

A Sampling Theorem for MLS Surfaces

Peer-Timo Bremer

John C. Hart

University of Illinois, Urbana-Champaign

Abstract

Recently, point set surfaces have been the focus of a large number of research efforts. Several different methods have been proposed to define surfaces from points and have been used in a variety of applications. However, so far little is known about the mathematical properties of the resulting surface. A central assumption for most algorithms is that the surface construction is well defined within a neighborhood of the samples. However, it is not clear that given an irregular sampling of a surface this is the case. The fundamental problem is that point based methods often use a weighted least squares fit of a plane to approximate a surface normal. If this minimization problem is ill-defined so is the surface construction. In this paper, we provide a proof that given reasonable sampling conditions the normal approximations are well defined within a neighborhood of the samples. Similar to methods in surface reconstruction, our sampling conditions are based on the local feature size and thus allow the sampling density to vary according to geometric complexity.

1. Introduction

The moving least squares (MLS) method filters a finite noisy collection of scattered surface points, projecting each data point onto a local approximated tangent plane [Lev03]. The robustness of this filtering has made MLS an attractive method for surface reconstruction [ABCO*03], and the MLS surface an attractive shape representation for editing [ZPKG02], modeling [PKKG03] and animation [MKN*04].

The simplicity of the MLS method belies an underlying complexity that hinders analysis. For example the original algorithms for projecting the data points onto the MLS surface [Lev03, ABCO*03] actually missed the surface, as previously reported [AK04]. The problem is that the MLS surface is defined as the stationary points of a projection, but the direction of this projection is dynamic across space, and can be undefined in some data point configurations. The robustness of the MLS method relies on ensuring this projection direction is well defined.

Adamson and Alexa [AA03] used this criterion to construct a sampling condition, labeling an MLS surface as *well sampled* if this projection direction is defined everywhere within a neighborhood of the MLS surface. This property suffices as a *smoothness criterion* as it allows the implicit function theorem to prove the differentiability of the MLS surface embedding. However this property fails as a *sam-*

pling criterion because it measures the relationship of the samples to the reconstructed MLS surface and not the original surface from which the samples were drawn.

This paper derives a new sampling condition that states a criterion on the samples based on the local feature size of the original surface sufficient to yield a well-defined (and smooth) MLS surface reconstruction. Our sampling condition is also adaptive since it is based on local feature size, which will eventually enable it to provide tighter bounds than one could with a uniform sampling.

This theorem relies on the local feature size of the sampled surface, which is in general not available when measuring real world data. Our review of previous work in Sec. 2 recalls that this was also the case for Voronoi-based surface reconstruction methods where a well-defined sampling criteria was needed first before stronger theorems on the topological and geometric accuracy of the reconstruction could be proven. A similar situation exists now for MLS surface reconstruction. Sec. 3 reviews properties of adaptive sampling relative to the local feature size that we will need for our sampling condition, and Sec. 4 reviews the MLS method using a convenient implicit formulation of the MLS surface.

Sec. 5 states and proves a theorem that shows that the normal vector needed by the MLS method does not vanish using a property of the MLS weighted variance of the sampling.

This preparation enables Sec. 6 to state and prove the main theorem of the paper, that given appropriate though overly restrictive conditions on the surface and its sampling, the normal is well defined (does not vanish) near the sampled surface S (as opposed to the reconstructed surface). Sec 7 concludes with a discussion of the sampling conditions with ideas on how they could be further relaxed.

2. Related Work

Our sampling criterion for an MLS surface is based on the local feature size of the original surface from which the samples were taken. This result resembles the first sampling criteria for the crust method [ABK98, AB99], which was shown to work when the number of samples was sufficiently dense relative to the local feature size. Even though the crust method worked well in practice with samples 0.5 times the local feature size apart, a very dense sampling of 0.06 times the local feature size was needed to prove the correctness of the crust method. Likewise, the constants used in the proof of our sampling theorem are similarly much more restrictive than what is used in practice for the MLS method, and will likely be relaxed by further, future analysis.

Similar to the methods described in this paper the original crust method is limited by requiring knowledge of the medial axis. Nevertheless, the crust algorithm forms is the ancestor of number of very successful combinatorial surface reconstruction algorithms such as power crust [ACK01], cocone [ACDL02] and the robust cocone [DG04]. All of which were eventually proven to yield topologically correct and geometrically accurate approximations to the original surface, under the necessary sampling conditions.

Recently, several papers have been published aiming at proving results similar to those of Voronoi-based algorithms. Kolluri [Kol05] proved that, under globally uniform sampling conditions, one particular implicit MLS reconstruction is homeomorphic and geometrically close to the original surface. However, Kolluri assumed input normals at sample points and suggested using the probabilistic method of Mitra and Nguyen [MN03] to estimate normals. By design these normal are likely but not guaranteed to be well behaved. Furthermore, the normals are fit to a uniformly weighted local neighborhood of sample points, whereas the MLS surface derives its normals from the non-uniform weighted average.

