University of Groningen

Computing refined skeletal features from medial point clouds

Kustra, Jacek; Jalba, Andrei; Telea, Alexandru

Published in:
Pattern Recognition Letters

DOI:
[10.1016/j.patrec.2015.05.007](https://doi.org/10.1016/j.patrec.2015.05.007)

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:
2016

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

Citation for published version (APA):

Kustra, J., Jalba, A., & Telea, A. (2016). Computing refined skeletal features from medial point clouds. *Pattern Recognition Letters*, 76, 13-21. <https://doi.org/10.1016/j.patrec.2015.05.007>

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.

Computing refined skeletal features from medial point clouds[☆]

Jacek Kustra^{a,*}, Andrei Jalba^b, Alexandru Telea^c

^a Institute Johann Bernoulli, University of Groningen, Nijenborgh 9, 9747AG Groningen, The Netherlands

^b Department of Mathematics and Computer Science, Eindhoven University of Technology, P.O. Box 513, 5600MB Eindhoven, The Netherlands

^c Department of Mathematics and Computing Science, University of Groningen, Nijenborgh 9, 9747 AG Groningen, The Netherlands

ARTICLE INFO

Article history:

Available online 30 May 2015

Keywords:

Feature extraction
Shape segmentation
Point cloud
Skeleton segmentation
Polygonal

ABSTRACT

Medial representations have been widely used for many shape analysis and processing tasks. Large and complex 3D shapes are, in this context, a challenging case. Recently, several methods have been proposed that extract point-based medial surfaces with high accuracy and computational scalability. However, the resulting medial clouds are of limited use for shape processing due to the difficulty of computing refined medial features from such clouds. In this paper, we show how to bridge the gap between having a raw medial cloud and enriching this cloud with feature points, medial-point classification, medial axis decomposition into sheets, robust regularization, and Y-network extraction. We further show how such properties can be used to support several shape processing sample applications including edge detection and shape segmentation, for a wide range of complex 3D shapes.

© 2015 Elsevier B.V. All rights reserved.

1. Introduction

3D shape processing is a topic of rising importance due to developments in 3D shape modeling and acquisition techniques, such as structured light scanners, time-of-flight cameras or surface extraction from 3D volumetric scans. Such shapes $\Omega \subset \mathbb{R}^3$ are typically represented by boundary sampling, by polygonal meshes or dense point clouds. Processing such shapes is a vast research area encompassing analysis (detecting shape features like edges, flat areas, or thin regions), shape matching and registration, fairing and denoising, and segmentation.

While the above-mentioned processing can be done by using the shape's surface or boundary $\partial\Omega$, an alternative way is given by medial representations [44]. Since their introduction by [4], such representations have received increasing attention due to their potential to simplify and generalize many shape processing operations.

Computing medial axes for 2D shapes is a well-understood topic, with efficient, robust, and easy-to-use methods [13,28,44,52]. The 2D medial axis transform (MAT) full encodes a shape's boundary $\partial\Omega$ by its medial axis. In contrast, 3D shapes admit two kinds of medial structures: *Surface skeletons* generalize 2D medial axes and their MATs to fully capture the shape's geometry and topology. *Curve skeletons* capture the shape topology as a set of 1D curves, but do not provide the full boundary encoding given by the MAT. Hence, curve

skeletons are good descriptors for tubular shapes such as vessel structures, but are less suited as analysis and processing tools for generic 3D shapes.

Recently, fast and accurate medial-surface computation was made possible for complex 3D models of millions of polygons [19,24]. Yet, such methods provide only a medial *point cloud*, which is far from sufficient for shape analysis and processing. Indeed, medial surfaces have a complex structure, consisting of several so-called medial *sheets*, or manifolds. Medial points have several *types*, which describe the kind of surface points they correspond to via the MAT [15]. *Boundaries* of these curves, mapped via the MAT to shape edges [33,44], can be used for shape segmentation [36]. Curves where sheets meet, also called *Y-intersection* curves, help with shape reconstruction and matching [7,10,22]. The individual *sheets* map to separate shape parts, enabling shape simplification and segmentation.

While the above medial features can be relatively easily computed for voxel-based representations [35,36], their computation for point-cloud medial surfaces is far from trivial [19,21,24,51]. Extracting features such as endpoints, branches, and junctions from curve skeletons is much easier than for medial surfaces, making the former more frequently used in applications, even though they encode less information than the latter. To become effectively useful for real-world applications, medial-surface point-cloud methods need enhancement in the sense of classifying medial points, computing separate medial sheets, extracting sheet boundaries and Y-intersection curves, and mapping all these higher-level medial features robustly and efficiently to the shape surface.

In this paper, we show how to efficiently and robustly construct all above features from medial surface point-clouds, by combining

[☆] This paper has been recommended for acceptance by Punam Kumar Saha.

* Corresponding author. Tel.: +31629014636.

E-mail address: jacek.kustra@gmail.com (J. Kustra).

several input-shape and medial properties. This makes 3D medial clouds as easy to use as the more commonly used curve skeletons, and opens the way for using such medial clouds for many shape analysis and processing applications.

The structure of this paper is as follows. Section 2 reviews related work, with a focus on computing high-level features for medial surfaces. Section 3 presents our methods to compute such features from raw medial clouds. Section 4 illustrates the applicability of our computed medial features for part-based and patch-based shape segmentation and classification. Section 5 discusses our techniques. Section 6 concludes the paper.

2. Related work

2.1. Basic definitions

We first outline the definitions required to introduce and discuss related work. Given a shape $\Omega \subset \mathbb{R}^3$ with boundary $\partial\Omega$, we first define its distance transform $DT_{\partial\Omega} : \mathbb{R}^3 \rightarrow \mathbb{R}^+$

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

The surface skeleton of Ω is next defined as

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

where $\mathbf{f}_1$ and $\mathbf{f}_2$ are two of the contact points with $\partial\Omega$ of the maximally inscribed ball in Ω centered at $\mathbf{x}$, also called *feature points* [17,34,48]. The vectors $\mathbf{v}_1 = \mathbf{f}_1 - \mathbf{x}$ and $\mathbf{v}_2 = \mathbf{f}_2 - \mathbf{x}$ are called *feature vectors* or *spoke vectors* [9], and are locally parallel to $\nabla DT_{\partial\Omega}$ [43]. The tuple $(S_{\Omega}, DT_{\partial\Omega})$ is called the MAT of Ω , and provides a full description of the boundary $\partial\Omega$, i.e. allows one to fully reconstruct $\partial\Omega$ given the MAT. For 3D shapes, S_{Ω} is a set of manifolds with boundaries which meet along a set of so called Y-intersection curves [7,10,22]. The set of all feature points of a skeleton point

$$F(\mathbf{x} \in S_{\Omega}) = \{\mathbf{f} \in \partial\Omega \mid \|\mathbf{x} - \mathbf{f}\| = DT_{\partial\Omega}(\mathbf{x})\} \quad (3)$$

is called the *feature transform* (FT) of $\mathbf{x}$ [17,19]. The cardinality $\|F(\mathbf{x})\|$ of $F(\mathbf{x})$ is related to the location of $\mathbf{x}$ on skeletal manifolds [15,30] (see also Fig. 3): In the continuous case, points $\mathbf{x}$ inside skeletal manifolds, also called A_2^1 points, have exactly two feature points; points on Y-intersection curves of $k \geq 3$ manifolds, also called A_k^1 points, have k feature points; and points on manifold boundaries, also called A_3 points, have for $F(\mathbf{x})$ an entire circle sector or spherical segment, whose size is proportional with $DT_{\partial\Omega}(\mathbf{x})$. F and its inverse F^{-1} establish a bidirectional mapping between S_{Ω} and $\partial\Omega$, and are thereby important shape analysis instruments [44].

