

What are Table Cartograms Good for Anyway? An Algebraic Analysis

A. McNutt 

Department of Computer Science, University of Chicago

Abstract

Unfamiliar or esoteric visual forms arise in many areas of visualization. While such forms can be intriguing, it can be unclear how to make effective use of them without long periods of practice or costly user studies. In this work we analyze the table cartogram—a graphic which visualizes tabular data by bringing the areas of a grid of quadrilaterals into correspondence with the input data, like a heat map that has been “area-ed” rather than colored. Despite having existed for several years, little is known about its appropriate usage. We mend this gap by using Algebraic Visualization Design to show that they are best suited to relatively small tables with ordinal axes for some comparison and outlier identification tasks. In doing so we demonstrate a discount theory-based analysis that can be used to cheaply determine best practices for unknown visualizations.

CCS Concepts

• *Human-centered computing* → *Visualization design and evaluation methods; Visualization theory, concepts and paradigms;*

1. Introduction

Understanding whether a chart has been used effectively is an important problem in the practice of data visualization. Possessing a clear notion of what is “good” or “bad” usage is critical, as it can guide designers towards impactful information-rich graphics and away from deceptive displays [CH17]. Yet even for those familiar with visualization best practices, it can be difficult to know if an unfamiliar chart type has been used effectively.

There are at least a dozen decades of advice [Wic13] on how to best use extant chart forms for known data types, yet this advice rarely applies to novel charts. This lack of guidance can impede effective usage by practitioners, can cause domain experts to question a design’s validity [VDBBC*18], and can impede automated analyses from making relevant suggestions [MCC20]. Developing a technical understanding of effective usage for a novel chart can be a quagmire of disentangling aesthetic and novelty responses [CM07] through often slow or costly user studies [AL20].

A potential salve is to utilize a theory-based analysis to generate guidelines. This would enable users without expertise or access to specialized analysis software to generate best practices for themselves. This is similar in spirit to discount usability studies [ZSN*06], which succinctly characterize a system’s usability through evaluation of heuristics by a small number of analysts. Prior efforts to apply theory in this way have focused on guiding design processes already situated in their task and context, rather than on understanding usage of a particular chart form [KS14, ZC06].

We demonstrate the potential of such an analysis by investigating the properties of an uncommon visualization—the table cartogram (TACO)—through a lens informed by Kindlmann and Scheidegger’s Algebraic Visualization Design (AVD) [KS14]. We study TACOs because, despite having existed for several years [EFK*13], appropriate usage is still unknown [MK20]. We focus on AVD because—in contrast to other frameworks—it provides concrete assertions about visualization quality which are human-operable and interpretable. Further, its methods are removed from the embodied response to visualizations, reducing the novelty effect that might be brought on by unusual visual forms.

We contribute an example of how AVD might be used to derive guidelines for novel visual forms, by considering what data (Sec. 3.1) and tasks (Sec. 3.2) are appropriate for one such form. While not every question can be answered using these tools (those related to perceptual quality are typically out of reach), we are able to construct cogent guidelines for most basic usage questions. This allows us to contribute recommendations on how TACOs might be effectively used. We argue that they are well suited to some comparison and outlier identification tasks for flat tables that have ordinal rows and columns. We suggest that they may be effective in contexts in which analytical insight is not the primary goal, as well as in discrete representations of time—such as month calendars. While not every property will be surprising to those familiar with similar graphics, our investigation offers a full picture of TACOs that explains their usage and helps form an agenda for their future study.

2. Related Work

Our work constructs a theory-driven analysis of an uncommon graphic, the table cartogram. We now ground this analysis by describing TACO's history and known properties, then prior theory-based evaluations, and finally our theory of interest, AVD.

2.1. Table Cartograms

Table cartograms (TACOs) are a specialized variation of value-by-area maps (cartograms) that shows tabular data rather than solely geographic data [EFK*13]. They depict a table of positive numbers as a grid of quadrilaterals, constrained to a rectangle, whose areas are brought into correspondence with the input data. The visual effect is that of a shaded matrix that has been “area-ed” rather than colored. They can be characterized as having *Planar Grid Topology* (none of the cells overlap) and an *Accurate Data Embedding* (data is represented accurately as area) [MK20]. While color is a consonant secondary encoding, its use is not definitionally required. Inoue and Li [IL20] refer to TACOs not constrained to a containing rectangle as *deformed table cartograms*. Although these are of interest, we instead focus on *undeformed table cartograms* as constraining the chart unambiguously defines the meaning of area and usefully limits the domain of our analysis.

TACOs were first described by Evans et al. [EFK*13]. They use a computational geometry-based algorithm to constructively demonstrate that all tables of positive numbers admit a TACO. Subsequent work [LI19, IL20, MK20] found that more expressive graphics could be produced with optimization-based techniques, utilizing the fact the problem is under-constrained [IL20]. This allows for multiple outputs for a single input, as in Fig. 1.

Prior studies primarily focus on TACO's construction without providing substantial considerations on how they might be usefully employed. Evans et al. [EFK*13] explore a series of designs, but provide no usage guidance. Inoue and Li [IL20] briefly touch on usage (which informs our discussion of data types in Sec. 3.1), however they do not consider task effectiveness or appropriate data domains. Confounding the development of usage guidelines is the common identification of TACOs as a form of geographic cartogram [AKV15, NK16]. While they can be used for such data, few geographies have topologies that can be mapped to a grid without substantial adjacency distortion (grid mapping schemes [EvKSS15] can reduce these distortions [NK16]). We focus instead on tabular data as it covers tabularized geographic data and is rarely supported among familiar chart forms (tables and shaded-matrices being the best-known exceptions). Despite these prior studies, little is known about effective TACO usage.

In this paper we seek to rectify this gap in understanding by answering the question: *What are table cartograms good for anyway?*

2.2. Visualization Analysis

A number of theories have been developed to analyze visualization quality or effectiveness, variously seeking to explain specific phenomena or enable particular applications. Here we review a set of evaluatory theories so as to inform our choice to focus on AVD.




Figure 1: Table cartograms admit multiple equally-accurate layouts for a given input. This is a **Hallucinator**: the multiplicity of correct solutions may yield varying interpretations.

		Visual Change ω	
		Obvious	None or illegible
Data Change α	Insignificant	Hallucinator	Invariant Representation
	Significant	Unambiguous Visualization	Confuser


Figure 2: The primary **success** and **failure** modes in AVD.

Some theories metricize quality, which enables the use of computational measurement as a way to automatically evaluate visualizations. Behrisch et al. [BBK*18] describe the state of the art for metric-based qualitative evaluation of visualizations, however each of these metrics is associated with particular visual forms. Mackinlay [Mac86] describes notions of effectiveness and expressiveness which informs the design of some recommendation systems [Lee20]. However, over-reliance on these notions can impede richer designs [BCF20]. Chen et al.'s [CJ10, CG15] work on using Information Theory to reason about visualization quality yields certain desirable properties—such as generating Shneiderman's mantra [Shn96] as an emergent property. Yet, these techniques are intractable to unassisted humans. Demiralp et al. [DSK*14] describe a visual embedding-based assertion system, however it requires substantial experimental data tuned to a particular chart. While metric-based analyses can yield useful insights, they can be non-trivial to deploy, are often tuned to specific chart types (forgoing novel forms), and can lack clear interpretation. We focus on AVD because its assertions are interpretable and can be evaluated without specialized software (i.e. it is human operable).

At the other end of the human-computer agency spectrum are theories that provide human-centered evaluation tools based on heuristic or critical under-pinnings. Some of these are predicated on personal or philosophical reflection [DFCC13, BBD20, DK16, VFR12, BZJ*20] which provide useful means for prompting the design process, but do not include testable assertions by which to judge quality. Adar and Lee [AL20] utilize learning objectives as a way to evaluate communicative visualizations, which while providing a definite heuristic by which to evaluate, does not help the designer actually conduct that evaluation. Wall et al. [WAM*18] build a set of evaluatory heuristics, however their approach requires several domain-experts and that task is pre-established. Zuk and Cappendale [ZC06] conduct heuristic analyses of uncertainty visualizations through the works of Bertin, Tufte, and Ware. While potentially informative, these theories are piecemeal: they do not provide a holistic description of quality. Our aims are aligned with these human-operable theories in that we wish to furnish analysts with easy-to-apply tools for analyzing visualizations. AVD circumvents these issues by providing a systematic framework made up of concrete assertions whose structure helps mitigate potentially biased analyses.

