Eurographics Symposium on Geometry Processing (2005)
M. Desbrun, H. Pottmann (Editors)

An Adaptive MLS Surface for Reconstruction with Guarantees

Tamal K. Dey and Jian Sun †

Abstract

Recent work have shown that moving least squares (MLS) surfaces can be used effectively to reconstruct surfaces from possibly noisy point cloud data. Several variants of MLS surfaces have been suggested, some of which have been analyzed theoretically for guarantees. These analyses, so far, have assumed uniform sampling density. We propose a new variant of the MLS surface that, for the first time, incorporates local feature sizes in its formulation, and we analyze it for reconstruction guarantees using a non-uniform sampling density. The proposed variant of the MLS surface has several computational advantages over existing MLS methods.

Categories and Subject Descriptors (according to ACM CCS): I.3.3 [Computer Graphics]: Line and Curve Generation

1. Introduction

In surface reconstruction, if the input point cloud is noisy, a surface fitting through the points can be too bumpy for practical use. A remedy to this problem is to define a target smooth surface and project or generate points on this smooth surface for reconstruction. Of course, the main problem is to choose a suitable smooth surface that resembles the original surface which the input point cloud presumably sampled. Several such target surfaces have been proposed recently with different algorithms for their computations. The radial basis function of Carr et al. [JCCCM*01], the multi-level partition of unity of Ohtake et al. [OBA*03], the natural neighbor surface of Boissonnat and Cazasls [BC00] and the moving least squares of Alexa et al. [ABCO*01] are examples of such surfaces, to name a few.

The moving least squares surfaces (MLS), originally proposed by Levin [Lev98] and later adopted by Alexa et al. [ABCO*01] for reconstruction have been widely used for modeling and rendering [AA03, MVdF03, PKKG03]. The popular open source software PointShop 3D [ZPKG02] implements the MLS surfaces and shows the effectiveness of the MLS surfaces on real world scanned data. Recently, the work of Amenta and Kil [AK04] and Kolluri [Kol05] have broaden the understanding of the MLS surfaces. Kolluri considered an implicit version of the MLS surfaces and proved, for the first time, theoretical guarantees about them. Subsequently, Dey, Goswami and Sun [DGS05] proved similar guarantees for the variant of MLS proposed by Amenta and Kil. These theoretical results assume a globally uniform sampling density which is quite restrictive. For example, in Figure 1, the globally uniform sampling condition needs more than 10^4 points to sample the arc *ab* because of a small feature somewhere else. Our aim is to prove guarantees about MLS surfaces under an adaptive sampling condition similar to the one used by Amenta and Bern [AB99] in the context of noise-free reconstruction. Under such adaptive sampling, one needs only 6 points to sample the arc in Figure 1. To accommodate an adaptive sampling, we come up with a new variant of the MLS surfaces, which incorporates the local feature sizes of the sampled surface in its formulation. Our results show that this new MLS surface has several advantages over other existing MLS methods.

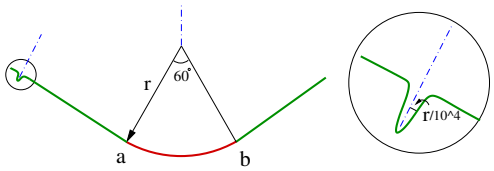

Figure 1: *The dash-dot lines represent the medial axis.*

† Dept. of CSE, The Ohio State University, Columbus, OH 43210.
{tamaldey,sunjia}@cse.ohio-state.edu

1.1. MLS surfaces

Moving least squares is an extension of the well known least squares technique to fit a surface to a given set of data points. The term 'moving' refers to the various weighting of the points in calculating their contributions to the solution at different locations. This unified view of the moving least squares is nicely explained by Shen et al. [SOS04]. There are mainly two types of MLS surfaces considered in the literature in the context of surface reconstructions.

1.1.1. Implicit MLS

Shen et al.[SOS04] defined a MLS surface implicitly by a function which they called IMLS surface. This surface is specified by the moving least-squares solution to a set of constraints that force the function to assume given values at the samples and also force its upward gradient to match the assigned normals at the samples. Each constraint is associated with a weight function. In the simplest case, the implicit function can be taken as

$$\mathcal{I}(x) = \frac{\sum_{p\in P}[(x-p)^T v_p]\theta_p(x)}{\sum_{p\in P}\theta_p(x)} \qquad (1.1)$$

where θ_p is a weighting function and v_p is the normal assigned to a sample p. Kolluri [Kol05] considered this surface and showed that the IMLS surface is indeed isotopic (stronger than homeomorphic) to the sampled surface if the input sample is sufficiently dense.

1.1.2. Projection MLS

Levin [Lev98] pioneered a MLS surface that is defined as a set of stationary points of a projection operator. We call this projection based MLS surfaces as PMLS surfaces. Amenta and Kil [AK04] gave a more explicit definition for PMLS surfaces as the local minima of an energy function along the directions given by a vector field. Based on this explicit definition, they gave an implicit form for PMLS surfaces. Specifically, they showed that the PMLS surface defined by Levin [Lev98] is actually the implicit surface given by the zero-level set of the implicit function

$$\mathcal{J}(x) = n(x)^T\left(\frac{\partial \mathcal{E}(y,n(x))}{\partial y}|_x\right)$$

where $n\colon \mathbb{R}^3 \to \mathbb{R}^3$ is a given vector field, $\mathcal{E}\colon \mathbb{R}^3\times\mathbb{R}^3 \to \mathbb{R}$ is an energy function given by $\mathcal{E}(y,n(x)) = \frac{1}{2}\sum_{p\in P}[(y-p)^T n(x)]^2\theta_p(y)$ with θ_p as a weighting function. If θ_p is a Gaussian with width h then

$$\mathcal{J}(x) = \sum_{p\in P}(x-p)^T n(x)[1-(\frac{(x-p)^T n(x)}{h})^2]\theta_p(x). \quad (1.2)$$

The PMLS surface definition inherently leads to a projection method by which points can be projected onto the surface. Dey, Goswami and Sun [DGS05] prove theoretical guarantees for PMLS surfaces.

1.1.3. Variation of PMLS (VMLS)

If we define the energy function as $\mathcal{E}(y,n(x)) = \frac{1}{2}\sum_{p\in P}[(y-p)^T n(x)]^2\theta_p(x)$ where the weighting function θ_p varies with x instead of y, we obtain a variant definition of PMLS surfaces, which we call VMLS in short. Indeed, this is the MLS surface actually implemented in PointShop 3D by Zwicker et al. [ZPKG02]. It has a very simple implicit form

$$\mathcal{G}(x) = \sum_{p\in P}[(x-p)^T n(x)]\theta_p(x). \qquad (1.3)$$

An advantage of this definition is that, unlike the standard PMLS surfaces, its inherent projection procedure does not require any non-linear optimization, which makes the algorithm fast, stable and easy to implement.

1.1.4. Results

We adopt the IMLS form in equation 1.1 and modify it to be adaptive to the feature sizes of the sampled surface. This enables us to prove guarantees under an adaptive sampling condition as opposed to the uniform one of Kolluri [Kol05] and Dey et al. [DGS05]. The modification is in the choice of the weighting function for defining the implicit function $\mathcal{I}$.

