Three perceptual tools for seeing and understanding visualized data

Steven L. Franconeri
Department of Psychology, Northwestern University

Address correspondence to:
franconeri@northwestern.edu
2029 Sheridan Rd, Evanston, IL, 60208

In press, *Current Directions in Psychological Science*

Abstract

The visual system evolved and typically develops to process the scenes, faces, and objects of the natural world. We have adapted this powerful system to process data within an artificial world of visualizations. Viewers appear to use at least three types of perceptual tools to extract patterns in data from these artificial displays, including a tool that extracts global statistics, one that extracts shapes within the data, and one that produces sentence-like comparisons. A better understanding of the power, limits, and deployment of these tools would lead to better guidelines for display design and pedagogy, across classrooms, information dashboards, presentations, and public policy explanations.

Keywords: Data visualization, Graphical Perception, Visual communication, Graph comprehension, Ensemble perception

Introduction

The table at the top of Figure 1 shows four types of data for four countries in 2015: GDP per capita, life expectancy, population, and CO2 emissions per capita. Searching for patterns in those data feels sluggish, because you can only read in a handful of numbers per second. And as you compare one handful of numbers to another, your limited working memory capacity causes the original numbers seem to slip away, or to be overwritten by the new values. Now inspect the data for those same countries in the visualization at the center of Figure 1, which depicts all of those data, plus additional data for GDP and Life Expectancy for the year 1985 (via the 'comet' lines). You should feel the processing power of a visual system that occupies around 40% of your brain (Van Essen, 1992), allowing you to rapidly extract statistics across those countries (e.g., China and India have the highest populations; average life expectancy across all countries is around 74 years), to note the roughly linear shape of the arrangement of the bubbles (i.e., GDP is positively correlated with life expectancy), and to efficiently browse specific comparisons among them (e.g., Bangladesh shows a slightly faster rise in life expectancy compared to India).

Data visualizations are powerful tools for thinking and communicating with data. Understanding and creating these visualizations are now critical skills for an educated public, even arguably similar in importance to reading and writing (Börner et al., 2019). To harness that power, a viewer must understand how a set of visual marks like bars, lines, or bubbles can depict data, and how patterns among those marks reflect patterns in the data. The following sections argue that viewers extract patterns and relationships from visualized data using at least three core perceptual tools – statistics, shapes and comparisons – and demonstrate how the choice of perceptual tool can influence which patterns are seen and how they are understood.

Year	Country	GDP/person	Life Expect.	Pop. (M)	CO2/person (T)
2015	Bangladesh	$3,230	73	156	0
2015	Brazil	$14,800	76	204	2
2015	China	$13,500	77	1,410	7
2015	India	$5,740	68	1,310	2
2015	Indonesia	$10,400	71	258	2
2015	Mexico	$17,500	76	122	4

