

# The ladder of abstraction in statistical graphics\*

Andrew Gelman and Kaiser Fung

15 Dec 2025

## Abstract

Graphical forms such as scatterplots, line plots, and histograms are so familiar that it can be easy to forget how abstract they are. As a result, we often produce graphs that are difficult to follow. We propose a strategy for graphical communication by climbing a ladder of abstraction (a term from linguistics that we borrow from Hayakawa, 1939), starting with simple plots of special cases and then at each step embedding a graph into a more general framework. We demonstrate with two examples, first graphing a set of equations related to a modeled trajectory and then graphing data from an analysis of income and voting.

## 1. Motivation

A primary goal of statistics is generalizing from data. Even the most basic statistical construct, the statistical average, is an abstract concept. Modern statistical models contain layers of abstraction. We think such abstraction poses robust challenges to the effective communication of statistical findings. While graphs are intended to facilitate understanding of data, they are often difficult to follow because they too are abstractions.

We propose a strategy for graphical communication by climbing a ladder of abstraction, a term from linguistics that we borrow from Hayakawa (1939) who explains how words exist at different levels of abstraction and proposes that an awareness of such structure enables meaningful language. His framework has proven practical and popular, initially among education researchers adapting itbase as a teaching tool for narrative writing (e.g., Moffett, 1968). More broadly, social scientists use Hayakawa's semantic principle to bring structure to text treated as data (e.g., Hurst and Walker, 1972). Later applications appeared in more analytical fields, such as accounting (Guthrie, 1972), safety management (Metzgar, 1978), cartography (Jahnke et al., 2011), and program evaluation (Leviton, 2015).

In a section about defining words, Hayakawa argued that proper definitions point down the ladder of abstraction. This accords with our view that an effective graphical communication strategy should start from simple plots of special cases and then, in calculated steps, introduce layers of abstraction. We will demonstrate this strategy using two examples.

In this article we limit our discussion to scatterplots, line plots, and grids of plots with quantitative axes on the Cartesian plane. We imagine that similar principles should apply to other graphical forms, both static and dynamic, in which complexity can be built up, step by step, using color and other visual cues along with verbal descriptions that explain the connections among these steps.

## 2. Visualizing a mathematical model: basketball shooting

Consider the following simple problem in applied mathematics: What is the optimal angle for shooting a basketball, if you would like to be able to throw the ball as softly as possible and still reach the hoop? The solution will be a function of your distance from the hoop and the altitude at which you release the ball.

---

\*We thank Howard Wainer, Ron Yurko, Dianne Cook, Alessandra Casella, and two anonymous reviewers for helpful comments.



Figure 1: Diagram of a person standing 10 meters away from a basketball hoop, with the ball released at an altitude of 1.7 meters and the hoop being 3.05 meters high. This is the first step of a series of graphs that will culminate in displaying a optimal shooting angle.




Figure 2: (Left:) Trajectory of a shot taken at a  $30^\circ$  angle at a velocity of 15 m/sec, as computed using Newtonian mechanics ignoring air resistance. The shot goes too high and misses the hoop. (Right:) Trajectories of a series of shots with different initial velocities. The velocity required to reach the hoop must be somewhere between 10 and 15 m/sec. The left graph is at a low level of abstraction in that the red line directly shows the path of the ball. The graph on the right is more abstract in that it shows several trajectories at once.




Figure 3: A shot at angle  $30^\circ$  from altitude 1.7 m, from a distance of 10 m, will reach the hoop if its initial velocity is 12.2 m/sec.

We would like to reach a general solution but we start with the specific case in Figure 1: a player standing  $d = 10$  meters from the basketball hoop, releasing the ball at an altitude of  $a = 1.7$  meters. The hoop itself is  $h = 3.05$  meters off the ground.

The left plot in Figure 2 shows the path of a ball shot at a  $\theta = 30^\circ$  angle at an initial velocity of  $v = 15$  m/sec. From Newton's laws of motion and assuming no air resistance, the trajectory is  $x(t) = v \cos(\theta) t$ ,  $y(t) = a + v \sin(\theta) t - 0.5gt^2$ , where  $g = 9.8$  m/sec<sup>2</sup> is the gravitational constant. The red line on this graph directly shows the ball's trajectory. The only way to make it more vivid and less abstract would be an animated presentation that dynamically shows the moving ball.

