Table cartogram

Table Cartogram[☆]

William Evans

Department of Computer Science, University of British Columbia

Stefan Felsner

Institut für Mathematik, Technische Universität Berlin

Michael Kaufmann

Wilhelm-Schickard-Institut für Informatik, Universität Tübingen

Stephen G. Kobourov

Department of Computer Science, University of Arizona

Debajyoti Mondal

Department of Computer Science, University of Manitoba

Rahnuma Islam Nishat

Department of Computer Science, University of Victoria

Kevin Verbeek

Department of Mathematics and Computer Science, TU Eindhoven

Abstract

A table cartogram of a two dimensional $m \times n$ table A of non-negative weights in a rectangle R, whose area equals the sum of the weights, is a partition of R into convex quadrilateral faces corresponding to the cells of A such that each face has the same adjacency as its corresponding cell and has area equal to the cell's weight. Such a partition acts as a natural way to visualize table data arising in various fields of research. In this paper, we give a $O(mn)$-time algorithm to find a table cartogram in a rectangle. We then generalize our algorithm to obtain table cartograms inside arbitrary convex quadrangles, circles, and finally, on the surface of cylinders and spheres.

[☆]A preliminary version of the paper appeared in the Proceedings of the 21st Annual European Symposium on Algorithms (ESA 2013) [1].

Email addresses: will@cs.ubc.ca (William Evans), felsner@math.tu-berlin.de (Stefan Felsner), mk@informatik.uni-tuebingen.de (Michael Kaufmann), kobourov@cs.arizona.edu (Stephen G. Kobourov), jyoti@cs.umanitoba.ca (Debajyoti Mondal), rnishat@uvic.ca (Rahnuma Islam Nishat), k.a.b.verbeek@tue.nl (Kevin Verbeek)

Keywords: Cartogram, Data Visualization, Grid Map, Tree Map.

1. Introduction

A *cartogram*, or *value-by-area diagram*, is a thematic cartographic visualization, in which the areas of countries are modified in order to represent a given set of values, such as population, gross-domestic product, or other geo-referenced statistical data. Red-and-blue population cartograms of the United States were often used to illustrate the results in the 2000 and 2004 presidential elections. While geographically accurate maps seemed to show an overwhelming victory for George W. Bush, population cartograms effectively communicated the near 50-50 split, by deflating the rural and suburban central states.

The challenge in creating a good cartogram is thus to shrink or grow the regions in a map so that they faithfully reflect the set of pre-specified area values, while still retaining their characteristic shapes, relative positions, and adjacencies as much as possible. In this paper we introduce a new *table cartogram* model, where the input is a two dimensional $m \times n$ table of non-negative weights, and the output is a rectangle with area equal to the sum of the input weights partitioned into $m \times n$ convex quadrilateral faces each with area equal to the corresponding input weight. Figure 1 and Figure 2 show two such examples. Such a visualization preserves both area and adjacencies. Besides, it is simple, visually attractive, and applicable to many fields that require visualization of data table.

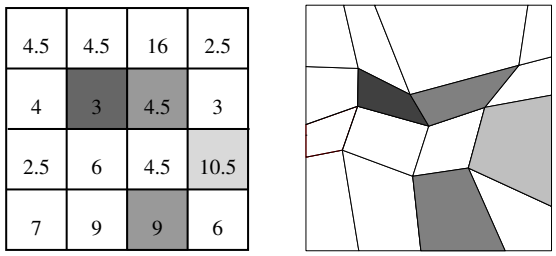

Figure 1: A 4×4 table, its table cartogram.

The solution to the problem is not obvious even for a 2×2 table. For example, Figures 3(a) and (b) show a table A and a unit square R, respectively. One attempt to find the cartogram of A in R may be to first split R horizontally according to the sum of each row, and then to find a good split in each subrectangle to realize the correct areas. But this approach does not work, because the first split prevents the creation of the two convex quadrilaterals with area ϵ in opposite corners that share a boundary vertex, see Figure 3(c). Figure 3(d) shows a possible cartogram.

The following little argument shows that 2×2 table cartograms exist. The argument contains some elements that will be reused for the general case. The

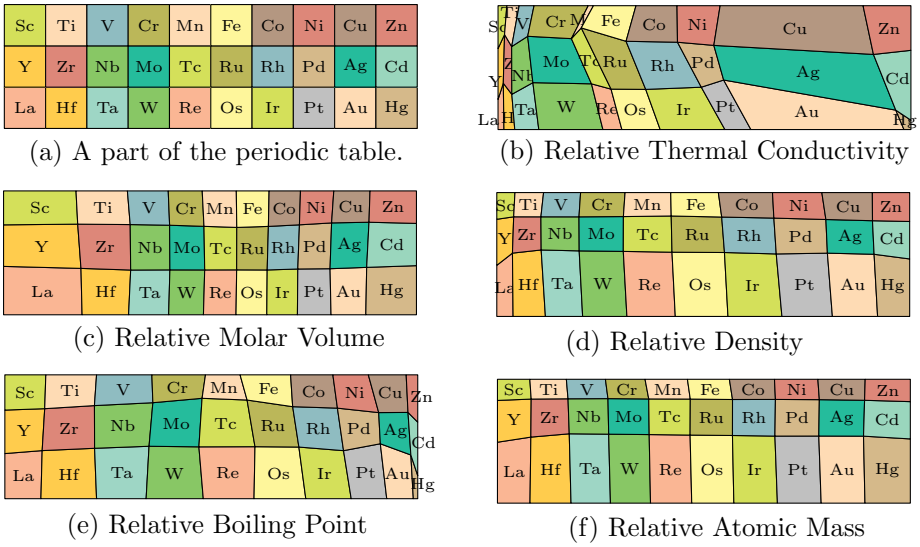

(a) A part of the periodic table. (b) Relative Thermal Conductivity

(c) Relative Molar Volume (d) Relative Density

(e) Relative Boiling Point (f) Relative Atomic Mass

Figure 2: (a) A part of the periodic table of chemical elements. (b)–(f) Cartograms of different chemical properties.

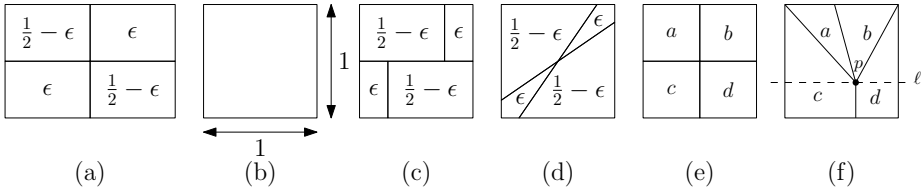

Figure 3: (a) A 2×2 table A. (b) R. (c) An attempt to find a cartogram. (d) A cartogram of A in R. (e) A 2×2 table A. (f) The cartogram showing ℓ as a dashed line.

input is a 2×2 table with four positive reals a, b, c, d with $a+b+c+d = 1$, as shown in Figure 3(e). Rotational symmetry of the problem allows us to assume that $a+b \leq 1/2$. Fix the unit square R with corners $(0,0), (0,1), (1,1), (1,0)$ as the frame for the table cartogram.

