

Practical Medial Axis Filtering for Occlusion-Aware Contours

M. Livesu and R. Scateni

Department of Mathematics and Computer Science, University of Cagliari, Italy


Abstract

We propose a filtering system for occlusion-aware contours. Given a point of view, we use the silhouette of a 3D shape from that point of view, its medial axis and a map of the occluded areas. Our filter is able to select the points of the medial axis which are projections of the curve-skeleton of the 3D shape, discarding all the points affected by occlusions. Our algorithm is easy to implement and works in real time. It can be plugged as is into existing methods for curve-skeleton extraction from 2D images; it can be used to robustly rank silhouettes according to how much they are representative of the 3D shape that generated them and can also be used for shape recognition from images or video sequences.

1. Introduction

Inferring information about 3D shapes through the analysis of its 2D projections has become a common paradigm in geometry processing, often referred to as *3D-from-2D*. Moving the analysis from the space to the plane has several advantages, one of the most important being a drastic reduction of computational costs. In fact, complex problems like surface reconstruction, curve-skeleton extraction, shape segmentation and many more can be efficiently solved in 2D [WGW*13, GSP12, BBP08, CK01].

Despite the many advantages, the reduction of dimensionality carries an intrinsic problem, that can be expressed in the form of a simple question: is a silhouette really representative of the 3D shape that generated it? Sometimes the shadow of an object cast on a plane is able to really capture its whole shape (see Figure 1); other times not only silhouettes have no apparent relation with the objects they were generated from, but they are even misleading, in the sense that they seem to be generated from a totally different object, as it happens for *shadowgraphs* (see inset above).



The reason why a contour may not be representative of the generating shape is that, depending on the point of view, parts of the shape may overlap along the direction of projection, i.e they are *occluded*. Occlusions cause loss of data when we examine a 3D object using its 2D projections. To compensate this lack of data most algorithms collect multiple images, observing the same shape from different points of view so as to accumulate enough information for the subsequent processing.

In this paper we target algorithms that employ a 3D-from-2D approach to compute the curve-skeleton of a 3D shape, e.g. [KJT13, LS13, LGS12]. Curve-skeletons are popular graph-like abstractions of 3D shapes widely used in many applications (see [CSM07] and references therein). These methods start with the medial axis of a number of silhouettes of a 3D shape as observed from different viewpoints, and combine them together to infer information about the curve-skeleton. The selection of reliable data to extract from each silhouette is therefore a crucial factor to ensure quality skeletons. Ranking each silhouette according to some global quality metric [SLF*11] can lead to wrong results, as the way occlusions affect the medial axis is purely local. Measuring the amount of occlusions affecting a silhouette can be misleading as well. In fact, even a tiny occlusion may have a dramatic impact on the medial axis (see Figure 3, top line). Ideally, for each medial axis point of each silhouette, the question to answer would be: is this point the projection of a curve-skeleton's point? Surprisingly enough, state of the

art algorithms spend little to none effort to answer this question.

We bridge the gap focusing our attention on the relation between the medial axis of a silhouette of a 3D shape and its curve-skeleton. Our contributions are twofold:

- we propose a novel filtering algorithm that inputs the silhouette of a 3D shape as observed from a point of view, its medial axis and a map of the occluded areas, and outputs the subset of 2D medial points over which projects the curve-skeleton of the depicted shape;
- we discuss an interesting link between our filtered medial axis and the Medial Geodesic Function [DS06]. Dey and Sun were the firsts to formally characterize curve-skeletons of 3D shapes; medial geodesic skeletons are usually considered the ground truth.

Our filter can be easily plugged into any existing pipeline for curve-skeleton extraction. It can be used to robustly rank silhouette silhouettes according to how much they are representative of the 3D shape that generated them and can also be used to perform shape recognition from occlusion-aware images or video sequences [ZL12, HY10].

2. Related Works

Computing curve-skeletons by means of 2D image analysis is a relatively new approach. In [LGS12] Livesu and colleagues pioneered the idea of using the 2D medial axis of a set of silhouettes to reconstruct the curve-skeleton of general shapes, observing that in absence of occlusions the curve-skeleton of a 3D shape projects over the medial axis of its silhouettes (Figure 1). Their method was soon extended by [KJT13] that employed a more sophisticated silhouette matching system based on a probabilistic approach. Both [KJT13] and [LGS12] somehow assume that each contour is representative enough of the input 3D shape, regardless the amount of occlusions affecting it. As this assumption is in general false, to compensate the error introduced by spurious data, noise reduction and smoothing techniques are heavily used in the final stages of the computations, affecting the centeredness of the skeleton paths. In [LS13] the input silhouettes are enriched with additional information regarding occlusions. The authors exploit this extra information to discard all the medial points that lie on (or are affected from) occluded parts of the silhouette. Being aware of the location and amount of occlusions affecting each contour this method is able to generate better quality results if compared to the previous ones. We successfully process medial axis on which this filter fails (see Section 5).

