Pattern Recognition Letters xxx (2014) xxx–xxx

Contents lists available at ScienceDirect

Pattern Recognition Letters

journal homepage: www.elsevier.com/locate/patrec

Comparison of curve and surface skeletonization methods for voxel shapes ☆

André Sobiecki [a,*], Andrei Jalba [b], Alexandru Telea [a]

[a] *Institute Johann Bernoulli, Univ. of Groningen, Nijenborgh 9, 9747 AG Groningen, The Netherlands*
[b] *Department of Mathematics and Computer Science, TU Eindhoven, Den Dolech 2, 5600 MB, Eindhoven, The Netherlands*

ARTICLE INFO

Article history:
Available online xxxx

Keywords:
Surface and curve skeletons
Voxel shapes
Medial axes

ABSTRACT

Surface and curve skeletons are important shape descriptors with applications in shape matching, simplification, retrieval, and animation. In recent years, many surface and curve skeletonization methods for 3D shapes have been proposed. However, practical comparisons of such methods against each other and against given quality criteria are quite limited in the literature. In this paper, we compare 4 surface and 6 recent curve skeletonization methods that operate on voxel shapes. We first compare the selected methods from a global perspective, using the quality criteria established by a reference paper in the field. Next, we propose a detailed comparison that refines the gained insights by highlighting small-scale differences between skeletons obtained by different methods.

1. Introduction

Skeletons are shape descriptors with many applications in shape processing, registration, retrieval, matching, animation, and compression [1]. 3D shapes admit two types of skeletons: *Surface* skeletons are 2D manifolds formed by the loci of maximally-inscribed balls within a shape [1,2]. *Curve* skeletons are 1D curves which are locally centered in the shape and capture the shape's part-whole structure [3].

Since the early skeleton definition by Blum [4], many methods have been proposed to compute the two skeleton types. Such methods differ in theoretical aspects, *e.g.* the exact definition for curve skeletons, and practical aspects, *e.g.* space discretization (voxels *vs* meshes); the various approximations being used; and the actual skeleton extraction algorithm. These aspects, and the inherent sensitivity of skeletons to boundary noise, makes different methods produce widely different skeletons for the same input. This causes challenges for the users of skeletons in both research and practical contexts.

Recognizing these challenges, Cornea et al. [3] have presented a taxonomy of curve skeletonization methods and the way these satisfy a set of desirable skeletal properties, and illustrated these for four such methods. Since this publication, several new skeletonization methods have been proposed. A recent study extended the work in [3] by comparing methods of a particular class (contraction-based methods for meshed shapes) against Cornea's criteria [5]. These two studies however cover only a very limited fraction of the current skeletonization methods. Also, papers introducing new skeletonization methods typically present only few additional comparisons. Table 1 illustrates this for a selection of methods, which is by no means exhaustive. Finally, very few comparisons of surface skeletons with curve skeletons are presented, so there are still many open questions on the relationships of the two skeleton types. As computational advances allow implementing increasingly complex skeletonization methods, the challenge of understanding the relative pro's and con's of such new methods only grows.

In this paper, we address the above challenges by presenting a comparison of 4 surface and 6 curve skeletonization methods. In contrast to [5], we focus here on voxel-based methods. In addition to [3], we cover here methods having emerged after their study was published. We use in our comparison the same desirable criteria as in [3]. In addition, we also propose a detailed comparison that aims to provide a fine-grained detail view on the subtle differences between skeletons computed by different methods, including comparisons of curve with surface skeletons. Our results offer additional insight in limitations and challenges of current methods which, to our knowledge, have not been highlighted so far. These results represent further support for the quest of designing better skeletonization methods.

☆ This paper has been recommended for acceptance by Cris L. Luengo Hendriks.
* Corresponding author. Tel.: +31 503633711.
E-mail addresses: a.sobiecki@rug.nl (A. Sobiecki), a.c.jalba@tue.nl (A. Jalba), a.c.telea@rug.nl (A. Telea).

http://dx.doi.org/10.1016/j.patrec.2014.01.012
0167-8655/

Table 1
Recent 3D skeletonization papers. For each method, we show its type (**V**olume or **M**esh), and the surface- and/or curve-skeletonization methods it is compared with. Dashes show that a method does not compute the respective (curve or surface) skeleton type. Last three rows are survey papers.

Method		Compared with	
Type	Name	Surface skeleton	Curve skeleton
M	Jalba et al. [8]	[19]	[7,14]
M	Giesen et al. [19]	[12,47,48]	–
M	Huang et al. [18]	–	[15]
M	Au et al. [14]	–	[49,37,30,34]
M	Dey and Sun [34]	–	[30]
M	Tagliasacchi et al. [50]	–	[34,14]
V	Arcelli et al. [45]	–	0
V	Reniers et al. [7]	0	0
V	Hesselink et al. [26]	0	–
V	Siddiqi et al. [25]	0	–
V	Liu et al. [51]	–	[7,37]
V	Ju et al. [52]	[53]	[53]
M + V	Livesu et al. [35]	–	[30,34,14,51]
M	Sobiecki et al. [5]	–	[14–16,54]
M + V	Cornea et al. [3]	–	[37,49,30,13]
V	**Our contribution**	[25,26,7,52]	[25,7,52,51,45,37]

This paper is organized as follows. Section 2 reviews related work. Section 3 presents the compared methods and comparison criteria. Section 4 presents the comparison methodology. Section 5 presents our comparison results. Section 6 discusses these results. Section 7 concludes the paper.

