists of multiple Jordan curves, is handled by computing the minimum σ among Jordan curves.

4.2. Placing Part Cuts

We use the junctions J_τ as computed in Sec. 3 to search for critical points whose Jordan curves make for accurate part cuts, based on the geodesicness measure σ as defined in Sec. 4.1. Higher values of σ indicate a smooth Jordan curve and thus a good candidate for critical point placement. Figure 3c depicts σ . Looking at the branches emanating from the junctions into the tubes, we see that σ is low at the junction itself and reaches a maximum inside the tubes. We thus proceed as follows. Starting from a junction $j \in J_\tau$, we search along the $\mathcal{C}$ -branches emanating from j for the first point q which has a sufficiently high value of σ . The search for q on a branch stops when another junction or the end of the branch is reached. In practice, σ does not reach the maximum of 1.0 inside each protrusion that can be considered a part, so we tolerate lower values. Extensive experimental study suggests that placing critical points at $\sigma \geq 0.8$ gives

Figure 3: The curve skeleton of an artificial shape consisting of a box and three attached tubes. The curve skeleton has a 2-junction j , a 1-junction k , and a 0-junction m (a). Undesirable segmentation using junctions (b). Candidate critical-points based on σ (encoded by color map) (c). Final critical-points after filtering (d). Desirable segmentation using the final critical-points (e).

Figure 5: Four junction types (a-d). The Jordan curves $\Gamma(q_i)$ for the candidate critical-points q_i are shown as magenta line-segments. Retained part-cuts shown as thick magenta lines, non-retained ones are stippled. Shortest-path sets of the junctions are shown using stippled line-segments $\Gamma(j)$.

good results. For some branches, the very first (regular) point of the branch might have high enough σ to be marked a critical point. For other branches, we might find no such point: no critical point makes for a suitable part cut.

To summarize, we obtain for each junction j a set of candidate critical-points Q_j , whose cardinality is at most the number of branches at j . For branches whose measure σ stays low until the end we find no critical points, so that $|Q_j|$ is lower than the number of branches. Figure 5 schematically shows the detected candidate critical-points q_i that we would detect in this manner for the four different junction types (cf. Fig. 2). We observe that according to intuition, not all critical points that were found make for valid part-cuts, and some points should be removed from Q_j . In Fig. 5a, all critical points are valid: no points should be removed from Q_j . In Fig. 5b, point q_2 should be removed, in Fig. 5c, points q_2 and q_3 should be removed. In Fig. 5d the only candidate critical-point q_1 that was found should be removed.

In characterizing which candidate critical-points do not

make for valid part-cuts, we make the assumption that each intersection involves one *body* part, which is the “largest” intersecting part, and several *limb* parts, which are the smaller parts. In Fig. 5a for example, there is one disc-shaped body and three limbs. In Fig. 5c, there is one body and one limb, for which q_1 is the part cut. We thus want to distinguish the body part from the limb parts by their sizes. We can express the size of a part in at least two ways: by means of the part’s surface-area, or by means of its circumference at the point of intersection. The problem with the former is that it is a non-local property. For example, the limb due to q_2 in Fig. 5a can have smaller or larger surface-area than the disc-like part depending on what the hidden part of the shape looks like. A more local property is the circumference of the limb at the intersection, which can be measured by ρ_S , which has already been computed in the curve-skeleton computation (see Sec. 3). We further note that the junction is always located inside the body, and never in one of the limbs. Hence, the idea is to compare the local circumference of the limb, given by $\rho_S(q_i)$ of candidate critical-point q_i , with the circumference of the body, given by $\rho_S(j)$ of the junction j . If the circumferences are similar enough, that is, $1.0 - \frac{|\rho_S(j) - \rho_S(q_i)|}{\rho_S(j)} \geq 0.6$, where 0.6 is an empirically defined threshold, the candidate critical-point should be removed from Q_j as it is considered to lie on the body. However, we should never remove more than two critical points, as the body should not contain a bifurcation. Indeed, using this criterion we would remove all three critical points in Fig. 5b, because all three intersecting parts have a similar circumference. In case we cannot distinguish between body and limbs using the local circumference, we apply the global criterion and remove the critical point representing the largest surface area ρ_C , which is q_2 in case of Fig. 5b.

