Generalized Filleting and Blending Operations toward Functional and Decorative Applications

Gershon Elber
Computer Science Department
Technion
Haifa 32000, Israel
Email: gershon@cs.technion.ac.il

Abstract

The use of blending and filleting operations in solid modeling and computer-aided geometric design is well established. The question of filling a gap between two (or more) surface boundaries or rounding a sharp edge has been extensively investigated. The vast majority of the prior work on blending and filleting concentrated on a wide variety of fitting schemes as well as attempts to establish and guarantee better continuity conditions.

This work extends the notion of filleting and blending modeling tools and elevates them into shaping operations that are either functional or ornamental in nature. The extended shaping operations can be conducted between two boundaries of two adjacent surfaces, much like traditional blending or filleting methods. Furthermore, the presented extended forms can also be applied to the interior of a single surface, guided by arbitrary parametric curves in the domain of the patch.

Additional Key Words and Phrases: *Curves & Surfaces, Decoration, Ornaments, Geometric Modeling Operations, Rounding, Cubic Hermite.*

1 Introduction

Blending and filleting operations are important tools in any contemporary geometric modeling environment. The ability to smooth corners or to continuously connect two adjacent surfaces into a single object is considered mandatory in any computer-based modeling environment. Blending and filleting surfaces are typically required to be G^1. That is, the blending surface must share a tangent plane with the surfaces whose gap it fills, along their shared boundary.

Different authors denote these type of surfaces differently. Blending surfaces, filleting surfaces, or even joining and rounding surfaces are the terms typically used. For the sake of consistency, we only employ the term blending surface in this work.

Classically, a blending surface is constructed along a shared seam or an intersection between two other already specified surfaces, called the *primary surfaces*. Typically, two of the boundary curves of the blending surface lie on the primary surfaces. These two boundary curves will be referred to as *rail curves*, following [8].

Blending surfaces may be subdivided into two types. The first is the exact type, in which the cross section must follow a prescribed form, typically circular; the second is the imprecise type, in which the cross section is only loosely specified. The former is known as constant radius blending and the latter as a "thumb fill", having a thumb running along the filleted area creating a non circular yet G^1, or tangent plane, continuously smooth filling.

Given a seam along which a blending surface needs to be constructed, the standard approach for deriving the rail curves is based on offsetting both primary surfaces [5, 15, 16]. This approach also guarantees the exact radius distance needed for constant radius blending. Another alternative to constructing the rail curve is by forming an intermediate circular sweep surface around the seam and intersecting it against the two primary surfaces [4].

In [2, 8, 14], blending approaches that do not use a constant radius are presented. These approaches are appealing for their simplicity and the elimination of the requirement of an offset surface computation. Blending algorithms for algebraic surfaces, including [1, 18], have been developed

separately. See [17] for a nice survey of the different existing blending techniques.

Another type of related work, such as [3], includes feature pasting over spline surfaces. These works add details to the spline surface in the form of a displacement map. This approach is an extended view of the hierarchical B-spline refinement scheme presented in [11].

In our work, the blending operations are extended into three-dimensional decorative features on the surface, for ornamental purposes. Hence, we are also interested in related previous work on three-dimensional texture and displacement mapping [10]. In computer graphics, texture mapping is considered an important and successful tool toward photorealism in image synthesis and rendering. A variant of texture mapping that creates the illusive perception of three-dimensional surface details is known as bump mapping. This variant is a special type of texture mapping where the normals that are fed into the shader are perturbed, instead of the colors, creating the effect of a bumpy shape.

A major drawback of the bump mapping technique is inherent in the fact that this bumpiness is illusive and the surface remains smooth. Along the silhouette curves the shape is not affected by the assigned bumps, as it should be. Some attempts at creating real three-dimensional texture over a surface were made, employing, for example, cellular texture generators [9]. A three-dimensional texture was constructed in [9] using a simulation of natural cellular development toward the modeling of surface details such as scales or thorns.

The graphics community frequently uses displacement mapping to alleviate some of the difficulties in the use of bump mapping. Here, a height map above the surface is defined and applied to the surface. The height map is typically an explicit function that is described as an image, with the pixels denoting height. The end result is a surface with a modified geometry.

This work also offers a variant of the methodology of cellular and/or particle systems based on three-dimensional textures and combines it with blending operations. The three-dimensional texture of [9] is formed out of zero-dimensional entities, or points. Here, we also examine a similar application of three-dimensional texture but use one-dimensional entities, or curves. These univariate texture detail curves could be represented using techniques similar to ones that are employed in blending operations.

Attempts were also made to glue geometry as ornaments on prescribed surfaces. One recent example is [12], where a primary-orientation pair of curves is used to guide the warping function along the base surface. Nevertheless, no continuity conditions are guaranteed in [12].

The rest of this paper is organized as follows. In Section 2, we review background material on the cubic Hermite interpolation scheme and in Section 3 we introduce our variant of the blending constructing methodology, a methodology that offers the freedom to select arbitrary cross sections as blending curves. In Section 4, we show one way of deriving the positional as well as the tangential constraints for curves on surfaces. Section 5 presents several examples that were created using the proposed scheme and, finally, we conclude in Section 6.

2 Background

The cubic Hermite interpolation scheme offers C^1 continuity [6]. Given an ordered set of points and tangent vectors at these locations, each consecutive pair of points is fitted using a single cubic Hermite curve that satisfies four constraints: the two locations or the end points, P_i and P_{i+1}, and the two end points' tangent vectors, T_i and T_{i+1}:

$$C(t) = P_i h_{00}(t) + P_{i+1} h_{01}(t) + T_i h_{10}(t) + T_{i+1} h_{11}(t), \tag{1}$$

where $h_{ji}(t)$ are the four cubic Hermite basis functions, with i being the point index and j the derivative's order:

$$\begin{aligned} h_{00}(t) &= t^2(2t-3)+1, \\ h_{01}(t) &= -t^2(2t-3), \\ h_{10}(t) &= t(t-1)^2, \\ h_{11}(t) &= t^2(t-1). \end{aligned}$$

These four cubic Hermite functions not only form a basis for the cubic polynomials but also satisfy the following orthogonality constraints:

Curve	$C(0)$	$C(1)$	$C'(0)$	$C'(1)$
$h_{00}(t)$	1	0	0	0
$h_{01}(t)$	0	1	0	0
$h_{10}(t)$	0	0	1	0
$h_{11}(t)$	0	0	0	1

By construction, the functions $h_{1i}(t)$ are zero at P_i and P_{i+1} and similarly, $h'_{0i}(t)$ are zero at P_i and P_{i+1}. This careful construction of the $h_{ji}(t)$ functions guarantees complete independence between the positional and tangential constraints.

These cubic Hermite functions can also serve to interpolate two curves instead of two points. Given two rail curves

$C_1(u)$ and $C_2(u)$ and two tangent fields along these two curves $T_1(u)$ and $T_2(u)$, the surface

$$\begin{aligned} S(u,v) &= C_1(u)h_{00}(v) + C_2(u)h_{01}(v) \\ &\quad + T_1(u)h_{10}(v) + T_2(u)h_{11}(v), \end{aligned} \tag{2}$$

interpolates both $C_1(u)$ and $C_2(u)$ at two of its boundary curves. Moreover, it is easy to verify that the cross derivatives, $\frac{\partial S}{\partial v}$, along these two boundaries are equal to $T_1(u)$ and $T_2(u)$.

