

On-the-fly curve-skeleton computation for 3d shapes

ANDREI SHARF¹, THOMAS LEWINER^{1,2}, ARIEL SHAMIR³ AND LEIF KOBBELT⁴

¹ School of Computer Science — Tel Aviv University — Israel

² Department of Mathematics — Pontifícia Universidade Católica — Rio de Janeiro — Brazil

³ Efi Arazi School of Computer Science — The Interdisciplinary Center — Herzliya — Israel

⁴ Computer Graphics Group — RWTH — Aachen — Germany

www.cs.tau.ac.il/~asharf. www.mat.puc-rio.br/~tomlew. www.faculty.idc.ac.il/arik.

www-i8.informatik.rwth-aachen.de.

Abstract. The curve-skeleton of a 3D object is an abstract geometrical and topological representation of its 3D shape. It maps the spatial relation of geometrically meaningful parts to a graph structure. Each arc of this graph represents a part of the object with roughly constant diameter or thickness, and approximates its centerline. This makes the curve-skeleton suitable to describe and handle articulated objects such as characters for animation. We present an algorithm to extract such a skeleton on-the-fly, both from point clouds and polygonal meshes. The algorithm is based on a deformable model evolution that captures the object's volumetric shape. The deformable model involves multiple competing fronts which evolve inside the object in a coarse-to-fine manner. We first track these fronts' centers, and then merge and filter the resulting arcs to obtain a curve-skeleton of the object. The process inherits the robustness of the reconstruction technique, being able to cope with noisy input, intricate geometry and complex topology. It creates a natural segmentation of the object and computes a center curve for each segment while maintaining a full correspondence between the skeleton and the boundary of the object.

Keywords: *3D curve-skeleton. deformable models. skinning.*




Figure 1: On-the-fly curve-skeleton computation for a real scan of a woman model. Even in the presence large missing data (left), the deformable model accurately interprets the shape to generate a thin curve skeleton with a meaningful segmentation.

1 Introduction

Computer graphics applications handle huge models representing complex 3D objects. The most common rep-

resentations for such objects are boundary meshes or point-sets. However, applications such as editing, animation, morphing or shape matching often need a higher level understanding of the shape and its structure. Such an understanding can be conveyed through the use of an inner curve-skeleton for the object. This geometrical and topological

Preprint MAT. 03/07, communicated on February 10th, 2007 to the Department of Mathematics, Pontifícia Universidade Católica — Rio de Janeiro, Brazil.

abstraction represents the 3D shape on an object, by mapping the relation of geometrically meaningful parts to a graph structure with univariate arcs. Each arc represents one of these meaningful parts by its centerline, assuming it has roughly constant diameter around it. The curve-skeleton is both concise and expressive enough to represent an abstraction of the 3D object. Moreover, it is easier to handle than the medial axis since it avoids surface elements, while its topological structure still represents the shape efficiently.

The main difficulty in computing curve-skeletons for complex objects is to correctly interpret their shape. Objects may include complex local topology, large missing parts and noise, and this requires a robust and accurate interpretation mechanism. In this work we present an algorithm to extract the curve skeleton from 3D objects represented as boundary meshes as well as point-sets. The algorithm relies on a surface reconstruction technique that can handle both meshes and point-sets surfaces, while robustly interpreting their shape.

The key idea is to track the reconstruction of a given object using a deformable model similar to [Sharf *et al.* 2006]. It uses competing fronts that evolve inside the object until they achieve a sufficient approximation of the shape. Since the deformable model reconstructs the original object, we compute the center curve of this model to define the center skeleton of the object itself. To do so, during the model deformation we follow the fronts' centers, generating a first approximation for the curves of the skeleton. These are filtered on-the-fly according to their geometry, their branching structure and the front competition performance, obtaining the final curve-skeleton (see Figure 1).

Furthermore, using the reconstructed model we can maintain a full correspondence between the curve-skeleton and the object boundary. For curve-skeleton of point sets, this mapping is explicit on the reconstructed mesh, while for 3D meshes it is defined by a simple closest-point projection. Such curve-skeleton computation inherits the robustness of the deformable model technique, while performing on-the-fly during the reconstruction process.

