Learning a 2-Manifold with a Boundary in $\mathbb{R}^3$

Christian Scheffer and Jan Vahrenhold
Faculty of Computer Science, Technische Universität Dortmund, 44227 Dortmund, Germany
{christian.scheffer,jan.vahrenhold}@cs.tu-dortmund.de

Abstract

We present an algorithm for approximating the local feature size of point samples used for reconstructing a 2-manifold in $\mathbb{R}^3$. Our algorithm improves previous results by simultaneously achieving the following two goals: it computes an approximation of the local feature size as well as the number of sample points needed for this approximation using local information only and at the same time is able to deal with 2-manifolds with a smooth boundary.

1 Introduction

Reconstructing a manifold from point samples in $\mathbb{R}^3$ is a fundamental problem that has attracted considerable interest both in Computer Graphics and Computational Geometry. Most algorithms follow the general framework introduced by Amenta and Bern [2] and perform the reconstruction based upon the Voronoi diagram for the sample points. Computing a Voronoi diagram in $\mathbb{R}^3$, however, has a quadratic worst-case complexity, and thus Funke and Ramos [7] present a careful analysis of their variant of the CoCone algorithm to show that a reconstruction is possible in near-linear time. Building upon this, Dumitriu *et al.* [5, 6] propose to first decimate the point set in a preprocessing step using information locally available before computing the (graph) Voronoi diagram for the reduced point set; this algorithm does not rely on CoCone. Recently, Dey *et al.* [4] considered the case of reconstructing a 2-manifold in $\mathbb{R}^3$ that has a smooth boundary.

We demonstrate how to modify the algorithm by Dumitriu *et al.* [5, 6] to be able to reconstruct a 2-manifold in $\mathbb{R}^3$ that has a smooth boundary. As it turns out, we only need to replace one single subroutine, namely the approximation of the so-called *local feature size*, by a variant that can handle (smooth) boundaries. Our algorithm builds upon ideas presented by Funke and Ramos [7] and we derive sufficient conditions for the correctness of the reconstruction and comment on the approximation quality of our algorithm. Unlike the algorithm of Dey *et al.* [4], our algorithm does not rely on a global sampling condition with a sampling constant known to the algorithm but retains the property that the density of the sampling may be locally different.

2 Preliminaries

Amenta and Bern [2] introduced the *local feature size* as central concept used for reconstructing smooth (closed) surfaces. For any point x on the manifold Γ, the local feature size $\mathit{lfs}\,(x)$ is defined as the distance of x to the medial axis of Γ. In different terms, the local feature size in a point $x \in \Gamma$ is the radius of the smaller of the two largest balls touching Γ in x from the inside resp. from the outside and not containing any other point of Γ in their interior. Thus, the local feature size captures the curvature and the folding of Γ. It is known that the (topological) correctness of a reconstruction algorithm depends on the density of the set $S \subset \Gamma$ of sample points used for the reconstruction relative to the local feature size.

Definition 1 *A discrete subset S of a smooth 2-manifold $\Gamma \subset \mathbb{R}^3$ is an ε-sample for Γ iff for every point $x \in \Gamma$ there is a point $s \in S$ with $|xs| \leq \varepsilon \cdot \mathit{lfs}\,(x)$.*

Dumitriu *et al.*'s algorithm [5, 6] (Algorithm 1) first uses a subroutine from Funke and Ramos [7] to compute an approximation of $\varepsilon \cdot \mathit{lfs}\,(s)$ for each $s \in S$ and then connects two sample points $s_1, s_2 \in S$ by an edge if and only if $|s_1 s_2| \leq c \cdot \varepsilon \cdot \max\{\mathit{lfs}\,(s_1), \mathit{lfs}\,(s_2)\}$ where c is a constant used for compensating for the approximation error. Using the resulting neighborhood graph, the algorithm computes a k-hop-stable subsample, locally identifies the connectivity of the induced subgraph, and computes a topologically correct reconstruction of the manifold. Finally, all sample points not in the subsample are reinserted to refine the reconstruction.

Figure 1(a) shows a situation where the presence of the boundary of Γ (drawn as a thick line) results in an illegal edge $\overline{s_1 s_2}$ to be constructed. Even a higher sampling density close to the boundary does not alleviate this problem–see Figure 1(b).

