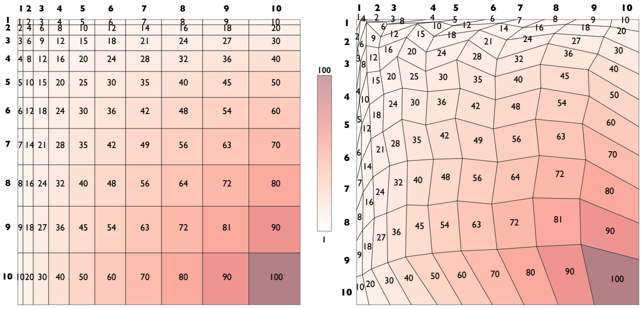

Figure 1: *The multiplicity of accurate layouts for a single input table yields a prominent table cartogram* **Hallucinator** *as in this pair of multiplication tables.*

Appendix A: Supplemental Material

In this supplement we provide a collection of asides and graphic studies demonstrating various TACO properties that did not fit in the main paper. Among these are Fig. 1 which provides a supplementary visual explanation of the layout ***Hallucinator***, and Fig. 2 which compares presentations of a hierarchical nominal dataset.

Non-axial Table Cartogram Arrangements

Tabular renderings of data are often organized by the meaning of the axes, that is by rows and columns—however this need not always be the case. Such non-axial layouts can be rendered as table cartograms, as in Fig. 3 and Fig. 4. These possess a ***Hallucinator***, as any arrangement of the data in a 2D layout will be non-reflective of the data and thus can be manipulated to make certain data classes appear larger or smaller than their actual data may warrant. These structures abandon one of the key affordances of the table cartogram: the rigid maintenance of adjacency, which makes it substantially less clear why one might select this design over an equivalent one (such as any of those shown in Fig. 3). We suggest that such structures might only be appropriate in cases where novelty is an essential design component.

Area Embedding in Familiar Forms

Sometimes tabular data describes entities that already canonically inhabit a table, such as slices of the periodic table of elements (as exemplified by Evans et al. [EFK*13]). The reader's familiarity with that tabular layout simplifies understanding the table cartogram, analogous to how geographic cartograms leverage the reader's prior knowledge of geography. Fig. 5 shows a table cartogram based on a Dungeons and Dragons *moral alignment chart*, a standard and frequently satirized 3×3 table. Readers familiar with that form can quickly read that players tend to prefer to be good and chaotic, since the visualization respects the standardized placement of those properties. While a similar effect could be achieved using a slice-dice tree map or mosaic plot, it may have the effect of breaking from the canonical form as in Fig. 5c. This also aligns with our suggestion that calendar displays pair effectively with TACOs, as their form is well understood and commonly used.

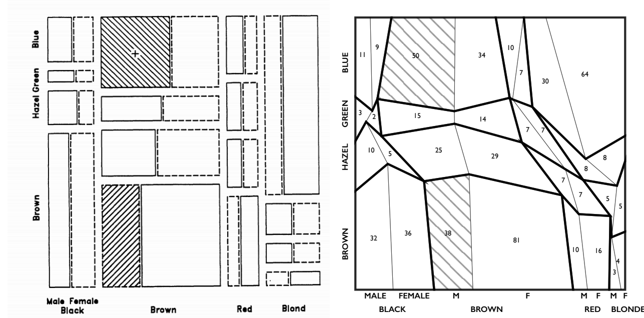

Figure 2: *A comparison between a mosaic diagram created by Friendly [Fri94], and the same data constructed as a TACO. We add bold outlines to the TACO to mirror the spaces that Friendly denotes with whitespace, and numbers (though they are not featured in the original image) for context. While the TACO rigidly maintains adjacency of the input data it does so at the expense of legibility, as comparisons of area in rectangles are easier to make than comparisons between* "blobs" *(as the quadrilaterals of the TACO might be interpreted) [CM86]. This exchange may be appropriate when the axes are ordinal, which is not the case here.*

On Data with Zeros

Rather than merely arguing data with zeros this should be avoided entirely, we now briefly consider a design strategy to address tabular data with values in $\mathbb{R}^{\geq 0}$. Despite recommendations to the contrary, a designer may feel that it is necessary to construct a table cartogram including such data. In these cases a reasonable approach is to treat the data values as interval (as discussed in Sec. **??**, can be used if the task is appropriate, such as FIND EXTREMUM or RETRIEVE VALUE), as this allows for arbitrary shifts and rescaling to numerical data. The zero value can then be rescaled to a visually appropriate value, which will necessarily be design dependent.