2 Related Works

Techniques to compute curve-skeletons of 3D models can be classified into three categories, according to [Cornea *et al.* 2005]: voxel topology, computational geometry and continuous implicit. Some of the most relevant works of each category are summarized next. We refer the reader to [Cornea *et al.* 2005] for a more complete survey.

Voxel topology. The computation of the curve-skeleton can be derived from computer vision techniques, such as topological thinning [Borgefors *et al.* 1999, Gagvani and Silver 1999]. These methods iteratively remove *simple points* from the boundary of a voxel set. They differ mainly by the definition of these simple points.

Geometry. Curve-skeleton is a natural structure for 2D shapes, where it is usually computed as a subset of the medial axis [Shaked and Bruckstein 1998]. For discrete shapes, this medial axis can be extracted from the Voronoi diagram of some points on or near the shape [Dey and Zhao 2004, Attali *et al.* 2007]. However, for 3D shapes, the medial axis may contain surface elements. The curve skeleton can be extracted from these 2D elements using their own medial axis [Attali and Montanvert 1997, Amenta *et al.* 2001], a distance field [Wu *et al.* 2003] or a more refined geodesic field [Dey and Sun 2006]. This set of techniques allows a definition of a curve-skeleton, but requires some delicate sampling condition to reach the correct interpretation. Our technique correctly interprets the data in various situations including both noise and large missing pieces.

Implicit. Another set of techniques compute the curve-skeleton from the ridge points of a 3D field [Schirmacher *et al.* 1998, Palagyi and Kuba 1999, Bitter *et al.* 2001]. These techniques generally detect and track the ridge points using an implicit method such as fast marching [Zhou and Toga 1999, Hassouna and Farag 2005, Cornea *et al.* 2005*] or active contours [Golland and Grimson 2000]. This uniform integration and the complex geometry tracking used in these work typically assume that the shape thickness and complexity decreases during tracking, which may lead to shape misinterpretation in complex topology and noisy case. Moreover, this implicit process does not maintain the correspondence between the curve-skeleton and the shape. Our technique shares some similarity with these approaches. However, we do not extract the curve skeleton directly from a field, but rather interpret it in a coarse-to-fine manner using explicit evolving fronts.




Figure 2: The deformable model with competing fronts. The fronts move in a coarse to fine manner, and may split to form sub-fronts, inducing the branching structure of our skeleton.



Figure 3: Tracking the center of fronts during reconstruction defines the inner curve-skeleton of a multi-torus object. The topological events on the deformable model are mapped as loops in the curve-skeleton.

3 Curve-skeleton by Reconstruction

Complex local topology and large missing parts of an object require a robust and accurate interpretation of its shape. Following this observation, our technique for curve-skeleton extraction uses strategies similar to the reconstruction method of [Sharf *et al.* 2006]. In particular, we utilize a deformable model to interpret the shape. We first extract from the deformation of the model a set of center curves. We filter this set during the evolution to generate a clean skeleton, while maintaining a full mapping between the model and the skeleton. This whole process is performed on-the-fly using inexpensive operations (see Algorithm 1 for pseudo-code of the whole process). In particular, the front center's tracking relies on an inexpensive marking technique, avoiding delicate geometry calculus or discrete integration.

Algorithm 1 On-the-fly curve-skeleton.

```

1: while evolving fronts do
2:   counter ++
3:   for all evolving front f do
4:     create skeleton node  $s_f$  from f
5:      $s_f.\text{arcid} \leftarrow (\text{counter}, f)$ 
6:     for all vertex v in front f do
7:        $(\text{old}_c, \text{old}_f) \leftarrow v.\text{arcid}$ 
8:       NewArcs+ ←
           $\text{Arc}[(\text{old}_c, \text{old}_f), (\text{counter}, f)]$ 
9:        $v.\text{arcid} \leftarrow (\text{counter}, f)$ 
10:    end for
11:  end for
12:  filter(NewArcs)
13:  remap vertices to skeleton nodes
14: end while

```

The deformable model. Similarly to [Sharf *et al.* 2006], our deformable model includes multiple evolving fronts that reconstruct the fine features of the shape only after the coarse ones are done. This coarse-to-fine interpretation of

the shape is achieved through a competition mechanism between the fronts on an explicit, dynamic mesh representation. Each front is a connected set of mesh vertices, which moves in outward normal direction of the mesh (see Figure 2). This movement is proportional to the distance to the surface, and constrained by a Laplacian system that enforces the smoothness of the mesh. This system is weighted in order to control the tension of the mesh. A small weight provides more elasticity, allowing it to evolve fast, while a large tension slows its evolution. The competition mechanism sets the tension of a front according to its geometry and previous evolutions.

