# The Curse of Knowledge in Visual Data Communication

Cindy Xiong*
Northwestern University

Lisanne van Weelden†
Utrecht University, Utrecht, Netherland

Steven Franconeri‡
Northwestern University

**ABSTRACT**

The curse of knowledge is an inability to separate one's own knowledge or expertise from that of an audience. We test the idea that this curse can substantially impair visual communication of data, and has the potential to fixate an analyst on a given pattern in data. Because a viewer can extract many potential relationships and patterns from any set of visualized data values, a viewer may see one pattern in the data as more visually salient than others. We demonstrate this phenomenon in the laboratory, showing that when people are given background information, they see the pattern in the data corresponding to the background information as more visually salient. Critically, they also believe that other viewers will experience the same visual salience, even when they are explicitly told that other viewers are naïve to the background information. The present findings suggest that the curse of knowledge affects the visual perception of data, explaining why presenters, paper authors, and data analysts can fail to connect with audiences when they communicate patterns in those data. Because the curse of knowledge may be difficult for a viewer to inhibit or even detect, analysts making decisions may benefit from visualizing their data a variety of formats, and soliciting perspectives of others.

**Keywords**: Information visualization, data communication, cognitive biases, perception and cognition, evaluation, expertise.

## 1 INTRODUCTION

The curse of knowledge refers to the inability for experts to imagine the mindset of a novice. It is a well-studied psychological phenomenon that appears in many domains. Well-informed business decision makers fail to predict the judgments of less-informed decision makers [6]. People given disambiguating information about ambiguous sentences, like "the daughter of the man and the woman arrived," assume that the sentence would no longer be ambiguous to other naïve listeners [19]. In one particularly powerful demonstration, people were asked to tap the rhythm of a set of well-known songs, such as "Happy Birthday," on a table. The listeners had to guess the songs based on the rhythm tapped by the tappers. Tappers were then asked to estimate at what percentage those listeners would be able to correctly identify the songs. The tappers were confident, estimating that around 50% of the songs would be identifiable. In reality, listeners could only identify 2.5% of the songs, revealing a vast overconfidence in tapper estimates [27]. When people tap songs, their percussion does not include pitch, yet the auditory system fills in those pitches based on previous experience and knowledge. Critically, it seems impossible to 'turn off' this filling-in process, and people assume that others will have the same experience [30] such that simulating the experience of being naïve can be literally inconceivable.

While the curse of knowledge is well-studied in the psychology of decision making, language and education, there is less direct research on potential consequences for the processing of data visualizations. We nevertheless see manifestations of this cognitive bias in visual data communication. For example, imagine a scientist showing some graphs of experimental results at a conference or a colloquium, or a data analyst updating company leadership on recent customer feedback with snapshots of graphs from a dashboard. These people are undoubtedly experts in their respective fields, but nevertheless they overwhelm their audiences with overly complex graphs delivered too quickly.

Compared to numerical and textual formats, data visualizations are effective in highlighting the relationships and patterns in data to facilitate understanding [7]. But at the same time, understanding complex visualizations can be similar in time and effort to reading a paragraph [33,12,21]. Moreover, just like one can also read many possible sentences from a paragraph and interpret from it many different meanings, a graph or figure can be interpreted in multiple ways depending on where the viewer is fixating or selectively attending to [17,35]. These different readings and interpretations of visualizations (and text) are triggered by the knowledge of the communicator and the addressees. Given the primary role that visualizations play in the communication of analytic data, across science, education and industry [24,22,20], and the possibilities to see visualizations in multiple ways [36,33], it is important to demonstrate how knowledge influences visual perception of data visualizations and cause communication failures. We suspect that the inability to separate one's own knowledge and expertise from that of his/her audience can make visual data communication more difficult and less clear than one realizes. This may be especially true when the visualization contains potentially complex patterns. If certain data points on a graph are more likely to draw attention, this will also impair subsequent decision making processes [9,29].

In this paper, we demonstrate in the lab that the 'curse of knowledge' indeed exists in data visualization – knowledge makes an expert recognize a given pattern in data as more visually salient, and the expert assumes that it is also visually salient to naïve observers. This provides practical significance and theoretical importance to information visualization research, especially in visual data communication and decision-making.

