Computer Aided Geometric Design

An elementary and superficial course

Malcolm Sabin

Section 8: more surfaces

Chapter 27

Bezier Triangles

27.1 Problems with the Tensor Product Formulation

There are many circumstances where the rectangular structure of the tensor product surface is most appropriate.

It can give very good results when the surface being designed has a natural flow. Lots of surfaces do.

– aircraft fuselages and wings

– many parts of cars

– handles of consumer products

However, this same inherent rectangularity can give real problems when the surface has a region where natural flows come together awkwardly. Lots of surfaces do.

– aircraft wing roots

– top of the windscreen pillar on cars

– where the handle meets the body in consumer products.

– bulbous bow on ships

– suitcase corner

The tensor product surface has a much higher order for curves crossing the patches diagonally than for those running along the generators.

If a feature runs skew to the natural flow, it is prone to ripple.

This higher order means that the patch is actually higher order than the original univariate basis functions, (the bicubic has 6th order u^3v^3 terms in it), and this impairs the conditioning and reduces the accuracy of computations.

27.2 Triangular Lattice Approach

The high order can be dealt with by using basis functions of given total order. These form naturally over triangular partitionings of the parameter plane.

If we deal with polynomials we are talking about terms using complete rows in the triangle

$$
\begin{array}{ccccccccc}
 & & & & 1 & & & & \\
 & & & u & & v & & & \\
 & & u^2 & & uv & & v^2 & & \\
 & u^3 & & u^2v & & uv^2 & & v^3 & \\
\dots & & \dots & & \dots & & \dots & & \dots
\end{array}
$$

They are most elegantly handled by using homogeneous (barycentric) coordinates u, v, w, $u+v+w=1$ over the triangle $u, v, w \geq 0$

Raise the normalisation expression to the nth power

$$(u+v+w)^n$$

The terms of this are the basis functions, just as the terms of $(\bar{t}+t)^n$ are the basis functions for the univariate Bezier curve.

They obviously satisfy

$$\phi(u,v,w) \geq 0$$

and

$$\sum \phi(u,v,w) = 1$$

The cubic Bezier triangle has terms

$$
\begin{array}{ccccccc}
 & & & w^3 & & & \\
 & & 3uw^2 & & 3vw^2 & & \\
 & 3u^2w & & 6uvw & & 3v^2w & \\
u^3 & & 3u^2v & & 3uv^2 & & v^3
\end{array}
$$

Exercise

A Expand $(u+v+w)^5$ to create the basis functions of a Bezier quintic triangle.

27.3 Handles for a Bezier Triangle

Each basis function has a maximum at a different place within the parameter triangle, and we use these maxima to visualise the structure of the coefficients.

They form a regular triangular lattice.

$$\begin{array}{ccccccc} & & & P_{003} & & & \\ & & P_{102} & & P_{012} & & \\ & P_{201} & & P_{111} & & P_{021} & \\ P_{300} & & P_{210} & & P_{120} & & P_{030} \end{array}$$

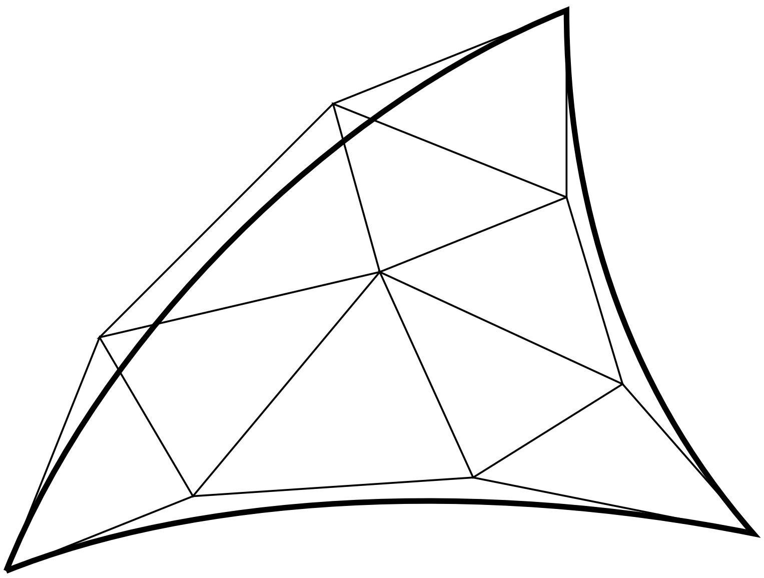

27.4 Boundary Conditions

At the point $u = 1$, $v = w = 0$ the only nonzero function is u^n, and so its coefficient (P_{300}) is the corresponding corner of the triangular patch.