2.3. Algebraic Visualization Design

At the core of our analysis is Kindlmann and Scheidegger’s Algebraic Visualization Design [KS14] (AVD). AVD is a framework for reasoning about the design of data visualizations through their intrinsic symmetries. Similar to how one may understand the properties of a triangle by identifying which rotations and reflections yield symmetries or asymmetries, AVD understands a visualization by exploring the effect of changes in data (referred to as α s) and the corresponding changes in the resulting image (ω s). This is mechanized by asserting that every α should have a corresponding ω , that is, these changes should commute:



Hallucinators and **Confusers**, two important failure modes (or asymmetries), occur when this commutativity assertion is not fulfilled (Fig. 2). If the image changes significantly as a result of only a small or superficial data change, then it has a **Hallucinator**: non-data is depicted in a way that risks appearing meaningful. Merely reordering columns of a radar chart can dramatically change its enclosed area and visual appearance, a **Hallucinator**. If the image does not appear to change with a significant data change, then it has a **Confuser**: a way in which the viewer will be blind to the data. Standard summary statistics (e.g. mean and variance) for Anscombe’s quartet [Ans73] yield **Confusers** across column changes (α s). Visualizations are probed for these states by adversarially selecting α s and ω s that might surface them.

All visualizations have **Hallucinators** and **Confusers**. Just as no chart is perfect for all occasions, no graphic is a panacea under algebraic analysis. Instead, **Confusers** must be chosen according to task, and **Hallucinators** generally minimized. These asymmetries help identify properties of the visualizations under consideration in a manner that is only minimally reliant on the observer’s experience; thereby providing distance from the human responses which might otherwise cloud an analysis of a novel graphic.

3. Algebraic Visualization Analysis of the Table Cartogram

We now carry out our algebraic analysis of the table cartogram. We organize our discussion by noting, per Hu et al. [HGH*19], that every visualization can be viewed as a triplet of (data, visual form, task). Only by correctly matching all three components will a particular visualization have value. A dataset and visual form might be well matched, but if their combination lacks details necessary to complete the task, then the resulting graphic will be ineffective.

As we are interested in a particular chart form (TACOs), we are left with two elements of this triplet to consider: data and task. Thus, we form our study around two corresponding questions: *What data can be used in this chart?* (Sec. 3.1) and *Which tasks are appropriate?* (Sec. 3.2). These questions capture many possible usage concerns, informing when and how TACOs might be used—although, as we discuss, there are properties that cannot be understood with this approach.

	Fix	Vary	Observe	Representative Question	Relevant Area or Work
Fix Data				What can be done with this data?	Exploratory Data Analysis
				Does this visual form work for this data?	Unknown
Fix Visual Form				What data can be used in this chart?	Sec 3.1
				Which tasks are appropriate for this encoding?	Sec 3.2
Fix Task				Does this visual form show the data clearly?	AVD Confusers , Lineups
				Does this visual form show non-data?	AVD Hallucinators

Triplet Model of Visualization: Data Visual Form Task

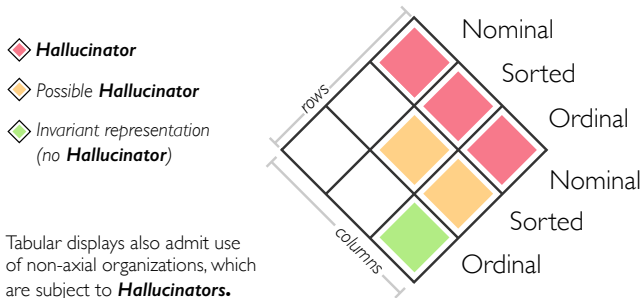
Figure 3: Many questions can be answered about visualization quality by taking a variational view of Hu et al.’s [HGH*19] triplet model. We focus on one of the many analysis families in this formulation by deriving properties of TACOs by varying task and data.

Methods. We perform this analysis by using a reframing of AVD. In past work AVD has usually been employed to guide the design process, wherein designs are invalidated for a fixed task by variation of data [KS14, WKD18]. We invert this procedure by using AVD as a way to guide data selection and task design for a given visual form, which we refer to as *Algebraic Visualization Analysis*.

We illustrate this variational scheme in Fig. 3. This locates our work among both that of AVD as well as Exploratory Data Analysis (which fits into this framing by viewing it as a process in which data is fixed and encoding is varied to see what tasks can be fulfilled). Wickham et al.’s [WCHB10] lineup protocol is closely related to past usage of AVD, fixing task and varying data to invalidate a design (though it can only be used once per dataset). Just as in the original AVD study, we select adversarial α s and ω s, potentially yielding one of the states in Fig. 2. We diverge from past usage by employing these methods not as a design tool, but as a way to better understand a chart form. Not all questions need to be answered by explicit variation: some failure modes can be identified by simply applying definitions. We forgo questions of interactivity as AVD is unable to reason about it, which, while unfortunate for real world applicability, usefully limits the scope of our discussion.

Warm up: A Hallucinator. As a warm up, we consider a prominent TACO **Hallucinator**. The under-constrained nature of the TACO allows there to be multiple equally accurate TACO layouts from the same data. Fig. 1 shows an example, garnered from changing algorithms, however significant variance can result from seemingly-innocuous parameter changes [MK20]. This multiplicity is a **Hallucinator**. While selecting Fig. 1(b) over Fig. 1(a) may be perceptually motivated (as it is easier to compare rectangular shapes), the value of selecting a particular starting condition in an optimization algorithm can be unpredictable and ambiguous. This contrasts tree maps, where algorithm design is well understood, predictable, and usually has meaningful design implications [SLD20].

One way to address this failure mode is to impose additional constraints—such as by minimizing bearing angle differences—which causes there to be a single “correct” layout [IL20]. While these criteria can yield more rectangular displays, their selection is arbitrary and an artifact only of designers’ preferences, an ambi-



Tabular displays also admit use of non-axial organizations, which are subject to **Hallucinators**.

Figure 4: The tabular form of TACO input data gives way to seven transpose unique axial data types.

guity which may in turn hold another **Hallucinator**. A reader comparing TACOs for different datasets prepared according to different heuristics (but who is unaware of this selection) may be deceived.

This flexibility is both a blight and a boon. The multiplicity of outputs creates space for uncertainty: *how can the reader know their interpretation of the chart is correct?* Yet this same property offers a great measure of freedom to chart designers to create visually interesting effects. The tension found between designer freedom and potential reader mistrust suggests that TACOs should not be used in decision-making contexts. Just as geographic cartograms are typically used to give big-picture summaries [NK16] rather than in data analysis, we argue that TACOs are best applied in situations where the task involves the readers awareness of the presentation medium itself (in what might be called *autotelic visualizations*) or in casual consumption contexts (such as in *enjoy* tasks [BM13]).

3.1. What Data Can Be Used?

Understanding TACOs' potential utility starts with identifying their valid data inputs. Here we begin our discussion in earnest by analyzing the space of possible data types, values, and sizes.

3.1.1. Data Type

We first focus on understanding the types of inputs that can be appropriately visualized by TACOs. Definitionally a TACO can only be meaningfully computed on a 2D-table of scalars. Yet this still leaves a large space of potentially allowable tables.

We argue that the scalars that make up such tables must share a single unit of measurement. Inputs with heterogeneous units would multiply define the meaning of area, and hence undermine the interpretability of the output. In particular, by changing the relative definition of the potentially unrelated units, arbitrary changes to the visualization would be induced; a **Hallucinator**. A similar **Hallucinator** is found in dual-y-axis charts [MKC20], which correlate unrelated units by an arbitrary choice of normalization [Ros18].

Next, to better understand this remaining space of inputs we organize it into a typology defined by row and column data (Fig. 4). We form this space by taking all pairwise combinations of elements in a slight expansion of Stevens' discrete types [Ste46]: nominal, sorted nominal (per Inoue and Li [IL20]), and ordinal. Beyond these simple types, tables afford a large space of hierarchical organizations [Hur06, BKM13], such as nested-pivot tables. We simplify these higher-order types by classifying them as nominal or




Figure 5: TACOs intermingle data with layout. Here 2016 US regional migration [Uni16] is permuted across axial orderings. Each ordering gives the bolded columns visually different layouts, yet these transformations (α) are not meaningful: a **Hallucinator**.

ordinal by whether their hierarchical order is nominal or ordinal. This yields six transpose unique axial combinations. Finally there are non-axial orderings, which use a layout algorithm to mold non-tabular data into a grid, such as in a waffle plot (see Appendix). While this case merits further study, we focus on the more common case of tables whose meaning is described by their axes.

With this model in mind, we argue that tables featuring either an intrinsic bidirectional ordering (as in calendars) or those that are sorted are preferable to those with nominal axes. In Fig. 5 we show region to region migration in the US across row and column permutations (AVD α s), each of which have equal accuracy. Each ordering is equally valid, as this data does not have an intrinsic order. However, the visual form of the cells and highlighted columns is inconsistent across these options which reveals that this data type is a **Hallucinator**. This implies that only tables with non-nominal orderings should be used. This excludes, for instance, data found in CSVs or tidy tables whose row order is usually intended to be non-meaningful, as well as data that can be coherently expressed under multiple table projections (such as tidy's `pivot_long` and `pivot_wide` [WAB⁺19]). Further, we suggest that flat tables should be preferred as tabular hierarchies are often non-ordinal.