## 2 GENERAL PROCEDURE & DESIGN

Participants read a story that conveyed background knowledge about a graph depicting political polling data. They were told that the experimenters will show the same graph they saw to 100 people, along with only a short description – "in the months before the elections of 2014 in a small European country, a polling organization asked citizens about their voting intentions on a daily basis."

They then predicted what uninformed viewers (with no knowledge of the story) would find to be the most visually salient features or patterns in the graph. The participants then predicted a second most salient feature, up to a fifth most salient feature. After writing down each feature they predicted, the participants also circled regions on the graph corresponding to each feature on a physical paper copy of the graph. They then reported how salient they think their five predicted features are on a scale from one to five, one being not at all visually salient and five being very visually salient. Finally, they matched their five predictions

* e-mail: cxiong@u.northwestern.edu
† e-mail: l.vanweelden@uu.nl
‡ e-mail: franconeri@northwestern.edu

as best as possible with five pre-determined features, as shown in Figure 1.

The experiments are within-subject, with the independent variable as feature congruency. A feature highlighted in the story is congruent, and un-highlighted incongruent. The participants are randomized to read one of three stories, each featuring one (set of) feature(s) among the five pre-determined features.

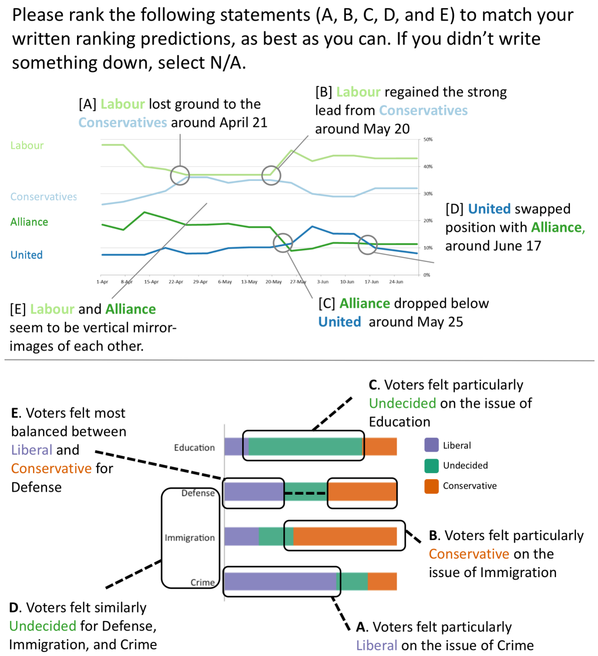


Fig. 1. Matching five pre-determined features in line graph experiment (top) and bar graph experiment (bottom).

## 3 Experiment 1 Line Graph

The participants read a story highlighting a competition between two among four political parties, illustrating how citizen voting intentions fluctuated with current events. Initially, between the two highlighted parties, one had a healthy lead in the polls. During an initial debate, the leading party lost voters to the less popular party and eventually lost the lead. In a later debate, however, the originally leading party was able to take back the votes the candidate lost and take the lead back again after a bad debate performance by his opponent. The three stories all describe this same competition, but ascribing it to the top two parties (Top-Prime Story), the top and third party (Middle-Prime Story) or the

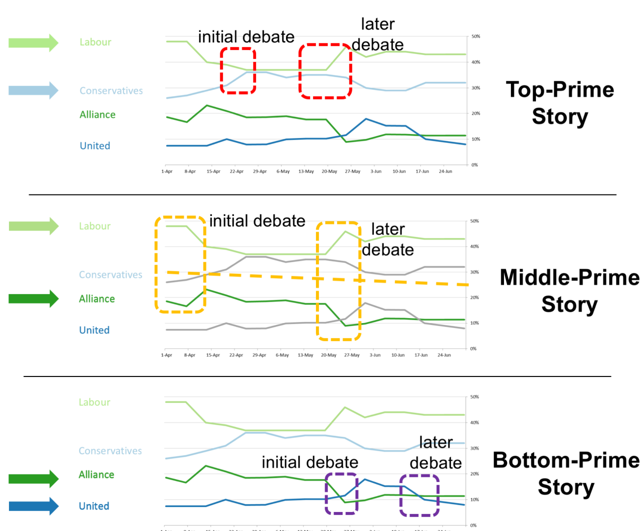


Fig. 2. Graph highlighting different features in line graph experiment.

bottom two parties (Bottom-Prime Story), highlighting the corresponding fluctuations. As shown in Figure 2, participants were shown polling data after reading the story. In each pair of lines, the party with the higher line cedes votes to the party with the lower line (initial debate), and then the higher line gains back that ground (later debate).

