Reading a graph is like reading a paragraph

Tal Boger[a,*], Steven Franconeri[b,*]

[a] *Johns Hopkins University, Baltimore, MD 21218*
[b] *Northwestern University, Evanston, IL 60208*

Abstract

Vision provides rapid processing for some tasks, but encounters strong constraints for others. Although many tasks encounter a capacity limit of processing 4 visual objects at once, some evidence suggests far lower limits for processing *relationships* among objects. What is our capacity limit for relational processing? If it is indeed limited, then people may miss important relationships between data values in a graph. To test this question, we asked people to explore graphs of trivially simple 2x2 datasets, and found that half of viewers missed surprising and improbable relationships (e.g., a child's height *decreasing* over time). These relationships were spotted easily in a control condition, which implicitly directed viewers to prioritize inspecting the key relationships. Thus, a severe limit on relational processing, combined with a cascade of other capacity-limited operations (e.g., linking values to semantic content), makes understanding a graph more like slowly reading a paragraph than immediately recognizing an image. These results also highlight the practical importance of 'data storytelling' techniques, where communicators design graphs that help their audience prioritize the most important relationships in data.

Materials: `https://osf.io/tjbyp/`

Keywords: data visualization, capacity limits, communication

---

*To whom correspondence should be addressed.

Email addresses: `tboger1@jhu.edu` (Tal Boger), `franconeri@northwestern.edu` (Steven Franconeri)

Public significance statement

This study shows that visual processing is drastically limited when processing relationships between objects, such as the bars in a graph. In graphs containing only 4 values, participants miss surprising relationships at high rates. This highlights the practical importance of designing graphs that guide an audience to process important relationships.

Introduction

The graph in Figure 1a depicts the height of two children over two different ages. What patterns do you notice? If you are like our participants, then you first noticed that River is taller than Charlie at age 10. Then, Charlie catches up and surpasses River at age 12. But did you notice that River *shrinks* between ages 10 and 12, which seems improbable? If not, you are like our participants in that regard as well — approximately half of which missed this surprising relationship, even after viewing the graph for 15 seconds. Why is this relationship so easy to miss? Processing only 4 visual values is well within capacity estimates of short-term attention and memory, so we might expect to notice such patterns. But extracting those relations requires a cascade of capacity-limited operations that makes reading even a trivially simple graph as slow as reading a paragraph of text.

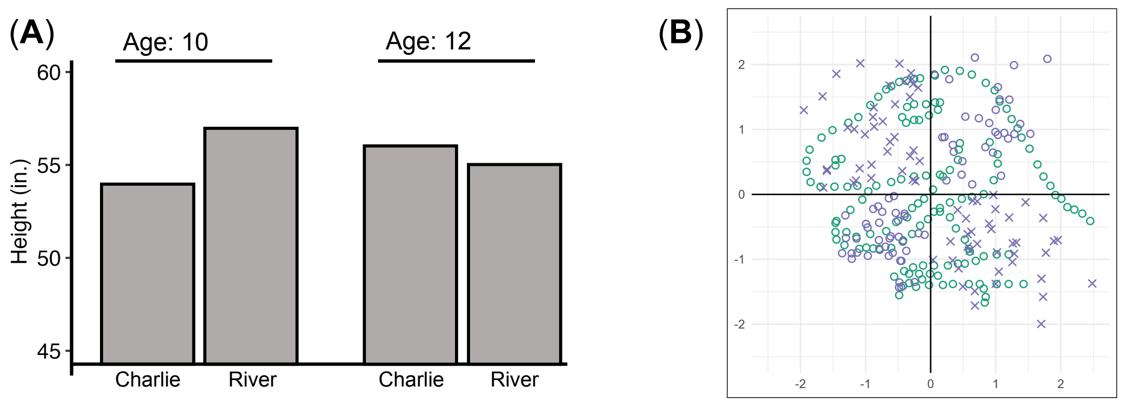

Figure 1: (A) What patterns do you notice in this graph? Did you notice that River shrinks between ages 10 and 12? Even in graphs like this containing only 4 values, our participants missed improbable relationships at rates of up to 58%. (B) Reproduction of Boger et al., 2021. When participants performed a task that required filtering for the blue X's in scatterplots like this, more than half missed the appearance of a conspicuous dinosaur formed by the green circles (93% for 1s presentation, 61% for 2.5s) in one of the plots.

