The λ-complex and surface reconstruction *

Udo Adamy, Joachim Giesen, Matthias John
Institute of Theoretical Computer Science
ETH Zürich, Switzerland

1 Abstract

We introduce the λ-complex, which combines features from α-shapes [4] and β-skeletons [6]. The λ-complex is a $(d-1)$-dimensional subcomplex of the d-dimensional Delaunay complex of a finite point set. We propose the λ-complex for surface reconstruction and discuss theoretical and experimental results.

2 Introduction

A *shape* is a geometric notion that can be assigned to finite point sets in Euclidean space. As Edelsbrunner et al. [4] point out it is difficult to give a formal meaning to the notion of *shape*. They offer α-shapes to give a concrete and formal definition that can be computed and applied.

In contrast to the general *shape* problem in the *reconstruction* problem it is known that the finite set of points stems from a concrete geometric object with concrete geometric properties, e.g. dimension or area. In the *surface reconstruction* problem the point set stems from a $(d-1)$-dimensional submanifold of $\mathbb{R}^d$. Such manifolds are called surfaces for short. The problem now asks for the *shape* that best approximates geometric quantities like area or curvatures of the original surface. The dimension need not to be reconstructed since it is known in advance. In fact we are looking for a surface complex that approximates the point set.

A $(d-1)$-dimensional surface complex is a $(d-1)$-dimensional complex K with the following properties:

(1) Every k-simplex of K, $0 \leq k < d-1$, is in the boundary of an $(k+1)$-simplex of K.

(2) Every $(d-2)$-simplex of K is in the boundary of at most two $(d-1)$-simplices of K.

(3) For every k-simplex σ of K, $0 \leq k < d-2$, let $K(\sigma)$ be the set of all $(d-1)$-simplices incident to σ. The $(d-2)$-simplices incident to simplices of $K(\sigma)$ and not incident to σ have to build a connected complex.

We briefly recapitulate two concepts to solve the *shape* or *surface reconstruction* problem for a finite point set $S \subset \mathbb{R}^d$.

We start with the definition of the α-shape of S for $\alpha \in [0, \infty]$, see [4]. For $k \leq d-1$, a k-simplex with vertices $v_0, \ldots, v_k$ belongs to the α-shape of S iff there is an empty open ball b with radius α and $\partial b \cap S = \{v_0, \ldots, v_k\}$. In the following we slightly abuse the notation and consider only $(d-1)$-simplices to belong to the α-shape.

The β-skeleton was introduced by Kirkpatrick and Radke [6]. For $\beta \geq 1$, the β-skeleton of S is the $(d-1)$-dimensional complex given by the following *forbidden region condition*. A $(d-1)$-simplex σ with vertices $v_0, \ldots, v_{d-1} \in S$ belongs to the β-skeleton of S iff the union of the two d-dimensional balls tangent to $v_0, \ldots, v_{d-1}$ with diameter $\beta \cdot \mathrm{diam}(\sigma)$ is empty. Here $\mathrm{diam}(\sigma)$ denotes the diameter of the unique $(d-1)$-dimensional ball which has $v_0, \ldots, v_{d-1}$ on its boundary.

Both concepts have their drawbacks in surface reconstruction. α-shapes cannot adopt to varying sampling density along the surface, since the global value α has to be chosen in such a way that it fits for the sparsest sampled parts of the surface. Such a value might not be adequate for densely sampled parts. In contrast to that, the β-skeleton adopts nicely to varying sampling densities, but the *forbidden region condition* is more restrictive than the condition for α-shapes. Hence many simplices one wants to have in the reconstruction might not be in the β-skeleton. This is especially the case for non-smooth surfaces.

We want to combine the advantages and, at the same

*This work was supported by grants from the Swiss Federal Office for Education and Science (Projects ESPRIT IV LTR No. 21957 CGAL and N0. 28155 GALIA).

time, avoid the disadvantages of both concepts. The λ-complex is a complex that adopts to varying sampling density but is not too restrictive.