3. Medial Axis Filtering

In this section we introduce our filtering operator. Before diving into the details we fix the notation.




Figure 1: When there are no occlusions along the direction of projection the curve-skeleton of a 3D shape projects on the medial axis of its silhouette.

Let $\mathcal{S} \subset \mathbb{R}^3$ be a 3D shape, $\Omega_{\mathcal{S}} \subset \mathbb{R}^2$ its silhouette as observed from some point of view and $\mathbf{MAT}(\Omega_{\mathcal{S}})$ the medial axis transform of $\Omega_{\mathcal{S}}$ [CCM97]. Let us also suppose to have an occlusion map, in the form of a boolean function $f : \Omega_{\mathcal{S}} \rightarrow \{0, 1\}$ defined as follows:

$$f(p) = \begin{cases} 1 & \text{if } p \text{ belongs to an occluded area} \\ 0 & \text{otherwise} \end{cases}$$

A color-coded example of the function f is given in Figure 3, where light gray means $f = 0$ (i.e. no occlusion) and dark gray $f = 1$ (i.e. occluded area).

It is known from the perception theory that in absence of occlusions the medial curves of \mathcal{S} project over the medial lines of $\Omega_{\mathcal{S}}$ [Mar10]. However, if a silhouette contains occlusions this relation ceases to exist. Our goal is to implement a filter that leverages the occlusion map to retain only the medial axis points that do not depend on the occluded areas of the silhouette, that is, the ones that have a clear relation with the curve-skeleton of \mathcal{S} .

The first observation is trivial: any medial point $p \in \mathbf{MAT}(\Omega_{\mathcal{S}})$ inside an occluded area is affected by occlusions, hence is not reliable and should be discarded. However, this is not sufficient. As can be noticed in Figure 3, medial paths that do not traverse occluded areas can be still affected by occlusions. We therefore need a more powerful criterion to

understand whether a medial axis point has a relation with the curve-skeleton of the original shape or not.

The key observation is that the medial axis is a *local* operator, that is, each point $p \in \mathbf{MAT}(\Omega_S)$ depends only on a fraction of the domain Ω_S , bounded by the maximal disk $D(p)$ centered at it. In fact, p would belong to the medial axis of any other silhouette Ω'_S in which $D(p)$ is still a maximal disk, regardless how much Ω_S and Ω'_S would differ outside of $D(p)$. This is a consequence of the *Domain Decomposition Lemma* [CCM97] that states that any domain Ω can be split in sub-domains $\Omega^1, \dots, \Omega^n$ around a maximal disk $D(p)$, and the medial axis of Ω is the union of the medial axis of each sub-domain, that is, $\mathbf{MAT}(\Omega) = \bigcup_{i=1}^n \mathbf{MAT}(\Omega^i)$.

Consequently, to understand whether a medial point p has to be filtered or not we can restrict our analysis to its local neighborhood, whose size is defined by the radius of the maximal disk centered at it. In particular, for any medial point p and maximal disk $D(p)$, if none of the points in the interior of $D(p)$ are occluded, that is, if

$$\sum_{q \in D(p)} f(q) = 0 \quad (1)$$

then p is *occlusion free*, in the sense that $p \in \mathbf{MAT}(\Omega_S)$ regardless the occlusions that may affect Ω_S .

From a practical point of view this means that $D(p)$ is the projection of a maximal sphere inscribed in the 3D shape \mathcal{O} , therefore its center is **guaranteed** to belong to the 3D medial axis of \mathcal{O} and p is its projection.

Note that not all the medial axis points of a 3D shape project onto the medial axis of its shadows. In absence of overlapping this is certainly true for a medial curve, for example for the fingers of a hand, but is false for a medial sheet, for example the palm of a hand. In the palm case, only the medial points that are central with respect to the medial sheet they belong to will project over the medial axis of the shadow of the hand. This relates with the concept of Medial Geodesic Function (MGF) formalized by Dey and Sun in [DS06] and explains why p is the projection a curve-skeleton point of the 3D shape.

An algorithmic description of our filter is given in Algorithm 1. Examples of our filtering operator can be seen in Figures 1, 3 and 4.

Algorithm 1 Occlusion-Aware Medial Axis Filter

Require: a contour Ω , its medial axis $\mathbf{MAT}(\Omega)$ and its occlusion map

- 1: **for** each $p \in \mathbf{MAT}(\Omega)$ **do**
 - 2: let $D(p)$ be the maximal disk centered at p
 - 3: **for** each $q \in D(p)$ **do**
 - 4: **if** $\text{IsOccluded}(q)$ **then**
 - 5: discard p
-




Figure 2: Restricting the area of influence of our filter to the segments $\overline{pc_1}$ and $\overline{pc_2}$ may lead to wrong results. In the image above such segments do not cross any occluded area of the silhouette, nevertheless the maximal disk $D(p)$ is not the projection of a maximal sphere of the triple torus (left) nor p is a projection of a curve-skeleton point.