Besides surface skeletons (Eq. (2)), 3D shapes also admit so-called *curve skeletons*. These are curvilinear descriptors locally centered in the shape, according to various definitions [46]. Curve skeletons are effective in capturing the *topology* of tubular 3D shapes and also allow a limited reconstruction and part-based segmentation of such shapes [2,35]. However, they do not provide the full MAT given by surface skeletons, nor do they capture features such as shape edges or enable patch-based segmentation, fairing, or simplification for the more general non-tubular shape family [16,25,35]. A recent comparison of curve and surface skeletons is given in [45]. Given the above, we focus our discussion next on surface skeletons.

A key component of any skeletonization method is *regularization*, i.e. removing (pruning) from S_{Ω} all points which encode small-scale surface features, such as noise. This is usually done by computing an importance measure $\mu : S_{\Omega} \rightarrow \mathbb{R}^+$ and next upper-thresholding it [10,41]. Both local and global measures exist. *Local measures* cannot separate locally-identical, yet globally-different, contexts (see e.g. Fig. 1 in [38]). Thresholding local measures can disconnect skeletons [30,43,49] and makes pruning less intuitive [41]. Local measures include the angle between feature vectors, the distance-to-boundary

[1,14,49], divergence metrics [5,43], first-order moments [40], and indicators of the multi-valuedness of $\nabla DT_{\partial\Omega}$ [9,47]. A good survey of such methods is given in [44]. *Global measures* monotonically increase from the skeleton boundary ∂S_{Ω} inward. Thresholding them yields connected skeletons which capture details at a user-given scale. Miklos et al. [25] propose a measure based on the geometric properties of an union of balls (UoB) approximation of Ω [16]. Dey and Sun propose the medial geodesic function (MGF), equal to the length of the shortest geodesic between feature points [11,31]. Reniers et al. [38] extend the MGF for surface and curve skeletons using geodesic lengths and surface areas between geodesics, respectively, inspired by collapse metrics used to extract 2D multiscale skeletons [13,28,52].

2.2. Computing medial surfaces and their features

Many methods have been proposed in the last decade to compute medial surfaces. These can be classified in thinning, field, and geometric (mesh-based) methods [2,19,25]. Thinning removes $\partial\Omega$ voxels while preserving homotopy [2,3,29,32]. Thinning is simple and fast, but can be sensitive to rigid Euclidean transformations. Field methods find S_{Ω} along singularities of $DT_{\partial\Omega}$ or related fields [17,20,23,39,40,43,53] and can be efficiently implemented on GPUs [6,49]. Mesh methods use a surface sampling of $\partial\Omega$ only, which is cheaper and faster than the volumetric sampling of Ω used by field and thinning methods. They include Voronoi diagrams [1,12] and finding maximally inscribed balls [19,22,24,25]. Given their scalability, we focus next on mesh-based methods only.

Recent mesh-based surface skeletonization methods [19,24] can extract highly accurate medial surfaces of complex meshed surfaces of millions of polygons in subsecond time using GPU acceleration techniques, and thereby make medial surfaces usable in interactive contexts. However, such methods deliver only a *medial point cloud*, i.e., an unorganized set of medial points with two corresponding feature points. The practical usefulness of such descriptors is quite limited, as mentioned already in [24] and [19], since one cannot directly reason about the relationships between several medial-cloud points or the relationship of such points with the surface points. We identify two classes of challenges in this respect, as follows.

Feature points: Computing the correct feature transform F (Eq. (3)) for each medial point is crucial for several subsequent operations. For instance, any application that uses Giblin's classification [15] requires exactly estimating $\|F(\mathbf{x})\|$ for all $\mathbf{x} \in S_{\Omega}$. While this is relatively easy for voxel-based skeletons [33,36] due to the regular sampling of both $\partial\Omega$ and S_{Ω} , this is far from trivial for medial clouds. Incorrect or incomplete estimation of F creates challenges for all regularization methods that use feature points or feature vectors. Consider for instance the class of MGF methods, which measure the geodesic distance μ between two given feature points $\mathbf{f}_1 \in F(\mathbf{x})$ and $\mathbf{f}_2 \in F(\mathbf{x})$ of a medial point $\mathbf{x}$ [11,19,38]. A_3 and $A_{k,k \neq 2}^1$ points, in Giblin's terminology, have more than two feature points. Which pair $\{\mathbf{f}_1, \mathbf{f}_2\} \subset F(\mathbf{x})$ should we use here? Using *any* such pair breaks off the continuity of μ over S_{Ω} , which in turn means that we cannot obtain noise-free and compact medial surfaces by simply thresholding μ .

Skeleton decomposition: Decomposing the medial surface into separate sheets is a key step for using this descriptor for further shape analysis or classification [44,47]. The same applies to finding Y-intersection curves and sheet boundaries, i.e., the medial scaffold [22]. This decomposition has been previously attempted by using local medial geometry properties. This is relatively easy to do for finely-sampled voxel skeletons [37], leading to using medial surfaces to create compelling multiscale shape segmentations [36]. Doing all above for medial clouds is, however, far from trivial. As shown in [21], using generic point-cloud segmentation methods for 3D skeletons is doable, but extremely challenging, since medial surfaces consist of numerous intersecting manifolds with boundaries, which are hard to

Fig. 1. Refined skeletal features computed from medial point clouds (top row) and subsequently enabled applications (bottom row).

capture even by very dense point clouds. In addition, the parameter settings in [21] are far from trivial, as they are related to local medial geometrical properties, such as *maximum local connectivity angle*. This adds difficulty in decomposing the medial axis, and requires a case-by-case parameter followed by a visual check to ensure the medial axis is decomposed in a satisfactory manner. In this paper, a different approach is taken to compute these abstractions. Rather than using the local geometric properties, the relationships between the medial points $\mathbf{x}_i$ feature vectors $\mathbf{f}_1$, $\mathbf{f}_2$ and their correspondence to surface properties is used. The next sections address this method.

