ELSEVIER

Computer-Aided Design 36 (2004) 1089–1100

COMPUTER-AIDED DESIGN

www.elsevier.com/locate/cad

Isotopic approximations and interval solids

Takis Sakkalis[a,b], Thomas J. Peters[c,d,*], Justin Bisceglio[c]

[a]*Agricultural University of Athens, Athens 118 55, Greece*
[b]*Massachusetts Institute of Technology, Cambridge, MA 02139, USA*
[c]*Department of Computer Science and Engineering, University of Connecticut, Storrs, CT 06269-1155, USA*
[d]*Department of Mathematics, University of Connecticut, Storrs, CT 06269-1155, USA*

Accepted 9 January 2004

Abstract

Given a nonsingular compact two-manifold F without boundary, we present methods for establishing a family of surfaces which can approximate F so that each approximant is ambient isotopic to F. The methods presented here offer broad theoretical guidance for a rich class of ambient isotopic approximations, for applications in graphics, animation and surface reconstruction. They are also used to establish sufficient conditions for an interval solid to be ambient isotopic to the solid it is approximating. Furthermore, the normals of the approximant are compared to the normals of the original surface, as these approximating normals play prominent roles in many graphics algorithms.

The methods are based on *global* theoretical considerations and are compared to existing *local* methods. Practical implications of these methods are also presented. For the global case, a differential surface analysis is performed to find a positive number ρ so that the offsets $F_o(\pm\rho)$ of F at distances $\pm\rho$ are nonsingular. In doing so, a normal tubular neighborhood, $F(\rho)$, of F is constructed. Then, each approximant of F lies inside $F(\rho)$. Comparisons between these global and local constraints are given.

Keywords: Ambient isotopy; Computational topology; Surface reconstruction; Interval solids; Offsets and deformations; Reverse engineering

1. Introduction and motivation

The problem of approximation of surfaces is of fundamental importance both in theoretical as well as in applied mathematics. In particular, in computer-aided geometric design (CAGD) it plays a crucial role in the discretization of the data for meshing applications. These applications generate, in general, piecewise linear (PL) approximations of the actual surface. In order for these approximations to be of practical, as well as of theoretical, value, it is often desirable to be within a tolerance given by the user. Simultaneously, it is also important to have topological equivalence via an ambient isotopy between the approximant and the original surface.

In this paper we propose global methods for creating a family of approximating surfaces to a given nonsingular compact surface F, while ensuring that each approximant is ambient isotopic to the original surface F. This, for example, is useful in graphics and animation, where lower degree surface approximants are often used for performance reasons, but the visual imperatives demand preservation of topological form. There can also be applications to surface reconstruction, as existing reconstruction methods only provide for PL approximations, while the proposed method is suitable for higher order approximations. Our primary tool is creating the offset of a surface. This is motivated by the recent work of Wallner et al. [30] in which some of the geometric and algebraic properties of offsets are explored. An *offset surface* is used in the construction of a normal tubular neighborhood of a manifold. Offset surfaces also have diverse potential applications in geometric modeling, such as in the construction of tolerance zones, the generation of tool paths for numerical control machining, etc. A second application is to interval solids, where sufficient conditions are given for an interval solid to be ambient isotopic to the solid that is being approximated.

The paper is organized as follows: in Section 2 we summarize related work. In Section 3 we present some fundamental definitions and basic results from differential topology and offset surfaces. The methods presented are

* Corresponding author. Address: Department of Computer Science and Engineering, University of Connecticut, Storrs, CT 06269-1155, USA.
E-mail addresses: tpeters@cse.uconn.edu (T.J. Peters), stp@aua.gr (T. Sakkalis).

0010-4485/$ - see front matter © 2004 Elsevier Ltd. All rights reserved.
doi:10.1016/j.cad.2004.01.008

based upon constraining the approximant to lie within a bounded offset of the surface, as determined by considerations from differential topology. To do so, a focal point is defined in Section 3.1, along with an intuitive explanation. Section 3.2 provides the constraints that are used to define a nonsingular offset surface and a tubular neighborhood, as tools in the approximation process. In particular, we give conditions on ρ so that the offset(s) $F_o(\pm\rho)$ of F are nonsingular. Section 4 gives a formal definition of ambient isotopy and contains the main theorem, giving criteria for a family of approximating surfaces to be ambient isotopic to a single surface. This section concludes with a discussion of how these results can be applied to improve the state of the art for surface reconstruction. Section 5 expands the topic to solids, by consideration of a surface bounding a finite volume. The particular focus is upon interval solids. In Section 6 a comparison between our methods and previously published techniques is presented. The error bounds on approximating normals for the PL case is presented in Section 7. Closing remarks are given in Section 8.

2. Related work

Classical aspects of topology [31] are emerging as valuable tools in solid modeling. As approximation is unavoidable in solid modeling, the question of whether an approximation is 'good enough' to preserve the essential features of the object is of central importance. Previous publications by the first two authors have invoked the concept of ambient isotopy, a topological notion of equivalence for admissibility of approximations to curves [21] and surfaces [7].

The notion of 'computational topology' [6] has been proposed primarily as the merging of combinatorial topology and computational geometry. Most work in computational topology to date [6,15] has ignored differentiability and approximation. To the contrary, the present work emphasizes the integration of general topology, differential topology and approximation.

The following description is the basis for preferring ambient isotopy for topological equivalence versus the more traditional equivalence by homeomorphism [31], and summarizes the justification previously presented [7]. While a formal statement of ambient isotopy is provided in Section 4, for present purposes the following informal notion will be sufficient. 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. Although any two simple closed planar curves are ambient isotopic, Fig. 1 shows two simple homeomorphic space curves, which are not ambient isotopic, because they describe different knots.[1]

[1] The different knot classifications of 0_1 and 4_1^m are indicated.

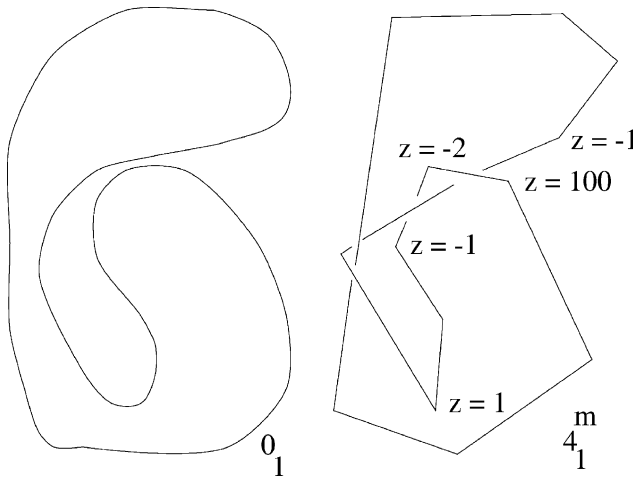

Fig. 1. Nonequivalent knots.

The smooth curve depicts the simplest closed curve, known as the unknot. The PL curve is an approximation of the smooth curve. In the right half of Fig. 1 the z-coordinates of some vertices are indicated to emphasize the knot crossings in $\mathbf{R}^3$ (All other end points have $z = 0$). All end points of the line segments in the approximation are also points on the original curve. Having this knotted curve as an approximant to the original unknot would be undesirable in many circumstances, such as graphics and engineering simulations. Similar pathologies can happen in approximating surfaces, but the work presented here can prevent these problems by appropriately constraining the approximations produced.

Earlier surface reconstruction algorithms guaranteed topological equivalence to the original surface by means of homeomorphisms [3,4,19].

A common technical tool for demonstrating an ambient isotopy of compact support is a function known as a 'push' [11]. A generalization of a push is used in the proof given in this paper.