An example of this strategy is Pierebean's analysis of frequency of Cantonese characters by pronunciation [pie]. This data is organized into a table by international phonetic alphabet initials (rows) and finals (columns), with each cell made into a scalar by counting the number of characters that match that combination. After the table cartogram layout is computed, a word cloud of the characters in each cell is placed into each cell. There are a number of combinations that do not occur in the Cantonese language, yielding zeros. To address this the designer treated these counts as interval and offset them to give the desired form to the design (in particular setting the zero value to be 25).

While this selection is perhaps algebraically unsound—as ratio data should be treated as ratio data while interval data should be treated as interval data—it does serve a different goal of being visually interesting and giving granular access to the data. Alternatives explored elsewhere in this supplement are unable to capture this effect. Mosaics would destroy the legibility of the axes. A heatmap would not leave room for the literal representation of the characters (which appears to be an important part of the graphic's appeal). Something more exotic, such as a gridded beeswarm chart would present the data as ratio and would still expose the literal encoding [Bra20] but at expense of the vertical and horizontal space

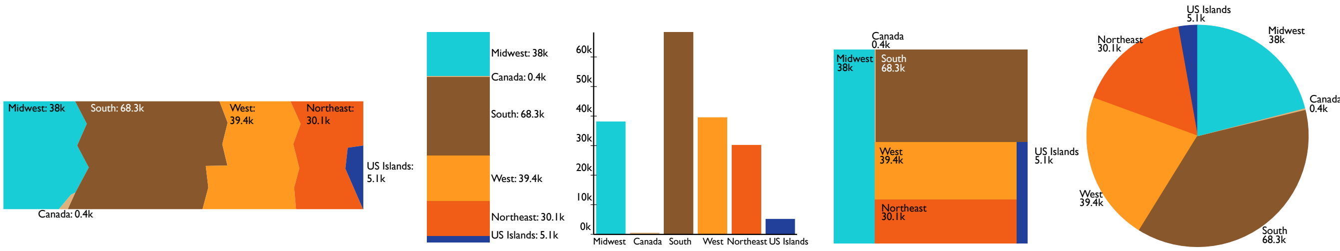

Figure 3: *Bird-airplane collisions [FAAb] by region of flight origin rendered as a table cartogram by treating a waffle plot of original data as a table, as well as some common alternatives. (Left to right: TACO, stacked bar chart, bar chart, tree map, pie chart).*

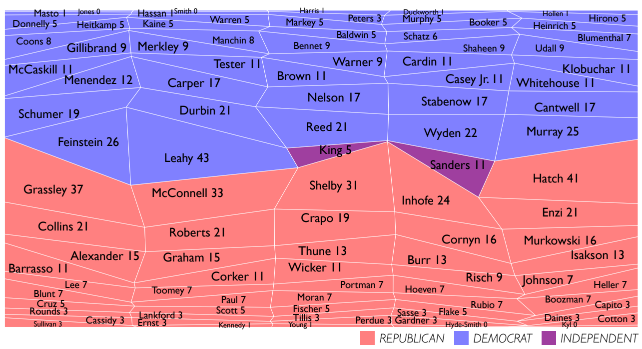

Figure 4: *Political party and length of incumbency in US senate, fall 2018 [FAAa]. Cells are organized into a waffle plot layout by party and by incumbency length and then mapped into a TACO.*

(though the Pierebean's graphic is already quite large). Each of these alternatives should be considered before using a TACO to represent ratio data with zeros. If these are insufficient then we recommend that the designer *"interval-ize"* their data, use annotations and labels to prompt ordinal comparisons, and to annotate zeros through any available channel (such as texture or color).

The tension between this potentially problematic usage and design desire is exactly akin to the problem of specifying normalization conditions in dual-y-axis charts (which is a ***Hallucinator***). We include this recommendation here because practitioners break formal rules (again just as in dual-y-axes), and we would prefer to give comprehensive guidance rather than piecemeal.