In [DS05b] Dey and Sun based on some earlier results of Dey and Goswami [DG04] show that an MLS surface similar to the one used here is isotopic to the original surface under uniform sampling conditions. The method relies heavily on well approximated normals and they provide a combinatorial algorithm to estimate provably good normals. That result complements our own as it provides good normals, which we lack, but requires a less desirable uniform sampling.

In a parallel paper Dey and Sun [DS05a] propose a new variant of an MLS surface which also reconstructs a surface

isomorphic to the original but under adaptive sampling conditions. Unlike the assumptions of Section 3, the new construction does not require a bounded range of local feature sizes, but leads to a result based on a different MLS surface than the one commonly used.

3. Preliminaries

Before we describe MLS surfaces, it will first be useful to review some useful definitions and results regarding the sampling of a surface.

A sampling condition should guarantee that a surface is “appropriately” sampled everywhere, which at least intuitively means that the local sampling density should depend on the local geometric complexity, e.g. some indication of curvature. For a point s on a closed manifold surface $S \subset \mathbb{R}^3$, we can define its *local feature size* $\rho(s)$ as the radius of the largest closed ball whose boundary contains s but whose interior does not intersect S .

The *medial axis* M of the surface S is the collection of points that have exactly two closest points in S . It is not difficult to see that the above definition of local feature size matches its usual (e.g. [AB99]) definition $\rho(p) = d(p, M)$ where $d(p, M)$ is the distance from a point p to its closest point in the closure of M . The local feature size is Lipschitz [AB99] with constant 1 so for any $a, b \in S$,

$$|\rho(a) - \rho(b)| \leq \|a - b\|. \quad (1)$$

For any point $q \in \mathbb{R}^3$ let $\text{proj}(q)$ denote the projection of q onto a closest point in S . For $\tau < 1$ we define a *tubular neighborhood* T_τ of the surface S as

$$T_\tau = \{q \in \mathbb{R}^3 \mid d(q, S) \leq \tau \rho(\text{proj}(q))\} \quad (2)$$

where as before $d(q, S)$ is the distance from q to the nearest point in S . In other words, the tubular neighborhood consists of points whose projection w onto S is no further away than τ times the local feature size of w . w is unique since $\tau < 1$.

The analysis of the MLS surface is eased by imposing both an upper and lower bound on the sampling density. The lower bound ensures that no part of S remains completely unsampled while the upper bound is necessary to avoid an arbitrary oversampling of a small region which could distort the analysis of the weighted averages.

We impose these bounds by requiring an (ϵ, δ) -sampling of the surface. A set $P = \{p_i\}_{i=1}^N, p_i \in S$, is an (ϵ, δ) -sampling of the surface S if (1) the collection of balls centered at the p_i of radius $\epsilon \rho(p_i)$ covers S , and (2) balls center at the p_i of radius $\delta \rho(p_i)$ do not intersect.

The density of an (ϵ, δ) -sampling is dependent on the local feature size. If the local feature size could become arbitrarily small the sampling density could grow arbitrarily

large which would further complicate the analysis of covariance. Let

$$\rho_{\min}(S) = \min_{s \in S} \rho(s), \quad \rho_{\max}(S) = \max_{s \in S} \rho(s) \quad (3)$$

represent the minimum and maximum local feature size of the surface S . We use these bounds to construct a bound α on the “bandwidth” of the surface S as $\alpha > \rho_{\max}(S)/\rho_{\min}(S)$.

Given an (ϵ, δ) -sample of a surface one can establish upper and lower bounds on the number of samples.

Lemma 1 (Sampling Density) Given an (ϵ, δ) -sampling P of surface S , let $F \subset \mathbb{R}^2$ indicate an arbitrary planar projection of S . Furthermore, let $Q = \{q_i\}_{i=1}^N$ be any set of samples such that balls centered at of radius $\delta\rho_{\min}(S)$ cover S . Then the number of samples in P is bounded from below and above by

$$\frac{\text{Area}(F)}{\pi\epsilon^2\rho_{\max}^2(S)} \leq |P| \leq |Q|. \quad (4)$$

Proof The lower bound holds because the cover of S with balls of radius $\epsilon\rho(S)$ when projected is a cover of the projection F of S and the planar area the projections of these balls could cover is less than the denominator. The upper bound follows from the fact that any covering of S with spheres of size $\delta\rho_{\min}(S)$ has more samples than allowed. \square

4. The MLS Surface

This section reviews the moving least squares method, using an implicit formulation based on the weighted covariance matrix [AA03].

Let $S \subset \mathbb{R}^3$ be a smooth manifold and $P = p_{i=0}^N$ be a set of points used to discretely sample S . The MLS method localizes its least squares fit using a *weight function* $\theta: \mathbb{R}^+ \rightarrow \mathbb{R}^+$ that decreases monotonically to zero as its parameter increases. As do most others, we use the Gaussian $\theta(r) = \exp(-r^2/\sigma^2)$.