The issue of rigorous proofs for the preservation of topological form in geometric modeling appears to have been initially raised within engineering design in problems regarding tolerances [12,13,29], but these papers did not directly propose ambient isotopy as a criterion. The class of geometric objects considered was appreciably expanded by theorems for ambient isotopic perturbations of models with spline boundaries [9,10]. For the simpler case of polygonal models, similar topology preserving approximations had been presented earlier under different technical terminology [8,14].

In response to the example of Fig. 1, a theorem was published that provided for ambient isotopic PL approximations of one-manifolds [21]. The proof utilizes 'pipe surfaces' from classical differential geometry [24]. The improved approximation is shown in Fig. 2. There is also

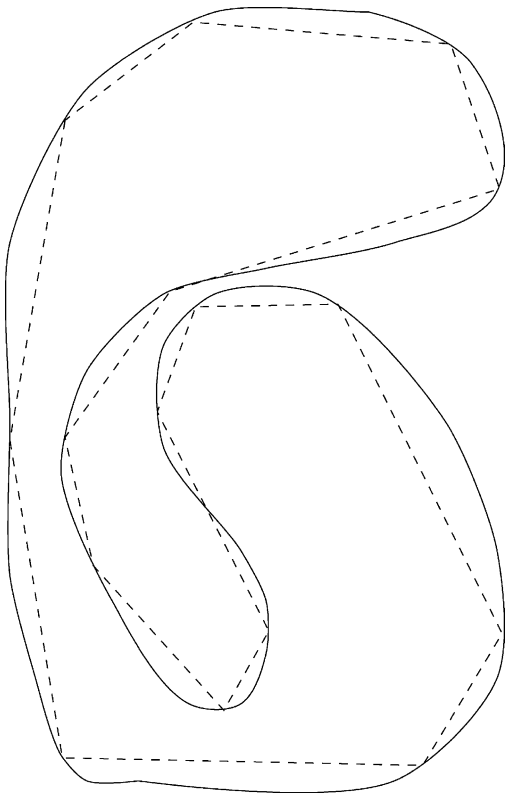

Fig. 2. Ambient isotopic approximation.

a publication of comparison of curves to α-shapes [16] via ambient isotopies [27].

3. Preliminaries

In this section we review some elementary facts about manifolds and offsets, taken from Refs. [22,30], that we will need in Section 4. In particular, we prove our initial basic result, Theorem 6, which is the key to the proof of Corollary 8. The theoretical results presented here were motivated by the pragmatic view that other surface reconstruction methods [2,5,7] depend upon a supporting computation of the medial axis, which must also be approximated from sampled data. The methods given here eliminate this intermediate approximation of the medial axis, which is often a formidable task. This may result in a need for denser sampling, but it is beyond the scope of the present paper to make a comprehensive comparison of the detailed implementations of these competing methods. The sampling criteria presented here remains admittedly implicit. Namely, if one wishes to reconstruct a surface from sampled data, then the surface must be sampled densely enough that the reconstructed approximant lies within $F_o(\pm\rho)$. This is consonant with methods previously developed for one-manifolds [21].

3.1. Differential topology: critical points

Let $f(x_1,\ldots,x_n):\mathbf{R}^n\to\mathbf{R}$ be a smooth function. A point $a\in\mathbf{R}^n$ is called a *critical* point of f if all first partial derivatives of f are zero at a; that is $(\partial f/\partial x_i)(a)=0$, for all i. Furthermore, a is called a *nondegenerate* critical point of f, if

1. a is a critical point of f, and
2. The matrix $Hf(a)=(\partial^2 f/\partial x_i\partial x_j)$ of the second partial derivatives of f at a is invertible. (This matrix is also known as the *Hessian* matrix of f).

Remark 1. A critical point of f that is not nondegenerate is called degenerate.

The notions of degenerate/nondegenerate critical points can be carried out on a real function on a manifold W of any dimension. Here, however, we shall specialize on manifolds of dimension 2. Let then $F\subset\mathbf{R}^3$ be a manifold of dimension 2, and let $g:F\to\mathbf{R}$ be a smooth function. A point $a\in F$ shall be called a critical point of g if the gradient vector $((\partial g/\partial x_1),(\partial g/\partial x_2),(\partial g/\partial x_3))$ is parallel to the unit normal vector $\mathbf{n}_a$ of F at a. Let a be a critical point of g, and let u_1,u_2 be local coordinates on F around a. Let $Hg(u)=(\partial^2 g/\partial u_i\partial u_j)$ be the Hessian matrix of g at a with respect to the coordinates u_1,u_2. Then, a is called degenerate/nondegenerate if $Hg(u)$ is singular/nonsingular, respectively.

Let $F\subset\mathbf{R}^3$ be as above and let $N\subset F\times\mathbf{R}^3$ be defined as

$$N=\{(q,v)\mid q\in F,\ v \text{ is perpendicular to } F \text{ at } q\}$$

It is easy to see that N is a three-dimensional manifold differentiably embedded in $\mathbf{R}^6$.

Let $E:N\to\mathbf{R}^3$ be defined as $E(q,v)=q+v$. (E is called the *endpoint* map.)

Definition 2. A point $e\in\mathbf{R}^3$ is *a* focal point of (F,q) with multiplicity μ if $e=q+v$ where $(q,v)\in N$ and the Jacobian of E at (q,v) has nullity $\mu>0$. The point e will be called a focal point of F if e is a focal point of (F,q) for some $q\in F$.

Intuitively, a focal point of F is a point of $\mathbf{R}^3$ where nearby normals intersect. Now consider the normal line L to F that goes through q and consists of all points $q+t\mathbf{n}_q$, where $\mathbf{n}_q$ is the unit normal of F at q and $t\in\mathbf{R}$. We then have:

Lemma 3. ([22], lemma 6.3, p. 34) The set of focal points of (F,q) along L is the set of points $q\pm K_i^{-1}\mathbf{n}_q$, where $i=1,2,K_i\neq 0$ are the principal curvatures at q.

3.2. Offset surfaces

Let $F \subset \mathbf{R}^3$ be an orientable[2] two-manifold which is C^2 (at each point of the manifold, the second derivative exists and is continuous); we will also call F a nonsingular surface in $\mathbf{R}^3$. We shall choose an orientation on F as follows: let $v \in F$ and let v_1, v_2 be a positively oriented basis for the tangent space TF_v, regarded as a subspace of $\mathbf{R}^3$. We say that F has the *positive* orientation at v if $\det[\mathbf{n}_v, v_1, v_2] > 0$, where $\mathbf{n}_v$ is the unit surface *normal* of F at v. Then F has positive orientation, if it has positive orientation at each of its points. (The term of *negative* orientation is defined similarly, with the obvious change of sign.) Let $\rho \in \mathbf{R}$. Then, the positive offset $F_o(\rho)$ of F and the negative offset $F_o(-\rho)$ of F are defined, respectively, as:

$$F_o(\rho) = \{x + \rho\mathbf{n}_x \mid x \in F\} \text{ and}$$

$$F_o(-\rho) = \{x - \rho\mathbf{n}_x \mid x \in F\}. \tag{1}$$

A simple geometric interpretation of the offset $F_o(\rho)$ is the surface locus swept out by the center $x + \mathbf{n}_x$ of a sphere of radius $|\rho|$ as the sphere rolls over every point x of F, where the $\mathbf{n}_x$ are consistently oriented relative to F.

In this subsection we will prove several results concerning offsets. In particular, we will give *necessary* and *sufficient* conditions on ρ so that $F_o(\rho)$ is *nonsingular* in terms of certain differential and geometric considerations of the surface F. (Clearly, these results apply similarly to $F_o(-\rho)$.) We will restrict our attention to compact manifolds *without* boundary.

Our first result comes as a direct application of Lemma 3.

