Provably Good Implicit MLS Surfaces

Nikola Milosavljevic

CS 468, Fall 2005

Implicit MLS Surfaces

- **Zero-level-set of a function $\mathcal{I}(x)$ over $\mathbb{R}^3$**

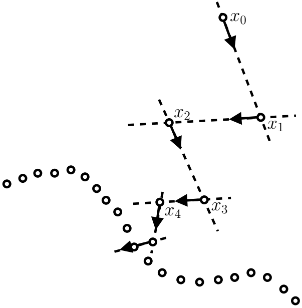

- Fixed points of a projection operator
- Weighted sum of signed distances

$$\mathcal{I}(x) = \frac{\sum_p n_p^T (x - p) \cdot \theta_p(x)}{\sum_p \theta_p(x)}$$

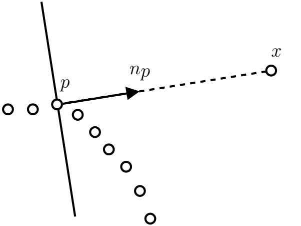

- Extremal surfaces

$$\mathcal{I}(x) = n(x) \cdot \left. \frac{\partial e_{MLS}(y, n(x))}{\partial y} \right|_x$$

MLS Surfaces with Guarantees

- Not discussed so far
  - What is the "ground truth" surface?
  - How do the samples arise?
  - How good is the reconstruction?

- In this talk
  - A notion of the **original surface**
  - A model of **sampling** and **noise**
  - Study the behavior of MLS surfaces

- The Holy Grail
  - Geometric accuracy
  - Correct topology
  - Smoothness

  - Fast (quadratic) convergence
  - Efficient (local) computation

Problem Statement

- Original smooth, closed surface Σ
- Given conditions on
  - Sampling density
  - Normal estimates
  - Noise
- Design an implicit function $\mathcal{I}(x)$ whose zero set recovers Σ

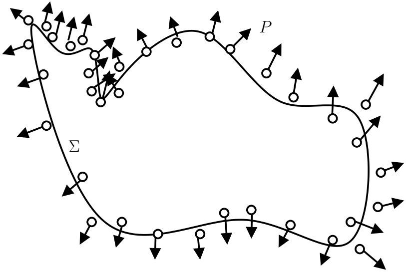

Outline

R. Kolluri (U.C. Berkeley),
“Provably Good Moving Least Squares”, SODA 2005.

- Globally uniform sampling + normals
- Correct topology
- Smoothness

T. Dey, J. Sun (Ohio State),
“An Adaptive MLS Surface for Reconstruction with Guarantees”, SGP 2005.

- Feature-sensitive sampling + normals
- Correct topology
- “Smoothness”

Sampling Conditions

- Medial axis
  - Points with multiple closest points on Σ

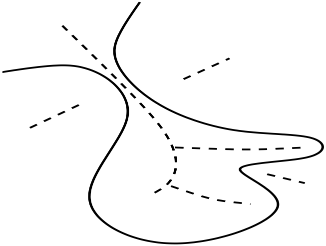

- Local feature size, $\mathrm{lfs}(\cdot)$
  - Distance from $x \in \Sigma$ to the medial axis
- ϵ-sampling
  - Every $x \in \mathbb{R}^3$ has a sample p at most $\epsilon\, \mathrm{lfs}(\tilde{x})$

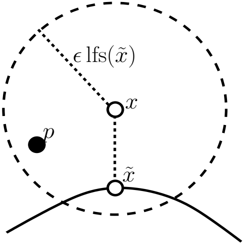

- No oversampling, $|B(x, \epsilon\, \mathrm{lfs}(\tilde{x}))| \leq \alpha$

Typical Proof Outline

- **Step 1:**
  - Analyze $\mathcal{I}(x)$
  - Localize the zero-set
  - Spurious zero-crossings?