Now consider the horizontal line ℓ with the property that every triangle $T(p)$ with top side equal to the top side of R and one corner p on ℓ has area $a+b$. Since $a+b \leq 1/2$, the line ℓ intersects R in a horizontal segment. For $p \in \ell \cap R$, the vertical line through p partitions $R \setminus T(p)$ into a left 4-gon S^- and a right 4-gon S^+. The areas of these two 4-gons depend continuously on the position of point p but their sum is always $c+d$. If p is on the left boundary, $Area\,(S^+) = c+d$, and if p is on the right boundary, $Area\,(S^+) = 0$. Hence it follows from the intermediate value theorem that there is a position for p on $\ell \cap R$ such that $Area\,(S^-) = c$ and $Area\,(S^+) = d$. We can now use the intermediate value theorem to find a partition of $T(p)$, i.e., we find a line by rotating a line around p such that the left triangle has area a, and the right triangle has area

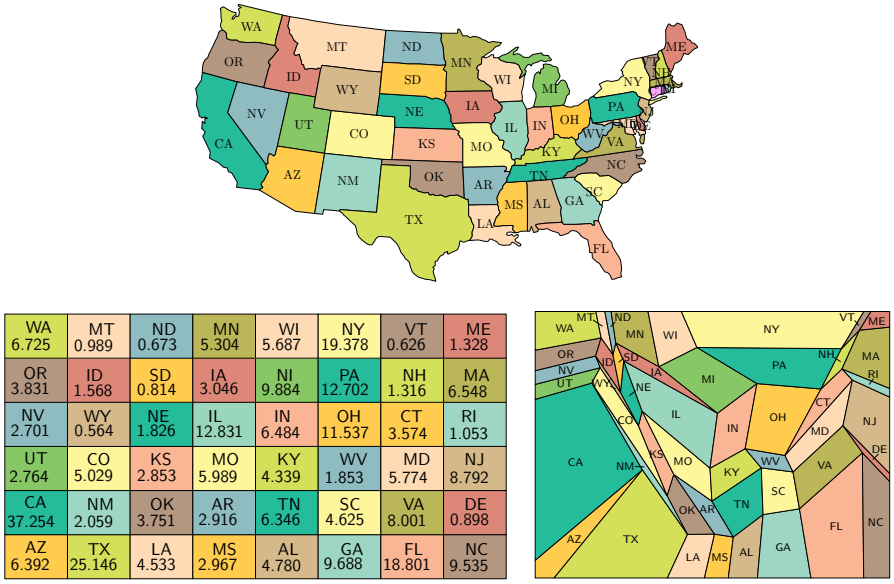

Figure 4: A grid map of the USA [6] and our corresponding table cartogram of the population of the states in 2010.

b. The resulting partition of R into four parts is a table cartogram for the input table, as shown in Figure 3(f). The critical reader may object that two of the 4-gons have a degenerate side. This can be avoided by perturbing the cartogram slightly to make a very short edge instead of a point. The result is an ε approximate cartogram without degeneracies. Another approach is to modify the construction rules so that degeneracies are avoided. We take this approach in Section 2 to show the existence of non-degenerate table cartograms in general.

1.1. Related Work

The problem of representing additional information on top of a geographic map dates back to the 19th century, and highly schematized rectangular cartograms can be found in the 1934 work of Raisz [2]. Recently, van Kreveld and Speckmann describe automated methods to produce rectangular cartograms [3]. With such rectangular cartograms it is not always possible to represent all adjacencies and areas accurately [4, 3]. However, in many "simple" cases, such as France, Italy and the USA, rectangular cartograms and even table cartograms offer a practical and straightforward schematization, e.g., Figure 4. *Grid maps* are a special case of single-level spatial treemaps: the input is a geographic map mapped onto a grid of equal-sized rectangles, in such a way as to preserve as well as possible the relative positions of the corresponding regions [5, 6]. As we show, such maps can always be visualized as table cartograms.

Eppstein *et al.* studied area-universal rectangular layouts and characterized

the class of rectangular layouts for which all area-assignments can be achieved with combinatorially equivalent layouts [7]. If the requirement that rectangles are used is relaxed to allow the use of rectilinear regions then de Berg *et al.* [8] showed that all adjacencies can be preserved and all areas can be realized with 40-sided regions. In a series of papers the polygon complexity that is sufficient to realize any rectilinear cartogram was decreased from 40 sides down to 8 sides [9], which is best possible due to an earlier lower bound [10].

More general cartograms without restrictions to rectangular or rectilinear shapes have also been studied. For example adjacencies can be preserved and areas represented perfectly using convex quadrilaterals if the dual of the map is an outerplanar graph [11]. Dougenik *et al.* introduced a method based on force fields where the map is divided into cells and every cell has a force related to its data value which affects the other cells [12]. Dorling used a cellular automaton approach, where regions exchange cells until an equilibrium has been achieved, i.e., each region has attained the desired number of cells [13]. This technique can result in significant distortions, thereby reducing readability and recognizability. Keim *et al.* defined a distance between the original map and the cartogram with a metric based on Fourier transforms, and then used a scan-line algorithm to reposition the edges so as to optimize the metric [14]. Gastner and Newman [15] project the original map onto a distorted grid, calculated so that cell areas match the pre-defined values. The desired areas are then achieved via an iterative diffusion process inspired by physical intuition. The cartograms produced this way are mostly readable but the complexity of the polygons can increase significantly. Edelsbrunner and Waupotitsch [16] generated cartograms using a sequence of homeomorphic deformations. Kocmoud and House [17] described a technique that combines the cell-based approach of Dorling [13] with the homeomorphic deformations of Edelsbrunner and Waupotitsch [16].

There are thousands of papers, spanning over a century, and covering various aspects of cartograms, from geography to geometry and from interactive visualization to graph theory and topology. The above brief review is woefully incomplete; the survey by Tobler [18] provides a more comprehensive overview.

1.2. Our Results

The main construction is presented in Section 2. We start with a simple constructive algorithm that realizes any table inside a rectangle in which each cell is represented by a convex quadrilateral with its prescribed weight. The approach relies on making many of the regions be triangles. We then modify the method to remove such degeneracies. The construction can be implemented to run in $O(mn)$ time, i.e., in time linear in the input size.

In Section 3 we find table cartograms inside arbitrary triangles or convex quadrilaterals, which is best possible, because regular n-gons, $n \geq 5$, do not always support table cartograms (e.g., consider a table with some cell value larger than the maximum-area convex quadrangle that can be drawn inside the n-gon). We also realize table cartograms inside circles, using circular-arcs, and on the surface of a sphere via a transformation from a realization on the cylinder.

2. Table cartograms in rectangles

We first construct a cartogram with degenerate 4-gons, where we allow the input table to have non-negative numbers. Later we show that one can remove the degeneracies if the entries of the table are strictly positive.

2.1. Cartogram with Degenerate 4-gons

The input is a table A with m rows and n columns of non-negative numbers $A_{i,j}$. Let $S = \sum_{i,j} A_{i,j}$ and let S_i be the sum of the numbers in row i, i.e., $S_i = \sum_{1 \le j \le n} A_{i,j}$. Assume, by scaling, that $S > 4$. Let R be the rectangle with corners $(0,0)$, $(S/2,0)$, $(S/2,2)$, $(0,2)$. We construct the cartogram within R and later generalize the construction to all rectangles with area S, as stated in the following theorem.

