

An Empire Built On Sand: Reexamining What We Think We Know About Visualization

Robert Kosara
Tableau Research
rkosara@tableau.com

ABSTRACT

If we were to design Information Visualization from scratch, we would start with the basics: understand the principles of perception, test how they apply to different data encodings, build up those encodings to see if the principles still apply, etc. Instead, visualization was created from the other end: by building visual displays without an idea of how or if they worked, and then finding the relevant perceptual and other basics here and there.

This approach has the problem that we end up with a very patchy understanding of the foundations of our field. More than that, there is a good amount of unproven assumptions, aesthetic judgments, etc. mixed in with the evidence. We often don't even realize how much we rely on the latter, and can't easily identify them because they have been so deeply incorporated into the fabric of our field.

In this paper, I attempt to tease apart what we know and what we only think we know, using a few examples. The goal is to point out specific gaps in our knowledge, and to encourage researchers in the field to start questioning the underlying assumptions. Some of them are probably sound and will hold up to scrutiny. But some of them will not. We need to find out which is which and systematically build up a better foundation for our field. If we intend to develop ever more and better techniques and systems, we can't keep ignoring the base, or it will all come tumbling down sooner or later.

1. INTRODUCTION

Everyone knows that pie charts are read by looking at the angle at the center of the chart. Most InfoVis courses teach it. Many books state it. It's clearly true.

Except it's wrong. The claim is based on research that was published in 1926 – 90 years ago. In a study conducted long before visualization as a field existed, Walter Crosby Eells asked about 100 students in an introductory psychology course to read percentages off of a collection of pie and bar charts [11]. In addition, he asked them to indicate which

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

BELIV '16, October 24 2016, Baltimore, MD, USA

© 2016 Copyright held by the owner/author(s). Publication rights licensed to ACM. ISBN 978-1-4503-4818-8/16/10...\$15.00

DOI: <http://dx.doi.org/10.1145/2993901.2993909>

visual cue they had used: area, central angle, arc length, or chord length. Just over half of his participants said they used central angle, and thus it was established as fact that angle was the way we read pie charts.

Why is this important? The visual cue used to read these charts is a key piece of information if we want to be able to understand how pie charts work. It also allows us to make predictions, like to expect that donut charts are worse than pie charts since they are missing the crucial center that is important for accurate angle judgments.

In two recent papers, Drew Skau and I cast doubt on the central-angle theory. We found that when presented with only angle, study participants did significantly worse than using either just area or just arc, or when presented with a full pie or donut chart [30]. We also found no difference between donut and pie charts, but did find that a slice with a larger radius in a pie chart leads to systematic overestimation of its percentage (Figure 1) – a phenomenon that is inconsistent with angle being the key visual cue, but consistent with both area and arc length [24].

Pie charts may not be liked by the visualization community, but they are extremely popular in business presentations, reports, and information graphics. Understanding them better means being able to provide advice based on evidence – rather than hearsay – that affects a huge number of people.

The larger point here is about research in visualization, though: allowing an assumption to become ingrained to the point of being the basis for predictions (*donut charts are worse than pie charts because they make angle judgment harder*) is dangerous. That is especially true if neither the assumption nor the implication are ever actually tested. Visualization as a field can't operate on the basis of assumptions, it needs evidence. And it needs a critical attitude where we question our assumptions.

2. WHAT WE THOUGHT WE KNEW

The pie charts example above may seem like an outlier. After all, how many other things are there that we assume but don't actually know? It turns out that there are quite a few examples.

Bar Charts

Cleveland and McGill's classic study of graphical perception [10] is often cited, and for good reason. It covered a large number of chart types, asked some important questions, and was the first systematic study of a body of questions about how well we can read different kinds of charts.




Figure 1: Error (difference between the true percent shown and the participant’s guess) by pie chart variation type [24]. Participants’ guesses systematically overestimated the value of the larger slice, which is consistent with them using area or arc length to read the chart, but not angle.

It is not without its issues, though. A healthy scientific process does not simply keep reiterating the same facts without questioning them, but actively tries to find limitations, asks deeper questions, and sometimes finds contradicting evidence. A paper by Talbot et al. [33] is a great example of this: they dug deeper into the comparison between bars and found a number of interesting new hypotheses about why certain tasks (like comparing stacked bars) are difficult – and that those may partly be caused by the design of the charts used in the experiments. Many of the original findings are still valid, but we now have a more nuanced understanding of the results. We should also realize that we need take some of them with a grain of salt.