We restricted our discussion to junctions having three branches, our approach also works for higher-order junctions: the assumption of multiple limbs intersecting with a single body is still valid. Fig. 6 shows for two shapes the curve skeletons $\mathcal{C}$, junctions J_t , the geodesicness measure σ encoded as color map on $\mathcal{C}$, the candidate critical-points Q_j , and the part cuts.

Figure 7: Segmentations of several shapes obtained using our method.

Figure 6: Part cuts (red loops) for the Hand (a) and Noisy Dino shape (b). Curve skeletons $\mathcal{C}$ with rainbow color-map encoding σ (red=1, blue=0), junctions J_τ (black balls), candidate critical-points (colored balls).

5. Results and Discussion

Figure 7 shows several shapes segmented using our approach, with a minimum coloring scheme applied (smooth iso-surfaces have been computed from the voxel output for clearer display). Most segmentations can be colored using only two colors, which reflects that most junctions are 1-junctions. As can be seen, our method handles organic, geometric, and noisy shapes, and shapes with tunnels. We next evaluate our method using the criteria proposed in [AKM*06].

Correctness of part-type segmentation is subjective and application specific. We thus have to assess the correctness of our method (and thereby the junction rule) visually. As can be seen in Figs. 7,8, the segmentations agree with intuition and do not suffer from over-segmentation. Fine structures such as fingers and toes are captured well, provided the grid resolution is high enough and the curve skeleton has not been simplified too much by the user.

The method is *multiscale* through the use of the user-

parameter τ [RVT08], which controls the simplification of the curve skeleton. The user-parameter is intuitive as it thresholds the importance measure ρ_C which has a geometric meaning: it assigns to a point the minimum surface-area induced by the Jordan curve at the point. By further increasing τ high-level abstractions of the shape are obtained. Figs. 6b,7 show that the noisy Dino is correctly segmented. If desired, a hierarchical segmentation could be produced by implementing the scheme proposed in [RT07].

The *boundaries* between segments are found directly as the Jordan curves associated with the critical points. The part cuts are closed by default (unlike e.g. [LLS*05]), and no refinement of part cuts is necessary (unlike e.g. [KT03]), keeping the algorithm straightforward to implement and intuitive to use. Furthermore, the boundaries are smooth by construction, as the Jordan curves are piecewise geodesic and critical points are only placed at points with a high total geodesicness. The saliency of the part cuts is guaranteed because the curve skeleton contains only significant branches, guaranteed by the scale parameter τ .

The segmentations are *pose-invariant* (see Fig. 8). That is, the number and location of segments do not change under different poses of the same shape. The reason is that the structure of the underlying curve skeleton does not change significantly under deformations of the shape. Even if a junction is introduced by a deformation, it will not immediately result in extra part cuts because the geodesicness measure σ on this new $\mathcal{C}$ -branch has a low value. Only if the new part really protrudes, σ reaches a high enough value. Methods using surface curvature might have problems here, as the curvatures change considerably under deformations.

The *asymptotic complexity* of our method is dominated by the curve-skeleton computation, which involves computing the shortest-path (using Dijkstra's algorithm) between each two feature voxels in $F(p)$ (Sec. 3) for each object

Figure 8: Pose-invariance of the segmentations.

voxel p , which amounts to $O(m(n \log n))$ where $m = |\Omega|$ and $n = |\partial\Omega|$. Fortunately, on articulated shapes for which our method is intended, to compute the shortest path between two feature voxels only a relatively small part of the boundary needs to be visited. Our method, implemented in C++ takes up to 5 minutes on a Pentium IV 3 GHz for curve-skeletonization and segmentation, for the shapes shown in this paper, with resolutions up to 400^3 voxels.