The constructed surface, $S(u,v)$, has no free degrees of freedom remaining. Being a cubic and having four degrees of freedom, they are all exploited by the four (two positional and two tangential) constraints. In Section 3, we examine ways to relax this limitation and allow us to employ an arbitrary cross section curve in the interpolation, in a natural and intuitive way.

3 The Basic Algorithm

We now examine the question of how one can employ an arbitrary open planar curve as a cross section, $C_s(v) = (x_s(v), y_s(v))$, to construct a shaped blending surface, S, that interpolates between the two space curves $C_1(u)$ and $C_2(u)$, and follows the two tangent fields $T_1(u)$ and $T_2(u)$ at these boundaries.

Our solution to this shaped blending surface S is constructive. We seek to preserve the complete separation between the positional constraints and the tangential constraints. Consider the surface $S_1(u,v)$:

$$S_1(u,v) = T_1(u)h_{10}(v) + T_2(u)h_{11}(v). \tag{3}$$

Clearly, $S_1(u,v)$ satisfies the given tangential constraints along the two boundaries. Yet, $S_1(u,v)$ identically vanishes for $v = 0$ or for $v = 1$.

Now assume we are given a surface $S_2(u,v)$ that satisfies the prescribed positional constraint. That is, $S_2(u,0) = C_1(u)$ and $S_2(u,1) = C_2(u)$. Moreover, assume that $\frac{\partial S_2(u,v)}{\partial v}$ identically vanishes for $v = 0$ or for $v = 1$. Then,

$$S(u,v) = S_1(u,v) + S_2(u,v), \tag{4}$$

satisfies all four constraints.

We are now ready to introduce our construction scheme for $S_2(u,v)$. Let

$$D(u) = \frac{C_2(u) - C_1(u)}{2}, \quad A(u) = \frac{C_2(u) + C_1(u)}{2}, \tag{5}$$

and let $V(u)$ be an arbitrary 3-space vector field. Using linear transformation of rigid motion and scaling, map the prescribed cross section $C_s(v) = (x_s(v), y_s(v))$, $t \in [0,1]$ so that $C_s(0) = (-1,0)$ and $C_s(1) = (1,0)$ (see Figure 1 for a few examples).

Now consider the surface

$$\begin{aligned} S_2(u,v) &= A(u) + D(u)x_s(v) + V(u)y_s(v) \\ &= \frac{C_2(u) + C_1(u)}{2} + \frac{C_2(u) - C_1(u)}{2}x_s(v) \\ &\quad + V(u)y_s(v), \end{aligned} \tag{6}$$

for $v = 0$, $x_s(0) = -1$ and $y_s(0) = 0$ by construction, or

$$S_2(u,0) = \frac{C_2(u) + C_1(u)}{2} - \frac{C_2(u) - C_1(u)}{2} = C_1(u).$$

Similarly, for $v = 1$, $x_s(0) = 1$ and $y_s(0) = 0$ we have

$$S_2(u,1) = \frac{C_2(u) + C_1(u)}{2} + \frac{C_2(u) - C_1(u)}{2} = C_2(u),$$

or $S_2(u,v)$ is interpolating the given two curves, $C_i(u)$, $i = 1,2$. Consider now the cross derivative of $S_2(u,v)$:

$$\frac{\partial S_2(u,v)}{\partial v} = D(u)x_s'(v) + V(u)y_s'(v). \tag{7}$$

At this point we impose one additional constraint on $C_s(t)$:

$$C_s'(0) = C_s'(1) = (0,0).$$

Then, inspecting Equation (7), $\frac{\partial S_2(u,v)}{\partial v}$ vanishes for both $v = 0$ and $v = 1$. In other words, $S(u,v) = S_1(u,v) + S_2(u,v)$ satisfies all four constraints originally imposed for $S(u,v)$ in Equation (4).

Consider a Bézier or NURBS curve, $C_s(v)$, $t \in [0,1]$. $C_s'(0) = (0,0)$ if $P_1 = P_0$ or the first two control points are identical. Similarly, if the last two control points of $C_s(v)$ are identical, $C_s'(1) = (0,0)$. In practice and in all the examples presented as part of this work, including the examples shown in Figure 1, the first and last two control points are set to be identical, coercing a vanishing speed at the two curves' end points.

We have complete freedom in selecting $V(u)$. One reasonable selection for this vector field could be the direction orthogonal to both $D(u)$ and $A'(u)$ (see Equation (5)). If the blending surface is (almost) planar, $D(u)$ and $A'(u)$ (approximately) span its tangent space and $V(u)$ will be in the direction of the surface's normal field. Hence, a default plausible selection for $V(u)$ could be

$$V(u) = D(u) \times A'(u).$$

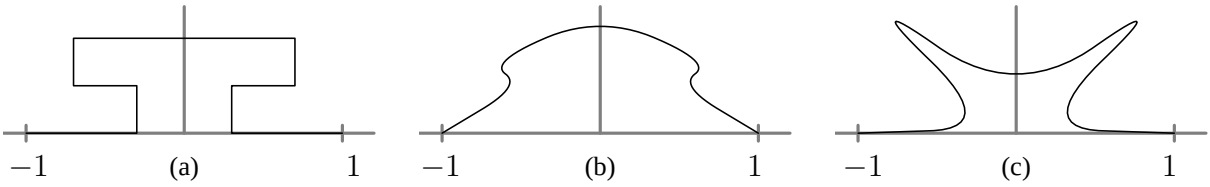

Figure 1. Examples of cross sections used in shaped surface blending. These curves are constructed so that their end points span -1 to +1 and furthermore, their speed at these end points is zero, by duplicating the end control points.

Figure 2 presents several examples of the proposed scheme with $V(u)$ selected in a direction that is approximately orthogonal to both $D(u)$ and $A'(u)$. Figure 2 (a) shows the result of using the traditional cubic Hermite form, while Figures 2 (b)-(d) present three examples of employing three different cross sectional curves, $C_s(v)$. The three cross sections exploited in Figures 2 (b)-(d) are shown in Figures 1 (a)-(c), respectively. The cross section of Figure 1 (a) is piecewise linear and hence is C^0 continuous. As a result, the reconstructed blending surface shown in Figure 2 (b) is also C^0 continuous (but is still G^1 along the shared boundaries!). The other two cross sections are C^2 continuous and the resulting blending surfaces are G^1 along the boundaries.

This shaped blending construction scheme has a few degrees of freedom that could be further employed, intuitively. Clearly, all the cross sections, $C_s(v)$, (see also Figure 1) could be scaled in height, effectively employing a different $V(u)$ vector field, while preserving all the continuity constraints. Figure 3 shows the same example as in Figure 2 (c) but with different vertical scales of $C_s(v)$.