Proposition 4. *Let $F_o(\rho)$ be as in Eq.* (1) *and let*

$$g : F \to F_o(\rho) \subset \mathbf{R}^3, \qquad g(x) = x + \rho\mathbf{n}_x \tag{2}$$

Then, the Jacobian of g at x has nullity $\mu > 0$ if and only if $x + \rho\mathbf{n}_x$ is a focal point of F.

Proof. The proof is a slight modification of the proof of Lemma 3. □

Proposition 4 shows that $F_o(\rho)$ is locally a two-dimensional manifold at $x + \rho\mathbf{n}_x$ precisely when $x + \rho\mathbf{n}_x$ is not a focal point of F.

We will proceed with the question when g is *globally* 1–1. For this, we define the map

$$G : F \times F \to \mathbf{R}, \qquad G(x, y) = \|x - y\|^2. \tag{3}$$

[2] For the context of this paper, we are considering two-manifolds embedded in $\mathbf{R}^3$. Since nonorientable two-manifolds without boundary can only be embedded in $\mathbf{R}^n$, for $n \geq 4$ [18, theorem 4.7], the additional assumption of orientability leads to no loss of generalization in the present proofs.

Obviously, $G(x, y) > 0$, for $x \neq y$. Second, note that G has a critical value $r > 0$ since $F \times F$ is compact. Let $x, y \in F$ with $G(x, y) = r$. Then, it is easy to see that $\mathbf{n}_x = \pm\mathbf{n}_y$ and the vector $x - y$ is parallel to both $\mathbf{n}_x$ and $\mathbf{n}_y$. Third, we claim that if

$$c = \inf\{r > 0 \mid r, \ \ r \text{ is a critical value of } G\} \tag{4}$$

then c is positive. For if c were to be zero, we would then have that for each arbitrarily small positive δ there exists a pair of points (x_δ, y_δ) so that

- $\|x_\delta - y_\delta\| < \delta$ and
- the vector $x_\delta - y_\delta$ is parallel to both $\mathbf{n}_{x_\delta}, \mathbf{n}_{y_\delta}$.

Since δ can be arbitrarily small, then x_δ, y_δ have to belong to the same component of F. Now let a be any point of F. Using the implicit function theorem, we may choose local coordinates u_1, u_2 around a so that F becomes the graph of a smooth function $v = h(u_1, u_2)$ with $(0, 0, h(0, 0)) = a$ and $\partial h/\partial u_1 = \partial h/\partial u_2 = 0$ at $(0, 0)$. For a point $b \in F$ near a we choose a corresponding point (b_1, b_2) close to $(0, 0)$ so that $b = (b_1, b_2, h(b_1, b_2))$. The normal vector of M at a is $(0, 0, 1)$ while the normal of F at b is $((\partial h/\partial u_1), (\partial h/\partial u_2), 1)$, where the partials are evaluated at (b_1, b_2). Then,

$$b - a = (b_1, b_2, h(b_1, b_2) - h(0, 0)),$$

and thus the vector $b - a$ is *not* parallel to $(0, 0, 1)$, which is the normal of F at a, for $a \neq b$. This proves the claim.

Definition 5. Let $x, y \in \mathbf{R}^3$, and $X, Y \subset \mathbf{R}^3$. We define $d_M(X, Y) = \max_{x \in X} \min_{y \in Y} \|x - y\|$,
For $a \in \mathbf{R}^3$, a point $s \in X$ is a nearest point on X to a if

$$\|a - s\| = \min\{\|a - t\| \mid t \in X\}.$$

We now consider a $\rho > 0$ so that

C1 For each point x of F neither principal curvature of F at x is equal to $\pm 1/\rho$, and
C1 $2\rho < c$, where c is as in Eq. (4).

We then have the following:

Theorem 6. *For ρ as above, and g as in Eq.* (2), *g is an* 1–1 *map.*

Proof. For a positive number ϵ we may define the open set

$$F(\epsilon) = \{x \in \mathbf{R}^3 \mid d(\{x\}, F) < \epsilon\}.$$

Using the ϵ-Neighborhood Theorem, [17, p. 69], we may find an ϵ with the following properties:

- Each point $w \in F(\epsilon)$ possesses a *unique* nearest point in F, denoted by $\pi(w)$, and
- the map $\pi : F(\epsilon) \to F, w \to \pi(w)$, is a submersion.

Now note that if ρ is chosen to satisfy **C1**, **C2** and is less than ϵ, then g is 1–1. For if $z \in F_o(\rho)$ then its distance from F is exactly ρ. Thus, if $z = x + \rho\mathbf{n}_x = y + \rho\mathbf{n}_y$, then $\|z - x\| = \|z - y\| = \rho$. Therefore, $\pi(z) = x = y$.
Suppose now that ρ is the *first* positive number for which g is not 1–1. Let then $x \neq y$ for which $g(x) = g(y)$; that is $x + \rho\mathbf{n}_x = y + \rho\mathbf{n}_y$. Then, $x - y = \rho(\mathbf{n}_y - \mathbf{n}_x)$. We claim that in this case (x, y) is a critical point of G. For if not, the locally two-dimensional manifold $F_o(\rho)$ has a nontangential self-intersection at $g(x)$, and that contradicts the choice of ρ. Thus, $\|x - y\| = 2\rho$. But $\|x - y\| \geq c$ and this is a contradiction to **C2**. □

Finally, we have:

Corollary 7. *Let F be a compact orientable surface without boundary. Then $F_o(\rho)$ is a nonsingular surface if ρ satisfies conditions* **C1** *and* **C2**.

Corollary 8. Let $\rho \in (0, \infty)$ *be so that ρ satisfies conditions* **C1** *and* **C2**. *Then, the open set $F(\rho)$ is a normal tubular neighborhood of F. In addition, if r is such that $0 \leq r \leq \rho$, then,*

- *The offsets $F_o(r)$ and $F_o(-r)$ are nonsingular.*
- *Every point $w \in F(\rho)$ has a unique nearest point $\pi(w)$ in F.*
- *Let π_r be the nearest point function $\pi_r : F(r) \to F$, $\pi_r(v) = \pi(v)$. Then, for every point $x \in F$ the set $\pi_r^{-1}(x)$ is equal to $(x - \mathbf{n}_x, x + \mathbf{n}_x)$.*

The concept of a tubular neighborhood of a submanifold without boundary is not new; a result similar to the above, stating its existence, appears in [18, theorem 5.2]. Here, however, an explicit numerical bound on the size of the neighborhood is presented, which is useful in computational applications. One can visualize a tubular neighborhood of F as follows: suppose that F is made out of thin rubber. Then, by uniformly inflating and deflating the *interior* of F so that no singularities occur in F, a tubular neighborhood is nothing but the union of the volumes created by the inflation and deflation of F.

4. Ambient isotopic approximations

Many approximation schemes for surfaces are concerned with the existence of a homeomorphism between the actual surface and its approximant. However, the latter does not guarantee that the surface and its approximant have the same embedding within $\mathbf{R}^3$. As an example of different embeddings in $\mathbf{R}^3$, consider the standard torus and a variant of it. Let T denote the regular torus $T = S^1 \times S^1$, where S^1 is the unit circle. Let KT denote a knotted torus, $KT = S^1 \times K$, where K is a trefoil knot, chosen so that KT is homeomorphic to T. However, within $\mathbf{R}^3$, one cannot continuously deform T into KT. This is demonstrated by showing that the spaces $\mathbf{R}^3 - T$ and $\mathbf{R}^3 - KT$ do not have the same homotopy type [23, p. 103, theorem 1].

Two related notations are defined here for use within the rest of the paper. Both concepts are standard in general topology and can be found in any basic topology text [31].

For any topological space X and any subset A in X, the notation $\mathrm{cl}_X A$ refers to the closure of A in X. For any topological space X and any subset A in X, the notation $\mathrm{Int}_X A$ refers to the interior of A in X. In cases where the particular space X being used is obvious from context, then the subscript is typically deleted, as is done in this paper, where $X = \mathbf{R}^3$.