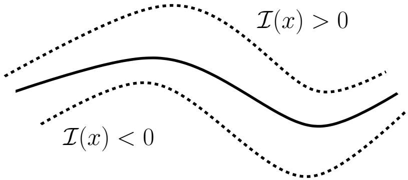

- **Implication:** Reconstruction is geometrically close

Typical Proof Outline

- **Step 2:**
  - Analyze $\nabla\mathcal{I}(x)$ close to the surface
  - Show that $\nabla\mathcal{I}(x) \neq 0$
  - Show that $\mathcal{I}(x)$ is **strictly** monotonic in the direction normal to Σ

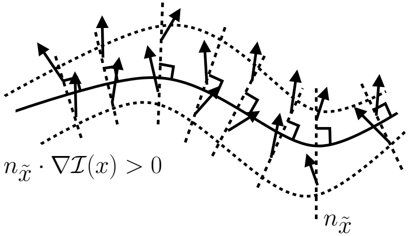

- **Implications:**
  - Reconstructed surface is a manifold
  - Normal directions define a homeomorphism

Typical Proof Outline

- **A common technique**
  - Bounding the influence of points farther than a suitably chosen threshold
  - The actual radius depends on the quantity that is being evaluated
  - Inside — reliable
  - Outside — negligible

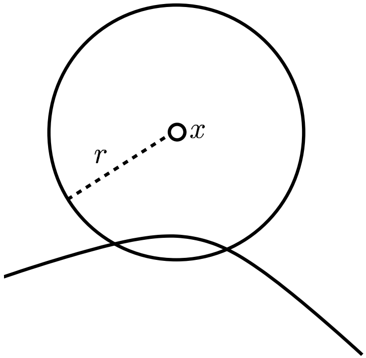

An MLS Surface for a Uniformly Sampled PCD

Assumptions

- **Uniform sampling**, $||x - p|| \leq \epsilon$
  - Assume $\text{lfs}(x) \geq 1$ everywhere
  - Smallest features determine density
- No oversampling, $|B(x, \epsilon)| \leq \alpha$
- Noise, $||p - \tilde{p}|| \leq \epsilon^2$
- Normal estimates, $\angle(n_p, n_{\tilde{p}}) \leq \epsilon$

Proposed MLS Surface

- Weighted sum of signed distances

$$\mathcal{I}(x) = \frac{\sum_p n_p^T (x - p) \cdot \theta_p(x)}{\sum_p \theta_p(x)} \qquad \theta_p(x) = \frac{1}{\alpha_p} e^{-||x-p||^2/\epsilon^2}$$

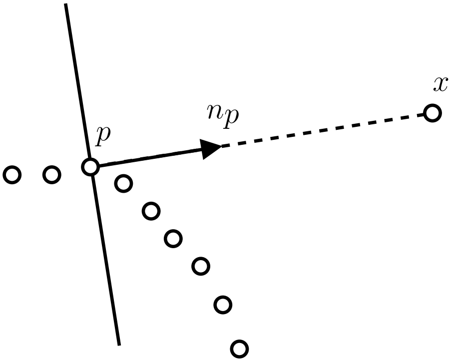

Analysis of $\mathcal{I}(x)$

- Can show that all zero-crossings are within δ from Σ

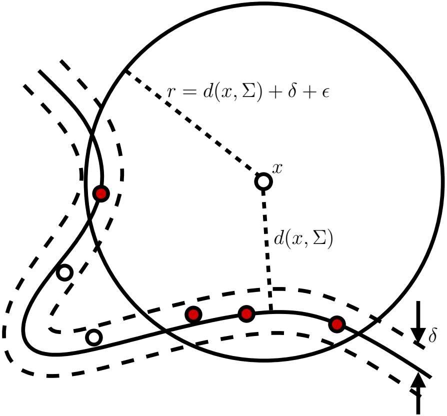

- Fix x far away (farther than δ) from the boundary

- Influence threshold

$$r = d(x,\Sigma) + \delta + \epsilon$$

  - If p is a **nearby sample**, $n_p^T(x-p) = d(x,\Sigma)\cdot(1+O(\epsilon)) + O(\epsilon^2)$