The vertices move while *active*, and are deactivated when they are close enough to the target shape. Deactivating vertices eventually disconnects a front. In such case, the front is split into connected sub-fronts, and each sub-front continues to evolve separately. Collisions between fronts are prevented during the evolution. At the end of the deformation, fronts collisions are interpreted as topological events and are applied as handle attachments by connecting the fronts that still collide. Even for complex topology (see Figure 3), this reconstruction strategy leads to a robust high level interpretation of the shape which we utilize for our skeleton extraction

Skeleton nodes computation. We use the same technique for both point-sets and meshes. The evolving fronts reconstruction algorithm is used to interpret the shape from coarse to fine, while we follow the centers of the fronts to define the centerline of the shape. Each front evolution is tracked separately. At each iteration we insert a new skeleton node for every front at the barycenter of its mesh vertices. For instance in Figure 5, the whole palm is conquered with a single front, whose center does not move away from the hand's center.

Model to skeleton correspondence. During the evolution, both the skeleton node (front barycenter) and the mesh vertices of the fronts are assigned an *arc identifier*, which is a pair composed of an iteration counter and the front




Figure 4: Filtering of the curve-skeleton of a dragon mesh (left) from an initial curve set (center left). Filtering with loose (center right) or tight (right) threshold removes the skeleton’s noise (in the dragon’s front leg) and fine details (spine’s bumps), while preserving the main shape elements.

identifier (see Algorithm 1: lines 5 and 9). This identifier links the vertex directly to the curve-skeleton nodes having the same arc identifier. When a vertex is deactivated during the model deformation, it keeps its last valid arc identifier. Since the model guarantees a close reconstruction of the original shape, we simply project the model onto the shape to obtain the shape-to-skeleton correspondence (Figure 5).

Curve-skeleton structure. The skeleton nodes are connected according to the fronts’ evolution that generates an initial curve set (see Figure 4). We compute these connections based on an inexpensive marking technique using arc identifiers in a union-find structure. At each iteration, the arc identifiers of some front vertices are updated. Skeleton nodes corresponding to the old and new arc identifiers are connected (see Algorithm 1: lines 7-9). Straight connections without branches belong to the same curve. Note that the front identifier does not need to be coherent from one iteration to the next one, since the vertices of the mesh both define the movement and the skeleton connections.




Figure 5: Color mapping of the evolution tension parameter (left). The initial skeleton structure (left) is filtered using the evolution tension parameter (center left), simplifying the skeleton (center right) while preserving the skeleton/model correspondence and segmentation (right).

Branching and topological events. During the evolution, fronts eventually split and merge. Whenever a front splits, we generate a branching node for the corresponding arc and continue to track the different sub-fronts separately. For example on the hand of Figure 5, the splitting of the initial front at entering the fingers generates a branch for each finger and separates a sub-front for the wrist. Fronts can merge either due to topological events or to the dynamic remeshing of the deformable model. When fronts merge due to a topological events, we connect the two arcs to close a loop in the curve graph which ensures that the skeleton has the same homotopy as the deformable model (see Figure 3). However, for local merging of close fronts we do not link the two curve branches together, preserving the correct topology of the skeleton. Instead, we merge the two branches into the branch of smallest curvature.

Filtering. The final curve-skeleton is a subset of the initial curve set. Nevertheless, very similar to the medial axis, it can become very noisy and contain spurious branching especially near the object’s boundary. Hence, we allow to simplify the initial curve graph connectivity by pruning and merging; Depending on a user specified parameter, we remove from the curve skeleton arcs whose length or corresponding front tension are smaller than the given threshold. Removing end arcs is a simple pruning operation, while removing intermediate arcs merge branchings. In both cases, we remap corresponding vertices to the new nodes using the union-find structure. Next, the tree (or graph with the topological event’s loop) geometry can be further smoothed to produce our curve-skeleton (see Figure 5) using a spline filter but keeping the branching nodes fixed. Note that both the filtered curve-skeleton and the model-skeleton mapping can be computed on-the-fly during the evolution or eventually at the end of the process, without complex geometrical calculii.