Banking to 45 Degrees

One of the few well-established rules in visualization is *banking to 45°*: the aspect ratio for line charts should be chosen so that the average angle is 45° [9, 10]. Several approaches have been proposed to achieve this [15, 31], all based on the assumption that banking to 45° was an established good idea.

And yet, it turns out that the recommendation was based on a faulty analysis [32]. The 45° recommendation was the result of a model that was run on a constrained set of data. Including more angles in the study would have revealed a much lower optimal angle. The 45° rule is still useful because it leads to aesthetically pleasing results, but the real maximum precision in line slope comparisons happens at a much shallower angle (which is not practical for most charts and presents other problems).

Deceptive Visualization

A commonly-held belief is that certain chart manipulations cause systematic error and thus make it possible to deceive the viewer. Among them are such commonly-cited issues as cropped vertical axes in bar charts, bubble charts sized by radius rather than area, inverted axes, changes in aspect ratio, etc.

In a recent paper, Pandey et al. tested some of these beliefs and found them to be true [27]. Assumptions can of course

be either true or false, so it is crucial to do the research and find out which is which. We now have evidence that some of the things we thought were deceiving really are. This may seem less than exciting, but it is crucial for building further studies on a solid foundation.

The findings in this study also don’t mean that other unproven assumptions will turn out to be true – they need to be tested one by one.

Embellishments and Chart Junk

Embellished charts have received some attention lately, after being almost completely ignored by the academic community for decades. They are very common in information graphics, but *chart junk* is considered a sin without clear proof that it actually hurts the reading of the charts.

The papers by Bateman et al. [1] and Borgo et al. [2], and more recently Borkin et al. [3, 4], as well as some of our own work [13], have looked at the effects of embellishments on memory. They have largely found that their negative effects are limited, and that they help with memory. Other studies have looked at embellishments purely in terms of effectiveness [29] and not found them to be impacted significantly.

This recent surge of interest is not the first time this has been tested, however. A largely-ignored paper in a journalism journal already did so in 1989 [20] and also found embellishments to not impact viewers of newspaper-style charts.

And yet the dangers of chart junk are still considered part of the visualization canon. The idea is definitely attractive: remove all that is unnecessary, show only the data. But our views need to reflect reality, not preconceived notions or ideas about purity. There are still many gaps in the research, and extraneous elements probably get in the way more for some tasks than for others. But we need to start acknowledging what the research is finding for situations where we have research findings.

Staggered Animation

The idea that staged and staggered animations help users track what is going on is attractive and was presented in a well-written paper by Heer and Robertson [16]. It intuitively makes sense based on common tricks used in animation to cue movements, to make cuts in movies easier to follow, etc.

A more recent paper found, however, that the staggered element of the animation had little to no effect [8]. A number of phenomena that are understood well in psychology, like crowding and occlusion, have a much more pronounced effect on our ability to follow moving objects.

Visual Metaphors

We tend to think of the plethora of different visualization techniques as purely a function of the data’s properties and the user’s task. But different representations of the exact same data can lead to different understanding and, more importantly, to different decisions. Caroline Ziemkiewicz and I found that people interpreted seemingly irrelevant cues like the balance of colors as indicators of stability, rigidity, etc., of data [38], that matching the metaphor used in phrasing the question to the visual metaphor allowed people to answer questions faster [37], and that the visual metaphors used in network diagrams seem to be based on our inherent ideas of physical connection between objects [39].

A study by Elting et al. [12] found the choice of visual representation to make a difference in a setting of particu-

lar importance: clinical studies and whether to continue or abort them. Showing individual visual objects for each patient led to better decisions than summary charts like bars or pies.

We tend to think of visual representations as a function of data properties: discrete vs. continuous, etc. But how much do we know about the way visualization techniques work for our perception of magnitude, number, etc.? How much do we know about how they influence decisions? The truth is, not much.