Theorem 1. *Let A be an $m \times n$ table of non-negative numbers $A_{i,j}$. Let R be a rectangle with width w, height h and area equal to the sum of the numbers of A. Then there exists a cartogram of A in R such that every face in the cartogram is convex. The construction requires $O(mn)$ arithmetic operations.*

In the rest of this section we prove Theorem 1. Let k be the largest index such that the sum of the numbers in rows $1, 2, \ldots, k-1$ is less than $S/2$. We may then choose $\lambda \in (0,1]$ such that $\sum_{1 \le i \le k-1} S_i + \lambda S_k = S/2$. We split the table A into two tables A^t and A^b. Table A^t consists of k rows and n columns. The first $k-1$ rows are taken from A, i.e., $A^t_{i,j} = A_{i,j}$ for $1 \le i \le k-1$ and $1 \le j \le n$. The last row is a λ-fraction of row k from A, i.e., $A^t_{k,j} = \lambda \cdot A_{k,j}$ for all j. Table A^b consists of $m-k+1$ rows and n columns. The first row accommodates the remaining portion of row k from A, i.e., $A^b_{1,j} = (1-\lambda) \cdot A_{k,j}$. All the other rows are taken from A, i.e., $A^b_{i,j} = A_{i+k-1,j}$ for $i > 1$ and all j. An example is shown in Figure 5(a). If $\lambda = 1$, then A^b contains a top row of zeros.

Let D^t_j be the sum of entries in columns $2j-2$ and $2j-1$ from A^t, where $1 \le j \le \lfloor m/2 + 1 \rfloor$. Note that D^t_1 is only responsible for one column. The same may hold for the last D^t_j depending on the parity of m. Similarly, D^b_l is the sum of entries in columns $2l-1$ and $2l$ from A^b, where $1 \le l \le \lceil m/2 \rceil$. Again, depending on the parity of m the last D^b_l may only be responsible for one column.

We now define a zig-zag Z in R (formally, Z is a polygonal line) such that the areas of the triangles defined by Z are the numbers $D^t_1, D^b_1, D^t_2, D^b_2, D^t_3, \ldots$ in this order. The zig-zag starts at $z_0 = (0,0)$. Since the height of R is 2, the first segment ends at $z_1 = (D^t_1, 2)$ and the second segment goes down to $z_2 = (D^b_1, 0)$. In general, for i odd, $z_i = (\sum_{j=1}^{\lceil i/2 \rceil} D^t_j, 2)$ and for i even, $z_i = (\sum_{l=1}^{i/2} D^b_l, 0)$. An important property of Z is that it ends at one of the two corners on the right side of R. This is because $\sum_j D^t_j = S/2 = \sum_l D^b_l$.

Lemma 2 shows that we can partition each triangle created by the zig-zag Z into triangles whose areas are the corresponding entries in A^t or A^b. It relies

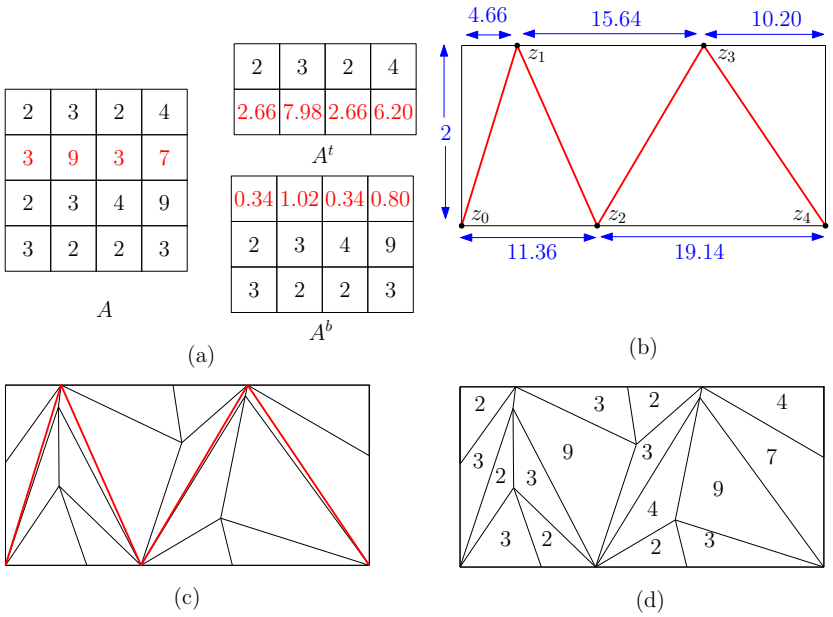

Figure 5: (a) Illustration for A, A^t and A^b, where $k = 2$ and $\lambda \approx 0.886$. (b) The zigzag path Z. We have distorted the aspect ratio of the figure to increase readability. (c) The subdivision of triangles, where Z is shown in red, and (d) the complete cartogram.

on the following lemma which is a consequence of properties of barycentric coordinates. However, we include a proof here because later variants of the lemma would require generalizing this proof.

Lemma 1 (Triangle Lemma). *Let $\triangle abc$ be a triangle and let α, β, γ be non-negative numbers, where $\alpha + \beta + \gamma = Area(\triangle abc)$. Then we can find a point p in $\triangle abc$, where $Area(\triangle pbc) = \alpha$, $Area(\triangle apc) = \beta$, $Area(\triangle abp) = \gamma$, in $O(1)$ arithmetic operations.*

Proof. Let ℓ_a be the line such that a triangle with side bc and a corner on ℓ_a has area α. This line intersects segments ab and ac. Let ℓ_b be the line such that a triangle with side ac and a corner on ℓ_b has area β. This line intersects segments ab and bc. Let q_a be the intersection point of segment ab with line ℓ_a and let q_b be the intersection point of segment ab with line ℓ_b. Assume that the order of points on segment ab is a, q_a, q_b, b. Now the triangles $\triangle q_a bc$ and $\triangle a q_b c$ cover the triangle $\triangle abc$ so that $Area(\triangle abc) < \alpha + \beta$. This contradiction shows that the order of points on segment ab is a, q_b, q_a, b. Hence ℓ_a and ℓ_b intersect in a point p inside of $\triangle abc$. This point p has the desired properties.

Note that p can be computed with $O(1)$ arithmetic operations. □

Lemma 2. *Let A be an $m \times 2$ table such that each cell is assigned a non-negative number. Let $\triangle abc$ be a triangle such that the area of $\triangle abc$ is equal to the sum of the numbers of A. Then A admits a cartogram inside $\triangle abc$ such that all*

cells of A are represented by triangles and the boundary between those triangles representing cells in the left column and those representing cells in the right column is a polygonal path connecting point a to some point on the segment bc.