### 3.1 Qualitative Results

All stories, experimental stimuli, and data files are available at http://viscog.psych.northwestern.edu/VisualizationCurse2017/

In order to get a sense of what participants truly thought as salient without external suggestion, we examined the freely identified salient features they drew on physical copies of the unannotated graph, see Figure 3.

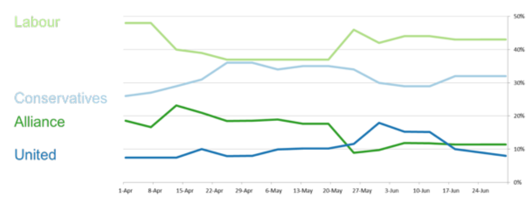


Fig. 3. The unannotated graph of the line graph experiment.

Figure 4 shows the top three predicted most salient features for each story. In the first three rows, within each feature block, the six graphs represent the responses from the six subjects for that feature in that story. The top left corner graphs correspond to the markings of subject one, the middle left graphs correspond to that of subject two, and so on. The individual markings can be collapsed for the three predicted most salient features by all six participants for each story, shown in the first three columns of the fourth row, from top to bottom are the most to third salient feature. These predictions can be further collapsed for each condition to illustrate their differences, shown in fourth column of the fourth row. Given that darker color represents more marking overlaps, we observed that participants who read the top-prime story mostly marked the top features as their top three visually

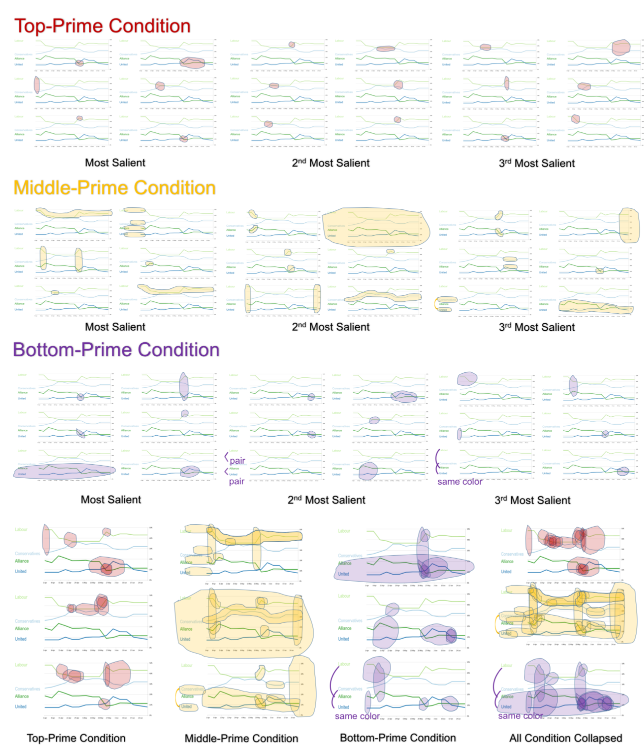


Fig. 4. Participants' prediction drawn on the graph.

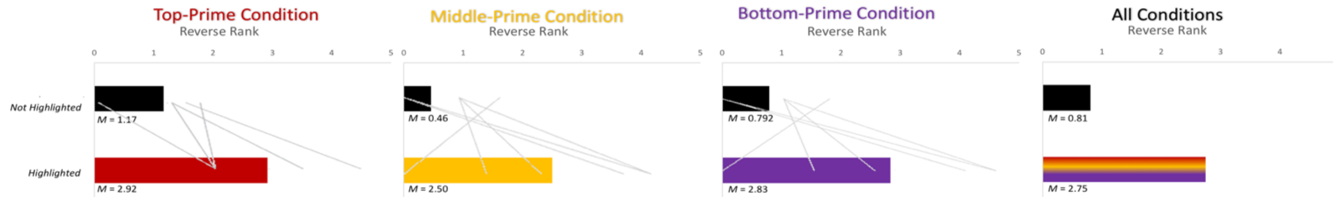


Fig. 5. Highlighted versus Not-Highlighted feature rankings (reversely coded for the figure only, reverse rank of 5 = actual rank of 1). The grey oriented lines represent individual subjects.