Along the edge $u = 0$, $v + w = 1$, all the terms containing u are zero, and so that edge of the patch is the Bezier curve defined by the coefficients (P_{003}, P_{012}, P_{021} and P_{030}) of the corresponding edge of the control point array.

Similarly, the variation of tangent plane along each edge is determined only by the two rows of control points.

Exercise

A The tangent vector along the edge is easily determined from the edge being a Bezier curve. The tangent vector across the edge is less easily nailed down, because it could be the tangent to either of the other two sets of iso-parametric lines.

In fact

$$\left(\frac{dP}{du}\right)_{w=const} = -\left(\frac{dP}{dv}\right)_{w=const} = \left(\frac{\partial P}{\partial u}\right) - \left(\frac{\partial P}{\partial v}\right)$$

and

$$\left(\frac{dP}{du}\right)_{v=const} = -\left(\frac{dP}{dw}\right)_{v=const} = \left(\frac{\partial P}{\partial u}\right) - \left(\frac{\partial P}{\partial w}\right)$$

Using these results, express each of the three possible tangent vectors as Bezier functions in terms of the original triangle's control points.

27.5 Joining Bezier Triangles Together

27.5.1 Affine transformations in the parameter plane

In the case of tensor product surfaces, we are stuck with keeping the alignment of u and v. We can, however, apply independent affine transformations to the two axes separately.

In the case of triangular patches described by barycentric coordinates we can apply arbitrary affine transformations of the form

$$u = au' + bv' + cw'$$
$$v = du' + ev' + fw'$$
$$w = gu' + hv' + jw'$$

(provided that u, v and w remain independent, so that the determinant of

$$\begin{bmatrix} a & b & c \\ d & e & f \\ g & h & j \end{bmatrix}$$

remains non-zero), thus mapping any triangle into our standard triangle.

The actual mapping merely depends on which triangle in parameter space we are mapping from.

27.5.2 Piecewise Basis Functions

It is therefore possible to characterize a piecewise polynomial basis function by just drawing the collection of triangles over which it is defined, and writing the Bezier control values as numbers at the right places in the diagram.

27.5.3 Achieving continuity

This is the first point at which we find a problem with Bezier triangles.

There is no difficulty with C0: just share the edge control points.

To get C1 we need to make sure that the quadrilaterals along the seam, formed by the little control point triangles, are all affine transformations of the same plane quadrilateral.

This means that the ring of control points around the common corner of a number of original triangles must be planar.

It also means that the parameter triangles must all be affine maps from some triangulation of (a part of) a single plane.

You have to go up to at least quartic triangles before this all becomes practicable for a general triangulation. Even then it is not easy to get them all to sew together. Vertices with an even number of triangles have a big problem, and you have to massage the difficulties towards vertices with an odd number of triangles.

27.6 Subtriangulations

However, there are three patterns worth knowing about, because they work out easy.

- Clough–Tucker cubics.
- Powell–Sabin 6–way quadratics.
- Powell–Sabin 12–way quadratics.

All three lend themselves to defining C1 surfaces over an arbitrary triangulation.

All three work by subdividing each triangle into smaller triangles in a particular way.

27.6.1 Clough Tucker cubics

Given a triangle in which the values and first derivatives are known at the three corners.

To Construct a function over the triangle continuous in position and first derivative internally and across all the boundaries with adjacent triangles sharing the same vertex data.

Construction Take a point in the interior of each triangle, and join it to the three vertices. In each of the three pieces put a cubic function.

The actual function for each piece is conveniently expressed in terms of the Bezier representations of the subtriangles.

