Locally equiangular triangulations

R. Sibson

School of Mathematics, University of Bath, Claverton Down, Bath, Avon BA2 7AY

Lawson (1972) has given a criterion of local equiangularity of a triangulation of the convex hull of a finite set of distinct points in the plane. In this note it is shown that (with suitable modifications to deal with degeneracy) there is only one such triangulation, and that it is the Delaunay triangulation (Rogers, 1964).

(Received March 1977)

The initial step in the construction of interpolating functions in two dimensions is the construction of the triangular grid on which the interpolation is based. Usually this grid is not given as part of the data in the problem; all that is available is a finite set of distinct points (*data sites*) at each of which the function and any relevant derivatives are evaluated. Triangulating a set of data sites means selecting a triangulation of the convex hull of the set, the vertices of that triangulation being all and only the data sites. This triangulation may then need further refinement, with the addition of extra vertices, to give the regions on which the interpolating function takes a simple form. This problem of further subdivision, which will not concern us here, is discussed by Powell and Sabin (1977).

A triangulation is regarded as 'good' for the purposes of interpolation if its triangles are nearly equiangular. When the data sites are placed almost as part of a regular triangular lattice, there is little doubt about which triangulation to choose and no difficulty over constructing it. But in practice this even spacing may well fail to hold. With arbitrarily placed points a close approach to equiangularity is seldom possible, and a criterion is needed for assessing the acceptability of a triangulation. Lawson (1972) has suggested such a criterion, which he calls the *max-min angle criterion*. This criterion requires that the diagonal of every convex quadrilateral occurring in the triangulation should be well chosen, in the sense that it should make the two resultant triangles as nearly equiangular as possible. A formal statement of the criterion is as follows:

> If two triangles in the triangulation share a common edge, they define a quadrilateral with that common edge as diagonal. If that quadrilateral is strictly convex (that is, each vertex is an extremal point of it) then replacement of the chosen diagonal by the alternative one must not increase the minimum of the six angles in the two triangles making up the quadrilateral, and this must hold for all such strictly convex quadrilaterals.

We shall call such triangulations *locally equiangular*.

It is the purpose of this note to show that there is only one locally equiangular triangulation of the convex hull of a finite set of distinct data sites, and to identify that triangulation as the Delaunay triangulation, the dual of the Dirichlet/Voronoi/Thiessen tessellation. These constructs are described in detail by Rogers (1964); the diversity of names is a consequence of their independent development in various different applications. The Dirichlet tessellation of a finite set of distinct data sites is obtained by associating with each data site a *tile* consisting of that part of the plane strictly closer to its generating data site than to any other. Clearly the tile of the data site P is the intersection of the open half-planes containing P and bounded by the perpendicular bisectors of lines PQ for all other data sites Q. Thus tiles are open convex polygonal regions. The tiles of data sites lying on the boundary of the convex hull of the set extend to infinity; other tiles are bounded. Only a few perpendicular bisectors are actually effective in delimiting each tile. Two tiles which share a boundary segment are said to be *contiguous*, as are their generating data sites. Normally tiles meet in threes; thus if contiguous data sites are joined by edges, a triangulation of the convex hull results. This is the Delaunay triangulation. Each point where three tiles meet is the circumcentre of the corresponding Delaunay triangle. **Fig. 1**, taken from Green and Sibson (1978), illustrates the constructs. Occasionally a degeneracy may occur: four or more tiles may meet at a common point. The generating data sites of such multiple points must be concyclic, with the multiple point as centre, and the Delaunay triangle becomes a cyclic polygon. It is here and here only that any ambiguity can arise. We shall call the Delaunay construct a *pretriangulation*, and shall speak of triangulations which are *completions*