A different degree of freedom for shaping the constructed surface is found in the magnitude of the prescribed tangent vector fields, $T_i(u)$. A change in the magnitude of these fields affects neither the positional continuity nor the tangential plane continuity. Figure 4 shows the example in Figure 2 (c) with different magnitudes of the tangent fields.

4 Curves on Surfaces

So far we have considered the constructed surfaces toward a generalized functional blending application. Nonetheless, one could equally employ this scheme over the interior of a surface, placing three-dimensional univariate ornaments over it. In essence, the end result is similar in nature to the three-dimensional texture mapping presented by [9]. Here, instead of using particles and/or cells, we attempt to use arbitrary parametric curves.

Let $c(t) = (u(t), v(t))$ be an arbitrary curve in the parametric domain of surface $S(u, v)$. We seek to prescribe an ornamental feature, $O(r, t)$, along $c(t)$ on $S(u, v)$, following the guidelines of Section 3. Towards this end, one needs to derive from $c(t)$ and $S(u, v)$ two positional constraints and two tangential constraints. An additional cross section's specification will complete the prescription of the shape of $O(r, t)$.

We start by attempting to establish the width of the ornament laid on the surface, in Euclidean space. The width d of the ornament will be derived by computing two offsets, $c_d(t)$ and $c_{-d}(t)$, from $c(t)$, in the parametric domain of S. The magnitude of the first order partials of $S(u, v)$ varies across the surface. Put differently, $S(u, v)$ is far from an Isometry, in general. A curve parallel to $c(t)$ in the parametric space of S, $c_d(t)$, at some fixed distance, d, will rarely retain this constant distance in the Euclidean space. That is, if $c_d(t)$ is an exact offset of $c(t)$ by an amount d, we have $||c_d(t) - c(t)|| = d$. Yet, $||S(c(t)) - S(c_d(t))||$ is rarely a constant function.

That said, and since we seek this estimated width for mostly decorative purposes, the exact and fixed width in the Euclidean space is of no real importance. By estimating the magnitude of the partials at one location, t_0, along $c(t)$, we could estimate the required width in the parametric space, d, for achieving a certain width, Δ, in Euclidean space. Granted, this width will be exact for that location only. Given d, the two positional constraints of $O(r, t)$ are computed from $c(t)$ as offsets $c_d(t)$ and $c_{-d}(t)$, where d is estimated from the magnitudes of $\frac{\partial S}{\partial u}$ and $\frac{\partial S}{\partial v}$ at location $c(t_0)$. Let $\tau(t) = (\tau_x(t), \tau_y(t))$ be the unit tangent field of $c(t)$. Then, d is estimated by

$$d = \frac{\Delta}{\left|\left| \tau_y(t_0) \left.\frac{\partial S}{\partial u}\right|_{\substack{u = u(t_0) \\ v = v(t_0)}} - \tau_x(t_0) \left.\frac{\partial S}{\partial v}\right|_{\substack{u = u(t_0) \\ v = v(t_0)}} \right|\right|}.$$

Let $C_d(t) = S(c_d(t)) = S(u_d(t), v_d(t))$. If both $S(u, v)$

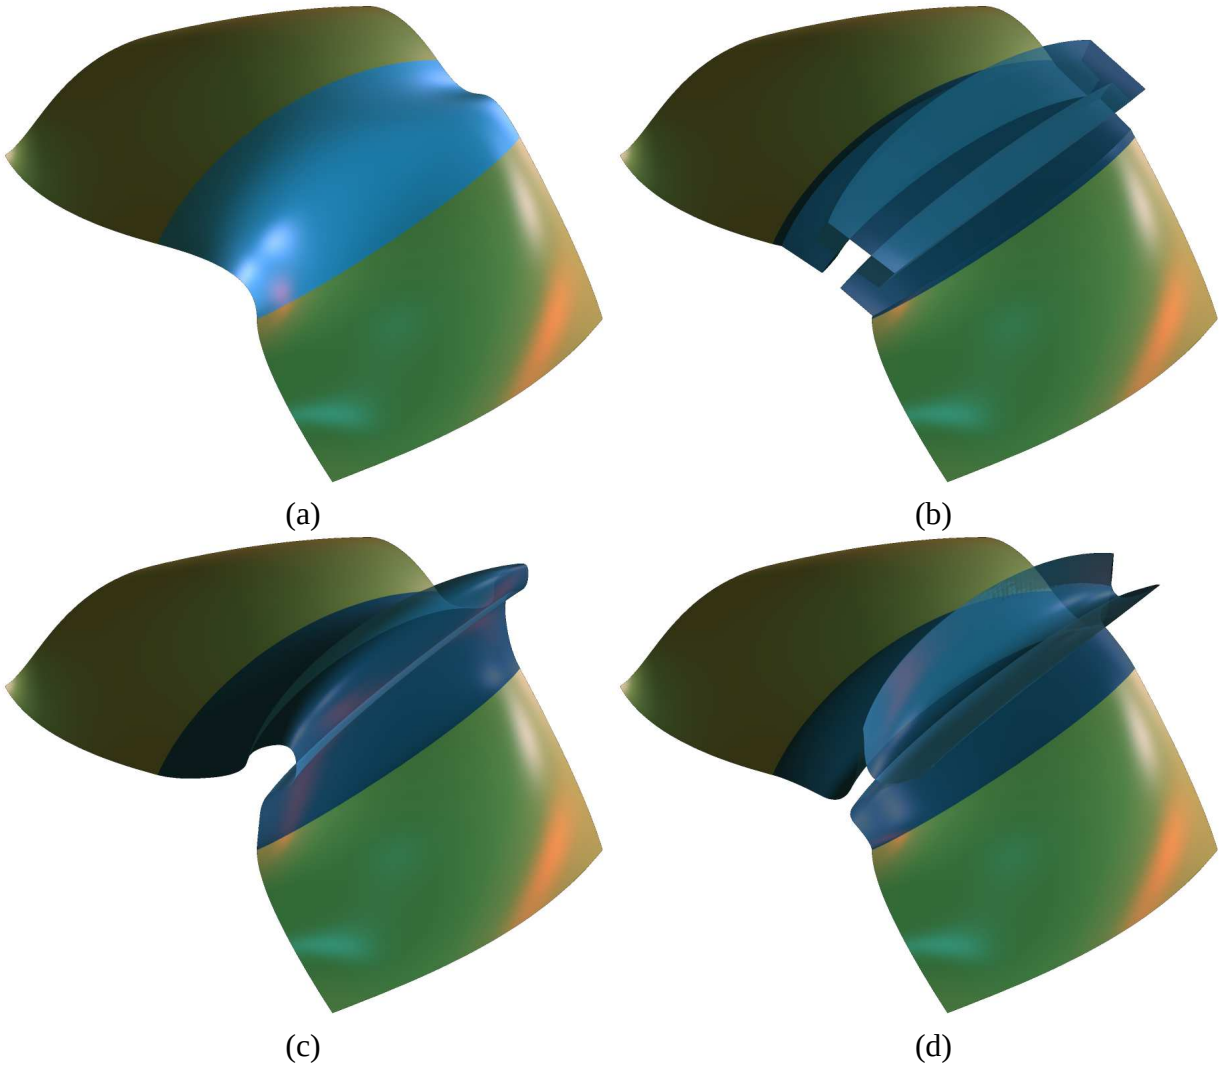

Figure 2. A functional blend. (a) presents the traditional G^1 continuous cubic Hermite surface blending while (b)-(d) show the blending using the different cross sections that are shown in Figure 1. Cross sections, similar to those used in (b)-(d), could serve some functional purpose in the object, such as a T slot along which to slide.

and $c_d(t)$ are rational, so is $C_d(t)$. $C_d(t)$ could be derived via a composition operation that could be evaluated symbolically [7].