Visual processing can unfold in rapid and powerful ways. We perform impressive operations broadly across the visual field when we visually extract statistics (Haberman and Whitney, 2012), recognize simple features (Wolfe and Horowitz, 2017), or even learn about complex objects (Ahissar and Hochstein, 2004). These operations are similarly powerful when we look at data visualizations (Franconeri et al., 2021; Healey and Enns, 2011; Szafir et al., 2016). But many visual processes encounter strong capacity limitations, such that processing too many objects or features leads to slower or less accurate performance, forcing observers to filter visual input as smaller subsets of information at a time (Franconeri, 2013; Serences and Yantis, 2006). This filtering is guided by not only goal-directed heuristics, but also bottom-up cues to process unique objects, or to simultaneously process objects that are similar in their color, size, spatial proximity, or connectivity (Yu et al., 2019a,b). Such bottom-up cues can also guide which relationships are prioritized within data visualizations (Xiong et al., 2021; Bearfield et al., 2023).

When information does not match a visual filter, people can fail to notice salient objects, such as motorcycles on a collision course, people holding umbrellas, or even a gorilla (Simons and Chabris, 1999, for review, see Jensen et al., 2011; Most, 2013). A visual filter can have similar effects within data visualizations, with one study finding that 93% of viewers focused on detecting patterns within a scatterplot's blue points failed to notice a salient dinosaur formed by an ignored set of green points (Boger et al., 2021; Figure 1b).

Given these strong capacity limitations, we ask: How limited is relational processing in a realistic case study of processing a simple graph? We choose graph processing because of its critical and ubiquitous role in helping people think about and communicate with quantitative information across science, education, and organizations (Franconeri et al., 2021).

Estimates of visual processing capacity vary, depending on what observers are asked to do. One study proposes a limit of approximately four variables when extracting interactions among variables in a bar graph (Halford et al., 2007). Similarly, when memorizing a list of visual features like colors or tracking a set of objects, estimates of visual capacity hover around 4 objects (Brady et al., 2011; Scimeca and Franconeri, 2015). However, when people are asked to remember features not as a simple list, but rather as features linked to specific locations or moving objects, some work shows even lower limits of 1-2 objects or features (Scimeca and Franconeri, 2015; Huang and Pashler, 2007; Xu and Franconeri, 2015; Saiki, 2002). This limit should apply even for simple static bar graphs, where each object (a bar) might need to be linked to several potential features, including its relative size (the data value), and/or relative spatial position. For those relative sizes or spatial positions alone, one model of visual relationship processing predicts that only a single relationship can be judged at a time between two objects (or statistical summaries of two groups of objects; Franconeri et al., 2012). This may be further limited to relations in a single direction (e.g., extracting that A is larger than B is different than extracting B being smaller than A; Michal et al., 2016) or to within a single feature dimension at a time (e.g., size or contrast, but not both; Michal and Franconeri, 2017). If a 2x2 graph presents 4 data values, that presents 6 possible pairwise relations to its viewer, plus at least 2 main effects and 2 interactions. All of those numbers would double if relations are interpreted in a 'directional' fashion. While this large space of possibilities should also be drastically reduced by strategic prioritization of which relations to process first, as a product of both current goals and previous experience with graphs, the number of potential relations is daunting.

Understanding a graph requires not only these lower-level perceptual operations, but also connecting those extracted features and relations to their semantic labels and problem context, which should additionally strain capacity limits along a cascade of cognitive operations (Franconeri, 2013). Indeed, in the graph comprehension literature, there is an influential mantra that understanding a graph is more like 'reading a paragraph.' In other words, reading a graph is a slow and serial extraction of relational 'sentences' from the data (Carpenter and Shah, 1998; Shah et al., 2005) rather than an instantaneous and parallel process like recognizing a picture (Li et al., 2002). Thus, is it possible that our exploration of visualized data could be limited to perhaps a *single* relationship at any given moment?