  - $n_p^T(x-p)$ and $d(x,\Sigma)$ have the same sign, provided $\delta = O(\epsilon)$

Analysis of $\mathcal{I}(x)$

Far away points

- Divide the "distant space" into spherical shells of thickness ϵ
- The number of samples in the i-th shell is $O(i^2)$ (uniform sampling)
- The influence decays as $O(e^{-i^2})$
- The overall influence

$$O\left(r \cdot \frac{r^2}{\epsilon^2} \cdot e^{-r^2/\epsilon^2}\right)$$

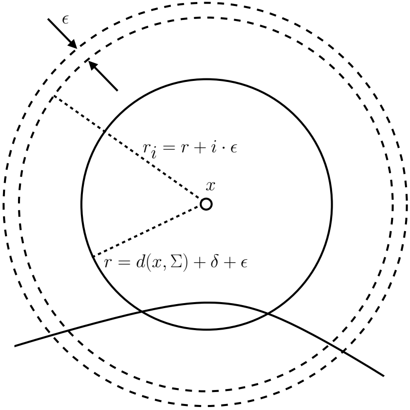

Analysis of $\nabla\mathcal{I}(x)$

- Fix x close (within δ) to the boundary
- Show $n_{\tilde{x}} \cdot \nabla\mathcal{I}(x) > 0$ where $\tilde{x} \in \Sigma$ is closest to x.
- Influence threshold

$$r = \sqrt{(d(x,\Sigma)+\epsilon)^2+\epsilon^2} = O(\epsilon)$$

- Far away points negligible

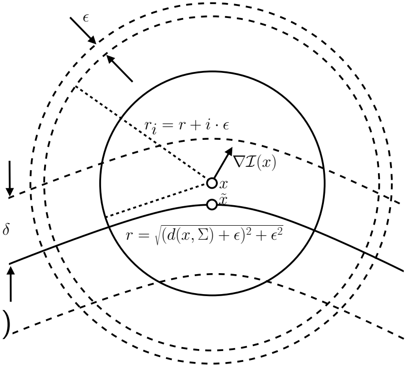

Analysis of $\nabla\mathcal{I}(x)$

- Nearby points contribution to the gradient vector
  - Signed distance functions

$$\sum_p \sum_q \theta_p(x)\theta_q(x) \cdot n_p$$

  - Change of weights

$$\sum_p \sum_q \theta_p(x)\theta_q(x) \cdot O(n_p^T(x-p)) \cdot (p-q)$$

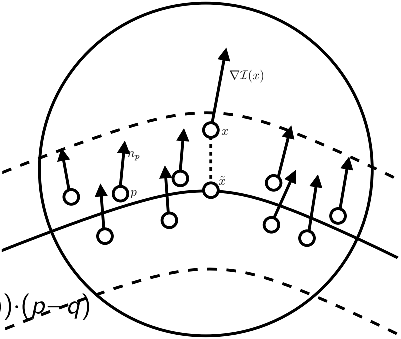

- The normals n_p are close to $n_{\tilde{x}}$!

Uniform Case: Conclusions

- The zero set of $\mathcal{I}$ is confined to the $\delta = O(\epsilon)$ thickening
  - The reconstruction is geometrically accurate

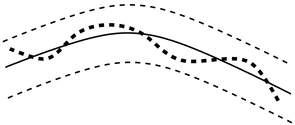

- Whenever $\mathcal{I}(x) = 0$, $\nabla\mathcal{I}(x) \neq 0$
  - The reconstruction is locally flat

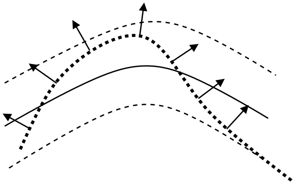

- The gradient lines provide a "morphing function"
  - The reconstruction is topologically correct