The following definition (Fig. 3) is central to the rest of the paper and it gives precise meaning to when two objects are both homeomorphic and have the same embedding within $\mathbf{R}^3$.

Definition 9. Let X, Y be subsets of $\mathbf{R}^3$. Then we say that X and Y are *ambient isotopic* if there is a continuous function

$$H : \mathbf{R}^3 \times [0, 1] \to \mathbf{R}^3$$

such that for each $t \in [0, 1], H(\cdot, t)$ is a homeomorphism from $\mathbf{R}^3$ onto $\mathbf{R}^3, H(\cdot, 0)$ is the identity and $H(X, 1) = Y$.

In the illustrative Fig. 3, the set X is continuously deformed into Y, while simultaneously, the region bounded by X and the horizontal dashed line segment is being deformed into the region bounded by Y and the horizontal dashed line segment. Thus, the smaller region is being stretched into a larger one. At the same time the region bounded by X and the two nonhorizontal line segments is being deformed into the region bounded by Y and the two nonhorizontal line segments, contracting a larger region into a smaller one. These complementary expansions and contractions allow for the definition of a continuous map that fixes all points on the boundary of and exterior to the triangle.

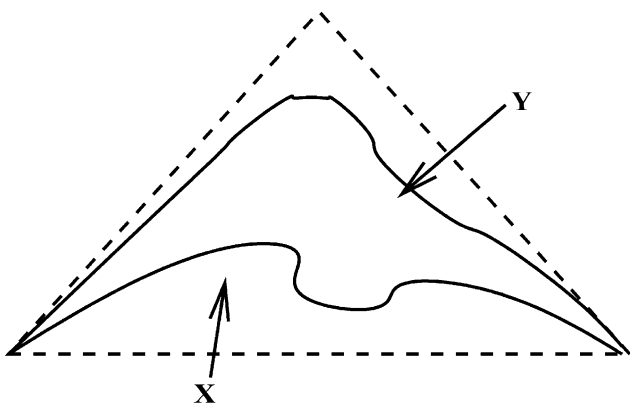

Fig. 3. Ambient isotopy.

The techniques of offsets and tubular neighborhoods given in Section 3 enable us to establish sufficient criteria for ambient isotopy of certain surfaces. In that regard, we term our results as *global* in the sense that the surface is perturbed globally.

Prior to the statement and proof of the main theorem of this paper, it is helpful to define two functions on intervals. These will be combined to form an ambient isotopy. Both are similar to the push function [11]. Let F be our surface and let ρ be as in Corollary 8. For a point $x \in F$, let $I(x,\rho)$ be the closed line segment that connects the points $x_\rho^- = x - \rho\mathbf{n}_x$ and $x_\rho^+ = x + \rho\mathbf{n}_x$.

Let $x \in F$ and let $q \in I(x,\rho)$ be a point different than x_ρ^- and x_ρ^+. Then we define the homotopy $\check{f}_q : [x_\rho^-, x] \times [0,1] \to I(x,\rho)$ that keeps x_ρ^- fixed, by

$$\check{f}_q([x_\rho^-, x], 0) = [x_\rho^-, x],$$

$$\check{f}_q([x_\rho^-, x], 1) = [x_\rho^-, q],$$

and $\forall t \in (0,1)$, the image of $[x_\rho^-, x]$ is of the form $[x_\rho^-, (1-t)x + tq]$. Similarly, we define the homotopy $\hat{f}_q : [x, x_\rho^+] \times [0,1] \to I(x,\rho)$ that keeps x_ρ^+ fixed, by

$$\hat{f}_q([x, x_\rho^+], 0) = [x, x_\rho^+],$$

$$\hat{f}_q([x, x_\rho^+], 1) = [q, x_\rho^+],$$

and $\forall t \in (0,1)$, the image of $[x, x_\rho^+]$ is of the form $[(1-t)x + tq, x_\rho^+]$.

The following is the main result of this section, and provides a means of establishing a family of surfaces ambient isotopic to F.

Theorem 10. *Let $F \subset \mathbf{R}^3$ be a compact nonsingular two-manifold without boundary, and ρ as in* Corollary 8. *Let also $W \subset \mathbf{R}^3$ be a compact two-manifold without boundary that satisfies the following*:

- *$W \subset F(\rho)$, and*
- *For every $x \in F$, $I(x,\rho)$ intersects W at precisely one point, denoted as $w(x)$.*

Then, W is ambient isotopic to F and $d_M(F,W) < \rho$ (Fig. 4).

Proof. Since F is compact, it is clear that $\partial F(\rho)$ is compact. Furthermore, since $F(\rho)$ is open and $W \subset F(\rho)$, it is also clear that $d(W, \partial F(\rho)) > 0$.

To define the ambient isotopy, it is sufficient to consider, for each $x \in F$, the behavior of the ambient isotopy along the interval originating at x_ρ^- and terminating at x_ρ^+. This is done simply by the following piecewise definition.

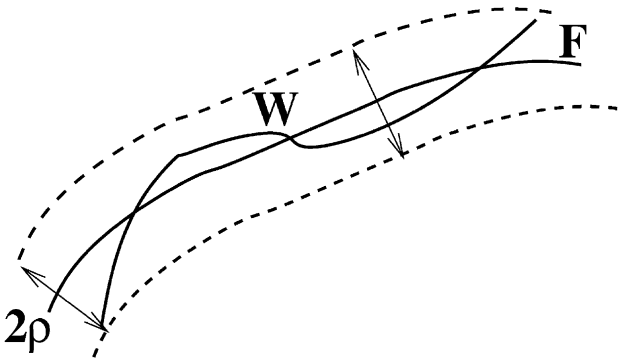

Fig. 4. Intersecting W in precisely one point.

Define the function $H : \mathbf{R}^3 \times [0,1] \to \mathbf{R}^3$ by

$$H(z,t) = \begin{cases} z & \text{if } z \in \mathbf{R}^3 - F(\rho),\ \forall t \in [0,1], \\ \check{f}_{w(x)}(z, 2t), & \text{if } x \in F,\ z \in [x_\rho^-, x],\ t \in [0, 1/2], \\ \hat{f}_{w(x)}(z, 2t-1), & \text{if } x \in F,\ z \in [x, x_\rho^+],\ t \in [1/2, 1]. \end{cases}$$

Note that the piecewise definition of H agrees on $\partial F(\rho)$. It is obvious that H is a homotopy such that $H(z,0) = z$, $\forall z \in \mathbf{R}^3$ and $H(F,1) = W$. To complete the proof that H is an ambient isotopy, it suffices to show that $H(\cdot,t)$ is a homeomorphism for all $t \in [0,1]$. However, since for each $t \in [0,1]$, the continuous function $H(\cdot,t)$ is the identity outside the compact set $\mathrm{cl}_{\mathbf{R}^3} F(\rho)$ (and is also the identity along the boundary of $\mathrm{cl}_{\mathbf{R}^3} F(\rho)$), it only remains to show that $H(\cdot,t)$ is $1-1$ for all $x \in \mathrm{cl}_{\mathbf{R}^3} F(\rho)$. But, this follows easily, since for all $x, y \in F$, with $x \neq y$, we have, by hypothesis, that the corresponding intervals are disjoint, namely,

$$[x_\rho^-, x_\rho^+] \cap [y_\rho^-, y_\rho^+] = \varnothing.$$

The bound given on the distance between W and F follows directly from the containment of W within $F(\rho)$. □

This completes the proof of a demonstration of a particular ambient isotopy, where the proof presented is similar to classical arguments [18, chapter 5], with the additional information provided here of a numeric bound on the size of the neighborhood containing the ambient isotopic images, in contrast to the classical arguments merely proving the existence of some neighborhood.