On Rendering Table Cartograms

A possible factor for the current rarity of TACOs is a lack of easily accessible implementations. At the time of this writing there are two known publicly available implementations, both of which utilize optimization-based schemes for computation [Has21, MK20]. There are several more known versions, which are not available publicly, including Inoue and Li's [IL20] optimization-based approach, and Evans et al.'s [EFK*13] geometry-based and optimization-based approaches. The latter of which they reference as producing their figures, although details are not provided. The figures in this paper are created using our typescript implementation [MK20], as its permissive interface design suited our need to explore the possible output space generated by the various parameter configurations. In future work it would be interesting to compare the characteristics of the TACOs created by these implementations.

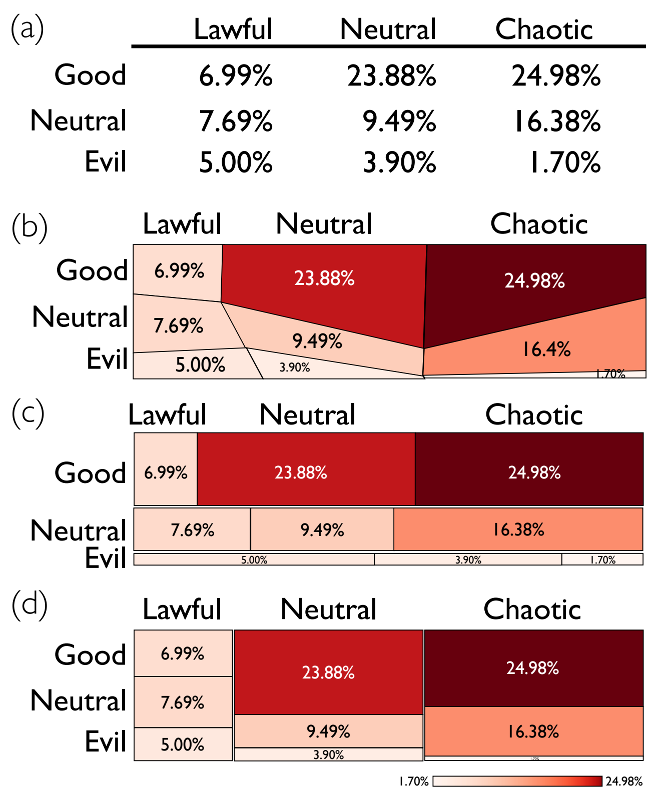

Figure 5: *The popularity of the character alignments available to players in Dungeons and Dragons found in an informal online survey [Red13] , across a standard table (a), a TACO (b), and two mosaics (c, d). While the table surfaces maxima via the number of digits, the other forms afford the same task (for both maxima and minima) via area. The TACO gives access to both marginal values, while the tree maps afford an accurate reading of a single marginal value (as comparing rectangular areas is more accurate than comparing blobs [CM86]).*

Additional Design Studies

A larger calendar. As we saw in Sec. **??** year and month calendars can be usefully combined to create more complex compositions. In Fig. 7 we show a slightly simpler one, that reveals how these traffic violations are highly correlated with the Chicago public schools calendar (e.g., the April 10-14 spring break accompanies an unusual dip in violations [Chib]). Fig. **??** is drawn from this figure.