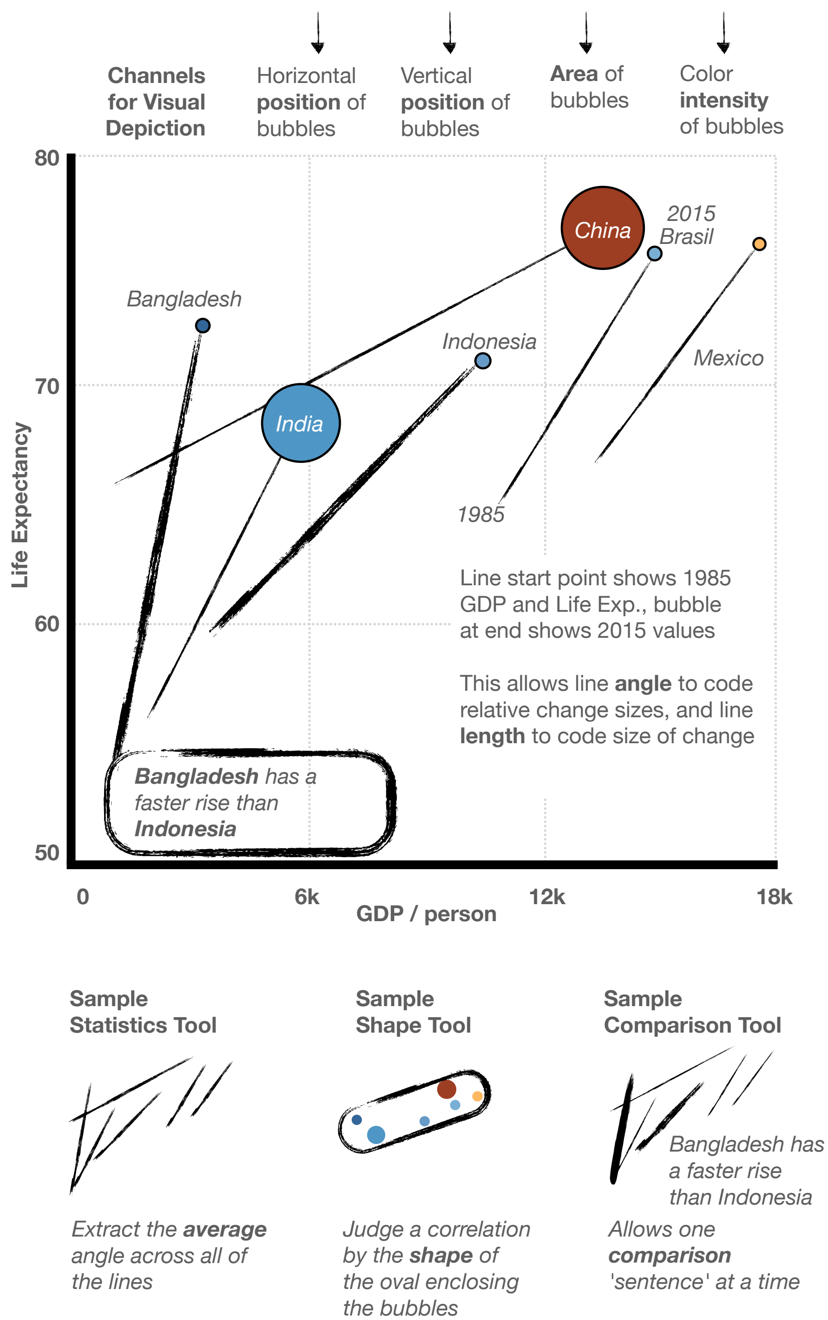

Figure 1: At top, browsing patterns in data are slow within a table. The arrows below show how each column of the 2015 data is mapped to a visual channel in the visualization, which also adds more data about the GDP and Life Expectancy of each country for an earlier year of 1985, via the 'comet tail' lines. At bottom, examples of how each of the three perceptual tools could help a viewer extract patterns. Data courtesy of Gapminder.org. Inspired by the work of Hans Rosling, https://www.ted.com/talks/Hans_rosling_the_best_stats_you_ve_ever_seen/

Perceptual Tool 1: Visual Statistics

Data visualizations map numeric data to a set of visual channels, including position, length, area, numerosity, angle, and intensity (Munzner, 2014). Figure 1 highlights how a single visualization can recruit every one of these channels to depict several types of data simultaneously. Each of these channels has a natural mapping between its visual appearance and the data value that it reflects – for example, higher positions or longer lines are easily processed and understood as reflecting larger values (Tversky, 2001).

In Figure 1, you can quickly extract statistics across any of the channels. Among all of the plotted countries, what is the lowest GDP coded as horizontal position, the largest population coded as each bubble's area, or the average angle of all of the tails of each 'comet' (Figure 1, lower left)? Within a split second, the visual system can process any of these channels in parallel across a two-dimensional page or screen, allowing a viewer to find values or regions in a visualization that differ from others (Wolfe & Horowitz, 2004) or to extract global statistics across those values with impressive precision (Alvarez, 2010). These feats are possible because they adapt the existing ability of the visual system to quickly extract similar statistics from natural scenes, helping you identify the lowest apple on the tree, the largest slice of pie, or the average angle of ocean waves. These visual statistics, therefore, serve as a proxy for actual statistics, allowing you to extract statistical information from data values that have been mapped to visual channels (for review, see Szafir, et al., 2014).

