

Bounds on the k -Neighborhood for Locally Uniformly Sampled Surfaces

Mattias Andersson¹

Joachim Giesen^{2†}

Mark Pauly³

Bettina Speckmann⁴

¹ Department of Computer Science, Lund Institute of Technology, S-22100 Lund, Sweden

² Department of Computer Science, ETH Zurich, CH-8092 Zurich, Switzerland

³ Computer Science Department, Stanford University, Stanford CA 94305, USA

⁴ Department of Mathematics and Computer Science, TU Eindhoven, 5600 MB Eindhoven, The Netherlands

Abstract

Given a locally uniform sample set P of a smooth surface S . We derive upper and lower bounds on the number k of nearest neighbors of a sample point p that have to be chosen from P such that this neighborhood contains all restricted Delaunay neighbors of p . In contrast to the trivial lower bound, the upper bound indicates that a sampling condition that is used in many computational geometry proofs is quite reasonable from a practical point of view.

Categories and Subject Descriptors (according to ACM CCS): I.3.5 [Computer Graphics]: Curve, surface, solid, and object representations

1. Introduction

Point primitives have recently become popular as a means of surface representation in computer graphics and geometric modeling. With an abundance of 3D acquisition and sampling methods that create point samples from surfaces, numerous algorithms for direct processing and rendering of these data sets have been proposed. These methods typically exploit the structural simplicity of point clouds for compact storage [KV03], fast re-sampling [PGK02], and efficient rendering [RL00, BWK02].

In the most general case, 3D acquisition or sampling methods create a finite set of points that only sample the 3D position of the underlying surface. However, subsequent processing or rendering algorithms [ZPvG01, KV01, PKKG03] require additional geometric information on the surface, such as surface normals or local curvatures, which have to be computed from the positional data. To obtain such local geometric information from the point cloud itself, we

need to define a neighborhood relation for every sample point. There are two objectives in obtaining such a neighborhood. On the one hand, the computational overhead to compute it should be as small as possible, to allow efficient processing of large data sets. On the other hand, the neighborhood should be such that local geometric information can be approximated provably well. Both objectives are contradictory to a certain extent.

Two neighborhood definitions gained popularity in different communities, putting different emphasis on the two objectives. The first definition considers for every sample point its k -nearest neighbors where k is a parameter that has to be adjusted to the point sample and application [MN03]. k -nearest neighbors are very popular in the graphics community, since they are efficient to compute or approximate [AMN*98] and have proven to be sufficiently reliable for estimating local surface properties for uniformly and densely sampled models [PKKG03]. The second neighborhood relation considers for every sample point its restricted Delaunay neighbors. Restricted Delaunay neighbors are not directly accessible, since they are defined via the unknown surface from which the point sample origi-

[†] Partially supported by the Swiss National Science Foundation under the project “Non-linear manifold learning”.

nates. But they are always a subset of the Delaunay neighbors, i.e., samples connected by a Delaunay edge, which are computable. Even for locally non-uniform but sufficiently dense sample sets it is possible to approximately filter the restricted Delaunay neighbors from the Delaunay neighbors. The restricted Delaunay neighborhood is popular in the computational geometry community, since it is possible to prove many geometric and topological results about it under certain sampling conditions [DGGZ02]. However, since the Delaunay triangulation is a global data structure, the restricted Delaunay neighborhood is not as efficiently computable as the *k*-nearest neighbors.

It is easy to construct examples where estimating local surface properties using the *k*-nearest neighbors will completely fail, whereas the restricted Delaunay neighborhood allows to estimate these properties faithfully. Though examples where the restricted Delaunay neighborhood fails are also easily constructed, the observation does not hold vice versa, i.e., if one succeeds using the *k*-nearest neighbors, one should also succeed using the restricted Delaunay neighborhood. The common characteristic where the *k*-nearest neighbors are not sufficient, but the restricted Delaunay neighbors are, is that the sample set is not locally uniform. In practice, however, most acquired point sets encountered in graphics applications are locally uniform. A result from Funke and Ramos [FR02] even shows that it is always possible to get a dense, locally uniform sample from a dense sample by removing points that are not essential for the description of the surface. Thus it should be possible to approximate the restricted Delaunay neighborhood from the *k*-nearest neighbors for a locally uniform sample and the right value for *k*.