The weighting function θ_p for defining $\mathcal{I}$ is chosen to be a Gaussian whose width depends on the local feature sizes as defined by Amenta and Bern [AB99]. The particular choice of this feature dependence is new and reasoned in section 3.1. We call the MLS surface given by $\mathcal{I}$ with these feature dependencies the *adaptive* MLS surface or AMLS surface in short. As an implicit surface, one can project a point onto AMLS surfaces by the Newton iteration method which strikingly outperforms the well known projection procedure for PMLS or VMLS surfaces.

We prove guarantees for AMLS surfaces under an *adaptive sampling condition* and note its advantages over other MLS surfaces. Specifically our results can be summarized as follows.

(i) The subset W of $\mathcal{I}^{-1}(0)$ onto which points are projected is indeed isotopic to the sampled surface, i.e., one can be continuously deformed to the other always maintaining a homeomorphism between them.

(ii) The above guarantee requires that the assigned vectors to the samples closely approximate the normals of the sampled surface. We present a provable algorithm to estimate the normals at the samples even when they are noisy.

(iii) We present an implementation to approximate the AMLS surface which establishes its effectiveness and discuss several of its advantages over other MLS surfaces. Specifically, our results show that the standard Newton iteration used for projection on AMLS surfaces is faster than the projections used for PMLS surfaces.

The proofs of some of our results cannot be included in the ten-page limit of the paper. An extended version of this paper including all the missing proofs is available from the authors' web pages [DS05].

2. Preliminaries

2.1. Surface and thickening

Let $\Sigma \subset \mathbb{R}^3$ be a compact C^2 smooth surface without boundary. For simplicity assume that Σ has a single connected component. Let Ω_I and Ω_O denote the bounded and unbounded components of $\mathbb{R}^3 \setminus \Sigma$ respectively. For a point $z \in \Sigma$, let $\tilde{n}_z$ denote the oriented normal of Σ at z where $\tilde{n}_z$ points locally toward the unbounded component Ω_O. Let a line through a point $z \in \Sigma$ along the normal $\tilde{n}_z$ be denoted $\ell_{z,\tilde{n}(z)}$.

For a point $x \in \mathbb{R}^3$ and a set $X \subset \mathbb{R}^3$, let $d(x,X)$ denote the distance of x to X, i.e., $d(x,X) = \inf_{y \in X} \|x-y\|$. The medial axis M of Σ is the closure of the set $Y \subset \mathbb{R}^3$ where for each $y \in Y$ the distance $d(y,\Sigma)$ is realized by two or more points. In other words, M is the locus of the centers of the maximal balls whose interiors are empty of any point from Σ. The local feature size at a point $z \in \Sigma$, $\mathrm{lfs}(z)$, is defined as the distance $d(z,M)$ of z to the medial axis M. The function $\mathrm{lfs}()$ is 1-Lipschitz.

Let $\nu : \mathbb{R}^3 \to \Sigma$ map a point $x \in \mathbb{R}^3$ to its closest point in Σ. It is known that ν is well defined if its domain avoids M which will be the case for our use of ν. Denote $\tilde{x} = \nu(x)$. Let $\phi(x)$ denote the signed distance of a point x to Σ, i.e., $\phi(x) = (x-\tilde{x})^T \tilde{n}_{\tilde{x}}$. For a real $\delta \geq 0$, define offset surfaces $\Sigma_{+\delta}$ and $\Sigma_{-\delta}$ where

$$\Sigma_{+\delta} = \{x \in \mathbb{R}^3 \,|\, \phi(x) = +\delta \mathrm{lfs}(\tilde{x})\}$$

$$\Sigma_{-\delta} = \{x \in \mathbb{R}^3 \,|\, \phi(x) = -\delta \mathrm{lfs}(\tilde{x})\}.$$

Let $\delta\Sigma$ be the region between $\Sigma_{-\delta}$ and $\Sigma_{+\delta}$, i.e.,

$$\delta\Sigma = \{x \in \mathbb{R}^3 \,|\, -\delta \mathrm{lfs}(\tilde{x}) \leq \phi(x) \leq \delta \mathrm{lfs}(\tilde{x})\},$$

Figure 2 illustrates the above concepts.

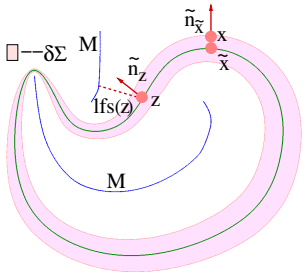

Figure 2: *The set* $\delta\Sigma$*, medial axis, normals and* $\mathrm{lfs}()$.

2.2. Sampling

An adaptive sampling density based on the local feature size called ε-sampling has been used for proving the correctness of several surface reconstruction algorithms [AB99, ACDL02, BC00]. For this work we assume a similar sampling condition but modify it for possible noisy input. Apart from condition (iii) justification of all other conditions can be found in [DG04]. We say $P \subset \mathbb{R}^3$ is a *noisy* (ε,α)*-sample* of the surface Σ if the following sampling conditions hold.

(i) The distance from each point $z \in \Sigma$ to its closest sample is less than $\varepsilon \mathrm{lfs}(z)$.
(ii) The distance from each sample $p \in P$ to its closest point $\tilde{p}$ on Σ is less than $\varepsilon^2 \mathrm{lfs}(\tilde{p})$.
(iii) Each sample p is equipped with a normal v_p where the angle between v_p and the normal $\tilde{n}_{\tilde{p}}$ at its closest point $\tilde{p}$ on Σ is less than ε.
(iv) The number of the samples inside $B(x, \varepsilon \mathrm{lfs}(\tilde{x}))$ is less than a small number α, for any $x \in \mathbb{R}^3$. In this paper α is set to be 5.

Obviously the less the ε is, the better Σ is sampled by P. Usually P is considered to be a good sample of the surface if $\varepsilon \leq 0.1$.

For our proofs we need a result that all samples near a point z on Σ lie within a small slab centering z. Kolluri [Kol05] made a similar observation assuming uniform sampling density. Here we extend it for adaptive sampling. A ball with the center x and radius r is denoted as $B(x,r)$. Denote $S(x,r)$ to be the boundary of $B(x,r)$. Consider any point z on Σ and a ball $B(z, \rho \mathrm{lfs}(z))$ with a small radius, i.e., $\rho > 0$ is small. Let PL_+ and PL_- be two planes perpendicular to $\tilde{n}_z$ and at a small distance $\omega \mathrm{lfs}(z)$ from z. We show that if ω is of the order of $\varepsilon^2 + \rho^2$, all points of P within the ball $B(z, \rho \mathrm{lfs}(z))$ lie within the slab made by PL_+ and PL_-.

Lemma 1 For $\rho \leq 1$ and $\varepsilon \leq 0.1$, any sample inside $B(z, \rho \mathrm{lfs}(z))$ lies inside the slab bounded by the planes PL_+ and PL_- where

$$\omega = \frac{(\varepsilon^2+\rho)^2}{2(1-\varepsilon^2)^2} + \frac{(1+\rho)}{1-\varepsilon^2}\varepsilon^2.$$

In addition we show that a small ball $B(x, \frac{\rho}{2}\mathrm{lfs}(\tilde{x}))$ centering any point x in $\mathbb{R}^3$ contains a small number of points from P.