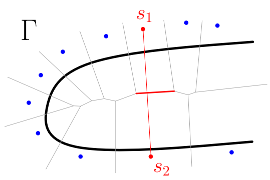

(a) Illegal edge $\overline{s_1 s_2}$.

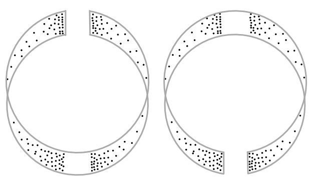

(b) Non-unique reconstruction.

Figure 1: Problems induced by boundaries.

Algorithm 1 Learning a 2-manifold in $\mathbb{R}^3$ [5, 6].

1: Compute an approximation $\phi(s)$ for $\varepsilon \cdot \mathit{lfs}\,(s)$ for $s \in S$ using a routine by Funke and Ramos [7].
2: Construct a neighborhood graph on S by connecting $s_1, s_2 \in S$ if $|s_1 s_2| \in O(\max\{\phi(s_1), \phi(s_2)\})$.
3: Compute $S_{\mathrm{sub}} \subseteq S$ as a maximal k-hop stable set in the neighborhood graph.
4: Construct the graph Voronoi diagram of S with respect to S_{sub}.
5: Use the graph Voronoi diagram to identify certified adjacencies between points in S_{sub}.
6: Identify faces and triangulate non-triangular faces.
7: Reinsert points from $S \setminus S_{\mathrm{sub}}$ using a routine by Funke and Ramos [7].

One of the key observations by Dumitriu *et al.* [5], however, is that Steps 1 and 2 guarantee that one can reconstruct a manifold from an ε-sample using only information locally available during the remainder of the algorithm (Steps 3–7). The contribution of this work is a replacement for Step 1, i.e. an approximation algorithm for the local feature size that uses only information locally available and can handle the case where the manifold to be reconstructed has a smooth boundary. This, however, requires the sampling condition to be modified.

3 Local feature size for 2-manifolds with boundary

In the remainder of this paper, we will assume that Γ is a smooth 2-manifold in $\mathbb{R}^3$ and that $\partial\Gamma$ is a smooth 1-manifold. We extend the definition of the local feature size to the case of 2-manifolds with a boundary.

Definition 2 *For the medial axis $\mathcal{M}_{\partial\Gamma}$ of the boundary $\partial\Gamma$ and the medial axis $\mathcal{M}_{\Gamma^\circ}$ of the interior Γ° we define*

$$\begin{array}{lcll} \mathit{lfs}_{\partial\Gamma}(x) & = & \min\{\mathrm{dist}(x, \mathcal{M}_{\partial\Gamma})\} & \text{for } x \in \partial\Gamma \\ \mathit{lfs}_{\partial\Gamma}(x) & = & \min_{y\in\partial\Gamma}\{\mathit{lfs}_{\partial\Gamma}(y) + |xy|\} & \text{for } x \in \Gamma^\circ \\ \mathit{lfs}_{\Gamma^\circ}(x) & = & \min\{\mathrm{dist}(x, \mathcal{M}_{\Gamma^\circ})\} & \text{for } x \in \Gamma \end{array}$$

The local feature size *of a point $x \in \Gamma$ then is defined as $\mathit{lfs}\,(x) := \min\{\mathit{lfs}_{\partial\Gamma}(x), \mathit{lfs}_{\Gamma^\circ}(x)\}$.*

Dey *et al.* also define a (similar) variant of the local feature size but then use it to define "a global measure $\rho = \inf_{x\in\Gamma} \mathit{lfs}\,(x)$" [4, p. 1373] that is know to the algorithm and used for defining the (global) quality of the sample. In contrast, we follow Funke and Ramos [7] and approximate the local feature size locally. The correctness of the algorithms of Funke and Ramos [7] and Dumitriu *et al.* [5, 6] is based upon the observation that the local feature size is a *1-Lipschitz function*, i.e., a non-negative function f with $f(x) \leq f(x') + |xx'|$ for all x, x'. Lemma 1 allows us to reuse their (algorithm and) correctness proof for points "far away" from $\partial\Gamma$:

Lemma 1 *The local feature size as defined in Definition 2 is a 1-Lipschitz function.*

Following Funke and Ramos' approach, we consider the Voronoi cell $\mathit{Vor}\,(s)$ of a sample point s restricted to the (unknown) manifold Γ. One then can show that the distance from s to the furthest vertex v in $\mathit{Vor}\,(s) \cap \Gamma$ is an approximate lower bound for $\frac{\varepsilon}{1-\varepsilon} \cdot \mathit{lfs}\,(s)$. Since Γ is unknown to the algorithm, we appoximate $\mathit{Vor}\,(s) \cap \Gamma$ by $\mathit{Vor}\,(s) \cap \tilde{T}_s$ where $\tilde{T}_s$ is an approximation of the plane T_s tangent to Γ in s. Since Γ has a boundary, $\partial\Gamma$ can cross $\mathit{Vor}\,(s)$ such that some Voronoi vertices in $\mathit{Vor}\,(s) \cap \tilde{T}_s$ do not have a corresponding vertex *on* Γ—see Figure 2.

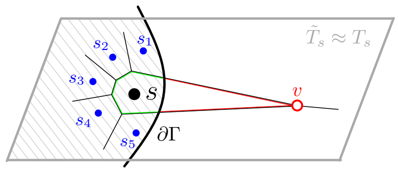

Figure 2: $\partial\Gamma$ crosses $\mathit{Vor}\,(s)$.

To avoid the lower bound for $\frac{\varepsilon}{1-\varepsilon} \cdot \mathit{lfs}\,(s)$ to be (arbitrarily) biased by such a cut-off vertex v, the main challenge is to restrict the "search space" for the Voronoi vertices on $\mathit{Vor}\,(s) \cap \tilde{T}_s$ to areas certified not to contain any part of the boundary or – if it crosses the cell – to exclude cut-off Voronoi vertices while computing the approximation.

4 Safeguarding Voronoi cells from $\partial\Gamma$

To simplify notation, we define what we call a *bicone*; bicones will be shown to contain the locally relevant part of the manifold (including the boundary).

Definition 3 *Let $x_1, x_2 \in \mathbb{R}$, $x_1 \neq x_2$, $r \geq |x_1 x_2|/2$, and let E be a plane containing both x_1 and x_2. Let $C_{E,1}$ and $C_{E,2}$ denote the two circles in E that have radius r and pass through x_1 and x_2. The bicone $C^{cone}_{x_1,x_2,r}$ of x_1 and x_2 with parameter r then is defined as the union of the intersections of $C_{E,1}$ and $C_{E,2}$ for all planes E containing x_1 and x_2 (see Figure 3):*

$$C^{cone}_{x_1,x_2,r} := \bigcup_{E \ni x_1, x_2} (C_{E,1} \cap C_{E,2})$$

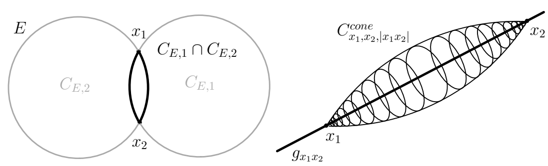

Figure 3: Bicone $C^{cone}_{x_1,x_2,r}$ of x_1 and x_2 w.r.t. r.

In Lemma 5, we name the conditions that guarantee the existence of a smooth curve connecting two points x_1 and x_2 inside a bicone $C^{cone}_{x_1,x_2,r}$ for a properly defined r. For Corollary 2 and Lemmas 3 and 4, we consider the three points $x, x_1, x_2 \in \Gamma$ where $|xx_1|, |xx_2|, |x_1x_2| \leq \frac{1}{2} \cdot \mathit{lfs}(x)$ and where x_1 and x_2 realize an angle of at most $\frac{1}{3} \cdot \pi$ at apex x. Based upon these points, we define $r := \max\{|xx_1|, |xx_2|\}$.