The present experiments demonstrate that, even when given time to explore a trivially simple 4-value bar graph (within typical estimates of short-term visual processing capacity), participants judge only a small subset of possible relations. We designed graphs that induce viewers to prioritize some relational comparisons over others, by placing some values closer to each other in space, making participants unlikely to compare other values that are farther apart. We then embedded surprising relationships among the far-apart values, and asked if participants would notice these surprising relationships. Participants missed these salient patterns at rates as high as 58%. Importantly, participants noticed these patterns at far higher rates when the graph designs highlighted them (again by spatial grouping), suggesting that participants missed the patterns due to limits on relational processing. If reading a graph is like reading a paragraph, then this manipulation should implicitly re-order the sentences of that paragraph, so that in the limited time available, the participant should extract some relations, but not others.

Previous work in data visualization suggests that a demonstration for such a trivially simple graph might be possible. When people were shown relatively complex bar graphs (3 values in one factor, 3 or 4 in the other), they were around 3 times more likely to produce descriptions that offered comparisons within local spatial groups, as opposed to within the factor that was interleaved across the local groupings (Shah et al., 1999; Shah and Freedman, 2011, see also Carpenter and Shah, 1998; Shah and Carpenter, 1995). However, these experiments relied on graphs with 9-12 total data values, which already lay outside typical limits of visual attention and memory, even before considering limits on processing relationships or semantic identities. In other words, any failures to process relationships in these experiments could be due to capacity limits of attention and memory, rather than capacity limits of relational processing. Previous studies also typically use graph designs containing a legend (instead of direct labeling, as used here) which presents an additional working memory load (Lohse, 1993). The graphs used also presented relatively complex topics and potential relationships (e.g., metric, ordinal, or nominal interactions among temperatures and noise levels on test scores). In contrast, the present experiments rely on simple datasets consisting of either one metric and one nominal independent variable, or an even simpler case of two nominal variables, each with only two values. Furthermore, previous work relied on less familiar topics (e.g., population changes across 3 years for 4 geographic regions), which require higher graphical literacy levels (Shah and Freedman, 2011), as opposed to the present experiments that rely on highly familiar contexts (e.g., children getting taller, comparing battery life for phones).

Finally, in previous work, participants were asked to generate statements about the graph, with the assumption being that more salient descriptions are given first. But a comparison might still have been made and not reported or remembered, even if it is not prioritized (Wolfe, 1999). One solution to this problem is to ensure that a comparison is so surprising that it would surely be reported if it were noticed, such as having a gorilla walk through the middle of a scene (Simons and Chabris, 1999) or a dinosaur appear in a scatterplot (Boger et al., 2021).

Methods

Transparency and openness

All materials, code, and data are available on our OSF repository here: `https://osf.io/tjbyp/`. Interested readers may view our experiments — exactly as participants did — here: `https://perceptionstudies.github.io/graphs`. Our experiments were pre-registered, with the exception of a final, exploratory experiment. This is specified later in our methods and in our results.

Participants

For each of our first three, pre-registered vignettes, we recruited 60 unique participants — 30 for each graph type. All participants were unique and could only view one type of graph or vignette, meaning we collected data from 180 participants total, all from the online recruiting platform Prolific (for a discussion of the reliability of this subject pool, see Peer et al., 2017). Participants were excluded if they did not submit a full data set, or if they claimed to see a relationship that was in fact not present in the graph (as per their binary responses to such a question).

For our final exploratory vignette (2-by-2 'age'), we recruited 40 additional unique participants (20 for each graph type).

Stimuli

We created 3 vignettes and 2 graph types for each vignette (except for the 'age' vignette, which had 2 graph types for both a 2-by-3 and a 2-by-2 vignettes). The vignettes presented either 2-by-2 or 2-by-3 (in the case of the third vignette) datasets. The graphs in the first two vignettes were identical, except that the labels and descriptions were interleaved in two ways, one that highlighted and one that hid the improbable relationship.