We then ask what would happen at this same launch angle but with different initial speeds. The right plot in Figure 2 shows trajectories corresponding to  $v = 5, 10, 15,$  and  $20$  m/sec. This graph is slightly more abstract than the one on the left in that it expands the coverage from a single case to a range of possible paths. In addition, because all the curves start at the same point, we moved the label of the initial velocity at the end of each path, which is not quite intuitive. This is an example of the conceptual distancing that can be required when increasing the information content of a plot. Adding lines to a plot does not necessarily make it more abstract, but in this case the additional lines represent another dimension in the analysis and another possible direction of comparison. We find it helpful to present both graphs of Figure 2 in order to guide the reader through the process of abstraction.

Our next step is to figure out the initial velocity needed to reach the basket, given the shooting angle. This is a straightforward algebra problem. First we figure out the time  $T$  required for the ball to reach the horizontal position of the hoop:  $d = v \cos(\theta) T$ , thus  $T = \frac{d}{v \cos(\theta)}$ . Then we solve for the initial velocity  $v$  so that the ball's vertical position at time  $T$  is that of the hoop:  $h = a + v \sin(\theta) \frac{d}{v \cos(\theta)} - 0.5g \frac{d^2}{v^2 (\cos(\theta))^2} = a + d \tan(\theta) - 0.5 \frac{gd^2}{v^2 (\cos(\theta))^2}$ , so

$$v = \sqrt{\frac{0.5gd^2}{(\cos(\theta))^2(d \tan(\theta) + a - h)}}. \quad (1)$$

Plugging  $a = 1.70, \theta = 30^\circ, d = 10, h = 3.05, g = 9.8$  into this equation yields  $v = 12.2$ . The resulting trajectory is the blue line in Figure 3. It is helpful to see it in the context of the red curves that show other velocities.

Now that we have an equation for the required initial velocity for the ball to reach the hoop, we can plot it as a function of the shooting angle,  $\theta$ . The left plot of Figure 4 shows the result.

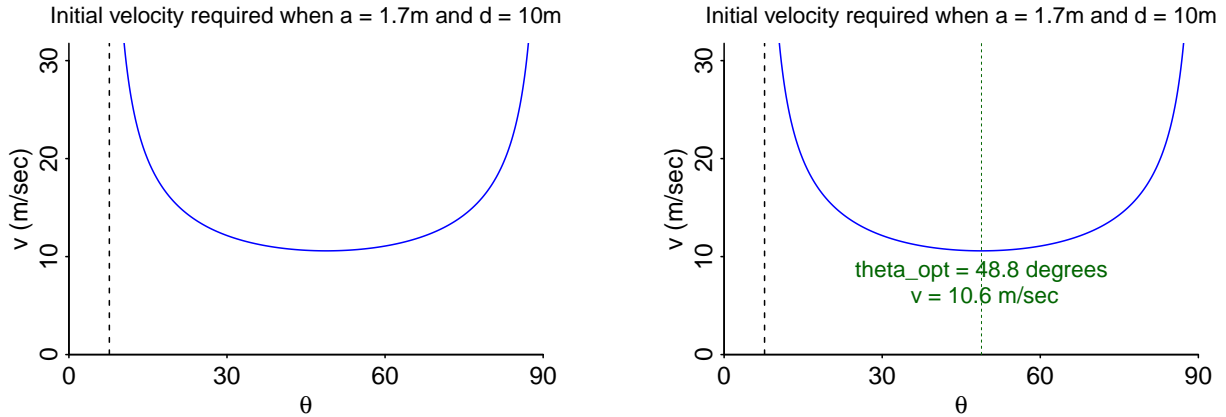


Figure 4: (Left:) Velocity required to reach the basketball hoop for a shot released at an altitude of 1.7 meters, 10 meters away from the basket. (Right:) On this graph, the required velocity  $v$  is minimized at an angle  $\theta = 48.8^\circ$ , at which point  $v = 10.6$  m/sec. The blue color of the lines matches the blue used to display the hoop-reaching solution in Figure 3.