Corollary 2 *There exist three continuous curves $\gamma_{xx_1}, \gamma_{xx_2}, \gamma_{x_1x_2} \subset \Gamma$ connecting x and x_1, x and x_2, and x_1 and x_2 with $\gamma_{xx_1} \subset C^{cone}_{x,x_1,r}$, $\gamma_{xx_2} \subset C^{cone}_{x,x_2,r}$ and $\gamma_{x_1x_2} \subset C^{cone}_{x_1,x_2,r}$.*

Lemma 3 (see Figure 4(a)) *The Voronoi region of x with respect to x_1 and x_2 lies outside of $C^{cone}_{x_1,x_2,r}$.*

Lemma 4 (see Figure 4(b)) *The patch of Γ bounded by $\gamma := \gamma_{xx_1}\gamma_{x_1x_2}\gamma_{x_2x}$ does not contain any point of $\partial\Gamma$ in its interior.*

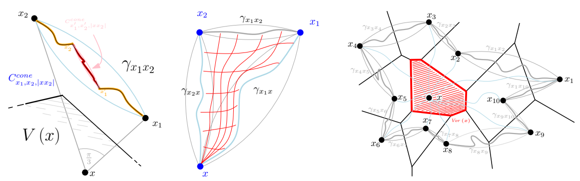

(a) $\gamma_{x_1x_2}$ does not cross $\mathit{Vor}\,(x)$. (b) Patch bounded by $\gamma_{xx_1}\gamma_{x_1x_2}\gamma_{x_2x}$. (c) Closed curve γ does not cross $\mathit{Vor}\,(s)$.

Figure 4: Configurations of Lemmas 3, 4, and 5

Combining the above results, we can show that the Voronoi cell of some (sample) point $x \in \Gamma$ restricted to the intersection with Γ is guaranteed not to contain any point from $\partial\Gamma$ if there are enough (sample) points distributed around it. More precisely, the guarantee holds if every wedge of a ball with radius $\frac{1}{2} \cdot \mathit{lfs}(x)$ centered at x contains at least one other sample point.

Lemma 5 (See Figure 4(c)) *Let $x, x_0, \ldots, x_{\ell-1}$ be points on Γ for which the following holds:*

- $|xx_i| < \frac{1}{2} \cdot \mathit{lfs}\,(x)$ *for all* $i \in \{0, \ldots, \ell-1\}$.
- $\forall i \in \{0, \ldots, \ell-1\} : \exists j, k \in \{0, \ldots, \ell-1\} \setminus \{i\} : \angle(x_j, x, x_i) \leq \frac{1}{3} \cdot \pi \;\wedge\; \angle(x_k, x, x_i) \geq \frac{5}{3} \cdot \pi$

Then there is no point of $\partial\Gamma$ inside the intersection of Γ and the Voronoi cell of x, i.e.

$$(\Gamma \cap \mathit{Vor}(x))^{\circ} \cap \partial\Gamma = \emptyset$$

5 Sampling 2-manifolds with a boundary

In Lemma 5, we considered a situation, which allows us to account the (restricted) intersection of $\mathit{Vor}\,(s)$ with T_s for approximating $\varepsilon \cdot \mathit{lfs}(s)$ by the largest distance of a point $v \in \mathit{Vor}\,(s) \cap T_s$ to s. This configuration of points describe the distribution of "close" sample points around s. The closeness of these sample points is guaranteed by the following theorem:

Theorem 6 *Let S be an ε-sample of Γ. Then, for each $s \in S$, we have at least $k_{\frac{1}{2},\varepsilon} \in \mathcal{O}(1/\varepsilon^2)$ nearest neighbors $\{s_1, \ldots, s_{k_{\frac{1}{2},\varepsilon}}\} \subset S$ such that $\operatorname{dist}(s, s_i) < \frac{1}{2} \cdot \mathit{lfs}(s)$ for each $s_i \in \{s_1, \ldots, s_{k_{\frac{1}{2},\varepsilon}}\}$.*

The final property of the sample points has to describe a "good" distribution of these $k_{\frac{1}{2},\varepsilon} \in \mathcal{O}(1/\varepsilon^2)$ nearest neighbors around s. Figure 1(b) illustrates that the standard sampling condition does not guarantee a unique reconstruction.