Design and procedure

At the start of the experiment, participants read a short vignette describing context for the graph (these descriptions are available as part of our experimental code in our OSF repository here: https://osf.io/tjbyp/). They were told that the graph would appear for 15 seconds and that they should try to remember at least "two interesting comparisons or patterns that they noticed in the data." Before the graph appeared, the axes were also visible on the screen (with no data present), such that participants could take as much time as they needed to understand the axes and labels. After participants revealed the graph (by pressing the right arrow key on their keyboard) and observed it for 15 seconds, the graph disappeared, at which point participants were asked to write at least 10 words describing patterns they found to be "most interesting" in the graph.

Following this description, participants had the opportunity to answer a free-response question asking if they saw anything "that didn't make sense in the plot." After these two free-response questions, participants moved on to binary questions. First, they reported whether they noticed the improbable relationship. Then, they reported whether they noticed a second unlikely relationship that was in fact not present in the data (this question was used to exclude participants). For example, this second question (asking about the non-present relationship) in the case of the phones vignette asked whether the participant noticed that one of the phones had the same battery life across the two conditions (audio only vs. audio + video). However, the initial graph did not depict this relationship. Therefore, a participant who responds positively (i.e., claims to see this relationship) may falsely report seeing the real, present improbable relationship (probed in the first question) due to a simple bias to respond "yes" to the binary questions. In other words, then, excluding participants in this manner provides an important attention check and ensures that we only analyze participants with trustworthy answers to these binary questions.

The present experiments

We presented participants with four different data vignettes. Two of the vignettes depict simple 2-by-2 relationships with four total values, and the third and fourth depict either a 2-by-3 relationship with six values or a 2-by-2 relationship with four values[1] (one such vignette is depicted in Figure 1a). Within each vignette, participants saw the same bar graph arranged to either implicitly highlight the improbable relationship — by placing that comparison in nearby bars in the graph — or implicitly hide the relationship — by placing that comparison across bars that were more distant from each other. Including a condition where the improbable relation is highlighted serves as evidence that the relationship is salient enough to be detected and reported, and not merely discarded in favor of reporting other potential relationships.

Unique participants viewed the graphs for 15 seconds before providing typed descriptions of "...the patterns that you see as most interesting within the graphed data." They were then asked: "Did you notice anything that didn't make sense in the plot?" After entering a free-response description to this question, they gave two binary responses. The first asked explicitly if they saw the improbable pattern, and the second asked whether they noticed a pattern that was not actually present; participants who claimed to notice the pattern that was not present were excluded from our analyses (as they may have falsely reported seeing the critical pattern), as per our pre-registered analysis plans. Interested readers may try our experiments for themselves here: `https://perceptionstudies.github.io/graphs`. Furthermore, all experimental code, data, and analysis is available on our OSF repository here: `https://osf.io/tjbyp/`

When the improbable relationship was implicitly highlighted, people were 1.8x-3.4x times more likely to find the improbable relationship (Figure 2), $\chi^2(1, 163) = 14.08$, $p < .001$. In a post-hoc exploratory analysis, we categorized the sentences that participants typed (before the binary responses) to describe the relationships that they noticed. This revealed that the improbable relationships were far less likely to be described in graph arrangements that hid these relationships, and that the arrangement also appeared to highlight or hide other relationships in similar ways,

[1] Initially, we pre-registered the two 2-by-2 relationships (phones and restaurants) and one 2-by-3 relationship (child age). We ran another exploratory experiment in which the same 'age' vignette depicts only a 2-by-2 relationship of the crucial comparison from the 2-by-3 graphs, and found similar results. However, because this result was not pre-registered, we do not include it in our subsequent χ^2 test.

such that values that were closer to one another were more likely to be compared (see Fig. S1).