Let $n(u, v)$ be the normal field of $S(u, v)$, $n(u, v) = \frac{\partial S}{\partial u} \times \frac{\partial S}{\partial v}$, which is another rational function. Then, and following [14], define

$$T_d(t) = C'_d(t) \times n(u, v), \qquad T_{-d}(t) = C'_{-d}(t) \times n(u, v), \tag{8}$$

as the two tangential field constraints along $C_d(t)$ and $C_{-d}(t)$. These two new tangential fields are clearly tangent to S as they are orthogonal to the normal field of S. Furthermore, they are also orthogonal to the tangential fields of the rail curves themselves, $C'_d(t)$ and $C'_{-d}(t)$. Finally, given a rational form of $n(u, v)$ and $C'_d(t)$, $T_d(t)$ is rational as well.

Given a curve in the parametric space of surface S, $c(t)$, we define the decorative ornament prescribed by $c(t)$ on S as

$$\begin{aligned} O(r, t) \quad = \quad & A_o(t) + D_o(t) * x_s(r) + n(u(t), v(t)) * y_s(r) \\ & + T_{-d}(t) h_{10}(r) + T_d(t) h_{11}(r), \end{aligned} \tag{9}$$

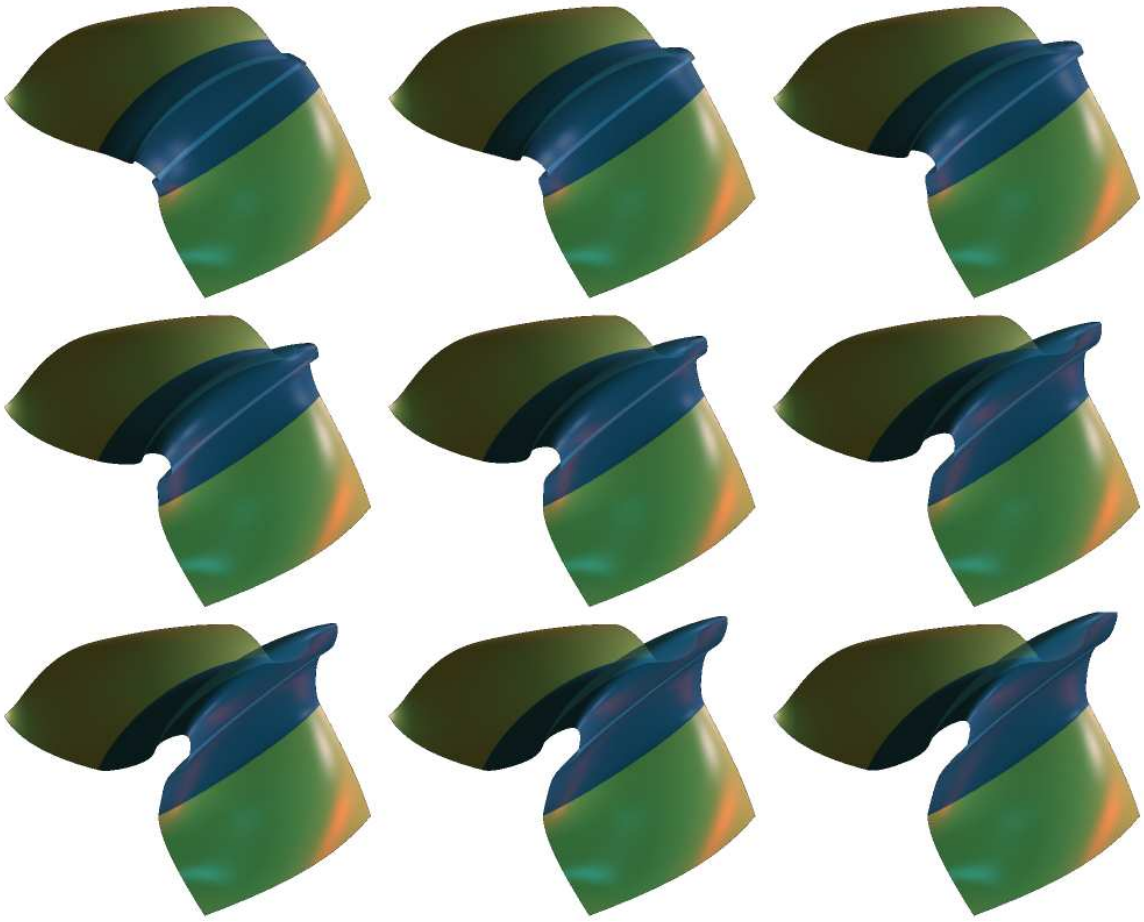

Figure 3. The height of the cross section function, set via scaling function $y_s(v)$, could be modified while all the continuity constraints are fully preserved. These examples employ the cross section seen in Figure 1 (b). Compare with Figure 2 (c).

where

$$D_o(t) = \frac{C_d(t) - C_{-d}(t)}{2}, \quad A_o(t) = \frac{C_d(t) + C_{-d}(t)}{2}.$$

Note that we employ the normal field in Equation (9) to orient the decorative ornament. This is clearly one plausible option to prescribe the free vector field $V(u)$ in Equation (6), out of many such feasible options, as the ornament will follow the surface orientation. In the Euclidean space, $C_d(t)$, $C(t)$, and $C_{-d}(t)$ are not collinear, in general. Hence, and while $A_o(t)$ is the average of the two opposite offsets, $A_o(t)$ does not equal $C(t)$.

So far, our algebraic manipulations ignored one potential difficulty. The surfaces prescribed by our above construction schemes will indeed interpolate the two boundaries and will certainly be tangent plane continuous there. Yet, these surfaces also suffer from a significant limitation by being dependent upon the original surface's parameterization. Re-examining Equations (8) and (9), the vector fields of $T_d(t)$ and $T_{-d}(t)$ depend upon the normal field, $n(u, v)$, which is not normalized, and hence is affected by a reparameterization of $S(u, v)$. Moreover, the unnormalized vector field of $n(u, v)$ also contributes directly to Equation (9), affecting the height of the constructed ornament above the surface.

A remedy to these unnormalized vector fields may be sought at two different levels, depending upon the required accuracy of that vector field. The tangent vector fields, $T_d(t)$ and $T_{-d}(t)$, must be kept precise for proper tangent plane continuity. In contrast, the vector field of $V(u)$ in Equation (6) ($n(u(t), v(t))$, in Equation (9)) need not be as precise.