Figure 5: (Left:) Shooting angle requiring the softest shot and the corresponding release velocity, as a function of the distance from the hoop, for a shooter releasing the ball at an altitude of 1.7 meters. (Right:) Same graphs, also showing the solutions for release altitudes of 1.2 and 2.2 meters. The green color of the lines matches the green used to display the minimum-initial-velocity solution in Figure 4.

For the lowest angles ( $\theta \leq \tan^{-1}(\frac{h-a}{d})$ , indicated by the dashed vertical line on the plot), the ball will go below the hoop no matter how hard it is thrown. The lowest initial velocities are needed for shots taken from between  $30^\circ$  and  $60^\circ$ , with harder shots required when the angle goes outside that range.

We can then use a general-purpose numerical optimizer to find the value of  $\theta$  that minimizes (1), given  $a = 1.70, d = 10, h = 3.05, g = 9.8$ . The result is  $\theta = 48.8^\circ$ , yielding  $v = 10.6$ . This result is displayed as a green dotted line in the right plot of Figure 4. Again, we find it helpful to display both graphs, as they show the development of the analysis: first the curve, then the minimum, which we label as the “optimum” in the sense of requiring the softest throw to reach the hoop.

Compared to Figure 3, Figure 4 represents a leap in abstraction: the curve is not a trajectory, and it embodies the solution to an equation. By seeing both charts, readers can observe the switch from physical space to an abstract space: the blue line in Figure 3 is the single dot at  $\theta = 30^\circ$  in Figure 4, in which the angle  $\theta$  and the initial velocity  $v$  are now represented by linear coordinates. Again, in statistics we are so familiar with this sort of graph that we don’t always remember how abstract it is.

The punch line of Figure 4 is the recommended shooting angle,  $\theta_{\text{opt}}$ , and corresponding initial velocity  $v$ , for a shot taken at an altitude of 1.7 meters, 10 meters from the basket. Naturally, we should ask how these parameters change as a function of distance from the basket. Obviously  $v$  will increase with distance, but at what rate? And what will happen to  $\theta_{\text{opt}}$ ?

We answer this question by unfixing  $d = 10$  and instead considering a grid of values of  $d$ , ranging from 1 to 15 meters from the hoop. For each value of  $d$ , we numerically optimize and find the value of  $\theta$  that minimizes (1), along with the corresponding value of  $v$ .

The left panel of Figure 5 shows the results. When you are standing close to the basket you need to shoot at a steep angle to reach the hoop; as the distance increases, the optimal angle approaches  $45^\circ$ . The required release velocity increases gradually, at a rate slightly less than linearly with distance. We have lined up the two plots vertically so they can share an  $x$ -axis and the reader can see how both  $\theta_{\text{opt}}$  and  $v$  vary together.

Finally, we gain one more level of generality by considering different release points: 1.2 m, 1.7 m, and 2.2 m off the ground, which might correspond to shooting from chest level, head level, or above the head. The right panel of Figure 5 shows curves for each of these three heights. Unsurprisingly, releasing from a greater altitude requires a lower shooting angle and a lower initial velocity.

Readers experience a similar step up in abstraction going from Figure 4 to Figure 5 as they did from Figure 3 to Figure 4. The curve in Figure 4 is now one pair of dots in the paired charts of Figure 5. The coupling of the angle and speed parameters also carries over from Figure 4.

The right panel of Figure 5 represents the greatest level of abstraction in this series of graphs, and, again, we find it helpful to build up to it rather than simply presenting these results and then expecting readers to follow.

We have also used graphical cues as much as possible to make connections between the graphs in the series. When two graphs share the same graphical space, as defined by the variables in the axes, the axis ranges and aspect ratios are kept constant. When the space definition shifts, we include sufficient information to convey the relationships between the two spaces.