3. Computing refined medial features

To address the challenges outlined in Section 2.2, we next present several new methods for computing the above-mentioned higher-level features from skeletal point clouds. We start by showing how to robustly classify medial points following [15], find skeletal boundaries and Y-curves, and robustly regularize the medial surface (Section 3.1). We next use these features to robustly segment the medial surface into separate manifolds (Section 3.2). Fig. 1 shows the steps of our feature-computation pipeline and also the applications enabled by it. As input for all these operations, we only assume a surface skeleton represented by an unstructured and unoriented point-cloud having exactly one skeleton point per surface point and exactly two feature-points per skeleton point, as computed e.g. by [17,19,24].

3.1. Medial points classification

3.1.1. Estimating the feature transform

To classify unstructured medial clouds following [15], we first need to estimate the feature transform $F(\mathbf{x} \in S_\Omega)$ (Eq. (3)). As explained earlier, F is not directly available in most skeletonization methods; in particular, our point-cloud methods [19,24] only compute two feature points per skeleton point. To find all feature points, we proceed as follows. Let $DT_{\mathbf{x}}$ be next a shorthand for $DT_{\partial\Omega}(\mathbf{x})$. For each skeleton point $\mathbf{x} \in S_\Omega$, we first find the closest points $F_\tau(\mathbf{x}) \subset \partial\Omega$ in a radius $DT_{\mathbf{x}} + \tau$, where τ is defined as a fraction of $\epsilon_{\rho_{\partial\Omega}(\mathbf{x})}$, where $\rho_{\partial\Omega}(\mathbf{x})$ is the average point density on $\partial\Omega$ in a small neighborhood around $\mathbf{f}_1(\mathbf{x}) \cup \mathbf{f}_2(\mathbf{x})$, and ϵ is a small constant set to 0.1. The slightly increased radius determines that the set $F_\tau(\mathbf{x})$ will conservatively contain all feature points of $\mathbf{x}$, i.e. $F(\mathbf{x}) \subset F_\tau(\mathbf{x})$. Setting τ to track the local sampling density of $\partial\Omega$ allows us to conservatively estimate F_τ for non-uniformly sampled meshes without introducing too many false-positives, i.e., minimizing the set $F_\tau \setminus F$.

Given the finite tolerance τ and the discrete sampling of $\partial\Omega$, $F_\tau(\mathbf{x})$ will also contain surface points which are slightly further from $\mathbf{x}$ than feature points; this is especially salient for points $\mathbf{x}$ of type A_3 , that

Fig. 2. Skeleton point classification based on fuzzy F_τ analysis. The figure is drawn for a 2D skeleton, for illustration simplicity. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

map to circular or spherical sectors on $\partial\Omega$ via the feature transform. However, as we shall see next, the conservative estimation of $F(\mathbf{x})$ given by $F_\tau(\mathbf{x})$ does not pose any problems to our medial attribute computation.

3.1.2. Classification of medial points

Since $F_\tau(\mathbf{x})$ is essentially a dilated, or fuzzy, version of $F(\mathbf{x})$, it consists of one or several point clusters centered around actual feature points. A cluster $C_i(\mathbf{x}) \subset F_\tau(\mathbf{x})$ can be defined as

$$C_i(\mathbf{x}) = \{\mathbf{f} \in F_\tau(\mathbf{x}) \mid \max_{\mathbf{f} \in C_i} \min_{\mathbf{g} \in F_\tau(\mathbf{x})} \|\mathbf{f} - \mathbf{g}\| < \min_{\mathbf{f} \in C_i} \min_{\mathbf{h} \notin C_i} \|\mathbf{f} - \mathbf{h}\|\} \quad (4)$$

i.e., all points which are closer to each other than to any point from another cluster $C_{j \neq i}(\mathbf{x})$ for the same skeleton-point $\mathbf{x}$.

We observed that the number of these clusters is a good indicator of the type of the medial point $\mathbf{x}$: For A_3 points, there is one such cluster, whose diameter is proportional to $DT_{\mathbf{x}}$; for A_2^1 points, we find two clusters; and for $A_{k,k \geq 3}^1$ points, we find k clusters. To compute k , we cluster the point-set $F_\tau(\mathbf{x})$ by a single-linkage hierarchical agglomerative method based on the Euclidean distance between the points. Next, we cut the resulting dendrogram, or cluster-tree, at a distance value equal to the average local sampling density $\rho_{\partial\Omega}$. This results in k clusters. The value of k gives us the type of point $\mathbf{x}$, as explained above.

Let us justify why k is a good point-type indicator. Fig. 2a shows an *incorrect* classification of medial point $\mathbf{p}$ which is on the skeleton S_Ω branch ended by point $\mathbf{e}$. Since the intersection of $\partial\Omega$ with a ball $B_{\mathbf{p}}$ of radius $DT_{\mathbf{p}} + \tau$ and center $\mathbf{p}$ (dotted red circle) yields a single cluster (thick red line), $\mathbf{p}$ is incorrectly marked as A_3 rather than as

Fig. 3. Medial cloud classification into different point types.

A_2^1 . This is caused by (1) the value τ used to compute F_τ being too large; (2) $\mathbf{p}$ being close to $\mathbf{e}$; and (3) the bump on $\partial\Omega$ corresponding to $\mathbf{e}$ being too shallow. Consider now the minimal distance $d_{\min}$ from $\mathbf{p}$ that we have to move on S_Ω away from $\mathbf{e}$ to find a point $\mathbf{q}$ which is correctly classified as A_2^1 (Fig. 2b). This happens when the intersection of the ball $B_{\mathbf{q}}$ of radius $DT_{\mathbf{q}} + \tau$ and center $\mathbf{q}$ (dotted blue circle) yields two disconnected clusters on $\partial\Omega$ (marked thick blue). To find $d_{\min}$, note that the maximal ‘inward shift’ between the upper parts of $B_{\mathbf{p}}$ and $B_{\mathbf{q}}$ equals $\sigma = DT_{\mathbf{p}} - DT_{\mathbf{q}} + d_{\min}$. To cause the disconnection of the compact cluster in Fig. 2a, σ must be larger than the maximal bump height on $\partial\Omega$ that fits in the sphere-shell of thickness τ , i.e., $DT_{\mathbf{p}} - DT_{\mathbf{q}} + d_{\min} > \tau$. Since $DT_{\mathbf{q}} - DT_{\mathbf{p}} = d_{\min} \cos \alpha$, where α is the angle between a feature vector and the tangent plane to S_Ω , it follows that

