Computational Topology for Reconstruction of Surfaces with Boundary: Integrating Experiments and Theory

K. Abe , J. Bisceglio,
T. J. Peters
(tpeters@cse.uconn.edu),
A. C. Russell, D. R. Ferguson,
University of Connecticut
Storrs, CT 06269 USA

T. Sakkalis
Agricultural University of Athens
Athens, 118 55 Greece
stp@aua.gr

Abstract

We report new techniques and theory in computational topology for reconstructing surfaces with boundary. This complements and extends known techniques for surfaces without boundary. Our approach is motivated by differential geometry and differential topology. We have also conducted significant experimental work to test our resultant implementations. We discuss some problematic issues that can arise regarding the roles of the medial axis and sampling density. The crucial topics for C^2 manifolds are

1. *important defining properties of C^2 manifolds with boundary,*
2. *creation of auxiliary surfaces, with emphasis near the boundary,*
3. *sampling density, and*
4. *successful practical algorithms and examples.*

Keywords: surface reconstruction, twice-differentiable manifold, differential topology, differential geometry, medial axis.

1. Introduction and Motivation

The primary contribution of this paper is to present new theory and techniques for topology-preserving reconstruction and approximation of surfaces. This paper reports new results on

1. computational topology properties of C^2 manifolds with boundary, and
2. successful practical algorithms and examples.

Partial results are reported, regarding

1. approximations based on auxiliary surfaces, with attention to boundary characteristics and
2. sampling density.

For each surface M with boundary, we construct an auxiliary surface, called an *envelope*. We use the envelope to perform approximation and reconstruction of surfaces, specifically inclusive of those with boundary. We discuss our implementation of code that enriches the class of surfaces that can be considered, and we articulate the supporting practical algorithms that are derived from our theoretical and experimental investigations.

A comparison between the images of Figure 1 and Figure 2 show the value of our method. An original trefoil surface was created for these experiments by a linear extrusion of a trefoil curve, so as to produce no self-intersections. This trefoil surface is like a ribbon with the corresponding knot for its two boundary curves. Figure 1 is a series of reconstructions using the envelope technique with different sample densities, with the density decreasing from left to right. Figure 2 is the same using direct reconstruction from the PowerCrust algorithm. We note that, for the same sampling densities, the method based upon use of the envelopes in Figure 1 appears to converge nicely to the desired trefoil surface, while the direct reconstruction of Figure 2 does not. We note that the well-known Power Crust algorithm was not designed to accept surfaces with boundary. However, in practice, the Power Crust algorithm has been found to be useful for reconstruction of some surfaces with boundary, with possible reliance upon *ad hoc* modifications [5]. Hence, our interest in our comparisons is to begin to formalize a rich admissible class of input surfaces for

a provable implementation of ambient isotopic approximations. A related example is presented for an unknot surface in Section 7.2.

Several recent approaches to topology-preserving surface approximation have been restricted to C^2 2-manifolds without boundary [4, 7, 14, 23]. The approach taken in the CoCone software specifically attempts to recognize undersampling near boundaries to improve reconstruction for surfaces with boundaries [13]. For some practical surface reconstructions, a heuristic method has proven generally successful for approximation of some manifolds with boundary [5]. In a related article [14] on surface reconstruction for computer-aided geometric design, questions were posed about the possibility of creating algorithms for surfaces with boundary. The approach offered here is responsive, postulating new assumptions for input to be sufficient for approximation of surfaces with boundary. While some methods [3, 5, 9] have had some success reconstructing surfaces with boundary, the scope of the class of admissible surfaces with boundary was not definitively articulated and was presented as a result of experimental observation. This work specifically provides new theory that all C^2 surfaces with boundary can have arbitrarily close ambient isotopic reconstructions (dependant, of course, upon sufficient sampling density, corroborating other published results [9]) and shows resultant practical implementations. These considerations about boundary are critical for practical engineering application.

As initial motivation for the practical value of our results, we refer to Figure 3, below. The object to be reconstructed is a cylinder with boundary curves at both ends, having no top or bottom. The left side of Figure 3 is representative of the results that could be expected from many contemporary algorithms, whereas the right side of Figure 3 shows our significantly improved boundary reconstruction. The normals and tangents of the surface are used to define sample points on the envelope around the boundaries. The reader will also note the differing tessellations between the left and right images, and this remains the object of further study. (Similar differences for the tessellations are also evident in Figures 4.)

The paper is organized as follows: In Section 2, we summarize related work. Section 3 provides an overview of the theory and extensions. Implementation and graphical experiments of these new techniques are presented in Section 4. Section 5 presents our algorithm. Section 6 contains a discussion of hypotheses about admissible input data for our techniques. Section 7 presents observations about the influence of normal approximations and sampling densities. Closing remarks are given in Section 8.