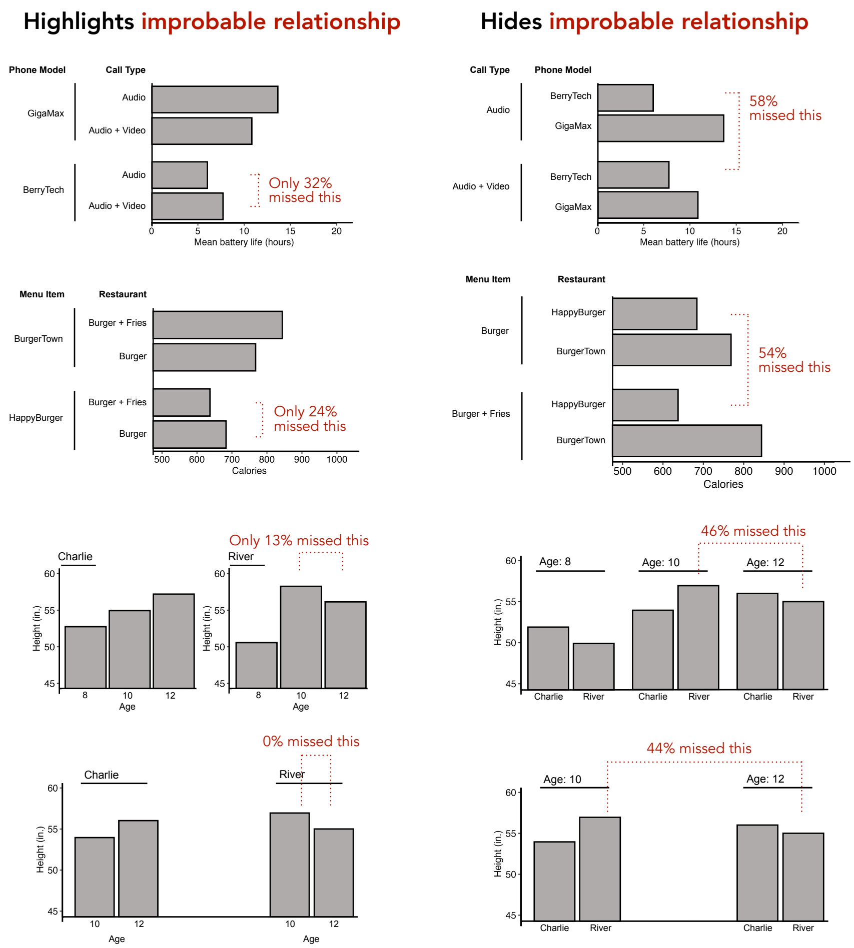

Figure 2: Participants saw one of these eight graphs for 15 seconds and were asked to remember at least two "interesting" relationships shown in the plots. Each contained an improbable relationship: a phone battery lasting longer under a tougher task, a larger meal containing fewer calories at a restaurant, or a child shrinking over time. The left column shows the values in a spatial arrangement that was predicted to highlight the improbable relationship by placing the relevant values closer to each other, while the right shows an arrangement that should hide it by placing them farther away. The annotations report the percent of participants that missed the relationship when asked in a binary response whether they noticed the improbable relationship in the plot.

Conclusion

We show that people miss improbable relationships in trivially simple graphs containing only 4 objects, even after 15 seconds of study. Visual processing can be capacity-limited, but a common limit is still around 4 objects. But making relational judgments in graphs appears to have far more restrictive limits. These limits start at the perceptual stage, where some models suggest that people can process very few (Hummel, 2000; Wolfe, 1999) or perhaps only one (Franconeri et al., 2012) relationship at a time. These limits should be compounded by the cascade of other cognitive operations that should also be capacity limited, including tying the extracted relations to their verbals labels and the meaning of those richer relations in context. The present results strongly support the mantra that understanding a graph is not immediate, like seeing a picture. Rather, it is a slow process that is more akin to reading a paragraph (Carpenter and Shah, 1998; Shah et al., 2005).

These results strongly support an emerging set of guidelines in the practitioner and research literatures on effective data communication. First, it strengthens the demonstrations that grouping factors like spatial proximity, connection (e.g. line graphs), or featural similarity (similar colors or shapes) guide the comparisons that people make in data visualizations (Shah and Carpenter, 1995; Shah et al., 1999, 2005; Shah and Freedman, 2011; Xiong et al., 2021; Bearfield et al., 2023). But because visualization authors often assume that a naive viewer will see the same relationship as they do (Xiong et al., 2019), designers should help viewers notice the 'right' pattern in a dataset by using data storytelling techniques, including highlighting values to be compared and annotating those values with the conclusions drawn from them (Ajani et al., 2021). These steps are important even for trivially simple visualizations. Finally, if a visualization designer can guide a viewer's capacity-limited relationship processing to the 'right' pattern, then graphical literacy training should include monitoring for bad actors who use the same technique to guide people to the 'wrong pattern' (Ge et al., 2023), much like a magician might use subtle attentional misdirection to hide an action from an audience (Kuhn et al., 2008).