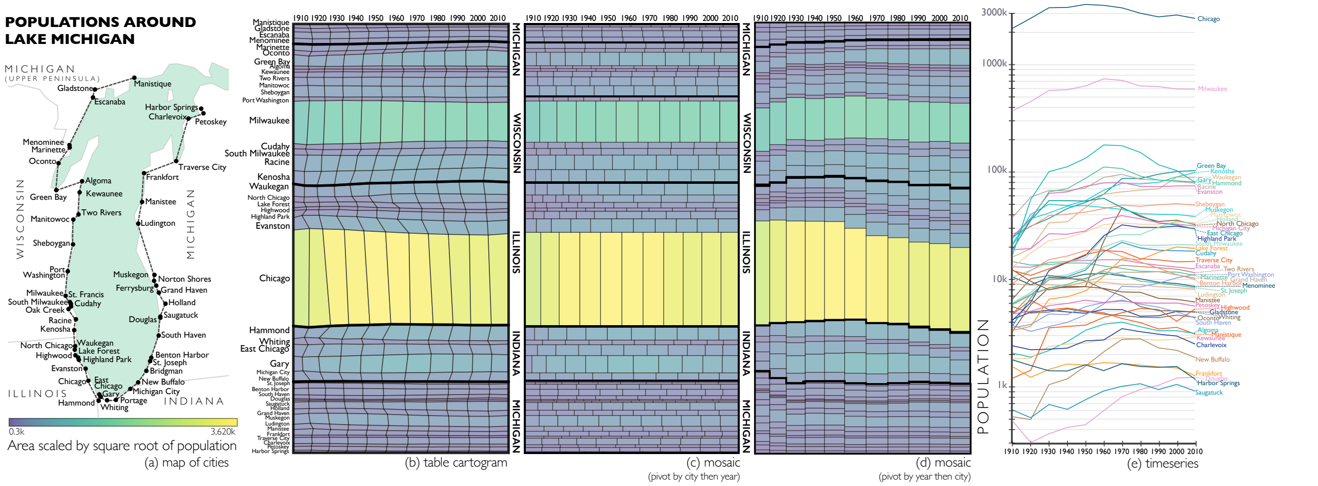

Figure 6: *Several displays of changes in populations of cities around Lake Michigan during the 20th century. (b)-(d) utilize an ordering created by a counter-clockwise lake (a) traversal starting at the break between Michigan's upper and lower peninsulas.*

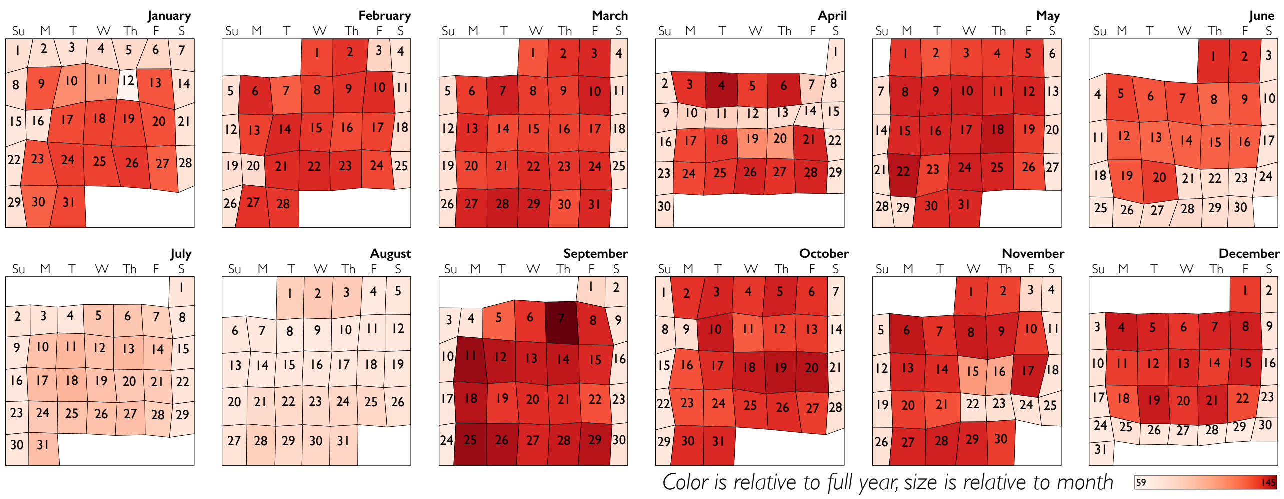

Figure 7: *This visualization reveals how speed camera violations in Children's Safety Zones in Chicago [Chia] are correlated with the Chicago public schools calendar (e.g., the April 10-14 spring break accompanies an unusual dip in violations [Chib]).*

We already use calendars to structure our understanding of events and periods in the year, so just as other familiar forms may aid our understanding of graphics possessing an area-embedded structure, so too can these forms promote understanding of data presented in this manner.