Given a query point in space $q \in \mathbb{R}^3$ the weighting function θ defines its *weighted average*

$$a(q) = \frac{\sum_i \theta(\|q - p_i\|) p_i}{\sum_i \theta(\|q - p_i\|)}, \quad (5)$$

and its 3×3 *weighted covariance matrix*

$$\text{cov}_\theta(q) = \sum_i \theta(\|q - p_i\|) (q - p_i)(q - p_i)^T \quad (6)$$

where q and p_i are represented by 3-element column vectors. If $\text{cov}_\theta(q)$ has a unique smallest eigenvalue λ_1 then its corresponding eigenvector v_1 is called the *normal direction* at q and denoted $n(q)$ (or simply n when the context is clear.)

This normal direction allows the formulation of the MLS surface reconstruction an implicit surface $\hat{S} = f^{-1}(0)$ of

$$f(q) = n(q)^T (q - a(q)). \quad (7)$$

This implicit formulation requires that the normal directions

$n(q)$ be well defined (non-vanishing), which lead to a Adamson and Alexa’s sampling criterion: a surface S is *well sampled* by a point set P if the normal directions $n(q)$ are defined (in other words the weighted covariance matrix has a unique minimum eigenvalue) inside a neighborhood of \hat{S} .

This definition has also been used by others, e.g. [AK04], to ensure that the local frame used in the projection operator can be fitted uniquely. While this definition of a “well sampled” surface is valid, it is problematic for both practical and theoretical reasons. The MLS surface construction is initialized by evaluating $f(q)$ or at least $n(q)$ in the neighborhood of the original surface S and its sample points P , and not the reconstruction \hat{S} which could be a significant distance from S and P . Hence according to this definition, the surface S could be “well sampled” by P and nevertheless result in an undefined normal direction at one or more of the sample points.

Furthermore, this definition of “well sampled” is circular in that it is impossible to test whether a surface is well sampled unless the sampling suffices to generate a singularity free MLS reconstruction. We would prefer instead a guarantee that under specified sampling conditions relative to S and P (and not \hat{S}) that a proper MLS surface results. We thus call a sampling P of a surface S *well defined* if the normal directions $n(q)$ do not vanish over a neighborhood of S that includes the samples P . We can prove a sampling to be well defined by ensuring the weighted covariance matrix has a unique smallest eigenvalue over a neighborhood of S .

5. Uniqueness of the Smallest Covariance Eigenvalue

This section states and proves a condition for the uniqueness of the smallest covariance eigenvector in terms of the weighted variance. The *weighted variance* of the samples P at a query position q is defined

$$\text{var}_\theta(q) = \sum_{i=1}^N \theta(\|q - p_i\|) \|q - p_i\|^2. \quad (8)$$

We can also define the *directional weighted variance* by examining the variance in a specific direction indicated by the unit vector $n \in S^2$ as

$$\text{var}_n(q) = \sum_{i=1}^N \theta(\|q - p_i\|) (n^T (q - p_i))^2. \quad (9)$$

Combining (9) with (6) relates the directional weighted variance to the weighted covariance as

$$\text{var}_n(q) = n^T \text{cov}_\theta(q) n, \quad (10)$$

which allows us to express the (real, symmetric and positive definite) directional weighted variance in terms of the decomposition of the weighted covariance matrix into eigenvalues $\lambda_1, \lambda_2, \lambda_3$ and corresponding unit-length mutually perpendicular eigenvectors v_1, v_2, v_3 as

$$\text{var}_n(q) = a^2 \lambda_1 + b^2 \lambda_2 + c^2 \lambda_3 \quad (11)$$




Figure 1: The query point q projects as far as $\tau\rho(w)$ to $w \in S$ which yields two osculating balls B_+, B_- of radius $\rho(w)$ whose complement bounds the surface. The value τ indicates the magnitude of this projection in units of local feature radius.




Figure 2: The two mutually perpendicular planes pass through q parallel to the y -axis, intersecting the ball B of radius $\rho(w)$ in an hourglass shape. Data points p_i in the hourglass region below (or above) the two planes increase var_n more, whereas data points above only one of the planes increase var_x more.




Figure 3: Top view of S around w . The regions R_i cover the hourglass region between the two dashed lines where data points contribute more to var_n than var_x . Points outside the hourglass are only partially measured by the regions G_i but nevertheless overcome even the overestimate due to R_i overlap.

where $a = n^T v_1, b = n^T v_2$ and $c = n^T v_3$.

We can use this relationship to show that when the smallest eigenvalue of the weighted covariance matrix is not unique, then the direction of least variance is also not unique.

Lemma 2 Let $\text{cov}_\theta(q)$ be a weighted covariance matrix with unit eigenvectors v_1, v_2 , and v_3 and eigenvalues $\lambda_1 = \lambda_2 \leq \lambda_3$. Then for any unit vector v there exists a perpendicular unit vector w such that $\text{var}_v(q) \geq \text{var}_w(q)$.