PROOF. The proof is by induction on m. The case $m = 1$ is obvious, recall the example illustrated in Figure 3. If $m > 1$ we define $\alpha = \sum_{1 \le i \le m-1} A_{i,1} + A_{i,2}$, $\beta = A_{m,1}$ and $\gamma = A_{m,2}$. Using Lemma 1 we find a point p in $\triangle abc$ that partitions the triangle into triangles of areas α, β and γ. We keep the triangles $\triangle apc$ and $\triangle abp$ as representatives for $A_{m,1}$ and $A_{m,2}$ and construct the cartogram for the first $m-1$ rows of A in the triangle $\triangle pbc$ by induction. □

To partition triangle $\triangle z_{2j-2}, z_{2j-3}, z_{2j-1}$, for $1 \le j \le \lfloor m/2 + 1 \rfloor$ (where $z_{-1} = (0, 2)$ and $z_{m+1} = (S/2, 2)$ if needed), we appeal to Lemma 2 with A (in the lemma) being the two columns from A^t whose sum is D^t_j. To make Lemma 2 applicable to cases like D^t_1 which represents only one column from A^t, we simply add a column of zeros to A. Similarly, we can partition triangle $\triangle z_{2l-1}, z_{2l-2}, z_{2l}$, for $1 \le l \le \lceil m/2 \rceil$.

This yields a table cartogram of the $(m+1) \times n$ table A^+ that is obtained by stacking A^t on A^b. Note, however, that all triangles representing cells from the last row of A^t have a side that equals one of the edges of Z. Symmetrically, all triangles representing cells from the first row of A^b have a side on Z. Hence by removing the edge of Z we glue two triangles of area $\lambda A_{k,j}$ and $(1-\lambda)A_{k,j}$ into a 4-gon of area $A_{k,j}$. The 4-gons obtained by removing edges of Z are convex because they have crossing diagonals. This completes the construction.

To complete the proof of Theorem 1, in which R is a $w \times h$ rectangle with area S, we scale the above cartogram by a factor of $h/2$ vertically and a factor of $2/h$ horizontally.

2.2. Removing Degeneracies

The construction of the proof of Theorem 1 creates faces of degenerate shape, i.e., some faces may not be perfect quadrangles, as shown in Figure 5(d). We modify this construction to avoid the degeneracies. Of course we have to make a stronger assumption on the input: All entries $A_{i,j}$ of the table are strictly positive. The first part of the construction remains unaltered.

- Determine k and λ such that $\sum_{1 \le i \le k-1} S_i + \lambda S_k = S/2$.
- Define A^t and A^b and the two-column sums D^t_j and D^b_l for these tables.
- Compute the zig-zag in the rectangle R of height 2 and width $S/2$.

Let $z_0, z_1, \ldots, z_n$ be the corner points of the zig-zag Z. For i even we define $z'_i = z_i + (0, v)$ and for i odd $z'_i = z_i - (0, v)$, i.e., z'_i is obtained by shifting z_i vertically a distance of v into R. We will choose this positive value v to obey inequalities (B1) and (B2) required by the construction, which will be specified later (after Figure 6). However, one can observe that by choosing a sufficiently

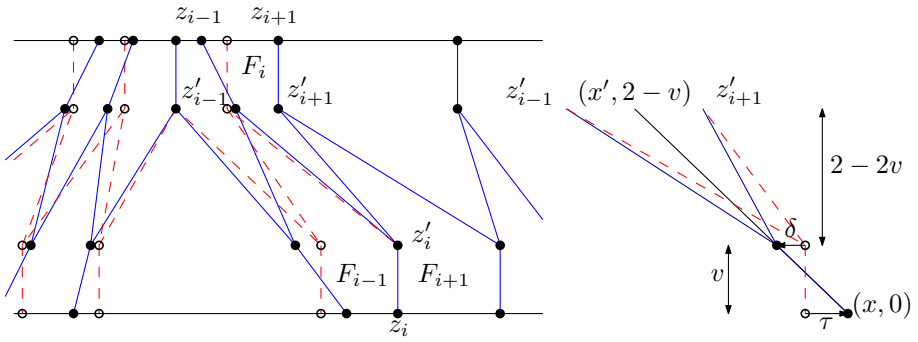

Figure 6: Before and after convexifying z'_i. The dashed lines represent the original Z'.

small value for v will cause both of those inequalities to be satisfied. Let Z' be the zig-zag with corners $z'_0, z'_1, \ldots, z'_n$. The segment z'_i, z_i is the *leg at* z'_i. The union of all the legs and Z' is the *skeleton* G' of a partition of R into 5-gons. We refer to the 5-gons with corners $z_{i-1}, z_{i+1}, z'_{i+1}, z'_i, z'_{i-1}$ as F_i. We refrain from introducing extra notation for the two 4-gons at the ends of Z' and just think of them as degenerate 5-gons.

Lemma 3. *A 5-gon in R with vertices $(x_1, 0), (x_3, 0), (x_3, v), (x_2, 2-v), (x_1, v)$ has the same area $x_3 - x_1$ as the triangle with corners $(x_1, 0), (x_3, 0), (x_2, 2)$.*

Proof. First note that changing the value of x_2 (shear) preserves the area of the 5-gon and of the triangle. Hence we may assume that $x_2 = x_3$. Now let P be the parallelogram with corners $(x_1, 0), (x_1, v), (x_2, 2), (x_2, 2-v)$. Both, the 5-gon and the triangle can be partitioned into the triangle $(x_1, 0), (x_2, 2-v), (x_3, 0)$ and a triangle that makes a half of P. □

Some of the 5-gons F_i may not be convex. However, concave corners can only be at z'_{i+1} or z'_{i-1}. To get rid of concave corners we deal with corners at $z'_1, z'_2, \ldots, z'_{n-1}$ in this order. At each z'_i we may slightly shift z'_i horizontally and bend the leg to rebalance the areas. This can be done so that the concave corner at z'_i is resolved. We then say that z'_i has been *convexified*. Figure 6 shows an example of the process.

The vertex z'_i has a concave corner in at most one of F_{i-1} and F_{i+1}. In the first case we move z'_i to the right, in the second case, we move z'_i to the left. By symmetry, we only detail the second case, i.e, z'_i has a concave corner in F_{i+1}.

Shifting z'_i horizontally keeps the area of F_i invariant. Only the areas of F_{i-1} and F_{i+1} are affected by the shift. By shifting z'_i a distance of δ to the left while keeping z_i at its place the increase in area of F_{i+1} is $\delta(2-v)/2$. To balance the increase we move z_i, the other end of the leg, to the right by an amount τ, where $\tau v/2 = \delta(2-v)/2$. To make sure that the corners at z'_i after shifting are convex we choose δ and τ so that the line connecting the new positions of z_i and z'_i contains the midpoint of z'_{i-1} and z'_{i+1}. If the new position of z_i is $(x, 0)$ and $(x', 2-v)$ is the midpoint of z'_{i-1} and z'_{i+1} then $v/(\tau + \delta) = 2/(x - x')$.

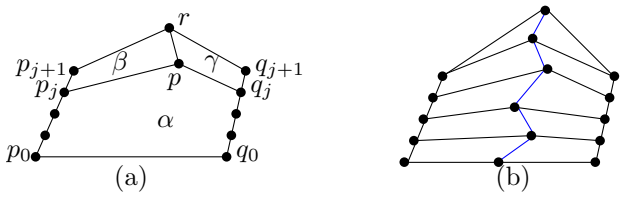

Figure 7: (a) The α, β and γ partition of F. (b) A final partition of F_i.