Spatio-temporal data. We argued that table cartograms are best suited to datasets with ordinal axes, which manifested as calendars in several examples. Yet this is by no-means the entirety of the possibilities for this category of data. In Fig. 6 we consider one such dataset: the populations of the cities around Lake Michigan over the 20th century [Uni19]. In (b-d) the ordering of the cities around the lake are linearized by traversing from counter-clockwise from the natural break at the top of Michigan (a), this is necessary to construct a relevant 1D ordering that might be used by one of the axes of a table cartogram. (b-e) all facilitate basic summaries, like that Chicago always has been the most populous city on the lake, but (c-e) lose temporal or ordinal detail. In (c) population changes over time are lost, in (d) the notion of a city is illegible due to the number of cities being pivoted across, and in (e) the ordering of the cities around the lake is lost, which precludes making observations about the geographic distribution of populations. The table cartogram (b) provides these properties by maintaining order and adjacency of the cells in the table. However it does so at the price of precise comparison between years or cities (in d and c respectively). This might not be a significant deficit if the task prefers individual comparisons instead of marginal ones. These graphics might be interestingly combined through Ritchie et al.'s [RWC19] quasi-modes in which alternate views are presented on-demand. This graphic features a minor ***Hallucinator*** in that the break-point that generates the linearization of the cities is arbitrarily chosen, as other values would lead to differing layouts.

References

[Bra20] BRATH R.: Literal Encoding: Text is a first-class data encoding. *Visualization for the Digital Humanities (VIS4DH)* (2020). 1

[Chia] CHICAGO DATA PORTAL: Speed Camera Violations. https://data.cityofchicago.org/Transportation/Speed-Camera-Violations/hhkd-xvj4. Accessed 09/23/2018. 3

[Chib] CHICAGO PUBLIC SCHOOLS: School Year Calendar. www.cps.edu. Accessed 10/04/2018. 2, 3

[CM86] CLEVELAND W. S., MCGILL R.: An Experiment in Graphical Perception. *International Journal of Man-Machine Studies 25*, 5 (1986), 491–500. 1, 2

[EFK*13] EVANS W., FELSNER S., KAUFMANN M., KOBOUROV S. G., MONDAL D., NISHAT R. I., VERBEEK K.: Table Cartograms. In *European Symposium on Algorithms* (2013), Springer, pp. 421–432. doi:10.1007/978-3-642-40450-4_36. 1, 2

[FAAa] FAA: United States Senate. https://www.senate.gov. 2

[FAAb] FAA: Wildlife Strike Database. https://wildlife.faa.gov. 2

[Fri94] FRIENDLY M.: Mosaic Displays for Multi-Way Contingency Tables. *Journal of the American Statistical Association 89*, 425 (1994), 190–200. 1

[Has21] HASAN M. R.: tcarto_applications. https://github.com/rakib045/tcarto_applications, 2021. Accessed 2/21/2021. 2

[IL20] INOUE R., LI M.: Optimization-Based Construction of Quadrilateral Table Cartograms. *ISPRS International Journal of Geo-Information 9*, 1 (2020), 43. doi:10.3390/ijgi9010043. 2

[MK20] MCNUTT A., KINDLMANN G.: A Minimally Constrained Optimization Algorithm for Table Cartograms. *IEEEVIS InfoVis Posters* (2020). 2

[pie] PIEREBEAN: [OC] >19k chinese characters sorted by pronunciation. www.reddit.com/r/dataisbeautiful/comments/jm042v/oc_19k_chinese_characters_sorted_by_pronunciation/. Accessed 11/2/2020. 1

[Red13] REDDIT USERS: Alignment Survey - Results. https://www.reddit.com/r/DnD/comments/1ejnft/alignment_survey_results/, 2013. Accessed 11/2/2020. 2

[RWC19] RITCHIE J., WIGDOR D., CHEVALIER F.: A Lie Reveals the Truth: Quasimodes for Task-Aligned Data Presentation. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (2019), pp. 1–13. doi:10.1145/3290605.3300423. 3

[Uni19] UNITED STATES CENSUS BUREAU: Decennial census of population and housing. https://www.census.gov/programs-surveys/decennial-census/data.html, 2019. Accessed 3/15/2019. 3