2. Related Work

An emphasis upon topological guarantees for surface approximants has recently appeared in the literature on surface reconstruction [4, 7, 14]. For surface reconstruction, it is typical that only point cloud data is assumed to be available, while the methods presented here rely formally on maximal curvature and minimal separation distances. For our foot and bunny examples, which are based upon point cloud data, we estimated these geometric values in order to perform an improved surface reconstruction. These curvature and separation values will often be available for reverse engineering of manufactured objects [19]. While our reliance upon these values remains the subject of further study, the methods presented here will be directly applicable in many graphics applications, when the surface definitions will already be given and the relevant problem will be to produce a topologically correct approximation. For instance, related earlier work by some of the present authors [8] has been used to prevent undesirable topological changes during object deformations [16] for animations. The methods presented here will provide even more general criteria for an animator to preserve the critical topological characteristics of an object as it changes across successive frames. Some previous topological guarantees relied upon knowledge of the medial axis [3, 4, 11, 14, 15], which implicitly captures this curvature and separation information. We continue to look for unifying themes, but it should be clear that some estimation of a bound on surface curvature is crucial to any well-defined surface approximation method.

The theoretical concerns in providing topological guarantees for surface approximations near boundaries have been presented in the literature [5, 13, 14, 17] within the context of approximants created during surface reconstruction. The paper [13] presents a theoretical approach for using Voronoi diagrams on sampled data to detect boundary points and this is used as the basis for the CoCone algorithm and implementation. The CoCone software is available for download and it is shown to be effective on examples with boundary. The paper [5] presents a heuristic argument to reconstruct a surface with boundary, with a relevant example being the reconstruction of a foot. In a different approach [14], a similar example of a foot is reconstructed as a manifold without boundary to avoid undersampling problems often experienced near the boundary. Both of these approaches for boundaries [5, 14] were pragmatic responses to the known difficulties of reconstruction of boundaries from unorganized sample points. The CoCone approach [13] is specifically designed to handle boundaries and this code is also demonstrated on a foot reconstruction. As all the theory for all three approaches depends on sampling relative to the medial axis, it is worth noting that it has also been shown [17] that the typical sampling input for surface re-

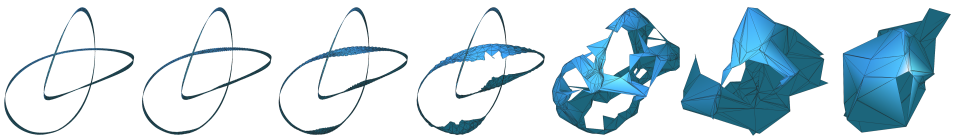

Figure 1. Trefoil A

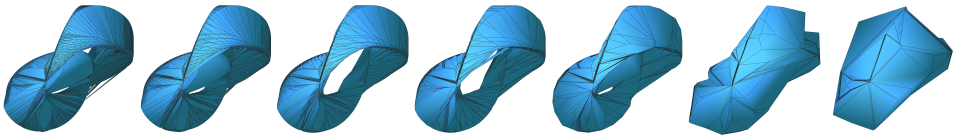

Figure 2. Trefoil B

construction is not sufficient, in general, to permit a topologically faithful reconstruction of the medial axis of the surface with boundary.

The value in preferring ambient isotopy for topological equivalence versus the more traditional equivalence by homeomorphism [26] has previously been presented [7, 23] and the interested reader is referred to those papers or to a standard mathematical text [18] for formal definitions. Intuitively, two closed curves will not be ambient isotopic if they form different knots, which can only be converted into each other by "untying" one knot and retying it to conform to the other, even while all knots are homeomorphic. For curves, a theorem has been published that provides for ambient isotopic piecewise linear (PL) approximations of a specifically described class of curves [20], inclusive of both those with and without boundary points, motivating the present investigation to surfaces with boundary. For our methodology, we define the terminology *the envelope of a surface*. The use of the term *envelope* has previously appeared with applications to tool-path generation for a specialized class of parametric surfaces [22]. The definition given there [22] is based upon isoparametric curves of the surface and is different from our usage. We expect that the contexts are sufficiently different that no confusion will result.

The present work emphasizes the integration of concepts from low-dimensional topology and differential geometry into the emerging sub-discipline of computational topology, as a complementary contribution to the incorporation of combinatorial topology and computational geometry formalisms that have already appeared [6, 12].

3. Preliminaries and Theory

In order to keep this section short, we refer the reader to standard definitions of a manifold with boundary [10], which are also summarized by the present authors in technical reports available on-line [1, 2]. As an intuitive overview, it suffices to observe that the differentiable properties along the boundary must follow as continuous limits of the corresponding differentiable properties within any neighborhood of a point on the boundary. In essence, this means that each compact C^2 manifold M, with boundary, can be considered as a submanifold of a compact C^2 manifold N without boundary.

The following definition of an envelope of a surface is central to our approach. Its use was motivated by a careful examination of the proofs previously presented [7, 23, 24] for reconstructing C^2 manifolds without boundary, which revealed a critical reliance upon a positive minimum distance between a surface and its medial axis. This has previously been proven for C^2 surfaces [7], but the extension here required showing that there also be a positive minimum between the envelope of a C^2 surface (as defined below) and the medial axis of this envelope. It is easy to show that this envelope is a surface without boundary, but, in general, the envelope will not be C^2. However, we were able to show that this envelope had sufficient smoothness to still conclude that there was a positive minimum distance between the envelope of a C^2 surface and the medial axis of this envelope. (The smoothness condition is stronger than C^1 and is known as $C^{1,1}$. For more details the interested reader is referred to our theory pre-print [2] .) While algorithms for computing this lower bound are still evolving, our prototype software suggests that these algorithms will have many performance and stability advantages over algorithms to approximate the medial axis. In the definition of the envelope, below, the value of ρ is less than the positive lower bound on the distance between the envelope and its medial axis. A subtle distinction about the new theory presented here is that it does *not* depend upon an *explicit* calculation of the medial axis.