In addition, throughout this example we have attempted to guide the reader using a succession of colors:

- *Black* for the basketball court in Figures 1–3 and the axes in Figures 4–5,
- *Red* lines for the trajectories in Figures 1–3.

- *Blue* line for the trajectory in Figure 3 that reaches the hoop, and also for the curves in Figure 4 that represent the initial velocity as function of angle for these shots,
- *Green* for the optimal angle that minimizes the required initial velocity in the right plot of Figure 4, and also for the curves in Figure 5 that show this optimal angle and corresponding initial velocity as a function of distance and release altitude.

The colors climb the ladder of abstraction, from direct mapping of trajectories to the solutions of a series of equations. When presenting the series of graphs we do not draw attention to the colors—that might feel like overkill and distract readers from the content being graphed—but we hope the colors help, in the same way that good design can make a tool more effective even for users who are unaware of these choices.

### 3. Visualizing data: income and voting within and between states

In a book on quantitative trends in politics by one of the co-authors (Gelman et al., 2009), we were particularly attuned to the challenges of communicating to a general audience. It was important to us to use graphs and not just text because we wanted to actively engage readers in the process of understanding and discovery. Rather than simply *saying*, for example, that richer people were more likely than poorer people to vote for Republicans, but Democrats did better in rich states, we wanted to *show* it.

Figure 6 juxtaposes the two patterns for George W. Bush in 2004, showing he did better among richer voters but worse within poorer states. Despite using standard chart forms, both these graphs are abstract, but in different ways: in the left plot, each dot represents the abstract construct of an income subgroup, a set of survey respondents with a common response to the income question; in the right plot, the horizontal and vertical axes establish an abstract space, in which each state is represented by a point, given by its average income and vote share.

The contrasting patterns for individuals and states motivated us to look at income and voting within states. Again, we climb the ladder of abstraction in Figure 7, by choosing three states (Mississippi, Ohio and Connecticut) to represent poor, middle-income, and rich states respectively. Without such generalization, Figure 7 would contain 50 lines and would be close to unreadable.

While the axis labels suggest income subgroups, you might notice the smooth lines rather than the connected points in the left panel of Figure 6. In Figure 7, the lines display summaries of a fitted hierarchical logistic regression model—each of the lines is a part of a logistic curve, with each being on a narrow enough range that the lines appear nearly straight. It might seem like cheating in Figure 7 to display fitted curves rather than raw data, but you can think of these as smoothed data or as an alternative to presenting logistic regression coefficients.

While Figure 7 draws out the income and voting patterns within state subgroups, in Figure 8 we add another feature to illustrate the between-state pattern, finally achieving our goal of juxtaposing divergent patterns within and between states. On top of each state’s line, open circles locate the five income subgroups and their respective probabilities of voting Republican, with the sizes of the circles proportional to the number of voters in each category, relative to the income distribution of the electorate as a whole, while the solid circle plots the Republican candidate’s vote share in the state against the average income in the state on this five-point scale. The startling insight from Figure 8 is the opposite directions of correlation depending on whether one looks between or within states. The pattern of the three solid circles in each graph shows how poorer states are more Republican-leaning, even while the open circles for each line indicates a positive correlation between income and voting Republican within each state. To draw another connection between




Figure 6: (Left:) From exit polls, national voting patterns showing that George W. Bush, the Republican candidate for president in 2004, did better among richer voters. (Right:) At the same time, Bush was more successful in poorer states.




Figure 7: From exit polls, George W. Bush's vote share in 2004, as a function of income, in a poor state, a middle-income state, and a rich state.