Slingsby et al. [SDW09] describe a set of guidelines for organizing hierarchical datasets in tree maps, arguing that nominal data should be treated with a consistent ordering to facilitate legibility. This agrees with our assertions that non-nominal data is preferable, as well Inoue and Li's [IL20] notion that sorted-by-similarity axes are more effective for TACOs. To this latter point: the validity of sorted nominals is complicated by the arbitrariness of choice of sorting algorithm —while possibly motivated by the maintenance of particular metrics or aesthetics, may yield a **Hallucinator**.

3.1.2. Data Value

Next, we continue to limit the space of possible inputs by considering appropriate cell values. A natural question to ask is whether TACOs support cell values beyond the positive scalars ascribed to them in our definition. Consider an α in which a particular value is negated. Our definition of TACOs does not provide meaning for a negative area, and as such this is undefined behavior. The result will be implementation dependent, perhaps being treated as a positive value (yielding a **Confuser**) or collapsing the planar topology. Next, consider an α in which a value is set to zero. The collapsed grid topology would then present non-adjacent cells next to each other. These collapses break our premise that we are keeping the visual form fixed. This leaves us with positive scalars.

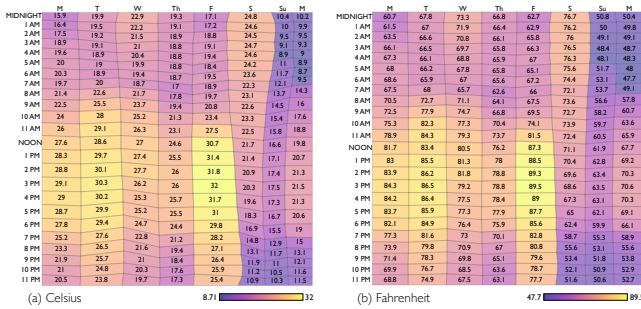


Figure 6: Temperature at O’Hare Airport Sept. 16-23, 2018 [Met]. Despite the underlying data equality, these TACOs are visually different: (a) shows the changes more dramatically. This is a **Hallucinator**, but can be tamed to serve particular tasks.

Within positive scalars there are two possible Stevens’ data types [Ste46]: ratio and interval. Area encoding has a natural root of zero (zero area indicating zero value), which pairs well with ratio encodings (which are defined by having a meaningful zero). The TACOs pliability to interval data on the other hand is less clear. Data for which an affine transform, such as $\hat{x} \rightarrow mx + b$, is merely a re-representation (such as changing between Celsius and Fahrenheit) belongs to the *interval* scale. That is, interval data are defined by having a zero whose meaning is not intrinsic to the type.

The incongruity between interval data’s non-meaningful zero and TACOs’ rooted zero would seem to indicate that TACOs should not be used with interval data. Yet, we suggest that a TACO of interval data can nonetheless be informative, depending on the task. Fig. 6 shows a set of temperatures as TACOs in different units. The differences between these charts is due to the (arbitrary) choice of unit selection. For interval data, this choice can yield a **Hallucinator** (if units are chosen to overly magnify differences in value) or a **Confuser** (as might be the case if Kelvin had been represented as well). Yet, what could be considered a critical failure mode for TACOs can also be wielded as an intentional design choice. *If the goal is to detect extrema or trends*, the Celsius units may be a better design choice, because this data is mapped to a wider variety of areas, even though there is nothing essential about areas being proportional to degrees Celsius. More extreme variations (to sub-freezing temperatures) would require a different affine transform to map the values to legible area variations, analogous to intentionally setting the axes bounds of scatter plots in a data-dependent way.

3.1.3. Data Size

All visualizations have limitations in the volume and range of values they can support while remaining comprehensible [Mun14]. Many guidelines place limits on these ranges, however there are typically exceptions or special cases to such bounds. We suggest that broadly applicable guidelines are defined with a measure of malleability or ambiguity. In this vein we construct a loose bound on appropriate data range and size.

We first look at range. Consider a table of values $x_{ij} \in \mathbb{R}^+$ with sum $\sigma = \sum_{ij} x_{ij}$. We induce a normalizing transform $\hat{x}_{ij} = x_{ij}\sigma^{-1}$. The area of a cell with no error is $a_{ij} = whx_{ij}\sigma^{-1} = wh\hat{x}_{ij}$, where




Figure 7: TACOs containing data ranges with magnitudes of 10^5 . Despite doubling of the smallest value in this toy dataset (an α), the images appear nearly identical: a **Confuser**. The visual subtlety of the difference might cause it to be overlooked.

w and h are the graphic’s height and width. We select a minimum legible area for a cell to be 1 pixel, as sizes below this are hypothetically not-representable in a pixel based system. Other reasonable values might be informed by notions such as area JND [RTAA20], however this simplistic choice serves our simplistic goals. We identify the minimum value, $\min x_{ij} = x_{min}$, for which changes would *not be invisible* (all α s yield a **Confuser**).

$$a_{min} = 1 = \hat{x}_{min}wh \implies \hat{x}_{min} = (wh)^{-1} \quad (2)$$

For a chart with $w = 500$ and $h = 500$ —as many TACOs in this paper are—then $\hat{x}_{min} = 2.5 \times 10^{-5}$. That is, TACOs with a range wider than 10^5 will possess a **Confuser** as any changes to its smallest values will be illegible, as in Fig. 7. This bound is conservative, as it describes the ratio between min and sum, and not the min and max (which would be tighter). However, as we sought a rough bound, this is sufficient. In practice we find that it is usually better to use a range smaller than 10^3 , however we do not argue that wider ranges (but $< 10^5$) are unsound. Such ranges may be legible if the data has a regular structure (as in Appendix Fig. ??) and the task is not dependent on individual values.

Appropriate row and column cardinality is more intertwined with task for TACOs than range. For instance, a table consisting of a small number of rows and columns are preferable when those individual values have relevance to the task under-consideration (such as RETRIEVE VALUE, cf. Sec. 3.2.1). Again, larger tables can be shown if the goal relates to aggregate relationships or smoothly varying distributions, since their interpretation involves trends rather than individual values. A rudimentary assertion would hold that each column or row should possess at least one pixel; suggesting a naive upper bound of $|rows| \leq h$ and $|columns| \leq w$. Yet, even this simple guideline has holes: adjacent cells of a common color with partial pixel areas can be understood as constituting larger units, allowing comparison of aggregates.

These approximations highlight a weakness in our approach: not every question can be answered by these analytic means. However, we believe that these approximations are sufficient to aid effective usage. These guidelines could be more closely examined (and made more precise) through a user study.

	Task	Description	Example Task for Fig. 12	AVD α	Confuser
Getting Values	Retrieve Value	Find a particular value	How many crossings occurred on July 15? What was the weather on Sept. 27?	Change value of interest	Yes, if label is non-numeric or value is too relatively small
	Compute Derived Value	Extract aggregate value	What fraction of bike crossings does summer account for?	Double the size of target group	Yes, if statistic requires absolute comparison
Making Comparisons	Find Extremum	Identify the largest or smallest cell	What day had the greatest number of crossings? What day had the least?	Change Extremum	No
	Sort	Order elements by an ordinal metric	Which day had the second most crossings? The third least?	Cell wise inverse ($x \rightarrow x^{-1}$)	No
	Determine Range	Find magnitude of min and max values	What are the minimum and maximum daily bike crossings?	Multiply all points by a scalar ($x \rightarrow \gamma x$)	Yes
Understanding Distributions	Characterize Distribution	Describe the distribution or identify normalcy	What parts of the week do people ride? What parts of the month or the year?	Change Distribution	Yes, if change is not in-phase with row & column partitions
	Correlate	Determine relationships between pairs of variables	Are Mondays in Jan. more popular than Fridays in Nov.? Do people ride in the rain?	Remove correlation in correlated data	Possibly, if distribution is too narrow
Finding Subgroups	Find Anomalies	Find items that are unusual	Do any particular days have especially low counts? High counts?	Induce or remove an outlier	Possibly, if rendering is noisy
	Cluster	Find items having related properties or similar size	What days have similar ridership to Mar. 3? Which months are similar?	Make all cells be similar or dissimilar	Possibly, if distribution is too narrow
	Filter	Find cells matching a particular predicate	Which days are bigger than average? Which months are smaller than average?	Change value or statistic of interest	Possibly, depending on data distribution

Task contains a Confuser. Never, Sometimes, or Always

Figure 8: The TACO’s performance on a low-level task typology [AES05] guided by whether those tasks exhibit Confusers for relevant α s.

3.2. What Tasks Can Be Used?

Next we conduct an algebraic task analysis by utilizing Amar et al.’s task typology [AES05], organized into four themes (Fig. 8). This typology describes a set of low-level functions which users might perform on visualizations. We focus on this typology (among others [BCF20, BM13]) because of its simplicity and ubiquity. We use this general typology—rather than one focused specifically on cartograms [NK15]—because it is domain agnostic, which enables comparison with non-geographic charts.