Complexities of Color

Color is an integral part of visualization, but color perception is highly complex. The rainbow colormap has been maligned to the point of it being a tired cliché [5], and yet it is widely used in many real-world applications. We have no idea why. This strikes me as very similar to the pie charts situation: both are widely used but hated by academics, and in turn they are poorly understood and not studied. What if the medical community decided that the common cold was too banal to be worthy of their time, and decided to leave people with this boring disease to fend for themselves?

To illustrate the complexities of color with just one example: a clever study found that our perception of color is influenced by the words we have for different colors [35]. What do we know about the influence of language on our perception of categorical color in visualization? What about continuous color?

What Do We Know?

All of these examples are meant to show the kinds of questions we need to ask. If we believe something to be true, we need to ask if there is a study showing it to be so. If there isn't, we have to acknowledge this gap in our knowledge and strive to close it.

As it is, many of the things we believe to be true are not based on research. This is not because we *can't* know what is true and what is not, but because we trust our beliefs too much and don't even ask the questions – we *choose* not to know. We don't even need to question the evidence, but simply ask if there even is evidence in the first place.

This isn't just a matter of knowing the evidence – it's a huge opportunity. Beliefs and ideas are the first step towards more formal study and, thus, science. Other sciences understand the difference between ideas and established facts, embrace it, and use it to drive their work forward [6].

3. THE DANGERS OF SEMINAL WORK

Landmark papers and historical examples are important pillars on which we build our research field. But there is a danger in relying on them too much and considering them as absolute truths or unassailable standards, when they should be the basis for critical reflection and new work.

3.1 Seminal Papers

Scientific knowledge is not static, but in constant flux. Seminal papers provide more support and for a longer time than others. They can and should serve as the starting point for many more projects than the average paper. At the same time, they need to be scrutinized for errors and limitations. While that is true for any piece of published work, it is even more true for the central papers that a field draws heavily from.

A seminal paper does not lose its importance if limitations are found, or if parts of it are found to not or no longer be true. The value of the paper is not in providing facts, but in establishing a way of thinking, showing a new direction, introducing a methodology, etc.

Slavishly following a seminal paper's every decision leads to problems down the road. Cleveland and McGill decided to use log error in their studies, with the justification that “[a] log scale seemed appropriate to measure relative error” [10, p. 540]. This has been copied by many researchers in many follow-up papers (including this author). The log error is difficult to interpret and arguably the wrong choice in many cases, however (in particular in pie charts, where the error is a difference in percentages, not proportional error). We need to treat these papers equally with respect and critical distance unless we consider their authors infallible – a dangerous assumption.

Visualization has been treating its seminal papers with too much deference and too little openness to criticism. I and others have argued that a culture of criticism in visualization (similar to how it is done in art and some other fields) would be helpful, and have laid out some ideas for it [21, 23]. Embracing a more critical attitude would help breathe some fresh air into the field and inspire many exciting new research projects.

3.2 Historically Significant Visualizations

In a similar vein, we tend to rely on a small number of classic examples when teaching visualization. Students taking introductory visualization courses are practically guaranteed to see Nightingale's coxcomb plot, Playfair's imports and exports chart, Jon Snow's Cholera map, and Minard's *Napoleon's March*.

These pieces are all historically important, but we need to understand their contexts and limitations. In particular, Snow's map is often misrepresented (it was created as a tool for persuasion, not to find the source of the disease [18]), and Nightingale's chart hides the real pattern in a confusing representation. Another often-misrepresented chart shows O-ring tests before the Challenger disaster [28].

We rely on these examples so much because of the short and sparse history of our field: there are only a handful known visual representations of numerical and/or statistical data that are more than 100 years old. Do we need to rely on them quite as much, though? Do we need to insist on a history that barely exists? It seems much more fruitful to use more recent examples that are more relevant to how data is used today, and simply better visualizations. Data journalism has produced a large amount of visualization pieces over the last ten to fifteen years, many with better interaction and aesthetics than what the academic field is producing.

Another important point to note is that most of the historic charts we use to teach were made for presentation rather than analysis. I have argued [22] that this is a crucial difference that we need to be aware of, both when creating visualizations and when assessing them. A lack of awareness of this difference adds to the confusion these examples create.

4. A SYSTEMATIC APPROACH

In many ways, visualization today resembles the early days of research into electricity [7]: there are some established facts, but also many unfounded beliefs and miscon-

ceptions. The history of electricity is quite chaotic: many discoveries were accidental, often based on wrong assumptions or preconceived ideas, and only understood later.