3.1. Local support

Let $C(p) = D(p) \cap \partial\Omega_S$ be the set of contact points of the maximal disk centered at p with the boundary of the silhouette. From the definition of medial axis $D(p)$ is at least bi-tangent to $\partial\Omega_S$, that is, it touches the boundary in at least two distinct points. To restrict more the local support of our filter one could be tempted to consider only the segments joining each contact point $c \in C(p)$ to p . However, considering only a subset of the area delimited by the maximal disk may lead to wrong results, the reason being that if $D(p)$ contains some occlusions it cannot be proven to be the projection of a maximal sphere of the original object \mathcal{O} . A clear example of this is given in Figure 3: although the paths $\overline{pc_1}$ and $\overline{pc_2}$ do not intersect any occluded area, the maximal disk $D(p)$ is clearly not the projection of a maximal sphere inscribed in the triple torus. Consequently, there is no guarantee that p is the projection of a medial point of the torus nor that it belongs to its curve-skeleton (in this case it does not).

4. Implementation and results

The implementation of our filter is very simple and can easily run in real time on images of moderate size. In our tests we always used 500×500 pixel silhouettes; this resolution proved to be more than sufficient for the application we tar-




Figure 3: Visual comparison with the filter proposed in [LS13]. On the left we show the reference 3D object we used to generate the contours; on the right we show the corresponding silhouettes and medial axis. Occluded areas are dark grey; the medial axis points are: white if they passed the filter test, red if they didn't. As can be clearly noticed from the top line: a medial axis with no apparent relation with the curve-skeleton is almost entirely validated by the filter proposed in [LS13] whereas our filter is able to detect all the medial points that are affected by occlusions and discard them.

get. As the classical medial axis is very sensitive to noise-level perturbations along the boundary we instead used the Scale Axis Transform [GMPW09]. To do so we first computed the distance transform of the contour, scaled it by a factor of 1.5, then computed the medial axis over the scaled DT. The distance transform is also used to estimate the radius of the maximal disk centered at each medial point. We used the Bresenham's algorithm to raster the outer circle and flooded the interior starting from the center to evaluate occlusions. For the computation of both the distance transform and the medial axis we relied on the Qgar C++ library [QGA] which proved to be the fastest and best accurate among the libraries we tested. For the computation of the occlusion map we counted how many primitives project over each image pixel, to do so we used the OpenGL Stencil buffer. All the filtered medial axis shown throughout the paper have been produced using this configuration.

We tested our filter on several silhouettes produced by different 3D shapes. Figure 4 showcases a number of silhouettes of the hand dataset as observed from many points of view. Next to each contour there is the reference 3D shape and relative skeleton, computed with [LGS12]. The white portions of each medial axis correspond to the paths over which projects the curve skeleton, while the red portions are the ones discarded by our filter. As can be noticed, only portions of the medial axis that overlap with the curve skeleton survive our filtering operator whereas the parts of the medial axis that are affected by occlusions and have an uncertain relation with the curve-skeleton are successfully discarded.

5. Comparisons

We compared our filter with the filtering system proposed in [LS13]. In Figure 3 we show two silhouettes of the triple torus. In both cases the curve-skeleton of the object does not project over the medial axis of the contours because of the occlusions affecting them. As can be noticed, even if there is no relation between the curve-skeleton and the medial axis the filter proposed in [LS13] validates nearly all the medial points of the silhouette, whereas our filter is able to retain all the medial points that are affected by occlusions, discarding the majority of the medial axis and recognizing the two silhouettes as poorly representative of the depicted shape.

6. Conclusions

We presented a new method to filter the medial axis of the silhouette of a 3D shape as observed from some point of view, in order to retain only the 2D medial curves over which projects the curve-skeleton of the 3D shape the silhouettes were generated from. We also discussed an interesting relation between our filter and the Medial Geodesic Function (MGF) which is the basis of the most adopted formal definition of curve-skeleton [DS06]. Finally, we showed that the local neighborhood to consider cannot be smaller than the maximal disk centered in a medial axis point (Section 3.1).

Our algorithm can be easily plugged into any existing framework for curve-skeleton extraction from 2D images. Being able to retain all the medial points that are meaningless to the curve-skeleton reconstruction problem is of great benefit for such algorithms because noise and outliers are removed a priori, thus increasing the robustness of the skeletonization process.

Counting how many points survive our filter can also be used as a metric to measure how much a silhouette is representative of the object that generated it.