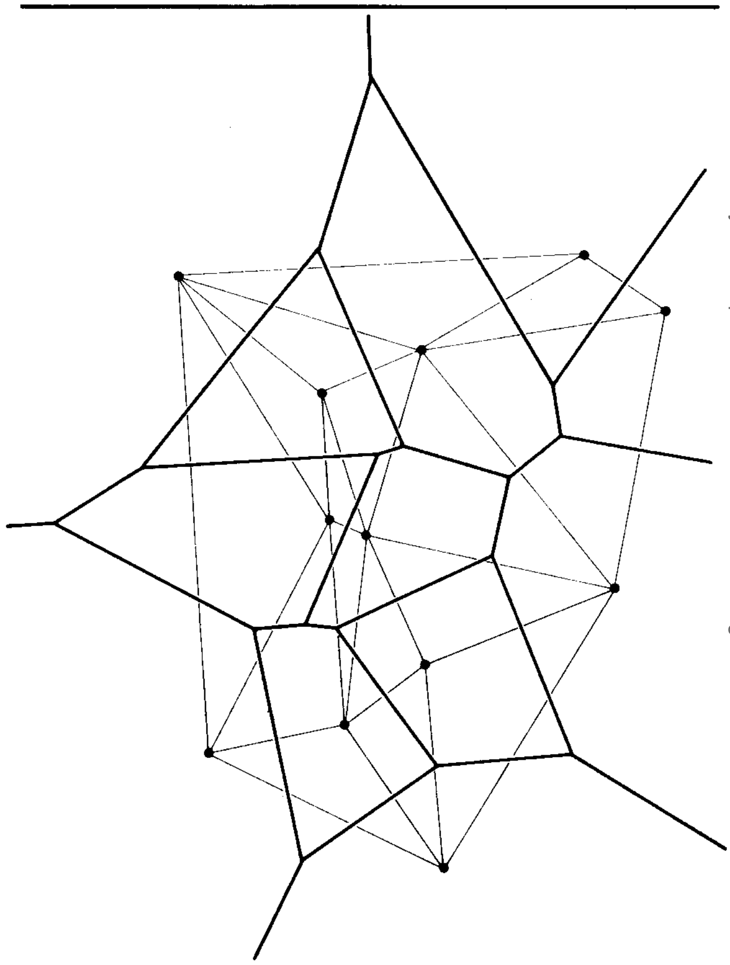

Fig. 1 The Dirichlet tessellation (bold lines) and Delaunay triangulation (fine lines) for a small-scale configuration

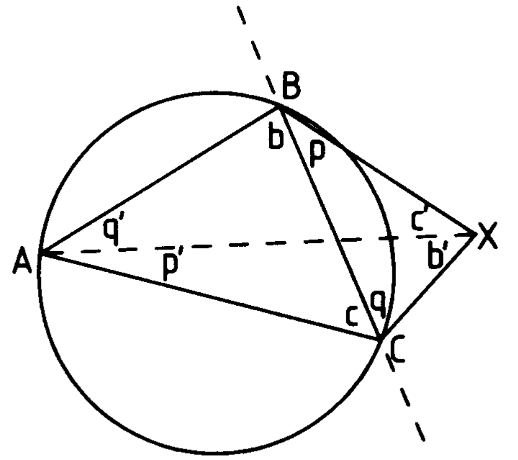

Fig. 2 Reformulation of Lawson's Criterion

of it, meaning that they are obtained by adding to the Delaunay pretriangulation edges which triangulate arbitrarily each cyclic polygon of more than three sides occurring in it. A precise statement of our main result is then as follows.

Theorem
A triangulation of the convex hull of a finite set of distinct data sites is locally equiangular if and only if it is a completion of the Delaunay pretriangulation.

Corollary
If the Delaunay pretriangulation is nondegenerate, and is thus a triangulation, it is uniquely characterised by local equiangularity. This result was asserted by Green and Sibson (1978) and the present note validates that assertion. Green and Sibson describe an algorithm for constructing what is essentially the Delaunay pretriangulation; it is an algorithm which has been implemented and run successfully on a very large scale—in excess of 10,000 data sites—and it is economical in storage and very fast. The present theorem allows it to be regarded also as a means of constructing locally equiangular triangulations for interpolation. It is unlikely that algorithms (such as that suggested by Lawson, 1972 and described below) which work directly in terms of the max-min angle criterion could be competitive with the direct recursive approach of the Green-Sibson algorithm.

The proof of the theorem comprises three lemmas. First, the max-min angle criterion is reformulated; secondly, the completions of the Delaunay pretriangulation are shown to satisfy it; thirdly—and this is the only nontrivial part of the argument—it is shown that no other triangulations do so.

Lemma 1
Let ABC be a triangle, X a point strictly on the opposite side of BC from A. The max-min angle criterion selects BC as the diagonal of the quadrilateral $ABXC$ if and only if X lies strictly outside the circumcircle of ABC; it selects AX as the diagonal if and only if X lies strictly inside the circumcircle; and it allows either BC or AX to be selected if and only if X lies on the circumcircle.

Proof
First suppose X is outside the circumcircle. Then either the quadrilateral is not strictly convex (in which case BC must be chosen as the diagonal although AX may possibly also be selected as an edge in the triangulation), or the situation is as in **Fig. 2.** Elementary geometry gives primed angles as strictly smaller than unprimed ones, paired as shown, and so BC is selected. If X is inside the circumcircle, the quadrilateral is necessarily strictly convex and the strict inequalities are all reversed, so AX is selected. If X is on the circumcircle the primed and unprimed angles are equal in pairs, and either BC or AX can be chosen. These are the only three cases.