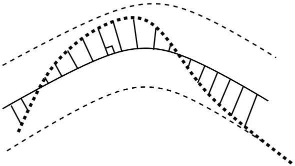

An MLS Surface for an Adaptively Sampled PCD

Motivation

- Allow variations in sampling density according to local feature size

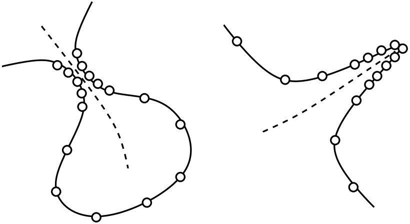

- Requires adaptive Gaussian kernel
- Uniform sampling: kernel width ϵ
  - Does not work for adaptive sampling

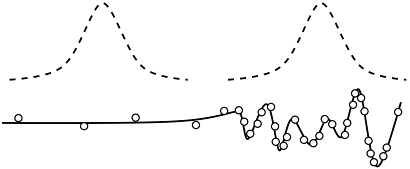

Adaptive Gaussian Kernel

- Adapt to $\text{lfs}(\tilde{x})$?

$$\theta_p(x) \sim e^{-O\left(\frac{||x-p||^2}{\epsilon^2 \text{ lfs}(\tilde{x})^2}\right)}$$

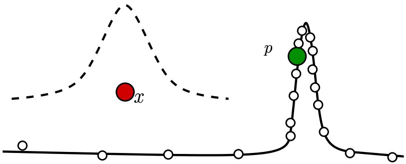

  - Bias toward small features

- Adapt to $\text{lfs}(\tilde{p})$?

$$\theta_p(x) \sim e^{-O\left(\frac{||x-p||^2}{\epsilon^2 \text{ lfs}(\tilde{p})^2}\right)}$$

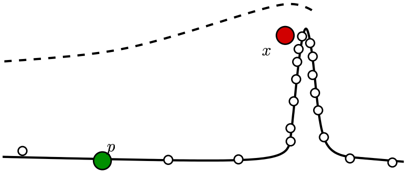

  - Influence may not decrease with distance

Adaptive Gaussian Kernel

- Solution: Adapt to $\sqrt{\operatorname{lfs}(p)\operatorname{lfs}(x)}$

$$\theta_p(x) = \exp\left(-\frac{\sqrt{2}\cdot||x-p||^2}{\epsilon^2\operatorname{lfs}(\tilde{x})\operatorname{lfs}(\tilde{p})}\right)$$

- Note: not smooth!

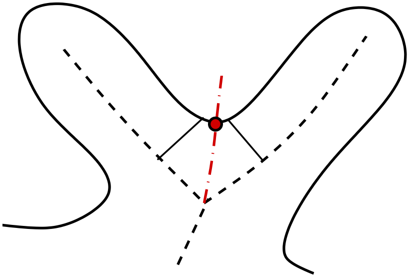

Other Assumptions

- No oversampling, $|B(x, \epsilon \operatorname{lfs}(\tilde{x}))| \leq \alpha$
- Noise magnitude at most $\epsilon^2 \operatorname{lfs}(\tilde{x})$
- Good normal estimates, $\angle(n_p, n_{\tilde{p}}) \leq \epsilon$

Analysis

- Extend the proofs for adaptive thickening, width $\delta\,\mathrm{lfs}(\tilde{x})$
- Fix a point x and a threshold radius r

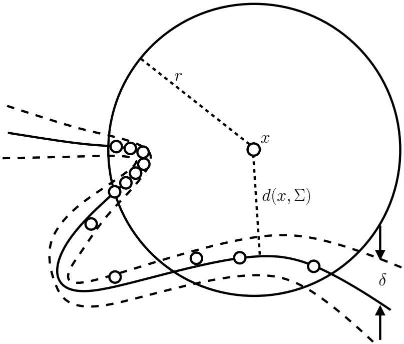

- **Bounding the influence of far away points**
  - Small distant features
- **Reliability of nearby points**
  - Small nearby features