Fig. 8 summarizes this analysis, categorizing each task by whether a relevant α yields a Confuser. We focus here on Confusers as they are more closely related to task than Hallucinators [KS16]. To ground this analysis we include a specific task for Fig. 12, which is discussed in greater length in Sec. 3.3.

3.2.1. Getting Values

Tasks: RETRIEVE VALUE, COMPUTE DERIVED VALUE
Discussion: TACOs, like unadjusted tables, appear to be well-suited to the RETRIEVE VALUE, but can exhibit a Confuser if not labeled appropriately or if the relevant value is too small to visually resolve. The TACO’s tabular structure allows looking values up by rows and columns, even in the presence of distortion—although this may be impeded by layouts that are highly dissimilar to grids. Changes to these values might be invisible (yielding a Confuser) if the change is small or the cells are not labeled with their corresponding value. This suggests that, like both geographic cartograms [NK16] and tree maps [Fri94], TACOs are more effective when used in conjunction with a secondary visual encoding, such as color or text, as it facilitates easier value retrieval than simply using area alone [IL20]. We argue that, given the low accuracy of the perceptual system for understanding numeric values through area [Mac86], that value retrieval hinges on the presence of secondary annotations, and possesses a Confuser otherwise. Beyond identifying their necessity, AVD’s coarse-grained assertions offer




Figure 9: Visitorship to Zion National Park [Nat] across α s. (b) demonstrates that TACOs accurately maintain data order across an order inverting transformation, while (c) highlights a Confuser: changes to scale are invisible.

little insight into the perceptual role that secondary encodings play. Subsequent work should investigate the effect that secondary encodings have on the completion of rudimentary visualization tasks.

Less dependent on the specifics of secondary encodings is COMPUTE DERIVED VALUE which involves the visual addition of cells to form an aggregate value. There are a variety of statistics that may be evaluated, however only those involving relative values are supported by TACOs, as those related to absolute magnitudes will possess a Confuser (per Fig. 9c rescalings are invisible). This preference for relative comparison in conjunction with the availability of visual addition suggests that they support part-to-whole and part-to-part relationships, just as in pie charts. We suggest that—also similarly to pie charts [Kos10]—TACOs only support the visual summation of adjacent parts, however this is beyond the scope of AVD’s simple assertions, so we do not verify it.

3.2.2. Making Comparisons

Tasks: FIND EXTREMUM, SORT, DETERMINE RANGE
Discussion: As we have seen, areas in TACOs are proportional to data values relative to the sum of all values. This enables some types of comparisons—although those dependent on non-relative




Figure 10: Strong correlations between axes can be apparent, however it may be impaired by the multi-layout **Hallucinator**: which can suggest non-existent correlations (as in (c)’s top and bottom rows). (a)’s axes are correlated, while in (b) and (c) that correlation is removed (an α) by making the x distribution uniform.

comparison exhibit **Confusers**. The relative nature of area comparison enables identification of a consistent visual ordering of the data values, as it is typically easy to identify which of two shapes is larger [CM84], though the irregular shapes of some cells can affect this ordering, as can small differences in areas [MD73, KRK01]. It is notable that area perception of blobs (as these quadrilaterals might be interpreted) yield more accurate comparison than area comparisons in circles [CM86], which suggests that TACOs may be preferable to those encodings in some contexts.

When per-cell data values are replaced by their reciprocals (an α) all relative orderings are reversed in a consistent way, as in Fig. 9(b). This suggests that TACOs are consonant with ordinal comparisons as the rendered visual ordering of the entire visualization is consistent across transformation, and thus pliable to SORT and FIND EXTREMUM, which both rely on relative comparisons (e.g which of two cells is bigger). Yet, not all comparisons are supported. Fig. 9 (c) shows that a uniform scaling (another α) is invisible, and hence a **Confuser** for DETERMINE RANGE, which asks viewers to find absolute values of extrema. This is consistent with familiar graphics such as tree maps or pie charts, which have **Confusers** for tasks depending on *absolute* scale [KS14].

3.2.3. Understanding Distributions

Tasks: CHARACTERIZE DISTRIBUTION, CORRELATE
Discussion: Displays which privilege individual values (e.g. cells in TACOs) prompt questions about those value’s distribution. Despite their un-aggregated form TACOs may have little value in this context, as they may exhibit **Confusers** (depending on distribution).

We can probe CHARACTERIZE DISTRIBUTION with an α that changes the distribution. Changes will only be visible if they are in-phase with the partitions of the rows and columns. For instance, shifting all visits to a park (such as in Fig. 9) to happen at night will be invisible on a month calendar, but changing the number of visitors on a particular day (such as by inducing a holiday) will be visible. For interval data, CHARACTERIZE DISTRIBUTION can yield a **Hallucinator** or a **Confuser**, depending on the selection of units (Sec. 3.1.2).

To investigate CORRELATE we can employ an α that removes an extant correlation, such as by shifting to a uniform distribution, as in Fig. 10. Again, the legibility of correlations is contingent on




Figure 11: TACOs highlight outliers by distorting the grid around those cells. Here frequency of speed camera violations in Chicago during Jan. 2016 [Chi] is altered (an α) to explore this property.

their relationship being visible under the selected row and column partitions. This task also involves consideration of correlation magnitude; for which the TACO is poorly tuned, as it’s variable layouts can hallucinate weak or non-existent correlations. For instance, the top and bottom rows of Fig. 10(c) appear to be marginally larger than the rest of the rows; implying a non-uniform relationship between value and vertical position. Beyond correlation between axes, TACOs can present correlation between size and secondary encoding (such as color). We forgo considering this property as it necessitates more precise tools than are available in AVD.

3.2.4. Finding Subgroups

Tasks: FIND ANOMALIES, CLUSTER, FILTER
Discussion: The TACO’s layout forms contiguous paths in rows and columns which can aid investigation of subsets of interest. The fabric of the table itself is distorted in a coherent manner across contiguous paths for all rows and columns. We suggest that this distortion facilitates identification of anomalous extrema (FIND ANOMALIES), while impeding identification of similarly valued cells or groups of cells—as FILTER and CLUSTER require.

We can observe TACOs handling of FIND ANOMALIES in Fig. 11, which removes outliers from smoothly distributed data (an α) by setting them to the weekly average. The visibility of this change indicates TACOs are not ineffective for tasks involving identifying outliers. Those now-average values are more difficult to differentiate from their neighbors. This can thus impede both FILTER and CLUSTER for narrow distributions. This is in agreement with prior work [MD73, KRK01] which notes that very different comparisons are easier to make than similar ones.

A notable caveat: the TACOs’ irregularity can cause false positives when trying to detect outliers as cells which may have the same value can be represented as quadrilaterals with greatly differing shapes. Yet, under an appropriate parameter configuration outliers can be clearly seen (Fig. 11). Unlike FIND ANOMALIES, the ease of FILTER or CLUSTER seems to be a function of both the separability of subclasses within the distribution, as well as the selection of axis units or partitions. As in other cases, this suggests that there is a **Confuser** when the distribution is narrow, or if the statistic that FILTER is predicated upon possesses a **Confuser**.