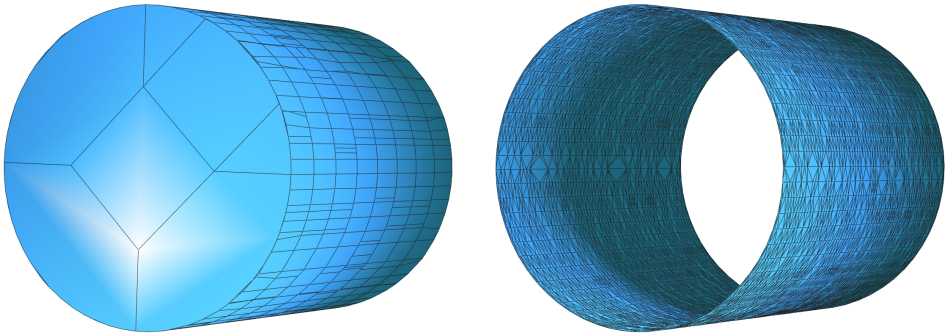

Figure 3. Comparison of Methods for Cylinder

Definition 3.1 *For suitably chosen values of $\rho > 0$, the* ρ-envelope *of M, denoted $E_\rho(M)$ is defined as*

$$E_\rho(M) = \{p \in \mathbf{R}^3 : d(p, M) = \rho\} \ .$$

It is *not* necessary to assume that M is orientable for our definition of the ρ-envelope, as given here. (A typical example of a unorientable surface with boundary is a Möbius strip.) The following theorem justifies the role of the envelope and its proof is presented in related pre-prints [1, 2], which also provide the bounds on ρ.

Theorem 3.1 *If M is C^2, then, for any $\epsilon > 0$, there exists a sufficiently small value of ρ such that its ρ-envelope has a minimum positive distance to its medial axis so that it is possible to explicitly define an ambient isotopic PL approximation to M via the nearest point mapping, where the distance between M and its approximation will be strictly less than ϵ.*

4. Computational Examples

The details of our theory presented in our pre-print [2] show how to create approximants that are ambient isotopic to $E_\rho(M)$, as well as approximants that are ambient isotopic to M. The examples presented here were motivated by that theory. They were created with new code as a pre-processing interface to the Power Crust algorithm in order to produce ambient isotopic approximations to $E_\rho(M)$. Complete adherence to the theory of our companion paper would have also required implementation of post-processing code to extract a subset of the Power Crust output to be ambient isotopic to M. This additional code has not yet been implemented. The examples presented here demonstrate a viable alternative to that full implementation. Namely, the component of the medial axis of $E_\rho(M)$ that lies interior to $E_\rho(M)$ is equal to M. Since the Power Crust also produces an approximation of this interior component of the medial axis of $E_\rho(M)$, this approximation is taken as an approximation of M. For the examples presented that compare our results with direct reconstructions from the Power Crust, we reiterate, in fairness, that the Power Crust algorithm was not designed to accept surfaces with boundary.

To explain our technique, experiments that were performed on NURBS surfaces are presented first in this paper. The techniques developed on the NURBS surfaces were then applied to challenging sets of point cloud data [14, 25] and our improvements are discussed. This approach permitted a controlled environment to analyze the results obtain by the envelope technique. All the information necessary to produce an envelope may be found analytically in a NURBS surface representation. The normals, partial derivatives and maximum curvature can be readily obtained to produce a precise envelope. This information, together with an estimate of the minimum feature size [3], then can guide the sampling rate to guarantee an ambient isotopic approximation similar to techniques already discussed in the literature [7] which are extended in our companion theory paper [2]. The examples presented here show that an accurate envelope construction will yield a faithful and desirable reconstruction.

4.1. Foot Example

The example presented here is a challenging one already seen in the literature. One method specifically is designed to reconstruct its boundary [13]. Another used an heuristic approach to respect the boundary [5] and an alternate method was presented to close off that boundary [14]. Here, no surface definitions were known in advance (in contrast to the other examples presented) and the point cloud data was provided by the previously cited author [14].

Figure 4 has two images. On the left is a direct reconstruction of the foot from the sample points provided using the Power Crust algorithm. This image also has an enlarged view of the boundary region near the ankle, where there are many artifacts which result in a closed surface.