Constraints on generality

In our experiments, we recruited participants from Prolific, an online recruiting platform (see Peer et al., 2017). The studies were open only to US adults. We do not take it for granted that our findings generalize beyond this group. However, our studies rely on simple questions to probe visual capacity limits that we believe generalize more broadly.

Author note

The authors have no conflicts of interest. Pre-registrations, stimuli, data and other materials are available on our OSF page: https://osf.io/tjbyp/. This work was supported by NSF IIS-CHS-1901485 and NSF IIS-HCC-2107490 awarded to SF.

Author contributions

Both authors developed the study concept and contributed to the study design and testing. TB collected and analyzed the data under the supervision of SF. Both authors wrote and edited the manuscript together.

Acknowledgments

We thank Amritha Anupindi and members of the Northwestern Visual Thinking Lab for helpful comments and feedback.

Supplementary figure (Fig. S1):

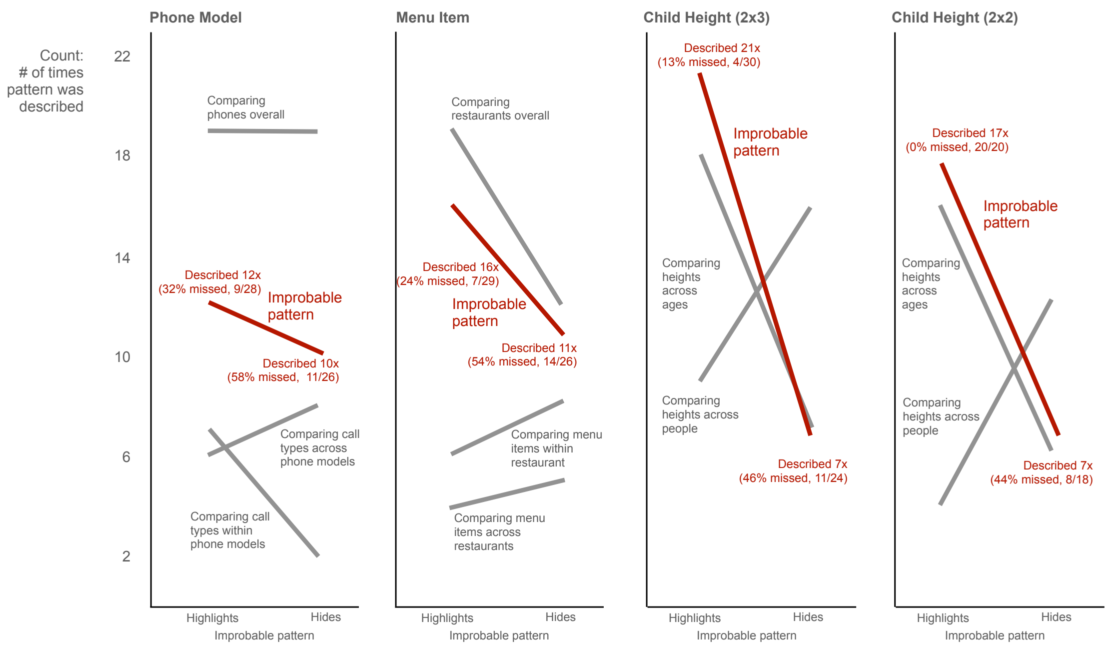

Along with a binary response to whether they saw the improbable relationship (shown in Fig. 2), participants typed descriptions of what they saw in the plot. We conducted an exploratory analysis of these descriptions by coding them as categories of comparisons for each combination of vignette and graph design. Note that, because participants were asked to describe multiple patterns, there are differences between the number of times a comparison is described and the number of times it is seen or missed in the binary responses. The left end of each line shows results from graphs with a spatial arrangement of values predicted to highlight that relationship, while the right end of each line shows an arrangement that should hide it.