Proof Without loss of generality, we rotate the coordinate system so the eigenvectors v_i are the coordinate axes. Then given any unit vector $v = (a, b, c)$ such that $|c| \neq 1$, the vector $w = 1/\sqrt{a^2 + b^2}(-b, a, 0)$ is perpendicular. (If $|c| = 1$ then w can be any unit vector of the form $(a, b, 0)$.) Then

$$\begin{aligned} \text{var}_v(q) &= a^2\lambda_1 + b^2\lambda_2 + c^2\lambda_3 \\ &\geq a^2\lambda_1 + b^2\lambda_2 + c^2\lambda_1 \\ &= \lambda_1 \\ &= \text{var}_w(q). \quad \square \end{aligned}$$

Thus a weighted covariance with two equal smallest eigenvalues indicates the existence of a plane where the variance is minimal. The contrapositive of this statement provides a sufficiency condition to guarantee the smallest eigenvalue of a covariance matrix is unique.

Theorem 1 If the directional weighted variance $\text{var}_v(q)$ for some unit vector v , is strictly less than the directional weighted variance of any perpendicular unit vector w , then the smallest eigenvalue of $\text{cov}_\theta(q)$ is unique.

Our MLS sampling condition now relies on showing for

each query point q in an adaptive tubular neighborhood of the surface S that we can find a normal direction whose weighted variance is less than that of any perpendicular tangent direction.

6. The Sampling Theorem

In this section, we project the query point q onto a surface point $w \in S$ in a normal direction n extending from w to the center of an osculating ball of radius $\rho(w)$, as illustrated by Fig. 1. Theorem 1 shows that the weighted normal does not vanish at a given location if we can find a direction whose weighted variance is strictly less than the weighted variance in any perpendicular direction. We show that for the point w , the variance in the normal direction n is strictly less than in any direction tangent to S at w . We derive an upper bound of the weighted variance in the normal direction n , and a lower bound of the weighted variance in an arbitrary chosen tangent direction x . We then determine sampling conditions under which the least possible variance in the x -direction still exceeds the greatest possible variance in the normal direction, which by Theorem 1 shows the normal direction is unique and well defined at q .

In order to show $\text{var}_n(q) < \text{var}_x(q)$, we need to partition our data points p_i into ones that increase var_n more versus ones that increase var_x more. From the definition of directional variance (9), this competition boils down to the magnitude of the n - or x - component of $(p_i - q)$. We thus construct two separating planes

$$P_\pm = \{p \in \mathbb{R}^3 | n^T(p - q) = \pm x^T(p - q)\} \quad (12)$$

through the query point q parallel to the $y = n \times x$ axis, as drawn in Fig. 2. Data points below (or above) both planes will increase $\text{var}_n(q)$ more than $\text{var}_x(q)$.

We allow the query point q to be as much as τ times the local feature size (at $w = \text{proj}(q)$) above (or below) the surface S , and the surface S can recede away from these planes by as much as is allowed by the curvature of w 's medial ball as shown in Fig. 2. Due to this recession, the region on this ball below both planes (corresponding to points that increase var_n more than var_x forms an hourglass shape (bounded by a pair of opposing circular arcs) whose x -width grows as the magnitude of y increases.

We project this hourglass region onto the lower medial ball B_- , and bound it with the union of a collection of "risky" regions R_i , indicated in red in Fig. 3. To be on the safe side we seek to find a limit of the maximum possible effect of these "risky" samples. Hence Sec. 6.2 defines the regions R_i and uses them to overestimate the number of samples and their contribution to var_n . The contribution of these "risky" points are countered by "good" samples that do not fall below (or above) both planes. Sec. 6.3 underestimates their number and contribution to both var_x and var_n by examining only their subsets in the "good" regions G_i , shown in green in Fig. 3.

The upper bound of var_n computed via R_i and G_i and the lower bound of var_x computed via G_i are only computed within a radius of $1/4$ of the local feature size of w . "Far" samples outside this radius are discussed in Sec. 6.4 where their contribution is ignored by the lower bound of var_x and an overestimate of all of their contributions (regardless of their position with respect to the separating planes) are added to the upper bound of var_n .

6.1. Upper Bounds of Data Point Displacement

Before we can construct the R_i regions bounding the data points used in the upper bound of var_n , we first need to determine the geometry of the hourglass shape obtained by slicing the lower (or upper) medial ball by the two separating planes. This geometry is revealed by computing how far a point descends from the tangent plane of a sphere as a function either of the distance along the tangent plane or the Euclidean distance from the descended point to the plane's point of tangency. Because the surface S lies between the two medial balls of radius $\rho(w)$ that osculate the surface at w , this spherical displacement serves as an upper bound on the distance from the surface to the tangent plane.

Let w be a surface point with local feature size ρ . Then a surface point $p \in S$ projecting perpendicularly to w 's tangent plane to a point x units away from w can be no farther than

$$\max |z| = \rho - \sqrt{\rho^2 - x^2} \quad (13)$$

from the tangent plane, using the Pythagorean theorem on the configuration shown in Fig. 4.




