Cubic precision Clough-Tocher interpolation

Computer Science Department
University of Waterloo
Research Report CS-98-15

Stephen Mann

Supported by the Natural Science and Engineering Research Council of Canada

Abstract

The standard Clough-Tocher split-domain scheme constructs a surface element with quadratic precision. In this paper, I will look at methods for improving the degrees of freedom in Clough-Tocher schemes. In particular, I will discuss modifications to the cross-boundary construction that improve the interpolant from quadratic precision to cubic precision.

Keywords: Scalar Data Fitting

Abbreviated title: Cubic precision Clough-Tocher interpolation

The Clough-Tocher method [1] is a standard technique for constructing a piecewise polynomial, C^1 surface that interpolates the positions and normals of a function above a triangulation of the plane. To obtain a consistent mixed partial derivative at the vertices of the triangulation, each domain triangle (macro-triangle) is split into three sub-triangles (mini-triangles). For each mini-triangle, the Clough-Tocher scheme constructs a cubic Bézier patch that meets the interpolation and continuity conditions.

After reviewing some background material, I will present the Clough-Tocher technique and then give Farin's modification to this method. Both of these schemes have quadratic precision. Next, I will state the cross-boundary technique of Foley and Opitz, and show how to integrate it into a Clough-Tocher scheme to yield a cubic precision method. I will then present a new variation of Farin's method, which also has cubic precision. After presenting these four schemes, I use isophote plots and shaded to compare them. Finally, I will discuss some iterative techniques, and mention some other Clough-Tocher like schemes.

1 Background

In this section, I will give the relevant background on functional Bézier patches. I will present the barycentric form of Bézier patches, and I will treat the control points as if they were points in a three space (although in the functional case, you only need to work with the z-coordinate of each point). For a more complete introduction to triangular Bézier patches, see, for example, Farin's book [4].

A triangular Bézier patch of degree n specified in barycentric coordinates relative to a domain triangle $\triangle D_0 D_1 D_2$ is given by

$$P(t) = \sum_{\substack{i,j,k\geq 0\\ i+j+k=n}} P_{ijk} B^n_{i,j,k}(t),$$

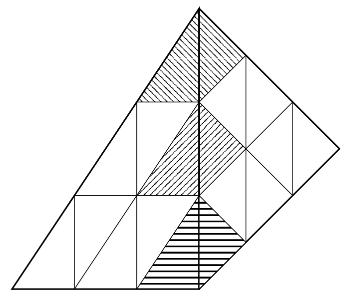

Figure 1: C^1 continuity for two cubics: the shaded panels must be coplanar.

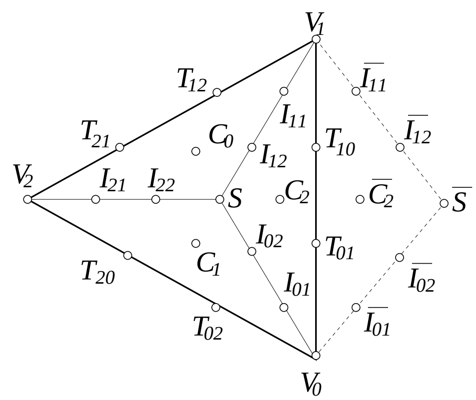

Figure 2: Clough-Tocher control points. The dashed line segments and the "barred" points are on the mini-triangle of the neighboring macro-triangle.

where

$$B^n_{i,j,k}(t) = \frac{n!}{i!j!k!} u^i v^j w^k.$$

Here (u, v, w) are the barycentric coordinates of t relative to the domain triangle. The points P_{ijk} are the control points of the patch. The x and y coordinates of these points are given by

$$P_{ijk}[x, y] = \frac{iD_0 + jD_1 + kD_2}{n},$$

where $P[x, y]$ refers to the x, y coordinates of the three space point P. Thus, the only degrees of freedom are the z-coordinates of the control points.