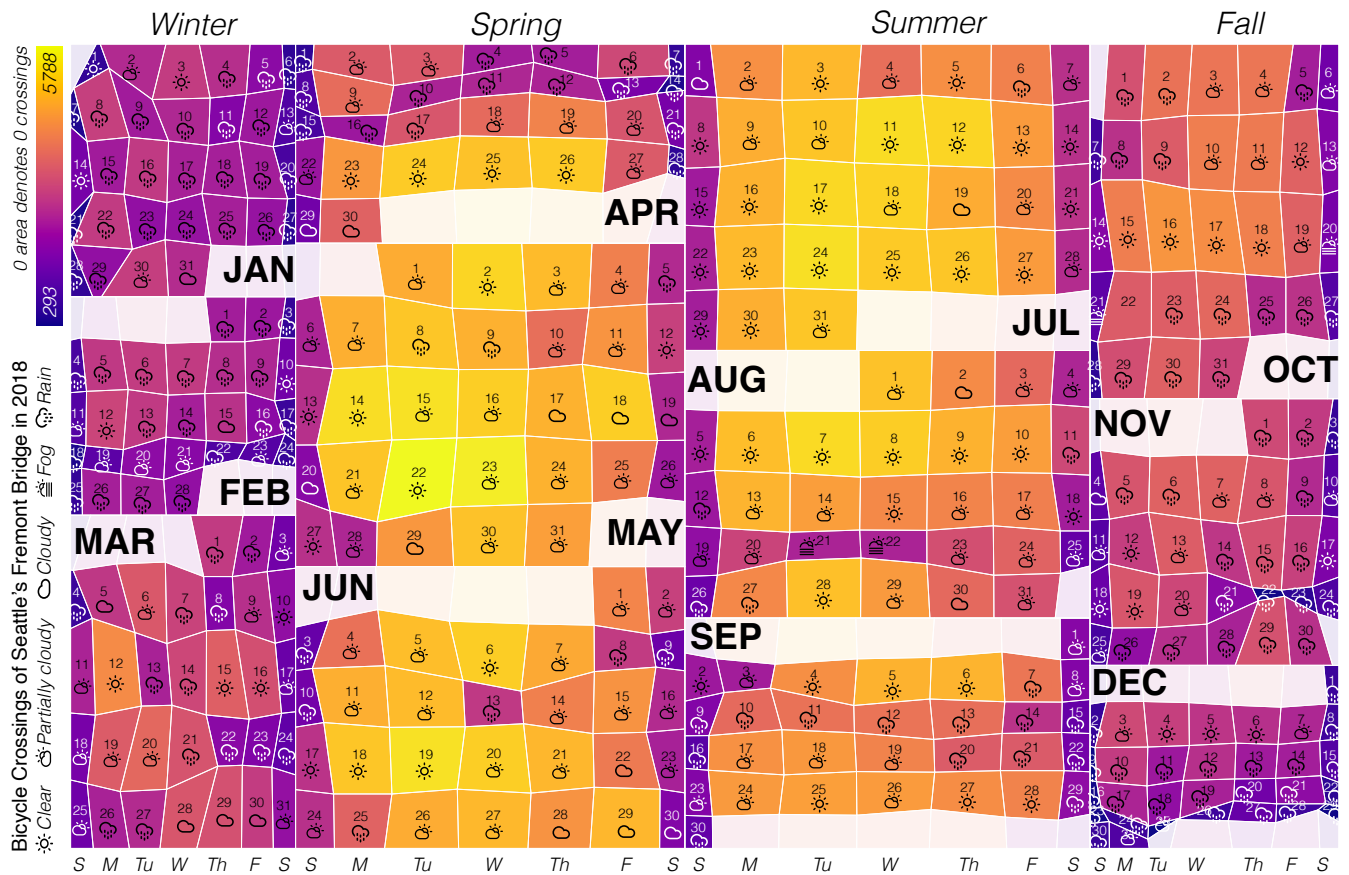


Figure 12: Bicycle crossings of the Fremont Bridge in Seattle in 2018 [McN19]. Months of the year are arranged by quarter in a slice-dice tree map, while days of the month are shown as TACOs such that each day is sized and colored according to the crossing volume.

3.3. Case Study

We conclude our analysis with a discussion of a larger example TACO [McN19], reproduced in Fig. 12. We focus on this design because it exemplifies a more complex composition than the examples we have seen so far, thereby showing more of the TACO’s strengths, as well as its weaknesses. It shows the bicycle crossings on a particular bridge in 2018 annotated with daily weather. The TACO is used here, not just as a graphic unto itself, but as an ingredient in a larger composition, showing each of the months as an independent TACO arranged in a slice-dice tree map. Crossing volume is doubly encoded through area and color. The area encoding provides a ratio representation of the value (zero area corresponding to zero crossings), while color shows that quantity through an interval rendering, facilitating ordinal readings—like that yellow cells have a greater number of crossings than others. This design highlights the TACO’s ability to blend with other chart forms, such as tree maps. In the appendix we discuss another real-world example [pie], which combines TACOs and word clouds.

Some tasks are eased by the visual structures of this graphic. As might be expected from our discussion of FIND ANOMALIES, some outliers are readily visible, such as near American Thanksgiving (Nov. 22), as well as seasonal and monthly trends (which are imparted by the containing tree map). Some correlations are

more readily apparent than others. Weather condition appears to be loosely correlated with increases in ridership (as in the third week of October)—although as we noted in Sec. 3.2.3 magnitude of correlation is hard to judge with TACOs. In contrast, while it is clear that more crossings occur during weekdays than on weekends, it is completely invisible that more crossings occur during the middle of the day on the weekends, as well as before and after business hours on weekdays (suggesting commuters). This graphic values ordinal comparisons over magnitude judgments: while values can be looked up, doing so requires utilizing color, as magnitude is illegible as area, and text is used to describe the date and weather. This is inline with our comments about the data, the types that can be appropriately described (ratio), and the tasks facilitated therein.

The purposes of this graphic might be equally well served by a number of alternatives. For instance, the TACOs could be simply replaced by shaded calendars, rather than being dual-encoded with area, the entire plot might be replaced with a line series, or with something specifically tuned to calendar displays [VWVS99, Har08]. Yet the novelty, and hence possibly the visual appeal, imparted to this graphic by the use of an unusual and somewhat perceptually difficult chart form may be diminished through the use of more pedestrian graphics. This is emblematic of the TACOs general use case: as a mechanism for supporting entertainment and engagement, and not as a tool for decision making.

4. Discussion

We now synthesize our findings, discuss our methodology, highlight limitations, and note opportunities for future work.

4.1. What are Table Cartograms Good For?

TACOs offer an intriguing combination of affordances. We argued that they support SORT and FIND EXTREMUM tasks, as well as FIND ANOMALIES and distribution tasks under some conditions. We saw that they can be effective for RETRIEVE VALUE tasks, which can be obscured if the cardinality of rows and columns is too large (suggesting that tables with a small axial cardinality are preferable). We claimed that they are best suited to tables with ordinal axes with a limited range and typically with ratio data. While not every insight in our analysis will be surprising to those familiar with geographic cartograms or tree maps, the way in which we arrived at them (through the consideration of the TACO's symmetries) facilitated a broad and self-consistent description of those conclusions, as well as a pragmatically-organized account of relevant guidelines. Yet, as we saw, TACOs are not without deficits.

Weaknesses. TACOs possess a prominent collection of weaknesses, which when coupled with their unusual and inconsistent visual form, severely limits their applicability. Their primary method of encoding data (as area) is far from the most accurate perceptual channel for encoding quantitative information [Mac86]. They face all of the same troubles as geographic cartograms, including that they are difficult to interpret, and possess area perception challenges [Nus17]. TACOs add further perceptual difficulties as they allow convex quadrilaterals just as readily as concave ones, such that cells that have a common area may not appear to be identical. While their *Hallucinators* allow a great degree of aesthetic control, they do not provide a consistent visual bedrock upon which to conduct analysis, suggesting that they are ineffective for exploring data. They are unsuitable for many tasks and datasets, such as comparisons of absolute values and datasets with wide value ranges. In most contexts, there are an array of alternatives which may be more effective, such as tables, shaded matrices, mosaics, and tree maps. These forms often offer preferable variations of the TACO's strengths. For instance, TACOs support FIND ANOMALIES, yet tree maps facilitate this task as well [Mun14] while affording the easier area comparisons. This suggests that effective usage may involve leveraging TACOs' more unusual maintenance of adjacency. However, in most contexts adjacency is not an important property to preserve at the expense of legibility.

Potential Opportunities. Despite these weaknesses, TACOs may yet be useful in an appropriate context. Visual complexity can increase viewer engagement or enjoyment [Lim11, HRC15]. Hullman et al. [HAS11] argue that graphics that are visually difficult can be usefully deployed as a design element to prompt engagement. Kosara [Kos16] suggests that presentation-only tasks can leverage aesthetic appeal for reader engagement, as with pie charts. The atypical shapes assumed by TACOs may help capture readers' attention, although further study is required to evaluate this assertion. To this end, they may be effective at capturing *enjoy* tasks [BM13], which is inline with the advocacy and education applications that geographic cartograms typically serve [NK16, Tob04].

The naturally bidirectional ordinal domain of many displays of time offers a compelling application for TACOs. In Sec. 3.1 we argued that only flat tables with ordinal data should be used as inputs to TACOs, which are well matched displays of time, just as in calendars. The ordered nature of calendars allows TACOs to highlight temporal anomalies and trends, such as in Fig. 11 and Sec. 3.3. This is inline with Drucker's call for "graphical displays that emphasize the relational and co-dependent quality of temporal events" [Dru11], as TACOs intermingle the structure of space with data. However, even this algebraically-sound application should be handled with care, as month calendars have a small *Hallucinator* due to the ambiguity of which day of the week is first.

Future work. Several questions remain about TACOs, including the human response to them and how our recommendations perform in real contexts, each of which would benefit from a user study. We ignored questions of interactivity (a limitation imposed by AVD), but future work might explore how animation and interactivity expand the design space—an active area or research for geographic cartograms [DTPG21]. Many of the geographic cartogram's ills can be addressed through "good design choices" [NK16], such as legends and annotations. The same may be true for TACOs, although further design exploration is necessary to understand the role these components could play.

4.2. How Might We Examine Unknown Visual Forms?

Understanding appropriate usage for novel visualizations is an intrinsically difficult, yet important task. Previous works [ZC06] argued for theory-based analysis methods which would allow for the discount evaluation of graphics that have not become sufficiently developed or well-known to prompt user studies. Such economical methods have been shown to have great value in other contexts, such as usability testing [BKM09] or evaluation [Gre89].

