Reconstruction Submanifolds of Euclidean Space

David Letscher *

Abstract

A generalization of the crust algorithm is presented that will reconstruct a smooth d-dimensional submanifold of $\mathbf{R}^k$. When the point sample meets satisfy a minimal density requirement this reconstruction is homeomorphic to the original submanifold. In fact the reconstructed manifold is ambiently isotopic to the original via an isotopy that moves points a small distance. Also, bounds are given comparing the metric of the source and reconstructed manifolds.

1 Introduction

The manifold reconstruction problem consists of finding a piecewise linear approximation to a smooth d-dimensional submanifold of $\mathbb{R}^k$ from a set of n points on the submanifold. Reconstructing a surface in $\mathbb{R}^3$ has applications in areas including computer graphics, medical imaging, computer vision and computer aided design. Higher dimension applications include manifold learning [4] and medical imaging where time is taken into account to work in 4 dimensions.

Topologically correct algorithms for surface reconstruction include the crust algorithm [1] and, a simplified version, the cocone algorithm [3]. Both of these algorithms return a surface that is collection of faces in the Delaunay triangulation and return a homeomorphic copy of the surface when certain minimum density sampling requirements are met. More recently improvements on these algorithms have been made that run in $O(n \log n)$ time [7]. In arbitrary dimensions an algorithm with running time $O(n \log n)$ in the size of the point samples and exponential in the dimension has been introduced [5]. This algorithm requires an exponentially dense sample of points and meet a minimum sampling density to correctly reconstruct a submanifold. If the dimension of the submanifold is not known in advance, an algorithm in [6] can provide an estimate.

In this paper we present an algorithm for reconstructing smooth submanifolds $M^d \subset \mathbb{R}^k$ from an arbitrary point sample. It generalizes the crust algorithm in [1]. For point sets that meet a minimum sample density that is dependent on the dimension of the submanifold, the resulting surface is an accurate reconstruction: it is both homeomorphic and isotopic to the original manifold, the approximation converges pointwise to the original manifold and induced metric approximates the Riemannian metric for the original. Like the algorithm in [5], the running time is exponential in the dimension of the ambient spaces as the complete Delaunay triangulation needs to be found.

In section 2 preliminary details and terminology are introduced for the problem. Section 3 presents an algorithm for estimating both the normal and tangent planes is presented that is used in the main algorithm in section 4. The topological guarantees on the quality of the reconstruction are given in sections 5. And section 6 discusses the issue of the manifold extraction problem in the main algorithm.

*Department of Mathematics and Computer Science, St. Louis University, letscher@slu.edu

2 Background

Some terminology and notation needs to be established before we can continue.

Medial axis and local feature size For any submanifold M of $\mathbf{R}^k$, the closure of the set of points that do not have a unique closest point on M is referred to as the *medial axis* of M. The *local feature size* of a point $p \in M$ is the distance from p to the medial axis and will be denoted $LFS(p)$. This provides a measure of the scale on which the geometry and topology of the manifold are interesting in a neighborhood of the point.

Restricted Delaunay triangulation For sample point $S \subset M$, consider $Vor(S) \cap M$. For a dense sampling this will be a cell decomposition of M and its dual triangulation will be referred to as the *restricted Delaunay triangulation* for S and M.

Sampling condition For any $\epsilon > 0$, a set of points $S \subset M$ is called ϵ-*dense* if for any point $p \in M$ there exists a sample point $s \in S$ such that $d(p, s) \leq \epsilon LFS(p)$. There are no minimums on how close the points in the sample can be to each other or other requirements on how evenly distributed the sample points must be.

Tangent and normal planes For a smooth d-dimensional submanifold M of $\mathbb{R}^k$, then for any point $p \in M$ then tangent plane through p, denoted T_pM, is the the d-dimensional plane though p containing all

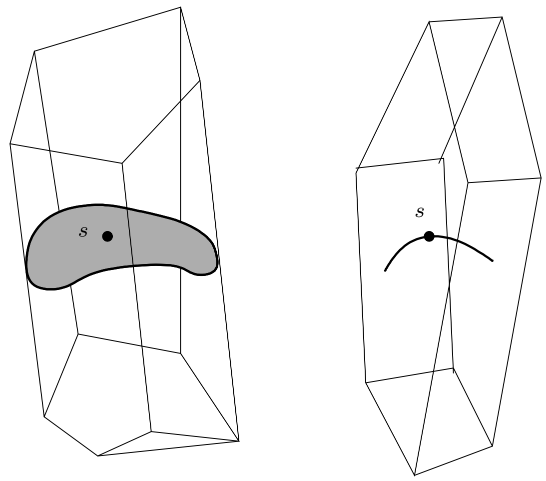

Figure 1: Voronoi cells for a surface and a curve in $\mathbb{R}^3$.

points x such that the vector $x-p$ is tangent to M at the point p. Similarly, the normal plane, N_pM is the set of points x such that $x-p$ is perpendicular to M at p.