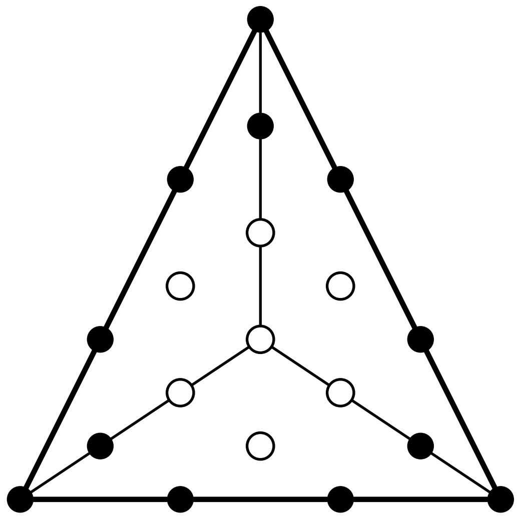

The control points adjacent to the original vertices are defined by the vertex data.

The control points in the interior of each piece are defined by making some arbitrary choice of cross-tangent vector at the mid-point of each edge. This arbitrary choice must obviously be made consistently in adjacent triangles. A possible choice (which loses invariance under affine transformation of the parameter plane) is to assume a linear variation of perpendicular cross slope along the edge.

We now have only the central control point and the edge control points adjacent to it to fix. The latter are fixed from the continuity of slope condition between adjacent subtriangles, and then the central control point is just coplanar with its three neighbours.

Exercise

A Prove that this construction satisfies the requirements.

27.6.2 Powell Sabin 6-way quadratics

Given a triangulation of the plane in which the values and first derivatives are known at all the vertices.

To Construct a function over each triangle continuous in position and first derivative internally and across all the boundaries with adjacent triangles sharing the same vertex data.

Construction Take a point in the interior of each triangle, and join it to the three vertices, and to the internal points of adjacent triangles. The choice of internal points must be such that each join between interior points cuts the common edge of the two triangles. This can always be done by taking the points at the incentres of the triangles.

In each of the six pieces put a quadratic function.

The actual function for each piece is conveniently expressed in terms of the Bezier representations of the subtriangles.

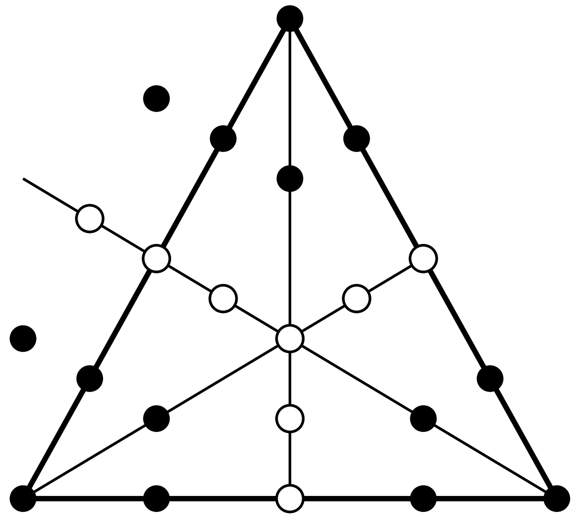

The control points adjacent to the original vertices are defined by the vertex data.

All the remaining control points except the central one are fixed by the slope continuity constraint, and the final central point by coplanarity.

Exercise

A Prove that this construction satisfies the requirements.

27.6.3 Powell Sabin 12-way quadratics

Given a triangle in which the values and first derivatives are known at the three corners.

To Construct a function over the triangle continuous in position and first derivative internally and across all the boundaries with adjacent triangles sharing the same vertex data.

Construction Cut the triangle into 12 pieces by the medians and the lines joining the midpoints of the sides.

In each of the 12 pieces put a quadratic function.

The actual function for each piece is conveniently expressed in terms of the Bezier representations of the subtriangles.