Clough-Tocher schemes construct networks of patches that meet with C^1 continuity. To achieve C^1 continuity, we must first have C^0 continuity, which is easily obtained by constructing neighboring patches to have the same boundary control points. For C^1 continuity, consider Figure 1. Two cubics meet with C^1 continuity if each pair of shaded panels in this figure are coplanar.

2 Clough-Tocher

Given a triangle $\triangle P_0P_1P_2$ (whose projection into the x-y plane forms the domain triangle $\triangle D_0D_1D_2$) with normals N_0, N_1, and N_2, and using the labeling of the control points from Figure 2, the Clough-Tocher construction fits three cubic patches to the data, setting the z-values of the control points with the following steps: For $ijk \in \{012, 120, 201\}$,

1. Set V_i to P_i
2. Set T_{ij} to lie in tangent plane at V_i.
3. Set I_{i1} to lie in tangent plane at V_i, or equivalently to $I_{i1} = (V_i + T_{ij} + T_{ik})/3$.
4. Set C_i to be coplanar with T_{jk}, T_{kj} and the corresponding C point on the other side of the V_jV_k boundary.
5. Set I_{i2} to lie in the plane spanned by C_j, C_k, and I_{i1}.
6. Set S to $(I_{02} + I_{12} + I_{22})/3$.

The first three steps are required to interpolate the data (P_i and N_i). Step (iv) creates a C^1 join across the external boundaries. And steps (iii), (v), and (vi) are required to obtain C^1 continuity across the interior boundaries. The only degrees of freedom in this construction are in step (iv).

The standard Clough-Tocher scheme sets each C_i by placing it in the plane spanned by T_{jk}, T_{kj} and the vector $I_{j1} - \overline{I_{j1}} + I_{k1} - \overline{I_{k1}}$, where $\overline{I_{j1}}$ and $\overline{I_{k1}}$ are the corresponding points on the other side of the boundary. This creates a cross-boundary tangent vector field that is linearly varying in one domain direction, and quadratically varying in the remaining directions. This choice of direction is not unique; we can choose any domain direction (other than the one parallel to the boundary) in which to have linear variation; we just have to ensure that the same direction is chosen for both patches. Also note that when the data are sampled from a quadratic function, the cross boundary variation is linear in all domain directions.

Using this standard Clough-Tocher cross-boundary field, it is easy to deduce that the Clough-Tocher technique has quadratic precision from the following observations: (1) the boundaries are constructed with cubic precision; (2) the C_i are set with quadratic precision; and (3) the patches constructed are the only cubic patches meeting each other C^1 continuity having a linearly varying cross-boundary derivative.

3 Farin

Farin used the degrees of freedom in the Clough-Tocher scheme to minimize the C^2 discontinuity across the macro-triangle boundaries by solving a small linear system of equations for each boundary [3]. Farin's approach is to initially fit a single cubic to the macro-triangle (Figure 3). Since the boundary control points of this triangle are determined by the data at the corners, only the center point C needs to be set. Farin sets C to obtain quadratic precision point [2]:

$$C = (T_{20} + T_{02} + T_{21} + T_{12} + T_{10} + T_{01})/4 - (V_0 + V_1 + V_2)/6.$$

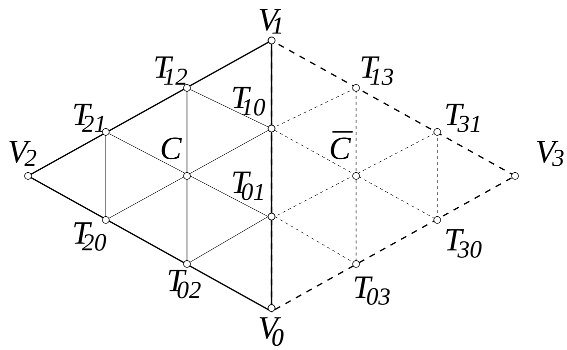

Figure 3: Control points of a cubic. The dashed line segments show a neighboring patch.