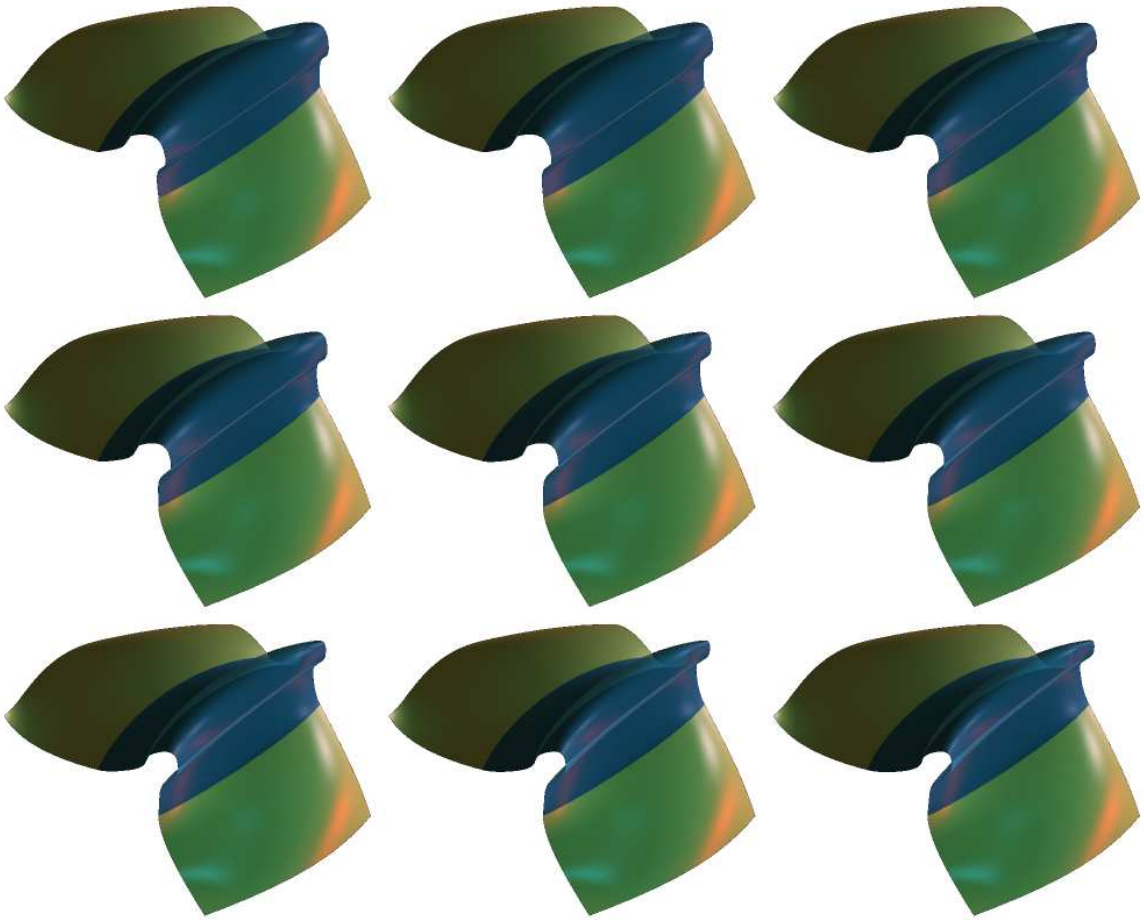

Figure 4. The magnitude of the tangent fields, $T_1(u)$ and $T_2(u)$, could be modified while the tangent plane continuity is fully preserved. These examples employ the cross section seen in Figure 1 (b). Compare with Figures 2 (c) and 3.

Assume all vector field curves are given as piecewise rational NURBS forms. The computation of an approximation to a unit vector field out of a given unnormalized vector field could be handled by simply normalizing all the control points of this B-spline vector field. This approach was employed in all the examples presented in this work for the imprecise vector fields that were used, specifically $V(u)$ and $n(u(t), v(t))$ in Equations (6) and (9).

In [14], a refinement-based [6] method was employed to approximate a unit vector field from a given vector field $W(t)$, and to an arbitrary precision. In [14], a polynomial approximation of the magnitude of the given vector field, w(t) = $\sqrt{\langle W(t), W(t)\rangle}$, is constructed. Then, the unit size vector field is approximated as the rational form of $\overline{W}(t) = \frac{W(t)}{w(t)}$. This second approach is not only guaranteed to converge under refinement but, and more importantly, is also safe in the sense that at no time is the direction of the vector field modified, only its magnitude.

Section 5 presents some examples that were constructed using the presented scheme of univariate ornamental texturing using parametric curves on freeform surfaces.

5 Examples and Extensions

This section presents several examples of blending and placing univariate decorative ornamental features along arbitrary parametric curves on the surface. All the examples presented in the section were created using an implementation based on the IRIT [13] solid modeling environment, developed at the Technion. In all examples, the end result

is a surface that is tangent plane continuous to the original surface. Before demonstrating the proposed blending capabilities, Figure 5, shows some additional cross sections that we are about to employ in this section.

Figure 6 shows several G^1 continuous blends between the body of the Utah teapot and its spout. Figure 6 (a) shows one possible result of using regular cubic Hermite to produce this blend. Figures 6 (b) and (c) use the simple cross section shown in 5 (d). The difference between 6 (b) and (c) is in the normal field used. In (b) the normal field is a radial vector field around the spout. In (c), the normal field has an additional component, in a direction away from the body. In Figures 6 (d) to (i), a more complex cross section, of the shape shown in Figure 5 (e), is used. Figures 6 (d) to (f) show the resulting geometry whereas Figures 6 (g) to (i) present side views of the cross sections of these blends, respectively. The differences in Figures 6 (d) to (f) stem for the use of three different magnitudes for the vector field of $V(u)$, with the smallest at (d) and the largest at (f).

In Figure 7, a simple, single bump, cross section in the shape of a Gaussian function (see Figure 5 (d)) is laid on the freeform surface of the body of the Utah teapot as general ornaments in the shape of the letters 'GMOD'. These ornaments are laid along general curves in the parametric space of the given surface S. The offsets of the curves are computed in the parametric domain of S, and then composed with S in order to derive the rail curves and the necessary generalized Hermite constraints, as is described in Section 4.

Figure 8 shows the Utah teapot with ornaments placed over all its surfaces. In Figures 8 (a) and (b), the teapot's body employs the cross section of Figure 5 (a) as ornaments along one parametric direction of the surface of the body, whereas in Figures 8 (c) and (d) the teapot's body employs the cross section of Figure 5 (b) along the other parametric direction. The teapot's lid in all four figures of Figure 8 employs the C^0 continuous cross section shown in Figure 1 (a) along the two parametric directions of the surface of the lid whereas the spout and the handle use the cross section shown in Figure 1 (b). Figure 8 shows the body, the lid, the handle, and the spout with ornaments along u or v isoparametric curves.