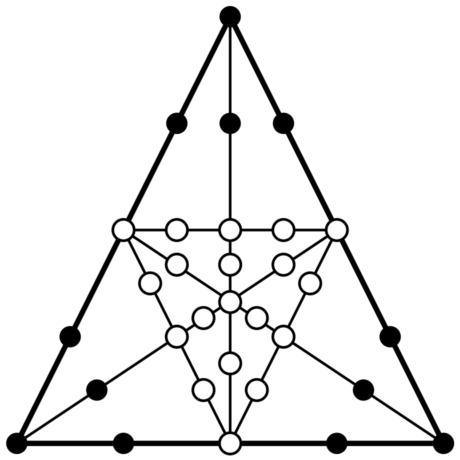

The control points adjacent to the original vertices are defined by the vertex data.

The control points near the original edges are defined by making some arbitrary choice of cross-tangent vector at the mid-point of each edge. This arbitrary choice must obviously be made consistently in adjacent triangles. A possible choice (which loses invariance under affine transformation of the parameter plane) is to assume a linear variation of perpendicular cross slope along the edge.

The remaining control points are easily constructed using the continuity condition, and coplanarity for the central point.

Exercise

A Prove that this construction satisfies the requirements.

27.7 Affine transforms in the parameter line

Any finite interval of the real line can be mapped to the interval by an affine mapping of the form

$$t = at' + b$$

When talking about piecewise curves as collections of bits, each bit can be parametrised from 0 to 1, but when we think of the whole thing as a single piecewise curve, we have to have a single parametrisation running the whole length of the curve. We therefore use remapping of the parameter line implicitly. Each piece has its own values of a and b.

If all the pieces share the same a value, we call the piecewise subdivision **equal interval**.

27.7.1 Piecewise basis Functions

When we talk about piecewise basis functions, we are joining together not the bits of curve themselves, but bits of basis function. However, if we use Hermite or Bernstein to describe the curve bits, these methods are linear in the coefficients, and therefore work for any dimension, including 1 ! We can play exactly the same affine mapping trick to describe scalar-valued basis functions. These are easier to draw, because we can just plot the basis-function value against the parameter value. We can also use the coefficients in the 1D equation of the basis function to characterise it. We call these **control values** rather than control points to stress that they are only scalars.

In the equal interval case, a sequence of Bezier control-values, or a sequence of Hermite value/derivative pairs, is also a complete characterisation of what is important about a basis function.

Chapter 28

Simplex Splines

We saw earlier that the univariate Bspline basis can be generated by partitioning the tensor product of a simplex and a line into subsimplices, and then associating with each subsimplex a function over the points of the line.

The same theory generalises to splines over spaces of higher dimensions.

28.1 Construction

Take a simplex of unit measure in a space of s dimensions, and a (polytopal) parameter region in a space of n dimensions.

The tensor product of these two is in $s + n$ dimensions.

Divide the resulting object into simplices. Each vertex should lie at the product of one of the original simplex vertices and a **knot point** in the parameter region.

Each piece of this partitioning induces a scalar function over the parameter region from the measure (area, volume) of the part of the piece which maps from a particular point in the parameter region.

28.2 Properties

There is a function ψ_i for each simplex piece, and these functions have the following properties

$$\psi_i \geq 0$$

$$\sum_i \psi_i = 1$$

Support is typically finite, though a partitioning can sometimes be devised in which every piece covers the whole parameter region.

Each ψ is piecewise polynomial of order s, and with continuity of derivative $s-1$

The amount of continuity will typically drop when more than two knot points of a single piece are collinear or more than three coplanar, etc. (other than at the boundaries of the parameter region).

28.3 Example

All this is very hard to visualise, except for the most trivial case, when the tensor product is between an interval (a 1-simplex) and a polygonal parameter region in 2 dimensions.

This gives a slab, which is partitioned into tetrahedra.