We do not want the shift of z_i to introduce a crossing. We ensure this with a bound on v. For all j, let $T_j = Area(F_j)$. Since the height of the strip is 2, one can infer from Lemma 3 that T_j is also the distance between z_{j-1} and z_{j+1} before the shift. If $\tau \leq T_{i+1}$, then the leg z'_i, z_i does not intersect leg z'_{i+2}, z_{i+2}. The absolute value of the slope of the leg z'_i, z_i after convexifying z'_i is less than v/τ. The slope of the leg is also between the slopes of z'_i, z'_{i-1} and z'_i, z'_{i+1}. The absolute value of these slopes is larger than $(2-2v)/(S/2)$ which is the minimum possible slope of a segment of Z' in R. Define $T = \min_j T_j$. Hence if $v/T < (2-2v)/(S/2) = 4(1-v)/S$, then $\tau < T$. We thus have an inequality that we want to be true for v:

$$v \leq \frac{4T}{S+4T}. \tag{B1}$$

Observe that convexifying z'_{i+1} may require a shift of z'_{i+1} by δ' (and a compensating shift of z_{i+1} by τ') after z'_i has been convexified. However, if $v \leq 1/4$, then balancing area and (B1) imply $\frac{1}{4}T > v\tau' = \delta'(2-v) \geq \delta'\frac{7}{4}$ whence $\delta' \leq T/7$. This shows that z'_{i+1} stays on the right side of the old midpoint of z'_{i-1} and z'_{i+1} so that the corners at z'_i stay convex.

The next step of the construction is to place equidistant points on each of the legs. The segments between two consecutive points on the leg z'_i, z_i will serve as sides for quadrangles of the quadrangular subdivision of F_{i-1} and F_{i+1}. Specifically, a leg z'_i, z_i with i odd is subdivided into $k-1$ segments of equal length and a leg z'_i, z_i with i even is subdivided into $m-k$ segments. Recall that k is the number of rows in A^t. For the partition of F_i into 4-gons with the prescribed areas we proceed inductively as in Lemma 2. We again need a partition lemma, as follows.

Lemma 4. *Consider a convex 5-gon F as shown in Figure 7(a). Let α, β, γ be positive numbers with $\alpha+\beta+\gamma = Area(F)$. If $\alpha > Area(\square p_0, q_0, q_j, p_j)$, $\beta > Area(\triangle p_j, r, p_{j+1})$, $\gamma > Area(\triangle q_j, q_{j+1}, r)$, then there exists $p \in F$ such that $\alpha = Area(\hexagon p_0, q_0, q_j, p, p_j)$, $\beta = Area(\square p_j, p, r, p_{j+1}), \gamma = Area(\square q_j, q_{j+1}, r, p)$.*

Proof. The assumptions imply that if p exists it has to be in the interior of the triangle $\triangle p_j, r, q_j$. The existence follows from Lemma 1. □

To ensure that the conditions for Lemma 4 are satisfied throughout the inductive partition of the regions F_i we need to bound v. Let $M = \min_{i,j} A_{i,j}$

be the minimum value in the table. Recall that $S/2$ is the width of R, and the y-distance of p_j and p_{j+1} is at most v. Hence $vS/2$ is a generous upper bound on $Area(\triangle p_j, r, p_{j+1})$, $Area(\triangle q_j, q_{j+1}, r)$, and $Area(\square p_0, q_0, q_j, p_j)$. We ensure that these areas are less than β, γ, and α respectively by requiring

$$v < \frac{2M}{S} \tag{B2}$$

Theorem 2. *Let A be an $m \times n$ table of non-negative numbers $A_{i,j}$. Let R be a rectangle of width w and height h such that $w \cdot h = \sum_{i,j} A_{i,j}$. Then there exists a non-degenerate cartogram of A in R such that every face in the cartogram is convex. The construction requires $O(mn)$ arithmetic operations.*

Proof. The steps of the construction are:

- Construct the table cartogram with degeneracies.
- Compute the bounds and fix an appropriate value for v, compute the skeleton G' and its regions F_i, and convexify the legs in order of increasing index.
- Subdivide each of the regions F_i into convex 4-gons (and two triangles).
- Remove the edges of the zig-zag to get the cells of the middle row as unions of two triangles, which must generate convex quadrangles, since these triangles are contained in rectangle R and the common side of a pair of triangles connects opposite sides of R.

All can be done with $O(mn)$ arithmetic operations. Regarding the degeneracies, however, there is an issue that remains. To break A into A^t and A^b, we split row k so that the last row of A^t is a λ-fraction of row k from A while the rest of this row becomes the first row of A^b. Degeneracies occur if $\lambda = 1$. However, rather than splitting row k in this case, we can treat cells of row k as generic cells and assign a section of a leg to each of them. The construction is almost as before. Two details have to be changed. The first partition of each F_i into three pieces now produces two 4-gons and a 5-gon, while previously we had two triangles and a 5-gon in this step, as shown in Figure 7(b). The other change is that we don't remove zig-zag edges belonging to Z' to merge triangles to 4-gons at the end of the construction. □

Instead of just knowing that there are no degeneracies, it would be nice to have a lower bound on the *feature size*, that is the minimum side-length of a 4-gon in the table cartogram. The segments subdividing the legs have length at least v/m. Because these leg segments have length at most v and $vS/2 < M$ (by (B2)), the opposite edges in a generic 4-gon (the blue edges in Figure 7(b)) have length at least v. However, the triangles whose composition creates the 4-gons representing cells of row k can have area smaller than M. These triangles may have area $\hat{\lambda}M$ where $\hat{\lambda} = \min\{\lambda, 1-\lambda\}$. This may lead to a very small feature size. To improve on this, another degree of freedom in the construction can be

used. Instead of breaking each cell $A_{k,j}$ into a λ and a $1-\lambda$ fraction, we can use individual values λ_j to define $A^t_{k,j} = \lambda_j A_{k,j}$. The choice of the values λ_j must satisfy two conditions: (1) $\sum_j \lambda_j A_{k,j} = \lambda S_k = \lambda \sum_j A_{k,j}$ and (2) if $\lambda_i = 0$ and $\lambda_j = 1$ then $|i-j| > 1$. By choosing most of the λ_j to be 0 or 1, and avoiding degeneracies, we may be able to have a substantial improvement in feature size.

3. Generalizations

In this section we present several generalizations of the results of Section 2.1. We first introduce the concept of cartograms on weighted area. We then show how to generalize table cartograms to other shapes and surfaces.

3.1. Cartograms on Weighted Area

One direction for generalizations is to use a broader notion of "area". This can be done by specifying the weight of a region as an integral over some density function $w : R \to \mathbb{R}^+$. The density function should have the property that the integrals over triangular regions with nonempty interiors exist and are positive. We call such a density function positive. The following lemma generalizes Lemma 1, which can be used to generalize the results of Section 2.1 for arbitrary positive density functions.

Lemma 5 (Weighted Triangle Lemma). *Let $\triangle abc$ and $w : \triangle abc \to \mathbb{R}^+$ be a triangle and a positive density function on $\triangle abc$, respectively. Let $Area(\triangle abc)$ be the w-weighted area of the triangle $\triangle abc$. Given three non-negative real numbers α, β, γ, where $\alpha+\beta+\gamma = Area(\triangle abc)$, there exists a unique point p inside $\triangle abc$ such that $Area(\triangle pbc) = \alpha$, $Area(\triangle apc) = \beta$, and $Area(\triangle abp) = \gamma$.*

Proof. Let ρ be a ray starting at a and intersecting the segment bc. Since w is a positive density function, $Area_w(\triangle xbc)$ is strictly decreasing as x moves from a along ρ. Hence there is a unique point $x(\rho)$ on ρ such that $Area_w(\triangle x(\rho)bc) = \alpha$. The points $x(\rho)$ for all different ρ trace a simple curve P_a in $\triangle abc$ that separates a from bc. This curve has a point on ab and a point on ac. Similarly, we get a curve P_b that intersects every ray from b to ac at a point x with $Area_w(\triangle axc) = \beta$.