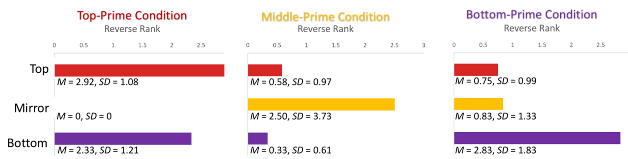


Fig. 6. Ranking details for each story in line graph experiment.

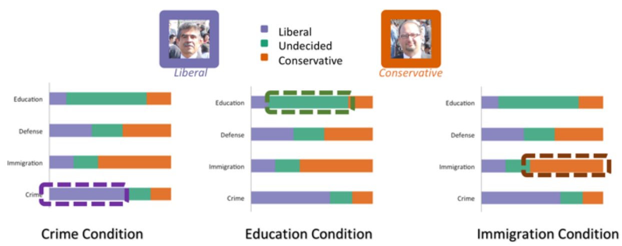


Fig. 7. Graph highlighting different features for bar graph experiment.

salient features to uninformed graph viewers; participants who read the middle-prime story mostly marked the mirroring features as their predicted top visually salient features; participants who read the bottom-prime story mostly marked the bottom features.

Overall, we find observable differences among the three conditions. These qualitative results seem to support our hypothesis such that across all conditions, participants predicted features that were depicted and highlighted in the story to be the most visually salient to uninformed viewers. To further support our hypothesis and our qualitative data, we conducted quantitative analysis.

### 3.2 Quantitative Results and Discussion

We conducted Wilcoxon Signed-Rank tests comparing the rankings between the highlighted and not highlighted features. We reversed the rankings for our analysis such that rank one (most salient) would get a score of five. Overall, as shown in Figure 5, the eighteen participants ranked highlighted features, $M$ = 2.75, significantly higher compared to the not highlighted features, $M$ = 0.81, W = 117, r = 0.76, $p < 0.01$. In Figure 6, detailed descriptive statistics are shown for all three stories (six subjects per story).

Additionally, using Spearman's Correlation, we found a strong relationship ($r_s = 0.55$, $p < 0.001$) between the self-rated salience of a feature, and the predicted salience rating for other naive observers. This indicates that the more visually salient a feature participants rated to an uninformed viewer, the more visually salient the participants think the feature was to themselves.

We then replicated this line graph experiment. We also conducted a follow-up study where no features depicted in the story was annotated for the participants (not described due to space limits). We found similar results in both.

## 4 Experiment 2 Bar Graph

We evaluated the generalizability of this curse of knowledge by replicating our findings using a bar graph. Participants read one of

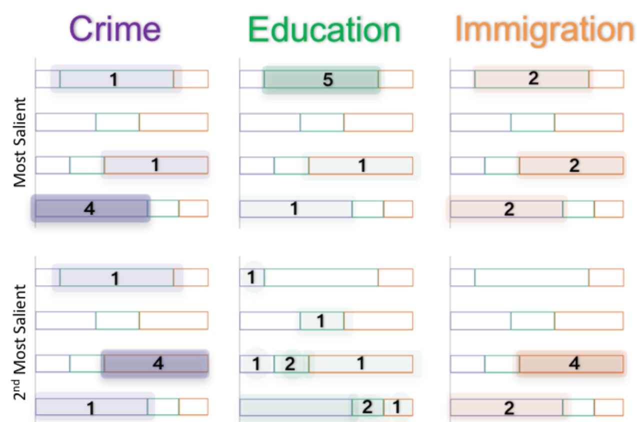


Fig. 8. Qualitative result of bar graph experiment.

three different backstories describing events leading to a presidential election between the Liberal and the Conservative parties. After the story, they were shown a public polling data highlighting the public opinion that eventually led to the victory of the winning candidate, as shown in Figure 7.

Participants read either a crime story, an immigration story, or an education story. The crime story portrays police brutality toward specific minority groups. The graph the participants saw corresponded to the story highlighting the majority's liberal public opinion of crime, explaining it as the reason behind the Liberal Party's victory. The immigration story describes a terrorist attack. The graph the participants saw corresponded to the story highlighting the majority's conservative public opinion on immigration, explaining it as the reason behind the Conservative Party's victory. The education story illustrates a debate between the Liberal and Conservative Parties on the country's education system. Neither candidate could come up with a clear vision on how to improve the system. This opened an opportunity for a third candidate, who was an expert on education. The graph the participants saw highlights the fact that most people in the country had been undecided (neither liberal nor conservatives) on the issue of education, opening the opportunity for the third candidate.