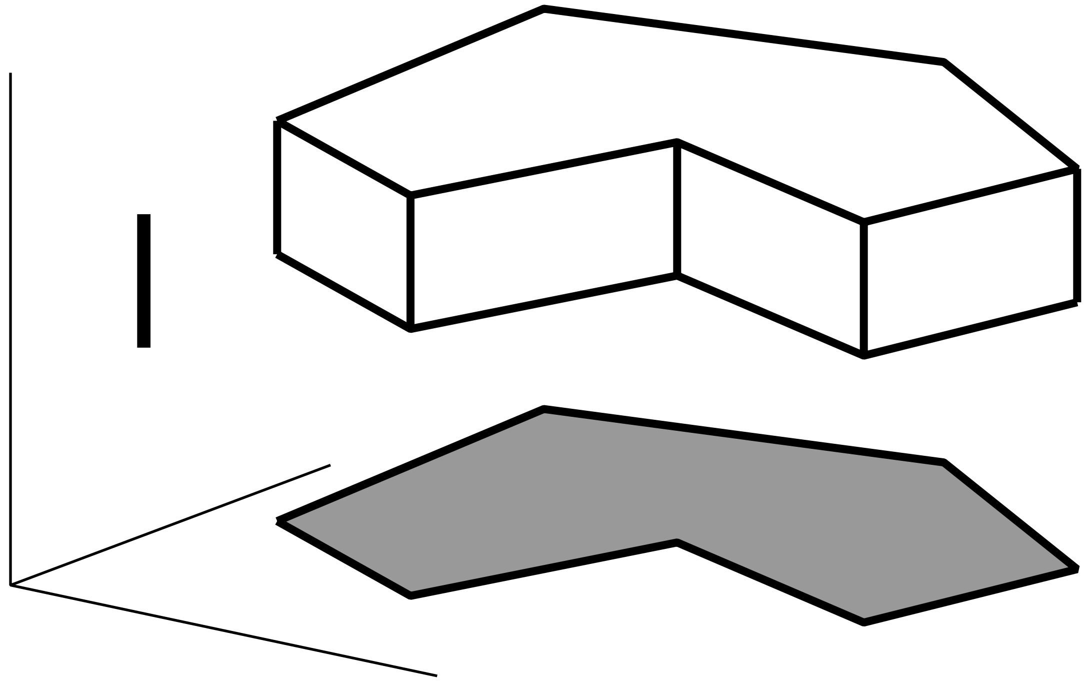

Each tetrahedron is then projected down on to the parameter space, to give either a 3-sided or a 4-sided shadow. Within each shadow there is a function defined as the height of the intercept of the vertical line there with the tetrahedron.

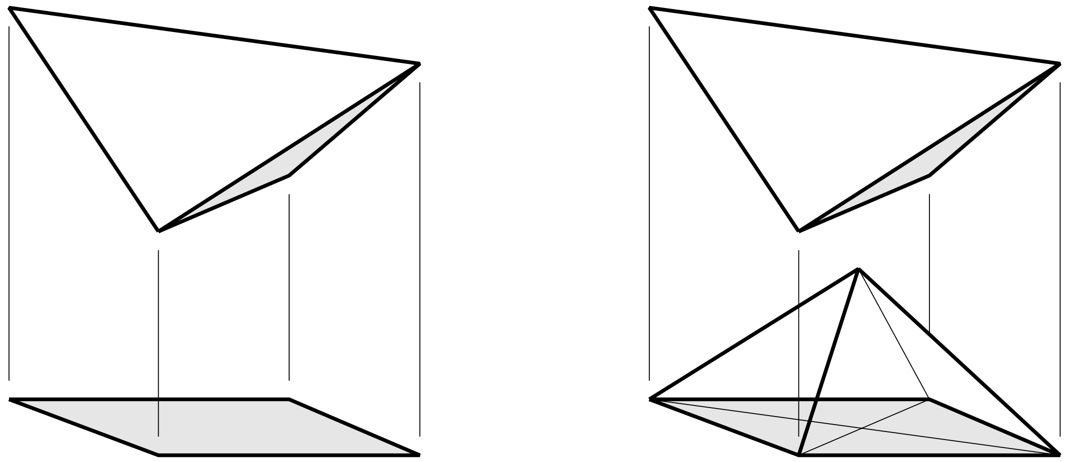

These functions are position continuous, and consist of linear pieces. No big deal, but getting slope continuity from quadratic pieces and curvature continuity from cubic pieces in as many dimensions as you like is really amazing. Compare the tensor product which has 6th order pieces for bicubics or 9th order for tricubics with only curvature continuity.