Corollary 11. *Let F, ρ be as in* Theorem 10. *Then for every r such that $0 \le r \le \rho$, $F_o(r)$ is ambient isotopic to F and $d_M(F, F_o(r)) \le \rho$.*

The above results provide an abundance of surfaces W ambient isotopic to F. Note that any such W is inside the tubular neighborhood $F(\rho)$ and within tolerance ρ from F. Existing methods [7] produce one PL approximant, whereas Corollary 11 provides for the existence of infinitely many ambient isotopic approximants, each with bounded

deviation from F. The new methods presented here may be valuable for applications in reverse engineering, where the desired output is often a CAGD B-rep model with spline boundary surfaces rather than just a PL approximant and the data to compute ρ may be available within the engineering specification [26].

5. Application to interval solids

Let F be our surface, as given in Theorem 10. It is intuitively obvious that if one generates another surface by continuous local perturbations of F, which neither create new self-intersections nor remove any existing self-intersection, the perturbed surface will be ambient isotopic to F. In this application, we will make use of this idea and approximate F by another surface which is generated by local perturbations of F. This construction yields a PL approximation and is motivated by the recent work of Sakkalis et al. [28], in which the concept of an *interval solid* is defined and some of its fundamental topological and geometric properties are proved.

Throughout this section, a *box* is a rectangular, closed parallelpiped in $\mathbf{R}^3$ with positive volume, whose edges are parallel to the co-ordinate axes. Such boxes are used to create 'interval solids', as defined and discussed in Ref. [28]. Some critical supporting material from Ref. [28] will be summarized here, to keep the paper reasonably self-contained.

Let F be our surface, and assume that F is connected. Then the Jordan Surface Separation Theorem asserts that the complement of F in $\mathbf{R}^3$ has precisely two connected components, F_I, F_O; we may assume that F_I is bounded and F_O is unbounded. Let also $\mathbf{B} = \{b_j, j \in J\}$ be a finite collection of boxes which meets the following conditions:

C3 $\{\mathrm{Int}(b_j), j \in J\}$ is a cover of F.

C4 Each member b of $\mathbf{B}$ intersects F generically; that is, $b \cap F$ is a nonempty closed disk that separates b into two (closed) balls, B_b^+ and B_b^-, with $B_b^+, (B_b^-)$ lying in $F_I \cup F(F_O \cup F)$, respectively.

C5 For any $b_i, b_j \in \mathbf{B}$, let $b_{ij} = b_i \cap b_j$. If $\mathrm{Int}(b_i) \cap \mathrm{Int}(b_j) \neq \emptyset$, then b_{ij} is also a box having property **C4**.

In Ref. [28], conditions **C3–C5** were assumed. Notice that condition **C4** indicates that every $b \in \mathbf{B}$ intersects F in a natural way (Fig. 5).

The following result summarizes several previously appearing results.

Theorem 12. [28, corollary 2.1, p. 165] *If F is connected and B satisfies* **C3–C5**, *then $F \cap U_{j\in J} b_j$ is a solid.*

To this end, we will show that whenever F is a connected surface satisfying the hypotheses of Theorem 10, then F is ambient isotopic to the boundary of an interval solid constructed from F, as described in Theorem 12. To do so, we introduce some well known results from the literature.

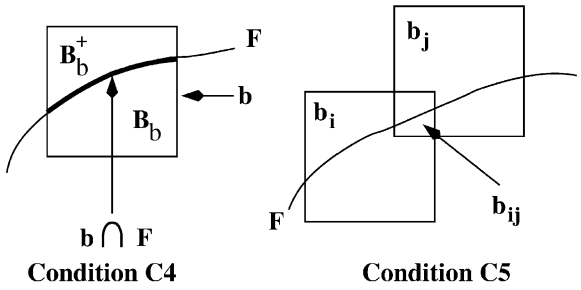

Fig. 5. 2D versions of conditions **C4** and **C5**.

Definition 13. [11, p. 214] A closed subset K of a PL three-manifold-with-boundary M is *tame* if there is a homeomorphism $h : M \to M$ such that $h(K)$ is a polyhedron.

Definition 14. [20, p. 371] Let M be a manifold with boundary, under the Euclidean metric d. Denote by $\mathbf{H}(M)$ the set of all homeomorphisms of M onto itself. Define a function α of $\mathbf{H} \times \mathbf{H}$ into the real extended number system as follows: $\alpha(f, g) = \sup_{x\in M} d(f(x), g(x))$. Then, if $\epsilon > 0, f$ and g are ϵ-isotopic if there is an isotopy H_t such that $H_0 = f, H_1 = g$ and if $t_1, t_2 \in [0, 1]$, then $\alpha(H_{t_1}, H_{t_2}) \leq \epsilon$.

The following Theorem 15 has previously appeared as a corollary [20] to a broader result which we do not need in this paper.

Theorem 15. [20, corollary 2, p. 372] *If K is a tame compact two-manifold in any three-manifold M and $\epsilon > 0$, there is a $\delta > 0$ so that if h is any homeomorphism of K into M moving no point more than δ and if $h(K)$ is tame, then there is an ϵ-isotopy of M taking $h(K)$ onto K pointwise and moving no point outside an ϵ-neighborhood of K.*

We are now in a position to present our main result of this section.

Theorem 18. *Let F be tame and connected. For each $\epsilon > 0$, there exists γ, with $0 < \gamma < \rho$ so that whenever a family of boxes* $\mathbf{B}$ *satisfies conditions* **C3–C5**, *and for each b of $\mathbf{B}$, b is a proper subset of $F(\gamma)$* (Fig. 6) *then, for $S = F \cup F_I$ and $S^{\mathbf{B}} = S \cup \bigcup_{j\in J} b_j$, the sets F and $\partial S^{\mathbf{B}}$* are *ϵ-isotopic with compact support. Hence, they are also ambient isotopic.*

Proof. It has previously been shown [28] that F and $\partial S^{\mathbf{B}}$ are homeomorphic by construction of an explicit homeomorphism. It has also been shown that $M = S^{\mathbf{B}}$ is a compact

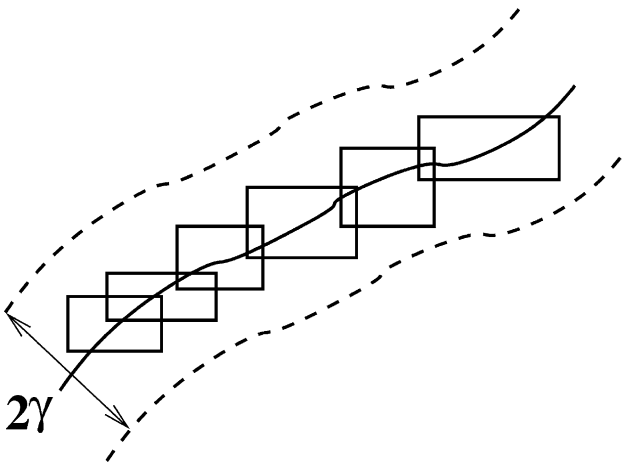

Fig. 6. 2D versions of proper subset condition.

three-manifold with boundary [28]. Furthermore, it is clear that M is PL. Consider now $K = \partial S^{\mathbf{B}}$. Since K is already a polyhedron, K is tame under the identity map on M and K is also a compact three-manifold within M. For the given ϵ, let δ be as given in Theorem 15. Furthermore, the homeomorphism $h : K \rightarrow F$ created [28] such that $h(K) = F$ relied upon a projection from the boundaries of the boxes onto F, where all of the boxes lie within $F(\rho)$. Hence, no point can be moved by h more than the maximal distance of the boundary of any box from F. Let now $\gamma = \min\{\rho, \delta\}$. Then h complies with the distance constraint of Theorem 15 and $h(K) = F$, where F is tame. Since M is compact, the ϵ-neighborhood provided by Theorem 15 implies that the ϵ-isotopy has compact support. Hence K and F are ambient isotopic. □

6. Global versus local methods

With an offset surface, there is a fixed value for the offset, resulting in a normal tubular neighborhood where the points generated as images of the endpoint map are all the same distance from the manifold. However, ambient isotopic approximations can be created where these distances need not all be equal [7], but are defined subject to local surface characteristics.