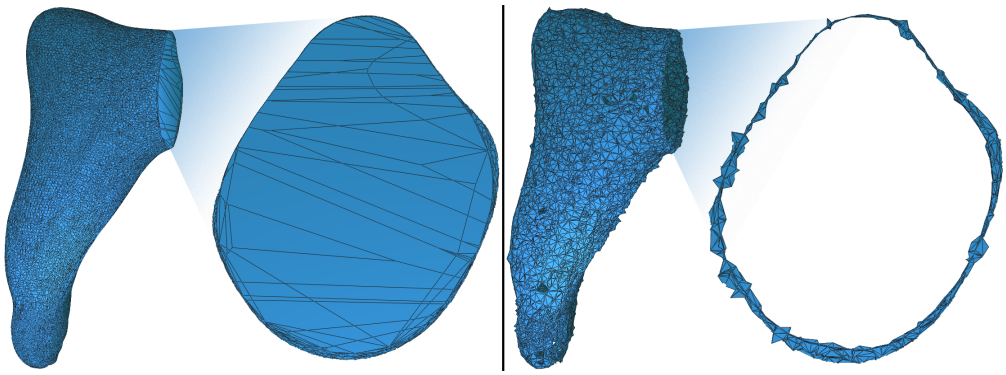

Figure 4. Comparison of Methods: Foot Data

The right image shows a reconstruction of the foot using the same original sample points as input to the pre-process that builds the envelope of this data. Again, there is a closer view of the boundary region near the ankle, showing that the boundary is more faithfully reconstructed.

Figure 5 has four images. The top left shows the original sample points for the foot, where these points were measured by a laser scan of the actual foot, and then their (x, y, z) co-ordinates were recorded in a text file. The top right shows the polar balls produced by the Power Crust, representing the radial field of the approximated medial axis of the point cloud. In this top-right image, noise is evident near the toes. The bottom left shows a sampling of poles determined from the Power Crust algorithm. The poles approximate normals to the original surface. The bottom right shows a point cloud representation for the envelope enclosure for the original point cloud. Since surface normals are central to the definition of the envelope and none are explicitly available here, envelope points are determined along the poles at a distance from the medial axis that is equivalent to the radius of the polar balls and offset in both directions. This foot envelope was constructed adaptively, where we experimented with varying the radius with location of the sample point in a modification of our definition of the envelope. This results in a tighter envelope around the toes and a slightly more generous envelope around the ankle. Along the boundary, additional points are created with user specified normals and tangents, appropriate to the envelope construction. This aspect currently remains within the judgement of the user, but the success of these experiments leads us to further investigate the theoretical constraints that would be involved formalizing this adaptive technique.

4.2. Split Bunny Example

It was also of interest to see how the techniques reported here would scale to large sets. Hence, the point cloud from the 'Stanford bunny' was considered [25]. In order to test our software, the bunny data set was partitioned to produce two surfaces, each having a very visible boundary. Our envelope techniques were then invoked along these newly introduced boundaries (The five holes in the base were considered to be small enough to be ignored in our analysis and reconstruction.) Each piece was then reconstructed with our approach, as shown in Figure 6. One can observe some missing triangles, for instance, near the tail. Further experimentation with our method and detailed comparisons to other techniques may result in future refinements.

4.3. Discussion of Input Needed and Final Output

Our foot reconstruction presented here significantly improves the foot boundary (near the ankle) and generally compares well with previous methods [5, 13]. Our results are generally comparable to those for CoCone [13]. These comparisons are subtle, though, and remain the object of future study. Our primary advantage is a reliance upon provable techniques over a well-defined class of permissible input surfaces. The CoCone software has the flexibility of variable ranges for user adjustments to the reconstruction. Ours relies on one parameter, which is defined in terms of curvature and minimal separation distance between points on the original manifold. Even when these cannot be directly computed, it is often possible to find positive lower bounds, which is sufficient. The relative advantages of these two differing approaches to reconstruction on manifolds with boundaries warrants further investigation.

In a further comparison between our methods and CoCone, we also attempted to reconstruct the trefoil knot depicted in Figure 1. There is an input parameter to CoCone for "flatness" which guides analysis of the Voronoi cells for determining the boundaries of the surface [13]. However, the authors of this current paper have not been able to choose any set of CoCone input parameter values to produce a completely correct topological result of the trefoil. There are still other input variables that may be adjusted, but