In this paper we give upper and lower bounds on the size of *k* in order to well approximate the restricted Delaunay neighborhood of a sample point provided the point sample is locally uniform. This allows efficient approximation of the restricted Delaunay neighbors without requiring a global structure such as the Delaunay triangulation. The bounds provide some guideline how to choose the value of *k* in theory. But more importantly it allows to test from a practical perspective the reasonability of sampling assumptions typically made in computational geometry proofs. If the upper bound on *k* that we derive under such a sampling condition is far off from what is needed in practice, then also the sampling condition is likely to be far off. It turns out that this is not the case.

The paper is organized as follows: In the next section we provide necessary definitions on samplings and neighborhoods. In the third section we prove our upper and lower bounds.

2. Definitions

In the following *S* always denotes a smooth surface with or without boundary embedded in \mathbb{R}^3 and $P \subset S$ denotes a finite point sample from *S*. In this section we are going to define structures derived from *S* and *P*, respectively.

Medial axis. A ball in \mathbb{R}^3 is called empty, if it does not contain any point from *S* in its interior. An empty ball is called maximal, if it is not contained in another empty ball. The medial axis $M(S)$ of *S* is the set of all center points of maximal empty balls in \mathbb{R}^3 . Maximal empty balls touch *S* in at least two points tangentially.

Local feature size. The local feature size is a function $f : S \rightarrow \mathbb{R}$ that assigns to every point in *S* its least distance to the medial axis of *S*, i.e., $f(x) = \min_{y \in M(S)} \|x - y\|$. In [AB98] it is shown that the feature size is 1-Lipschitz continuous, i.e., $f(x) \leq f(y) + \|x - y\|$.

ε -sample. The point sample *P* is called an ε -sample of *S*, if every point $x \in S$ has a point in *P* at distance at most $\varepsilon f(x)$.

In [AB98] it is shown that for an ε -sample with $\varepsilon \leq 0.08$ it is always possible to compute a triangle mesh from *P* that is homeomorphic to *S*. The algorithm that does so is based on the Delaunay triangulation of *P*. The proof of homeomorphy makes use of a result obtained by Edelsbrunner and Shah [ES94] that states that the Delaunay triangulation of *P* restricted to *S* is homeomorphic to *S* provided its dual, the restricted Voronoi diagram, fulfills the so-called closed ball property. We are now going to define these concepts.

Voronoi diagram. The Voronoi diagram $V(P)$ of *P* is a cell decomposition of \mathbb{R}^3 in convex polytopes. Every *Voronoi cell* corresponds to exactly one sample point and contains all points of \mathbb{R}^3 that do not have a smaller distance to any other sample point, i.e., the Voronoi cell corresponding to $p \in P$ is given as

$$V_p = \{x \in \mathbb{R}^3 : \forall q \in P \quad \|x - p\| \leq \|x - q\|\}.$$

Closed facets shared by two Voronoi cells are called *Voronoi facets*, closed edges shared by three or more Voronoi cells are called *Voronoi edges* and the points shared by four or more Voronoi cells are called *Voronoi vertices*. The term *Voronoi object* can denote either a Voronoi cell, facet, edge or vertex. The Voronoi diagram is the collection of all Voronoi objects.

Delaunay diagram. The Delaunay diagram $D(P)$ of *P* is the dual of the Voronoi diagram. It is a cell complex that decomposes the convex hull of the points in *P*. The convex hull of four or more points in *P* defines a *Delaunay cell*, if the intersection of the corresponding Voronoi cells is not empty and there exists

no superset of points in P with the same property. Analogously, the convex hull of three or two points defines a *Delaunay face* or *Delaunay edge*, respectively, if the intersection of their corresponding Voronoi cells is not empty. Every point in P is a *Delaunay vertex*. The term *Delaunay object* can denote either a Delaunay cell, face, edge or vertex. If there are no five or more points whose corresponding Voronoi cells have a non-empty intersection, then all Delaunay cells are tetrahedra and the Delaunay diagram is called *Delaunay triangulation*.