While we do not conduct formal statistics on these post-hoc exploratory analyses, the slopes of the lines are consistent with the grouping effect that causes difference in binary responses reported in the manuscript. For each red line, the direction of the difference for the comparison coding (the slope of the lines) is the same as the difference for the binary responses (lower noticing rates in the 'hides' condition) — the spatial grouping designed to highlight or hide the improbably pattern had a similar effect on rate of describing that pattern. The coded descriptions shown in the grey lines also appear to have been influenced by the same spatial grouping effect that highlighted or hid the improbably pattern, such that people are more prone to make comparisons of nearby values, and less likely to compare values that are spatially distant and have other values interleaved between them. The directions of the differences are all congruent with that bias, except for the equal rates for 'Comparing phones overall' in the 'Phone Model' condition and the 'Comparing menu items within restaurant' in the 'Menu Item' condition.

References

Ahissar, M. and Hochstein, S. (2004). The reverse hierarchy theory of visual perceptual learning. *Trends in Cognitive Sciences*, 8(10):457–464.

Ajani, K., Lee, E., Xiong, C., Knaflic, C. N., Kemper, W., and Franconeri, S. (2021). Declutter and focus: Empirically evaluating design guidelines for effective data communication. *IEEE Transactions on Visualization and Computer Graphics*, 28(10):3351–3364.

Bearfield, C. X., Stokes, C., Lovett, A., and Franconeri, S. (2023). What does the chart say? grouping cues guide viewer comparisons and conclusions in bar charts. *IEEE Transactions on Visualization and Computer Graphics*.

Boger, T., Most, S. B., and Franconeri, S. L. (2021). Jurassic mark: Inattentional blindness for a datasaurus reveals that visualizations are explored, not seen. In *2021 IEEE Visualization Conference (VIS)*, pages 71–75. IEEE.

Brady, T. F., Konkle, T., and Alvarez, G. A. (2011). A review of visual memory capacity: Beyond individual items and toward structured representations. *Journal of Vision*, 11(5):4–4.

Carpenter, P. A. and Shah, P. (1998). A model of the perceptual and conceptual processes in graph comprehension. *Journal of Experimental Psychology: Applied*, 4(2):75.

Franconeri, S. L. (2013). The nature and status of visual resources. *Oxford Handbook of Cognitive Psychology*, 8481:147–162.

Franconeri, S. L., Padilla, L. M., Shah, P., Zacks, J. M., and Hullman, J. (2021). The science of visual data communication: What works. *Psychological Science in the Public Interest*, 22(3):110–161.

Franconeri, S. L., Scimeca, J. M., Roth, J. C., Helseth, S. A., and Kahn, L. E. (2012). Flexible visual processing of spatial relationships. *Cognition*, 122(2):210–227.

Ge, L. W., Cui, Y., and Kay, M. (2023). Calvi: Critical thinking assessment for literacy in visualizations. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*, pages 1–18.

Haberman, J. and Whitney, D. (2012). Ensemble perception: Summarizing the scene and broadening the limits of visual processing. *From Perception to Consciousness: Searching with Anne Treisman*, 33.

Halford, G. S., Phillips, S., Wilson, W. H., McCredden, J., Andrews, G., Birney, D., Baker, R., and Bain, J. D. (2007). Relational processing is fundamental to the central executive and is limited to four variables. *The Cognitive Neuroscience of Working Memory*, pages 261–280.

Healey, C. and Enns, J. (2011). Attention and visual memory in visualization and computer graphics. *IEEE Transactions on Visualization and Computer Graphics*, 18(7):1170–1188.

Huang, L. and Pashler, H. (2007). A boolean map theory of visual attention. *Psychological Review*, 114(3):599.

Hummel, J. E. (2000). Where view-based theories break down: The role of structure in shape perception and object recognition. *Cognitive Dynamics: Conceptual Change in Humans and Machines*, pages 157–185.

Jensen, M. S., Yao, R., Street, W. N., and Simons, D. J. (2011). Change blindness and inattentional blindness. *Wiley Interdisciplinary Reviews: Cognitive Science*, 2(5):529–546.