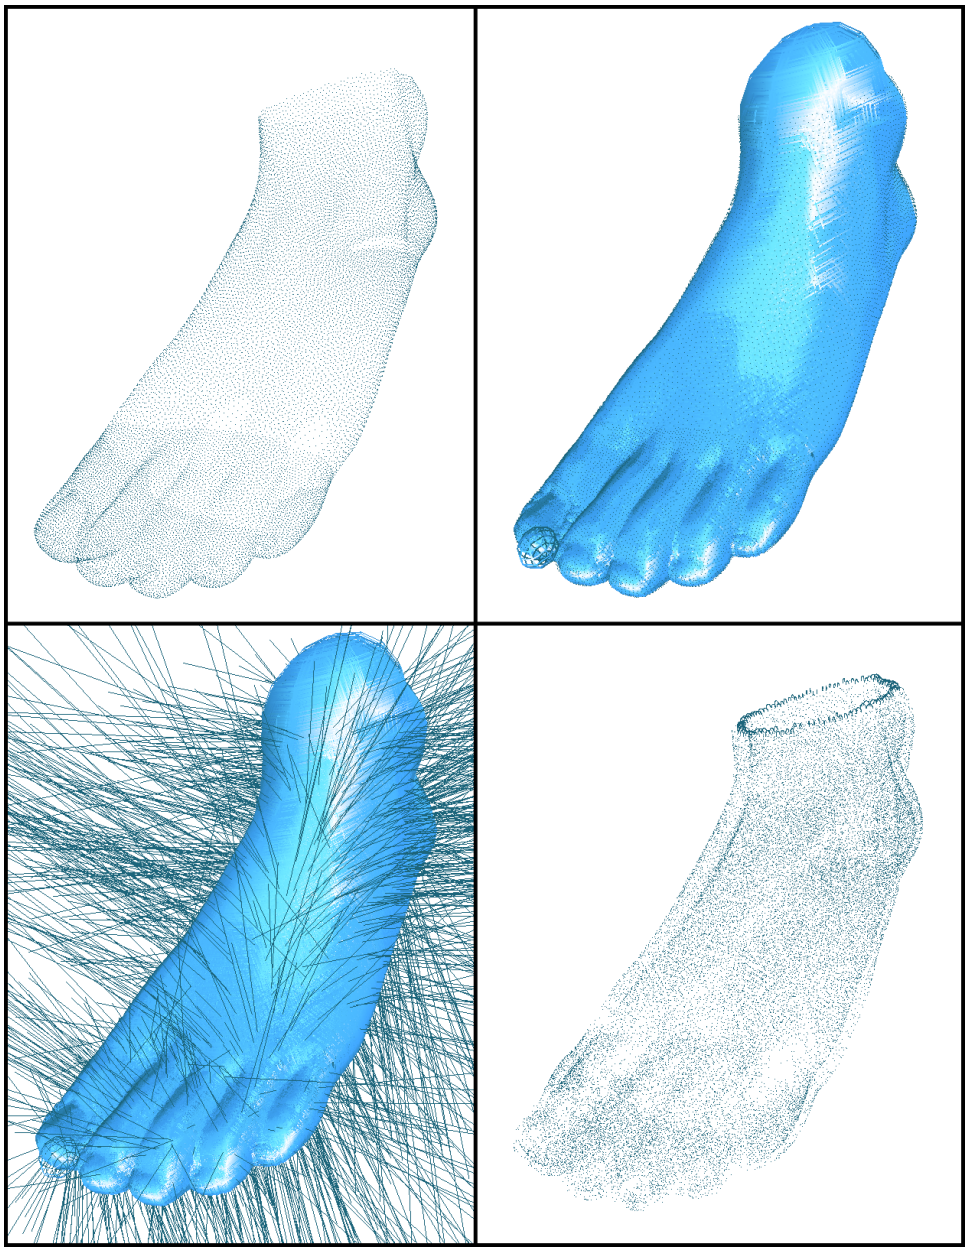

Figure 5. Stages of Method: Foot Data

it remains uncertain whether this input combinatorial problem can be easily solved to produce a correct result, even in this simple case. Further study is warranted, but an advantage of our work is that our algorithm does not require any user inter-action. A representative set of images from our CoCone studies are given in Figure 7.

This Figure 7 was created with the densest sample set that was used with the earlier experiments for Figures 1 and 2. The variation from left to right was in terms of input parameters to the CoCone software. The parameters were radius and flatness. The radius was fixed at 0.01 and the flatness varied from left to right, as follows: 1.2, 2.4, 4.8, 9.6, 28.4 This appears to show some advantage for our techniques on objects which have topologically complicated embeddings, like knots. In all the images shown here, there are occurrences of spurious geometry near crossings that do not occur in the left-most image of Figure 1. In experiments undertaken, the number of artifacts decreases as the flatness value is increased; however, these artifacts were not eliminated and at the highest value holes begin to appear in the reconstruction.

The other primary approach to this problem [14] is not

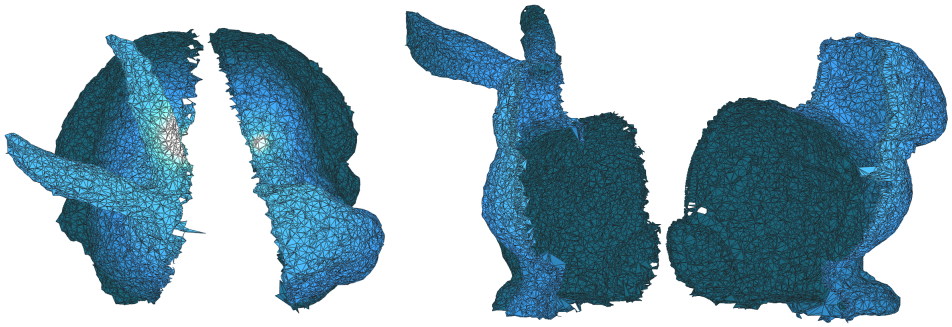

Figure 6. Split Bunny

Figure 7. Trefoil by CoCone

directly comparable, as it eliminated this boundary, whereas we preserve it. However, the results of Subsection 4.1 can be compared to the other experiments to show that our reconstruction is sensitive to the absence of critical geometric data, which we attribute primarily to the need to approximate normals and ball radii to use in our envelope construction technique. Those approximated normals are shown in the lower left of Figure 5 in addition to the polar balls which indicate the accuracy of the medial axis approximation. Artifacts in the foot reconstruction appear in the form of holes and local maximum/minimum that are inconsistent with the original geometry, as seen in Figure 4.