The height of the cross section could be modulated as a function of the parameter of the rail curves. Recalling Equation (9), we can add a modulating scaling function, $s(t)$, to the normal field:

$$S_o(r,t) = A_o(t) + D_o(t)x_s(r) + n(u(t),v(t))s(t)y_s(r) + T_d(t)h_{10}(r) + T_{-d}(t)h_{11}(r).$$

Figure 9 (a) shows the Utah teapot with no modulating function and a cross section in the shape of Figure 5 (c). Figure 9 (b) is identical to Figure 9 (a) except for a wavy modulating function $s(t)$. In Figure 9 (c), the modulating function is in the shape of a discontinuous square wave.

So far we have seen examples that modulate the height of the constructed blending surface. Clearly, one can add similar modulating functions to other degrees of freedom and, for example, vary the magnitude of the tangent fields $T_i(t)$ or even the width of the computed offset, d, making it a function of t as well, as $d(t)$. In Figure 10, we present three examples of univariate ornaments laid over a torus freeform surface with vaying width and height. The green ornament in Figure 10 uses a scaling field of varying width of four sine cycles. The red ornament uses the same four sine cycles to control the height of the geometry whereas the cyan ornament exploits the scaling field to modulate both the width and height of the shape.

6 Conclusions and Future Work

This work presented a generalization of the concept of blending surfaces toward either functional or decorative purposes. The G^1 blending capabilities are augmented with intuitive and powerful control over the interior cross section shape of these constructed surfaces.

Ornaments placed on the surface can be either closed loops or open ended. In the latter case, either the ornament should be scaled down to an infinitesimal size or, alternatively, caps should be placed at the two ends. Careful reexamination of Figure 7 would reveal that the end points of the digits are, in fact, open. Two caps at the two end points could easily resolve this issue.

As demonstrated, with the aid of cubic Hermite functions, ornamental shapes with G^1 continuity can be placed on or connected to different surfaces. G^2 or higher order continuity can be achieved in a similar manner, by employing quintic Hermite functions [6] or even higher order Hermite functions.

References

[1] C. L. Bajaj and I. Ihm. Algebraic Surface Design with Hermite Interpolation. ACM Transactions on Graphics, Vol 11, No 1, pp 61-91, January 1992.

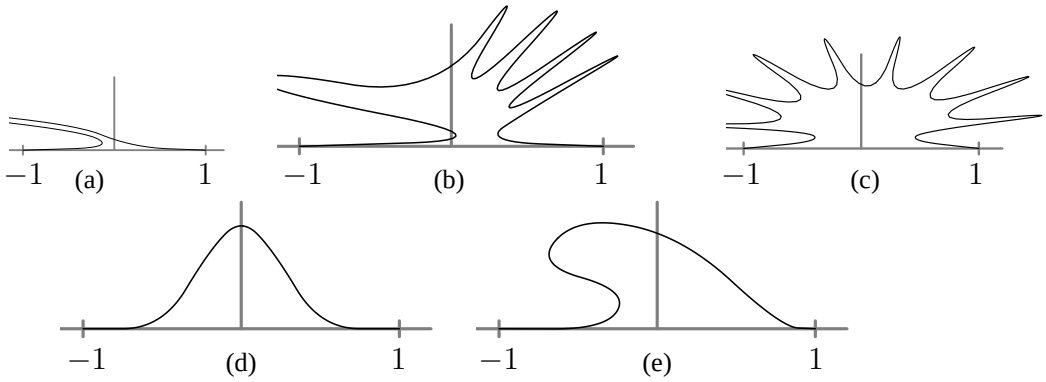

Figure 5. Cross sections used in the decoration of the Utah Teapot in Figures 6, 7, 8, and 9.

[2] A. P. Bien and F. Cheng. A Blending Model for Parametrically Defined Geometric Objects. Symposium on Solid Modeling Foundation and CAD/CAM Applications, pp 339-347, 1991.

[3] C. Barghiel, R. Bartels, and D. Forsey. Pasting Spline Surfaces. Mathematical Methods for Curves and Surfaces, M. Daehlen, T. Lyche, L. Schumaker (eds.), Vanderbilt University Press, pp 31-40, 1995.

[4] M. Bloomenthal, University of Utah, 1990. Private communications.

[5] B. K. Choi and S. Y. Ju. Constant Radius Blending in Surface Modeling. Computer Aided Design, Vol 21, No 4, pp 213-220, May 1989.

[6] E. Cohen, R. F. Riesenfeld, and G. Elber. Geometric Modeling with Splines: An Introduction. A. K. Peters, Natick Massachusetts, 2001.

[7] G. Elber. Symbolic and numeric computation in curve interrogation. Computer Graphics forum, Vol 14, pp 25-34, 1995.

[8] D. J. Fillip. Blending Parametric Surfaces. ACM Transaction on Graphics, Vol 8, No 3, pp 165-173, July 1989.

[9] K. Fleischer, D. Laidlaw, B. Currin, and A. Barr. Cellular Texture Generation. SIGGRAPH 95 Conference Proceedings, pp 239–248, Aug. 1995.

[10] J. D. Foley, A. van Dam, S. K. Feiner, and J. F. Hughes. Fundamentals of Interactive Computer Graphics. Addison-Wesley Publishing Company, second edition, 1990.

[11] D. R. Forsey and R. H. Bartels. Hierarchical B-Spline Refinement. SIGGRAPH 88 Conference Proceedings, pp 205-212, Aug. 1988.

[12] K. C. Hui. Freeform design using axis curve-pairs. Computer Aided Design, Vol 34, No 8, pp 583-596, 2002.

[13] IRIT 8.0 User's Manual. The Technion—IIT, Haifa, Israel, 2000. Available at http://www.cs.technion.ac.il/∼irit.

[14] K. Kim and G. Elber. New Approaches to Freeform Surface Fillets. The Journal of Visualization and Computer Animation, Vol 8, No 2, pp 69-80, 1997. Also the Third Pacific Graphics Conference on Computer Graphics and Applications, Seoul, Korea, pp 348-360, August 1995.

[15] R. K. Klass and B. Kuhn. Fillet and Surface Intersections Defined by Rolling Balls. Computer Aided Geometric Design 9, pp 185-193, 1992.

[16] J. R. Rossignac and A. A. G. Requicha. Constant Radius Blending in Solid Modelling. Computers in Mechanical Engineering, pp 65-73, July 1984.

[17] K. Vida, R. R. Martin, and T. Varady. A Survey of Blending Methods that Use Parametric Surfaces. Computer Aided Design, Vol 26, No 5, pp 341-365, May 1994.

[18] J. Warren. Blending Algebraic Surfaces. ACM Transactions on Graphics, Vol 8, No 4, pp 263-278, August 1989.