Other notation For any point $x \in \mathbb{R}^k$, we will denote the open ball of radius r centered at x by $B_r(x)$. For a finite set of points X, P_X will denote the plane spanned by the points and $P_X^\perp$ the plane perpendicular to P_X.

3 Approximating the Normal and Tangent Planes

For a dense set of sample points on a submanifold the Voronoi cell have a large aspect ratio. They are small in the tangential directions and large in the normal directions as shown by the following theorem.

Theorem 1 *Assume S is an ϵ-sample for $\epsilon \leq .25$ of $M^d \subset \mathbb{R}^k$. If $s \in S$ and $Vor(s)$ is the Voronoi cell containing s then*

1. $Vor(s) \cap N_sM \supset B_{LFS(s)}(s) \cap N_sM$
2. $Vor(s) \cap M \subset B_{\frac{\epsilon}{1-\epsilon}}(s)$
3. $Vor(s) \cap T_sM \subset B_{6.36LFS(s)}(s)$

The lemma indicates that the Voronoi cells are relatively large in the normal directions and small in the tangential directions, see figure 1. This motivates how the tangent and normal planes can be accurately estimated.

The algorithm generalize the notion of "poles" used in the crust [1] and cocone [3] algorithms for surface reconstruction; the vector to the pole, or furthest Voronoi vertex, is close to the surface normal. By carefully choosing $k-d$ Voronoi vertices that are far away from the sample point and meet an orthogonality condition we obtain an accurate approximation to the normal.

Algorithm 1 Normal plane approximation

Input: a sample point s, the Voronoi diagram $\mathcal{V}\wr\nabla(S)$ and a dimension d
Output: a $(k-d)$-plane AN_sM

```
X := {s}
while |X| < n - d + 1 do
    Let v be the Voronoi vertex in the cell containing
    s such that the projection of sv⃗ to P_X^⊥ is longest
    X := X ∪ {v}
end while
return P_X
```

The angle between the normal and its approximation is $O(\epsilon\sqrt{k-d})$ as shown by the following theorem.

Theorem 2 *For an ϵ-sample S, $\epsilon \leq \frac{.25}{k-d}$, and $s \in S$ then the angle between N_sM and AN_sM is at most $\sin^{-1}(2\epsilon\sqrt{k-d})$.*

Note that both the density requirement and accuracy are function of the co-dimension of the manifold, $k-d$. This will be the case with most of the results that follow. This implies that lower density is required for higher dimensional submanifolds of $\mathbb{R}^k$. For example, curves in $\mathbb{R}^3$ require a denser point sampling that surfaces in $\mathbb{R}^3$.

4 Reconstruction Algorithm

Using the approximate normal and tangent planes, the reconstruction algorithm proceeds essentially the same as the crust algorithm [1]. Algorithm 2 shows the major steps of the algorithm. It finds a subcomplex T_S of the Delaunay triangulation and extracts a d-dimensional manifold from it. This extracted manifold will be the reconstruction of our original manifold and will be discussed in section 6.

Ideally, the reconstruction algorithm would always return the restricted Delaunay triangulation of the submanifold. This would only happen if the only manifold that can be extracted is the restricted Delaunay triangulation. The following shows that for dense enough point samples, the restricted Delaunay triangulation is one of the possible reconstructions. A consequence of this is that the manifold extraction steps always succeeds.