5. Summary of Algorithm

The algorithm, its input and some example output have been presented. The algorithm is now summarized in the following pseudocode. We re-emphasize that our reconstruction algorithm avoids direct reliance upon calculation of the medial axis. For a compact, C^2 manifold M, denote as follows:

- S = a set of sample points of M, with appropriately chosen density, $\epsilon > 0$. (The density requirement is that for every point $x \in M$, there exists a point $s \in S$ within ϵ of s, denoted as $d(x, s) < \epsilon$.)
- λ = the minimum positive distance between M and its medial axis,
- $B_r(x)$ = 3-ball of radius r centered at x, with $r > 0$,
- $\hat{K}$ = an approximation to $E_\rho(M)$, created within the algorithm, below.
- ψ denotes the nearest point mapping,

$$\psi : M \to E_\rho(M).$$

(This mapping will be *into* $E_\rho(M)$, but not *onto* $E_\rho(M)$).

- Ψ denotes the mapping defined on M to extend ψ as follows: For each $x \in M$, determine $\Psi(x)$ as that unique point along the ray from x through $\psi(x)$ that is the nearest point of $\hat{K}$ to x.

Reconstruction of Manifold with Boundary

Input: S
1.Choose ρ such that $\rho \in (2\epsilon, \lambda)$;
2.For each $x \in S$, create $B_\rho(x)$;
3.Let $D = \bigcup_{x \in S} B_\rho(x)$;
4.Find ∂D as an approximation to $E_\rho(M)$;
5.Set $K = \Psi(M)$
Output: K

Comments:

1. Both D and ∂D are created, in the algorithm, above, so that ∂D is an approximation of $E_\rho(M)$. These approximations are practical computational steps that are not specifically mentioned within the companion theory paper [2]. Additional care must be taken to ensure that these approximations of D and ∂D, as well $\hat{S}$, are sufficiently well chosen so that $\hat{K}$ will be ambient isotopic to $E_\rho(M)$. The experiments reported in this pa-

per demonstrate that the techniques used here are reasonably robust.

2. That K is an ambient isotopy of M relies upon the function $\psi : M \rightarrow E_\rho(M)$ being a homeomorphism onto the image of ψ, which is then used to show that M and the image of ψ are ambient isotopic.
3. The approximation K will generally not be PL, as the image of ∂M will not be PL under Ψ. However, as discussed in the companion theory paper [2], once K is obtained, further approximations are possible to create a PL ambient isotopic approximation of M.
4. Some subtleties remain for creating a computational representation of K as the image of M under Ψ, both for the above algorithm and for the immediately preceding item #3. While it is easy to state the existence of this image, any computational representation would be based upon some curve approximation, which remains the subject of further study, where many possible techniques are available, particularly in the spline literature [21]. The specific choice of technique will depend on a good balance between topological considerations and efficient algorithms for the specific input data. These additional investigations are beyond the scope of the present paper, but remain important topics for further consideration.

Discussion: The value for ρ is, of course, estimated. Note that this is applicable both to surface approximation, as well as surface reconstruction. When surface definitions are available, such as the widely-used splines, then ρ can be found quite accurately via computation of curvature on C^2 surfaces together with standard numerical methods to estimate minimal point separations, where these values are fully explained in the companion theory paper [2]. Of course, there remains some error associated with these numerical computations, but the expectation is that these computations will be much better conditioned and more stable than approximations of the medial axis.

We note the our development of a full implementation of this algorithm is not yet complete, as doing so would require the creation of significantly more code for proper approximation of the boundary. One of the delicate open issues is the efficient, accurate approximation of normals along a boundary when only point cloud data is available. This is discussed further in an expanded version of this paper [1].

6. Discussion of Hypotheses

This section presents an example that shows our reliance upon the C^2 hypothesis of Theorem 3.1. Let $x = y^2$, for $y \in [0, 2]$ be rotated about the x-axis. We note the importance of our smoothness assumptions to our implementation. For instance, consider the image shown in the left of Figure 8. It is defined as a surface of revolution of the curve $y = x^{3/2}$. It is a surface with boundary (at the top); note that artifacts appear near the base during reconstruction, as shown in the magnification at the right of Figure 8. It is tempting to speculate that these artifacts appear because the surface fails to satisfy even mild smoothness assumptions (namely $C^{1,1}$) at the base, although the underlying algorithmic causes remain the subject of further investigation.

7. Envelope Bounds and Sampling Density

This section shows typical data of experiments done to better understand the roles of accurate normals and sampling density, both relative to the original manifold.

7.1. Approximation with Envelopes

Figure 9 shows how the resultant surface approximation varies with the accuracy of the approximation. The progression from left to right is of decreasingly close approximations by the envelope. The suggestion to investigate this relationship further arose from our previously discussed reconstruction with the foot data, where improving the envelope approximations along the boundary resulted in significant improvements to the final surface approximation.