In order to validate the suitability of the introduced complexes we count the simplices that violate conditions (2) or (3) of the surface complex.

3 Definition and Computation

Before we define the λ-*complex* we have to introduce some notations. In the following we deal with a finite point set $S \subset \mathbb{R}^d$. For a simplex σ of dimension d let $\mathrm{diam}(\sigma)$ be the diameter of the unique circumscribed ball of the same dimension. With $\Delta(S)$ we denote the set of all $(d-1)$-dimensional simplices with vertices in S. Given $\lambda \in (0, 1]$ and a simplex $\sigma \in \Delta(S)$, a λ-ball $b_\lambda(\sigma)$ is an open d-dimensional ball with diameter $\mathrm{diam}(\sigma)/\lambda$ that has all vertices of σ on its boundary. For $\lambda = 0$, we define $b_\lambda(\sigma)$ to be an open hyperplane through the vertices of σ.

Definition 3.1 *The λ-interval $I(\sigma)$ assigned to a simplex $\sigma \in \Delta(S)$ is*

$$I(\sigma) = \mathrm{closure}(\{\lambda : b_\lambda(\sigma) \cap S = \emptyset\}).$$

With this definition of λ-intervals we are able to define a λ-complex.

Definition 3.2 *The λ-complex for a given interval $I \subseteq [0, 1]$ consists of all $(d-1)$-simplices $\sigma \in \Delta(S)$ for which the interval I is contained in the assigned λ-interval, i.e. $I \subseteq I(\sigma)$.*

The λ-complex is a subset of the Delaunay complex $\mathcal{D}(S)$. Let σ be a $(d-1)$-dimensional simplex of $\mathcal{D}(S)$. Its λ-interval $I(\sigma)$ can be computed efficiently as follows. Let σ_1 and σ_2 be the two d-dimensional simplices adjacent to σ in $\mathcal{D}(S)$. If σ_i does not exist, i.e. $\sigma \in \partial\mathrm{conv}(S)$, set $\lambda_i = 0$, otherwise set

$$\lambda_i = \frac{\mathrm{diam}(\sigma)}{\mathrm{diam}(\sigma_i)}.$$

Note that in the two-dimensional case $\lambda_i = \sin\alpha_i$, where α_i is the angle in σ_i opposite to σ.

The lower bound of the λ-interval $I(\sigma)$ then equals to $\min\{\lambda_1, \lambda_2\}$. The upper bound of $I(\sigma)$ is given by 1, iff the centers of the circumscribed balls of σ_1 and σ_2 lie on different sides of the hyperplane through σ. Otherwise the upper bound equals $\max\{\lambda_1, \lambda_2\}$.

4 Curve Reconstruction

Now we consider the case $d = 2$. A one-dimensional submanifold of $\mathbb{R}^2$ is a simple (closed) curve. In the following we deal with a finite point set sampled from a simple (closed) curve in the plane. We will examine how well the curve is approximated by a λ-complex, which now is a one-dimensional complex. To do so we need some further notions, which we summarize in the following definition:

Definition 4.1 *Let $\gamma : [0, 1] \to \mathbb{R}^2$ be a curve.*

(1) Let

$$T = \{(t_1, t_2) \,:\, t_1 < t_2, t_1, t_2 \in [0, 1]\}$$

and

$$\tau : T \to \mathbb{S}^1, (t_1, t_2) \mapsto \frac{\gamma(t_2) - \gamma(t_1)}{|\gamma(t_2) - \gamma(t_1)|}.$$

The curve γ is called left (right) regular at $\gamma(t_0)$ with left tangent $\mathrm{l}(\gamma(t_0))$ or right tangent $\mathrm{r}(\gamma(t_0))$, if for every sequence (ξ_n) in T which converges from left (right) to (t_0, t_0) in $\mathrm{closure}(T)$ the sequence $\tau(\xi_n)$ converges to $\mathrm{l}(\gamma(t_0))$ or $\mathrm{r}(\gamma(t_0))$, respectively. We call γ s-regular if it is left and right regular in all points $\gamma(t), t \in [0, 1]$.