Restricted diagrams. The Voronoi diagram $V_S(P)$ of P restricted to S consists of the common intersection of the Voronoi objects with S . The restricted Voronoi diagram fulfills the closed ball property, if every restricted Voronoi object is homeomorphic to a closed ball. The dimension of this ball has to be one less than the dimension of the original Voronoi object. The Delaunay diagram $D_S(P)$ of P restricted to S is defined via the restricted Voronoi diagram in the same way the Delaunay triangulation is defined via the Voronoi diagram.

Theorem 1 [Edelsbrunner, Shah] If $V_S(P)$ fulfills the closed ball property, then $D_S(P)$ and S are homeomorphic.

This theorem stresses the importance of the restricted Delaunay diagram though it is in general not accessible since S is unknown. Here we are not interested in the global restricted Delaunay diagram but only in the restricted Delaunay neighbors of a sample point $p \in P$. These neighbors are sample points in P connected to p by a restricted Delaunay edge and are also not directly accessible. But two observations help to at least approximate the set of restricted Delaunay neighbors. First, by definition, the restricted Delaunay neighbors are a subset of the Delaunay neighbors. Second, it was shown by Giesen and Wagner [GW03] that the restricted Delaunay neighbors cannot be too far away from p . This allows to filter the Delaunay neighbors just by their distance to the sample point p . Note that the filtered set in general still is a superset of the restricted Delaunay neighbors. This is due to the existence of flat tetrahedra called *slivers* in the Delaunay triangulation. But at least the filtered set gives a reasonable local neighborhood. Since we will need it later on, we restate the lemma by Giesen and Wagner here. The lemma bounds the extent of the restricted Delaunay neighborhood.

Lemma 1 [Giesen, Wagner] Let P be an ε -sample of S with $\varepsilon < \frac{1}{2}$. If $p, q \in P$ and pq is a restricted Delaunay edge, then

$$\|p - q\| \leq \frac{2\varepsilon}{1 - \varepsilon} \min\{f(p), f(q)\}.$$




Figure 1: An ε -sample of a cylinder, where the k -nearest neighbors do not adequately represent the surface patch intersected by the query ball.




Figure 2: A rotational laser scan of vase exhibits a very similar sampling pattern as the one shown in Figure 1.

It can be easily seen that an ε -sample is not sufficient to obtain a good neighborhood from taking the k nearest neighbors of a sample point (see Figures 1 and 2). A remedy is to strengthen the sampling condition. In [DGGZ02, FR02] the notion of an (ε, δ) -sample has been introduced.

(ε, δ) -sample. A subset $P \subseteq S$ is called an (ε, δ) -sample of S , if it is an ε -sample, and $\|p - q\| \geq \delta f(p)$, for all $p, q \in P$.

The ε -criterion gives a lower bound on the sampling density, while the δ -criterion provides a corresponding upper bound and controls the positions of the sample points to ensure a certain uniformity of the sampling. A drawback of this definition is that a very uniform

sample that is denser than allowed by the δ -criterion is ruled out though it should be very well suited for all theoretical and practical purposes. On the other hand from a theoretical point of view the δ -criterion is not a severe limitation since Funke and Ramos [FR02] presented an algorithm that computes in almost linear time an (ε', δ) -sample from a given ε -sample using point removal. Points that are not essential for the description of the surface are discarded, i.e., the redundancy of an ε -sample is diminished.

In the following we assume that we are given an (ε, δ) -sample and want to derive bounds on k such that the k -neighborhood contains the restricted Delaunay neighborhood.