$$d_{\min} > \frac{\tau}{1 - \cos \alpha}. \quad (5)$$

Separately, for a $\partial\Omega$ with local sampling density $\rho_{\partial\Omega}$, the corresponding skeleton sampling density is

$$\rho_S = \frac{\rho_{\partial\Omega}}{\sin \alpha}. \quad (6)$$

A correct classification should mark only a one sampling-point-thin ‘band’ of skeleton points as A_3 (skeleton boundary). If ρ_S is smaller than the minimal ball-shift $d_{\min}$ needed to change point type from A_3 to A_2^1 , this band gets thicker, leading to the incorrect classification in Fig. 2a. Substituting our value of $\tau = \epsilon \rho_{\partial\Omega}$ (Section 3.1.1) in Eq. (5), and next the values of $d_{\min}$ and ρ_S from Eqs. (5) and (6) in the inequality $\rho_S < d_{\min}$, it follows that incorrect classification can only appear if $\epsilon \geq \frac{1 - \cos \alpha}{\sin \alpha}$. For our chosen value of $\epsilon = 0.1$, this implies $\alpha \lesssim 11.4^\circ$. In other words, for any medial sheets except those corresponding to highly obtuse angles on $\partial\Omega$, our method finds skeleton boundaries (A_3 points) which are precisely one sample-point thick.

Following the above, if $k = 1$ or $k = 2$, we can confidently say that we have found an A_3 , respectively A_2^1 , skeleton point. As k increases, the spatial separation of the clusters decreases too, so k does not reflect accurately the skeleton point type. We have empirically verified that the cluster count k accurately finds A_3^1 up to A_4^1 points for densely-sampled surfaces $\partial\Omega$. A more robust way to find such Y-curve points, that is far less sensitive on the sampling density of $\partial\Omega$, is described further in Section 3.2.2, based on the segmentation of S_Ω into individual medial sheets.

Related to our work, Reniers et al. [33] found A_3 points by computing the set difference between the full medial surface S and a simplified version S_τ of S , where τ is a small fixed value and simplification uses the MGF metric (Section 2.1). Compared to our approach, their method does not find $A_{k,k>1}^1$ points, and does not give an analysis of how to set parameter values.

Fig. 4. Skeleton regularization. (a) Rounded spleen shape with feature vectors shown for A_3 points. (b) Skeleton regularized by filtering A_3 points. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

3.1.3. Skeleton regularization using A_3 edge filtering

As outlined in Section 2.1, the MGF metric [11,38] provides very good regularization properties such as separating spurious skeleton points from important ones while maintaining skeleton connectivity. The MGF importance $\mu(\mathbf{x})$ of a medial point $\mathbf{x}$ equals the length of the longest shortest-geodesic on $\partial\Omega$ between any two feature points of $F(\mathbf{x})$. Hence, the MGF requires an accurate computation of the feature transform F (Eq. (3)). As discussed in Section 3.1.1, we compute a conservative F_τ which may contain tens of feature points for A_3 -type points. Computing shortest-geodesics between all such point-pairs is very expensive. Given this cost, Jalba et al. [19] and Reniers et al. [38] compute the MGF using only two feature points per skeleton point, i.e. implicitly consider all medial points to be of type A_2^1 . This has two problems. First, the importance μ for A_3 points will be typically underestimated, since one has no guarantee of finding the longest shortest-geodesic connecting any two feature points. This, in turn, creates a relatively jagged appearance of the simplified skeleton. Secondly, computing the MGF is expensive for large models, even when using only two feature points per medial point and highly optimized GPU implementations [19].

We propose here an alternative way to regularize medial surfaces by simply filtering A_3 points found by our classification. Fig. 4 shows this for a shape having highly rounded edges, i.e. whose A_3 points have many feature points. This is the kind of shape where the aforementioned problem of the MGF metric occurs. Fig. 4a shows the medial cloud with feature vectors (in red) for the A_3 points. Fig. 4b shows our regularized skeleton, with all noisy points being removed. Since A_3 points appear only on the medial boundary by definition, our regularization does not create gaps or disconnect the medial surface. Since our method requires only a simple clustering of feature points based on their Euclidean distance, it is considerably faster than the MGF metric (see Section 5 for details). However, in contrast to the MGF, our method cannot deliver a multiscale of progressively simplified skeletons; we can only remove the finest scale of noisy boundary points. As such, our regularization is useful when one needs a clean and detail-rich surface skeleton for further processing, rather than a multiscale skeleton representation.

3.2. Surface skeleton decomposition

Besides classifying skeleton points, higher level features can be computed. One such feature is the decomposition of the medial surface into separate sheets, used in shape analysis and segmentation tasks [22,36]. While such decompositions can be computed relatively

Fig. 5. Medial sheet computation: (a) Distance function δ , illustrated in 2D. (b) Sparse distance matrix, used as an input for hierarchical point clustering. (c) Medial sheets found for a palatine bone shape (see Section 3.2.1).

easy for voxel skeletons [37], this is challenging for medial clouds, especially when these contain a large number of complex sheets [21].

We address this task by clustering the medial cloud based on a novel definition of medial sheets that uses the medial cloud properties (Section 3.2.1). Next, we use this decomposition to robustly find Y-intersection curves where several such sheets meet (Section 3.2.2). Finally, we use the feature transform to construct compact (meshed) sheet representations (Section 3.2.3).

3.2.1. Medial sheet computation

We first define the distance of two A_2^1 points $\mathbf{x}$ and $\mathbf{y}$ as

$$\delta(\mathbf{x}, \mathbf{y}) = \sum_{\mathbf{a} \in F(\mathbf{x})} \min_{\mathbf{b} \in F(\mathbf{y})} \text{MGF}(\mathbf{a}, \mathbf{b}), \quad (7)$$

where $\text{MGF}(\mathbf{a}, \mathbf{b})$ is the medial geodesic function, i.e. the length of the shortest geodesic on $\partial\Omega$ between feature points $\mathbf{a}$ and $\mathbf{b}$ [38]. Next, we define a medial sheet γ as all medial points having a distance δ lower than a threshold τ

$$\gamma = \{\mathbf{x} \in S_\Omega \mid \exists \mathbf{y} \in S_\Omega, \delta(\mathbf{x}, \mathbf{y}) < \tau\}. \quad (8)$$

The rationale behind Eq. (8) is that medial points $\mathbf{x}$ and $\mathbf{y}$ which are close and belong to the same sheet have small distances (along $\partial\Omega$) between their corresponding feature points. This statement is supported as follows: Since medial sheets are locally smooth and have a low curvature [30], their feature vectors vary smoothly and slowly locally; in turn, this implies that the corresponding feature points vary slowly and smoothly across $\partial\Omega$. Fig. 5 illustrates this for a 2D shape (for simplicity): Medial points $\mathbf{x}$ and $\mathbf{y}$ are on the same sheet, and have small MGF distances between their feature points, thus a small $\delta(\mathbf{x}, \mathbf{y})$. In contrast, medial points $\mathbf{w}$ and $\mathbf{z}$, which belong to different sheets, have at least two feature vectors pointing in different directions, thus a large $\delta(\mathbf{w}, \mathbf{z})$.