Lemma 2 For $\rho \leq 0.4$ and $\varepsilon \leq 0.1$, the number of samples inside $B(x, \frac{\rho}{2}\mathrm{lfs}(\tilde{x}))$ is less than λ where

$$\begin{aligned} \lambda &= \alpha \quad \text{if } \rho \leq 2\varepsilon \\ &= \frac{75\rho^3\alpha}{\varepsilon^3} \quad \text{otherwise.} \end{aligned}$$

3. Definition of AMLS

3.1. Weighting functions

The implicit function value $\mathcal{I}(x)$ at a point x should be primarily decided by the nearby samples. That is exactly the reason why different implicit MLS functions proposed so far weigh the samples differently in a sum instead of giving them equal weights. We have already seen that the samples within a sufficiently small neighborhood are predictably

distributed within a small slab (Lemma 1). However the surface outside the neighborhood could vary arbitrarily without changing its part inside. As a result, the samples outside the neighborhood could be arbitrarily distributed. Hence we should design a weighting function such that the samples outside the neighborhood have much less effect on the implicit function than those inside.

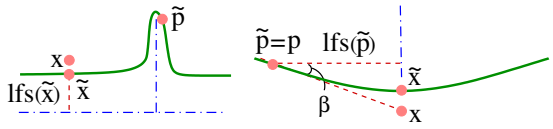

Figure 3: *The solid curves and the dash-dot lines represent part of the surface and its medial axis respectively.*

Gaussian functions are a good choice to meet the above requirements since their widths can control the influence of the samples. To make the implicit surface sensitive to features of the original sampled surface, we take the width of the Gaussian function to be a fraction of the local feature size. However, one needs to be more careful. If we simply take a fraction of $\mathrm{lfs}(\tilde{x})$ as the width, i.e., take $e^{-\frac{\|x-p\|^2}{[\rho \mathrm{lfs}(\tilde{x})]^2}}$ as the weighting function where $\rho < 1$, we cannot bound the effect of the far away samples. Consider the left picture in Figure 3. The local feature size at the point $\tilde{p}$ can be arbitrarily small and hence the number of samples around $\tilde{p}$ needs to be arbitrarily large to meet the sampling conditions. Consequently, the summation of the weights over those samples which are outside $B(x, \mathrm{lfs}(\tilde{x}))$ becomes too large to be dominated by the contributions of the samples in the neighborhood of x.

An alternative option is to take a fraction of $\mathrm{lfs}(\tilde{p})$ as the width, i.e., take $e^{-\frac{\|x-p\|^2}{[\rho \mathrm{lfs}(\tilde{p})]^2}}$ as the weighting function. However it also fails as illustrated in the right picture in Figure 3. The samples such as p has a constant weight $e^{-\frac{1}{[\rho \cos\beta]^2}}$. As the summation extends outside the neighborhood of x, the contribution of the samples remains constant instead of decreasing. As a result, one cannot hope to bound the outside contribution.

We overcome the difficulty by a novel combination of the above two options, i.e, by taking a fraction of $\sqrt{\mathrm{lfs}(\tilde{x})\mathrm{lfs}(\tilde{p})}$ as the width of the Gaussian weighting functions. This takes into account the effects from both members, the contribution sender p and the contribution receiver x. Unlike $e^{-\frac{\|x-p\|^2}{[\rho lfs(\tilde{p})]^2}}$, such form of weighting function decreases as p goes far away from x no matter how the surface looks like. In addition, such form of weighting function assigns a small value to the points that sample small features, which in turn cancels out the effect that small features require more samples.

There is still one more difficulty. The function lfs(), though continuous, is not smooth everywhere on Σ. The non-smoothness appears where Σ intersects the medial axis of its own medial axis M. To make the implicit function smooth, we use a smooth function $f()$ arbitrarily close to lfs() where

$$|f(x) - \mathrm{lfs}(x)| \leq \beta \mathrm{lfs}(x) \tag{3.4}$$

for arbitrarily small $\beta > 0$, say 10^{-100}. This is doable since the family of real valued smooth functions over smooth manifolds is dense in the family of continuous functions [Hir88] and the minimal feature size is strictly positive for any C^2 manifold [Wol92]. Finally we choose a fraction (given by ρ_e) of $\sqrt{f(\tilde{x})f(\tilde{p})}$ as the width of the Gaussian weighting functions. Specifically we take

$$\ln \theta_p(x) = -\frac{\sqrt{2}\|x-p\|^2}{\rho_e^2 f(\tilde{p}) f(\tilde{x})}. \tag{3.5}$$

The factor $\sqrt{2}$ in the exponent is for the convenience in proofs as we will see later. In general, it is known that larger values of ρ_e make the MLS surface smoother. To have a sense of appropriate values of ρ_e, consider the case where x is on the surface Σ. The samples such as p in Figure 4 across the medial axis to point x should have little effect on the implicit function value at x. Taking $\rho_e \leq 0.4$ makes the weight of p at x less than $e^{-25\sqrt{2}} \approx 5 \times 10^{-16}$ since $\|x-p\| \geq 2max\{\mathrm{lfs}(\tilde{x}), \mathrm{lfs}(\tilde{p})\}$.

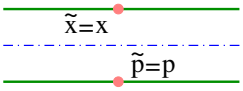

Figure 4: *The solid and the dash-dot lines represent part of the surface and its medial axis respectively.*

3.2. AMLS function

Define

$$\mathcal{N}(x) = \sum_{p \in P} [(x-p)^T v_p] \theta_p(x) \tag{3.6}$$

and

$$\mathcal{I}(x) = \frac{\mathcal{N}(x)}{\mathcal{W}(x)} \tag{3.7}$$

where $\mathcal{W}(x) = \sum_{p \in P} \theta_p(x)$. Obviously the implicit functions $\mathcal{N}$ and $\mathcal{I}$ have exactly the same 0-level set, i.e., $\mathcal{I}^{-1}(0) = \mathcal{N}^{-1}(0)$. Therefore, we could have taken $\mathcal{N}$ instead of $\mathcal{I}$ for AMLS, but we show in section 8 that $\mathcal{I}$ has a significant computational advantage since Newton iteration for $\mathcal{I}$ has a much larger convergent domain than the one for $\mathcal{N}$. However, the function $\mathcal{N}$ has a simpler form to analyze. Hence, we analyze the 0-level set of $\mathcal{I}$ via the function $\mathcal{N}$.

4. Contributions of distant samples

Theorem 1 is one of our main contributions in this paper, which says that the effect of distant samples on a point can be bounded. Once we prove this, the overall technique of Kolluri [Kol05] for IMLS surfaces with uniform sampling or of Dey et al. [DGS05] for PMLS surfaces with uniform sampling can be applied. However, the details need careful work.

In various claims, the contribution of a sample p to the implicit function or its derivative at point x will be bounded from above by an expression that involves the term

$$I_p(x) = e^{-\frac{\sqrt{2}\|x-p\|^2}{\rho^2 f(\tilde{p}) f(\tilde{x})}} \cdot \frac{\|x-p\|^s}{[\frac{\rho^2}{\sqrt{2}} f(\tilde{p}) f(\tilde{x})]^t}.$$

The values of s and t will vary between 0 to 2 and from 0 to 1 respectively in various equations where I_p is used. For instance, the contribution of a sample p to function $\mathcal{N}$ at x can be bounded by $I_p(x)$ with $s = 1$ and $t = 0$.