Exercise

A Define a set of linear simplex splines over a square parameter space with knots at the corners, by taking the tensor product with an interval. Express each of the functions in terms of the triangular bezier representation of its non-zero pieces. Hint: you will probably need five pieces, and to avoid singularity, you will probably need to use one diagonal in the top face and the other diagonal in the bottom face.

Chapter 29

Box Splines

Consider the regular triangular grid.

Although sewing triangular patches together was presented as being hard in an earlier chapter, we can in fact construct spline-like functions by using convolution.

29.1 Bivariate Convolution

- Shifting is well-defined in the lattice directions.
- Subtraction works because our whole system is linear.
- The derivative of a Bezier triangle in one of the three mesh directions can be constructed by taking differences. Integration is just the inverse operation.

The combination is a directional convolution.

Note that:-

– Shifting increases the support area

– Integration increases the order of each triangle by 1

– Integration increases the continuity across two of the three directions by 1.

Thus a combination of three convolutions, one in each direction, increases the order by 3, the continuity by 2, and adds one layer of triangles all the way round the support region.

Exercises

A Identify the numerical operation on the coefficients of a bezier triangle, corresponding to integration in a direction parallel to an edge. Hint: the constant of integration is an arbitrary polynomial of the right order along one of the edges.

B Why is the constant of integration not arbitrary when using integration inside convolution ?

29.2 Splines over regular triangulations

By applying three convolutions to a unit impulse at one vertex we get the linear function with a plane face in each triangle. This function is C0, and is defined over a hexagon of side 1 unit.

A second cycle of convolutions gives a C2 function, with quartic pieces, defined over a hexagon of side 2 units.

Starting from a function unity over a single triangle and zero elsewhere, gives a C1 function, with cubic pieces, defined over a slightly irregular hexagon with sides of alternately 1 and 2 units.

29.3 Box-splines

These functions are currently known as **Box-Splines**, because they can be viewed as the shadow of an n dimensional box, in just the same way as the simplex splines are shadows of simplices.

(Note that because a box can be partitioned into simplices, these functions must be sums of simplex spline basis functions. The simplex splines are without doubt the more fundamental, but the box-splines turn out to be a less spiky and uncooperative form with which to work.)

In particular, the function with six linear triangles is the shadow of a unit 3–dimensional box, viewed diagonally.

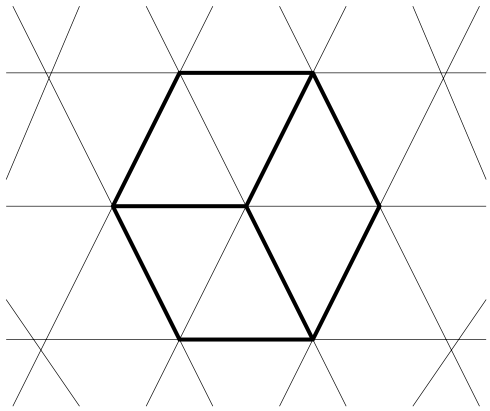

The quartic function is the shadow of a unit 6–dimensional box, two of whose edge directions project into each of the lattice directions.

29.3.1 The 4-direction mesh

There is also a slightly less useful function, defined over a rectangular mesh plus diagonals, which is the shadow of a unit 4–dimensional box one of whose edges projects into each of the four directions.

It has C1 continuity and quadratic pieces. It is defined over an octagon.

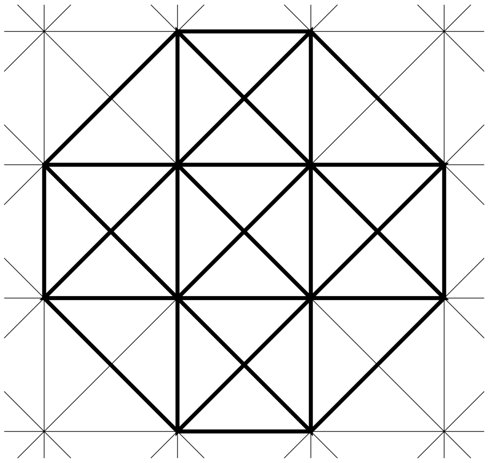

It is less useful because there are two subsets of the complete basis which each sum to 0.5, so that it is not possible to use this basis for interpolation on a regular lattice.