Only a single *control parameter* is needed: the curve-skeleton simplification threshold τ . The critical-point placement has two internal parameters: one for the minimum geodesicness of a critical point and the other for distinguishing between body and limbs. We have experimentally determined appropriate settings for these parameters and kept these fixed for the shapes shown in this paper.

Finally, we discuss the limitations of our approach. By using the curve skeleton and associated Jordan curves, each presenting a potential part-cut, we have essentially reduced the problem of segmentation from 2D to 1D. However, this also causes some limitations. In case a part is strongly tapered, we may obtain a low σ all along the respective C -branch, and we might not find a critical point. In practice, if a part is indeed so strongly tapered that σ does not come above our fixed threshold for the *whole* C -branch, then it must be a small part, and failing to segment it is not a serious problem. A second limitation is that we have to choose part cuts from a limited set such that they are smooth, but beyond that we cannot enforce complete smoothness. In Fig. 8 for example, the Horse's neck is tapered, and so its part-cut is not so smooth as it could be when placed more freely (using e.g. a snake-based approach such as in [LLS*05]). Another limitation occurs when the curve skeletons of two intersecting shapes do not generate a junction, which is the case when the two curve skeletons are perfectly aligned. However, the minima rule suffers from similar problems when two shape intersect non-transversally. Although our method works on voxel-shapes, we argue that it can be adapted to polyhedral shapes as well. Computation of the curve skeleton and asso-

ciated Jordan curves is available for meshes [DS06]. However, because the implementation is non-trivial and not available to us, we have chosen for a voxel-based implementation. The other two steps, computing the geodesicness measure and placing the part cuts, are also not inherently voxel-based.

5.1. Comparison

We now briefly compare the results of our method to seven other state-of-the-art methods by visual inspection of the results. Figure 9 shows the Homer shape segmented using four existing methods [KT03, KLT05, MPS*04, AFS06] (images from [AKM*06]) and our own method in the rightmost column. We notice that our method captures fine features as the fingers and the nose, whereas the other methods do not. Our method segments only the largest toe, as the others are not separated from each other by the outside volume. We further observe that our method, like [MPS*04], does not label the head as a separate segment. The reason for this is that the neck does not have a (significantly) smaller radius than the body. Hence, our method considers the body and head belonging to the same part. Finally, our segment borders for the legs are tight and smooth, whereas the other methods can be seen to have some problems for these areas. We compare segmentations for the Hand shape with another three state-of-the-art methods [LLS*05, LWTH01, RT07] (images from the respective papers), of which the last two are similar to our method in that they also use the curve skeleton. We observe that our method tightly splits off the fingers and thumb, whereas [LWTH01] and [RT07] do not. In [RT07] each finger contains a large part of the palm of the hand, caused by the fact that this method uses the junctions directly as critical points (Sec. 4). Finally, all other three methods have a segment border on the palm, whereas our method leaves the palm intact.

6. Conclusion

We have presented a new method for creating part-type segmentations of articulated shapes using the junction rule. A user-parameter is incorporated to filter the curve skeleton of spurious branches and ensure the saliency of parts. Each curve-skeleton point has an associated Jordan curve that presents a candidate part-cut. Suitable part-cuts are selected by measuring their smoothness, computed based on a geodesic criterion, and considering the body/limb relationships of the intersecting parts. Our method is straightforward to implement, yields smooth segment boundaries by construction, and is robust to boundary noise.

Acknowledgements

This work was supported by the Netherlands Organization for Scientific Research (NWO) under Grant No. 612.065.414. We used shapes from the AIM@SHAPE shape

Figure 9: Visual comparison of segmentations for the Homer and Hand shapes.

repository, McGill 3D Shape Benchmark, and the Princeton Shape Benchmark. The Horse poses are from [SP04].