This patch will meet its neighbors with only C^0 continuity. To achieve C^1 continuity, Farin subdivides this patch to get initial settings of all the control points of Figure 2, and then adjusts each C_i point to minimize the C^2 discontinuity across the corresponding macro-boundary.

Farin derived the following formulas, for minimizing the error in the two C^2 conditions.[1] First, express $V_2[x,y]$ and $V_3[x,y]$ in barycentric coordinates relative to the neighboring triangle:

$$
\begin{aligned}
V_2[x,y] &= \hat{u}V_3[x,y] + \hat{v}V_1[x,y] + \hat{w}V_0[x,y] \\
V_3[x,y] &= uV_2[x,y] + vV_0[x,y] + wV_1[x,y]
\end{aligned}
$$

Then compute the z-coordinate of C_2 and C_4 as follows (all control points should be interpreted as the z-value of the point):

$$
\begin{aligned}
r_1 &= \hat{u}\bar{I}_{12} + \hat{v}\bar{I}_{11} - uI_{12} - wI_{11} \\
r_2 &= \hat{u}\bar{I}_{02} + \hat{w}\bar{I}_{01} - uI_{02} - vI_{01} \\
r_3 &= vT_{01} + wT_{10} \\
a_{11} &= 2(v^2 + w^2) \\
a_{12} &= -2(v\hat{w} + w\hat{v}) \\
a_{22} &= 2(\hat{w}^2 + \hat{v}^2) \\
s_1 &= 2(vr_1 + wr_2) \\
s_2 &= -2(\hat{w}r_1 + \hat{v}r_2) \\
D &= 2ua_{12} + u^2a_{22} + a_{11} \\
\bar{C}_2 &= (us_1 + ua_{12}r_3 + u^2s_2 + r_3a_{11})/D \\
C_2 &= (\bar{C}_2 - vT_{01} - wT_{10})/u
\end{aligned}
$$

After computing C_0 and C_1 in a similar fashion, Farin then continues the Clough-Tocher construction given above from step (v).

[1] There are minor errors in the formulas in both the Farin paper [2] and the Farin-Kashyap paper [5] which are corrected here.

Although an improvement over the standard Clough-Tocher technique, Farin's scheme also has quadratic precision, since it first constructs the quadratic precision point, then splits, and then minimizes the C^2 discontinuity between the mini-triangles and associated quadratic precision patches (Figure 5(a)).

3.1 Kashyap

Kashyap [9] further analyzed the Clough-Tocher interpolant [9], and considered several variations of the method. Kashyap's paper provides a good survey of Clough-Tocher interpolants, discussing the following methods:

- The C^0 quadratic precision patch that fits a single cubic to each macro-triangle;
- The original C^1 Clough-Tocher interpolant;
- The Farin-Kashyap [5] C^0 interpolant that has cubic precision;
- The Farin [3] C^1 interpolant that attempts to minimize the C^2 discontinuity across macro-triangle boundaries;
- A new C^1 scheme for minimizing the C^2 discontinuities across mini-triangle boundaries; note that this scheme reproduces a subspace of cubic polynomials, but not all cubic polynomials;
- An iterative scheme that repeatedly minimizes the C^2 continuity across macro- and mini-triangle boundaries, using the previously constructed surface as a starting point at each step.

These schemes are all trying to achieve several goals: C^1 continuity, minimization of C^2 discontinuity, and cubic precision. However, none of the above schemes has both C^1 continuity and cubic precision. In the next two sections, I will present two methods for achieving both of these goals.

4 Foley-Opitz

Foley and Opitz developed an alternative cross-boundary field [7]. Their construction builds a *hybrid cubic*, which is a rational blend of three cubic Bézier surfaces, each of which interpolates the corner data. One patch is used to obtain the desired cross-boundary behavior along one of the boundaries. The rational blend they used ensures that the resulting surface will interpolate the data at all three corners, and have the desired cross-boundary behavior along all three boundaries.