Eqs. (7) and (8) define medial sheets without using any medial connectivity information, and thus allow us to segment a medial cloud into its sheets, as follows. First, we define a distance matrix M encoding the distances $\delta(\mathbf{x}, \mathbf{y})$ between all medial point pairs. For efficiency, we only compute matrix entries that correspond to δ values below our chosen threshold τ , since the sheet definition (Eq. (8)) only requires to know when $\delta < \tau$. Secondly, when computing $\delta(\mathbf{x}, \mathbf{y})$, if the length of the geodesic traced on $\partial\Omega$ from $\mathbf{x}$ to $\mathbf{y}$ exceeds τ , we stop tracing it and skip the respective matrix entry. Overall, this turns the computation and storage of M from a quadratic process in the number of medial points into a linear process, since any medial point $\mathbf{x}$ has only a limited number of points $\mathbf{y}$ at close distance δ from it. Finally, we use M as input for a single-linkage hierarchical clustering [18], which outputs a partition of S_Ω into a set of medial sheets γ_i , so that $\gamma_i \cap \gamma_{j \neq i} = \emptyset$ and $\cup_i \gamma_i = S_\Omega$. Fig. 5c illustrates the separated sheets of the medial surface of a palatine bone shape. Same-sheet points are

Fig. 6. Medial sheet extraction: (a) [37]; (b), (d), (e), (g), (h) [21]; (c), (f), (i) our method. For [21], its key parameter, the number of nearest neighbors NN, is indicated. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

marked by the same color. Although the medial cloud is quite complex, its sheets are cleanly separated. Such sheets can be processed to create compact representations thereof, as discussed next in Section 3.2.3.

Fig. 6 compares our method for extracting sheets from a medial cloud with two other methods. Fig. 6a shows the method of [37], which works in brief as follows: Given a (voxel) medial surface S , its Y-network voxels S_Y are found based on the cardinality of the feature transform for A_2^1 points (Section 2.1). Next, separate medial sheets are found as being the connected components of the voxel set $S \setminus S_Y$. This method is quite sensitive to the voxel sampling of the input shape. For example, the cog wheel detail in Fig. 6a (128³ voxels) shows two separate components c_1 (red) and c_2 (purple), which actually are part

Fig. 7. Y-network extraction with Y-curve points colored green. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

of the same sheet. These are wrongly separated since (1) the sampling resolution disconnects the detected medial sheet half-way and (2) sheet detection is based on connected component finding. Fig. 6b shows the method of [21], which is a general-purpose clustering of 3D point clouds based on a similarity metric that accounts for both distance (computed between nearest neighbors in the point cloud) and local sheet flatness. Fig. 6c shows our method. As visible, both Kustra et al. [21] and our method correctly detect a single sheet c instead of the two separate fragments c_1 and c_2 . Images (d)–(i) further compare our method with [21] for two shapes and two different values of the nearest-neighbor count NN, the key parameter of [21]. Two observations can be made here. First, we see how the results of [21] depend quite strongly on the NN choice. In contrast, our method does not need such parameter tuning. Secondly, and more importantly, the point-similarity used by Kustra et al. [21] is purely *local*, as it uses only inter-point distances and local flatness of the sheets. In contrast, our method uses a distance function (Eq. (7)) which captures *global* shape properties, due to the underlying MGF function. This makes the sheet computation far less sensitive to local shape variations or small-scale noise.

3.2.2. Y-intersection curve extraction

Once the medial sheets γ_i are found, the Y-intersection curves can be found as those points $\mathbf{x} \in S_\Omega$ that belong to at least two different sheets. However, performing this test directly on the medial sheet-set is not possible, since our sheets are disjoint, i.e. $\gamma_i \cap \gamma_{j \neq i} = \emptyset$. Hence, we find Y-curve points as those skeleton points $\mathbf{x}$ which have at least one k -nearest neighbor belonging to a different sheet than the one containing $\mathbf{x}$. Tuning k allows controlling the thickness of the Y network being computed. Fig. 7 shows three examples of Y networks, computed for $k = 3$.

3.2.3. Computing compact medial sheets

In Section 3.2.1, we computed medial sheets as unstructured point clouds. Many shape processing operations require *compact* sheets, e.g. in the form of a triangle mesh. We show next how such meshes can be created based on an analysis of the feature vectors $\mathbf{v}_1(\mathbf{x})$ and $\mathbf{v}_2(\mathbf{x})$ of each skeleton point $\mathbf{x}$ (Section 2.1). The key idea is to use the feature vectors to back-project the connectivity information captured by the $\partial\Omega$ mesh onto each sheet γ . The method has two steps, as follows.

Feature vector alignment: The projection

$$P(\gamma) = P_1(\gamma) \cup P_2(\gamma) = \{\mathbf{x} \in \partial\Omega \mid \exists \mathbf{y} \in \gamma, \mathbf{x} \in \{\mathbf{f}_1(\mathbf{y}), \mathbf{f}_2(\mathbf{y})\}\} \quad (9)$$

of a sheet γ consists of two compact areas $P_1(\gamma)$ and $P_2(\gamma)$ of $\partial\Omega$, one for each side of γ . If we can isolate any of these two areas, we can next simply transfer its connectivity information onto γ to obtain our desired sheet mesh. For this, we reorder, or align, the feature vectors $\mathbf{v}_1(\mathbf{x})$ and $\mathbf{v}_2(\mathbf{x})$ of all sheet points so that all $\mathbf{f}_1$ are included in $P_1(\gamma)$, and all $\mathbf{f}_2$ are included in $P_2(\gamma)$, as follows. First, we select an arbitrary reference point $\mathbf{x}_{\text{ref}} \in \gamma$ and mark it as visited. We next visit all other points $\mathbf{x} \in \gamma$ in order of increasing distance to $\mathbf{x}_{\text{ref}}$ and redefine their

feature points as

$$\mathbf{f}_i = \arg \min_{\mathbf{f} \in \{\mathbf{f}_1, \mathbf{f}_2\}} \text{MGF}(\mathbf{f}, \mathbf{f}_i^{\text{vis}}), \quad i \in \{1, 2\}, \quad (10)$$

where $\mathbf{f}_i^{\text{vis}}$ is the closest visited (aligned) feature point to $\mathbf{f}_i$, and mark $\mathbf{f}$ as visited. When all points of γ are visited, all feature vectors $\mathbf{v}_1$ will be on the same side of γ as $\mathbf{v}_1(\mathbf{x}_{\text{ref}})$, while all $\mathbf{v}_2$ will be on the other side. We can next find the projection of side $i \in \{1, 2\}$ of γ as $P_i(\gamma) = \{\mathbf{x} \in \partial\Omega \mid \exists \mathbf{y} \in \gamma, \mathbf{x} = \mathbf{f}_i(\mathbf{y})\}$.