Figure 4: Geometry for determining the maximum vertical displacement z a point p can extend from the tangent plane of a surface point w with local feature size ρ , in terms of the distance $d = \|p - w\|$ in space or the length x along the tangent plane.

To find the maximum z displacement given instead the distance $d = \|p - w\|$, we first use the law of cosines on the isosceles triangle to find $\cos \alpha = 1 - d^2/(2\rho^2)$ and angle sums to find $\alpha = 2\theta$. Then

$$\max |z| = d \sin \frac{\alpha}{2} = d \sqrt{\frac{1 - \cos \alpha}{2}} = \frac{d^2}{2\rho}. \quad (14)$$

6.2. Upper Bound on the Number of "Risky" Samples

The query point q can be as much as $\tau\rho(w)$ units above its projection $w \in S$. The planes P_{\pm} descend with unit slope, which creates an affine correspondence between altitude and horizontal offset of $|x| = \tau + |z|$. Using (14) to maximize $|z|$ given the Euclidean distance to w , the hourglass shape H is defined on the medial ball B_- as

$$H = \{p \in B_- \mid |x^T p| \leq \tau\rho(w) + \frac{\|p - w\|^2}{2\rho(w)}\}. \quad (15)$$

We cover H with the regions $R_i \subset B_-$ defined as

$$R_i = \left\{ p \in B_- \mid \begin{array}{l} (i-1)\Delta r \leq \|p - w\| \leq i\Delta r, \\ |x^T p| \leq \frac{(i\Delta r)^2}{2\rho(w)} + \tau\rho(w) \end{array} \right\} \quad (16)$$

for positive integers i .

We now need an upper bound on the number of samples in R_i which we will use to compute an upper bound of var_n .

Lemma 3 Let $P(R_i) \subset P$ be the subset of the (ϵ, δ) -samples of S lying between B_- and B_+ that project in the n -direction onto R_i . Then

$$\begin{aligned} |P(R_1)| &\leq \left\lceil \frac{2}{\Delta b^2} \right\rceil \lceil \Delta r \rceil \lceil \Delta r^2 + 2\tau \rceil, \\ \Delta b &= \frac{\sqrt{2}\delta(\rho(w) - i\Delta r)\sqrt{\rho(w)^2 - 2\rho(w)i\Delta r}}{i\Delta r} \end{aligned} \quad (17)$$

and for $i > 1$,

$$|P(R_i)| \leq \left[\frac{1}{\Delta b^2} \right] \times \left[2i\Delta r - 2\sqrt{\star^2 - (\tau + i^2\Delta r^2/2)^2} \right] \times [(i\Delta r)^2 + 2\tau], \quad (18)$$

$$\star = \sqrt{(i-1)^2\Delta r^2 - \frac{(i-1)^4\Delta r^4}{4\rho(w)^4}}. \quad (19)$$

Proof Let $S_i \subset S$ be the portion of the surface between B_- and B_+ that projects in the n -direction onto R_i . The set P is an (ϵ, δ) -sampling of S , so we seek to find a finer cover of the surface patch S_i with balls of size less than $\delta\rho_{\min}$. The number of balls in this finer cover will serve as an upper bound of the number of samples in $P(R_i)$.

If the region S_i was flat we could cover it with balls centered at the vertices of a rectilinear grid of cell size $\sqrt{2}\delta\rho_m$, with $\rho_m = \min_{s \in S_i} \rho(s)$. Since S_i may not be flat, we will need to compress the covering grid using a Lipschitz constant on its altitude function over the tangent plane at w .

From (16) we find $\|w - s\| \leq i\Delta r, \forall s \in S_i$. Since local feature size is Lipschitz with constant one, it follows that $\rho(s) \geq \rho(w) - i\Delta r, \forall s \in S_i$. We can construct from (13) an altitude equation with the most severe slope allowed given $\rho(s) = \rho(w) - i\Delta r$ as

$$z(x) = (\rho(w) - i\Delta r) - \sqrt{(\rho(w) - i\Delta r)^2 - x^2} \quad (20)$$

which yields the slope

$$z'(x) = \frac{-x}{\sqrt{(\rho(w) - i\Delta r)^2 - x^2}}. \quad (21)$$

Since $|z'|$ increases monotonically we find its maximum over S_i occurs when $x = i\Delta r$, yielding

$$\text{Lip}_{S_i} z(x) = \frac{i\Delta r}{\sqrt{\rho(w)^2 - 2\rho(w)i\Delta r}}. \quad (22)$$

Thus a uniform rectilinear grid of balls spaced

$$\Delta b = \frac{\sqrt{2}\delta(\rho(w) - i\Delta r)\sqrt{\rho(w)^2 - 2\rho(w)i\Delta r}}{i\Delta r} \quad (23)$$

apart suffices to cover S_i .