3. Bounds

We first give an upper bound on k in order for the k -neighborhood of a sample to contain its restricted Delaunay neighborhood. The statement of the lemma might be confusing at a first glance, since k seems to be bounded from below which does not look like an upper bound. The correct way to read the lemma is, k has to be at most as large as the stated bound to guarantee that the statement of the lemma holds.

Lemma 2 Let P be an (ε, δ) -sample of S . Choosing

$$k \geq \frac{(\delta(1+w) + w)^2}{\delta^2(1-w)^2 - w^4},$$

where $w = \frac{2\varepsilon}{1-\varepsilon}$, guarantees that the k nearest neighbors of a sample point p include all its restricted Delaunay neighbors.

Proof We want to derive this bound from a packing argument. From Lemma 1 we know that all restricted Delaunay neighbors of p are contained in a ball B of radius at most $wf(p)$ centered at p . An upper bound on the number of sample points that we can pack into B immediately gives us the desired upper bound on k . The δ -criterion makes sure that the sample points cannot be packed arbitrarily densely into B . In fact, a sample point q contained in B does not have another sample at distance less than $\delta f(q)$. Just packing such δ -balls into B would give worse bounds than the ones stated above. Thus we are going to exploit another fact proven by Giesen and Wagner [GW03] which allows us to go from a three dimensional to a two dimensional packing problem, see Figure 3.

Fact. Let q be a point in the $wf(p)$ neighborhood of p . Then the distance of q from its orthogonal projection q' onto the tangent plane at p is bounded by $w^2f(p)$, i.e., $\|q - q'\| \leq w^2f(p)$.

This fact implies that the ball of radius $\delta f(q)$ centered at q intersects the tangent plane in a disk whose

squared radius r^2 is at least $\delta^2 f^2(q) - w^4 f^2(p)$. Making use of the Lipschitz continuity of the local feature size, i.e., $f(p) \leq f(q) + \|p - q\|$, we get $f(q) \geq (1 - w)f(p)$. Plugging this into the formula for r^2 we get

$$r^2 \geq (\delta^2(1-w)^2 - w^4)f^2(p).$$

Thus we only have to consider packing disks of radius r non-overlapping into a disk of radius at most $R = (\delta(1+w) + w)f(p)$ in the tangent plane at p . The additional term of $\delta(1+w)f(p)$ in the definition of R accounts for the fact that the δ -ball for sample points at distance $wf(p)$ need not completely lie inside B . We get an upper bound on the packing number by dividing the area of a disk with radius R by the area of a disk with radius r . Thus the packing number can be bounded by

$$\frac{R^2}{r^2} = \frac{(\delta(1+w) + w)^2}{\delta^2(1-w)^2 - w^4}.$$

By construction this packing number also gives an upper bound on the number of sample points in B which itself is an upper bound on the number of restricted Delaunay neighbors of the sample point p . \square




Figure 3: Reducing the three dimensional packing problem to a two dimensional packing problem in the proof of Lemma 2.

Note that the bound in Lemma 2 is not tight since dividing the two surface areas should give only a rough estimate on the packing number. Choosing $\delta = 3w/8 \simeq 3\varepsilon/4$ which seems to be a reasonable choice gives the following upper bound on k :

$$\frac{(11 + 3w)^2}{9 - 18w - 55w^2}$$

In many computational geometry proofs ε is smaller than 0.08, i.e., w is smaller than 0.175. Plugging this into the bound gives that k has to be at most 32 which lies close to the range of k -values that are used in practice. See for example [PGK02, PKKG03], where values of k between 8 and 15 are successfully employed (our formulas yield $k \leq 15$ for $\varepsilon \leq 0.019$). Note though that the bound on k also depends on the ratio $\delta/w = \delta(1 - \varepsilon)/2\varepsilon$. Larger ratios are more

restrictive on the sampling but lead to smaller values for k .

The lower bound on k is the minimal number of restricted Delaunay neighbors for any point in P . This number trivially follows from Theorem 1.