4 Results

We have tested our algorithms on various inputs ranging from manifold meshes (Figures 7 and 6) to noisy points-




Figure 6: Comparing our technique on a horse model. From left to right: The horse mesh. The curve-skeleton computed using the distance field of [Gagvani and Silver 1999] (16 seconds). Using the topological thinning method of [Palagyi and Kuba 1999] (92 seconds). Using the potential field method of [Cornea *et al.* 2005*] (16 minutes). Our result, computed in 3.1 seconds.

sets with holes (Figure 1) including objects with complex topology (Figures 3 and 9) and geometry (Figure 4). Table 1 gives some timing results for the computation. We also demonstrate how using our method, the skeleton can assist in interpreting the shape and structure of an object. Figures 5 and 8 show a mapping of the skeleton curves to the boundary mesh using the skeleton mapping to the boundary. They also show how the skeleton segments the shape into meaningful parts. This method can also support a hierarchical decomposition following major and minor skeleton arcs Figure 4. In the illustrations of this work, the curve-skeleton is drawn using cylinders for visualization purposes.

We compare our method with implicit techniques using the code provided in [Cornea *et al.* 2005]. In Figure 6, we use the same volumetric field for guiding our model and for the other technique. The distance field of [Gagvani and Silver 1999] performs in 16 seconds, the topological thinning method of [Palagyi and Kuba 1999] in 92 seconds and the potential field method of [Cornea *et al.* 2005*] in 16 minutes, while our method last 3.1 seconds.




Figure 7: The result of our algorithm is a thin, 1D curve-skeleton even in degenerated cases such as a 3D box. When the initial model is not exactly centered (left), the skeleton quickly recovers the symmetry of the input (center). The filtering preserves the symmetry (right).

Our method guarantees the following properties (stated in [Cornea *et al.* 2005] as curve-skeleton qualitative criteria); **Homotopic:** Since we close a loop in the skeleton for

each topological event of the deformable model, the curve-skeleton has the same homotopy as the deformable model (Figures 3 and 9). **Thin:** As opposed to many topological thinning and distance transform methods, our skeleton is computed directly as a set of 1D curves, leading to a thin skeleton (Figure 7). **Robust component-wise differentiation:** The deformable model performs a robust shape interpretation of the data (Figures 3 and 9), which generates a meaningful decomposition of both the surface and the curve-skeleton (Figures 5 and 8). **Efficient:** Our technique performs on-the-fly in only a few seconds and is among the fastest curve-skeleton extraction algorithms (see Table 1). **Hierarchic:** As many of the distance field techniques, we use a filter to remove spurious branches. This filter can be tuned to generate a hierarchical skeleton (Figure 4).

Nevertheless our method does not fully guarantee the following two criteria: direct visibility of the shape from the skeleton and centeredness. The centeredness property is observed but not guaranteed in our method since we merely use a tracking heuristic for the centers.

Conclusions

We propose an efficient and robust method to extract a curve-skeleton from either meshes or point-set objects. This work illustrates the similarity of shape interpretation for surface reconstruction and curve-skeleton computation. In both problems, correct interpretation and segmentation must be performed to achieve high quality results. Our method uses a deformable model borrowed from reconstruction techniques to robustly interpret the shape and topology of the target object. The curve skeleton is then computed on-the-fly during the model deformation. Its robustness and speed make it attractive for many applications such as animation and matching, and we plan to utilize it to such problems in the future.




Figure 8: The front tension parameter from the skeleton to the boundary can distinguish between coarse and fine parts of the object (color-mapped on left figures). The curve skeleton can also be used to interpret the shape: the different branches of the skeleton impose a meaningful segmentation (color-mapped on right figures).