Fig. 3 shows each R_i consists of two components, one above and one below the x -axis. The projected area of an R_i 's upper component is included in the tangent plane rectangle

$$\left[-\frac{i^2\Delta r^2}{2} - \tau, \frac{i^2\Delta r^2}{2} + \tau \right] \times \left[\sqrt{\star^2 - \left(\frac{i^2\Delta r^2}{2} + \tau \right)^2}, i\Delta r \right],$$

as shown in Figure 5. Notice that unlike its near side (y_{\min}), we have not foreshortened the far side of the rectangle ($y_{\max} = i\Delta r$) because we are counting samples on its projection onto the surface S , which could slope less and perhaps even be flat. \square

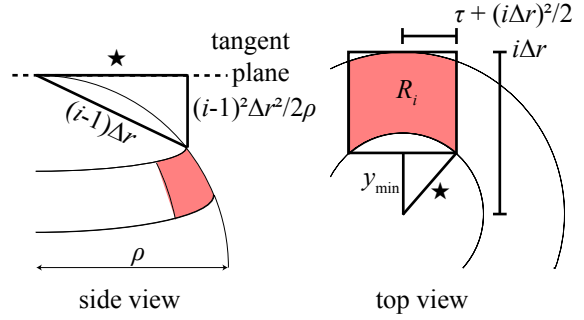


Figure 5: Derivation of a bounding rectangle of the projection of R_i in the z -direction onto the tangent plane at w .

6.3. Lower Bound on the Number of “Good” Samples

All samples outside the hourglass shape H contribute more to var_x than to var_n . Rather than considering the complete surface outside H we restrict ourselves to a smaller subset for which we bound the contribution to the variances with relative ease. As indicated in Figure 3 the G_i 's are defined as

$$G_i = \left\{ p \in B_- \mid \begin{array}{l} \|p - w\|^2 - (n^T p)^2 \leq 2^{i+1}\Delta r, \\ |x^T p| \geq 2^i\Delta r \end{array} \right\} \quad (24)$$

By definition R_i is outside H and we now need a lower bound on the number of samples in R_i . This bound is used in computing the lower bound on var_x and also must be considered in the upper bound of var_n .

Lemma 4 Let $P(G_i) \subset P$ be the subset of the (ϵ, δ) -samples of S lying between B_- and B_+ that project in the n -direction onto G_i . Then

$$|P(G_i)| \geq \left\lfloor \frac{2\pi-3\sqrt{3}(2^{i+1}\Delta r)^2}{\pi\epsilon^2\rho_M^2} \right\rfloor, \quad (25)$$

$$\rho_M = \rho + \sqrt{2\rho^2 - 2\rho\sqrt{\rho^2 - 4^{i+1}\Delta r^2}}, \quad (26)$$

where $\rho = \rho(w)$ and ρ_M denotes the largest possible feature size over the region G_i (not the entire surface S).

Proof The numerator is the area of G_i and the denominator includes an upper bound ρ_M of the local feature size of S over G_i using the Lipschitz bound on the local feature size and the results of Sec. 6.1. \square

We will later also need the following upper bounds.

Lemma 5 For a sample $p \in P(G_i)$,

$$|n^T(q-p)| \leq \tau + \rho - \sqrt{\rho^2 - 4^{i+1}\Delta r^2}, \quad (27)$$

$$\|w-p\|^2 \leq 2\rho - 2\sqrt{\rho^2 - 4^{i+1}\Delta r^2}, \quad (28)$$

$$\|q-p\|^2 \leq \left(\tau + \rho - \sqrt{\rho^2 - 4^{i+1}\Delta r^2} \right)^2 + 4^{i+1}\Delta r^2. \quad (29)$$

Proof The rather straightforward application of (13) to the geometry of G_i on B_- . \square

6.4. Upper Bound of var_n for “Far” Samples

We will use the regions R_i and G_i to analyze samples inside a ball about q of radius one-fourth the local feature size of w . Outside this ball, we will ignore the samples when computing var_x , but will use the following lemma to overestimate their effect on var_n .

Lemma 6 Let $B_{\frac{\rho}{4}} = B(q, \rho(w)/4)$. Then

$$\text{var}_n(P \setminus B_{\frac{\rho}{4}}) \leq \frac{2}{\delta^2 \rho_{\min}^3} \int_{\rho/4 - 2\rho_{\min}}^{\infty} r^4 \theta(r) dr. \quad (30)$$

Proof We fill space with a ridiculously large number of samples by placing them on a series of concentric spherical shells. The maximum number of samples $P(r)$ possible on one of these shells $\partial B(q, r)$ is

$$|P(r)| \leq \frac{4\pi r^2}{\pi \delta^2 \rho_{\min}^2} \quad (31)$$

whose numerator is the surface area of the shell and denominator is the surface area consumed by an (ϵ, δ) -sample (recall Lemma 1). We can thus compute an upper bound of var_n for the shell $\partial B(q, r)$ as

$$\text{var}_n P(r) \leq \frac{4r^2}{\delta^2 \rho_{\min}^2} r^2 \theta(r) \quad (32)$$

where the second factor r^2 bounds the $(n^T(q-p))^2$ factor of (9) by pretending all of the samples lie along the z -direction from q .