Connectivity projection: We finally construct a meshed version of γ by simply copying all triangle information from $P_i(\gamma)$ to γ , with i being either 1 or 2 (both sides are equally good). That is, for any triangle $t = \{\mathbf{x}^i\}_{1 \leq i \leq 3}$ in e.g. $P_1(\gamma)$, we construct a triangle $t_\gamma = \{\mathbf{y}^i\}_{1 \leq i \leq 3}$ where $\mathbf{x}^i = \mathbf{f}_1(\mathbf{y}^i)$. Fig. 8 illustrates the resulting meshed sheets for the surface skeletons of several complex anatomical shapes from the open database in [26], where neighbor sheets have different colors for illustration purposes. Given these meshed sheets, we can now use any polygon-based geometric algorithm to analyze or process them further, e.g., to estimate curvature, areas, elongation, or compute shortest paths or distance fields.

4. Applications

We next use our computed medial features (point classification, regularization, and medial surface decomposition into sheets) to support several shape analysis applications. These examples implicitly illustrate the quality and robustness of our feature computation methods. Secondly, they show how such features enhance the added-value of surface skeletons by allowing it to support the construction of surface processing tools.

4.1. Surface edge detection

Finding edges on a 3D surface has many applications in segmentation and classification. Most existing edge detectors are based on the surface's curvature tensor [8,27,50]. A problem of such detectors is that they operate at a given scale, i.e. find edges of a sharpness range which must be specified. Using skeletons allows finding both sharp and blunt edges, i.e., removes the need to specify an edge-sharpness range: Following the observation that medial surface boundaries (A_3 points) correspond to curvature maxima or edges on the input surface [30], Reniers et al. compute surface edges by finding A_3 points as explained in Section 3.1.2, and next back-projecting these on the input surface by the feature transform [33]. We propose here an alternative approach: For each A_3 skeletal point $\mathbf{x}$, detected as explained in Section 3.1.2, we find all surface points enclosed in a sphere of radius $DT_{\mathbf{x}} + \tau$, with τ set as explained in Section 3.1.1, and assign to each surface point the smallest $DT_{\mathbf{x}}$ value which encloses it. Remaining surface points are assigned a value of $\max(DT_{\mathbf{x}})$. Fig. 9(a)–(c) compares our method with the classical curvature detector of [50] and with [33] for a brain cortex surface. The goal is to find the sulcal brain structures, which correspond to (soft) convex surface edges, an important task in many structural and functional anatomic brain analyses. The presence of sulci is shown using a blue (concave) to red (sharp convex) rainbow colormap, mapping the three studied detectors: mean curvature [50] (Fig. 9a), geodesic distance to back-projected A_3 points [33] (Fig. 9b), and our sphere-radius metric (Fig. 9c). Our method achieves a sharper sulci separation than [33], which in turn performs better than [50]. Fig. 9(d)–(f) show our method applied to three additional shapes which exhibit a mix of sharp and blunt edges. As visible, our detector finds both sharp (and thus, thin) and blunt (and thus, thick) edges. The edge sharpness and thickness is also visible in the color mapping.

Fig. 8. Compact medial sheets computed for several anatomical models (Section 3.2.3).

Fig. 9. Soft edge detection using (a) curvature estimation [50]; (b) skeleton method of [33]; (c)–(f) our method. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

4.2. Patch-based segmentation

Patch-based segmentation (PBS) divides a shape $\partial\Omega$ into *patches*, i.e. quasi-flat areas which are separated by sharp creases. Most PBS methods work by clustering surface points using, as similarity metric, the surface curvature or similar quantities [42]. Since medial surfaces fully capture the surface information via the MAT (Section 2.1), these medial surfaces can be used for PBS. For this, Reniers and Telea [36] compute soft edges by using the feature transform of low-importance medial-surface points, and next use these thick edges to segment the shape. However, their method needs to handle a large number of special cases (and is thereby quite complicated), and only works for voxel shapes. We propose here a much simpler approach: We project all skeleton-boundary points $\mathbf{p}$ (type A_3) to $\partial\Omega$ via our extended feature transform F_τ , i.e. compute the set $E = \{\mathbf{x} \in F_\tau(\mathbf{p}) | \mathbf{p} \in S_\Omega \wedge \text{type}(\mathbf{p}) = A_3\} \subset \partial\Omega$. The set E consists of a thick version of the edges of $\partial\Omega$. Due to the conservativeness of F_τ (Section 3.1.1), E will contain connected edges, in contrast to e.g. a naive thresholding of the curvature of $\partial\Omega$ or

other similar local surface classifiers. Hence, we next find patches by simply computing connected components of $\partial\Omega \setminus E$. Finally, we add the points in E to their closest patch, thereby making the resulting patches become a partition of $\partial\Omega$.

Fig. 10 shows our results, using the same color scheme as Fig. 8. For models with clear, sharp, edges, we see how patches neatly follow these edges (e.g. Fig. 10a, rib sockets in Fig. 10g, skull concavity in Fig. 10i). More importantly, our method handles equally well models with soft edges (Fig. 10(b), (c), and (f)) and/or mixes of sharp and soft edges (Fig. 10(d), (g), and (h)).

4.3. Medial sheet mapping segmentation

In contrast to patch-based segmentation (Section 4.2), part-based segmentation (pBS) separates a shape $\partial\Omega$ into components that are perceived as being the natural ‘parts’ of the shape [42]. Among the many methods for pBS, *curve skeletons* are often used, as they readily capture the part-whole topology of shapes having elongated protrusions. One way to compute a pBS is to find the so-called junction points of curve skeletons (equivalent to Y-curves for surface skeletons), and then cut the shape with curves that go around these points [35]. Such methods are robust and relatively simple to implement, but work well only for shapes with a tubular structure, i.e., which have a meaningful curve skeleton. We propose here to use the surface skeleton for pBS. For this, we compute its medial sheets γ (Section 3.2.1), and next project these into $\partial\Omega$ using $P(\gamma)$ (Eq. (9)). Since all points on $\partial\Omega$ have a skeleton point by construction [19], the entire shape is covered by such projections, which give us the ‘parts’ of the shape. The borders separating two such neighbor parts are nothing but the projections of the Y-curves. Since such curves are smooth [44], and the feature-vector field used for projection is also smooth (since parallel to $\nabla \text{DT}_{\partial\Omega}$ which is divergence-free away from the skeleton, see [43]), the resulting part-borders will also be smooth. Fig. 11 show several part-based segmentation examples. Although many alternative pBS segmentations are possible, we argue that the found segments match well the perceived distinct shape parts. We note that such segmentations cannot be achieved using only curve-skeletons, since the shown shapes do not have a tubular structure.