The proofs given in the present paper rely upon creation of a normal tubular neighborhood about F. While this is a sufficient condition it is *not* necessary that the original surface be entirely contained within an open neighborhood. In particular, there can be fixed points for an ambient isotopy [7].

There are two primary advantages to using offsets versus existing work [2,7] relying upon the medial axis [7] to construct a piecewise-linear ambient isotopic approximation of a manifold:

1. approximating the medial axis is a difficult task [5] whereas the method presented here requires no such computation and

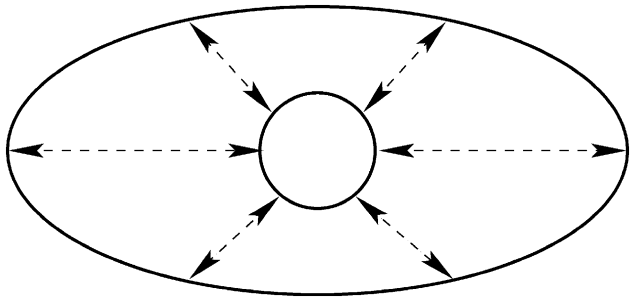

Fig. 7. Ambient isotopy by parameter matching.

2. the ambient isotopic approximants presented here need not be PL, and this may be valuable in engineering problems where the primary data representations are free-form surfaces.

Example. Let E be any nondegenerate, noncircular ellipse, which, without loss of generality, is assumed to be symmetric about the origin in the usual $x - y$ plane. There exists a minimal value of $\rho > 0$ such that the internal offset E by ρ is self-intersecting, and designate this offset by $E_o(-\rho)$. Now, consider any nondegenerate circle C centered about the origin so that C is inside $E_o(-\rho)$. There exists an ambient isotopy from E to C, which can easily be constructed by parametrizing C and E by $s \in [0, 2\pi]$ so that the mapping of points with the same parametric values is $1 - 1$, as depicted in Fig. 7. The obvious generalizations can be made in three dimensions, leading to the conclusion that if one only wishes to generate an ambient isotopic approximation, that the bounds previously given in this paper as well as in previous work [2,7,21] can be overly restrictive.

7. Comparisons of normals

This paper has, so far, emphasized two criteria for surface approximations

1. an upper bound on the distance between the approximating surface and the original surface, and
2. sufficient conditions for the approximant to be topologically equivalent to the original, via the strong criterion of ambient isotopy.

Other aspects, however, of a PL approximation are also of practical interest. In particular, one may wish to determine how far the normals of a PL approximant differ from the normals of the original, especially for applications in graphics. Note, that a PL surface reconstruction is guaranteed [7] to be ambient isotopic to the original surface under the condition of having a sampling density at each point $p \in F$ that is bounded by

$$r \times d(p, MA(F)), \text{ with } r = 0.08, \tag{5}$$

where $MA(F)$ denotes the medial axis of F.

Previous surface reconstruction [2] work has investigated the approximation of true surface normals by the normals of the reconstructed surface, but with an upper bound of 1/7 on the value of r defined in Eq. (5). Since the ambient approximations presented here are not constrained by the values of r presented in Eq. (5), it remains of interest to examine error bounds on these approximating normals, as will be demonstrated by adapting published comparison techniques [25] to an illustrative example. We further note that the approach presented here for comparison of normals may also extend to practical surface reconstruction problems, because, in practice, it has been found that a factor of r in Eq. (5) as high as 0.4 often produces acceptable algorithmic output [1].

We provide an example to illustrate the computation of error bounds on approximating normals. Both this paper and our guiding Ref. [25] share the following aspects in the construction of a PL approximation M of the original surface F. For each point $m \in M$, there exists a unique nearest point $\pi(m) \in F$. Then one considers the normal to M at m, denoted here as $\mathbf{n}_m^M$ and the normal to F at $\pi(m)$, denoted here as $\mathbf{n}_{\pi(m)}^F$. Then, after identifying the points m and $\pi(m)$, the angle between $\mathbf{n}_m^M$ and $\mathbf{n}_\pi(m)^F$ can be measured for each m. Then the upper bound is the supremum of these angular differences, taken over all points m in M.

First we created a NURBS curve C that is a trefoil knot. Then, we created a surface K by sweeping a circle of radius 0.01 with its center moving along the curve C (Fig. 8). This surface K was created so that the distance at each knot crossing was greater than 0.01, so that a nonsingular open normal tubular neighborhood about K could be created at any distance less than 0.01. Note that this distance exceeds the distance constraint imposed in Eq. (5).

The final two figures of this paper will illustrate how we employed the paper [25] to bound the normal deviations. The analysis [25] is critically dependant upon the length of the longest side of any triangle created in the triangulation.

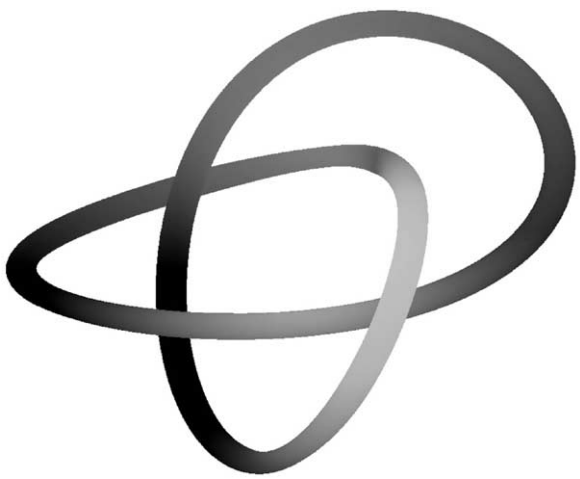

Fig. 8. Knotted surface.

To have control of the triangles generated on the surface, we triangulated the parametric domain. The triangles created in the parametric domain were all right triangles. For the particular example created, this led to corresponding triangles approximating K so that each such triangle was observed to be nearly a right triangle. This assumption that the triangles are approximately right triangles simplifies our computations. In particular, the upper bound [25, theorem 1] formula depends, for each triangle T in the triangulation, upon the variable str(T), which is known as the *straightness* of T, and while we do not repeat the details of that definition, here, we note that our assumption about right triangles permits us to approximate str(T) by the constant value of 1, for all triangles T.

An expression for an upper bound on the error of the approximating normals [25, theorem 1] is presented in our notation as

$$\sin \alpha_{\max} \leq ((\sqrt{10}/(2 \times \mathrm{str}(M))) + 1)(\pi_F(M))/(1 - \omega_F(M)), \tag{6}$$

We note that with the assumptions made about right triangles, the factor

$$((\sqrt{10}/(2 \times \mathrm{str}(M))) + 1)$$

immediately simplifies to the constant $((\sqrt{10}/2) + 1)$. Then, similar to our remarks about straightness, the reader is referred to Ref. [25] for the detailed definition of $\pi_F(M)$, known as the *relative height* of M to F. For our estimate of $\pi_F(M)$, we use

$$\pi_F(M) \approx L \times P_\xi, \tag{7}$$

where P_ξ is the maximum curvature of the original surface and L is the length of the longest side of any triangle created in the triangulation.