References

- [Amenta *et al.* 2001] N. Amenta, S. Choi and R. Kolluri. The Power Crust, unions of balls, and the medial axis transform. *Computational Geometry: Theory and Applications*, 19(2–3):127–153, 2001.
- [Attali *et al.* 2007] D. Attali, J.-D. Boissonnat and H. Edelsbrunner. Stability and computation of medial axes. to appear.
- [Attali and Montanvert 1997] D. Attali and A. Montanvert. Computing and simplifying 2D and 3D skeletons. *Comp. Vision and Image Understanding*, 67(3):156–169, 1997.
- [Bitter *et al.* 2001] I. Bitter, A. Kaufman and M. Sato. Penalized-distance volumetric skeleton algorithm. *Transactions on Visualization and Computer Graphics*, 7(3):195–206, 2001.
- [Borgefors *et al.* 1999] G. Borgefors, I. Nystrom and G. DiBaja. Computing skeletons in three dimensions. *Pattern Recognition*, 32(7):1225–1236, 1999.
- [Cornea *et al.* 2005] N. Cornea, D. Silver, X. Yuan and R. Balasubramanian. Curve-skeleton applications. In *Visualization*, pages 95–102. IEEE, 2005.
- [Cornea *et al.* 2005*] N. Cornea, D. Silver, X. Yuan and R. Balasubramanian. Computing hierarchical curve-skeletons of 3d objects. *The Visual Computer*, 21(11):945–955, 2005.
- [Dey and Zhao 2004] T. K. Dey and W. Zhao. Approximate medial axis as a Voronoï subcomplex. *Computer-Aided Design*, 36(2):195–202, 2004. Solid Modeling and Applications '04.
- [Dey and Sun 2006] T. K. Dey and J. Sun. Defining and computing curve-skeletons with medial geodesic func-

tion. In *Symposium on Geometry Processing*, pages 143–152. ACM/Eurographics, 2006.

[Gagvani and Silver 1999] N. Gagvani and D. Silver. Parameter-controlled volume thinning. *Graph. Models and Image Proc.*, 61(3):149–164, 1999.

[Golland and Grimson 2000] P. Golland and W. Grimson. Fixed topology skeleton. In *Comp. Vision and Pattern Recog.*, pages 1010–1017. IEEE, 2000.

[Hassouna and Farag 2005] M. Hassouna and A. A. Farag. Robust centerline extraction framework using level sets. In *Comp. Vision and Pattern Recog.*, pages 458–465. IEEE, 2005.

[Palagyi and Kuba 1999] K. Palagyi and A. Kuba. A parallel 3D 12-subiteration thinning algorithm. *Graph. Models and Image Proc.*, 61(4):199–221, 1999.

[Schirmacher *et al.* 1998] H. Schirmacher, M. Zockler, D. Stalling and H.-C. Hege. Boundary surface shrinking: a continuous approach to 3D center line extraction. In *Image and Multidimensional Digital Signal Processing*, pages 25–28. IEEE, 1998.

[Shaked and Bruckstein 1998] D. Shaked and A. Bruckstein. Pruning medial axis. *Comp. Vision and Image Understanding*, 69(2):156–169, 1998.

[Sharf *et al.* 2006] A. Sharf, T. Lewiner, A. Shamir, L. Kobbelt and D. Cohen-Or. Competing fronts for coarse-to-fine surface reconstruction. *Computer Graphics Forum*, 25(3):389–398, 2006.

[Wu *et al.* 2003] F.-C. Wu, W.-C. Ma, P.-C. Liou, R.-H. Laing and M. Ouhyoung. Skeleton extraction of 3D objects with visible repulsive force. In *Pacific Graphics*, pages 409–413. IEEE, 2003.

[Zhou and Toga 1999] Y. Zhou and A. Toga. Efficient skeletonization of volumetric objects. *Transactions on Visualization and Computer Graphics*, 5(3):196–209, 1999.

	Deformation	Skeleton	Filtering
foot	0.7 secs.	0.4 secs.	0.1 secs.
woman	4.9 secs.	3.5 secs.	0.4 secs.
horse	1.7 secs.	1.2 secs.	0.2 secs.
multi-torus	1.0 secs.	0.6 secs.	0.3 secs.
hand	0.5 secs.	0.1 secs.	0.2 secs.
CAD	1.9 secs.	1.2 secs.	0.1 secs.
dragon	8.6 secs.	6.1 secs.	2.3 secs.

Table 1: Computation time on a Pentium IV 1GHz / 1Gb.




Figure 9: Skeleton extraction for a complex CAD mesh.