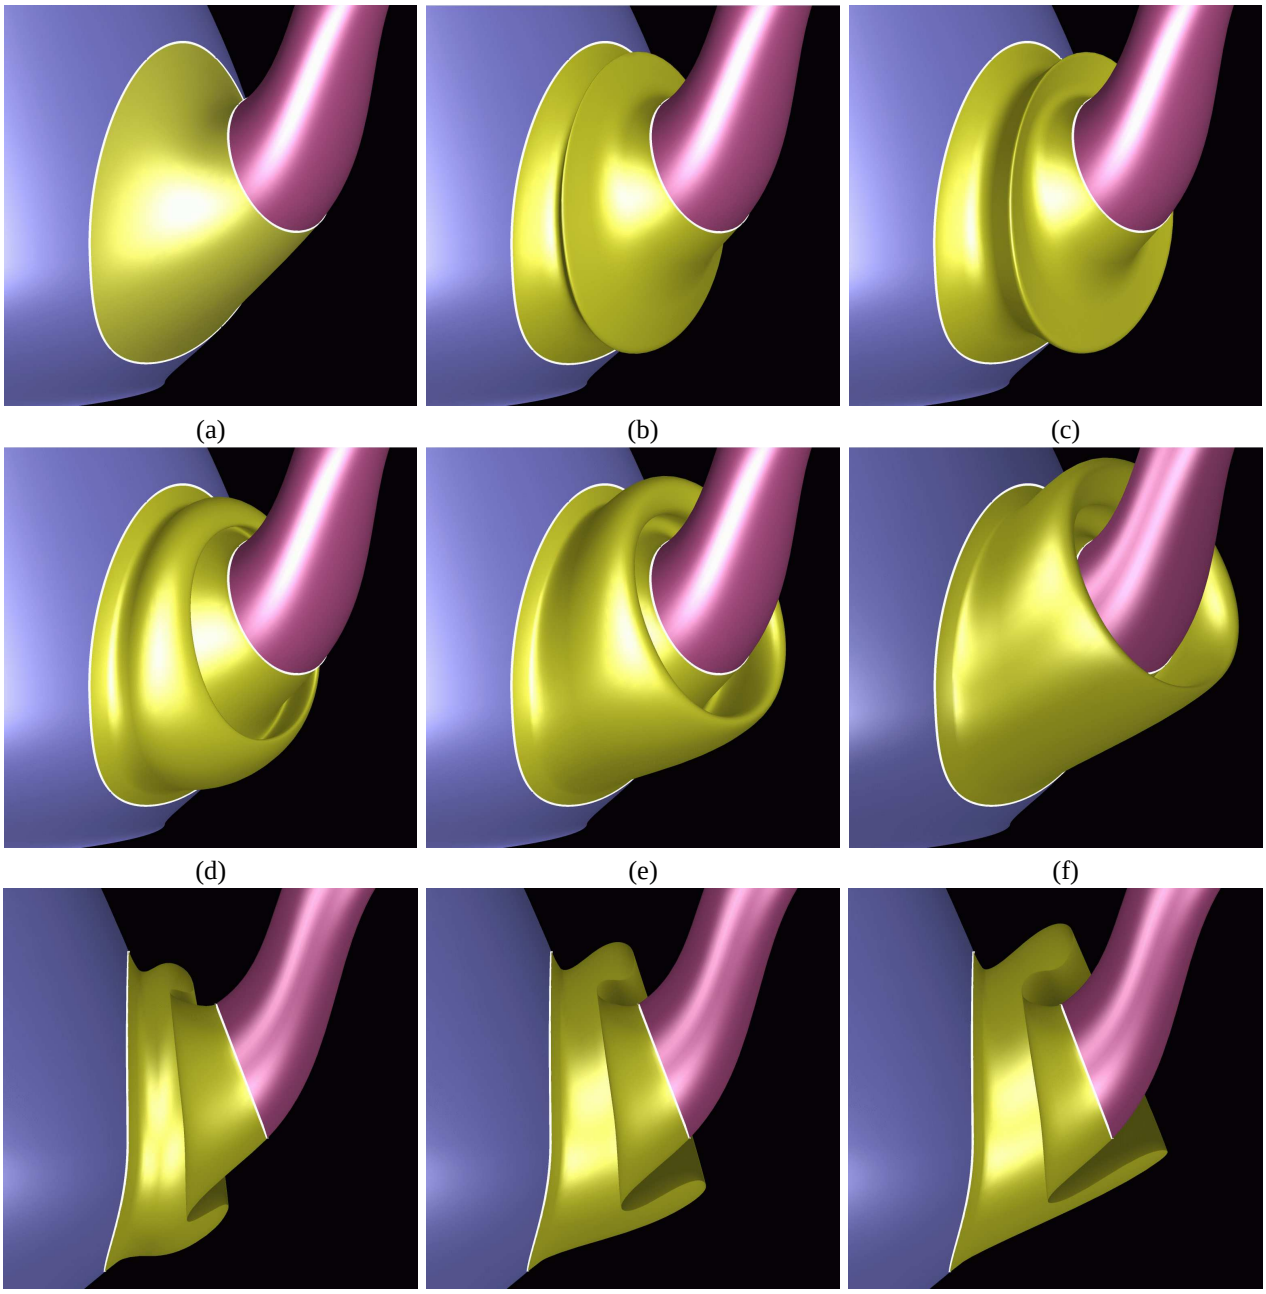

Figure 6. An example of a rounding between the body of the Utah teapot and its spout. (a) shows an example of a regular cubic Hermite blend. In (b) to (i), blending variants using different cross sections and parameters are shown. In (b) and (c), a simple cross section was used, a cross section that is shown in 5 (d). In (d) to (i), a more complex cross section was employed, a cross section that is shown in 5 (e). (d) to (f) show the geometry whereas (g) to (i) present their side cross sections, respectively.

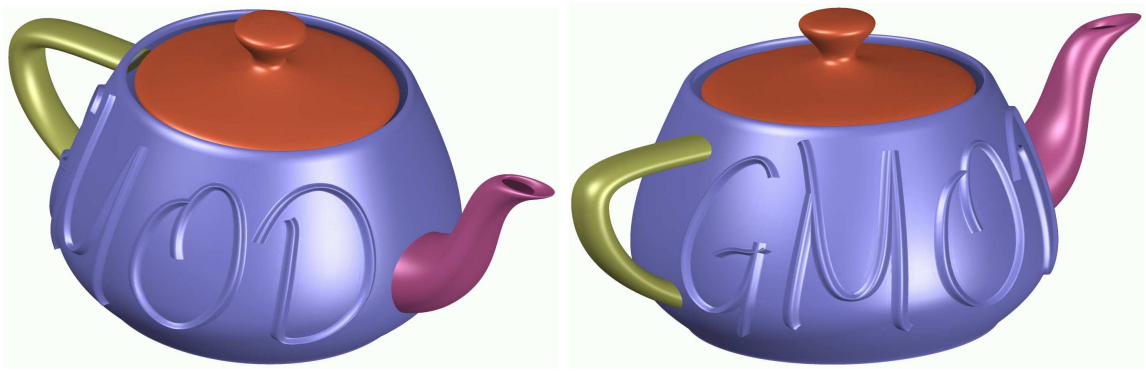

(a) (b)

Figure 7. Three-dimensional features over freeform surfaces. Shown are two views of the letters 'GMOD' etched over the body of the Utah teapot using a simple, single bump, cross section (see Figure 5 (d)).