Let q_a and q_b be the intersection points of P_a and P_b with segment ab. Assume that the order of points on segment ab is a, q_a, q_b, b, then the triangles $\triangle q_a bc$ and $\triangle a q_b c$ cover the triangle $\triangle abc$ so that $Area_w(\triangle abc) < \alpha + \beta$. This contradiction shows that the order of points on segment ab is a, q_b, q_a, b. Hence P_a and P_b intersect at a point p inside of $\triangle abc$. This point p has the desired properties.

Assume that p and p' are two points with the desired properties. Then there is a pair $\{y, z\}$ among $\{a, b, c\}$ such that $\triangle(pyz) \subset \triangle(p'yz)$ and hence $Area_w(\triangle pyz) < Area_w(\triangle p'yz)$. This contradiction proves uniqueness. □

Note that Lemma 5 can be deduced in several other ways, e.g., using the Knaster-Kuratowski-Mazurkiewicz lemma [19] (even for higher dimensions), or in the context of barycentric coordinate systems [20].

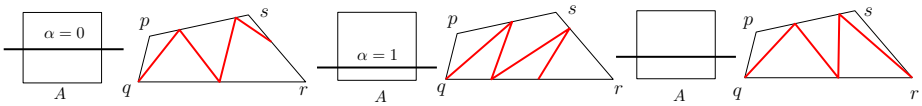

Figure 8: (a)–(b) Two zigzag paths Z_0 and Z_1 such that their endpoints define a continuous interval on the boundary of $\square pqrs$ that contains r. (c) By the intermediate value theorem, for some α the corresponding zigzag path must hit r.

3.2. Cartogram in an arbitrary convex quadrilateral

We show that every table can be represented as a cartogram in any arbitrary convex quadrilateral with area equal to the sum of the numbers in the table.

Theorem 3. *Let A be a table with m rows and n columns of non-negative numbers. Let $\square pqrs$ be an arbitrary convex quadrilateral with area equal to the sum, S, of the numbers of A. Then there exists a cartogram of A in $\square pqrs$ (with degeneracies).*

Proof. The proof for the case when $\square pqrs$ is a rectangle can be found from Theorem 1. We may thus assume that $\square pqrs$ is a convex quadrilateral that is not a rectangle, as shown in Figure 8(a). We now show that there must exist a zigzag path Z that starts at q and ends at r or s (depending on whether n is even or odd), and realizes the required areas for the intermediate triangles.

Note that such a zigzag path Z does not necessarily split A into two equal halves, rather, the row k of A that determines A^t and A^b satisfies $\sum_{1\leq i\leq k-1} S_i + \lambda S_k = \alpha S$ for some $\alpha \in (0, 1]$. We prove that such a zigzag path exists by using the intermediate value theorem, as follows. Without loss of generality assume that n is even. For each $k = 0, 1, \ldots, m$, we find the zigzag paths Z_0 and Z_1 by choosing $\alpha = 0$ and $\alpha = 1$. Observe that for some k, the end points of Z_0 and Z_1 define a continuous interval on the boundary of $\square pqrs$ that contains r. By the intermediate value theorem, for some α, the corresponding zigzag path must hit r. Figures 8(a)–(c) illustrate such a scenario.

Given the zigzag path Z, we can continue as in the proof of Theorem 1. □

Note that the proof of Theorem 3 is existential. A natural approach to find the required value of α would be to use a binary search in the interval $[0, 1]$. However, the running time in this case should depend on the precision required to define α. Hence applying a binary search would give only an approximate solution in practice.

3.3. Table Cartograms in a Circle

In this section we show that any table A of size $m \times n$ of non-negative weights admits a cartogram inside a circle with area equal to the sum of the numbers of A, where every edge in the cartogram is represented as a circular arc. To prove the existence of such a cartogram, we modify the proof of Lemma 5.

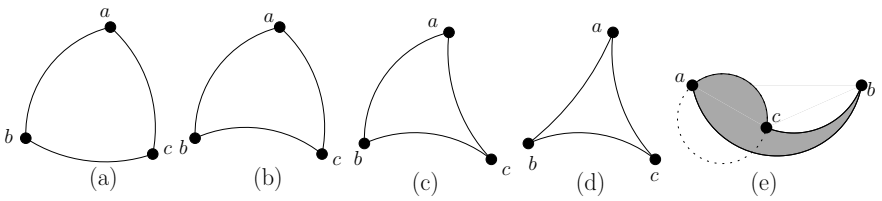

Figure 9: (a)–(d) Circular triangles. (e) Not a circular triangle.

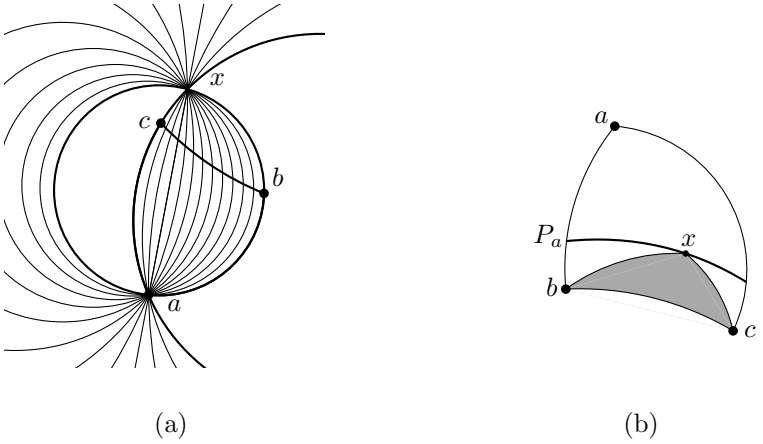

Figure 10: (a) The set of circles $\mathcal{C}_{ax}$ used to define $S_a(t)$ for $\triangle abc$. (b) The trace P_a of points x forming a circular triangle $\triangle xbc$ with area α.

A *circular triangle* $\triangle abc$ is (the bounded component of) the intersection of three closed regions, each of which is the interior or exterior of one of three pairwise intersecting circles. The region is bounded by three circular arcs (called *arms*) that pairwise meet at the points a, b, and c (called *vertices*). Figures 9(a)–(d) are circular triangles, but Figure 9(e) is not since the (shaded) region is contained in both the interior and exterior of the circle defined by arm ac. A closed disc with three vertices on its boundary is a circular triangle, where all three arms are the arcs of a single circle.