The Foley-Opitz cross-boundary construction is similar to that of Farin, except rather than minimize the C^2 discontinuity, they set up a system of equations that by construction will set the center point to have cubic precision. Letting (u, v, w) and $(\hat{u}, \hat{v}, \hat{w})$ be defined as in the previous section, and using the labeling in Figure 3, Foley and Opitz use the following formulas for C and $\bar{C}$ (again, all formulas give z-coordinates):

$$\begin{aligned} C &= (T_{30} - v^2 V_0 - 2vwT_{01} - 2vuT_{02} - w^2 T_{10} - u^2 T_{20} + \\ &\qquad T_{31} - v^2 T_{01} - 2vwT_{10} - 2wuT_{12} - w^2 V_1 - u^2 T_{21})/(2wu + 2vu) \\ \bar{C} &= (T_{20} - \hat{w}^2 V_0 - 2\hat{w}\hat{v}T_{01} - 2\hat{w}\hat{u}T_{03} - \hat{v}^2 T_{10} - \hat{u}^2 T_{30} + \\ &\qquad T_{21} - \hat{w}^2 T_{01} - 2\hat{w}\hat{v}T_{10} - 2\hat{v}\hat{u}T_{13} - \hat{v}^2 V_1 - \hat{u}^2 T_{31})/(2\hat{v}\hat{u} + 2\hat{w}\hat{u}) \end{aligned}$$

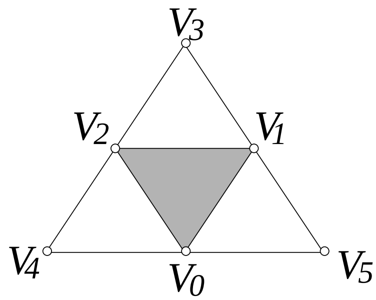

Figure 4: If the data at all six V_i comes from a cubic, then the Foley-Opitz and the modified Farin scheme have cubic precision over the shaded triangle.

Foley and Opitz show that if the four P_i and the four N_i are all sampled from a single cubic function, then the two Bézier patches constructed with their method will reproduce this cubic function.

The Foley-Opitz cross-boundary construction can be integrated into a Clough-Tocher scheme by first constructing a cubic patch using their technique, and then subdividing this patch to obtain one of the cross-boundary points. For efficiency reasons, we will not want to perform a complete subdivision; instead, we can just directly compute the C_i point that results from the subdivision. E.g., from the equations above, compute C_2 as

$$C_2 = (T_{10} + T_{01} + C)/3.$$

Once all three C_i have been computed, we continue the Clough-Tocher construction from step (v).

This technique has cubic precision for the same reason that the Foley-Opitz hybrid patch has cubic precision: If the data comes from a cubic, then by construction, this selection of C will reconstruct that cubic. If the construction of the C across all three macro-triangle boundaries produces the same cubic, then this technique will create the subdivided patches of this cubic. Note, however, that the data at six points of Figure 4 must come from a common cubic for the Foley-Opitz method to have cubic precision over the shaded region.

5 Modified Farin

Farin's scheme minimizes the C^2 discontinuity between the mini-triangles along neighboring macro-triangle boundaries (Figure 5(a)). We can modify Farin's scheme to have cubic precision by first fitting three patches to the *macro*-triangle, each of which minimizes the C^2 discontinuity across one boundary (Figure 5(b)). We then split each patch to get the needed cross-boundary control points for the mini-triangles, and then complete the Clough-Tocher construction.

The modification is simple. The only change is that in Farin's equations, one should use T_{21}, T_{12}, T_{20} and T_{02} instead of I_{12}, I_{11}, I_{02} and I_{01} respectively (with a similar change from $\bar{I}$ to $\bar{T}$).

This modified scheme has cubic precision, since (roughly speaking) if data shown in Figure 4 comes from a cubic, then the original cubic will have no C^2 discontinuity, and thus will be chosen by the minimization process.