(2) A sample S of γ is a finite sequence $S = \{p^1, \ldots, p^n\}$ of points where $p^i \in \gamma$. We assume that the sample points p^i are ordered according to the order of the $\gamma^{-1}(p^i) \in [0, 1]$. For every sample S its density is defined to be the inverse of the following number

$$\varepsilon(S) = \sup_{x\in\gamma} \min\{|p^i - x| \,:\, i = 0, \ldots n\}.$$

(3) We denote by $P(S)$ the polygon that connects the sample points S in their order along the curve γ.

Now we are prepared to prove some theorems on curve reconstruction with the λ-complex.

Theorem 4.2 *Let γ be a smooth curve in $\mathbb{R}^2$ and (S_n) be a sequence of samples from γ with $\lim_{n\to\infty} \varepsilon(S_n) = 0$. We have:*

(1) For every two sequences (p_n), (q_n) with $p_n, q_n \in S_n$ and $\mathrm{conv}\{p_n, q_n\} \subset P(S_n)$ one finds:

$$\lim_{n\to\infty} \sup I(\mathrm{conv}\{p_n, q_n\}) = 1.$$

(2) For every two sequences (p_n), (q_n) with $p_n, q_n \in S_n$ and $\mathrm{conv}\{p_n, q_n\} \not\subset P(S_n)$ one finds:

$$\limsup\, \sup I(\mathrm{conv}\{p_n, q_n\}) < 1$$

or

$$\liminf\, \inf I(\mathrm{conv}\{p_n, q_n\}) > 0.$$

PROOF. The proof is left to the full version of this paper. □

The next theorem is about a larger class of curves than the smooth ones. For this class of curves the correct reconstruction is a subgraph of the Delaunay graph and coincides with the Traveling salesman path (tour) provided the points are sampled densely enough [5]. So far no filter-based algorithm is known that can provably reconstruct all curves of this class.

Theorem 4.3 *Let γ be a s-regular curve in $\mathbb{R}^2$ with*

$$\sup\{\angle(\mathrm{l}(q),\mathrm{r}(q))\ :\ q\in\gamma\}<\pi$$

and let (S_n) be a sequence of samples from γ with $\lim_{n\to\infty}\varepsilon(S_n)=0$. For every two sequences (p_n), (q_n) with $p_n,q_n\in S_n$ and $\operatorname{conv}\{p_n,q_n\}\subset P(S_n)$ one finds:

$$\lim_{n\to\infty}\inf I(\operatorname{conv}\{p_n,q_n\})=0.$$

PROOF. Assume the contrary. That is, there exist two sequences (p_n), (q_n) with $p_n,q_n\in S_n$ and $\operatorname{conv}\{p_n,q_n\}\subset P(S_n)$ such that

$$\limsup\ \inf I(\operatorname{conv}\{p_n,q_n\})>0.$$

By turning to an appropriate subsequence we can assume that

$$\lim_{n\to\infty}\inf I(\operatorname{conv}\{p_n,q_n\})=c>0.$$

Observe that the edges $\operatorname{conv}\{p_n,q_n\}$ cannot be edges of the convex hull of S_n, because for convex hull edges e we have $\inf I(e)=0$. Furthermore we can assume, using the compactness of γ and $\lim_{n\to\infty}\varepsilon(S)=0$, that

$$\lim_{n\to\infty}p_n=\lim_{n\to\infty}q_n=p\in\gamma.$$

We know from [5] that for sufficiently dense samples the correct reconstruction is a subgraph of the Delaunay graph of the sample points. Hence for large n the edge $\operatorname{conv}\{p_n,q_n\}$ is an edge of the Delaunay graph, which has two incident triangles with vertices in S_n. Let r_n and s_n be the vertices opposite to the edge $\operatorname{conv}\{p_n,q_n\}$ in these triangles.