7.2. Sampling Density of Knots

The knot surface reconstructions of this subsection are, also, all of surfaces with boundary. The intent is to create surfaces based upon the unknot and the trefoil knot. Those original surfaces were created by drawing each knot as a curve and then these curves were extended into surfaces by a linear extrusion which produced no self-intersections. Each surface is like a ribbon with the corresponding knot for its two boundary curves.

Figure 10 shows the expected pattern of the reconstruction improving with higher sampling density, depicted for the unknot. In this unknot surface example, its envelope surface was then constructed at varying radii, $\lambda_1 < \lambda_2 < \ldots < \lambda_6$, increasing from left to right, while the sampling density of points from each envelope was kept constant. On the extreme right of the series of images, denote the radius of this envelope as λ_6. The value of λ_6 is sufficiently large that there is a perceptible artifact towards the center of this image, where there appears to be a self-intersection or an undercrossing in $\mathbf{R}^3$, although none should occur. Likely, this was caused by having the value of λ_6 exceed the value for ρ, as given in Theorem 3.1, but more precise numerical studies are needed to verify this condition. For comparison, if one views the images in middle of this sequence, the

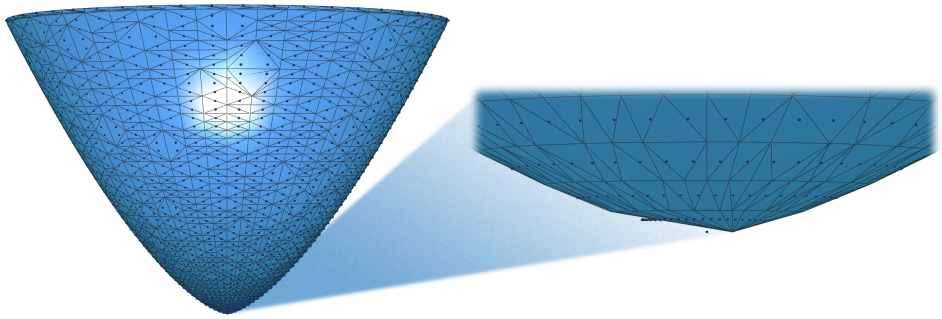

Figure 8. Reconstruction Artifacts Observed

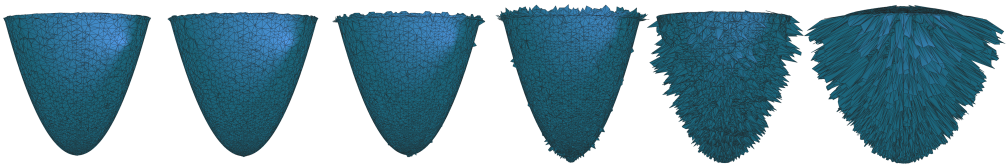

Figure 9. Envelope Proximity

smaller values of λ yield better images of the unknot. Proceeding to the left-most image, its radius of λ_1 is so small that the least feature size criterion of the Power Crust algorithm would require a much finer sampling density of the envelope than the constant density that is being maintained. Since the sampling density is no longer sufficient, holes and other visual artifacts begin to appear in the resulting reconstruction. Hence, this study shows the balance required between sampling density and radius chosen for the envelope surface.

This unknot study then led to consideration of the more challenging trefoil knot surface, as a comparison of our envelope reconstruction method versus techniques that have already appeared in the literature. That visual comparison has already been presented in Figures 1 and 2.

8. Concluding Remarks and Future Work

Our surface reconstruction technique is demonstrated for C^2 manifolds with boundary, where the method is dependant upon definition and implementation of an auxiliary surface, called the envelope. The promising results achieved here were by an effective expedient. We used the component of the medial axis of an envelope surface that is contained in the interior of the envelope as an approximation to the original manifold. An approximation to this component is already produced by the Power Crust algorithm and the results presented here show this to be a good approximation in practice.

It can be shown that a compact C^2 manifold M, with boundary, is equal to the interior component of the medial axis of the ρ-envelope of M, denoted as $E_\rho(M)$, for suitably chosen values of ρ. Hence, in principle, the expedient used here is well-founded, but there remain some issues for further investigation. Namely, the Power Crust necessarily produces an approximation of the medial axis, so, if there are any deviations of this approximation from the true medial axis, then no formal topological guarantees can be given for the examples presented here. This remains the subject of further investigation, but the results presented here are promising that more detailed investigation will be fruitful.

The experiments conducted provide interesting information about the role of bounds for the envelope in reconstructing surfaces with boundary. Furthermore, the images produced help to visualize the interplay between preserving topological characteristics and required sampling density. Further work needs to be done on both these subjects, towards optimal sampling criteria, which is a subject of broad ongoing interest.