Following Kolluri [Kol05] our strategy for bounding $I_p(x)$ will be to decompose the space into spherical shells centering x. Theorem 1 shows that the total contribution from all the samples in the shells decreases as their distances from x increase. Let $\mathbb{S}_x(w,\rho)$ be the shell region between the spheres $S(x, w\text{lfs}(\tilde{x}))$ and $S(x,(w+\rho)\text{lfs}(\tilde{x}))$. For $i = 0, 1, \ldots$ consider the nested shells given by $\mathbb{S}_x(w_i,\rho)$ where $w_i = r + i\rho$ (Figure 5). To prove Theorem 1 we need a result that bounds the total contribution of the samples lying within the intersection of a small ball of radius $\frac{\rho}{2}\text{lfs}(\tilde{x})$ and the shell $\mathbb{S}_x(w_i,\rho)$. Let $B_{\frac{\rho}{2}}$ be any such ball. We would like to bound the sum $\sum_{p \in B_{\frac{\rho}{2}} \cap \mathbb{S}_x(w_i,\rho)} I_p(x)$. The ball $B_{\frac{\rho}{2}}$ has a radius $\frac{\rho}{2}\text{lfs}(\tilde{x})$ though its center is not necessarily x. Therefore, one cannot use Lemma 2 to bound the number of samples inside $B_{\frac{\rho}{2}}$. We overcome this difficulty by using a hierarchical subdivision of the bounding box NC_1 of $B_{\frac{\rho}{2}}$. The subdivision divides a cube unless it can be covered with a ball $B(c,r)$ where r is a fraction of $\text{lfs}(\tilde{c})$. Then, one can call upon Lemma 2 to bound the number of samples in $B(c,r)$ and hence in the cubes of the subdivision. Therefore, we can bound the number of samples in $B_{\frac{\rho}{2}}$ using the number of the leaf nodes in its corresponding subdivision tree. Notice that we do not have an explicit bound for the number of samples in any $B_{\frac{\rho}{2}}$ since at different positions $B_{\frac{\rho}{2}}$ may have different subdivision trees adapting to the local geometry of the surface. However, we do have an explicit upper bound for the total weights from the samples inside any $B_{\frac{\rho}{2}}$ as proved in Lemma 3.

Assume a hierarchical subdivision tree HST of NC_1 as follows. Let c_1 be the center of the bounding cube NC_1. Subdivide NC_1 into 27 subcubes of size $\frac{\rho}{3}\text{lfs}(\tilde{x})$ if $\text{lfs}(\tilde{c}_1) < \text{lfs}(\tilde{x})$. Let NC_2 be any such subcube. It can be covered by a ball $B_{\frac{\rho}{2^2}} = B(c_2, \frac{\rho}{2^2}\text{lfs}(\tilde{x}))$ where c_2 is the center of NC_2. Subdivide NC_2 in the same way if $\text{lfs}(\tilde{c}_2) < \frac{1}{2}\text{lfs}(\tilde{x})$. In general, keep subdividing a subcube NC_k at the kth level if $\text{lfs}(\tilde{c}_k) < \frac{1}{2^{k-1}}\text{lfs}(\tilde{x})$ where c_k is the center of NC_k. Observe that NC_k is covered by $B_{\frac{\rho}{2^k}} = B(c_k, \frac{\rho}{2^k}\text{lfs}(\tilde{x}))$. Figure 5 shows a HST in 2D case. We use NC_k also to denote its intersection with $B_{\frac{\rho}{2^k}}$.

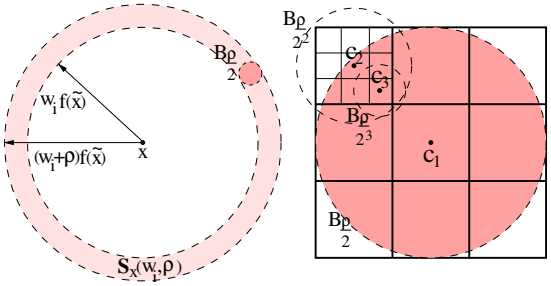

Figure 5: *The nested shells and the hierarchical subdivision tree*

Lemma 3 If $\rho \le 0.4$, $\varepsilon \le 0.1$ and $r \ge 5\rho$,

$$\sum_{p \in B_{\frac{\rho}{2}} \cap \mathbb{S}_x(w_i,\rho)} I_p(x) \le \lambda e^{-\frac{rw_i}{(1+2r)\rho^2}} \cdot \frac{w_i^s}{\rho^{2t}} f(\tilde{x})^{s-2t}$$

where $0 \le s \le 2$, $0 \le t \le 1$ and λ is defined in Lemma 2.

Proof Case 1: $\text{lfs}(\tilde{c}_1) \ge \text{lfs}(\tilde{x})$: HST has only one node NC_1. Let p be any sample in $B_{\frac{\rho}{2}}$. Observe that $\|\tilde{p} - \tilde{c}_1\| \le 2\|\tilde{p} - c_1\| \le 2(\|\tilde{p} - p\| + \|p - c_1\|) \le 2\varepsilon^2\text{lfs}(\tilde{p}) + \rho\text{lfs}(\tilde{c_1})$. By Lipschitz property of lfs(),

$$\text{lfs}(\tilde{p}) \ge \frac{1-\rho}{1+2\varepsilon^2}\text{lfs}(\tilde{x}).$$

From inequality 3.4 we have

$$f(\tilde{p}) \ge \frac{1-\rho}{\beta'(1+2\varepsilon^2)} f(\tilde{x})$$

where $\beta' = \frac{1+\beta}{1-\beta}$. Similarly from condition $\|x-p\| \ge r\text{lfs}(\tilde{x})$ (p lies in $\mathbb{S}_x(w_i,\rho)$) and the fact $\|\tilde{x} - \tilde{p}\| \le 2(\|x-p\| + \|p - \tilde{p}\|)$ we obtain

$$f(\tilde{p}) \le (1+\beta)\frac{1+2r}{r(1-2\varepsilon^2)}\|x-p\|.$$

Hence

$$I_p(x) \le e^{-\frac{\sqrt{2}(1-2\varepsilon^2)}{(1+\beta)(1+2r)}\frac{r\|x-p\|}{\rho^2 f(\tilde{x})}} \cdot [\frac{\sqrt{2}\beta'(1+2\varepsilon^2)}{1-\rho}]^t \cdot \frac{\|x-p\|^s}{[\rho f(\tilde{x})]^{2t}}$$

which is a decreasing function of $\|x-p\|$ when $\|x-p\| \ge 4\rho\text{lfs}(\tilde{x})$. Since $\|x-p\| \ge \frac{w_i f(\tilde{x})}{1+\beta}$, we have

$$\begin{aligned} I_p(x) &\le e^{-\frac{\sqrt{2}(1-2\varepsilon^2)}{(1+\beta)^2(1+2r)}\frac{rw_i}{\rho^2}} \cdot \frac{[\sqrt{2}\beta'(1+2\varepsilon^2)]^t}{(1-\rho)^t(1+\beta)^s} \cdot \frac{w_i^s}{\rho^{2t}} f(\tilde{x})^{s-2t} \\ &\le e^{-\frac{rw_i}{(1+2r)\rho^2}} \cdot \frac{w_i^s}{\rho^{2t}} f(\tilde{x})^{s-2t}. \end{aligned}$$

It is not hard to verify the second inequality under the

given conditions. The lemma follows from the fact that $B(c_1, \frac{\rho}{2} f(\tilde{c}_1))$ covers $B_{\frac{\rho}{2}}$ and hence the number of samples inside $B_{\frac{\rho}{2}}$ is less than λ from Lemma 2.