This modified-Farin scheme is similar to using Foley-Opitz cross-boundaries, since both construct a quadratically varying cross-boundary field, and make use of the degrees of freedom in the

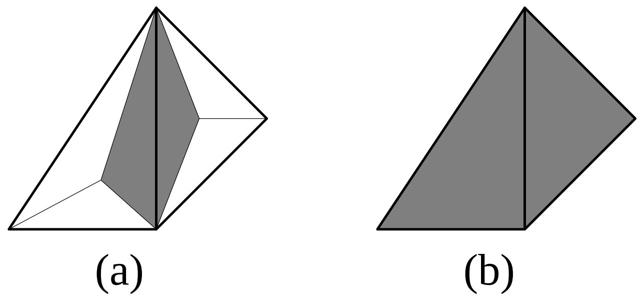

Figure 5: Minimizing the C^2 error; (a) Farin minimizes across mini-triangle boundaries; (b) minimization across macro-triangle boundaries.

Clough-Tocher scheme to obtain cubic precision. Note also that the modified Farin method can be used in the Foley-Opitz hybrid-patch construction.

6 Evaluation

As one test data set, I looked at the following Franke function [6]:

$$\begin{aligned} f_1(x,y) &= \frac{3}{4}e^{-((9x-2)^2+(9y-2)^2)/4} + \frac{3}{4}e^{-(9x+1)^2/49-(9y+1)^2/10} + \\ & \frac{1}{2}e^{-((9x-7)^2+(9y-3)^2)/4} - \frac{1}{5}e^{-(9x-4)^2-(9y-7)^2}. \end{aligned}$$

To visualize the constructed surfaces, I plotted isophotes [8]. If $F(u,v)$ is our surface, $N(u,v)$ is the normal to F, and v is a fixed direction, then the isophote condition is

$$< N(u,v),\ v > \ = c = \text{const.}$$

I.e., an isophote is a curve on the surface where the normal to the surface along the curve forms a fixed angle with a given direction. In Figure 6, I have graphed the projection of the isophotes of the surfaces constructed by each method for a sampling of f_1 over the region $[0,1] \times [0,1]$. The isophotes are plotted for angles of 5, 15, 25, 35, 45, 55, 65, 75, and 85, with v being parallel to the z-axis.

In this figure, I have overlayed the isophote plots with the sampling of f_1 that I used to fit the Clough-Tocher surfaces. Both positions and normals were sampled from f_1. This sampling was taken on a jittered grid (if taken on a uniform grid, the Foley-Opitz method and the modified Farin method produce identical surfaces). Since the goal of this work is to compare the cross-boundary techniques, no surface patches were computed for the outer layer of data (which would require a boundary condition to compute some of their interior points).

What we see in these isophote plots is that all of these Clough-Tocher methods tend to reveal the boundaries of the sampling triangles. While particularly true of the standard Clough-Tocher method, some of the triangle boundaries are apparent with the other three methods.

The following can be readily seen from these plots:

- All three improvements give significantly better results over the standard Clough-Tocher technique. The isophotes do not follow the triangle edges as closely, and are smoother.

- The Foley-Opitz method and the modified Farin method produce better results than the original Farin method. In particular, the isophotes are smoothed in several regions, and two false peaks have disappeared. See, however, the note below.

- The Foley-Opitz method and the modified Farin method produce nearly identical results. Although there are some differences, these are fairly minor.

While the isophotes reveal many of the flaws in a surface, there are other surface quality metrics that should be considered. In particular, from the isophote plots, the Foley-Opitz method and the modified Farin method appear to be always superior to the original Farin method. However, shaded images (Figure 7) reveal while overall the modified Farin's method is better than the the original Farin technique, there are spots on the surface where the original Farin method produces better shaped surfaces.

7 Iterative techniques

It is interesting to note that while the Foley-Opitz method and the modified Farin method produce surfaces that are generally of higher quality than the original Farin method, there are times when the original Farin method does a better job. This results from an oddity in the C^2 minimization methods: While both the original Farin method and the modified Farin method both attempt to reduce the C^2 discontinuity between patches, neither succeeds, since in their minimization equations, both methods use control points that are modified after the minimization is performed (e.g., Farin's method used the I_{i2} to compute the C_i to minimize the C^2 error, and it then uses the C_i to find new settings of the I_{i2} to satisfy the C^1 condition across the internal boundaries), and thus, the C^2 error is not actually minimized.