In this work we exemplified one such method of abstract evaluation based on the language of *Algebraic Visualization Design*, and, in doing so, we showed that practical guidelines can be generated through application of theory. AVD provided particular value in this pursuit, as it offered a systematic way to ask and answer questions about the visual form of interest. While not able to address every concern—most matters related to perceptual quality being too subtle for the coarse way in which we applied AVD—it was able to help us build a coherent set of usage suggestions. This approach inverts AVD's typical usage as guide in the design process, wherein designs are invalidated for a particular task [CLKS18, CC21, PK20, WKD18]. Similar analyses to ours have been used to consider novel graphics [CH16, WD08], although in an ad hoc and implicit way. We formalize this approach as *Algebraic Visualization Analysis* (AVA).

AVA addresses two key questions (Fig. 3). The first—*What data can be used?*—is answered by developing a model of possible inputs (across type, structure, and size), then pruning that space via adversarial examples, whose effects were considered through the lens of AVD's *Confusers* and *Hallucinators*. Our model of possible inputs was dictated by the form of data that TACOs accept. Analysis of another graphic, such as a pie chart, would require a model incorporating input data, aggregation, and other relevant

parameters. We used a variety of tools in this trimming (including Stevens' data types [Ste46] and JNDs), which were selected to probe the TACO's particular properties. Analysis of a different graphic may necessitate a different selection of tools. For instance, analysis of a pie chart may involve a calculation of the minimum perceptible arc area, as well as consideration of the effect of affine-versus-ratio value transformations.

The second question—*What tasks can be performed?*—is addressed by again segmenting a possibility space, this time using Amar et al.'s [AES05] task taxonomy and then constructing relevant α s for each task. AVA is composed of various interlocking theories, many of which could be swapped, should the need arise. For instance, this taxonomy could be exchanged for another—although we favor this one for its brevity. Analysis of a different chart form would involve identifying relevant α s for each task and visually evaluating them. Future work might develop tools that given a formal specification of allowed data derive α s pre-matched to tasks.

Following Hu et al.'s [HGH*19] triplet model of visualization a third potential question might ask *What visualization is appropriate?* However, as selection of a visual form is prior to AVA, this question might be refined to *How do this visual form's non-data parameters affect its usage?* This arises in our discussion of TACO's multiplicity of solutions for a given dataset. Such issues can be answered with AVD, such as in Correll et al.'s [CLKS18] study of the effect that bin-widths have on histograms. These matters are prior to AVA's key questions, as developing an understanding of what forms a chart might take on is vital to probing its usage.

The goal of answering these questions (*providing chart making guidance*) is one shared by a variety of mediums, such as guidelines [DKA*20], validations [MK18, HCS20], prompts for introspection [DFCC13, WKD18], and recommendations [Lee20]. We believe that advances in any of these interwoven modalities should lead to advances in the others, and that AVD is an opportune foundation on which to unify them, as it is both human and machine [CC21, VC19] operable. Yet, such research directions should be followed with caution as both AVA and AVD have limitations.

AVA Limitations. As with any theory, AVA has a set of limitations and imprecisions. Any investigation conducted by a single individual will likely exhibit bias. While this is problematic, the nature of our approach is somewhat self-regulating. This approach does not verify that a given visual form is good for a certain task; rather, it only invalidates particular tasks or inputs for that form. We believe that AVA's non-existence proofs offer a valuable picture of the graphic under consideration while usefully limiting the scope of claims that can be made. These assertions can be checked by simply comparing outputs across α s. AVA's iterated partitioning can inadvertently ignore errors outside of its framing. However we do not claim this analysis to be total, instead we claim that it upper bounds performance. Evocatively: we see a graphic as being as bad as AVA describes it to be, or worse. In addition to the *Confusers* and *Hallucinators*, AVD has two further failure modes: *Jumblers* and *Misleadings*, which test whether an α *appropriately* corresponds to an ω . While they can provide useful insights, they rely on more observer-dependent judgements than the modes we considered.

Visualizations are situated in their context and their value is de-

pendent on how they perform at tasks in those situations. In our analysis we considered a visual form devoid of such a context and reduced it to its most elementary components. Our hope is that by considering the form in the abstract, that the rendered guidelines will be generally applicable. Yet, this reductive approach may cause our suggestions to be invalid in some contexts. In future work we will compare it with other analysis methods for graphics whose properties are already well known, so as to develop a deeper understanding of what can and cannot be achieved with AVA.

We encountered phenomena in our analysis that could not be addressed through our limited set of tools. As noted in Sec. 3.1, while coherent guidelines can be produced using this framing, the precision of such recommendations may be limited. Echoing Chen et al. [CJ10], we believe that user studies are necessary for the rigorous study of visualizations (and therein construction of guidance), and that theory cannot stand alone. The dual of this statement—that user studies must be built atop theory—does not hold, as valid experiments can be conducted without an overarching theoretical foundation. That said, we believe that theory-based analysis has value. Beyond its demonstrated value for discount evaluations, it can, for instance, support hypothesis generation, the result of which can be used as the basis for experiments.

AVD Limitations. While AVD can be an advantageous framework on which to base analyses, its structural limitations can impede some analyses. The assertions it makes are coarse-grained, which—while being readily interpretable—can fail to provide an explanation for subtler phenomena (such as dual encodings). It can demonstrate that tasks *can be achieved*, yet it does not offer any explanation on *how easily* those tasks might be under-taken. This is analogous to how heuristic evaluation can highlight problems, but cannot suggest repairs [NM90]. It is unable to validate interactive or time varying graphics (such as HOP plots [KNKH18]). It cannot reason about design components that are not purely graphical, such as embodied affect [DK16] (as in chartjunk [BMG*10]) or value-sensitivity. Schwabish and Feng describe how a line chart for a race-based COVID chart gave way to deficiency framings of those races most affected [SF20]. While AVD might consider such an issue by generating α s related to race, judgement of the rendered change is not limited purely to the visual spectrum as other α s tend to be; instead it relies on the composition of biases in the viewer for validation. This may prevent the analyst from having sufficient distance to judge the sensitivity of their construction. Constructing analysis frameworks that are value-sensitive is an important task, and should be a component of future work. While there has been some examination of the connection between AVD's failure modes and human perception [CLKS18], questions remain. To accept a theory as a basis of analysis it is important to demonstrate that it has a clear connection with the real world. AVD's explanatory power seen here and elsewhere, suggests it can validly answer questions, although experimental data is still outstanding.

5. Acknowledgments

We thank our reviewers for constructive feedback. We also thank Will Brackenbury, Ravi Chugh, Brian Hempel, Gordon Kindmann, as well as the UChicago Vis Reading group for their insightful comments, thoughtful discussion, and splendid support.