Case 2: $\text{lfs}(\tilde{c}_1) < \text{lfs}(\tilde{x})$: Consider a leaf node NC_k at the kth level which is covered by $B_{\frac{\rho}{2^k}}$ in HST. We have $\text{lfs}(\tilde{c}_k) \geq \frac{1}{2^{k-1}}\text{lfs}(\tilde{x})$. Let p be any sample inside the node. Since $\|\tilde{p} - \tilde{c}_k\| \leq 2\|\tilde{p} - c_k\|$, we obtain

$$f(\tilde{p}) \geq \frac{1-\rho}{\beta'(1+2\varepsilon^2)} \cdot \frac{1}{2^{k-1}} f(\tilde{x}).$$

On the other hand, p is also inside the parent node NC_{k-1} covered by $B_{\frac{\rho}{2^{k-1}}}$ in HST. Since $\|\tilde{p} - \tilde{c}_{k-1}\| \leq 2\|\tilde{p} - c_{k-1}\|$ and $\text{lfs}(\tilde{c}_{k-1}) < \frac{1}{2^{k-2}}\text{lfs}(\tilde{x})$, we obtain

$$f(\tilde{p}) \leq \frac{\beta'(1+\rho)}{1-2\varepsilon^2} \cdot \frac{1}{2^{k-2}} f(\tilde{x}).$$

Hence for the given value of ρ and ε, we have

$$\begin{aligned} & I_p(x) \\ & \leq e^{-2^{k-2}\frac{\sqrt{2}(1-2\varepsilon^2)}{\beta'(1+\rho)}\frac{\|x-p\|^2}{[\rho f(\tilde{x})]^2}} \cdot 2^{t(k-2)} \left[\frac{2\sqrt{2}\beta'(1+2\varepsilon^2)}{1-\rho}\right]^t \cdot \frac{\|x-p\|^s}{[\rho f(\tilde{x})]^{2t}} \\ & \leq \frac{1}{27} e^{-2^{k-2}\frac{rw_i}{(1+2r)\rho^2}} \cdot 2^{t(k-2)} \cdot \frac{w_i^s}{\rho^{2t}} f(\tilde{x})^{s-2t} \end{aligned}$$

Since $B(c_k, \frac{\rho}{2} f(\tilde{c}_k))$ covers $B_{\frac{\rho}{2^k}}$ and hence the number of samples inside the leaf node NC_k is less than λ from Lemma 2, we have

$$\sum_{p \in NC_k} I_p(x) \leq \frac{1}{27} \cdot \lambda e^{-2^{k-2}\frac{rw_i}{(1+2r)\rho^2}} \cdot 2^{t(k-2)} \cdot \frac{w_i^s}{\rho^{2t}} f(\tilde{x})^{s-2t} \tag{4.8}$$

The above equation gives the bound for contributions of samples inside a single leaf node NC_k at any level $k \geq 2$. We use induction to establish that the bound also holds for any *internal* node. Let NC_k be an internal node. Then, by induction we can assume that each of the 27 children of NC_k satisfy equation 4.8 with $k = k+1$. Summing over this 27 children and replacing k with $k+1$ in equation 4.8, we get

$$\begin{aligned} \sum_{p \in NC_k} I_p(x) \; & \leq \; \lambda e^{-2^{k-1}\frac{rw_i}{(1+2r)\rho^2}} \cdot 2^{t(k-1)} \cdot \frac{w_i^s}{\rho^{2t}} f(\tilde{x})^{s-2t} \\ & \leq \; \frac{1}{27} \cdot \lambda e^{-2^{k-2}\frac{rw_i}{(1+2r)\rho^2}} \cdot 2^{t(k-2)} \cdot \frac{w_i^s}{\rho^{2t}} f(\tilde{x})^{s-2t}. \end{aligned}$$

The lemma follows from the fact that 27 NC_2s partition $B_{\frac{\rho}{2}}$. □

Theorem 1 If $\rho \leq 0.4$, $\varepsilon \leq 0.1$ and $r \geq 5\rho$, then for any $x \in \mathbb{R}^3$

$$\sum_{p \notin B(x, rf(\tilde{x}))} I_p(x) \leq C_1\lambda \cdot \frac{r^2 + r\rho + \rho^2}{\rho^2} e^{-\frac{r^2}{(1+2r)\rho^2}} \cdot \frac{r^s}{\rho^{2t}} f(\tilde{x})^{s-2t}$$

where $0 \leq s \leq 2$, $0 \leq t \leq 1$ and $C_1 = 180\sqrt{3}\pi$.

Proof The space outside $B(x, r\text{lfs}(\tilde{x}))$ can be decomposed by $(\mathbb{S}_x(w_i, \rho))_{i=0}^{\infty}$ where $w_i = r + i\rho$. Each $\mathbb{S}_x(w_i, \rho)$ can be covered by less than $\frac{36\sqrt{3}\pi(w_i^2 + w_i\rho + \rho^2)}{\rho^2}$ balls of radius $\frac{\rho}{2}\text{lfs}(\tilde{x})$ as in [Kol05]. From Lemma 3 the contribution from the samples inside each of these balls are bounded. Hence

$$\begin{aligned} & \sum_{p \notin B(x, rf(\tilde{x}))} I_p(x) = \sum_{i=0}^{\infty} \sum_{p \in \mathbb{S}_x(w_i, \rho)} I_p(x) \\ & \leq \frac{C_1\lambda}{5} \sum_{i=0}^{\infty} \frac{w_i^2 + w_i\rho + \rho^2}{\rho^2} e^{-\frac{rw_i}{(1+2r)\rho^2}} \cdot \frac{w_i^s}{\rho^{2t}} f(\tilde{x})^{s-2t} \\ & \leq C_1\lambda \cdot \frac{r^2 + r\rho + \rho^2}{\rho^2} e^{-\frac{r^2}{(1+2r)\rho^2}} \cdot \frac{r^s}{\rho^{2t}} f(\tilde{x})^{s-2t}. \end{aligned}$$

The last inequality holds because the series is bounded from above by a geometric series with common ratio less than 0.8. □

5. Isotopy

Although we prove Theroem 1 with hypothesis $\rho \leq 0.4$ and $\varepsilon \leq 0.1$ which is plausible in practice, our proof for isotopy uses the setting $\varepsilon \leq 4 \times 10^{-3}$ and $\rho_e = \varepsilon$. The requirement for such small ε is probably an artifact of our proof technique. There are rooms to improve these constants though the proofs become more complicated (see the discussion in the extended version [DS05]). We focused more on demonstrating the ideas behind the proofs rather than tightening the constants. In our experiments, the AMLS surfaces work well on sparse data sets as we show in section 7.

Let $W = \mathcal{N}^{-1}(0) \bigcap 0.1\Sigma$, the subset of $\mathcal{N}^{-1}(0)$ inside 0.1Σ. Lemma 4 shows that W is indeed within $0.3\varepsilon\Sigma$. In addition, Lemma 6 implies that $\bigtriangledown\mathcal{N}$ cannot vanish in $0.3\varepsilon\Sigma$ and hence 0 is a regular value. So, by implicit function theorem W is a compact, smooth surface. Recall that $\nu\colon \mathbb{R}^3 \to \Sigma$ takes a point to its closest point on Σ. Let $\nu|_W$ be the restriction of ν to W. We prove that $\nu|_W$ is a homeomorphism. Since W is included in a topological thickening $0.3\varepsilon\Sigma$ of Σ and W separates the sides of $0.3\varepsilon\Sigma$, we also have W and Σ isotopic in $\mathbb{R}^3$ due to a result of Chazal and Cohen-Steiner [CCS04]. So, to prove isotopy we only need to prove that W and Σ are homeomorphic.