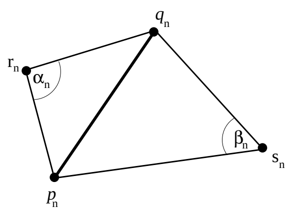

From our assumptions we find for the sequence (α_n) of angles α_n at vertex r_n:

$$\limsup\ \sin\alpha_n\geq c$$

Analogously we find for the sequence (β_n) of angles β_n at vertex s_n:

$$\limsup\ \sin\beta_n\geq c$$

By turning to appropriate subsequences we can assume that

$$(1)\qquad \lim_{n\to\infty}\sin\alpha_n\geq c\ \text{ and }\ \lim_{n\to\infty}\sin\beta_n\geq c$$

We find for the diameters $\operatorname{diam}(p_n,q_n,r_n)$ and $\operatorname{diam}(p_n,q_n,s_n)$ of the unique balls through the points p_n,q_n and r_n or p_n,q_n and s_n, respectively:

$$\begin{aligned}\lim_{n\to\infty}\operatorname{diam}(p_n,q_n,r_n)&=\lim_{n\to\infty}\frac{|p_n-q_n|}{\sin\alpha_n}=0,\\ \lim_{n\to\infty}\operatorname{diam}(p_n,q_n,s_n)&=\lim_{n\to\infty}\frac{|p_n-q_n|}{\sin\beta_n}=0\end{aligned}$$

That is,

$$\lim_{n\to\infty}r_n=\lim_{n\to\infty}s_n=p.$$

Assume that in p left- and right tangent coincide. That is, in p exists a unique tangent. Without loss of generality we can assume that $p_n<q_n<r_n$ for all $n\in\mathbb{N}$ in the order along γ locally around p. Consider the triangle with vertices p_n,q_n and r_n. From the definition of regularity the sequence of angles (α_n) converges to zero, which is a contradiction to inequalities (1). Hence it suffices to consider the case that in p left- and right tangent do not coincide.

By turning to appropriate subsequences one can consider without loss of generality two cases:

(1) $p_n<p<q_n$ along γ for all $n\in\mathbb{N}$.

(2) $p_n,q_n\leq p$ or $p\leq p_n,q_n$ along γ for all $n\in\mathbb{N}$.

First case: Let (h_n) be the sequence of hyperplanes through the points p_n and q_n. Since the sequences (r_n) and (s_n) both converge to p we find for sufficiently large $n\in\mathbb{N}$ that the points r_n and s_n lie both in the half space with boundary h_n that contains the point p. But this is impossible in a Delaunay triangulation.

Second case: Without loss of generality we assume that $p\leq p_n,q_n$ along γ for all $n\in\mathbb{N}$.

Now assume that $p\leq r_n$ or $p\leq s_n$ along γ for arbitrary large n. From the existence of the right tangent in p we find that the sequence of angles (α_n) or the sequence (β_n) has a subsequence that has to converge to zero. That is a contradiction to inequalities (1).

We are left with the case that there exists $N\in\mathbb{N}$ such that $r_n,s_n<p$ for all $n\geq N$. Consider the sequence (h_n) of hyperplanes through the points p_n and q_n. This sequence converges to the hyperplane defined by the right tangent at p. Since $\angle(\mathrm{l}(p),\mathrm{r}(p))<\pi$ and the sequences (r_n) and (s_n) both converge to p we find for large n that the points r_n and s_n have to lie both on the same side of the hyperplane h_n. Again, this is impossible in a Delaunay triangulation. □

Theorems 4.2 and 4.3 show that the λ-complex for large intervals $I \subset [0,1]$ is a reconstruction scheme for smooth curves. For the larger class of curves defined by the prerequisites of Theorem 4.3 the λ-complex for appropriate intervals $I \subset [0,1]$ is a conservative filter. That is, the λ-complex contains all edges that have to be in the reconstruction, but possibly contains some further edges. Our experience shows that for almost all curves of the latter class there is a λ-complex that gives the correct reconstruction. Hence, the λ-complex can be used as a preprocessing step for optimization techniques like in [1] for curve and especially surface reconstruction.