Note that for simple value ratio computations (Jacob, Vallentin & Nieder, 2012) some channels transmit data values more precisely than others (Cleveland & McGill, 1985). For example, in Figure 1 you can easily tell that the ratio of GDP/person between Bangladesh and Indonesia is around 1:3, because horizontal position is precisely extracted on a perceptual level. You are less sure of the ratio between the CO_2 production of Mexico and China, because the color intensity channel does not allow for precise perception of its underlying values. While this ranking may encourage visualization designers to prefer more precise channels (e.g., position) over less precise ones (e.g., grayscale intensity), this ranking was originally created based on performance in a single task: the precision of a ratio judgment task between values at a time. When the task is changed to statistical judgment that requires a viewer to isolate a single value (e.g., find the minimum or maximum) across data shown with a position channel (e.g., a line graph) or a color intensity channel (e.g., a heatmap), the ranking appears to hold. But for other more global statistical judgments of the mean or variance of data, the pattern of precision dominance disappears or even flips, with color intensity now beating position (Albers et al., 2014). Figure 2 depicts a real-world example of a design that purposely uses color intensity to focus the viewer on big-picture statistics, instead of a precise reading of individual values.

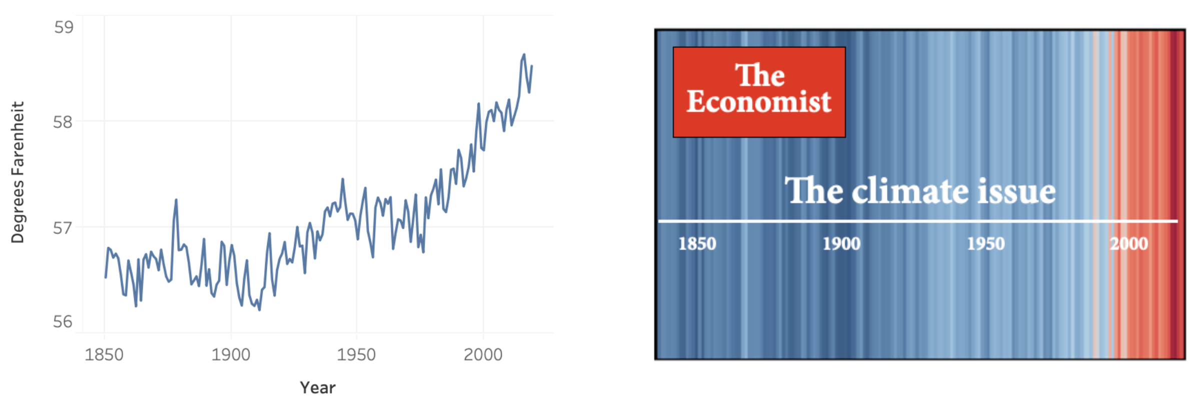

Figure 2: At left, global climate data precisely encoded as a line graph. At right, the September 2019 issue of The Economist used a color intensity channel to depict the same data with far less precision for each value, focusing the viewer toward the bigger picture.

Despite a large research literature on the processing of visual statistics (see Alvarez, 2010 for review), most previous studies focus on human performance within the natural world of scenes, objects, and faces. Recent work in the data visualization field has begun to adopt these research paradigms to test performance in the artificial world of data visualizations, with the eventual goal of producing a holistic understanding of the precision – and potential biases – of visually-recovered statistics (mean, variance, outliers) across actual visualization displays (Szafir, et al., 2014)

Perceptual Tool 2: Shape Processing

Viewers can also directly map known shapes to previously seen patterns in data. Figure 3 (left) depicts average temperatures across two cities. An experienced graph reader will immediately notice the X-shaped pattern in the data, which lets them quickly understand that the seasons are ‘flipped’ in relative temperature across the cities. Similarly, in a scatterplot (Figure 1, bottom center panel) people appear to use the shape of the oval that encloses most of the points as a proxy for correlation (Yang et al., 2020).

Shape processing can also lead less experienced graph readers astray. In Figure 3 (right), students were asked to give a real-world example that might produce the graphed data. A correct answer would be ‘an object does not move for a period of time, then it moves at a constant rate, and then it stops’. But shape processing led many students to see the graph as a natural scene of a hill, guiding them toward answers such as “An object rolls along a flat surface, then down a hill, and then it stops.’ (Kozhevnikov et al., 2002).