Theorem 3 *If $\epsilon \leq \frac{.25}{\sqrt{k-d}}$ and S is an ϵ-sample for $M^d \subset \mathbb{R}^k$ then*

1. *the restricted Delaunay triangulation is contained in T_S*
2. *the restricted Delaunay triangulation is homeomorphic to M*

This theorem is proved for curves in the plane for $\epsilon \leq .4$ in [2] and for surfaces in $\mathbb{R}^3$ for $\epsilon \leq .1$ in [1]. Ideally,

Algorithm 2 Submanifold reconstruction

Input: A finite set of sample points $S \in \mathbf{R}^k$ and an integer $d \in \{1, \ldots, n-1\}$

Output: A triangulation of a d-manifold

1: Compute the Voronoi diagram and Delaunay triangulation
2: For each $s \in S$, compute $AN_s(M)$ using algorithm 1 and calculate its perpendicular AT_sM
3: Let C be the set of $(k-d)$-cells of the Voronoi diagram such that for some sample point for an adjacent cell the angle between the edge and AN_sM is less than $\pi/8$
4: Let T_S be the set of d simplicies of the Delaunay triangulation so that their dual Voronoi cells in C
5: Extract a closed d-manifold, M_S, from T_S such that every point in S is a vertex of M_S and any two d-simplicies of M_S that intersect meet at an angle less than $\pi/2$
6: **return** M_S

no other triangulation would ever be returned. This is true for curves in the plane if $\epsilon \leq .25$ [2]. However, it is not true in general; so we must show any extracted triangulation meets our needs.

5 Topological Guarantees

The first property that is desired for the reconstructed manifold is that original manifold and the reconstructed one are topologically identical; that is, homeomorphic.

Theorem 4 *For $\epsilon \leq \frac{.25}{\sqrt{k-d}}$ and an ϵ-sample S for $M^d \subset \mathbb{R}^k$ the reconstructed manifold M_S is homeomorphic to M.*

Instead of proving this theorem, a stronger property will be established. Consider the neighborhood M_ϵ of M defined as $M_\epsilon = \bigcup_{x \in M} B_{\epsilon LFS(x)}(x)$ For ϵ at most 1, M_ϵ is disjoint from the medial axis. So the map $\pi : M_\epsilon \to M$ taking a point z to its unique closest point on M is well defined. And for any $x \in M$, $\pi^{-1}(x) \subset N_xM$. In fact, $\pi^{-1}(x)$ is the $(n-d)$-ball of radius $\epsilon LFS(x)$ in N_xM.

These balls in the normal planes provides M_ϵ with some additional structure as a *fiber bundle*, a generalization of a product space. See [8] for a full discussion.

Definition 5 *A* fiber bundle *is a quadruple (E, B, π, F) such that*

1. *E, B, F are topological spaces*
2. *$\pi : E \to F$ is a continuous surjection*
3. *For any $x \in B$ there exists an open neighborhood U of x such that $\pi^{-1}(U)$ is homeomorphic to the product $U \times F$. Furthermore, $\pi|_{\pi^{-1}(U)}$ is equal to $proj_1 \circ \phi$ where $proj_1 : U \times F$ is the projection to the first coordinate and $\phi : \pi^{-1}(U) \to U \times F$ is a homeomorphism.*

E is referred to as the total space*, B the* base space *and F the* fiber*.*

Lemma 6 *For $\epsilon \leq 1$, $(M_\epsilon, M, \pi, D^{n-d})$, with D^{n-d} the $n-d$ disk, is a fiber bundle.*

When the fibers are disks these bundles are called *disc bundles*. Topologically, in any fibration there are many manifolds that could be used as a base space and it turn out that M_S is one of them.

Definition 7 *A* section *of a fiber bundle (E, B, π, F) is a continuous map $f : B \to E$ such that $\pi(f(x)) = x$ for all $x \in B$.*

Theorem 8 *For $\epsilon \leq \frac{.25}{\sqrt{k-d}}$ and $2.65\epsilon \leq \delta \leq 1$, M_S is a section of the fiber bundle $(M_\delta, M, \pi, D^{n-d})$.*

This follows from the fact that M_S intersects each fiber transversely in a single point. Using this fibration allows us to prove that that M and M_S are more than just homeomorphic. See [8] for more details.

Definition 9 *Subspaces X, Y of Z are* ambiently isotopy *if there exists a continuous map $h : Z \times [0,1] \to Z$ such that*

1. *For every $z \in Z$, $h_t(z) = h(z,t)$ is a homeomorphism from Z to itself.*
2. *$h(x, 0) = x$ for all $x \in X$*
3. *$h(X, 1) = Y$*

This allows us to prove the main result about the quality of the reconstruction. In particular, we can show that the manifold and its reconstruction are equivalent in how they are embedded in $\mathbb{R}^k$ and that no point in M_S is very far from its equivalent point in M.