Farin notes this "non-minimization" in his paper, and suggests an iterative process where, after constructing the entire patch, you use the constructed control points (the I_{i2} of Figure 2) to re-minimize the C^2 equations, and reconstruct the interior points. Farin and Kashyap discuss this idea further, showing that a further reduction in the C^2 discontinuity is possible [5]. Note that we could also start this process by minimizing the C^2 discontinuity across the macro-triangles.

Later, Kashyap investigated minimizing the C^2 discontinuity between mini-triangles within a macro-triangle [9]. He then set up an iterative process, minimizing at each step the C^2 error across macro-triangle boundaries and across mini-triangle boundaries.

However, while this iterative process will reduce the C^2 discontinuity between mini-triangles, we lose the cubic precision of the modified Farin method, at least in a local sense. Cubic precision is retained in the global sense since if the entire data set is sampled from a cubic polynomial, then the modified Farin method will produce patches with no C^2 discontinuity across macro-boundaries (and the iteration terminates after one step). Whether cubic precision or minimization of the C^2 discontinuity is more important will probably depend upon the individual application.

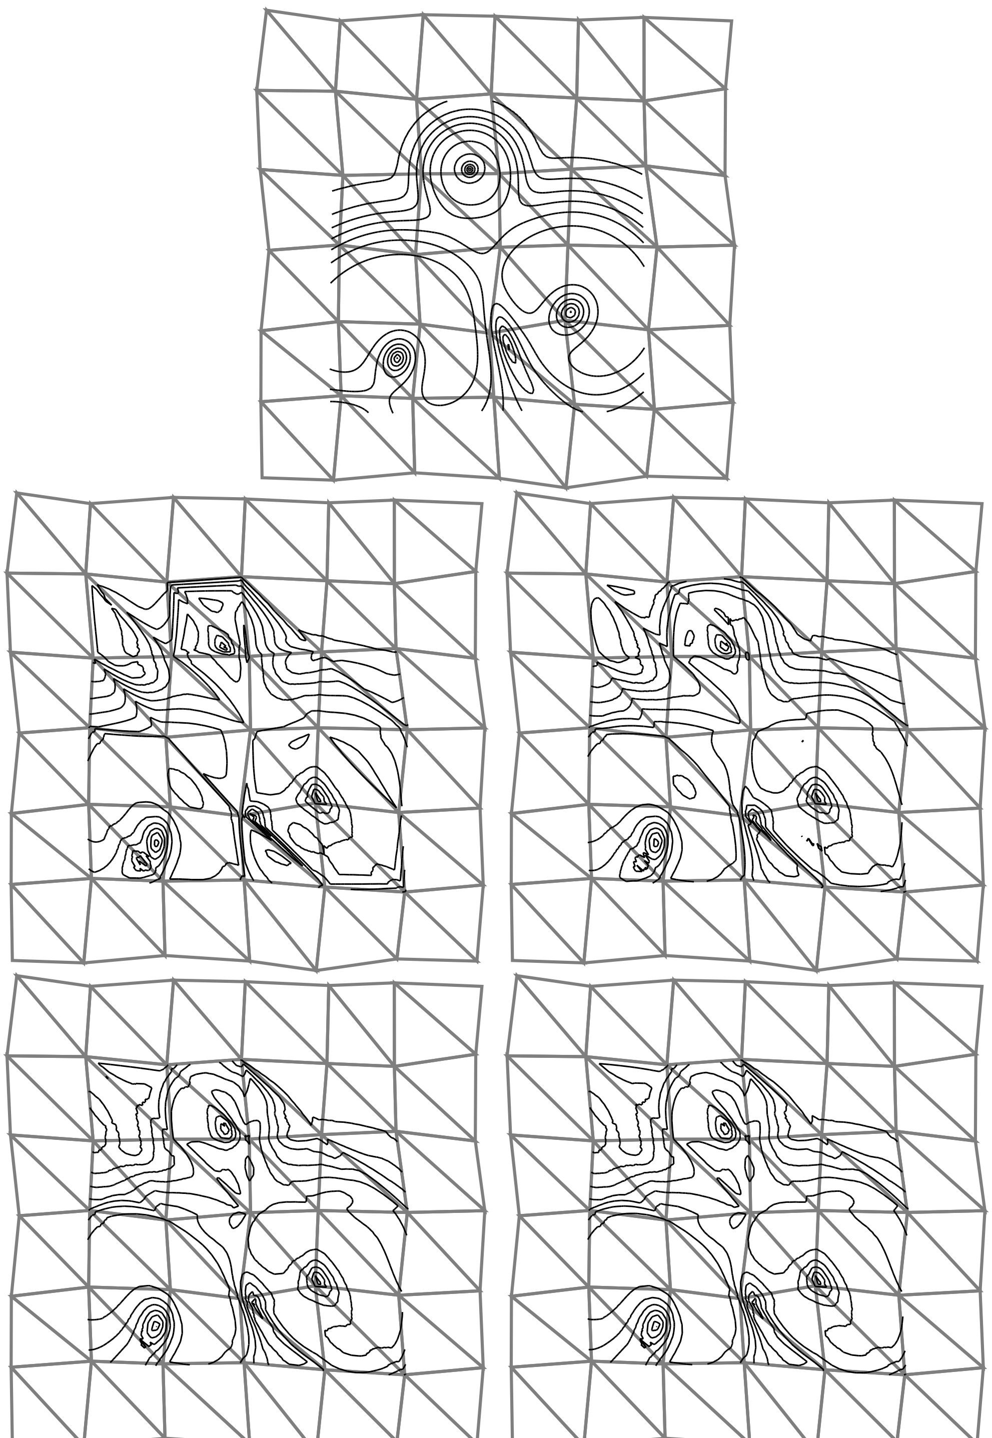

Figure 6: Isophotes. Top: The Franke function. Middle left: Standard Clough-Tocher. Middle right: Farin's method. Bottom left: Foley-Opitz cross-boundaries. Bottom right: Modified Farin method