Theorem 2 $\nu|_W$ is a homeomorphism.

Proof The function $\nu|_W$ is continuous since ν is. Since W is compact, it is sufficient to show that $\nu|_W$ is surjective and injective which are the statements of Lemma 5 and Lemma 7 respectively. □

To prove that $\nu|_W$ is surjective we use the following lemma which says that $\mathcal{N}$ crosses zero within $0.3\varepsilon\Sigma$, i.e. $W \in 0.3\varepsilon\Sigma$.

Lemma 4

$$\begin{aligned} \mathcal{N}(x) \; & > 0 \quad \text{if } x \in (0.1\Sigma \setminus 0.3\varepsilon\Sigma) \cap \Omega_O \\ & < 0 \quad \text{if } x \in (0.1\Sigma \setminus 0.3\varepsilon\Sigma) \cap \Omega_I \end{aligned}$$

Lemma 5 $\nu|_W$ is surjective.

Proof Let z be any point in Σ. The normal line $\ell_{z,\tilde{n}(z)}$, through z along the normal $\tilde{n}_z$, intersects $\mathcal{N}^{-1}(0)$ within $0.3\varepsilon\Sigma$, thanks to Lemma 4. By definition of W, it intersects W at point x such that $\|x-z\| \leq 0.3\varepsilon \mathrm{lfs}(z)$. This means $\nu|_W$ maps a point of W to z or another point y on Σ. We argue that $y \neq z$ does not exist. For if it does, the distance $\|y-x\|$ has to be less than $0.3\varepsilon\mathrm{lfs}(z)$. Hence we have $\|z-y\| \leq 0.6\varepsilon\mathrm{lfs}(z)$. On the other hand, $\|y-z\|$ has to be more than the distance of z to the medial axis, which is at least $\mathrm{lfs}(z)$. Therefore, for each point $z \in \Sigma$, there is a point in W which is mapped by $\nu|_W$ to z. □

In the following lemma we prove that the directional derivative of $\mathcal{N}$ along $\tilde{n}_{\tilde{x}}$ is always positive. This, in turn, helps us to prove that $\nu|_W$ is injective. We introduce a handy notation. Let $u_x[\mathcal{N}]$ be the directional derivative of $\mathcal{N}$ at a point x along the unit vector u.

Lemma 6 Let z be any point on Σ, then for any $x \in \ell_{z,\tilde{n}_z} \cap 0.3\varepsilon\Sigma$

$$(\tilde{n}_z)_x[\mathcal{N}] > 0$$

Lemma 7 $\nu|_W$ is injective.

Proof To prove the injectivity of $\nu|_W$, assume for contradiction that there are two points w and w' in W so that $\nu|_W(w) = \nu|_W(w') = z$. This means $\ell_{z,\tilde{n}_z}$ intersects W at w and w' within $0.3\varepsilon\Sigma$ (Lemma 4). Without loss of generality assume that w and w' are two such consecutive intersection points. Then, $\ell_{z,\tilde{n}_z}$ makes at least $\frac{\pi}{2}$ angle with one of the normals to W at w and w'. But, that is impossible since Lemma 6 implies that

$$\angle \tilde{n}_z, \bigtriangledown\mathcal{N}(x) < \frac{\pi}{2}$$

for any point $x \in \ell_{z,\tilde{n}_z} \cap 0.3\varepsilon\Sigma$. □

6. Normal estimation

The computation of the normal vector field n requires assigned normals at the points of P. These assigned normals should approximate the normals at the closest points of Σ. Recently Dey et al. [DGS05] proposed a Delaunay based algorithm for estimating normals. The proofs for this estimation works with uniform sampling condition. We can extend the proofs to the adaptive sampling condition [DS05].

A Delaunay ball is a circumscribing ball of a tetrahedron in the Delaunay triangulation of the input point set P. It is denoted $B(c,r)$ if its center is c and radius is r. We call $B(c,r)$ big if r is more than certain times the average nearest neighbor distances of the samples incident on $B(c,r)$. The vectors from the samples incident to such big Delaunay balls towards their centers indeed approximate the normals of Σ. Figure 6 shows an implementation of this concept.

The following lemma is the basis of our normal estimation.

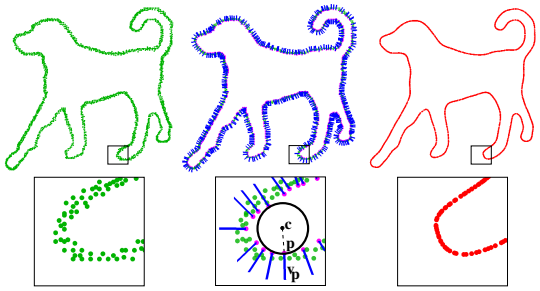

Figure 6: *Outward normals are estimated from big Delaunay balls at a subset of samples (middle); points after projection with these normals (right).*

Lemma 8 Let $p \in P$ be incident to a Delaunay ball $B(c,r)$ where $r > \frac{1}{5}\mathrm{lfs}(\tilde{p})$ and $c \in \Omega_O$. Then, $\angle \vec{pc}, \tilde{n}_{\tilde{p}} = O(\varepsilon)$ for a sufficiently small $\varepsilon > 0$.

Lemma 8 tells us that normals can be estimated from large Delaunay balls. The next lemma is a direct consequence of Lemma 5 in [DG04] which says that there are many big Delaunay balls.

Lemma 9 For each point $x \in \Sigma$, there is a Delaunay ball containing a medial axis point inside and a sample on the boundary within $O(\varepsilon)f(x)$ distance from x.

Lemma 8 and Lemma 9 together suggest an algorithm for estimating the normals of Σ from P. We compute the big Delaunay balls by comparing their radii with the nearest neighbor distances of the incident samples. For a point $p \in P$, let λ_p denote the average nearest distances to the five nearest neighbors of p in P. We determine all Delaunay balls incident to p whose radii are larger than $c\lambda_p$ where c is an user defined parameter. We take $c = 2.5$ in our experiments. Notice that some points in P may not satisfy this condition which means they do not contribute any big Delaunay balls. After obtaining the normals, we orient them using a consistent walk on inner and outer Delaunay balls as described in [DG04].

7. Algorithm and Implementation

In this section we summarize different steps of the algorithm for reconstruction and their implementations. In absence of Σ one cannot compute $\mathrm{lfs}(\tilde{x})$ and hence $f(\tilde{x})$ for a point x exactly. Due to this difficulty, our implementation as described below can only approximate the AMLS surface. However, the results of our implementation show that this approximation is effective in practice.

We have already discussed the normal estimation step. For feature estimation, we use the approximation of the medial axis with Voronoi diagrams. In noiseless case, Amenta, Choi and Kolluri [ACK01] showed that, for a dense sample, the Voronoi vertices furthest from their generating sites

AMLS(P)
  NORMAL ESTIMATION:
    Compute Del(P)
    for each point p with big Delaunay ball
      compute the normal n_p
  FEATURE ESTIMATION:
    for each $p \in P$ estimate $f(\tilde{p})$
  PROJECTION:
    for each $p \in P$
1   project p to p' by Newton iteration;
      if $\|p - p'\| > \tau$ go to 1 with $p := p'$
    endfor
  RECONSTRUCTION:
    Let P' be the projected point set;
    reconstruct with P'.