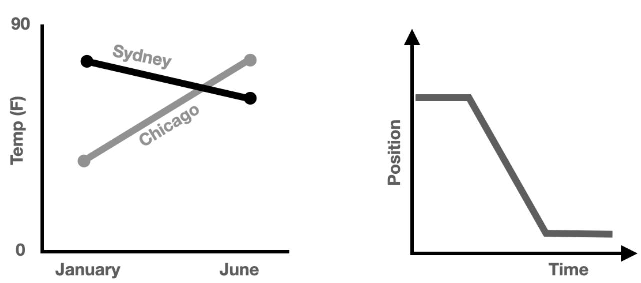

Figure 3: At left, graph depicting temperatures for Sydney and Chicago in January and June. Noticing an X-shape allows for quick judgment of a critical pattern in the data. At right, an example from Kozhevnikov et al., 2002 that shows how shape processing can lead students astray.

Figure 4 shows examples where shape processing can lead to misunderstandings of statistical uncertainty. The first column shows two examples where statistical estimates of uncertainty are depicted as shapes (a 'cone of uncertainty' for a hurricane; error bars for an election forecast). Translating those shapes into an intuitive sense of uncertainty for the viewer requires specialized statistical training. Without training, these depictions can lead to serious misunderstandings, such that the viewer might interpret the graph as illustrating that the hurricane is getting larger, instead of having a more uncertain path (Padilla, et al., in press). The second column shows graphical redesigns that allow the viewer to leverage their perceptual statistics tool, and that communicates that uncertainty in a more naturalistic way (Padilla, et al., in press).

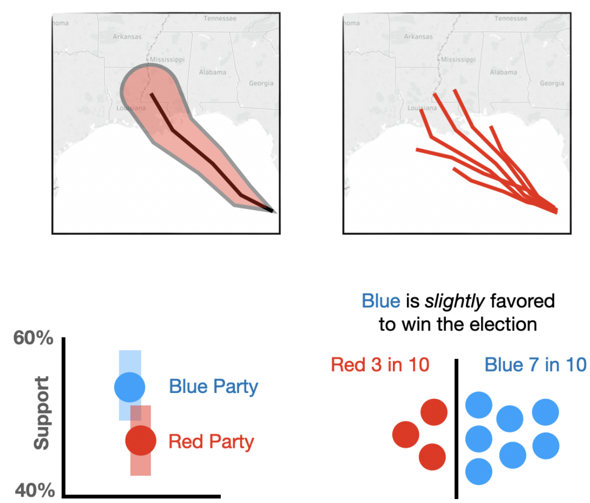

Figure 4: On the top row (left), viewers can substantially misinterpret hurricane forecasts that convey model uncertainty with a shape envelope reflecting statistical variation among different models. Allowing observers to extract naturalistic statistics from a set of example model paths (top right) can lead to better understanding of that uncertainty (Liu, et al., 2018). On the bottom row (left), uncertainty in election forecasts shown via similar shape envelopes (error bars) is tough to understand without statistical training. At bottom right, fivethirtyeight.com redesigned their forecast to depict a more naturalistic statistical metaphor of 'possible outcomes' for the 2020 US presidential election. This design leverages our natural ability to extract statistical summaries from collections of objects.

Perceptual Tool 3: Comparisons

Figure 1 allows you to efficiently browse specific comparisons among individual countries and groups of countries (e.g., Bangladesh has a faster rise in Life Expectancy compared to India; Mexico produces more GDP/person than India). Note that each of these comparisons is expressed as a sentence. The *direction* of each comparison sentence is psychologically important (Miller & Johnson-Laird, 1976). While logically equivalent, one is unlikely to say that 'the building is to the right of the bike' compared to 'the bike is to the left of the building'. Or compare the excitement of "I met *Barack Obama"*, to the hubris of "Barack Obama met *me*" (Gleitman, et al., 1996). Visual comparisons, including those within data visualizations, appear to have a similar directional structure. Moving your eyes or attention from Mexico to India might tell you that, in 2015, Mexico had a higher GDP/person, while the opposite pattern might tell you that India had a lower GDP/person. When a viewer's attention is dragged around a screen by animated cues (Roth & Franconeri, 2012), or freely roams but is monitored with an eyetracker (Michal & Franconeri, 2016, 2017), viewers are faster to verify the truth of sentences (e.g., Mexico has a higher GPD/person than India) if their attention or eyes had traveled between the values in that directional order.