Lemma 3 If P is an ε -sample of S with $\varepsilon < 0.08$, then every sample point has at least three restricted Delaunay neighbors.

Proof It was shown in [AB98] that for $\varepsilon < 0.08$ the restricted Voronoi diagram fulfills the closed ball property. That is, by Theorem 1 the restricted Delaunay triangulation of P is a simplicial surface homeomorphic to S . Since the minimum degree of a vertex in a simplicial surface is three also the minimum number of restricted Delaunay neighbors of any sample point has to be three. \square

4. Conclusion

We derived upper and lower bounds on the k -neighborhood size such that this neighborhood contains the restricted Delaunay neighborhood for every sample point. We derived these bounds under a sampling condition which is common in many computational geometry proofs. For reasonable values of sampling parameters we obtained that some value of k between 3 and 50 should provide a k -neighborhood that allows to faithfully approximate local geometric surface properties provided the sampling is locally uniform. That is, properties that can be proven for the restricted Delaunay neighborhood should also hold for the k -neighborhood. These bounds on k are in good accordance with the values for k that are used in practice.

References

- [AB98] AMENTA N., BERN M.: Surface reconstruction by Voronoi filtering. In *Proc. 14th symposium on Computational geometry* (1998), pp. 39–48.
- [AMN*98] ARYA S., MOUNT D.M., NETANYAHU N.S., SILVERMAN R., WU A.Y.: An optimal algorithm for approximate nearest neighbor searching fixed dimensions. *J. ACM* 45, 6 (1998), 891–923.
- [BWK02] BOTSCH M., WIRATANAYA A., KOBBELT L.: Efficient high quality rendering of point sampled geometry. In *Proc. 13th Eurographics workshop on Rendering* (2002), pp. 53–64.
- [DGGZ02] DEY T.K., GIESEN J., GOSWAMI S., ZHAO W.: Shape dimension and approximation from samples. In *Proc. 13th ACM-SIAM symposium on Discrete algorithms* (2002), pp. 772–780.
- [ES94] EDELSBRUNNER H., SHAH N.R.: Triangulating topological spaces. In *Proc. 10th symposium on Computational geometry* (1994), pp. 285–292.
- [FR02] FUNKE S., RAMOS E.A.: Smooth-surface reconstruction in near-linear time. In *Proc. 13th ACM-SIAM symposium on Discrete algorithms* (2002), pp. 781–790.
- [GW03] GIESEN J., WAGNER U.: Shape dimension and intrinsic metric from samples of manifolds with high co-dimension. In *Proc. 19th symposium on Computational geometry* (2003), pp. 329–337.
- [KV01] KALAI AH A., VARSHNEY A.: Differential point rendering. In *Proc. 12th Eurographics Workshop on Rendering Techniques* (2001), pp. 139–150.
- [KV03] KALAI AH A., VARSHNEY A.: Statistical point geometry. In *Proc. Eurographics/ACM SIGGRAPH symposium on Geometry processing* (2003), pp. 107–115.
- [MN03] MITRA N.J., NGUYEN A.: Estimating surface normals in noisy point cloud data. In *Proc. 19th symposium on Computational geometry* (2003), pp. 322–328.
- [PGK02] PAULY M., GROSS M., KOBBELT L.P.: Efficient simplification of point-sampled surfaces. In *Proc. conference on Visualization '02* (2002), pp. 163–170.
- [PKKG03] PAULY M., KEISER R., KOBBELT L.P., GROSS M.: Shape modeling with point-sampled geometry. *ACM Transactions on Graphics* 22, 3 (2003), 641–650.
- [RL00] RUSINKIEWICZ S., LEVOY M.: Qsplat: a multiresolution point rendering system for large meshes. In *Proc. 27th conference on Computer graphics and interactive techniques* (2000), pp. 343–352.
- [ZPvG01] ZWICKER M., PFISTER H., VAN BAAR J., GROSS M.: Surface splatting. In *Proc. 28th conference on Computer graphics and interactive techniques* (2001), pp. 371–378.