### 4.1 Qualitative Results

There are observable differences in the order of salient feature predictions for the three stories. Figure 8 shows the features the seventeen participants circled as salient to another viewer on an unannotated bar graph, shown in Figure 9.

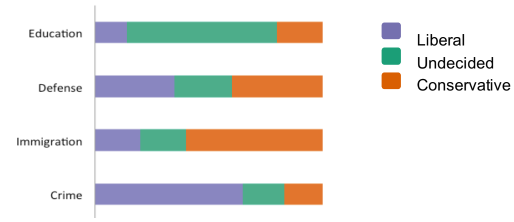


Fig. 9. Unannotated bar graph.

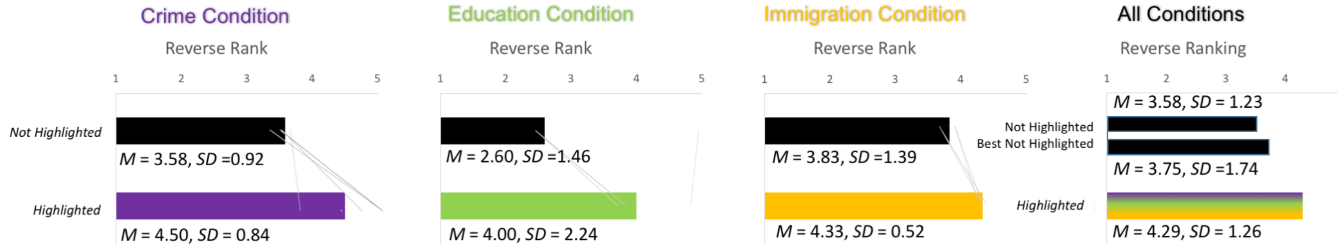


Fig. 10. Highlighted versus Not-Highlighted feature ranking for bar graph experiment, reverse coded. The grey oriented lines represent individual subjects.

In Figure 8, each column represents a different story. The first and second rows show the most and second most visually salient predicted features for the three stories respectively. The numbers on the graph represent the number of times the shaded region was chosen to be visually salient to another viewer. The darker the shading of a feature, the more frequently it was chosen to be visually salient. Overall, the participants generally circled the feature that has been highlighted in the story they read as the more salient feature. This indicates that they generally predict that other uninformed graph viewers think the feature they read about to also be more visually salient.

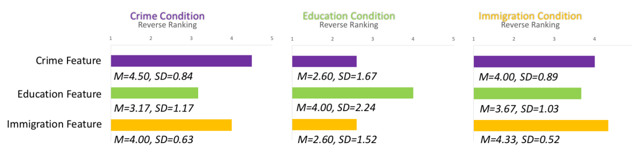


Fig. 11. Ranking details for each story in bar graph experiment.

### 4.2 Quantitative Results

Wilcoxon Signed-Rank Test indicated that the overall features highlighted in the story (reversely ranked), $M = 4.29$, were statistically significantly ranked higher and more salient than the overall not highlighted features, $M = 3.58$, W = 150, r = 0.98, $p < 0.001$, as shown in Figure 10. The descriptive statistics are shown for the three stories in Figure 11. These support that the highlighted features were predicted to be more visually salient to uninformed viewers than features not highlighted in the story.

We also found significant correlation between predicted features' saliency ranking and self-rated saliency of these features using Spearman's Correlation, $r_s = 0.52$, $p < 0.001$, indicating that participants predicted features visually salient to themselves were also more salient to uninformed viewers.

## 5 Conclusion

The experiment demonstrated that knowledge the participants obtained by reading the story biased their predictions such that, in general, they saw the features depicted in the story as more visually salient than features not depicted in the story. More importantly, after acquiring this background knowledge, participants were biased to predict that other uninformed graph viewers would rate those features as more visually salient as well. This cognitive bias occurred despite explicit instructions to ignore what they knew, and to take a naïve perspective. To our knowledge, this is the first empirical demonstration of the curse of knowledge in the realm of visual perception.