Definition 4 *A discrete subset $S \subset \Gamma$ is an (ε, k)-sample iff S is an ε-sample and the following condition holds for each $s \in S$: for any point $y \in \Gamma$ with $|ys| \in [\frac{\varepsilon}{1+\varepsilon} \cdot \mathit{lfs}(s), \frac{2\cdot\varepsilon}{1-\varepsilon} \cdot \mathit{lfs}(s)]$, there exists a point $s' \in S$ who is one of the k nearest neighbors of s and for which $\angle(y, s, s') \leq \frac{1}{6} \cdot \pi$ hold.*

For an (ε, k)-sample, we can prove that the premise of Lemma 5 holds for every sample point which has at least a distance of $\frac{\varepsilon}{1-\varepsilon} \cdot \mathit{lfs}(s)$ to $\partial\Gamma$. Also, we can compute a restriction H_s of $\mathit{Vor}\,(s) \cap T_s$ for all other sample points by analyzing the $k \leq k_{\frac{1}{2},\varepsilon}$ nearest neighbors of s, so that $\partial\Gamma$ do not intersect $H_s \cap \mathit{Vor}\,(s)$.

As mentioned in Section 3, the algorithm relies on an approximation $\tilde{T}_s$ of the plane T_s tangent to a sample point $s \in \Gamma$. Whereas the Funke and Ramos' base algorithm works with every (in an asymptotic sense) good approximation, our variant requires an absolute bound on the deviation of the planes' normals:

Lemma 7 *Let S be an (ε, k)-sample with $\varepsilon \leq \frac{1}{16}$ and $k \leq k_{\frac{1}{2},\varepsilon}$. We can compute an approximation $\tilde{T}_s$ of the plane T_s tangent to Γ in $s \in S$ with $\angle(T_s, \tilde{T}_s) \leq 17°$ by inspecting k nearest neighbours.*

6 Approximating the local feature size

Summarizing the results from Section 4, we obtain the following algorithm for approximating the local feature size for an (ε, k)-sample of a 2-manifold with a boundary (Algorithm 2).

Substituting this algorithm for Step 1 in Algorithm 1 and observing that the remaining steps work on information locally available in the neighborhood graph, we obtain an algorithm for reconstructing a smooth (orientable or non-orientable) 2-manifold in $\mathbb{R}^3$ with a boundary.

Algorithm 2 Approximating *lfs*(s) for $s \in S$.

Require: S $((\varepsilon, k)$-sample with $\varepsilon \le \frac{1}{16}$ and $k \le k_{\frac{1}{2},\varepsilon})$
1: **for** each $s \in S$ **do**
2: NH_s :=NN$(S,s,\ k)$. ▷ Find k NNs of s in S.
3: $\tilde{T}_s$:=TangPlane$(NH_s,\ s)$. ▷ See Lemma 7.
4: $Vor_{\tilde{T}_s}(s) := Vor(s) \cap \tilde{T}_s$.
5: **if** s does not fulfill Lemma 5's premise **then**
6: $Vor_{\tilde{T}_s}(s) := Vor_{\tilde{T}_s}(s) \cap H_s$.
7: v := Voronoi vertex of $Vor_{\tilde{T}_s}(s)$ furthest to s.
8: $\phi(s) := |vs|$.
9: Return ϕ.

7 Analysis of the algorithm

A direct implementation of Algorithm 2 results in a quadratic running time due to the complexity of computing the Voronoi diagram in $\mathbb{R}^3$. For this version, however, we can prove the following:

Theorem 8 *Algorithm 2 computes* $\phi(s)$ *for each* $s \in S$ *such that* $\phi(s) \le 1.135 \cdot \frac{\varepsilon}{1-\varepsilon} \cdot lfs(s)$.