2. Related work

2.1. Skeletonization methods

For a shape $\Omega \subset \mathbb{R}^3$ with boundary $\partial\Omega$, we first define its distance transform $DT_{\partial\Omega} : \mathbb{R}^3 \to \mathbb{R}^+$

$$DT_{\partial\Omega}(\boldsymbol{x} \in \Omega) = \min_{\boldsymbol{y} \in \partial\Omega} \|\boldsymbol{x} - \boldsymbol{y}\|. \tag{1}$$

The surface skeleton, also called the medial surface, S_Ω of Ω is next defined as

$$S_\Omega = \{\boldsymbol{x} \in \Omega \,|\, \exists \boldsymbol{f}_1, \boldsymbol{f}_2 \in \partial\Omega, \quad \boldsymbol{f}_1 \neq \boldsymbol{f}_2, \ \|\boldsymbol{x} - \boldsymbol{f}_1\| = \|\boldsymbol{x} - \boldsymbol{f}_2\| = DT_{\partial\Omega}(\boldsymbol{x})\} \tag{2}$$

where $\boldsymbol{f}_1$ and $\boldsymbol{f}_2$ are the contact points with $\partial\Omega$ of the maximally-inscribed ball in Ω centered at $\boldsymbol{x}$ [6,7], also called *feature transform* (FT) points [8]. When $\Omega \subset \mathbb{R}^2$, Eq. (2) yields the 2D skeleton, also called the medial axis, of the shape Ω. Surface skeletons contain several manifolds with boundaries which meet along a set of Y-intersection curves [9–11]. Curve skeletons are loosely defined as 1D structures locally centered within a shape $\Omega \subset \mathbb{R}^3$.

Surface and curve skeletons can be computed by geometric, distance field, general field, and thinning methods. **Geometric** methods include Voronoi diagrams [12] and subsets thereof [13], mesh contraction in normal direction [14–17], mean-shift-like clustering [18], and union-of-balls approaches [19,20,8]. Such methods use meshed shape representations and thus scale well to handle high-resolution models [20,8]. **Distance-field** methods find S_Ω along singularities of $DT_{\partial\Omega}$ [21–26] and can be efficiently done on GPUs [27,28]. **General-field** methods use fields smoother (with fewer singularities) than distance transforms [29–32]. Such methods are more robust for noisy shapes. Various regularization metrics, *e.g.* the angle between feature vectors [33,27], or the geodesic distance between feature points [34,7], are used to eliminate spurious skeleton details caused by noise on $\partial\Omega$. Field methods can also compute 3D curve skeletons by backprojecting 2D skeletons of 2D projections [35] or axis-aligned slices [36] of the shape back into 3D. Field methods are implemented for both voxel and mesh shapes. **Thinning** methods remove $\partial\Omega$ voxels while preserving connectivity [37,38]. Tools from mathematical morphology [39] were among the first used to compute curve skeletons by thinning. The residue of openings, based on Lantuéjoul's formula [40], usually leads to disconnected skeleton branches, whereas methods based on homotopic thinning transformations [40–42,37] yield connected skeletons. Constraining thinning by distance-to-boundary order [43–45] or flux-order [46] further enforces centeredness. Further details on related work and such methods are given in Section 3.2.

2.2. The challenge of comparison

Unsurprisingly, the wealth of existing skeletonization methods makes an exhaustive comparison hard. Aspects which contribute to this challenge are (a) different shape representations (voxels *vs* meshes *vs* point clouds), (b) the unavailability of several implementations, and (c) different skeleton definitions. The last aspect is particularly important: For *surface* skeletons, one could argue that Eq. (2) is a unique definition against which all methods can be checked. However, both spatial discretizations of Eq. (2) and heuristic regularizations that remove small-scale 'noise' details allow multiple weak forms of Eq. (2) [26,7,19]. For *curve* skeletons, the problem is even harder, as these have no unique definition, not even in the continuous $\mathbb{R}^3$ space.

Such aspects make it hard to analytically compare, and reason about, the properties of the produced skeletons. As such, qualitative comparisons have been proposed. In 2007, Cornea et al. compared four curve-skeletonization methods, one from each class listed in Section 2.1. To facilitate the comparison, they also propose several quality criteria that skeletons should obey. Six years later, this comparison was extended for six other contraction-based curve-skeletonization methods [5]. Schaap et al. proposed a quantitative comparison of 13 centerline extraction algorithms for coronary artery datasets [55]. As reference, they use a centerline constructed by manual annotation by expert users. However, in contrast to the tubular artery shapes considered in [55], manual construction of curve, and even more so of surface, skeletons for general 3D shapes not feasible. A cursory scan over many skeletonization papers shows that such method comparisons are very limited (see Table 1). As such, more comparison studies are strongly needed to better understand the strengths and limitations of existing methods.

3. Methods

We next describe a study that adds 4 surface and 6 curve skeletonization methods for voxel shapes to the existing comparison surveys mentioned in Section 2. Section 3.1 presents the desirable criteria that we compare against. Section 3.2 introduces the methods selected for comparison.

3.1. Comparison criteria

Following [3,1,5,8], we focus on the following well-known quality criteria for curve and surface skeletons:

Homotopy: The skeleton is topologically equivalent to the input shape (same number of connected components, cavities, and tunnels).

Thin: The skeleton should be as thin as the sampling model used allows it. Voxel-based skeletons should be one voxel thick, *i.e.*, no 2×2 foreground-voxel configurations should exist.

Centered: For surface skeletons, this is equivalent to Eq. (2). For curve skeletons, no unique centeredness definition exists. An