Because of the (ϵ, δ) -sampling, the shells cannot be spaced closer than $2\rho_{\min}$ from each other. We can thus bound the contribution to var_n of all samples outside $B_{\frac{\rho}{4}}$ as

$$\text{var}_n(P \setminus B_{\frac{\rho}{4}}) \leq \sum_{i=0}^{\infty} \text{var}_n P(\rho/4 + 2i\rho_{\min}) \quad (33)$$

$$= \frac{4}{\delta^2 \rho_{\min}^2} \sum_{i=0}^{\infty} (\rho/4 + 2i\rho_{\min})^4 \theta(\rho/4 + 2i\rho_{\min}), \quad (34)$$

$$\leq \frac{4}{\delta^2 \rho_{\min}^2} \int_0^{\infty} (\rho/4 + 2(i-1)\rho_{\min})^4 \theta(\rho/4 + 2(i-1)\rho_{\min}) di, \quad (35)$$

which by a change of variables equals (30). \square

6.5. Main Theorem

Theorem 2 Let S be a surface of bounded positive local feature size $\rho_{\max}/\rho_{\min} \leq \alpha$, sampled by an (ϵ, δ) -sampling of points $p \in P$ no farther than $\tau\rho(w)$ from S , where w is the closest point on S to p . Then for $\alpha \leq 1000$, $\epsilon \leq 1/200$, $\delta \geq 1/2000$, and $\tau \leq 1/250$, the MLS surface constructed with an adaptive Gaussian kernel of standard deviation $\sigma = \rho(w)/25$ on the samples P is well defined in that its normals never vanish over the τ neighborhood of S .

Proof Let $w = \text{proj}(q)$ be the closest point on S to a query point q . Without loss of generality assume $n = (0, 0, 1)$ to be the normal of S at w , $\rho(w) = 1$, and $x = (1, 0, 0)$ represent an arbitrarily chosen unit vector perpendicular to n .

We will prove the result by contradiction, by first assuming

$$\text{var}_x(q) \leq \text{var}_n(q). \quad (36)$$

We now split the evaluation of the variances inside $B_{\frac{\rho}{4}}$ and outside $B_{\frac{\rho}{4}}$

$$\text{var}_x(P \cap B_{\frac{\rho}{4}}) \leq \text{var}_n(P \cap B_{\frac{\rho}{4}}) + \text{var}_n(P \setminus B_{\frac{\rho}{4}}), \quad (37)$$

ignoring the samples outside $B_{\frac{\rho}{4}}$ for var_x , which makes the left side smaller and leaves the right side unchanged. We then restrict both sides to the regions $R = \{\dots, R_i, \dots\}$ and $G = \{\dots, G_i, \dots\}$ defined earlier

$$\sum_{i=0}^{|G|} \text{var}_x(G_i) \leq \sum_{i=0}^{|G|} \text{var}_n(G_i) + \sum_{i=1}^{|R|} \text{var}_n(R_i) + \text{var}_n(P \setminus B_{\frac{\rho}{4}}), \quad (38)$$

ignoring the samples not above R_i on the left side which makes it smaller yet. The number of regions in R is $|R| = \lceil \frac{1}{4\Delta r} \rceil$, and in G is $|G| = \lceil -\log_2(4\Delta r) - 1 \rceil$.

Now we bound the remaining terms. We assume a minimum number of samples in each G_i , since for each sample $p \in G_i$, $|n^T(q-p)| \leq |x^T(q-p)|$, which leads to the lower bound

$$\text{var}_x(G_i) \geq |P(G_i)| 4^i \Delta r^2 \theta(r_{\max}) \quad (39)$$

using the minimal number of samples, the minimal distance in x direction, and the smallest possible weight, obtained by setting

$$r_{\max} = \sqrt{\left(\tau + \rho - \sqrt{\rho^2 - 4^{i+1} \Delta r^2}\right)^2 + 4^{i+1} \Delta r^2}. \quad (40)$$

Alternatively the upper bound

$$\text{var}_n(G_i) \leq |P(G_i)| \left(\rho - \sqrt{\rho^2 - 4^{i+1} \Delta r^2} + \tau\right)^2 \theta(2^i \Delta r), \quad (41)$$

uses again the fewest possible samples, but the maximal distance in the n direction and the maximum weight. Likewise the upper bound

$$\text{var}_n(R_i) \leq |P(R_i)| \left(\frac{i\Delta r^2}{2} + \tau\right)^2 \theta((i-1)\Delta r), \quad (42)$$

uses the largest number of samples, maximal distance, and maximal weight.

Setting $\Delta r = \frac{1}{128}$ and using the parameter settings provided by the Theorem statement, the evaluation of the inequality (via a Mathematica program) leads to the contradiction

$$0.000849593 \leq 0.000707577. \quad \square \quad (43)$$

7. Conclusion

We have shown that given a dense enough sampling the normal estimations based on the covariance matrix are well de-

finer in a neighborhood around the original samples and so is the point set surface construction. The results suggest that normal directions should not be computed any distance away from the samples as they might not be well defined in which case the result is random noise.