Influence of Far Away Points

- Subdivide into cubes, accumulate the counts bottom-up
  - Size of top-level cubes $O(\epsilon \, \mathrm{lfs}(\tilde{x}))$
  - Stop subdivision when $\mathrm{lfs}(c_k) \geq (\epsilon/2^k) \, \mathrm{lfs}(\tilde{x})$
  - Apply "no oversampling" to the leaves

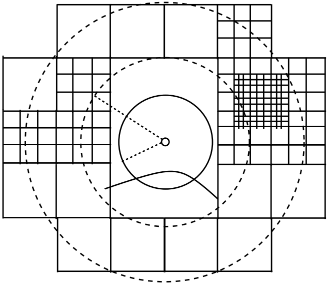

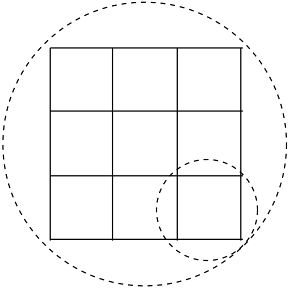

- The i-th shell may contain more than $O(i^2)$ samples, but still the total contribution is $O(i^2 e^{-i^2})$
- Total contribution to $\mathcal{I}(x)$ and $n_{\tilde{x}} \cdot \nabla \mathcal{I}(x)$

$$O(\mathrm{poly}(r/\epsilon) \cdot \exp(-r^2/\epsilon^2))$$

Influence of Nearby Points

- Works only for $d(x, \Sigma) \leq 0.1\,\mathrm{lfs}(\tilde{x})$ from the surface
- If x is at least $(\delta = 0.3\epsilon)\,\mathrm{lfs}(\tilde{x})$ away from the surface, the sign is correct

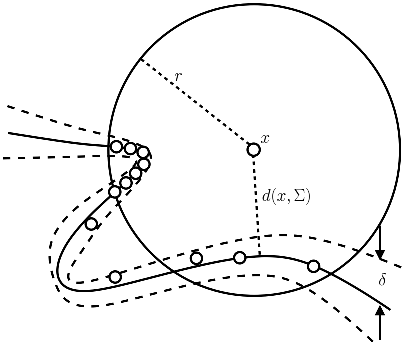

- Can claim $\mathcal{I}(x) \neq 0$ only for

$$0.3\epsilon \leq \frac{d(x, \Sigma)}{\mathrm{lfs}(\tilde{x})} \leq 0.1$$

Influence of Nearby Points

- Works only if x is not too far from the surface
  - May have fake zero-crossings outside $0.1\,\mathrm{lfs}(\tilde{x})$!

- Uniform sampling

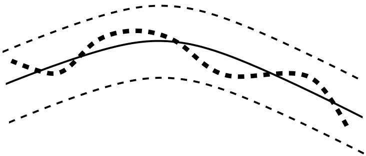

- Adaptive sampling

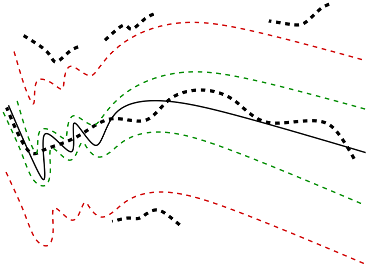

Remarks

- Delaunay-based estimation of normals and medial axis

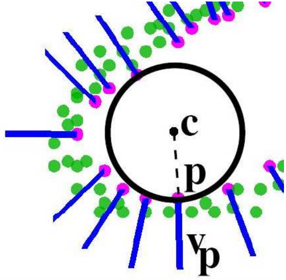

- Maxima layers for standard PMLS surfaces

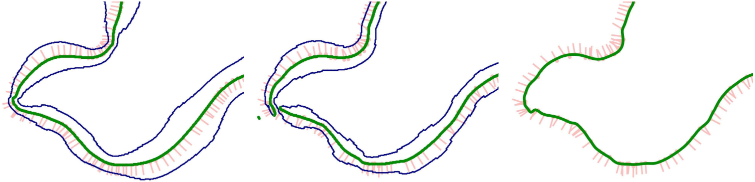

- Comparison with other projection methods