5 Experimental results

In Figures 1-3 several λ-complexes are shown. For their computation we used the C++ library CGAL [3].

Figure 1 (middle) gives the correct curve reconstruction for a smooth curve in the plane with a large interval (see Theorems 4.2 and 4.3). Smaller intervals select too many edges (Figure 1 (left and right)). Note that the intersection of the left and right λ-complexes gives the correct reconstruction.

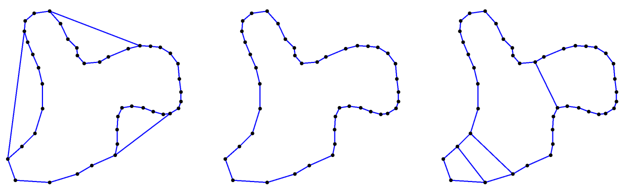

(a) $[0.3, 0.4]$ (b) $[0.3, 1.0]$ (c) $[0.9, 1.0]$

Figure 1: λ-complexes of a sample of a smooth curve

Figure 2 compares the λ-complex with the Crust algorithm [2] for a s-regular curve. While the λ-complex reconstructs the curve correctly, the Crust algorithm fails at non-regular points of the curve.

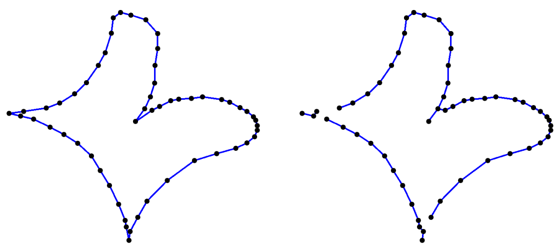

(a) $I = [0.3, 0.7]$ (b) Crust algorithm

Figure 2: Comparison between λ-complex and Crust algorithm for a sample of a s-regular curve

Figure 3 shows two λ-complexes of point sets in $\mathbb{R}^3$. While green areas violate the surface complex conditions (2) or (3), grey areas belong to the correctly reconstructed parts of the original surfaces. Details like engravings on the club are recognizable. Most of the violating facets can be removed by a simple topological postprocessing step.

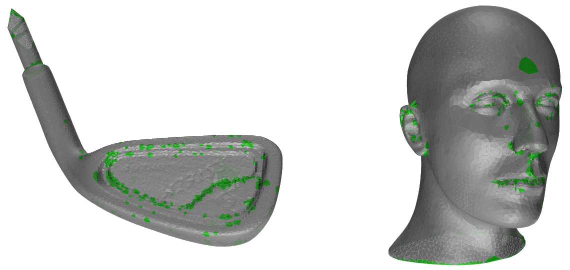

(a) club (16864 points) (b) mannequin (12772 points)

Figure 3: λ-complexes with interval $[0.9, 1.0]$

Acknowledgment We would like to thank our advisor Emo Welzl for helpful discussions.

References

[1] E. Althaus, K. Mehlhorn, *TSP-Based Curve Reconstruction in Polynomial Time*, Proc. 11th SIAM Sympos. Discrete Algorithms (2000)

[2] N. Amenta, M. Bern, D. Eppstein, *The Crust and the β-Skeleton: Combinatorial Curve Reconstruction*, Graphical Models and Image Processing 60/2:2, pp. 125–135 (1998)

[3] `http://www.cs.uu.nl/CGAL/`

[4] H. Edelsbrunner, E.P. Mücke, *Three Dimensional Alpha Shapes*, ACM Transactions on Graphics 13(1), pp. 43–72 (1994)

[5] J. Giesen, *Regularity Considerations in Curve Reconstruction*, submitted to Theoret. Comp. Science (1999)

[6] D.G. Kirkpatrick, J.D. Radke, *A framework for computational morphology*, in Computational Geometry , Elsevier North Holland, pp. 234–244 (1985)