How many unique sentences, and therefore unique comparisons, could you produce from Figure 1? Starting with only two countries, it is clear that there will be dozens of possible comparisons: (e.g., China is larger than Brazil, China has a lower GDP/person than Brazil) (see Michal & Franconeri, 2016 for 12 sample comparisons from a single 2-bar graph). The number of comparisons would be enormous, given the number of possible grouping combinations within the six displayed countries, the four data types, and the additional data for 1985.

Most models of such relational comparisons predict that you can only process a handful of these relational comparisons at a time (Hummel & Biederman, 1992) or some predict that limit is only *one* at a time (see Franconeri, 2012 for review). One recent study validated that this severe capacity limitation extends to data visualizations. Among pairs of 2-bar graphs (bar A & B) distinguishing bar pairs with one relationship (A>B) from the opposite relationship (A<B) was extremely slow, compared to processing the global statistics of similar displays (which single bar is largest?) which could be accomplished instantly (Nothelfer & Franconeri, 2019).

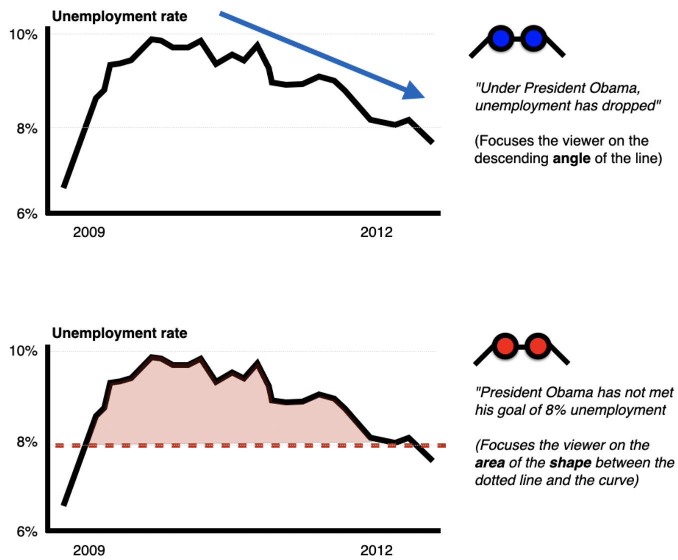

Figure 5: A viewer can make multiple comparisons within visualized data, and the choice of those comparisons may be as consequential as the choice of verbal descriptions for comparisons within data. These data can be seen from different perspectives, using unemployment rates under US president Obama as an example. A member of Obama's party (blue) might highlight a positive pattern, while a member of the opposite party (red) might highlight a negative pattern. Adapted from *https://archive.nytimes.com/www.nytimes.com/interactive/2012/10/05/business/economy/one-report-diverging-perspectives.html*

Because comparisons are so slow, visualization designers often try to help viewers compare the ‘right’ values. One strategy is to use visual grouping cues that encourage the viewer to process those values together (Shah & Hoeffner, 2002). Examples of grouping clues include placing the values to be compared close to one another, making them the same color, or connecting them with a line or contour (Brooks, 2015). Because this practice still leaves a viewer with many potential comparisons to make, practitioner guides also encourage designers to highlight critical values with a unique color to draw attention to them, and furthermore to annotate that pattern with a single comparison ‘sentence’ (as in the ‘Bangladesh / Indonesia’ comparison in Figure 1).

This practice is now widely used by professional data communicators across business and journalism. The power of these highlighting and annotation techniques should be clear from Figure 5, where competing designers might draw you toward different comparisons, which can lead you toward importantly different conclusions from the data. With this great power to frame the 'right' patterns for an audience comes great responsibility for doing so ethically. The research literature has examined the persuasive power of such manipulations, finding that highlighting and annotating a given pattern makes it far more likely to be remembered by a viewer (Adjani, et al., in press). A practitioner literature curates real-world examples, primarily from journalism, of more and less 'truthful' designs that vary across definitions of what 'truthful' means (Cairo, 2016).