9. Acknowledgements

Partial funding for K. Abe was from NSF grants CCF 0429477 and CCR 0226504; for J. Bisceglio and D. R. Ferguson was from NSF grant DMS-0138098; for T.J. Peters from NSF grants CCF 0429477, DMS 0138098 and CCR 0226504; for A.C. Russell from NSF grants CCF 0429477 and CCR 0226504; for T. Sakkalis from NSF

Figure 10. Unknot

grants DMS-0138098, CCR 0231511, CCR 0226504 and from the Kawasaki MIT Chair. All statements in this publication are the responsibility of the authors, *not* of these funding sources. We thank T. Dey for providing the point cloud data for the foot example and N. Amenta for some helpful discussions.

References

[1] K. Abe et al. Computational topology for reconstruction of surfaces with boundary, part i: Applications. *Pre-print, www.cse.uconn.edu/˜tpeters*, pages 1–23, 2004.

[2] K. Abe et al. Computational topology for reconstruction of surfaces with boundary, part ii: Mathematical foundations. *Pre-print, www.cse.uconn.edu/˜tpeters*, pages 1–15, 2004.

[3] N. Amenta, M. Bern, and M. Kamvysselis. A new voronoi-based surface reconstruction algorithm. In *Proc. ACM SIGGRAPH*, pages 415 – 421. ACM, 1998.

[4] N. Amenta, S. Choi, T. Dey, and N. Leekha. A simple algorithm for homeomorphic surface reconstruction. In *ACM Symposium on Computational Geometry*, pages 213–222, 2000.

[5] N. Amenta, S. Choi, and R. Kolluri. The power crust. In *Sixth ACM Symposium on Solid Modeling*, pages 249–260. ACM, June 2001.

[6] N. Amenta et al. Emerging challenges in computational topology. In *Workshop Report on Computational Topology*. NSF, June 1999.

[7] N. Amenta, T. J. Peters, and A. C. Russell. Computational topology: ambient isotopic approximation of 2-manifolds. *Theoretical Computer Science*, 305:3–15, 2003.

[8] L.-E. Andersson, T. J. Peters, and N. F. Stewart. Equivalence of topological form for curvilinear geometric objects. *International Journal of Computational Geometry and Applications*, 10(6):609–622, 2000.

[9] J.-D. Boissonnat and S. Oudot. Provably good surface sampling and approximation. In *Eurographics Symposium on Geometry Processing*, pages 9–18, 2003.

[10] W. M. Boothby. *An introduction to Differentiable Manifolds and Riemannian Geometry-Second Edition*. Academic Press, New York, 1986.

[11] T. Culver, J. Keyser, and D. Manocha. Accurate computation of the medial axis of a polyhedron. In *Proceedings of Fifth Symposium on Solid Modeling and Applications*, pages 179–190. ACM, June 1999.

[12] T. K. Dey, H. Edelsbrunner, and S. Guha. Computational topology. In *Advances in Discrete and Computational Geometry (Contemporary Mathematics 223*, pages 109–143. American Mathematical Society, 1999.

[13] T. K. Dey and J. Giesen. Detecting undersampling in surface reconstruction. *http://www.cse.ohio-state.edu/˜tamaldey*, 2002.

[14] T. K. Dey and S. Goswami. Tight cocone: a water-tight surface reconstructor. In *Eighth ACM Symposium on Solid Modeling and Applications*, pages 127–134. ACM, June 2003.

[15] T. K. Dey, H. Woo, and W. Zhao. Approximate medial axis for cad models. In *Eigth ACM Symposium on Solid Modeling and Applications*, pages 280–285. ACM, June 2003.

[16] A. Gain and A. Dodgson. Preventing self-intersection under free-form deformation. *IEEE Trans. on Visualization and Computer Graphics*, 7(4):289–298, 2001.

[17] M. Gopi. On sampling and reconstructing surfaces with boundaries. In S. Wismath, editor, *Proceedings of the Canadian Conference on Computational Geometry*, pages 49–53, August 2002.

[18] M. W. Hirsch. *Differential Topology*. Springer-Verlag, New York, 1976.

[19] W. Macy. Personal communication. 2003.

[20] T. Maekawa, N. M. Patrikalakis, T. Sakkalis, and G. Yu. Analysis and applications of pipe surfaces. *Computer Aided Geometric Design*, 15(5):437–458, 1998.

[21] L. Piegl and W. Tiller. *The* NURBS *Book, 2nd Edition*. Springer, New York, NY, 1997.

[22] H. Pottmann and M. Peternell. Envelopes – computational theory and applications. In *Spring Conference on Computer Graphics 2000*, pages 3–23. Comenius University, Bratislava, May 2000.

[23] T. Sakkalis and T. J. Peters. Ambient isotopic approximations for surface reconstruction and interval solids. In *Eighth ACM Symposium on Solid Modeling and Applications*, pages 176–184. ACM, June 2003.

[24] T. Sakkalis, T. J. Peters, and J. Bisceglio. Isotopic approximations and interval solids. *CAD*, 36 (11):1089–1100, 2004.

[25] G. Turk. Stanford university computer graphics laboratory. *Stanford bunny, point cloud scan data set*, `graphics.stanford.edu/data/3Dscanrep/`, 1993.

[26] S. Willard. *General Topology*. Addison-Wesley Publishing Company, Reading, MA, 1970.