If one is willing to trade approximation quality for running time, one can (prove and) use that the Voronoi cell of a sample point s with respect to its k nearest neighbors is a "good" approximation for the Voronoi cell used in Algorithm 2 in the sense that the constant in Theorem 8 increases from 1.135 to no more than 1.3. In this case, the dominating step of this algorithm is the computation of the k nearest neighbors, which (using a well-separated pair decomposition [3] and Theorem 6) can be done in $\mathcal{O}(\frac{n}{\varepsilon^2} \log \frac{n}{\varepsilon^2})$ time. With the exception of Step 2 (computing the neighborhood graph), the remaining steps of Algorithm 1 can be executed in the same time complexity, since the subgraph worked with is not too large [6] and re-inserting the excluded points can be done in near-linear time [7]. Step 2 boils down to answering n spherical range queries in $\mathbb{R}^3$ (i.e., n halfspace reporting queries in $\mathbb{R}^4$) which can be done in $\mathcal{O}(n^{4/3} \operatorname{polylog} n)$ time [1, 8].

8 Experimental Evaluation

Since we changed only the first step of Algorithm 1, we were interested in its output, i.e. the approximation of the local feature size. As it is common practice, we first assess the quality by visual inspection: Figure 5 shows a reconstruction of part of the dragon model used in [6] (the graph shows the union of all Delaunay duals for the (restricted) Voronoi cells from Step 7 of Algorithm 2.). We obtained a manifold with a boundary by extracting a slice of the dragon's head.

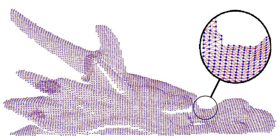

Figure 5: Dragon's head.

In addition to this visual assessment, we also evaluated the approximation quality of our algorithm for estimating the local feature size. To this end, we constructed samples S° and S^S of two manifolds Γ° and Γ^S in $\mathbb{R}^3$ in a way that we were able to derive the local feature size of the manifold from the construction algorithm and thus were able to also numerically assess the quality of our approximation algorithm. By construction, the correct values of ϕ° (for S°) and thus also of ϕ^S (for S^S) can be shown to be approximately 1.85. Our approximation of $\varepsilon \cdot lfs$ computes in both cases values in the range $[1.83, 1.87]$, that is very close to the actual value.

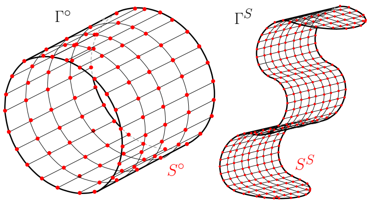

Figure 6: Manifolds and samples.

9 Conclusions

We have presented a new sampling condition that allows for approximating the local feature size of the points on a 2-manifold with a boundary embedded in $\mathbb{R}^3$. The resulting approximation algorithm builds upon previous work but has the advantage of simultaneously being able to handle manifolds with boundaries and not having to assume a global lower bound on the local feature size. The experimental evaluation tightly supports the theoretical analysis of the approximation quality.

References

[1] P. K. Agarwal and J. Erickson. Geometric range searching and its relatives. In *Advances in Discrete and Computational Geometry*, pages 1–56. AMS, 1999.

[2] N. Amenta and M. Bern. Surface reconstruction by Voronoi filtering. *Discrete & Computational Geometry*, 22(4):481–504, 1999.

[3] P. B. Callahan and S. R. Kosaraju. A decomposition of multidimensional point sets with applications to k-nearest-neighbors and n-body potential fields. *Journal of the ACM*, 42(1):67–90, 1995.

[4] T. K. Dey, K. Li, E. A. Ramos, and R. Wenger. Isotopic reconstruction of surfaces with boundaries. *Computer Graphics Forum*, 28(5):1371–1382, 2009.

[5] D. Dumitriu, S. Funke, M. Kutz, and N. Milosavljević. On the locality of extracting a 2-manifold in $\mathbb{R}^3$. In *Proc. Scandinavian Workshop on Algorithm Theory*, LNCS 5124, pages 270–281, Springer, 2008.

[6] D. Dumitriu, S. Funke, M. Kutz, and N. Milosavljević. How much geometry it takes to reconstruct a 2-manifold in $\mathbb{R}^3$. *ACM Journal on Experimental Algorithmics*, 14, May 2009. Article 2.2, 17 pages.

[7] S. Funke and E. A. Ramos. Smooth-surface reconstruction in near-linear time. In *Proc. Symposium on Discrete Algorithms*, pages 781–790. ACM, 2002.

[8] J. Matoušek. Reporting points in halfspaces. *Computational Geometry: Theory and Applications*, 2(3):169–186, 1992.