Exercise

C Elaborate the shape of the 4-direction basis function in terms of its bezier triangle coefficients in each of the pieces.

Hints: Don't worry about scale, you can always divide by the sum afterwards. Do use symmetry to reduce the number of separate numbers you have to think about. There are only 11 numbers in an octant of the function, and five of them are zero before you start.

29.3.2 Tensor Product splines as box-splines

Finally, the tensor product Bsplines can be viewed as being shadows of $2n$-dimensional boxes, with n edge directions projecting into each of the u and v directions.

29.4 Multivariate Unequal–interval Bsplines over triangulations

Sorry, this theory just doesn't exist. There is no equivalent of the deBoor recurrence over a general (irregular) triangulation. The only theories we have are:-

(i) simplex splines, where the triangulations are limited to those which emerge as shadows of the edges of the simplices. We do not have general ways of partitioning the $m+n$-dimensional slab even into triangulations with given vertices.

(ii) box splines, where there are strict regularity assumptions, and which generate functions over the regular and semiregular partitionings of the parameter plane.

There have been a few attempts at a more general theory, but these just do not feel 'right' yet. The main thrust has been to build a structure in which the Bezier triangles come out as special cases of something more general. There is nothing easily usable for general purposes except the subtriangulations covered in chapter 26.

Important names are:-

Wolfgang Dahmen

Paul Dierckx

Klaus Hollig

Charles Micchelli

Hans Peter Seidel

Chapter 30

N-sided Patches

We have seen how recursive subdivision works well to achieve topological freedom as long as only slope continuity is required. The quadratic recursive subdivision does have a small geometric drawback, that the distribution of curvature within the resulting surface is not ideal. Starting from a cube of 8 control points, one would expect a surface fairly close to a sphere to result. In fact the radius towards one of the control points is significantly greater than towards the middle of a face, and it is fairly easy to prove that some inequality must be the case with a quadratic subdivision.

The alternative approach is a rather ad hoc collection of n-sided patches which can be used, either to fill in the hole after a small number of recursive division steps, or else directly.

The patches currently known also tend to be limited to giving slope continuity, though there is no obvious reason why this is an inherent limit.

30.1 Stuffed triangles

The reason why the standard triangle theory is so rigid is that the tangent vector along an edge is the same order polynomial as the tangent vector across that edge. This means that the equation for tangent plane continuity has no spare freedoms.

We can reintroduce one by requiring that the actual edge of a patch should be one order lower than the interior of the patch, so that a cubic triangle, for example, could be required to have quadratic edges.

This reintroduces a freedom which the tensor product four sided patch has, at the expense of ensuring that the patch order is at least two higher than the continuity.

30.2 Non-Planar Parameter space

One blatant approach to representing the geometry of manifolds which do not map conveniently to a planar parametrisation, is to use a parameter manifold which is geometrically simpler than the target surface, but has the required topology.

One example of this is to use the triangle,

$$u + v + w - 2uvw = 1,\ u, v, w \geq 0$$

The effect of the extra term is to alter the basis so that an extra term appears 'in the middle'. The triangular surface becomes

$$\begin{gathered} P = Au^2(1 - 2vw) + Bv^2(1 - 2uw) + Cw^2(1 - 2uv) \\ +2Dvw(1 - u^2) + 2Ewu(1 - v^2) + 2Fuv(1 - w^2) \\ +4Guvw \end{gathered}$$

There exists a tangent across the $u = 0$ edge of this patch which is equal to $[F - B]v^2 + 2[G - D]vw + [E - C]w^2$, and so each edge can be made slope continuous easily with an adjacent biquadratic, just by lining up the control points. In fact a Bspline-like network can be built, containing only quad faces, but with 3-valent, 4-valent or 5-valent vertices.

The analogous 5-sided patch uses a parameter space in 5-dimensions, with the three normalisation equations

$$\begin{aligned} u_4 &= 1 - u_1 u_2 \\ u_3 &= (1 - u_1)/u_4 \\ u_5 &= (1 - u_2)/u_4 \end{aligned}$$

Exercise

A Determine u_4, u_5 and u_1 in terms of u_2 and u_3.