References

- [AFS06] ATTENE M., FALCIDIENO B., SPAGNUOLO M.: Hierarchical mesh segmentation based on fitting primitives. *The Visual Computer* 22, 3 (2006), 181–193.
- [AKM*06] ATTENE M., KATZ S., MORTARA M., PATANE G., SPAGNUOLO M., TAL A.: Mesh segmentation - a comparative study. In *IEEE Int. Conference on Shape Modeling and Applications (SMI)* (2006), p. 7.
- [CSYB05] CORNEA N. D., SILVER D., YUAN X., BALASUBRAMANIAN R.: Computing hierarchical curve-skeletons of 3d objects. *The Visual Computer* 21, 11 (2005), 945–955.
- [DS06] DEY T. K., SUN J.: Defining and computing curve-skeletons with medial geodesic function. In *Proc. of Eurographics Symposium on Geometry Processing* (2006), pp. 143–152.
- [HS97] HOFFMAN D. D., SINGH M.: Saliency of visual parts. *Cognition* 63, 1 (1997), 29–78.
- [HSKK01] HILAGA M., SHINAGAWA Y., KOHMURA T., KUNII T. L.: Topology matching for fully automatic similarity estimation of 3d shapes. In *Proc. of SIGGRAPH'01* (2001), ACM, pp. 203–212.
- [KLT05] KATZ S., LEIFMAN G., TAL A.: Mesh segmentation using feature point and core extraction. *The Visual Computer* 21 (2005), 649–658.
- [KT03] KATZ S., TAL A.: Hierarchical mesh decomposition using fuzzy clustering and cuts. *ACM Transactions on Graphics* 22, 3 (2003), 954–961.
- [LKA05] LIEN J.-M., KEYSER J., AMATO N. M.: Simultaneous shape decomposition and skeletonization. In *Proc. of ACM Solid and Physical Modeling Symposium (SPM)* (2005), pp. 219–228.
- [LLS*05] LEE Y., LEE S., SHAMIR A., COHEN-OR D., SEIDEL H. P.: Mesh scissoring with minima rule and part saliency. *Computer Aided Geometric Design* 22 (2005), 444–465.
- [LWTH01] LI X., WOON T. W., TAN T. S., HUANG Z.: Decomposing polygon meshes for interactive applications. In *Proc. of symposium on Interactive 3D graphics* (2001), pp. 35–42.
- [MPS*04] MORTARA M., PATANÈ G., SPAGNUOLO M., FALCIDIENO B., ROSSIGNAC J.: Plumber: a multiscale decomposition of 3d shapes into tubular primitives and bodies. In *Proc. of Solid Modeling and Applications* (2004), pp. 339–344.
- [PSS*03] PIZER S. M., SIDDIQI K., SZÉKELY G., DAMON J. N., ZUCKER S. W.: Multiscale medial loci and their properties. *Int. Journal of Computer Vision* 55, 2-3 (2003), 155–179.
- [RT07] RENIERS D., TELEA A.: Skeleton-based hierarchical shape segmentation. In *Proc. of the IEEE Int. Conf. on Shape Modeling and Applications (SMI)* (2007), pp. 179–188.
- [RVT08] RENIERS D., VAN WIJK J. J., TELEA A.: Computing multiscale curve and surface skeletons of genus 0 shapes using a global importance measure. *IEEE Transactions on Visualization and Computer Graphics* 14, 2 (2008), 355–368.
- [Sha04] SHAMIR A.: A formulation of boundary mesh segmentation. In *Proc. of the Second Int. Symp. on 3D Data Processing, Visualization and Transmission (3DPVT'04)* (2004), pp. 82–89.
- [SP04] SUMNER R. W., POPOVIC J.: Deformation transfer for triangle meshes. *ACM Transactions on Graphics* 23, 3 (2004), 399–405.
- [TVD07] TIERNY J., VANDEBORRE J.-P., DAOUDI M.: Topology driven 3d mesh hierarchical segmentation. In *Proc. of the IEEE Int. Conf. on Shape Modeling and Applications (SMI)* (2007), pp. 215–220.