While a certain amount of chaos is certainly part of the scientific process, we have a better understanding today, and also better technology and support. A more directed and systematic approach will allow us to make much faster progress and avoid dead ends. Efficiency is not typically a consideration in research, but it seems appropriate at this stage to close the gaps in our knowledge, or at least assess what we do and do not know, as quickly as possible.

4.1 Systematic Studies

The questions in visualization can often be stated in ways that make them apply across many different kinds of charts and situations. Those then can be studied one by one.

Systematic studies also have other benefits. They lead to better coverage of the field than one-offs, and they allow studies to be compared more easily. If different studies are run with different designs, different parameters, etc., they are much harder to compare than ones that are based on the same designs.

There can be redundancy in systematic studies, as perhaps questions are revisited that have already been answered. Those are not wasted efforts, though. Replication is valuable (see below), and it can be useful to be able to use consistent data and findings across many questions.

4.2 Modeling, Prediction, Reanalysis

Modeling phenomena requires a much deeper understanding than just observing. Modeling also allows us to make predictions, which we can then test, which leads to better models, etc.

Modeling and predictions are a central part of scientific work in the hard sciences, but are currently underexplored in visualization. Reporting on the results of a study is a good first step. Being able to make predictions based on a model, and then testing those predictions, is much more powerful, however. Models allow us to predict not just for the purposes of studies, but can be used in applications, can be combined for more complex uses, etc.

For visualization to become a real science, it not only needs to start questioning and confirming its basic assumptions, but build on top of them. Models are a key element of this.

Models are slowly starting to be developed in visualization, but they are still rare. One example is the reanalysis of a paper that reported on a study that found that high-level judgments of correlation followed the rather low-level *Weber's law* [14]. A more refined model and some changes in how individual performance was weighted led to a better fit and a better understanding of the phenomenon [19].

4.3 Replication

A cornerstone of many sciences is replication: can others repeat the experiment and get similar results? In fact, in many fields, a finding is not considered valid unless the study has been replicated at least once.

It speaks to the strong computer science roots of visualization that replication is not valued in this field. It looks like a waste of time to people coming from a world where running the same algorithm again will yield the same results. But when humans are involved, things are a lot more

complicated.

For a stronger visualization field, we need to be able to replicate studies and publish those replications. Only then will we be able to trust our results and learn from the replications. Any new study should start with the replication of an existing study and then build on it. That would create a strong network of verification and trust.

Failed replications would raise interesting questions that would lead to further studies and hopefully new insights. A seeming failure can be the starting point for a new research direction (or it can be a sign that there was a flaw in the original study).

Today, replications in visualization are done under the assumption that they will find the exact same results as the original study. This is not necessarily the case, however, even if the original study was perfectly sound and the observed effect was real. Statistical variation will lead to different outcomes, which will need to be analyzed properly to understand when new evidence supports the existing theories, and when it contradicts them.

This point is moot, however, until replications can be published. I am aware of several instances where papers were rejected for being “just replications.” Replications are valuable and necessary. Until visualization is able to publish them, it will be difficult for it to make progress towards being a real science.

4.4 Introspection and Best Practices

That there are issues with evaluation in visualization is not an original idea, of course. The BELIV conference was started to address this issue and explore new ideas. Munzner has proposed a very useful model [26] that captures common mistakes that I still see quite often when reviewing papers.

Lam et al. [25] have organized a large number of evaluation papers from information visualization into seven *scenarios*: combinations of research questions and evaluation methods. Isenberg et al. [17] expands on this by including scientific visualization and comparing the kinds of evaluation methods used in the different communities. Both papers provide a large number of pointers to work that can be used as the starting point for further research, and in particular classifications of methods that can form the basis for new work.

Tory [34] provides many useful insights and tips for running studies based on her own extensive experience.

4.5 Funding

All of the above points to the need for an organized, funded research program. Funding can be difficult to find, especially for fundamental work in visualization. It seems that funding agencies and the grant proposal reviewers view these proposals as unnecessary and too basic.