These results join other recent findings of the influence of perceptual and cognitive biases on interpretations of patterns in data visualization. Other work has shown an influence of the 'attraction effect' – a cognitive bias where irrelevant information can influence decisions about otherwise equal alternatives – can manifest in the perception of visualized data [8,13]. A preference for salient visuals and distinctive designs can determine whether a data visualization keeps people engaged [2,3,11] and is remembered as being previously viewed [4,5]. Storytelling techniques adapted from those employed for writing and more cognitive tasks can have affect the way that we extract data from visualizations [14,15,25,26,28,32]. Data visualizations are an ideal testbed for such biases, given their importance as a tool for information exploration, engagement, and understanding.

## 6 Future Directions

The curse of knowledge is tough to detect and inhibit. Critique provides a feedback loop for what is communicated, and what is not, making it a critical tool to help the data visualizers see more clearly the strength and shortcomings of their visual data communication and then make appropriate revisions [23,34].

The curse of knowledge may also lead viewers to become fixated on given patterns in a dataset, leaving them less likely to see new or alternative patterns. As the design of a visualization can strongly influence what comparisons are made (e.g., people are more likely to compare proximal values, or values that are depicted with the same line in a line graph [33]), using a variety of designs might help 'kick people out' of a given perspective on patterns in a dataset. For example, plotting data in different arrangements and formats might force the viewer to see new patterns in their own data.

Presenters, paper authors, and data analysts can fail to connect with audiences when they communicate patterns in data. The present results provide an empirical demonstration that the curse of knowledge may be largely to blame. There may be inspiration for attenuating this problem within research in perspective taking, which has shown that people predict strangers' reactions more accurately through projecting themselves onto the stranger [10,37]. Strengthening interactions between presenter and audience may help presenters gauge the most effective way to communicate without overwhelming their audiences [32,15,16].

## References


[1] D. W. Allbritton, G. McKoon, and R. Ratcliff. Reliability of prosodic cues for resolving syntactic ambiguity. Journal of Experimental Psychology: Learning, Memory, and Cognition, 22(3):714–735, 1996.

[2] S. Bateman, R. L. Mandryk, C. Gutwin, A. Genest, D. McDine, and C. Brooks. Useful junk?: the effects of visual embellishment on com- prehension and memorability of charts. In Proceedings of the 28th Inter- national Conference on Human Factors in Computing Systems, CHI '10, pages 2573–2582. ACM, 2010.