Kuhn, G., Amlani, A. A., and Rensink, R. A. (2008). Towards a science of magic. *Trends in Cognitive Sciences*, 12(9):349–354.

Li, F. F., VanRullen, R., Koch, C., and Perona, P. (2002). Rapid natural scene categorization in the near absence of attention. *Proceedings of the National Academy of Sciences*, 99(14):9596–9601.

Lohse, G. L. (1993). A cognitive model for understanding graphical perception. *Human-Computer Interaction*, 8(4):353–388.

Michal, A. L. and Franconeri, S. L. (2017). Visual routines are associated with specific graph interpretations. *Cognitive Research: Principles and Implications*, 2(1):1–10.

Michal, A. L., Uttal, D., Shah, P., and Franconeri, S. L. (2016). Visual routines for extracting magnitude relations. *Psychonomic Bulletin & Review*, 23:1802–1809.

Most, S. B. (2013). Setting sights higher: Category-level attentional set modulates sustained inattentional blindness. *Psychological Research*, 77(2):139–146.

Peer, E., Brandimarte, L., Samat, S., and Acquisti, A. (2017). Beyond the turk: Alternative platforms for crowdsourcing behavioral research. *Journal of Experimental Social Psychology*, 70:153–163.

Saiki, J. (2002). Multiple-object permanence tracking: Limitation in maintenance and transformation of perceptual objects. *Progress in Brain Research*, 140:133–148.

Scimeca, J. M. and Franconeri, S. L. (2015). Selecting and tracking multiple objects. *Wiley Interdisciplinary Reviews: Cognitive Science*, 6(2):109–118.

Serences, J. T. and Yantis, S. (2006). Selective visual attention and perceptual coherence. *Trends in Cognitive Sciences*, 10(1):38–45.

Shah, P. and Carpenter, P. A. (1995). Conceptual limitations in comprehending line graphs. *Journal of Experimental Psychology: General*, 124(1):43.

Shah, P. and Freedman, E. G. (2011). Bar and line graph comprehension: An interaction of top-down and bottom-up processes. *Topics in Cognitive Science*, 3(3):560–578.

Shah, P., Freedman, E. G., and Vekiri, I. (2005). The comprehension of quantitative information in graphical displays. *The Cambridge Handbook of Visuospatial Thinking*, pages 426–476.

Shah, P., Mayer, R. E., and Hegarty, M. (1999). Graphs as aids to knowledge construction: Signaling techniques for guiding the process of graph comprehension. *Journal of Educational Psychology*, 91(4):690.

Simons, D. J. and Chabris, C. F. (1999). Gorillas in our midst: Sustained inattentional blindness for dynamic events. *Perception*, 28(9):1059–1074.

Szafir, D. A., Haroz, S., Gleicher, M., and Franconeri, S. (2016). Four types of ensemble coding in data visualizations. *Journal of Vision*, 16(5):11–11.

Wolfe, J. M. (1999). Inattentional amnesia. *Fleeting Memories*, 17(5).

Wolfe, J. M. and Horowitz, T. S. (2017). Five factors that guide attention in visual search. *Nature Human Behaviour*, 1(3):0058.

Xiong, C., Setlur, V., Bach, B., Koh, E., Lin, K., and Franconeri, S. (2021). Visual arrangements of bar charts influence comparisons in viewer takeaways. *IEEE Transactions on Visualization and Computer Graphics*, 28(1):955–965.

Xiong, C., Van Weelden, L., and Franconeri, S. (2019). The curse of knowledge in visual data communication. *IEEE Transactions on Visualization and Computer Graphics*, 26(10):3051–3062.

Xu, Y. and Franconeri, S. L. (2015). Capacity for visual features in mental rotation. *Psychological Science*, 26(8):1241–1251.

Yu, D., Tam, D., and Franconeri, S. L. (2019a). Gestalt similarity groupings are not constructed in parallel. *Cognition*, 182:8–13.

Yu, D., Xiao, X., Bemis, D. K., and Franconeri, S. L. (2019b). Similarity grouping as feature-based selection. *Psychological Science*, 30(3):376–385.