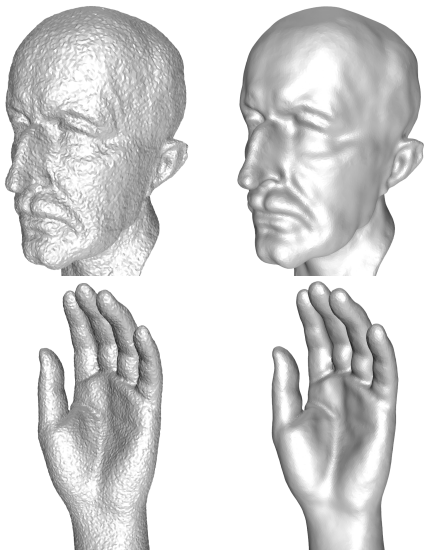

Figure 7: *Reconstruction results before (left) and after (right) smoothing with AMLS.* $\rho_e = 0.75$ *for both models. The reason we choose a bigger value for* ρ_e *than the one* (0.4) *we suggest in section 3.1 is that our feature estimation method tends to give a feature size slightly smaller than the exact one.*

also called *poles* approximate the medial axis. Then, by measuring the shortest distance of a sample p to the poles one can approximate $\text{lfs}(p)$. In case of noise, this does not work as poles do not necessarily approximate the medial axis. As a remedy we fall back on an observation of Dey and Goswami [DG04] which says that each medial axis point is covered with a big Delaunay ball. So, we consider each point $p \in P$ that has a big Delaunay ball incident to it. We take the center of the biggest Delaunay ball incident to p and also the center of the biggest Delaunay ball incident to p in the opposite direction. These centers act as poles in the noisy sample. Let L denote these set of poles. We approximate $f(\tilde{x})$ at any point x as $d(p,L)$ where p is the closest point to x in P. Actually, p approximates $\tilde{x}$ and $d(p,L)$ approximates the distance of $\tilde{x}$ to the medial axis.

In projection we move p to a new point p' by Newton iteration which can be described as follows. Project p along $\bigtriangledown\mathcal{I}(p)$ to a new position

$$p' = p - \frac{\mathcal{I}(p)}{\| \bigtriangledown\mathcal{I}(p)\|^2} \bigtriangledown\mathcal{I}(p) \tag{7.9}$$

and iterate until the distance between p and p' becomes smaller than a given threshold. To compute $\mathcal{I}$ and $\bigtriangledown\mathcal{I}$, we only take the samples inside the ball with radius 5 times the width of the Gaussian weighting function since the samples outside this ball have litte effect on the function. The convergence of the Newton iteration can be proved using standard numerical techniques. Due to the space limit, we skip a formal proof here. However, we show by examples that the Newton iteration for AMLS surface converges quickly and has a big convergent domain in section 8.

Finally, the projected set of points are fed to a reconstruction algorithm to produce the output. We used the COCONE software [COC] to reconstruct from the projected set of points. Figure 7 shows the results of our algorithm applied on MAX-PLANCK and BIGHAND point clouds.

8. Computational aspects of AMLS

In this section we discuss several properties of AMLS some of which provide it an edge over the others.

8.1. Normalizing weights

The difference between the functions $\mathcal{N}(x)$ and $\mathcal{I}(x)$ is that the weights in $\mathcal{I}(x)$ is normalized. Although normalizing weights does not change the implicit surface as we discussed in section 5, it does change the behavior of the Newton projection (NP). Specifically it increases the domain of convergence significantly. One can see from equation 7.9 that NP moves x along the direction $\bigtriangledown\mathcal{I}(x)$ when $\mathcal{I}(x) < 0$ and along the opposite direction when $\mathcal{I}(x) > 0$. Newton projection for the funtion $\mathcal{N}$ has a similar behavior. Figure 8 shows the gradient field and the results of NP for $\mathcal{I}$ and $\mathcal{N}$ applied on a noisy data.

8.2. AMLS vs. VMLS

VMLS surfaces as discussed in section 1.1.3 have an inherent projection procedure (PP) by which the points are projected onto the surfaces. PP can be described as follows. Project x along $n(x)$ to a new position

$$x' = x - \frac{\mathcal{G}(x)}{\sum_{p\in P} \theta_p(x)} n(x) \tag{8.10}$$

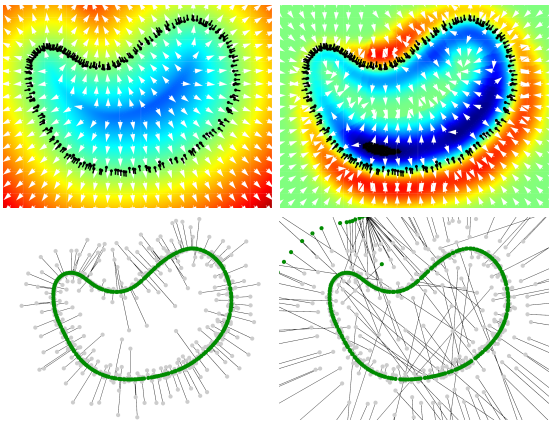

Figure 8: *The left and right columns show the result of NP on $\mathcal{I}$ and $\mathcal{N}$ respectively. The top row shows the function values and the gradient field. The values are negative at cold color and positive at warm color. In the bottom, the lines connect input gray points to their corresponding stationary green points of NP.*

Model	|P|	Method	#nb	#iter	Time
Max-planck	49137	NP	1000	3.1	94
		PP	1108	7.2	310
bighand	38214	NP	1392	3.2	109
		PP	1527	8.6	400

Table 1: *|P| is the number of points in the point cloud. We compute #iter in the average sense, i.e., we add up the number of iterations used to project all the input points and divide it by |P| to get #iter. Similarly #nb is the average number of points considered as neighbors. We choose the threshold $\tau = 10^{-25}$. Times (second) are for projecting all the input points (PC with a 2.8GHz P4 CPU and 1GB RAM).*

and iterate until a stationary point is reached. We argue that NP used for AMLS surface is better than the PP in two respects: convergence rate and timing. As Table 8.2 shows, NP, in general, uses less iterations to project a point onto the implicit surface. This is not surprising as $\bigtriangledown\mathcal{I}(x)$ with x close to the implicit surface can estimate the normal more accurately at its closest point on the implicit surface. In addition, one has to compute $n(x)$ before evaluating $\mathcal{G}(x)$. Hence to compute the new position using PP, one has to iterate twice over its neighboring points which makes PP slower than NP even in each iteration.

8.3. AMLS vs. PMLS

In the definition of standard PMLS, the actual PMLS surface is only a subset of the zero level set $\mathcal{J}^{-1}(0)$ where the energy function $\mathcal{E}$ reaches a minimum along the normal direction. As one can deduce from equation 1.2, there are two other layers of zero-level sets of the implicit function $\mathcal{J}$ on both sides of the PMLS surface, where the energy function $\mathcal{E}$ reaches the local maximum; see the left most picture in Figure 9. We refer to these two layers as *maxima layers*. The distance between these layers could be extremely small at places where either the local feature size is small or the noise level is high or both. In that case, computations on the PMLS surface become difficult.