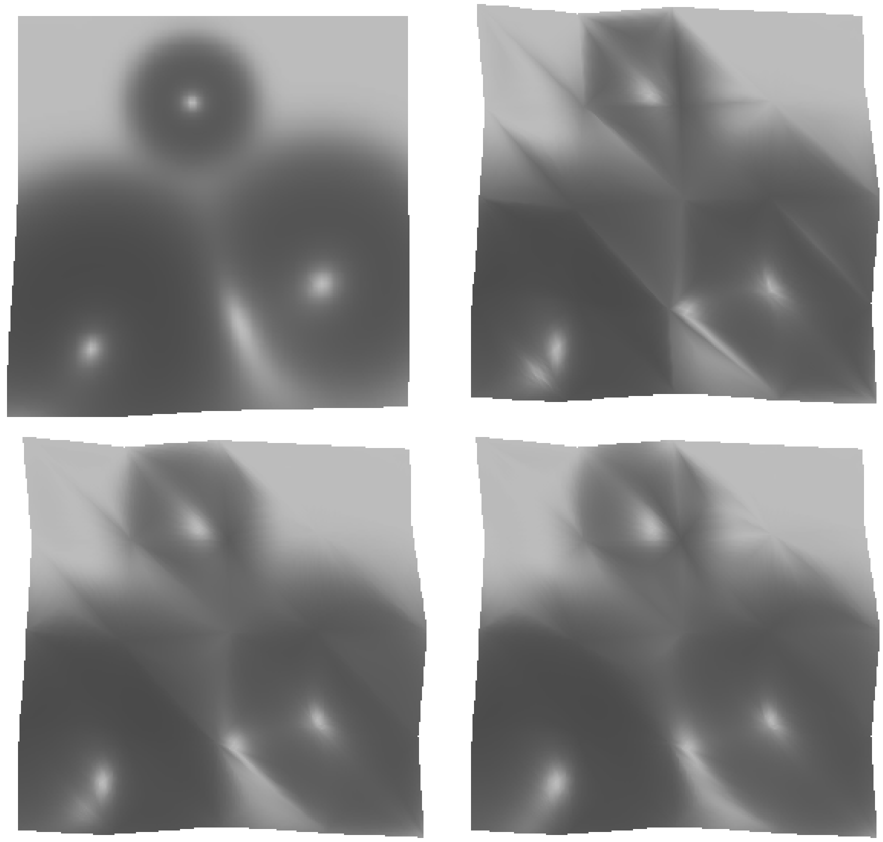

Figure 7: Shaded images. Upper left: Franke Function. Upper right: Standard Clough-Tocher. Lower left: Farin's method. Lower right: modified Farin's method.

8 Other Clough-Tocher Schemes

The Clough-Tocher 3:1 split idea has been used in non-functional settings. Several researchers have generalized the method to parametric data (see the paper of Mann et al. for a survey of such schemes [11]). Note that in the parametric setting, quartic patches must be used (instead of cubic), and these parametric methods suffer from shape defects similar to those in the functional Clough-Tocher scheme. Unfortunately, cubic precision is more difficult to obtain in the parametric setting, and the improvements discussed in this paper are not easy to generalize to the parametric setting.

The idea of using the free parameters to minimize the C^2 discontinuity was also used by Liu and Schumaker to create a Clough-Tocher interpolant over the sphere [10]. Essentially, this is the same technique as the modified Farin's, but it is applied to the spherical setting rather than the functional setting.

References

[1] R. Clough and J. Tocher, Finite element stiffness matrices for analysis of plates in bending. In *Proceedings of Conference on Matrix Methods in Structural Analysis*, 1965.

[2] Gerald Farin, Smooth interpolation to scattered 3D data, in: R.E. Barnhill and W. Böhm, eds, *Surfaces in CAGD*, North-Holland, Amsterdam (1983) pp 43–63.

[3] Gerald Farin, A Modified Clough-Tocher Interpolant, Computer Aided Geometric Design, 2 (1985) pp 19–27.

[4] Gerald Farin, "Curves and Surfaces for CAGD: A Practical Guide, third edition, Academic Press, 1992.

[5] Gerald Farin and Praveen Kashyap, An iterative Clough-Tocher interpolant, Mathematical Modeling and Numerical Analysis 26, pp 201–209.

[6] Richard Franke, A critical comparison of some methods for interpolation of scattered data, Technical report NPS-53-79-003, Naval Postgraduate School, Monterey, California.

[7] Thomas A. Foley and Karsten Opitz, Hybrid cubic Bézier triangle patches, in Mathematical Methods for Computer Aided Geometric Design II, T Lyche and L Schumaker (eds.), Academic Press, 1992, pp 275–286.

[8] Hans Hagen, Stephanie Hahmann, Thomas Schreiber, Yasuo Nakajima, Burkard Wördenweber, and Petra Hollemann-Grundetedt, Surface interrogation algorithms, Computer Graphics and Applications, 12 (1992) pp 53–60.

[9] Praveen Kashyap, Improving Clough-Tocher interpolants, Computer Aided Geometric Design, 13 (1996) pp 629–651.

[10] Xiaoyu Liu and Larry Schumaker, Hybrid Bézier patches on sphere-like surfaces, J. Comp. Appl. Math. 73 (1996), 157–172.

[11] Stephen Mann, Charles Loop, Michael Lounsbery, David Meyers, James Painter, Tony DeRose, and Kenneth Sloan, A survey of parametric scattered data fitting using triangular interpolants, in G. Farin, ed, *Curve and Surface Design*, SIAM (1992) pp 145–172.