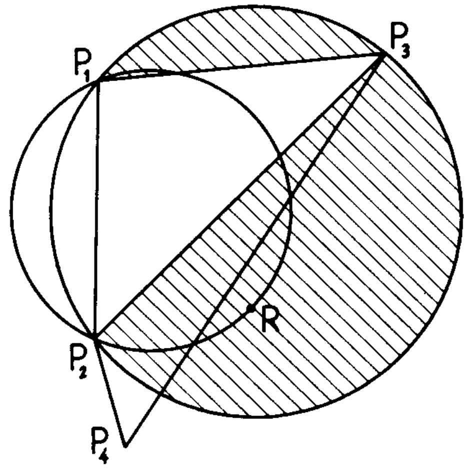

Fig. 3 Proof of uniqueness—the inductive step

Lemma 2
Let S be a circle, whose centre is a point at which three or more tiles of the Dirichlet tessellation meet, and which passes through the generating data sites of those tiles. Then every other data site lies strictly outside S.

Proof
We have already stated the obvious fact that the multiple point is the circumcentre of the generating data sites of the tiles that meet at it. If another data site lay on S, then it too would have a tile meeting at S, and if it lay strictly inside S, then the centre of S would be strictly closer to it than to the data sites on S, which is impossible.

Corollary
Every completion of the Delaunay pretriangulation is locally equiangular.

Lemma 3
Every locally equiangular triangulation of the convex hull of a finite set of distinct data sites is a completion of the Delaunay pretriangulation.

Proof
The edges forming the boundary of the convex hull are common to both triangulations. The proof consists of showing that the only way of 'growing' a triangle on an edge which is known to be common must be one of those permitted as part of a completion of the Delaunay pretriangulation. Let P_1P_2 be an edge of an arbitrary locally equiangular triangulation $\mathscr{T}$, and suppose that P_1P_2 also occurs as an edge in a completion of the Delaunay pretriangulation. Let the *positive* side of P_1P_2 be the side on which it is desired to construct the next triangle. At the boundary of the convex hull this will be the inwards side; otherwise it may be either side. There will be a finite nonempty set of data sites on the positive side of P_1P_2, each one defining a triangle on P_1P_2 and hence a circumcircle whose centre lies on the perpendicular bisector of P_1P_2. Among such circumcircles there will be one whose centre is most negatively placed; the data sites on the positive side of P_1P_2 and on this circle are the candidates for the third vertex of the triangle on the positive side of P_1P_2 in completions of the Delaunay pretriangulation. Suppose that the third vertex P_3 of the corresponding triangle in $\mathscr{T}$ is not one of these; let

R be any one of these data sites. We shall deduce a contradiction.

R must lie strictly outside $\Delta P_1P_2P_3$, because $\mathscr{T}$ is a triangulation; but also it must lie strictly inside the circumcircle $\bigcirc P_1P_2P_3$, because of the extremal property of $\bigcirc P_1P_2R$. Thus it lies in the shaded region in **Fig. 3**; without loss of generality it can be assumed to lie in the part between P_2 and P_3, as shown. Since R and P_1 lie on opposite sides of P_2P_3, there must be a $\mathscr{T}$-triangle on the opposite side of P_2P_3 from P_1, and with P_2P_3 as one of its edges—the position of R prevents P_2P_3 from being a facet of the convex hull. Let P_4 be the third vertex of this triangle. P_4 cannot be R, because R is strictly inside $\bigcirc P_1P_2P_3$ and, by Lemma 1 and the assumption that $\mathscr{T}$ is locally equiangular, P_4 is outside or on $\bigcirc P_1P_2P_3$. R lies strictly outside $\Delta P_2P_3P_4$, because $\mathscr{T}$ is a triangulation. And it lies strictly inside $\bigcirc P_2P_3P_4$, because P_4 is outside or on $\bigcirc P_1P_2P_3$, whence the region strictly to the opposite side of P_2P_3 from P_1, and strictly inside $\bigcirc P_1P_2P_3$ (in which region R lies) is contained in the region strictly inside $\bigcirc P_2P_3P_4$ and also on that side of P_2P_3. A possible position for P_4 is shown in Fig. 3. What we have done is to establish one step of an inductive construction, and this can be repeated indefinitely to obtain P_5,P_6, and so on. The perpendicular distance from R to P_iP_{i+1} is monotone strictly decreasing, so there must be infinitely many distinct points P_i. This is impossible. It follows that P_3 must after all have been one of the data sites on $\bigcirc P_1P_2R$ on the positive side of P_1P_2, and the proof is complete.