It may also be somewhat disingenuous to ask for support of research that provides the foundations of much of the work that is currently being done. Does that invalidate that other work? Does that mean such work should not have been funded? Any science has applied and theoretical elements though, so it should be easy to argue that while we have been doing good applied work, we need more support for the theoretical side.

If funding is hard to come by, there are always senior projects, master's theses, etc. Many of the studies described here are quite straight-forward and can be run in a sort of pipeline.

5. SOME RESEARCH QUESTIONS

It is not difficult to come up with a long list of questions that we need to answer. Many of these are quite fundamental, and all are easy to generalize to more than individual questions.

Aspect Ratios. Banking to 45° has been shown to not really be a real rule, so what can we find out about aspect ratio? What should be the aspect of a line chart, bar chart, scatter plot, parallel coordinates, Sankey diagram, etc.? The banking paper looked at a very specific task, what about other tasks? Are different aspect ratios better for different tasks? And are the optimal aspect ratios undesirable for other reasons, like the extremely shallow ratio that is theoretically best for slope comparison? Similarly, what about spacing between bars, width of bars, etc.?

Visual Metaphors. Our perceptual systems did not evolve by us staring at abstract shapes on computer screens. What tricks are we playing on our perception by using shapes that are nothing like the natural imagery we're presumably optimized for? Where does that have adverse effects, and where can we exploit it for better visualization techniques, easier comprehension, and perhaps even truly intuitive techniques? This might include elements that are generally ignored in information visualization, like textures and elements of 3D.

Animation. Beyond the staggered animation issue mentioned above, there has been some limited work on animation in visualization. There is a lot of room for more research, however. How effective are different kinds of transitions for different purposes? Those might include being able to follow individual items, or merely seeing an indication that the view has changed (because of a filter changing, etc.).

Data-induced Clutter. What are the limits of visualization techniques in terms of the number of values, the complexity of the data, etc.? Scatterplots and parallel coordinates can look very clean or very cluttered for the same number of data points depending on their values. Zacks et al. [36] found that noisy data between two bars has an equally adverse effect as adding gratuitous 3D to the bars. We can design visualization tools that don't show unnecessary depth cues, but how do we deal with the data itself causing issues?

Beyond Task Performance. What uses are there for visualization beyond exploration? The recent work on memorability and engagement is arguably just a first step. There is no doubt that visualization is often used to engage audiences and catch their interest. How is this done in an effective way? What are the trade-offs between eye-catchiness and effectiveness for deeper understanding? If we want people to remember what they have seen, what exactly do we mean by that? And how do we measure that? How do we assess what people have learned from a news visualization piece that isn't just repeating numbers? What are the ultimate goals of visualization in different contexts other than exploration and analysis of data?

Re-Evaluate Other Techniques. The issues listed above mostly center around pie and bar charts. What about the rest? Perhaps some of our assumptions about scatterplots are wrong. Have we figured out what effect the size of the space between axes in parallel coordinates has? How much do we know about dot plots, line charts, Gantt charts, Sankey diagrams, etc.? We know some basics about stacked bars now, what about grouped bars? What about stacked bars that always sum up to 100%, are those better or worse? How bad are the stacked bar components in a Sankey diagram?

Some of the above will be studied using the classical time and error metrics. But some require measures of engagement, memory, and deeper understanding than just reading individual numbers. Some will also depend on task much more than others.

6. CONCLUSIONS: HOW TO DO BETTER

Many of the supposed rules in visualization are tightly interwoven with aesthetics. It's easy to side with the idea of minimalist charts that lack the garish embellishments of infographics. Pie charts are easy to hate. Staggered animations and 45° line charts make intuitive sense.

The danger of these seemingly obvious rules and facts is that they are deceptive in their beauty and simplicity. What if reality is more complicated and doesn't adhere to modernist design aesthetics? What if our perception and memory are messier and work better when there are more decorative elements to hang on to? What if there is no single rule that tells us what aspect ratio to pick? What if our existing visual representations don't mesh well with our ways of thinking about the data? Etc.

A number of our assumptions are currently unproven. Many of them are undoubtedly true or close to the truth. We need to find out which ones those are, though. And the ones that do not hold up to scrutiny need to be revised. We have the opportunity to learn an enormous amount of new information from performing this work. And what is more, we will end up with a much stronger field that is based on solid foundations we can trust and build on.