Some creators of visualization might feel these additional steps are unnecessary – if they see the pattern in their own visualization, won't others see it too? Unfortunately, designers, authors, teachers, and presenters can be afflicted with 'curse of expertise', where they fail to distinguish between what they understand, and what an audience understands (Birch & Bloom, 2007). Once an author recognizes an important comparison in a dataset, they focus on that comparison automatically and assume that new viewers do too. As evidence, in a recent study we told participants an interesting story about an data pattern in a simulated election polling dataset, either for the top two lines, or for the bottom two lines (Figure 6, left). They were then told to forget the story and to annotate the visual pattern that a *naive* viewer (who did not know the background story) would find most salient. If inhibiting their own knowledge were possible, there should be no difference in the annotated patterns across the 'top' and 'bottom' story conditions. Instead, participants tended to report that naive viewers would find the same pattern (either top or bottom) most salient (actual participant drawings shown in Figure 6 at right) (Xiong, et al., 2019). Clearly, their expertise served as a 'curse' that affected their predictions of what others would see in the data.

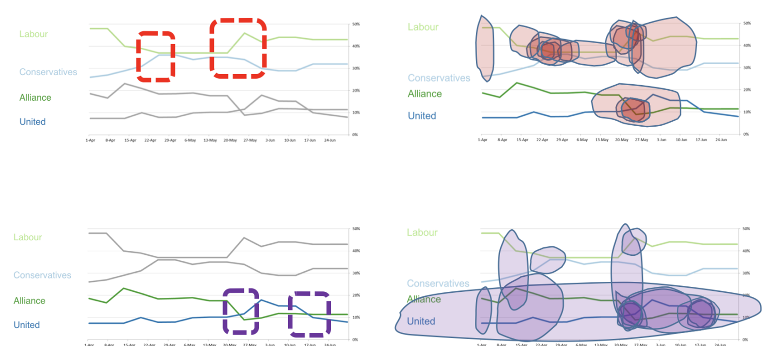

Figure 6: Stimuli (left) and Results (right) for Xiong et al.. (2019)

Conclusion

The research reviewed here merges psychological theory, evidence, and methods with research and practice in data visualization. It suggests a set of at least three tools that we rely on to process visualized data: statistical extraction, shape processing, and relational comparison. Viewers can effectively use these tools to explore data, and clearly communicate patterns within it – but only when the visualization is designed to effectively leverage these tools, and when the viewer is properly trained to know which tool to use, and when to use them.

Understanding the power, limits, and deployment of these tools can help inspire guidelines for display designs and pedagogy that should lead to faster perception and deeper understanding of the critical patterns depicted by a data visualization. A better understanding would help even more. The statistics tool allows viewers to rapidly extract means or outliers, but we lack a full understanding of the precision and potential biases of those operations. When deployed appropriately, the shape tool can provide fast and holistic view of the data for simple shapes, yet we currently have almost no understanding of how we process complex shapes like the dance of a line graph of a company's stock price. The comparison tool extracts single comparisons at a time, and understanding visual grouping should help guide those comparisons, in theory – but almost no work validates that advice across the complexities of real-world visualizations. These are important – and solvable – problems that are being addressed by a new generation of researchers meeting at the intersection of data visualization and cognitive science fields.

Acknowledgements

Thanks to Evan Anderson, Cristina Ceja, Ryan Smith, Chase Stokes, Cindy Xiong, Jiaying Zhao, and the reviewers for their helpful comments on this manuscript. S.F. was funded by NSF IIS-1901485

References

Adjani, K., Lee, E., Xiong, C., Knaflic, C., Kemper, W., & Franconeri, S. (in press). Declutter and Focus: Empirically Evaluating Design Guidelines for Effective Data Communication. *IEEE Transactions on Visualization and Computer Graphics.*

Albers, D., Correll, M., & Gleicher, M. (2014, April). Task-driven evaluation of aggregation in time series visualization. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (pp. 551-560).

Alvarez, G. A. (2011). Representing multiple objects as an ensemble enhances visual cognition. *Trends in Cognitive Sciences*, *15*(3), 122-131.

Birch, S. A., & Bloom, P. (2007). The curse of knowledge in reasoning about false beliefs. *Psychological Science*, *18*(5), 382-386.

Börner, K., Bueckle, A., & Ginda, M. (2019). Data visualization literacy: Definitions, conceptual frameworks, exercises, and assessments. *Proceedings of the National Academy of Sciences*, *116*(6), 1857-1864.