There is also a parameter space for a six-sided patch, but the equations relating the other parameters to two adjacent ones require the solution of a quadratic equation.

There is continuing work at Birmingham (J.J.Zheng), who has produced basis functions for the six-sided patch and is reputed to have produced a complete family of n-sided patches using the same principles.

30.3 Boundaries not straight in the parameter plane

The use of non planar boundaries can also often be expressed in terms of bits of parameter plane which have curved boundaries. For example, if we were mapping from an octant of the sphere, we could reduce the parametrisation to a locally planar one by making a stereographic projection from a point in the middle of the far octant. The triangular patch would now be a map from a region of the parameter plane bounded by three circles. Avoiding having to carry the normalisation around all the time could well mean that this would be worth while.

Recent work on transfinite interpolation (I invented a 2-sided patch, which you can only do if you have curved boundaries) suggests that this may turn out to be a fruitful way of looking at things.

30.4 Composite Patches

The additional freedom which exists in the 4-sided patch may be used, with care, to create an n-sided region built from n 4-sided patches, which behaves across its boundaries exactly like a single bicubic.

These depend on the freedom which the four sided patch has, that because the longitudinal derivative along an edge is one order lower in the parameter varying along that edge than the derivative across that edge, there can be a linear variation of cross-derivative.

Suppose that we have two patches P and Q meeting at a common edge u =constant. Then at that edge we can have

$$dP/du - \alpha dQ/du = (\beta + \gamma v)dP/dv$$

This approach has been explored in recent years by the late John Gregory and Vincent Lau at Brunel University, by Jorg Peters at Purdue and by Charles Loop at University of Washington.

30.5 Transfinite Patches with slope control

Within the literature there are a number of patches which define the interior not from a fixed finite number of control points or other characteristics, but from the boundary curves and possibly boundary cross-slopes.

The Coons Patch was the first of these, but Barnhill produced a triangle and Gregory a general formulation for $n > 4$

These patches mapped from a regular n-gon in the parameter plane.

Any such patch can be turned into a finite-dimensional patch matching adjacent bipolynomial patches by constraining the boundaries to have the required form. It is then trivial to show that the patch will match its neighbours C0.

The Gregory patch was extended to match cross-derivative boundary conditions, which thus give C1 continuity.

Unfortunately, the taking of linear combinations of estimators based on the ends of the common edge does not give a patch matching higher derivatives at the edge, and so it is not so easy to produce a C2 patch.

30.6 Warren patches

A recent development by Joe Warren, of Rice University, has been to observe that if all the coordinates of an end control point of a rational Bezier curve are zero, the limit to which the curve tends as the parameter approaches the end value is the next control point

Let O denote the 'point' whose homogeneous coordinates are $[0,0,0,0]$, and consider the cubic case.

$$\begin{aligned} P(t) &= O\bar{t}^3 + 3A\bar{t}^2 t + 3B\bar{t}t^2 + Ct^3 \\ &= t(3A\bar{t}^2 + 3B\bar{t}t + Ct^2) \end{aligned}$$

The factor of t can be divided out, because we are working in homogeneous space, and the parameter can be given the projective transformation $t = \sqrt{3}u/(1+(\sqrt{3}-1)u)$ which balances teh weights at the two ends.

The result differs from the quadratic with control points A, B and C only in a slight change of the weight associated with the central point, which is decreased by a factor of $\sqrt{3}/2$

If one vertex of a Bezier triangle is placed at $[0,0,0,0]$ the limits as that vertex is approached in different directions vary, so the parametric vertex 'blows up' into, in this case, a straight line, thus adding an edge to the patch.

If we take a 6th order Bezier triangle, and place 3 control points at each corner at $[0,0,0,0]$ the remaining control points define a hexagonal patch with six quadratic sides.

This looks a very promising approach, but what is not yet clear is how to build parametrisations which are well-conditioned everywhere.

30.7 Comparison

Clearly a composite patch is different from one which has all derivatives continuous in its interior. It may well turn out, however, that the other approaches are all different ways of thinking about the same thing. This has not yet been proven, but neither have the differences been clearly shown.