7. REFERENCES

- [1] S. Bateman, R. Mandryk, C. Gutwin, A. Genest, D. McDine, and C. Brooks. Useful Junk? The Effects of Visual Embellishment on Comprehension and Memorability of Charts. *ACM Conference on Human Factors in Computing Systems (CHI)*, pages 2573–2582, 2010.
- [2] R. Borgo, A. Abdul-Rahman, F. Mohamed, P. W. Grant, I. Reppa, L. Floridi, and M. Chen. An Empirical Study on Using Visual Embellishments in Visualization. *IEEE Transactions on Visualization and Computer Graphics*, 18(12):2759–2768, 2012.
- [3] M. A. Borkin, Z. Bylinskii, N. W. Kim, C. M. Bainbridge, C. S. Yeh, D. Borkin, H. Pfister, and A. Oliva. Beyond Memorability: Visualization Recognition and Recall. *IEEE Transactions on Visualization and Computer Graphics*, 22(1):519–528, 2015.
- [4] M. A. Borkin, A. A. Vo, Z. Bylinskii, P. Isola, S. Sunkavalli, A. Oliva, and H. Pfister. What Makes a

- Visualization Memorable? *IEEE Transactions on Visualization and Computer Graphics*, 19(12):2306–2315, 2013.
- [5] D. Borland and R. M. Taylor II. Rainbow Color Map (Still) Considered Harmful. *IEEE Computer Graphics and Applications*, 27(2):14–17, 2007.
- [6] J. Brockman, editor. *What We Believe but Cannot Prove*. Harper Perennial, 2006.
- [7] H. Camenzind. *Much Ado About Almost Nothing*. Man’s Encounter with the Electron. Booklocker.com, 2007.
- [8] F. Chevalier, P. Dragicevic, and S. Franconeri. The Not-so-Staggering Effect of Staggered Animated Transitions on Visual Tracking. *IEEE Transactions on Visualization and Computer Graphics*, 20(12):2241–2250, 2014.
- [9] W. S. Cleveland. A Model for Studying Display Methods of Statistical Graphics. *Journal of Computational and Graphical Statistics*, 2(4):323–343, 1993.
- [10] W. S. Cleveland and R. McGill. Graphical Perception: Theory, Experimentation, and Application to the Development of Graphical Methods. *Journal of the American Statistical Association*, 79(387):531–554, 1984.
- [11] W. C. Eells. The Relative Merits of Circles and Bars for Representing Component Parts. *Journal of the American Statistical Association*, 21(154):119–132, 1926.
- [12] L. S. Elting, C. G. Martin, S. B. Cantor, and E. B. Rubenstein. Influence of data display formats on physician investigators’ decisions to stop clinical trials: prospective trial with repeated measures. *British Medical Journal*, 318:1527–1531, 1999.
- [13] S. Haroz, R. Kosara, and S. L. Franconeri. ISOTYPE Visualization – Working Memory, Performance, and Engagement with Pictographs. In *Proceedings CHI*, pages 1191–1200, 2015.
- [14] L. Harrison, F. Yang, S. L. Franconeri, and R. Chang. Ranking Visualizations of Correlation Using Weber’s Law. *IEEE Transactions on Visualization and Computer Graphics*, 20(12):1943–1952, 2014.
- [15] J. Heer and M. Agrawala. Multi-Scale Banking to 45°. *IEEE Transactions on Visualization and Computer Graphics*, 12(5):701–708, 2006.
- [16] J. Heer and G. G. Robertson. Animated Transitions in Statistical Data Graphics. 13(6):1240–1247, 2007.
- [17] T. Isenberg, P. Isenberg, J. Chen, M. Sedlmair, and T. Möller. A Systematic Review on the Practice of Evaluating Visualization. *IEEE Transactions on Visualization and Computer Graphics*, 19(12):2818–2827, 2013.
- [18] S. Johnson. *The Ghost Map*. Riverhead Books, 2007.
- [19] M. Kay and J. Heer. Beyond Weber’s Law: A Second Look at Ranking Visualizations of Correlation. *IEEE Transactions on Visualization and Computer Graphics*, 22(1):469–478, 2016.