In the future we plan to plug our filter into shape recognition systems, for example for surveillance applications. Given a database of curve-skeletons and given an occlusion-aware contour (computed for example on top of a RGBD picture or video [ZL12, HY10]) our future goal is to find the




Figure 4: Different silhouettes of the hand dataset captured from several points of view. Our filter is able to understand how occlusions affect the medial axis of each contour and retains only the medial paths that are the real projection of the curve-skeleton of the 3D hand.

skeleton that best matches the filtered medial axis of the contour in order to discover the 3D object being represented in the image/video.

6.1. Limitations

The medial points selected by our system are **guaranteed** to be the projection of some skeleton points of the initial 3D shape (according to the definition of curve-skeleton presented in [DS06]). Unfortunately, the opposite is not always true. Imagine for example a big vase with a small handle: if the handle is not visible from a point of view it will not influence the silhouette cast from it and, consequently, its medial axis. In scenarios like this the filter will discard the medial points affected by occlusions whether they are biasing the medial axis or not. Although from a theoretical point of view this is correct (the system is not aware of what's hidden behind an occluded area; it could be a small handle but it could also be a hole) it would be interesting to enable a deeper analysis of the occlusions, for example considering the number of times a projective ray intersects the surface of the object, in order to validate more medial points per silhouette, whenever possible.

Acknowledgements

The authors thank the anonymous reviewers for their insightful comments. This work was developed while Marco Livesu was at University of British Columbia, Vancouver, Canada.

References

- [BBP08] BELLETTINI G., BEORCHIA V., PAOLINI M.: Topological and variational properties of a model for the reconstruction of three-dimensional transparent images with self-occlusions. *Journal of Mathematical Imaging and Vision* 32, 3 (2008), 265–291. 1
- [CCM97] CHOI H. I., CHOI S. W., MOON H. P.: Mathematical theory of medial axis transform. *pacific journal of mathematics* 181, 1 (1997), 57–88. 2, 3
- [CK01] CYR C., KIMIA B.: 3D object recognition using shape similarity-based aspect graph. *Proceedings Eighth IEEE International Conference on Computer Vision. ICCV 2001 1* (2001), 254–261. 1
- [CSM07] CORNEA N. D., SILVER D., MIN P.: Curve-skeleton properties, applications, and algorithms. *IEEE Transactions on Visualization & Computer Graphics*, 3 (2007), 530–548. 1
- [DS06] DEY T. K., SUN J.: Defining and computing curve-skeletons with medial geodesic function. In *Symposium on Geometry Processing* (2006), pp. 143–152. 2, 3, 4, 5
- [GMPW09] GIESEN J., MIKLOS B., PAULY M., WORMSER C.: The scale axis transform. In *Proceedings of the twenty-fifth annual symposium on Computational geometry* (2009), ACM, pp. 106–115. 4
- [GSP12] GUGGERI F., SCATENI R., PAJAROLA R.: Shape reconstruction from raw point clouds using depth carving. In *Eurographics (Short Papers)* (2012), pp. 33–36. 1
- [HY10] HE X., YUILLE A.: Occlusion boundary detection using pseudo-depth. In *Computer Vision—ECCV 2010*. Springer, 2010, pp. 539–552. 2, 4
- [KJT13] KUSTRA J., JALBA A., TELEA A.: Probabilistic view-based 3d curve skeleton computation on the gpu. In *VISAPP* (2013), pp. 237–246. 1, 2

- [LGS12] LIVESU M., GUGGERI F., SCATENI R.: Reconstructing the curve-skeletons of 3d shapes using the visual hull. *IEEE transactions on visualization and computer graphics* 18, 11 (2012), 1891–1901. [1](#), [2](#), [4](#)
- [LS13] LIVESU M., SCATENI R.: Extracting curve-skeletons from digital shapes using occluding contours. *The Visual Computer* 29, 9 (2013), 907–916. [1](#), [2](#), [4](#)
- [Mar10] MARR D.: *Vision: a Computational Investigation into the Human Representation and Processing of Visual Information*. The MIT Press, 2010. [2](#)
- [QGA] QGAR PROJECT LORIA: Qgar. www.qgar.org/. [4](#)
- [SLF*11] SECORD A., LU J., FINKELSTEIN A., SINGH M., NEALEN A.: Perceptual models of viewpoint preference. *ACM Transactions on Graphics* 30, 5 (Oct. 2011). [1](#)
- [WGW*13] WANG Y., GONG M., WANG T., COHEN-OR D., ZHANG H., CHEN B.: Projective analysis for 3d shape segmentation. *ACM Transactions on Graphics (Special Issue of SIGGRAPH Asia)* 32, 6 (2013), Article 192. [1](#)
- [ZL12] ZHANG S.-H., LIU J.-X.: A self-occlusion detection approach based on depth image using svm. *International Journal of Advanced Robotic Systems* 9 (2012). [2](#), [4](#)