References

- [AES05] AMAR R., EAGAN J., STASKO J.: Low-Level Components of Analytic Activity in Information Visualization. In *IEEE Symposium on Information Visualization* (2005), IEEE, pp. 111–117. doi:10.1109/INFVIS.2005.1532136. 6, 10
- [AKV15] ALAM M. J., KOBOUROV S. G., VEERAMONI S.: Quantitative Measures for Cartogram Generation Techniques. In *Computer Graphics Forum* (2015), vol. 34, Wiley Online Library, pp. 351–360. doi:10.1111/cgf.12647. 2
- [AL20] ADAR E., LEE E.: Communicative Visualizations as a Learning Problem. *IEEE Transactions on Visualization and Computer Graphics* (2020). 1, 2
- [Ans73] ANSCOMBE F.: Graphs in statistical analysis. *The American Statistician* 27, 1 (1973), 17–21. 3
- [BBD20] BRÜGGEMANN V., BLUDAU M.-J., DÖRK M.: The Fold: Rethinking Interactivity in Data Visualization. *Digital Humanities Quarterly* 14, 3 (2020). 2
- [BBK*18] BEHRISCH M., BLUMENSCHNEIN M., KIM N. W., SHAO L., EL-ASSADY M., FUCHS J., SEEBACHER D., DIEHL A., BRANDES U., PFISTER H., ET AL.: Quality Metrics for Information Visualization. In *Computer Graphics Forum* (2018), vol. 37, Wiley Online Library, pp. 625–662. doi:10.1111/cgf.13446. 2
- [BCF20] BERTINI E., CORRELL M., FRANCONERI S.: Why Shouldn't All Charts Be Scatter Plots? Beyond Precision-Driven Visualizations. In *IEEE Visualization Conference (VIS)* (2020), IEEE, pp. 1–5. doi:10.1109/VIS47514.2020.00048. 2, 6
- [BKM09] BANGOR A., KORTUM P., MILLER J.: Determining What Individual SUS Scores Mean: Adding an Adjective Rating Scale. *Journal of usability studies* 4, 3 (2009), 114–123. 9
- [BKM13] BAKKE E., KARGER D. R., MILLER R. C.: Automatic Layout of Structured Hierarchical Reports. *IEEE Transactions on Visualization and Computer Graphics* 19, 12 (2013), 2586–2595. doi:10.1109/TVCG.2013.137. 4
- [BM13] BREHMER M., MUNZNER T.: A Multi-Level Typology of Abstract Visualization Tasks. *IEEE Transactions on Visualization and Computer Graphics* 19, 12 (2013), 2376–2385. doi:10.1109/TVCG.2013.124. 4, 6, 9
- [BMG*10] BATEMAN S., MANDRYK R. L., GUTWIN C., GENEST A., MCDINE D., BROOKS C.: Useful Junk? The Effects of Visual Embellishment on Comprehension and Memorability of Charts. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (2010), pp. 2573–2582. doi:10.1145/1753326.1753716. 10
- [BZJ*20] BARES A., ZELLER S., JACKSON C. D., KEEFE D. F., SAMSEL F.: Using close reading as a method for evaluating visualizations. In *IEEE Workshop on Evaluation and Beyond - Methodological Approaches to Visualization, 2020* (2020), IEEE, pp. 29–37. doi:10.1109/BELIV51497.2020.00011. 2
- [CC21] CRISAN A., CORRELL M.: User Ex Machina: Simulation as a Design Probe in Human-in-the-Loop Text Analytics. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (2021), pp. 1–16. To Appear. 9, 10
- [CG15] CHEN M., GOLAN A.: What May Visualization Processes Optimize? *IEEE Transactions on Visualization and Computer Graphics* 22, 12 (2015), 2619–2632. doi:10.1109/TVCG.2015.2513410. 2
- [CH16] CORRELL M., HEER J.: Surprise! Bayesian Weighting for De-Biasing Thematic Maps. *IEEE Transactions on Visualization and Computer Graphics* 23, 1 (2016), 651–660. doi:10.1109/TVCG.2016.2598618. 9
- [CH17] CORRELL M., HEER J.: Black Hat Visualization. In *DECISive: Workshop on Dealing with Cognitive Biases in Visualisations* (2017). 1
- [Chi] CHICAGO DATA PORTAL: Speed Camera Violations. <https://data.cityofchicago.org/Transportation/Speed-Camera-Violations/hhk4-xvj4>. Accessed 09/23/2018. 7
- [CJ10] CHEN M., JAENICKE H.: An Information-theoretic Framework for Visualization. *IEEE Transactions on Visualization and Computer Graphics* 16, 6 (2010), 1206–1215. doi:10.1109/TVCG.2010.132. 2, 10
- [CLKS18] CORRELL M., LI M., KINDLMANN G., SCHEIDEGGER C.: Looks Good To Me: Visualizations As Sanity Checks. *IEEE Transactions on Visualization and Computer Graphics* 25, 1 (2018), 830–839. doi:10.1109/TVCG.2018.2864907. 9, 10
- [CM84] CLEVELAND W. S., MCGILL R.: Graphical Perception: Theory, Experimentation, and Application to the Development of Graphical Methods. *Journal of the American Statistical Association* (1984). 7
- [CM86] CLEVELAND W. S., MCGILL R.: An Experiment in Graphical Perception. *International Journal of Man-Machine Studies* 25, 5 (1986), 491–500. 7
- [CM07] CAWTHON N., MOERE A. V.: The Effect of Aesthetic on the Usability of Data Visualization. In *11th International Conference Information Visualization* (2007), IEEE, pp. 637–648. doi:10.1109/IV.2007.147. 1
- [DFCC13] DÖRK M., FENG P., COLLINS C., CARPENDALE S.: Critical InfoVis: exploring the politics of visualization. In *Extended Abstracts on Human Factors in Computing Systems*. ACM, 2013, pp. 2189–2198. doi:10.1145/2468356.2468739. 2, 10
- [DK16] D'IGNAZIO C., KLEIN L. F.: Feminist Data Visualization. In *Workshop on Visualization for the Digital Humanities (VIS4DH)*, Baltimore. IEEE (2016). 2, 10
- [DKA*20] DIEHL A., KRAUS M., ABDUL-RAHMAN A., EL-ASSADY M., BACH B., LARAMEE R. S., KEIM D. A., CHEN M.: Studying Visualization Guidelines According to Grounded Theory. URL: <https://arxiv.org/abs/2010.09040>, arXiv:2010.09040. 10
- [Dru11] DRUCKER J.: Humanities Approaches to Graphical Display. *Digital Humanities Quarterly* 5, 1 (2011), 1–21. 9
- [DSK*14] DEMIRALP Ç., SCHEIDEGGER C. E., KINDLMANN G. L., LAIDLAW D. H., HEER J.: Visual Embedding: A Model for Visualization. *IEEE Computer Graphics and Applications* 34, 1 (2014), 10–15. doi:10.1109/MCG.2014.18. 2
- [DTPG21] DUNCAN I. K., TINGSHENG S., PERRAULT S. T., GASTNER M. T.: Task-Based Effectiveness of Interactive Contiguous Area Cartograms. *IEEE Transactions on Visualization and Computer Graphics* 27, 3 (2021), 2136–2152. doi:10.1109/TVCG.2020.3041745. 9
- [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
- [EvKSS15] EPPSTEIN D., VAN KREVELD M., SPECKMANN B., STAALS F.: Improved Grid Map Layout by Point Set Matching. *International Journal of Computational Geometry & Applications* 25, 02 (2015), 101–122. 2
- [Fri94] FRIENDLY M.: Mosaic Displays for Multi-Way Contingency Tables. *Journal of the American Statistical Association* 89, 425 (1994), 190–200. 6
- [Gre89] GREEN T. R.: Cognitive Dimensions of Notations. *People and Computers V* (1989), 443–460. 9
- [Har08] HARTL P. R.: *Visualization of Calendar Data*. PhD thesis, Vienna University of Technology, 2008. 8
- [HAS11] HULLMAN J., ADAR E., SHAH P.: Benefitting Infovis with Visual Difficulties. *IEEE Transactions on Visualization and Computer Graphics* (2011). doi:10.1109/TVCG.2011.175. 9
- [HCS20] HOPKINS A. K., CORRELL M., SATYANARAYAN A.: VisualLint: Sketchy In Situ Annotations of Chart Construction Errors. In *Computer Graphics Forum* (2020), vol. 39, Wiley Online Library, pp. 219–228. doi:10.1111/cgf.13975. 10
- [HGH*19] HU K., GAIKWAD S., HULSEBOS M., BAKKER M. A., ZGRAGGEN E., HIDALGO C., KRASKA T., LI G., SATYANARAYAN