For $\triangle abc$, let C_{ab}, C_{bc}, and C_{ac} be the circles defined by arms ab, bc, and ac, respectively. A pair of these circles, say C_{ab} and C_{ac}, intersect at a. Let x be their additional intersection point (if it is unique), $x = a$ (if C_{ab} is tangent to C_{ac}), or x be any point on C_{ab} that is not contained in arm ab or ac (if $C_{ab} = C_{ac}$). If a and x are distinct points, let $\mathcal{C}_{ax}$ be the set of circles through a and x (see Figure 10(a)). Otherwise (if C_{ab} is tangent to C_{ac}) let $\mathcal{C}_{ax}$ be the set of circles C through a such that one of C_{ab} and C_{ac} is in the interior and the other in the exterior of C.

The set $\mathcal{C}_{ax}$ defines a set of circular arcs, $\{C \cap \triangle abc | C \in \mathcal{C}_{ax}\}$, each of which partitions $\triangle abc$ into two circular triangles. For example (as in Figure 10(a)), if $\triangle abc$ is the intersection of C_{ab}^{+} (the closed interior of C_{ab}), C_{bc}^{-} (the closed

exterior of C_{bc}), and C^+_{ac}, then the two circular triangles (for a circle $C \in \mathcal{C}_{ax}$) are $C^+_{ab} \cap C^-_{bc} \cap C^{\pm}$ and $C^+_{ac} \cap C^-_{bc} \cap \overline{C^{\pm}}$, where $C^{\pm}$ is the closed interior or exterior of C that contains arm ab and $\overline{C^{\pm}}$ is its closed complement.

Let $S_a(t)$ be a function from $t \in [0,1]$ to circular arcs $\{C \cap \triangle abc | C \in \mathcal{C}_{ax}\}$ that interpolates between $S_a(0) =$ arm ab and $S_a(1) =$ arm ac, so that the area of the left partition (in $C^{\pm}$) continuously increases from 0 to $Area(\triangle abc)$. Define $S_b(t)$ and $S_c(t)$ in a similar way.

Lemma 6 (Circular Triangle Lemma). *Let $\triangle abc$ be a circular triangle and let α, β, γ be non-negative numbers, where $\alpha + \beta + \gamma = Area(\triangle abc)$. Then there exists a point p in $\triangle abc$ and circular triangles $\triangle pbc$, $\triangle apc$, $\triangle abp$ such that $Area(\triangle pbc) = \alpha$, $Area(\triangle apc) = \beta$, and $Area(\triangle abp) = \gamma$.*

Proof. The proof is similar to the proof of Lemma 5. We first find the trace P_a of a point x inside $\triangle abc$ such that $Area(\triangle xbc) = \alpha$, as shown in Figure 10(b). Note that the arms bx and cx of $\triangle xbc$ are determined by the arcs of S_b and S_c that pass through x. Therefore, each ray in S_a can be intersected by P_a at most once. We define P_b with respect to S_b analogously. In a similar way to the proof of Lemma 5, we can prove that P_a and P_b intersect at a single point p such that $Area(\triangle pbc) = \alpha$, $Area(\triangle apc) = \beta$, and $Area(\triangle abp) = \gamma$. Note that each of $\triangle pbc$, $\triangle apc$, and $\triangle abp$ is determined by the intersection of the interior/exterior of three mutually intersecting circles. Therefore, $\triangle pbc$, $\triangle apc$, and $\triangle abp$ are circular triangles. □

A plane graph G with $n \geq 3$ vertices is called a plane 3-tree if G is a triangulated plane graph, and for $n > 3$, G has a vertex whose deletion gives a plane 3-tree of $n-1$ vertices. Let G be an arbitrary plane 3-tree such the internal faces of G are assigned non-negative weights. Any plane 3-tree admits a straight-line drawing inside a triangle, where the faces are drawn as triangles respecting the prescribed face weights [21]. One can use Lemma 6 to generalize this result to circular triangles, i.e., given a circular triangle $\triangle abc$ such that $Area(\triangle abc)$ is equal to the sum of the face weights of G, one can obtain a drawing of G inside $\triangle abc$ where the faces are drawn as circular triangles respecting the prescribed face weights.

Lemma 7. *Let G be an arbitrary plane 3-tree such the internal faces of G are assigned non-negative weights. Let $\triangle abc$ be a circular triangle with $Area(\triangle abc)$ equal to the sum of the face weights of G. Then there exists a drawing of G inside $\triangle abc$ such that each face of G is drawn as a circular triangle with area equal to its prescribed weight.*

We are now ready prove the main result of this section.

Theorem 4. *Let A be a table with m rows and n columns of non-negative numbers. Let C be an arbitrary circle with area equal to the sum of the numbers of A. Then there exists a cartogram of A in C where every face is a circular triangle.*

PROOF. Similar to the proof of Theorem 3, we first find the largest index k such that the sum of the numbers in rows $1, 2, \ldots, k-1$ is less than $S/2$. We may choose $\lambda \in (0, 1]$ such that $\sum_{1 \le i \le k-1} S_i + \lambda S_k = S/2$. We then define the tables A^t and A^b with respect to row k. We next find a zigzag path in G that divides the circle into circular triangles. Finally, we use Lemma 6 to compute the cartograms of the columns of A^t and A^b that correspond to those circular triangles, and remove the segments on the zigzag path to merge the cells corresponding to the row k of A. □

3.4. Table Cartogram on a Sphere

In this section we show that one can compute a cartogram of a table on the surface of a sphere. The idea is to first construct a cartogram[1] of the table on a cylinder, and then use the map projection techniques that preserve area [22, 23] to find a cartogram on the surface sphere. There are several well studied area preserving map projection techniques such as Lambert's cylindrical equal-area projection, Lambert azimuthal equal-area projection, or Hammer-Aitoff Equal-Area Projection. Here we use Lambert's cylindrical equal-area projection.

Lemma 8. *Let A be an $m \times n$ table of non-negative numbers. Let H be a cylinder with surface area equal to the sum of the numbers of A. Then there exists a cartogram of A on H.*

PROOF. The case when $n = 1$ is straightforward. Therefore, we assume that $n \ge 2$. Let h be the height and r be the radius of H. If n is even, then we first compute a cartogram on a rectangle H of height h and width $2\pi r$ in a similar way as in the proof of Theorem 1. Since n is even, the zig-zag path starts and ends at the same corner, when we wrap H into a cylinder. However, the vertices on the left side of H may not match the vertices on the right side of H. So we merge the leftmost and rightmost triangles formed by the zigzag path in H and find the cartogram of columns 1 and n of A^t in the resulting triangle using Lemma 2. This process is shown in Figures 11(a)–(d).

If n is odd, we equally divide the leftmost column into two columns and move one of these columns to the right of A, as shown in Figures 11(e)–(f). We then compute the cartogram, as in the case when n is even, and remove the segments that separate the divided leftmost column of A^t, and the (infinitesimal) segments that separate the divided leftmost column of A^b, as in Figures 11(g) and (h).

When n is even, the faces are convex quadrilaterals. However, when n is odd, the faces with areas from the leftmost column of A may be concave hexagons, as shown in Figure 11(h). □

We now use Lambert's cylindrical equal-area projection to find a cartogram of A on the surface of a sphere. Note that this technique only projects the two

[1]This cartogram may have non-convex faces.