Brooks, J. L. (2015). Traditional and new principles of perceptual grouping. In J. Wagemans (Ed.), *Oxford Handbook of Perceptual Organization*. Oxford, UK: Oxford University Press.

Cairo, A. (2016). The truthful art: Data, charts, and maps for communication. New Riders.

Cleveland, W. S., & McGill, R. (1985). Graphical perception and graphical methods for analyzing scientific data. *Science*, *229*(4716), 828-833.

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

Gleitman, L. R., Gleitman, H., Miller, C., and Ostrin, R. (1996). Similar, and similar concepts. *Cognition*, *58*, 321–376

Jacob, S. N., Vallentin, D., & Nieder, A. (2012). Relating magnitudes: the brain's code for proportions. *Trends in cognitive sciences, 16*(3), 157-166.

Kozhevnikov, M., Motes, M. A., & Hegarty, M. (2007). Spatial visualization in physics problem solving. *Cognitive Science*, *31*(4), 549-579.

Liu, L., Padilla, L., Creem-Regehr, S. H., & House, D. H. (2018). Visualizing uncertain tropical cyclone predictions using representative samples from ensembles of forecast tracks. *IEEE Transactions on Visualization and Computer Graphics*, *25*(1), 882-891.

Michal, A. L., 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., Franconeri, S. L. (2016). Visual routines for extracting magnitude relations. *Psychonomic Bulletin & Review*.

Miller, G. A., & Johnson-Laird, P. N. (1976). *Language and perception*. Belknap Press.

Munzner, T. (2014). *Visualization analysis and design*. AK Peters/CRC Press.

Nothelfer, C., & Franconeri, S. (2019). Measures of the benefit of direct encoding of data deltas for data pair relation perception. *IEEE Transactions on Visualization and Computer Graphics*, *26*(1), 311-320.

Padilla, L. M., Kay, M., & Hullman, J. (in press). Uncertainty Visualization. To appear in, *Handbook of Computational Statistics and Data Science*.

Roth, J., & Franconeri, S. (2012). Asymmetric coding of categorical spatial relations in both language and vision. *Frontiers in Psychology*, *3*, 464.

Shah, P., & Hoeffner, J. (2002). Review of graph comprehension research: Implications for instruction. *Educational Psychology Review*, *14*(1), 47-69.

Szafir, D. A., Haroz, S., Gleicher, M., & Franconeri, S. L. (2015). Four types of ensemble encoding in data visualizations. *Journal of Vision*.

Tversky, B. (2001). Spatial schemas in depictions. In Spatial schemas and abstract thought (pp. 79-111).

Wolfe, J. M., & Horowitz, T. S. (2017). Five factors that guide attention in visual search. *Nature Human Behaviour*, *1*(3), 1-8.

Van Essen, D. C., Anderson, C. H., & Felleman, D. J. (1992). Information processing in the primate visual system: an integrated systems perspective. *Science*, *255*(5043), 419-423.

Yang, F., Harrison, L. T., Rensink, R. A., Franconeri, S. L., & Chang, R. (2018). Correlation judgment and visualization features: A comparative study. *IEEE Transactions on Visualization and Computer Graphics*, *25*(3), 1474-1488.

Suggestions for Further Reading

For a review of the 'statistics' tool, see:

Szafir, D. A., Haroz, S., Gleicher, M., & Franconeri, S. L. (2015). Four types of ensemble encoding in data visualizations. *Journal of Vision*.

For more information on the 'comparison' tool, see:

Nothelfer, C., & Franconeri, S. (2019). Measures of the benefit of direct encoding of data deltas for data pair relation perception. *IEEE Transactions on Visualization and Computer Graphics*, *26*(1), 311-320.

For more information on how people correctly and incorrectly reason with graphs, see:

Shah, P., & Hoeffner, J. (2002). Review of graph comprehension research: Implications for instruction. *Educational Psychology Review*, *14*(1), 47-69.

For a review of visual uncertainty representation, see:

Padilla, L. M., Kay, M., & Hullman, J. (in press). Uncertainty Visualization. To appear in, *Handbook of Computational Statistics and Data Science*.

For a brief review of prescriptions for clear visual data communication, see:

Zacks, J. M., Franconeri, S. L. (2020). Designing Graphs for Decision-Makers. *Policy Insights from the Behavioral and Brain Sciences*.