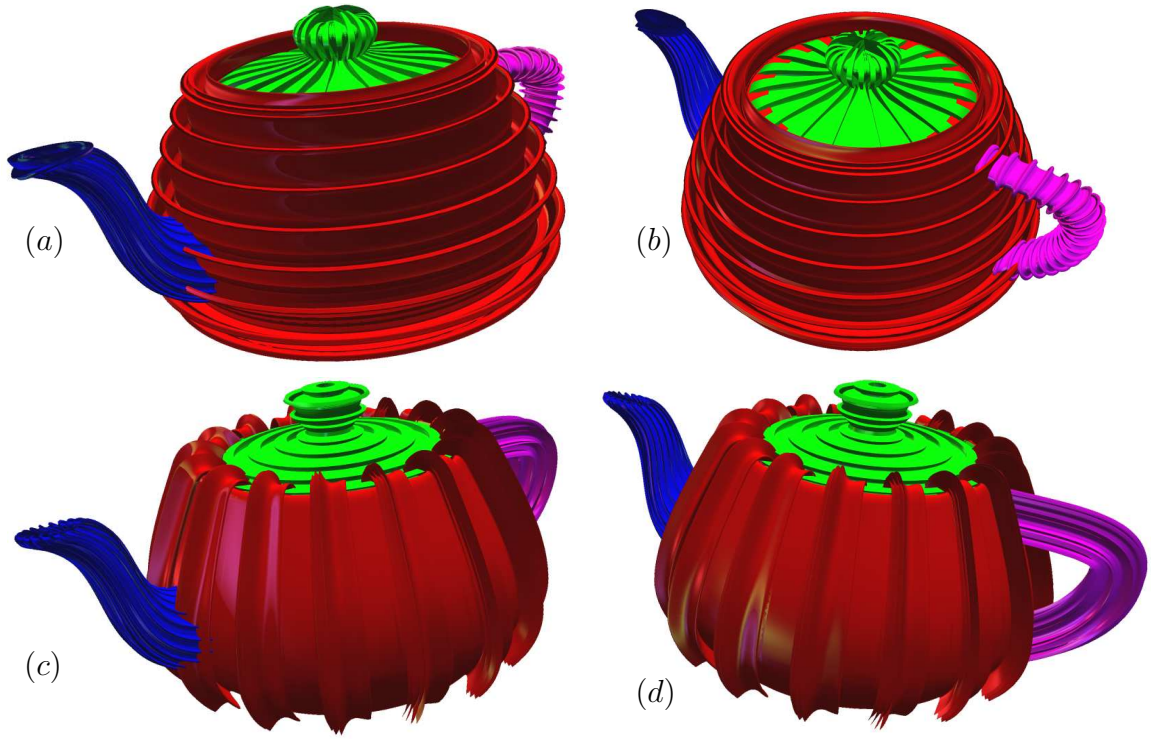

Figure 8. More examples of univariate surface details over freeform surfaces. The Utah teapot is presented with ornaments that employ the cross sections shown in Figures 1 and 5.

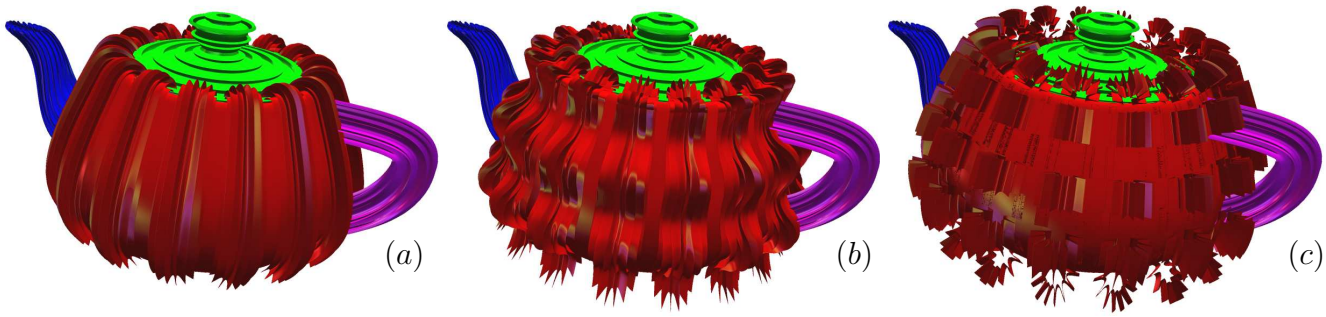

Figure 9. Modulating the univariate ornaments on the surface. The body's ornaments without modulation (a), with a continuous wavy modulation (b) and with a discontinuous square wave modulation (c).

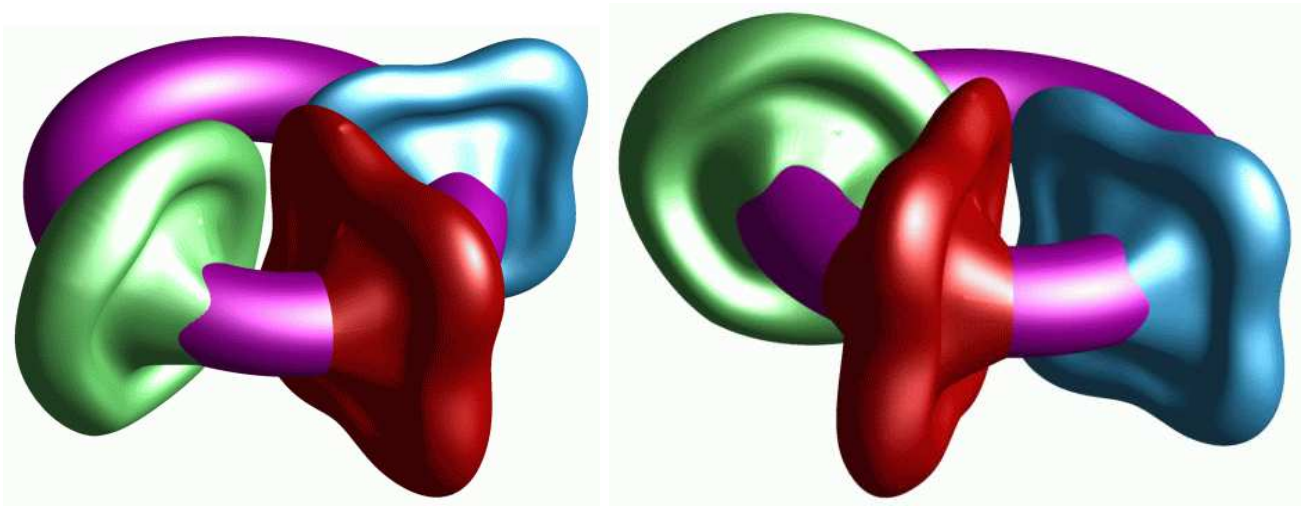

(a) (b)

Figure 10. Three-dimensional features over freeform surfaces with varying width and height. Shown are two views of three univariate ornaments laid over a torus. The green example uses varying width, the red example uses varying height and the cyan example uses both. All three examples use a simple, single bump, cross section (see Figure 1 (b)).