While the proof is nearly exclusively design to show the existence of proper sampling conditions, most approximations (except $\text{var}_n(P \setminus B_{\frac{\rho}{4}})$) are tight enough to suggest some practical implications. An important but unavoidable difference to traditional methods is that the point set surface used here relies on an adaptive σ . The rather small value of σ is mostly due to the crude approximation of $\text{var}_n(P \setminus B_{\frac{\rho}{4}})$. Ignoring the contribution outside $B_{\frac{\rho}{4}}$ could increase σ and decrease the sampling density. In practice θ is often evaluated over finite support and $P \setminus B_{\frac{\rho}{4}}$ will likely fall outside its domain. Finite support would also allow remove the bound on the local feature size.

Even more restrictive than σ is the small τ , which seems to be less of an artifact of the approximations and more of an indication of how unstable the covariance matrix is at estimating the normal. τ could be raised if σ was increase accordingly. For example $\tau = 0.01$ and $\sigma = 0.2$ are possible values if one ignores $\text{var}_n(P \setminus B_{\frac{\rho}{4}})$. But raising σ to nearly the radius of $B_{\frac{\rho}{4}}$ increases the contribution of points outside $B_{\frac{\rho}{4}}$ significantly. Even more sophisticated estimations of $\text{var}_n(P \setminus B_{\frac{\rho}{4}})$ would cause a comparatively large σ to likely invalidate the proof. One practical solution is to use $n(a(q))$ instead of $n(q)$ as it is implemented in the extremal surface extension of PointShop3d. This would require only $a(q)$ to lie inside T which would allow much greater distances of q to S . Unfortunately, we have yet to formally bound the distance of $a(q)$ from the surface.

Acknowledgments This work was funded in part by the NSF under the SGER grant SCI-0432257 and the ITR grant SCI-0121288. Thanks to Tamal Dey, Marc Alexa and the PBG reviewers for very useful comments in the preparation of this work.

References

- [AA03] ADAMSON A., ALEXA M.: Approximating and intersecting surfaces from points. In *Proc. Symp. Geom. Proc.* (2003), pp. 230–239.
- [AB99] AMENTA N., BERN M.: Surface reconstruction by voronoi filtering. *Discrete and Computational Geometry* 22 (1999), 481–504.
- [ABCO*03] ALEXA M., BEHR J., COHEN-OR D., FLEISHMAN S., LEVIN D., SILVA C. T.: Computing and rendering point set surfaces. *IEEE TVCG* 9, 1 (2003), 3–15.
- [ABK98] AMENTA N., BERN M., KAMVYSSELIS M.: A new voronoi-based surface reconstruction algorithm. In *Proc. SIGGRAPH* (1998), pp. 415–421.
- [ACDL02] AMENTA N., CHOI S., DEY T. K., LEEKHA N.: A simple algorithm for homeomorphic surface reconstruction. *Intl. J. Comp. Geom. and Apps.* 12, 1–2 (2002), 125–141.
- [ACK01] AMENTA N., CHOI S., KOLLURI R. K.: The power crust. In *Proc. Symp. Solid Modeling and Apps.* (2001), pp. 249–266.
- [AK04] AMENTA N., KIL Y. J.: Defining point-set surfaces. *ACM TOG* 23, 3 (2004), 264–270.
- [DG04] DEY T. K., GOSWAMI S.: Provable surface reconstruction from noisy samples. In *Proc. Symp. Comp. Geom.* (2004), pp. 330–339.
- [DS05a] DEY T. K., SUN J.: *An Adaptive MLS Surface for Reconstruction with Guarantees*. Tech. Rep. OSU-CISRC-4-05-TR26, Apr. 2005.
- [DS05b] DEY T. K., SUN J.: *Extremal Surface Based Projections Converge and Reconstruct with Isotopy*. Tech. Rep. OSU-CISRC-05-TR25, Apr. 2005.
- [Kol05] KOLLURI R.: Provably good moving least squares. In *Proc. SoDA* (2005).
- [Lev03] LEVIN D.: Mesh-independent surface interpolation. In *Geometric Modeling for Scientific Visualization*, Brunnett, Hamann, Mueller, (Eds.). Springer-Verlag, 2003, pp. 37–49.
- [MKN*04] MÜLLER M., KEISER R., NEALEN A., PAULY M., GROSS M., ALEXA M.: Point based animation of elastic, plastic and melting objects. In *Proc. Symp. Comp. Anim.* (2004), pp. 141–151.
- [MN03] MITRA N. J., NGUYEN A.: Estimating surface normals in noisy point cloud data. In *Symp. Comp. Geom.* (2003), pp. 322–328.
- [PKKG03] PAULY M., KEISER R., KOBBELT L. P., GROSS M.: Shape modeling with point-sampled geometry. *ACM TOG* 22, 3 (2003), 641–650.
- [ZPKG02] ZWICKER M., PAULY M., KNOLL O., GROSS M.: Pointshop 3d: an interactive system for point-based surface editing. In *Proc. SIGGRAPH* (2002), pp. 322–329.