First of all, many existing implicit surface techniques such as raytracing and polygonizing become hard to apply on the standard PMLS surface since one needs to distinguish different zero-level sets. When the maxima layers come close to the true PMLS surface, the marching step in a raytracer and the size of the cubes in a polygonizer may become impractically small. Actually, in the leftmost picture of Figure 9, although one edge of the red cube only intersects the curve (bold green line) once, a polygonizing algorithm [Blo94] misjudged the edge not intersecting the surface.

Second, the inherent projection procedure for the standard PMLS surface requires a non-linear optimization, specifically an one-dimensional minimization. The one-dimensional minimization algorithms usually begin with an interval known to contain a minimum guess m such that the function value at m must be less than the function values at the ends of the interval. Finding such a minimum guess m could be hard if the two maxima layers come close.

Third, the standard PMLS surface is more sensitive to the noise. When the noise level for position or normal or both increases, the three layers of the zero-level sets (one for minima and two for maxima) could easily interfere with each other. In the middle picture of Figure 9, the zero-level set for minima gets merged with those for maxima. As a result, the standard PMLS could give an implicit surface with holes or disconnectness. However under the same level of noise, the AMLS still gives the proper implicit surface, see the rightmost picture in Figure 9.

Acknowledgements.

We acknowledge the support of Army Research Office, USA under grant DAAD19-02-1-0347 and NSF, USA under grants DMS-0310642 and CCR-0430735.

References

[AA03] ADAMSON A., ALEXA M.: Ray tracing point set surfaces. In *Proc. Shape Modeling Internat.* (2003), pp. 272–279.

[AB99] AMENTA N., BERN M.: Surface reconstruction by voronoi filtering. *Discr. Comput. Geom. 22* (1999), 481–504.

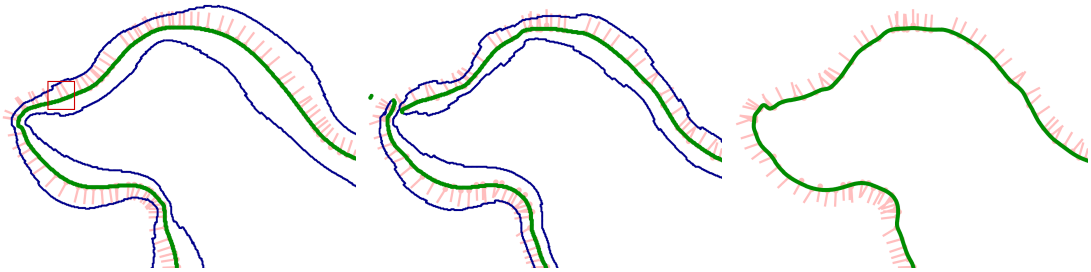

Figure 9: *The leftmost and the middle pictures show 0-level sets of the standard PMLS under two different noise levels. The noise level in the middle is higher. Bold green curves represent the 0-level set $\mathcal{J}^{-1}(0)$ where $\mathcal{E}$ reaches minima while the blue curves are 0-level sets where the energy function reaches maxima. The rightmost picture shows the 0-level set $\mathcal{I}^{-1}(0)$ under the same noise level as in the middle picture.*

[ABCO*01] ALEXA M., BEHR J., COHEN-OR D., FLEISHMAN S., LEVIN D., SILVA C.: Point set surfaces. In *Proc. IEEE Visualization* (2001), pp. 21–28.

[ACDL02] AMENTA N., CHOI S., DEY T. K., LEEKHA N.: A simple algorithm for homeomorphic surface reconstruction. *Internat. J. Comput. Geom. & Applications 12* (2002), 125–141.

[ACK01] AMENTA N., CHOI S., KOLLURI R. K.: The power crust, union of balls, and the medial axis transform. *Comput. Geom.: Theory Applications 19* (2001), 127–153.

[AK04] AMENTA N., KIL Y. J.: Defining point-set surfaces. In *Proceedings of ACM SIGGRAPH 2004* (Aug. 2004), ACM Press, pp. 264–270.

[BC00] BOISSONNAT J. D., CAZALS F.: Smooth surface reconstruction via natural neighbor interpolation of distance functions. In *Proc. 16th. Annu. Sympos. Comput. Geom.* (2000), pp. 223–232.

[Blo94] BLOOMENTHAL J.: An implicit surface polygonizer. *Graphics Gems 5* (1994), 324–349.

[CCS04] CHAZAL F., COHEN-STEINER D.: A condition for isotopic approximation. In *Proc. Ninth ACM Sympos. Solid Modeling Appl.* (2004).

[COC] COCONE: www.cse.ohio-state.edu/∼tamaldey. *The Ohio State University*.

[DG04] DEY T. K., GOSWAMI S.: Provable surface reconstruction from noisy samples. In *Proc. 20th Annu. Sympos. Comput. Geom.* (2004), pp. 330 – 339.

[DGS05] DEY T. K., GOSWAMI S., SUN J.: Extremal surface based projections converge and reconstruct with isotopy. *Technical Report OSU-CISRC-4-05-TR25* (April 2005).

[DS05] DEY T. K., SUN J.: An adaptive mls surface for reconstruction with guarantees. *Technical Report OSU-CISRC-4-05-TR26* (April 2005).

[Hir88] HIRSCH M. W.: *Differential Topology*. Springer Verlag, 1988.

[JCCCM*01] J. C. CARR R. K. B., CHERRIE J. B., MITCHELL T. J., FRIGHT W. R., MCCALLUM B. C., EVANS T. R.: Reconstruction and representation of 3d objects with radial basis functions. In *Proceedings of ACM SIGGRAPH 2001* (2001), ACM Press, pp. 67–76.

[Kol05] KOLLURI R.: Provably good moving least squares. *SODA 2005* (2005), 1008–1017.

[Lev98] LEVIN D.: The approximation power of moving least-squares. *Math. Computation 67* (1998), 1517–1531.

[MVdF03] MEDEROS B., VELHO L., DE FIGUEIREDO L. H.: Moving least squares multiresolution surface approximation. In *SIBGRAPI 2003-XVI Brazilian Sympos. Comput. Graphics Image Process.* (2003).

[OBA*03] OHTAKE Y., BELYAEV A., ALEXA M., TURK G., SEIDEL H.-P.: Multi-level partition of unity implicits. In *Proceedings of ACM SIGGRAPH 2003* (Aug. 2003), ACM Press, pp. 463–470.

[PKKG03] PAULY M., KEISER R., KOBBELT L., GROSS M.: Shape modeling with point-sampled geometry. In *Proceedings of ACM SIGGRAPH 2003* (2003), ACM Press, pp. 641–650.

[SOS04] SHEN C., O'BRIEN J. F., SHEWCHUK J. R.: Interpolating and approximating implicit surfaces from polygon soup. In *Proceedings of ACM SIGGRAPH 2004* (Aug. 2004), ACM Press, pp. 896–904.

[Wol92] WOLTER F.-E.: Cut locus and medial axis in global shape interrogation and representation. *MIT Design Laboratory Memorandum 92-2 and MIT Sea Grant Report* (1992).

[ZPKG02] ZWICKER M., PAULY M., KNOLL O., GROSS M.: Pointshop 3d: An interactive system for point-based surface editing. In *Proceedings of ACM SIGGRAPH 2002* (2002), ACM Press, pp. 322–329.