Provably Better Moving Least Squares

Ravikrishna Kolluri James F. O'Brien Jonathan R. Shewchuk
University of California, Berkeley

1 Introduction

We analyze a variant of the implicit *moving least squares* (MLS) algorithm proposed by Shen, O'Brien, and Shewchuk [4]. We show that under certain sampling conditions the surface reconstructed by the MLS algorithm is geometrically and topologically correct.

The input to the MLS algorithm is a set of sample points S near a surface F, with approximate normals. For each sample $s \in S$ we define a linear point function that approximates the signed distance function of F in the local neighborhood of s. These functions are blended together using Gaussian weight functions yielding a smooth function I whose zero set U is the reconstructed surface. We prove that I is a good approximation to the signed distance function of the sampled surface F, and that U is homeomorphic to F and geometrically close to F.

Shen, O'Brien, and Shewchuk originally proposed their MLS construction with different weight functions for building manifold surfaces from polygon soup. Kolluri [3] showed that for reconstructing surfaces from points sets, a variant of this algorithm is geometrically and topologically correct under uniform sampling conditions. In this work we extend the analysis to handle adaptively sampled point data in which the sampling density is proportional to the local surface complexity. Our sampling requirements defined in Section 2, are similar to the sampling requirements of Delaunay-based algorithms like Crust [1].

2 Sampling Requirements

The *local feature size* (lfs) at a point $p \in F$ is the distance from p to the nearest point of the medial axis of F, as shown in Figure 1. S is an ϵ-sample of the surface F if the distance from any point $p \in F$ to its closest sample in S is less than $\epsilon\,\mathrm{lfs}(p)$. Our results are valid for values of $\epsilon \leq 0.01$.

Amenta and Bern [1] show that the function lfs is 1-Lipschitz. We extend the definition of the function lfs beyond the points on the surface F. This extension is used in defining our sampling requirements and our MLS construction. We define the *extended local feature size* of a point p as

$$\mathrm{elfs}(x) = \min_{p \in F}\{\mathrm{lfs}(p) + d(x,p) - |\phi(x)|\}.$$

Here, $\phi(x)$ is the signed distance from x to the surface F and $d(x,p)$ is the distance between point x and point p. It is easy to show that the function elfs is 1-Lipschitz and reduces to the function lfs for points on the surface.

Observation 1 *For any two points, p and q, $|\mathrm{elfs}(p) - \mathrm{elfs}(q)| \leq d(p,q)$. For any point $p \in F$, $\mathrm{elfs}(p) = \mathrm{lfs}(p)$.*

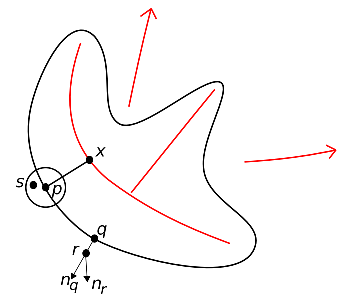

Figure 1: A closed curve along with its medial axis. The local feature size of p is the distance to the closest point x on the medial axis.

Our sampling requirements allow for noisy data when the amount of noise in the sample coordinates is small compared to the sample spacing. We assume that for each sample s, the distance to its closest surface point $p \in F$ is less than $\epsilon^2\mathrm{elfs}(s)$. We also allow a small amount of noise in the estimated sample normal. Consider a sample r with estimated normal $\vec{n}_r$, as shown in Figure 1, whose closest point in F is q with true normal $\vec{n}_q$. The angle between $\vec{n}_r$ and $\vec{n}_q$ should be less than ϵ.

Our MLS construction builds the function I by blending together functions associated with each sample point. Hence arbitrary oversampling in one region of the surface can distort the value of the function in other regions. To prohibit such oversampling, we require that local changes in the sampling density be bounded. Let α be the number of samples inside a ball of radius $\epsilon\,\mathrm{elfs}(p)$ centered at a point p. If $\alpha > 0$, the number of samples inside a ball of radius $2\epsilon\,\mathrm{elfs(p)}$ at p is at most 8α.

3 Surface Definition

The input to the MLS algorithm is a set of sample points S near the surface F. Each sample $s \in S$ has an associated vector $\vec{n}_s$ that approximates the outside normal of the surface near s.

We build a point function for each sample $s \in S$ that approximates the signed distance function of F near s. The point function $P_s(x)$ of sample point s with normal $\vec{n}_s$ is the signed distance from x to the tangent plane at s, $P_s(x) = (x - s) \cdot \vec{n}_s$. A weighted average of the point functions gives the function I whose zero set is the implicit surface we seek.

$$I(x) = \frac{\sum_{s \in S} W_s(x)((x-s)\cdot\vec{n}_s)}{\sum_{s \in S} W_s(x)}.$$