Fig. 10. Patch based segmentation (Section 4.2).

Fig. 11. Medial sheet mapping segmentation (Section 4.3).

Table 1
Medial sheet extraction times on the CPU.

Model	Points $\ \partial\Omega\ $	Sheet extraction (s)
Hip bone	24,852	46.87
Xiphoid	64,690	68.88
Vertebra	48,924	49.09
Manubrium	65,038	64.22
Spleen	29,507	55.35
Sternum	27,525	62.26
Scapula	233,856	436.57
Ear	24,900	49.58
Cortex	58,490	133.98
Gyrus	29,116	56.75
Frontal bone	48,171	131.54

Medial sheet extraction: Separating sheets from a raw medial cloud is a hard task, for which few methods exist, and for which generic point-cloud clustering tools cannot be used [7,21]. Our contribution—the similarity metric (Eq. (7))—combines both local and global shape information, and enables a medial cloud segmentation into sheets which is noise-resistant and has a simple parameter setting. In detail, Kustra et al. [21] require tuning three parameters: The number of nearest neighbors of each skeletal point, the maximal allowed local-flatness of each sheet, and the sheet similarity. In contrast, our method requires a single parameter, the maximum MGF distance between feature-points of two skeleton-points that are on the same sheet (τ in Eq. (8)). For all tested shapes, a value of τ equal to four times the local point-density $\rho_{\partial\Omega}$ on the input surface yielded optimal results such as shown in this paper.

Y-network extraction: Our Y-network extraction finds the points around the Y-network of the skeleton. Exact Y-network points are, by definition, not explicitly found by the core skeletonization method we build atop, since this method always assumes two contact points for each medial point (Eq. (2)).

Scalability: On an Intel Core i7 3.8 GHz computer, our method, implemented in single-threaded C++, performs all described steps (medial feature computation, edge detection, PBS and pBS) in under 3 s for all shapes in this paper, which range between 30K and 230K skeleton points. Memory used is linear with the input point count. Medial sheet extraction is more costly, as it uses the expensive MGF metric (Eq. (7)). Timings for this step are given in Table 1, for CPU-based MGF computation. If desired, higher performance can be easily obtained by GPU-based MGF computation ([19]).

5. Discussion

We discuss next several aspects related to our contribution—showing that we can efficiently and easily compute high-level medial features from large point-cloud skeletons, and that using such features in various applications is a practical proposition.

Generality: We require as input only a raw medial 3D cloud with two feature points per skeleton point. Such clouds can be very efficiently computed by recent GPU methods [19,24] or older CPU methods [17], for any type of 3D shape topology or geometry.

Point classification: To our knowledge, our work is the first attempt to compute Giblin's medial point classification [15] for raw medial clouds. Using this classification to regularize skeletons (Section 3.1.3) is much simpler to implement, and also much faster, than the MGF metric. Compared to local regularization metrics [30,43,49], we do not disconnect the skeleton, and only remove a thin layer of boundary points. This gives a simple, fast, automatic, and effective way to create clean medial surfaces that preserve relevant skeletal details.

Limitations: The quality of our medial features highly depends on the quality of the input medial cloud. This depends next on the sampling density of the input shape, since we require only two feature points per medial point (Section 3).

Applications: For the segmentation and classification applications in Section 4, we note that better specific techniques (not using medial descriptors) exist. Our sample applications are aimed at showing the possibilities that refined medial features open, as alternatives and in contrast to established approaches, and not as a definitive solution to the underlying use-cases.

6. Conclusions

We have presented a set of techniques for computing refined medial features from raw medial-surface point clouds. These features include medial point classification, skeleton regularization, Y-network extraction, separating medial sheets, and reconstructing meshed sheets. Such features enrich the level on which one can reason about medial surfaces, and open new ways for shape processing applications using medial clouds. We provide, for illustration, sample applications for edge detection and shape segmentation. Overall, our work shows that the more complex (and information-rich) surface skeletons can be, technically, used with the same ease and computational efficiency as the simpler, and so far more frequently used, curve skeletons.