Theorem 10 *For $\epsilon < \frac{.25}{\sqrt{k-d}}$, M and M_S are ambiently isotopy via an isotopy, $h(x,t)$. Furthermore, the isotopy moves a point $x \in M$ a distance at most $2.65\epsilon LFS(x)$.*

The proof of this theorem follows from the fact that for disk bundles any two sections are ambiently isotopic.

6 Manifold Extraction

A variety of strategies can be used to extract a manifold in algorithm 2 and any of them will satisfy the quality guarantees. However, there may be many possible manifolds to extract so some assistance is needed to do it efficiently. The original crust algorithm [1] relies on prior knowledge of the sample density to extract a manifold using the fact that the angle between adjacent triangles in $O(\epsilon)$ so knowing ϵ was needed for this step to work accurately. The cocone algorithm [3] to improve this stage of the reconstruction algorithm by first removing "sharp edges" and performing a depth-first search on the "outer boundary". A similar idea can be used in this context.

Algorithm 3 Manifold extraction

Input: A d-complex, T
Output: A triangulation of a d-manifold

```
while T is not the triangulation a closed d-manifold do
    Recursively remove d-simplicies from T that have a (d−1)-dimensional face without an adjacent d-simplex meeting it at an angle of 3π/2 or larger.
    Find a (d−1)-simplex τ of T such that there exist more than two d-simplicies containing τ.
    Let σ, σ′ be two of the faces that meet in an angle smaller than π/2.
    Remove all of the d-simplicies from T containing τ except σ and σ′.
end while
return T
```

Theorem 3 implies that the restricted Delaunay triangulation is in T_S. This ensures that it is possible to extract a closed manifold; however, we need to ensure that this algorithm will always construct such a manifold.

Theorem 11 *For $\epsilon \leq \frac{.25}{\sqrt{k-d}}$, when algorithm 3 is run with input T_S then it returns a closed d-manifold.*

7 Conclusion

We have presented an algorithm that can reconstruct any submanifold of Euclidean space. For sufficiently dense set of points this algorithm will always construct an ambiently isotopic copy of the original manifold. This density requirement needs the minimum distance to a sample point that is inversely proportional to the square root of the co-dimension of the manifold. As $\epsilon \to 0$, the manifold converges pointwise, the tangent planes converge and the metrics converge to those of the original manifold. All of the errors are $O(\epsilon)$ as $\epsilon \to 0$.

While this algorithm is not practical in high dimensions, the density requirements do not increase quickly with the dimensions of the spaces involved and strong quality guarantees are provided. In the future it may be possible to combine the $O(n \log n)$ running time of [5] with the techniques of this paper to weaken the requirements on the point sample.

References

[1] N. Amenta and M. W. Bern. Surface reconstruction by Voronoi filtering. *Discrete & Computational Geometry*, 22(4):481–504, 1999.

[2] N. Amenta, M. W. Bern, and D. Eppstein. The crust and the beta-skeleton: Combinatorial curve reconstruction. *Graphical Models and Image Processing*, 60(2):125–135, 1998.

[3] N. Amenta, S. Choi, T. K. Dey, and N. Leekha. A simple algorithm for homeomorphic surface reconstruction. *International Journal of Computational Geometry and Applications*, 12(1-2):125–141, 2002.

[4] C. Bregler and S. M. Omohundro. Nonlinear manifold learning for visual speech recognition. In *ICCV '95: Proceedings of the Fifth International Conference on Computer Vision*, page 494, Washington, DC, USA, 1995. IEEE Computer Society.

[5] S.-W. Cheng, T. K. Dey, and E. A. Ramos. Manifold reconstruction from point samples. In *SODA '05: Proceedings of the sixteenth annual ACM-SIAM symposium on Discrete algorithms*, pages 1018–1027, 2005.

[6] T. K. Dey, J. Giesen, S. Goswami, and W. Zhao. Shape dimension and approximation from samples. In *SODA '02: Proceedings of the thirteenth annual ACM-SIAM symposium on Discrete algorithms*, pages 772–780, Philadelphia, PA, USA, 2002. Society for Industrial and Applied Mathematics.

[7] S. Funke and E. A. Ramos. Smooth-surface reconstruction in near-linear time. In *SODA '02: Proceedings of the thirteenth annual ACM-SIAM symposium on Discrete algorithms*, pages 781–790, 2002.

[8] M. Hirsch. *Differential Topology.* Springer, 1976.