The weight functions are Gaussian functions modified by a

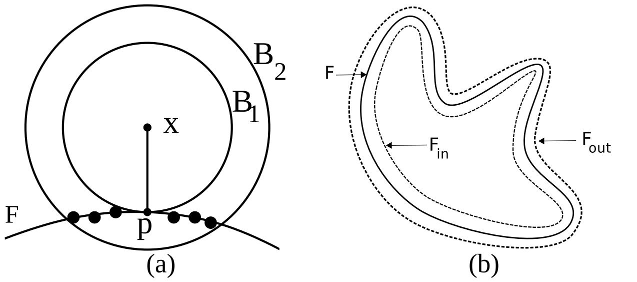

Figure 2: (a)The function I at point x is mostly determined by the point functions of the samples inside the thin shell bounded by B_1 and B_2. (b) The offset curves F_{in} and F_{out} of a curve F.

normalization factor associated with each sample point.

$$W_s(x) = e^{-\|x-s_i\|^2/\text{elfs}^2(x)}/A_s.$$

The normalization factor associated with each sample point s accounts for oversampling near s. Let $\alpha > 0$ be the number of samples inside a ball B_ϵ of radius $\epsilon\,\text{elfs}(s)$ centered at sample s, including s itself. The value of A_s is given by

$$A_s = \frac{\alpha}{\epsilon^3 \text{elfs}^3(s)}.$$

4 Results

Consider a point x whose closest point on the surface is p as shown in Figure 2(a). Let B_1(x) be a ball of radius $|\phi(x)|$ centered at x. Consider a second ball $B_2(x)$, that is slightly bigger than $B_1(x)$, also centered at x. The radius of $B_2(x)$ is $|\phi(x)| + \tau\,\text{elfs}(x)$. Here $\tau = 2\epsilon$ is a constant that depends on the sampling density. Our results are based on the observation that the value of the function at point x is mostly determined by the samples inside the thin shell bounded by $B_1(x)$ and $B_2(x)$.

Let F_{out} be the τ-offset surface outside of F that is obtained by moving each point $p \in F$ along the normal at p by a distance $\tau \cdot \text{elfs}(p)$. Similarly, let F_{in} be the τ-offset surface inside of F as shown in Figure 2(b). The τ-neighborhood is the region bounded by the inside and the outside offset surfaces. Our first geometric result is that the zero set U of I is inside the τ-neighborhood of F.

Theorem 2 *For each point x outside F_{out}, $I(x) > 0$ and for each point y inside F_{in}, $I(y) < 0$.*

Theorem 2 proves that the function I does not have any spurious zero crossings far away from the sample points. Our second geometric result is about the gradient of I at points in the zero set of I.

Theorem 3 *Let x be a point in the τ-neighborhood of F and let p be the point on F closest to x. Let $\vec{n}$ be the normal of p. Then, $\vec{n} \cdot \nabla I(x) > 0$.*

Theorem 3 proves that the gradient can never be zero inside the τ-neighborhood. From Theorem 2, the zero set of I is inside the τ-neighborhood of F. Hence, from the *implicit function theorem* [2], zero is a *regular* value of I and the zero set U is a compact, two-dimensional manifold.

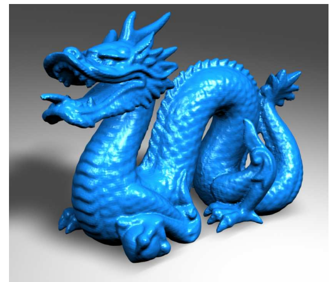

Figure 3: MLS reconstruction of the Stanford Dragon model from raw data.

We use these geometric results to define a homeomorphism between F and U. As F and U are compact, a one-to-one, onto, and continuous function from U to F defines a homeomorphism.

Definition: Let $\Gamma : \mathbb{R}^3 \to F$ map each point $q \in \mathbb{R}^3$ to its closest point on F.

Theorem 4 *The restriction of Γ to U defines a homeomorphism from U to F.*

5 Discussion

Our sampling requirement that $\epsilon \leq 0.01$ is probably an artifact of our proof technique. The MLS algorithm works quite well on data obtained from laser range, for which ϵ is much larger than 0.01 as shown in Figure 3.

Our definition of the MLS surface requires knowledge of the $\text{elfs}(x)$, function which is unknown. In our analysis, elfs can be replaced by any 1-Lipschitz function f such that $f(x) \leq \text{elfs}(x)$ at all points x, and the input sample is an ϵf-sample for $\epsilon \leq 0.01$. We can relax our requirements and assume that the elfs function is known only at the sample points. A 1-Lipschitz function function $f(x)$ can now be defined as

$$f(x) = \min_{s \in S}\{d(x,s) + \text{elfs(s)} - d(x, \text{nn}(x))\},$$

where $\text{nn}(x)$ is the sample nearest x in S.

References

[1] Nina Amenta and Marshall Bern. Surface Reconstruction by Voronoi Filtering. *Discrete & Computational Geometry*, 22:481–504, 1999.

[2] Jules Bloomenthal, editor. *Introduction to Implicit Surfaces*. Morgan Kaufman, 1997.

[3] Ravikrishna Kolluri. Provably Good Moving Least Squares. In *ACM-SIAM Symposium on Discrete Algorithms*, 2005.

[4] Chen Shen, James F. O'Brien, and Jonathan R. Shewchuk. Interpolating and Approximating Implicit Surfaces from Polygon Soup. *ACM Transactions on Graphics*, 23(3):896–904, 2004.