- A., DEMIRALP Ç.: VizNet: Towards a Large-Scale Visualization Learning and Benchmarking Repository. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (2019), pp. 1–12. doi:10.1145/3290605.3300892. 3, 10
- [HRC15] HARRISON L., REINECKE K., CHANG R.: Infographic Aesthetics: Designing for the First Impression. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (2015). doi:10.1145/2702123.2702545. 9
- [Hur06] HURST M.: Towards a Theory of Tables. *International Journal of Document Analysis and Recognition* 8, 2-3 (2006), 123–131. doi:10.1007/s10032-006-0016-y. 4
- [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, 3, 4, 6
- [KNKH18] KALE A., NGUYEN F., KAY M., HULLMAN J.: Hypothetical Outcome Plots Help Untrained Observers Judge Trends in Ambiguous Data. *IEEE Transactions on Visualization and Computer Graphics* 25, 1 (2018), 892–902. doi:10.1109/TVCG.2018.2864909. 10
- [Kos10] KOSARA R.: Understanding Pie Charts. <https://eagereyes.org/techniques/pie-charts>, 2010. Accessed 11/2/2020. 6
- [Kos16] KOSARA R.: Presentation-Oriented Visualization Techniques. *IEEE Computer Graphics and Applications* 36, 1 (2016), 80–85. doi:10.1109/MCG.2016.2. 9
- [KRK01] KRIDER R. E., RAGHUBIR P., KRISHNA A.: Pizzas: π or Square? Psychophysical Biases in Area Comparisons. *Marketing Science* 20, 4 (2001), 405–425. doi:10.1287/mksc.20.4.405.9756. 7
- [KS14] KINDLMANN G., SCHEIDEGGER C.: An Algebraic Process for Visualization Design. *IEEE Transactions on Visualization and Computer Graphics* 20, 12 (2014), 2181–2190. doi:10.1109/TVCG.2014.2346325. 1, 3, 7
- [KS16] KINDLMANN G., SCHEIDEGGER C.: Algebraic Visualization Design for Pedagogy. IEEE VIS Workshop on Pedagogy of Data Visualization, Oct. 2016. 6
- [Lee20] LEE D. J.-L.: Insight Machines: The Past, Present, and Future of Visualization Recommendation. <https://link.medium.com/fM5VhrSL5db>, Feb. 2020. Multiple Views: Visualization Research Explained. Accessed 2/21/2021. 2, 10
- [LI19] LI M., INOUE R.: Table Cartogram Generation as an Optimization Problem. *Abstracts of the International Cartographic Association 1* (2019), 1–2. 2
- [Lim11] LIMA M.: *Visual Complexity: Mapping Patterns of Information*. Princeton Architectural Press, 2011. 9
- [Mac86] MACKINLAY J.: Automating the Design of Graphical Presentations of Relational Information. *ACM Transactions On Graphics* (1986). doi:10.1145/22949.22950. 2, 6, 9
- [MCC20] MCNUTT A., CRISAN A., CORRELL M.: Divining insights: Visual analytics through cartomancy. In *Extended Abstracts of the CHI Conference on Human Factors in Computing Systems* (2020), pp. 1–16. doi:10.1145/3334480.3381814. 1
- [McN19] MCNUTT A.: Cycles Rain Seasons In Size. <https://www.informationisbeautifulawards.com/showcase/4499-cycles-rain-seasons-in-size>, 2019. Accessed 11/2/2020. 8
- [MD73] MARTINEZ N., DAWSON W. E.: Ranking of Apparent Area for Different Shapes of Equal Area. *Perceptual and Motor Skills* (1973). doi:10.1177/003151257303700319. 7
- [Met] METEOBLUE: O’hare International Airport Weather. www.meteoblue.com/en/weather/archive/export/chicago-o%27hare-international-airport-united-states-of-america_4887479. Accessed 09/23/2018. 5
- [MK18] MCNUTT A., KINDLMANN G.: Linting for Visualization: Towards a Practical Automated Visualization Guidance System. In *Vis-Guides: 2nd Workshop on the Creation, Curation, Critique and Conditioning of Principles and Guidelines in Visualization* (2018). 10
- [MK20] MCNUTT A., KINDLMANN G.: A Minimally Constrained Optimization Algorithm for Table Cartograms. *IEEEVIS InfoVis Posters* (2020). 1, 2, 3
- [MKC20] MCNUTT A. M., KINDLMANN G. L., CORRELL M.: Surfacing Visualization Mirages. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (2020), ACM, pp. 1–16. doi:10.1145/3313831.3376420. 4
- [Mun14] MUNZNER T.: *Visualization Analysis and Design*. AK Peters/CRC Press, 2014. 5, 9
- [Nat] NATIONAL PARK SERVICE: Zion Park Visitation Statistics. <https://www.nps.gov/zion/learn/management/park-visitation-statistics.htm>. Accessed 10/04/2018. 6
- [NK15] NUSRAT S., KOBOUROV S.: Task Taxonomy for Cartograms. In *Proceedings of the Eurographics / IEEE VGTC Conference on Visualization: Short Papers* (2015). doi:10.2312/eurovisshort.20151126. 6
- [NK16] NUSRAT S., KOBOUROV S.: The State of the Art in Cartograms. In *Computer Graphics Forum* (2016), vol. 35, pp. 619–642. doi:10.1111/cgf.12932. 2, 4, 6, 9
- [NM90] NIELSEN J., MOLICH R.: Heuristic Evaluation of User Interfaces. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (1990), ACM, pp. 249–256. doi:10.1145/97243.97281. 10
- [Nus17] NUSRAT S.: *Cartogram Visualization: Methods, Applications, and Effectiveness*. PhD thesis, The University of Arizona, 2017. 9
- [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. 8
- [PK20] PU X., KAY M.: A Probabilistic Grammar of Graphics. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (2020), pp. 1–13. doi:10.1145/3313831.3376466. 9
- [Ros18] ROST L. C.: Why not to use two axes, and what to use instead. <https://blog.datawrapper.de/dualaxis/>, 2018. Accessed 11/2/20. 4
- [RTAA20] RAIDVEE A., TOOM M., AVERIN K., ALLIK J.: Perception of means, sums, and areas. *Attention, Perception, & Psychophysics* (2020), 1–12. doi:10.3758/s13414-019-01938-7. 5
- [SDW09] SLINGSBY A., DYKES J., WOOD J.: Configuring Hierarchical Layouts to Address Research Questions. *IEEE Transactions on Visualization and Computer Graphics* 15, 6 (2009), 977–984. doi:10.1109/TVCG.2009.128. 4
- [SF20] SCHWABISH J., FENG A.: Applying Racial Equity Awareness in Data Visualization. *Visualization for Communication* (2020). 10
- [Shn96] SHNEIDERMAN B.: The Eyes Have It: A Task by Data Type Taxonomy for Information Visualizations. In *Visual Languages, 1996. Proceedings., IEEE Symposium on* (1996), IEEE, pp. 336–343. doi:10.1109/VL.1996.545307. 2
- [SLD20] SCHEIBEL W., LIMBERGER D., DÖLLNER J.: Survey of treemap layout algorithms. In *Proceedings of the 13th International Symposium on Visual Information Communication and Interaction* (2020), ACM, pp. 1:1–1:9. doi:10.1145/3430036.3430041. 3
- [Ste46] STEVENS S. S.: On the Theory of Scales of Measurement. *Science* (1946). 4, 5, 10
- [Tob04] TOBLER W.: Thirty Five Years of Computer Cartograms. *Annals of the Association of American Geographers* 94, 1 (2004), 58–73. doi:10.1111/j.1467-8306.2004.09401004.x. 9
- [Uni16] UNITED STATES CENSUS BUREAU: State-to-State Migration Flows. <https://www.census.gov/data/tables/time-series/demo/geographic-mobility/state-to-state-migration.html/>, 2016. Accessed 1/16/2019. 4

- [VC19] VERAS R., COLLINS C.: Discriminability Tests for Visualization Effectiveness and Scalability. *IEEE Transactions on Visualization and Computer Graphics* 26, 1 (2019), 749–758. doi:10.1109/TVCG.2019.2934432. 10
- [VDBBC*18] VAN DEN BERG H., BETTI A., CASTERMANS T., KOOPMAN R., SPECKMANN B., VERBEEK K., VAN DER WERF T., WANG S., WESTENBERG M. A., ET AL.: A Philosophical Perspective on Visualization for Digital Humanities. In *Visualization for the Digital Humanities (VIS4DH)* (2018). 1
- [VFR12] VICKERS P., FAITH J., ROSSITER N.: Understanding Visualization: A Formal Approach Using Category Theory and Semiotics. *IEEE Transactions on Visualization and Computer Graphics* 19, 6 (2012), 1048–1061. doi:10.1109/TVCG.2012.294. 2
- [VWVS99] VAN WIJK J. J., VAN SELOW E. R.: Cluster and Calendar Based Visualization of Time Series Data. In *IEEE Symposium on Information Visualization* (1999), IEEE, pp. 4–9. doi:10.1109/INFVIS.1999.801851. 8
- [WAB*19] WICKHAM H., AVERICK M., BRYAN J., CHANG W., MCGOWAN L. D., FRANÇOIS R., GROLEMUND G., HAYES A., HENRY L., HESTER J., KUHN M., PEDERSEN T. L., MILLER E., BACHE S. M., MÜLLER K., OOMS J., ROBINSON D., SEIDEL D. P., SPINU V., TAKAHASHI K., VAUGHAN D., WILKE C., WOO K., YUTANI H.: Welcome to the tidyverse. *Journal of Open Source Software* 4, 43 (2019), 1686. doi:10.21105/joss.01686. 4
- [WAM*18] WALL E., AGNIHOTRI M., MATZEN L., DIVIS K., HAASS M., ENDERT A., STASKO J.: A heuristic approach to value-driven evaluation of visualizations. *IEEE Transactions on Visualization and Computer Graphics* 25, 1 (2018), 491–500. doi:10.1109/TVCG.2018.2865146. 2
- [WCHB10] WICKHAM H., COOK D., HOFMANN H., BUJA A.: Graphical Inference for Infovis. *IEEE Transactions on Visualization and Computer Graphics* 16, 6 (2010), 973–979. doi:10.1109/TVCG.2010.161. 3
- [WD08] WOOD J., DYKES J.: Spatially Ordered Treemaps. *IEEE Transactions on Visualization and Computer Graphics* 14, 6 (2008). doi:10.1109/TVCG.2008.165. 9
- [Wic13] WICKHAM H.: Graphical Criticism: Some Historical Notes. *Journal of Computational and Graphical Statistics* 22, 1 (2013), 38–44. doi:10.1080/10618600.2012.761140. 1
- [WKD18] WOOD J., KACHKAEV A., DYKES J.: Design Exposition with Literate Visualization. *IEEE Transactions on Visualization and Computer Graphics* 25, 1 (2018), 759–768. doi:10.1109/TVCG.2018.2864836. 3, 9, 10
- [ZC06] ZUK T., CARPENDALE S.: Theoretical Analysis of Uncertainty Visualizations. In *Visualization and Data Analysis* (2006), vol. 6060, International Society for Optics and Photonics, p. 606007. doi:10.1117/12.643631. 1, 2, 9
- [ZSN*06] ZUK T., SCHLESIER L., NEUMANN P., HANCOCK M. S., CARPENDALE S.: Heuristics for information visualization evaluation. In *Proceedings of the Workshop on Beyond Time and Errors: Novel Evaluation Methods for Visualization, BELIV* (2006), ACM Press, pp. 1–6. doi:10.1145/1168149.1168162. 1