These three lemmas together establish the main result, some of the implications of which have already been discussed. We conclude by investigating some other criteria of equiangularity related to local equiangularity. The algorithm which Lawson (1972) suggested for achieving local equiangularity was to start with an arbitrary (but hopefully reasonably good) triangulation and make exchanges of diagonal in convex quadrilaterals according to the max-min angle criterion until no more such exchanges are required. To avoid the trivial possibility of cycling at a degeneracy, it is desirable to make only exchanges *required* by the criterion; ambiguous cases are left undisturbed. As mentioned above, algorithms of this kind are unlikely to be competitive with the Green-Sibson algorithm for computation, but they may, as here, prove useful as theoretical tools. It is not immediately obvious that Lawson's algorithm is guaranteed to terminate. One way of establishing that it does so is to look at the circumradii of the triangles. Each Lawson exchange replaces two triangles by two different ones, and it is easily seen that the new triangles both have strictly smaller circumradii than both of the ones they replace. Thus cycling cannot occur, and since there are only finitely many triangulations, the algorithm must terminate. It does so at a locally equiangular triangulation and hence, by the theorem, at a completion of the Delaunay pretriangulation. This argument also identifies a family of globally defined objective functions optimised by locally equiangular triangulations and by those only: every symmetric strictly isotone function of the circumradii—for example, their sum or the sum of their squares—is such a function. Another criterion, superficially an attractive one, is the maximisation of the minimum angle occurring in the entire triangulation. Such a triangulation might be called *globally equiangular*. If $\mathscr{U}$ is a globally equiangular triangulation, Lawson's algorithm can be applied to it to produce a locally equiangular triangulation $\mathscr{V}$. Now the minimal angle occurring in $\mathscr{V}$ is not greater than that in $\mathscr{U}$, by global equiangularity of $\mathscr{U}$; neither is it less, because Lawson exchanges certainly cannot reduce it. Thus $\mathscr{V}$ is also globally equiangular; we conclude that the completions of the Delaunay pretriangulation are all globally equiangular. Of course, the converse is false; in many cases it will be possible to apply a few anti-Lawson exchanges to a globally equiangular triangulation without reducing the minimal angle occurring in it as a whole.

References

GREEN, P. J. and SIBSON, R. (1978). Computing Dirichlet tessellations in the plane, *The Computer Journal*, Vol. 21, No. 2, pp. 168-173.

LAWSON, C. L. (1972). Generation of a triangular grid with application to contour plotting, *California Institute of Technology Jet Propulsion Laboratory, Technical Memorandum* 299.

POWELL, M. J. D. and SABIN, M. A. (1977). Piecewise quadratic approximations on triangles. *ACM Transactions on Mathematical Software*, Vol. 3, pp. 316-325.

ROGERS, C. A. (1964). *Packing and Covering*, Cambridge Mathematical Tracts 54, Cambridge University Press.

Book reviews

Methods for Statistical Data Analysis of Multivariate Observations, by R. Gnanadesikan, 1977; 311 pages. (*Wiley*, £15·00)

The development of multivariate techniques has been rapid in recent years and the use of computers has changed the emphasis from statistically rigorous procedures of the multiple regression type to data reduction and graphical exploration techniques. This trend is reflected in this new book which begins with Factor Analysis and related techniques and proceeds to discuss dependency methods, classification and clustering before turning to the difficult question of hypothesis testing in multivariate analysis. The final chapter, oddly entitled 'Summarization and Exposure' deals with multi-response samples, including methods using multidimensional residuals and outlier methods.

This is an advanced methodological book suited to graduate statistical students and other research workers involved in multivariate analysis. There are a good number of illustrative examples which use computer graphical facilities. The bibliography is reasonably comprehensive but the references to it appear in curiously condensed form without annotation at the ends of chapters. More within text reference would have helped. The book is an up-to-date and worthwhile addition to the applied multivariate analysis literature and will prove a useful reference book.

ROBERT W. HIORNS (Oxford)

Microprocessors: Infotech State of the Art Report (*Infotech*, £110)

It may be fairly claimed that these two volumes embrace all aspects of the subject they treat. There is material on the underlying LSI technology, on microprocessors design and on microprocessor applications with particular attention to control applications. There is something on software and even a contribution on the content of a university course on the subject.

The presentation follows that of earlier state of tne art reports from Infotech. Volume 2 contains the text of twenty invited papers; Volume 1, which is of similar size, surveys the subject systematically, topic by topic, and consists of quotations from the invited papers and other relevant Infotech publications with linking editorial material. The main topic headings in this volume are: microprocessors—definitions and evaluation; integrated circuit technologies; microprocessor architecture; microprocessor applications; multi-processors and distributed intelligence; evaluation and selection; system design and development; software; manning. There is also an annotated bibliography. In Volume 2 a paper that particularly interested me was one by K. Dixon on the Ferranti F100-L 16 bit microprocessor. I am sure that anyone who is at all concerned with microprocessors will find something to interest him.

M. V. WILKES (Cambridge)