The factor $\omega_F(M)$) is given by

$$\omega_F(M) = \sup_{m \in M} \|\xi(m) - m\| \times P_\xi.$$

While there are obvious performance and data volume advantages to having fewer sampling points, there are intuitively obvious tradeoffs among several criteria for the approximating surface. For instance, if one considers the following three approximation criteria

1. nearness to surface being approximated,
2. preservation of topological form, and
3. error bounds on approximating normals,

then, fewer sampling points typically provided worse approximations relative to items 1 and 3, above, even in the presence of ambient isotopy. The analyses and numerical results presented here explicate those inter-relationships in order to provide some guidance to the practitioner towards making optimally practical choices regarding the sampling set. This requires a judicious balancing amongst the criteria articulated, above, as well as

Table 1
Upper bound on error of approximating normal

r	L	$\alpha_{\max}$
$4.0 \times 10 - 1$	$4.0 \times 10 - 3$	90.0+
$2.05 \times 10 - 1$	$2.05 \times 10 - 3$	44.2
$1.0 \times 10 - 1$	$1.0 \times 10 - 3$	19.6
$8.0 \times 10 - 2$	$8.0 \times 10 - 4$	15.5
$1.0 \times 10 - 2$	$1.0 \times 10 - 4$	1.9
$5.0 \times 10 - 5$	$5.0 \times 10 - 5$	0.95

possibly some others (surface area, volume enclosed, center of mass, etc.…). The approach presented here should offer guidance for other analyses to be extensible to include additional criteria for optimality.

In Table 1 the notation r represents the various values of the multiplicative factor used in Eq. (5). The associated columns indicate the corresponding values of L (as in Eq. (7), being length[3] of the longest side of any triangle created in the triangulation) and representing the model-space sample distance, and $\alpha_{\max}$ the upper bound of deviation between the normals of vertices on the mesh and points on the surface (in degrees).

Figs. 9 and 10 show how the true normals to the surface deviate from the normals to the triangles approximating the surface. Fig. 9 corresponds to the top line of Table 1 and Fig. 10 corresponds to the next line of Table 1. In both Figs. 9 and 10 a true normal to the surface at a given point is indicated by a dashed line, whereas a solid line indicates a normal of an approximating triangle. The deviations shown appear to decrease from Figs. 9 and 10, as would be expected from the decrease in the upper bounds give in Table 1. After the second line of Table 1 the corresponding images were such that the true and approximating normals were so close that it was hard to discriminate between them.

To further illustrate the comparison between normals, we also include two figures to show a close-up of selected normals. Each of Figs. 11 and 12 shows three geometric objects

1. a dashed line for a true surface normal,
2. a solid line for a normal to an approximating triangle, and
3. a shaded image of a triangle, where this shaded triangle is to the right of the figure.

Both lines for the normals are placed at a vertex of the triangle shown, which is also a point on the surface. Fig. 11 has the coarser sampling, corresponding to the third line of Table 1. Fig. 12 has sampling which is an order of magnitude finer than that of Fig. 11 and this finer sampling corresponds to the fifth line of Table 1. As would be expected from the comparison of the upper bounds in Table 1, the angle between the normals in Fig. 11 is much larger than the corresponding angle in Fig. 12. Also, note that,

[3] These lengths were created by using Taylor's Theorem to compute a related sampling in parameter space to map to these points in model space.

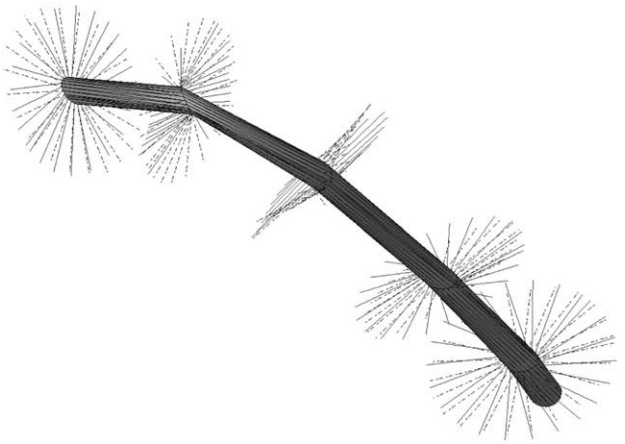

Fig. 9. Normal deviations, case 1.

consistent with the finer sampling rate, the triangle in Fig. 12 is smaller than that in Fig. 11.

8. Concluding remarks

In this paper we focused on a method for establishing surfaces which can approximate a given nonsingular compact manifold F without boundary so that each approximant is ambient isotopic to F. The approximants need not be PL and these nonPL approximations may be particularly useful in engineering applications where spline geometry dominates. The methods described are also directly applicable to creating approximants for graphics and animation, where the underlying surfaces already have a particular mathematical description stored in some standard format. They offer further theoretical insight into surface reconstruction problems, but at the expense of reliance upon data that is not typically present when only point cloud data is given. However, specifically for surface reconstructions in reverse engineering, the required local data of maximum curvature and the global data for creation of nonsingular tubular offset neighborhoods may be available.

The results presented here start from considerations of curvature. While local values of curvature and the medial

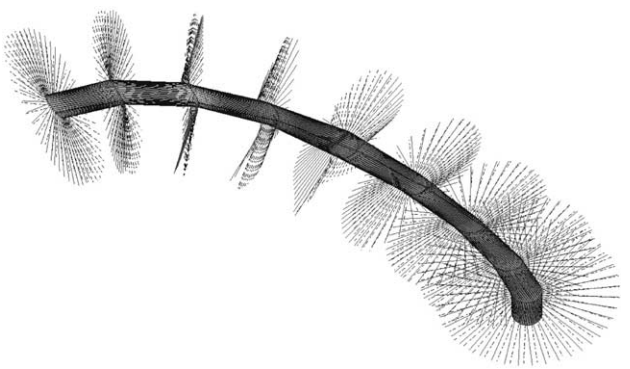

Fig. 10. Normal deviations, case 2.

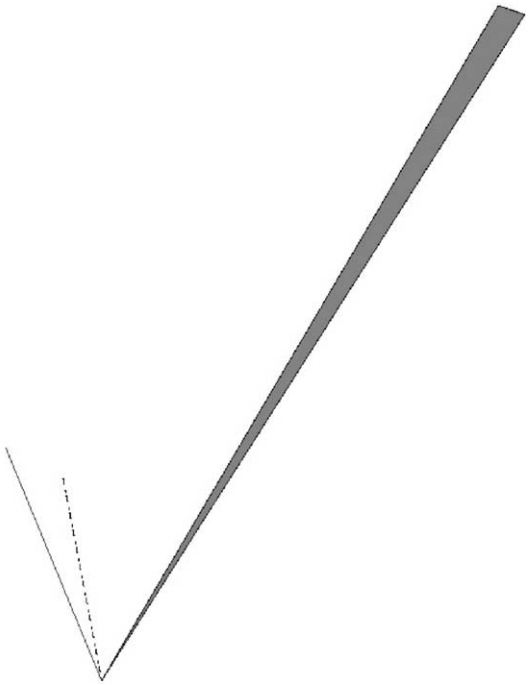

Fig. 11. Coarse sampling (0.1).

axis are closely related, examples are given to show that this primary attention to curvature affords advantages over previous methods which relied upon the medial axis, even when the objective is to create PL approximants. An application presented is the demonstration of sufficient conditions for an interval solid to be ambient isotopic to the solid it is approximating.

The results presented are restricted to smooth surfaces, even while real engineering parts often are only piecewise smooth, possibly having sharp features where surface derivatives are undefined. Consideration of such sharp discontinuities in developing ambient isotopies over multi-patch spline surfaces has appeared in the literature [10], but it is beyond the scope of the present investigation to integrate these two ideas. The theory presented here is offered as an important *a priori* step to the more challenging future extensions to real engineering parts, which the present authors are investigating in an ongoing project.

Fig. 12. Finer sampling (0.01).