[3] 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. Visualization and Computer Graphics, IEEE Transactions on, 18(12):2759–2768, 2012.
[4] B 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] M. A. Borkin, Z. Bylinskii, N. W. Kim, C. M. Bainbridge, C. S. Yeh, D. Borkin, ... and A. Oliva. Beyond memorability: Visualization recognition and recall. IEEE transactions on visualization and computer graphics, 22(1):519-528, 2016.
[6] C. Camerer, G. Loewenstein, and M. Weber. The curse of knowledge in economic settings: An experimental analysis. Journal of political Economy, 97(5):1232-1254, 1989.
[7] S. K. Card, J. D. Mackinlay, and B. Shneiderman, editors. Readings in Information Visualization: Using Vision to Think. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 1999.
[14] navigation in information visualization: A survey. IEEE Transactions on visualization and computer graphics, 6(1):24-43, 2000.
[15] J. Hullman, E. Adar, and P. Shah. Benefitting infovis with visual difficulties. Visualization and Computer Graphics, IEEE Transactions on, 17(12):2213–2222, 2011.
[16] J. Hullman and N. Diakopoulos. Visualization rhetoric: Framing effects in narrative visualization. IEEE transactions on visualization and computer graphics, 17(12):2231-2240, 2011.
[17] J. Hullman, S. Drucker, N. H. Riche, B. Lee, D. Fisher, and E. Adar. A deeper understanding of sequence in narrative visualization. IEEE Transactions on Visualization and Computer Graphics, 19(12):2406-2415, 2013.
[18] N. Kawabata. Attention and depth perception. Perception, 15(5):563–572, 1986.
[19] D. S. Kerby. The Simple Difference Formula: An Approach to Teaching Nonparametric Correlation 1. Comprehensive Psychology, 3(1):1-14, 2014.
[20] B. Keysar and A. S. Henly. Speakers' overestimation of their effectiveness. Psychological Science, 13(3):207-212, 2002.
[21] M. Khan and S. S. Khan. Data and information visualization methods, and interactive mechanisms: A survey. International Journal of Computer Applications, 34(1):1-14, 2011.
[22] W. Kintsch. The role of knowledge in discourse comprehension. A construction-integration model. Psychological Review, 95(2):163–182, 1988.
[23] C. N. Knaflic. Storytelling with data: a data visualization guide for business professionals. John Wiley & Sons, 2015.
[24] R. Kosara, F. Drury, L. E. Holmquist, and D. H. Laidlaw. Visualization criticism. IEEE Computer Graphics and Applications, 28(3):13-15, 2008.
[25] G. McKenzie, T. J. Barrett, M. Hegarty, W. Goodchild, and W. Thompson. Assessing the Effectiveness of Visualizations for Accurate Judgements of Geospatial Uncertainty. In Visually-Supported Reasoning with Uncertainty Workshop, Conference on Spatial Information Theory. COSIT '13; Scarborough, UK.
[26] A. L. Michal and S. L. Franconeri. Visual routines are associated with specific graph interpretations. Cognitive Research: Principles and Implications, 2(1):20, 2017.
[8] E. Dimara, A. Bezerianos, and P. Dragicevic. The attraction effect in information visualization. IEEE Transactions on Visualization and Computer Graphics, 23(1):471-480, 2017.
[9] R. Duclos. The psychology of investment behavior:(De) biasing financial decision-making one graph at a time. Journal of Consumer Psychology, *25*(2), 317-325,2 015.
[10] N. Epley and A. Waytz. Mind perception. In S. T. Fiske, D. T. Gilbert and G. Lindzey (Eds.), The handbook of social psychology (Vol. 2, 5th ed., pages 498–541). New York, NY: Wiley, 2010.
[11] S. Haroz, R. Kosara, and S. L. Franconeri. Isotype visualization: Working memory, performance, and engagement with pictographs. In Proceedings of the 33rd annual ACM conference on human factors in computing systems, ACM, pages 1191-1200, 2015.
[12] M. Hegarty. Multimedia learning about physical systems. In R. E. Mayer (Ed.), Handbook of multimedia, pages 447–465. New York: Cambridge University Press, 2005.
[13] I. Herman, G. Melançon, and M. S. Marshall. Graph visualization and
[27] A. V. Moere, M. Tomitsch, C. Wimmer, B. Christoph, and T. Grechenig. Evaluating the effect of style in information visualization. IEEE transactions on visualization and computer graphics, 18(12):2739-2748, 2012.
[28] L. Newton. Overconfidence in the communication of intent: Heard and unheard melodies. Unpublished doctoral dissertation, Stanford University, Stanford, CA 1990.
[29] A. V. Pandey, A. Manivannan, O. Nov, M. Satterthwaite, and E. Bertini. The persuasive power of data visualization. IEEE transactions on visualization and computer graphics, 20(12):2211-2220, 2014.
[30] P. Raghubir, S.R.Das. The long and short of it: Why are stocks with shorter runs preferred?. Journal of Consumer Research, *36*(6), 964-982, 2010.
[31] L. Ross, and A. Ward. Naïve realism in everyday life: Implications for social conflict and misunderstanding. In E. Reed, E. Turiel, & T. Brown (Eds.), Social cognition: The Ontario Symposium, pages 305–321. Hillsdale, NJ: Erlbaum, 1996.
[32] W.-M. Roth, C. J. McRobbie, K. B. Lucas, and S. Boutonné. Why May Students Fail to Learn from Demonstrations? A Social Practice Perspective on Learning in Physics. Journal of Research in Science Teaching. 34(5):509-533, 1997.
[33] E. Segel and J. Heer. Narrative visualization: Telling stories with data. IEEE transactions on visualization and computer graphics, 16(6):1139-1148, 2010.
[34] P. Shah and E. G. Freedman. Bar and line graph comprehension: An interaction of top-down and bottom-up processes. Topics in Cognitive Science, 3(3):560-578, 2011.
[35] T. Skog, S. Ljungblad, L.E. Holmquist, "Between Aesthetics and Utility: Designing Ambient Infor- mation Visualizations," Proc. Information Visualization, IEEE CS Press, 2001, pages 30–37.
[36] Y. Xu and S. L. Franconeri. The head of the table: Marketing the "front" of an object is tightly linked with selection. The Journal of Neuroscience, 32(4):1408-1412, 2012.
[37] A. L. Yarbus. Eye movements and vision. Springer, 1967.
[38] H. Zhou, E. A. Majka, and N. Epley. Inferring Perspective Versus Getting Perspective: Underestimating the Value of Being in Another Person's Shoes. Psychological Science, 2017. doi:10.1177/0956797616687124.