- [20] J. D. Kelly. The Data-Ink Ratio and Accuracy of Newspaper Graphs. *Journalism & Mass Communication Quarterly*, 66(3):632–639, 1989.
- [21] R. Kosara. Visualization Criticism – The Missing Link Between Information Visualization and Art. In *International Conference on Information Visualization (IV ’07)*, pages 631–636. IEEE, 2007.
- [22] R. Kosara. Presentation-Oriented Visualization Techniques. *IEEE Computer Graphics and Applications*, 36(1):80–85, 2016.
- [23] R. Kosara, F. Drury, L. E. Holmquist, and D. H. Laidlaw. Visualization Criticism. *IEEE Computer Graphics and Applications*, 28(3):13–15, 2008.
- [24] R. Kosara and D. Skau. Judgment Error in Pie Chart Variations. In *Short Paper Proceedings of the Eurographics/IEEE VGTC Symposium on Visualization (EuroVis)*, pages 91–95. The Eurographics Association, 2016.
- [25] H. Lam, E. Bertini, P. Isenberg, C. Plaisant, and S. Carpendale. Empirical Studies in Information Visualization: Seven Scenarios. *IEEE Transactions on Visualization and Computer Graphics*, 18(9):1520–1536, 2011.
- [26] T. Munzner. A Nested Model for Visualization Design and Validation. *IEEE Transactions on Visualization and Computer Graphics*, 15(6):921–928, 2009.
- [27] A. V. Pandey, K. Rall, M. L. Satterthwaite, O. Nov, and E. Bertini. How Deceptive are Deceptive Visualizations?: An Empirical Analysis of Common Distortion Techniques. *Proceedings CHI*, pages 1469–1478, 2015.
- [28] W. Robison, R. Boisjoly, D. Hoeker, and S. Young. Representation and Misrepresentation: Tufte and the Morton Thiokol Engineers on the Challenger. *Science and Engineering Ethics*, 8(1):59–81, 2002.
- [29] D. Skau, L. Harrison, and R. Kosara. An Evaluation of the Impact of Visual Embellishments in Bar Charts. *Computer Graphics Forum*, 34(3):221–230, 2015.
- [30] D. Skau and R. Kosara. Arcs, Angles, or Areas: Individual Data Encodings in Pie and Donut Charts. *Computer Graphics Forum*, 35(3):121–130, 2016.
- [31] J. Talbot, J. Gerth, and P. Hanrahan. Arc Length-Based Aspect Ratio Selection. *IEEE Transactions on Visualization and Computer Graphics*, 17(12):2276–2282, 2011.
- [32] J. Talbot, J. Gerth, and P. Hanrahan. An Empirical Model of Slope Ratio Comparisons. *IEEE Transactions on Visualization and Computer Graphics*, 18(12):2613–2620, 2012.
- [33] J. Talbot, V. Setlur, and A. Anand. Four Experiments on the Perception of Bar Charts. *IEEE Transactions on Visualization and Computer Graphics*, 20(12):2152–2160, 2014.
- [34] M. Tory. User Studies in Visualization: A Reflection on Methods. In W. Huang, editor, *Handbook of Human Centric Visualization*, pages 411–426. Springer, 2014.
- [35] J. Winawer, N. Witthoft, M. C. Frank, L. Wu, A. R. Wade, and L. Boroditsky. Russian blues reveal effects of language on color discrimination. *Proceedings of the National Academy of Sciences*, 104(19):7780–7785, 2007.
- [36] J. Zacks, B. Tversky, E. Levy, and D. J. Schiano. Reading Bar Graphs: Effects of Extraneous Depth Cues and Graphical Context. *Journal of Experimental Psychology: Applied*, 4(2):119–138, 1998.
- [37] C. Ziemkiewicz and R. Kosara. The Shaping of

Information by Visual Metaphors. *IEEE Transactions on Visualization and Computer Graphics*, 14(6):1269–1276, 2008.

[38] C. Ziemkiewicz and R. Kosara. Implied Dynamics in Information Visualization. In *Proceedings Advanced Visual Interfaces (AVI)*, pages 215–222. ACM Press,

2010.

[39] C. Ziemkiewicz and R. Kosara. Laws of Attraction: From Perceptual Forces to Conceptual Similarity. *IEEE Transactions on Visualization and Computer Graphics*, 16(6):1009–1016, 2010.