Acknowledgements

Partial funding for T. Sakkalis was obtained from NSF grants DMS-0138098, CCR 0231511, CCR 0226504 and from the Kawasaki chair endowment at MIT. Partial funding for T.J. Peters and J. Bisceglio was from NSF grants DMS-9985802, DMS-0138098 and CCR 0226504. All statements in this publication are the responsibility of the authors, not of these funding sources. The authors express their appreciation for this funding and would further like to thank Professors K. Abe, N. Amenta, N. M. Patrikalakis, J. A. Roulier and A. C. Russell for various discussions which helped to stimulate our collaboration on this manuscript.

Appendix A

Topology definitions

Definition 17. A function f on spaces X and Y is a *homeomorphism* if f is bi-continuous, 1–1 and onto.

Definition 18. A function f from X onto itself has compact support if there exists a compact set $A \subset X$ such that f is the identity except possibly on A.

Definition 13. Two functions f and g from a space X into a space Y are called *homotopic* if there exists a continuous function $F : X \times [0,1] \rightarrow Y$ such that for each point $x \in X$,

$$F(x,0) = f(x) \text{ and } F(x,1) = g(x).$$

Definition 18. Two functions f and g from a space X into a space Y are called *isotopic* if they are homotopic via a function F such that for each $t \in [0,1], F(\cdot,t)$ is a homeomorphism.

If the original functions f and g are both onto Y, then we will interchangeably refer to the functions being isotopic, as well as the spaces X and Y being isotopic. It is this latter usage that is adopted within the main body of this paper in the definition of *ambient isotopy* given in Definition 9.

References

[1] Amenta N. Personal communication. May 15; 2003.
[2] Amenta N, Bern M. Surface reconstruction by voronoi filtering. Discrete Computational Geometry 1999;22:481–504.

[3] Amenta N, Bern M, Kamvysselis M. A new voronoi-based surface reconstruction algorithm. Proc ACM SIGGRAPH, ACM; 1998. p. 415–21.
[4] Amenta N, Choi S, Dey T, Leekha N. A simple algorithm for homeomorphic surface reconstruction. ACM Symposium on Computational Geometry. 2000. p. 213–22.
[5] Amenta N, Choi S, Kolluri R. The power crust, union of balls and the medial axis transform. Computational Geometry: Theory Appl 2001; 19:127–73.
[6] N. Amenta et al. Emerging challenges in computational topology. In Workshop Report on Computational Topology. NSF; June 1999. http://xxx.lanl.gov/abs/cs/9909001
[7] Amenta N, Peters TJ, Russell AC. Computational topology: ambient isotopic approximation of 2-manifolds. Theor Comput Sci 2003;305: 3–15.
[8] Andersson L-E, Dorney SM, Peters TJ, Stewart NF. Polyhedral perturbations that preserve topological form. Comput Aided Geometric Des 1995;12:785–99.
[9] Andersson L-E, Peters TJ, Stewart NF. Selfintersection of composite curves and surfaces. Computer Aided Geometric Des 1998;15(5): 507–27.
[10] Andersson L-E, Peters TJ, Stewart NF. Equivalence of topological form for curvilinear geometric objects. Int J Computational Geometry Appl 2000;10(6):609–22.
[11] Bing RH. The Geometric topology of 3-manifolds. Providence, RI: American Mathematical Society; 1983.
[12] Boyer M, Stewart NF. Modeling spaces for toleranced objects. Int J Robotics Res 1991;10(5):570–82.
[13] Boyer M, Stewart NF. Imperfect form tolerancing on manifold objects: a metric approach. Int J Robotics Res 1992;11(5):482–90.
[14] Cohen J, et al. Simplification envelopes. Proc ACM SIGGRAPH 96, ACM; 1996. p. 119–28.
[15] Dey TK, Edelsbrunner H, Guha S. Computational topology. Advances in discrete and computational geometry (Contemporary Mathematics 223), American Mathematical Society; 1999. p. 109–43.
[16] Edelsbrunner H, Mücke EP. Three-dimensional alpha shapes. ACM Trans Graphics 1994;13(1):43–72.
[17] Guillemin V, Pollack A. Differential topology. Englewood Cliffs, NJ: Prentice-Hall; 1974.
[18] Hirsch MW. Differential topology. New York: Springer; 1976.
[19] Hu C-Y. Towards robust interval solid modeling for curved objects. PhD thesis, Massachusetts Institute of Technology, Cambridge, MA; 1995.
[20] Kister J. Small isotopies in euclidean spaces and 3-manifolds. Bull Am Mathematical Soc 1959;65:371–3.
[21] Maekawa T, Patrikalakis NM, Sakkalis T, Yu G. Analysis and applications of pipe surfaces. Computer Aided Geometric Des 1998; 15(5):437–58.
[22] Milnor J. Morse theory. Princeton, NJ: Princeton University Press; 1969.
[23] Moise E. Geometric topology in dimensions 2 and 3. New York: Springer; 1977.
[24] G. Monge, Application de l'Analys à la Geometrie, Bachelier, Paris; 1850.
[25] Morvan J-M, Thibert B. Smooth surface and triangular mesh: comparison of the area, the normals and the unfolding. Proceedings of the Symposium on Solid Modeling, ACM; 2002. p. 147–57.
[26] Macy W. Personal communication. July 14; 2003.
[27] Sakkalis T, Charitos C. Approximating curves via alpha shapes. Graphical Models Image Processing 1999;61:165–76.
[28] Sakkalis T, Shen G, Patrikalakis N. Topological and geometric properties of interval solid models. Graphical Models 2001;63:163–75.
[29] Stewart NF. Sufficient condition for correct topological form in tolerance specification. Computer Aided Des 1993;25(1):39–48.
[30] Wallner J, Sakkalis T, Maekawa T, Pottmann H, Yu G. Self-intersections of offset curves and surfaces. Int J Shape Modeling 2001; 7(1):1–21.
[31] Willard S. General topology. Reading, MA: Addison-Wesley; 1970.

Dr Sakkalis is an Associate Professor of Mathematics at the Agricultural University of Athens, Greece and a Research Affiliate at MIT. He received a Diploma in Mathematics in 1979 from the University of Athens, Greece, and a PhD in Mathematics in 1986 from the University of Rochester, NY. His research interests include computer aided geometric design, computer algebra, geometric modeling and computational geometry. Dr Sakkalis has been a Visiting Research Scientist at the IBM Thomas J. Watson Research Center and at the Design Laboratory at MIT. From 1986 to 1990 he was an Assistant Professor of Mathematics at New Mexico State University, NM, while from 1990 to 1994 was an Assistant Professor of Mathematics at Oakland University, MI.

T.J. Peters took his PhD in mathematics (topology) from Wesleyan University. He has held senior technical positions in the CAD/-CAM industry and at Charles Stark Draper Laboratory. He currently enjoys joint appointments in the Department of Mathematics and the Department of Computer Science and Engineering at the University of Connecticut, where he is an Associate Professor. His primary research interests are computational topology, computer graphics and computational geometry—viewed as an integrative whole in support of CAGD. He maintains active collaborative ties with industry and was recently a program co-chair for the inaugural SIAM conference, 'Mathematics in Industry'.

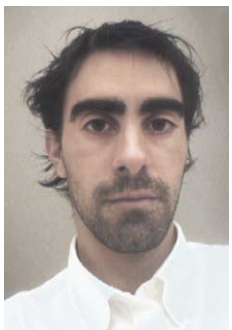

Justin Bisceglio recently completed a Master's degree in Computer Science from the University of Connecticut. He was a research assistant under an NSF funded CARGO collaboration on topological and numerical issues for surface intersections in CAGD. He led the technology transfer of theoretical developments to practical intersection software used at The Boeing Corporation. Prior to graduate school, Justin received his BS in film and photography from Ithaca College and worked in the commercial television industry. He is currently employed in the research and development department at Blue Sky Studios in White Plains, NY.