References

- [1] N. Amenta, S. Choi, R. Kolluri, The power crust, in: Proceedings of the SMA, ACM, 2001, pp. 65–73.
- [2] C. Arcelli, G. Sanniti di Baja, L. Serino, Distance-driven skeletonization in voxel images, IEEE TPAMI 33 (4) (2011) 709–720.
- [3] X. Bai, L. Latecki, W.Y. Liu, Skeleton pruning by contour partitioning with discrete curve evolution, IEEE TPAMI 3 (29) (2007) 449–462.
- [4] H. Blum, A transformation for extracting new descriptors of shape, in: Models for the Perception of Speech and Visual Form, MIT Press, 1967, pp. 362–380.
- [5] S. Bouix, K. Siddiqi, A. Tannenbaum, Flux driven automatic centerline extraction, Med. Imag. Anal. 9 (3) (2005) 209–221.
- [6] T. Cao, K. Tang, A. Mohamed, T. Tan, Parallel banding algorithm to compute exact distance transform with the GPU, in: Proceedings of SIGGRAPH I3D Symposium, ACM, 2010, pp. 134–141.
- [7] M. Chang, F. Leymarie, B. Kimia, Surface reconstruction from point clouds by transforming the medial scaffold, CVIU 113 (2009) 1130–1146.
- [8] U. Clarenz, M. Rumpf, A. Telea, Robust feature detection and local classification of surfaces based on moment analysis, IEEE TVCG 10 (5) (2004) 516–524.
- [9] S. Stolpner, S. Whitesides, K. Siddiqi, Sampled medial loci and boundary differential geometry, in: Proceedings of IEEE 3DIM, IEEE, 2009, pp. 87–95.
- [10] J. Damon, Global medial structure of regions in $\mathbb{R}^3$, Geom. Topol. 10 (2006) 2385–2429.
- [11] T. Dey, J. Sun, Defining and computing curve skeletons with medial geodesic functions, in: Proceedings of the SGP, IEEE, 2006, pp. 143–152.
- [12] T. Dey, W. Zhao, Approximating the medial axis from the Voronoi diagram with a convergence guarantee, Algorithmica 38 (2003) 179–200.
- [13] A.X. Falcão, J. Stolfi, R.A. Lotufo, The image foresting transform: Theory, algorithms, and applications, IEEE TPAMI 26 (1) (2004) 19–29.
- [14] M. Foskey, M. Lin, D. Manocha, Efficient computation of a simplified medial axis, in: Proceedings of IEEE SMA, ACM, 2003, pp. 135–142.
- [15] P. Giblin, B. Kimia, A formal classification of 3D medial axis points and their local geometry, IEEE TPAMI 26 (2) (2004) 238–251.
- [16] J. Giesen, B. Miklos, M. Pauly, C. Wormser, The scale axis transform, in: Proceedings of Annual Symposium on Computational Geometry, ACM, 2009, pp. 106–115.
- [17] W. Hesselink, J. Roerdink, Euclidean skeletons of digital image and volume data in linear time by the integer medial axis transform, IEEE TPAMI 30 (12) (2008) 2204–2217.
- [18] A.K. Jain, M.N. Murty, Data clustering: A review, ACM Comput. Survey 31 (3) (1999) 264–323.
- [19] A. Jalba, J. Kustra, A. Telea, Surface and curve skeletonization of large 3D models on the GPU, IEEE TPAMI 35 (6) (2013) 1495–1508.
- [20] R. Kimmel, D. Shaked, N. Kiryati, A. Bruckstein, Skeletonization via distance maps and level sets, CVIU 62 (3) (1995) 382–391.
- [21] J. Kustra, A. Jalba, A. Telea, Robust segmentation of multiple intersecting manifolds from unoriented noisy point clouds, CGF 33 (1) (2014) 73–87.
- [22] F. Leymarie, B. Kimia, The medial scaffold of 3D unorganized point clouds, IEEE TVCG 29 (2) (2007) 313–330.
- [23] F. Leymarie, M. Levine, Simulating the grassfire transform using an active contour model, IEEE TPAMI 14 (1) (1992) 56–75.
- [24] J. Ma, S. Bae, S. Choi, 3D medial axis point approximation using nearest neighbors and the normal field, Vis. Comput. 28 (1) (2012) 7–19.
- [25] B. Miklos, J. Giesen, M. Pauly, Discrete scale axis representations for 3D geometry, in: Proceedings of the ACM SIGGRAPH, ACM, 2010, pp. 394–493.
- [26] N. Mitsuhashi, K. Fujieda, T. Tamura, S. Kawamoto, T. Takagi, K. Okubo, Body-Parts3D: 3D structure database for anatomical concepts, Nucl. Acid Res. 37 (1) (2009) 782–785.
- [27] H. Moreton, C. Séquin, Functional optimization for fair surface design, in: Proceedings of ACM SIGGRAPH, 1992, pp. 167–176.
- [28] R.L. Ogniewicz, O. Kubler, Hierarchic Voronoi skeletons, Pattern Recog. 28 (3) (1995) 343–359.
- [29] K. Palagyi, A. Kuba, Directional 3D thinning using 8 subiterations, in: Proceedings of DGCI, 1568, Springer, 1999, pp. 325–336.
- [30] S. Pizer, K. Siddiqi, G. Szekely, J. Damon, S. Zucker, Multiscale medial loci and their properties, IJCV 55 (2–3) (2003) 155–179.
- [31] S. Prohaska, H.C. Hege, Fast visualization of plane-like structures in voxel data, in: Proceedings of IEEE Visualization, IEEE Computer Society, 2002, p. 2936.
- [32] C. Pudney, Distance-ordered homotopic thinning: a skeletonization algorithm for 3D digital images, CVIU 72 (3) (1998) 404–413.
- [33] D. Reniers, A. Jalba, A. Telea, Robust classification and analysis of anatomical surfaces using 3D skeletons, in: Proceedings of VCBM, 2008, pp. 61–68.
- [34] D. Reniers, A. Telea, Tolerance-based feature transforms, in: Advances in Computer Graphics and Computer Vision, Springer, 2007, pp. 187–200.
- [35] D. Reniers, A. Telea, Part-type segmentation of articulated voxel-shapes using the junction rule, CGF 27 (7) (2008) 1837–1844.
- [36] D. Reniers, A. Telea, Patch-type segmentation of voxel shapes using simplified surface skeletons, CGF 27 (7) (2008) 1954–1962.
- [37] D. Reniers, A. Telea, Segmenting simplified surface skeletons, in: Proceedings of DGCI, Springer, 2008, pp. 132–145.
- [38] D. Reniers, J.J. van Wijk, A. Telea, Computing multiscale skeletons of genus 0 objects using a global importance measure, IEEE TVCG 14 (2) (2008) 355–368.
- [39] L. Rossi, A. Torsello, An adaptive hierarchical approach to the extraction of high resolution medial surfaces, in: Proceedings of 3DIMPVT, IEEE Computer Society, 2012, pp. 371–378.
- [40] M. Rumpf, A. Telea, A continuous skeletonization method based on level sets, in: Proceedings of VisSym, Eurographics Association, 2002, pp. 151–158.
- [41] D. Shaked, A. Bruckstein, Pruning medial axes, CVIU 69 (2) (1998) 156–169.
- [42] A. Shamir, A survey on mesh segmentation techniques, CGF 27 (6) (2008) 1539–1556.
- [43] K. Siddiqi, S. Bouix, A. Tannenbaum, S. Zucker, Hamilton-Jacobi skeletons, IJCV 48 (3) (2002) 215–231.
- [44] K. Siddiqi, S. Pizer, Medial Representations: Mathematics, Algorithms and Applications, Springer, 2009.
- [45] A. Sobiecki, A. Jalba, A. Telea, Comparison of curve and surface skeletonization methods for voxel shapes, Pattern Recog. Lett. 47 (2014) 147–156.
- [46] N. Cornea, D. Silver, P. Min, Curve-skeleton properties, applications, and algorithms, IEEE TVCG 13 (3) (2007) 87–95.
- [47] S. Stolpner, S. Whitesides, K. Siddiqi, Sampled medial loci for 3D shape representation, CVIU 115 (5) (2011) 695–706.
- [48] R. Strzodka, A. Telea, Generalized distance transforms and skeletons in graphics hardware, in: Proceedings of VisSym, Eurographics Association, 2004, pp. 221–230.
- [49] A. Sud, M. Foskey, D. Manocha, Homotopy-preserving medial axis simplification, in: Proceedings of SPM, 2005, pp. 103–110.
- [50] G. Taubin, Estimating the tensor of curvature of a surface from a polyhedral approximation, in: Proceedings of ICCV, IEEE, 1995, pp. 902–907.
- [51] A. Telea, A. Jalba, Computing curve skeletons from medial surfaces of 3D shapes, in: Proceedings of TPCG, EG Press, 2012, pp. 273–280.
- [52] A. Telea, J.J. van Wijk, An augmented fast marching method for computing skeletons and centerlines, in: Proceedings of VisSym, Eurographics Association, 2002, pp. 251–259.
- [53] M. Wan, F. Dachille, A. Kaufman, Distance-field based skeletons for virtual navigation, in: Proceedings of IEEE Visualization, IEEE Computer Society, 2001, pp. 239–246.