OPAM 2017

The Power of Top-Down Salience in Data Visualizations

Cindy Xiong, Lisanne van Weelden, Steven Franconeri

Abstract

The duck-rabbit and Necker cube illusions reveal that the visual system can lock into a single view of a multi-stable percept. Such ambiguity is rare in the natural world, but it is ubiquitous in the artificial world of information visualizations – graphs, for example, contain many perceivable patterns. We told participants stories that explained fluctuations in simulated political polling data, which strongly increased the relevant pattern's visual salience ratings. Critically, they believed that naïve viewers would see the same salient patterns, revealing a 'curse of knowledge' that may underlie failures to effectively communicate data patterns to others.

Introduction

The visual system seeks meaning in the image presented by the world, by matching visual input to hypotheses about object identities and spatial relations. Sometimes, as demonstrated in the duck-rabbit and Necker cube illusions, multiple hypotheses might exist (Attneave, 1971), leading two people having different interpretations of the same image.

While such ambiguity is actually rare in the natural world, it is ubiquitous in the artificial world of information visualizations. A graph, for example, contains many potential perceivable patterns. After viewers extract a set of ensemble statistics about the dataset (Szafir et al., 2016), they must exercise top-down attentional control (Egeth et al., 2010) to extract a series of relationships and patterns from the data values (Michal et al., 2016; 2017), with an equivalent time and effort to reading a paragraph (Shah and Freedman, 2011). The diversity of percepts extractable from patterns in graphs might lead to two people seeing different patterns in the same graph - one the 'duck' and the other the 'rabbit'.

The present experiment shows that top-down cues (a verbal description of a particular pattern) can increase visual salience for that particular pattern, driving it so strongly that participants could not turn it 'off'. After giving a verbal description of a particular pattern, we asked participants to imagine the graph from a naïve perceptive, as when they first saw it, and to predict what a naïve viewer would find salient. Participants could not recover an unbiased percept of the image, and they strongly predicted that others would see the same newfound salience that they saw.

Method

Participants read a story that explained fluctuations in simulated political polling data. They then predicted what naïve viewers (with no knowledge of the story) would find to be the top five most visually salient features or patterns in the graph. After writing down each feature they predicted, participants also circled regions on the graph corresponding to each prediction on a paper copy of the graph. They then reported how salient they think their five predicted features are to *themselves*, on a scale from one to five, one being not at all visually salient and five very visually salient. Finally, to facilitate quantitative analyses, they categorized their sketched predictions to match five pre-determined features, as shown in Figure 1a.

The experiments were within-subject, with the independent variable as feature congruency. A feature highlighted in the story was congruent, and un-highlighted incongruent. The participants were randomized into three groups (top, bottom, or mirror), each highlighting one set of features among the five pre-determined features.

The story highlighted competition over time between two candidates among four political parties, depending on a participant's randomized group. It illustrated how citizen voting preferences fluctuated with current events, maintaining qualitatively consistent with all three groups. As shown in Figure 1b, initially, between the two highlighted parties, one had a healthy lead in the polls. During an initial debate, the leading party turned voters over to the less popular party and eventually lost the lead. In a later debate, however, the originally leading party was able to regain the votes the candidate lost and took the lead back after a bad debate performance by his opponent. 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).

Results

The regions participants circled on the graph are shown in Figure 1c, revealing how the salience predictions of participants were strongly driven by the stories they read – and were instructed to forget. A Wilcoxon Signed-Rank tests comparing the rankings between the highlighted and non-highlighted features showed that the eighteen participants predicted highlighted features significantly more salient compared to the non-highlighted features, W=117, r=0.76, $p<0.01$, for the eyes of naïve viewers, see Figure 1d. Spearman's Correlation revealed 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 naïve observers.

We replicated this line graph experiment, and conducted another conceptual replication using different response methods, and found similar results for all three instances. We then replicated our findings using bar graphs and a different set of stories. Overall congruent feature ranks were statistically significantly higher than the overall incongruent feature ranks, W=150, r=0.98, $p<0.001$. Spearman's Correlation, $r_s=0.52$, $p<0.001$, again indicated that participants predicted features visually salient to themselves to also be salient to naïve viewers.

Conclusion

A graph presents many interpretations. These experiments show that top-down cues can guide these interpretations – and more importantly, they are tough to turn off once a given pattern is highlighted, as shown by our participants' inability to imagine the mindset of a novice when predicting the graph salience ratings of a naïve viewer.

This phenomenon presents a novel example of the "curse of knowledge", but in a perceptual domain. The curse, a widely studied bias from decision making, communication, and education research communities, is an inability of experts to 'turn off' relevant knowledge, adopt or even simply predict the mindset of a novice. Knowledgeable business decision makers fail to predict the judgments of uninformed decision makers (Camerer et al, 1989). People given disambiguating information about syntactically ambiguous sentences, such as "Angela shot the man with the gun," assume that the sentence is not ambiguous to naïve listeners (Keysar and Henly, 2002). Song tappers overestimate the probability others would correctly guess the songs they tapped (Newton, 1990).

Like these results, the present data suggest that while people may indeed see different patterns in the same graphed dataset, they likely would not *realize* the difference in percepts. This may explain how conference speakers and paper writers can overwhelm audiences with graphical depictions that make sense to the presenter, yet leave the audience confused.

References

Attneave, F. (1971). Multistability in perception. *Scientific American*.

Camerer, C., Loewenstein, G., & Weber, M. (1989). The curse of knowledge in economic settings: An experimental analysis. *Journal of political Economy*, *97*(5), 1232-1254.

Egeth, H. E., Leonard, C. J., & Leber, A. B. (2010). Why salience is not enough: Reflections on top-down selection in vision. *Acta psychologica*, *135*(2), 130.

Keysar, B., & Henly, A. S. (2002). Speakers' overestimation of their effectiveness. *Psychological Science*, *13*(3), 207-212.

Newton, L. (1990). Overconfidence in the communication of intent: Heard and unheard melodies. Unpublished doctoral dissertation, Stanford University, Stanford, CA.

Michal, A. L., & Franconeri, S. L. (2017). Visual routines are associated with specific graph interpretations. *Cognitive Research: Principles and Implications*, *2*(1), 20.

Michal, A. L., Uttal, D., Shah, P., & Franconeri, S. L. (2016). Visual routines for extracting magnitude relations. *Psychonomic bulletin & review*, *23*(6), 1802-1809.

Shah, P., & 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.

Szafir, D. A., Haroz, S., Gleicher, M., & Franconeri, S. (2016). Four types of ensemble coding in data visualizations. *Journal of vision*, *16*(5), 11-11.