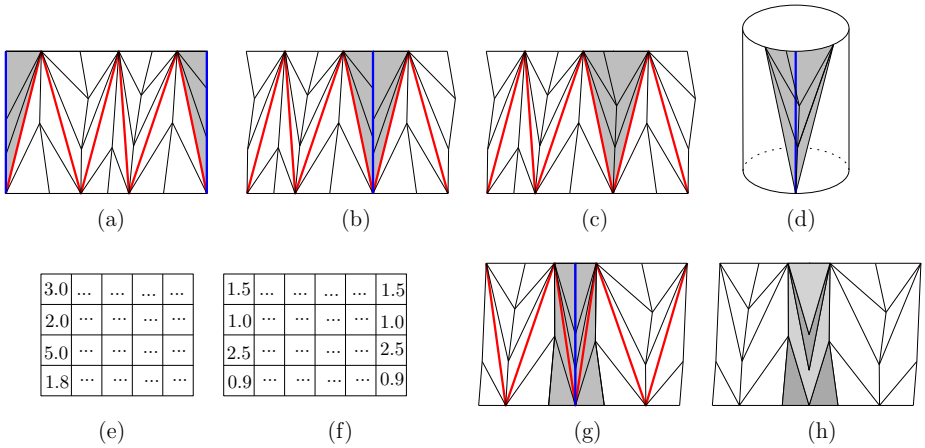

Figure 11: Cylindrical cartogram construction. (a)–(d) n is even, (e)–(h) n is odd. Note that when n is even, the faces are convex quadrilaterals. However, when n is odd, the faces with areas from the leftmost column of A may be concave hexagons.

dimensional cartogram on the sphere. Some parts of the sphere, for example the poles, will not be covered.

A major problem of using an equal area projection to compute cartograms on the surface of a sphere is the distortion of shapes. Therefore, it is worth investigating algorithms that can produce aesthetically pleasing cartograms on a sphere.

4. Conclusions and Future Work

We have presented a simple constructive algorithm that realizes any table inside a rectangle in which each cell is represented by a convex quadrilateral with its prescribed weight. If all weights are strictly positive, then we can also obtain non-degenerate realizations. This method can be further extended to realize any table inside an arbitrary convex quadrilateral, inside a circle using circular arcs, or even on the surface of a sphere. From a practical point of view, the cartograms obtained by our method may not be visually pleasing, but by using additional straightforward heuristics that improve the visual quality while keeping the areas the same, we can obtain cartograms of practical relevance, as shown in Figure 1 and Figure 4. Our theoretical solution plays a vital role in this context, since heuristics used directly may get stuck, being unable to obtain the correct areas. Whether there exists a method that can gradually change the areas to provably obtain the correct areas remains an interesting open problem. It would also be interesting to examine table cartograms for other types of tables, such as triangular or hexagonal grids. From a theoretical point of view, finding algorithms for table cartograms on a sphere with less distortion, and generalizing our result to 3D table cartograms (inside a box) are further interesting open problems.

Acknowledgments

Initial work on this problem began at Dagstuhl Seminar 12261 "Putting Data on the Map" in June 2012 and most of the results of this paper were obtained at the Barbados Computational Geometry workshop in February 2013. We would like to thank the organizers of these events, as well as many participants for fruitful discussions and suggestions.

Dedication

We fondly remember our interactions with Ferran Hurtado. His keen interest and enjoyment in geometric problems inspired an increased pleasure in working in this area.

References

[1] W. S. Evans, S. Felsner, M. Kaufmann, S. G. Kobourov, D. Mondal, R. I. Nishat, K. Verbeek, Table cartograms, in: Proceedings of the 21st Annual European Symposium on Algorithms (ESA), Vol. 8125 of LNCS, Springer, 2013, pp. 421–432.

[2] E. Raisz, The rectangular statistical cartogram, Geographical Review 24 (2) (1934) 292–296.

[3] M. van Kreveld, B. Speckmann, On rectangular cartograms, Computational Geometry: Theory and Applications 37 (3) (2007) 175–187.

[4] R. Heilmann, D. A. Keim, C. Panse, M. Sips, Recmap: Rectangular map approximations, in: Proceedings of the IEEE Symposium on Information Visualization (INFOVIS), 2004, pp. 33–40.

[5] J. Wood, J. Dykes, Spatially ordered treemaps, IEEE Transactions on Visualization and Computer Graphics 14 (6) (2008) 1348–1355.

[6] D. Eppstein, M. J. van Kreveld, B. Speckmann, F. Staals, Improved grid map layout by point set matching, International Journal of Computational Geometry and Application 25 (2) (2015) 101–122.

[7] D. Eppstein, E. Mumford, B. Speckmann, K. Verbeek, Area-universal and constrained rectangular layouts, SIAM Journal on Computing 41 (3) (2012) 537–564.

[8] M. de Berg, E. Mumford, B. Speckmann, On rectilinear duals for vertex-weighted plane graphs, Discrete Mathematics 309 (7) (2009) 1794–1812.

[9] M. Alam, T. Biedl, S. Felsner, M. Kaufmann, S. G. Kobourov, T. Ueckerdt, Computing cartograms with optimal complexity, Discrete & Computational Geometry 50 (3) (2013) 784–810.

[10] K.-H. Yeap, M. Sarrafzadeh, Floor-planning by graph dualization: 2-concave rectilinear modules, SIAM Journal on Computing 22 (3) (1993) 500–526.

[11] M. J. Alam, T. Biedl, S. Felsner, M. Kaufmann, S. G. Kobourov, Proportional contact representations of planar graphs, Journal of Graph Algorithms and Applications 16 (3) (2012) 701–728.

[12] J. A. Dougenik, N. R. Chrisman, D. R. Niemeyer, An algorithm to construct continuous area cartograms, The Professional Geographer 37 (1) (1985) 75–81.

[13] D. Dorling, Area cartograms: their use and creation, no. 59 in Concepts and Techniques in Modern Geography, University of East Anglia, 1996.

[14] D. A. Keim, S. C. North, C. Panse, Cartodraw: A fast algorithm for generating contiguous cartograms, IEEE Transactions on Visualization and Computer Graphics 10 (1) (2004) 95–110.

[15] M. T. Gastner, M. E. J. Newman, Diffusion-based method for producing density-equalizing maps, National Academy of Sciences 101 (20) (2004) 7499–7504.

[16] H. Edelsbrunner, R. Waupotitsch, A combinatorial approach to cartograms, Computational Geometry: Theory and Applications 7 (5–6) (1997) 343–360.

[17] D. H. House, C. J. Kocmoud, Continuous cartogram construction, in: Proceedings of the conference on Visualization (VIS), 1998, pp. 197–204.

[18] W. Tobler, Thirty five years of computer cartograms, Annals Assoc. American Geographers 94 (1) (2004) 58–73.

[19] B. Knaster, C. Kuratowski, S. Mazurkiewicz, Ein Beweis des Fixpunktsatzes für n-dimensionale Simplexe, Fundamenta Mathematicae 14 (1929) 132–137.

[20] H. S. M. Coxeter, Introduction to Geometry, 2nd Edition, Wiley, New York, 1969.

[21] T. C. Biedl, L. E. R. Velázquez, Drawing planar 3-trees with given face areas, Computational Geometry: Theory and Applications 46 (3) (2013) 276–285.

[22] T. G. Feeman, Portraits of the Earth: A Mathematician Looks at Maps, American Mathematical Society, 2002.

[23] J. P. Snyder, Map Projections–A Working Manual, U. S. Geological Survey Professional Paper 1935, 1987.