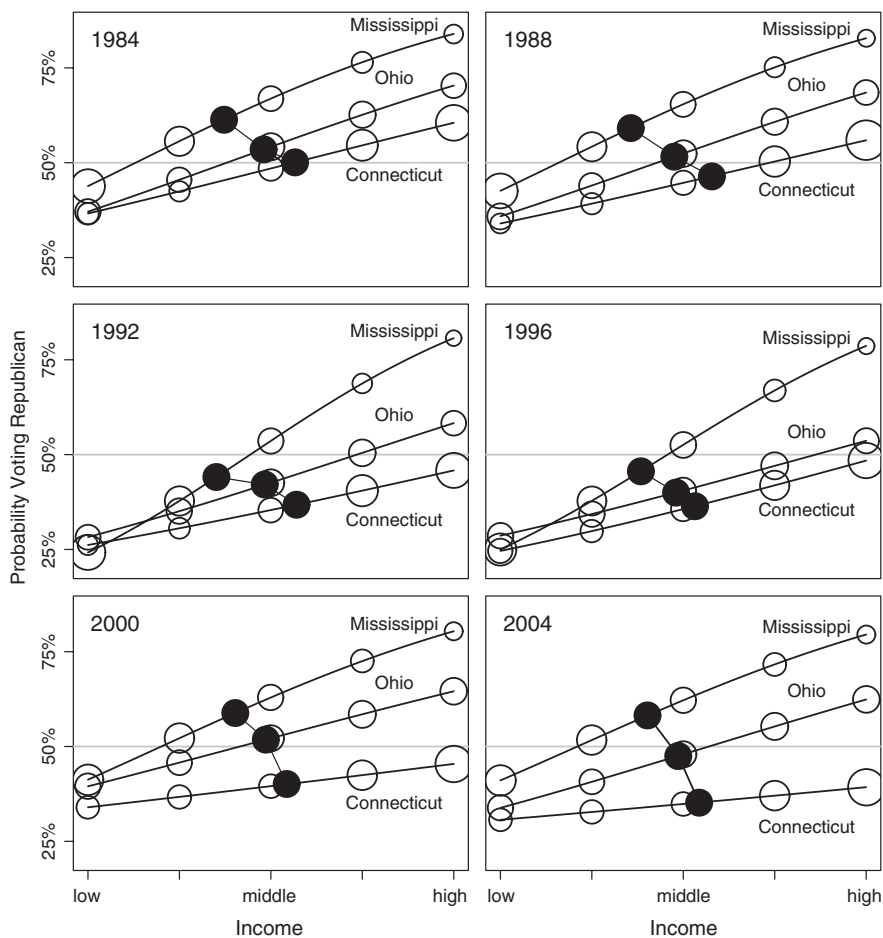


Figure 8: *Probability of Republican vote as a function of income in poor, middle-income, and rich states in six straight presidential elections. The open circles on the plot show the relative size of each income group in each state (thus, Mississippi has more poor people than average, and Connecticut has more rich people). The solid circles show the average Republican vote and average income in each state.*

graphs, the lower-right plot in Figure 8 corresponds to 2004, the same year as displayed in Figure 7, but the lines are slightly different because the estimates come from a different data source.

We considered options to include uncertainty estimates in these graphs, such as error bars or shaded uncertainty zones, but decided not to, as we were worried about cluttering that would make it more difficult to make visual comparisons. In addition, the variation of the estimates from year to year gives some sense of estimation uncertainty. In any case, there is no right answer here; there is an inevitable tradeoff between clarity and the presentation of more information.

As it stands, Figure 8 is a leap of abstraction compared to Figure 7. Each state is further represented by its average income, rather than continuous or categorical measures of income. By expanding from one to six presidential election cycles, we have to generalize from voting for Bush to voting for the Republican candidate.

Indeed, one could take another step of abstraction and summarize each fitted curve by an intercept and slope on the logistic scale (with the income predictor coded as  $-2, -1, 0, 1, 2$ ), so the intercept would be directly interpretable and then make a scatterplot of slope vs. intercept with two-letter state abbreviations. Such a graph could be difficult to explain on its own but could be




Figure 9: *From a large national survey in 2000, average positions of respondents on social and economic issues among voters in different income levels, within Republican-leaning, battleground, and Democratic-leaning states. Richer voters are consistently more conservative on economic issues, but the patterns on social issues are more complicated.*

a natural followup to Figure 7, showing fifty states instead of just three.

To better understand what was happening with income and voting, we went to polling data and created left-right scores on economic and social issues by combining several survey items in each area. The estimates were based on a large national poll from 2000 with enough data to get stable estimates for subgroups characterized by individual incomes and state political leanings. Figure 9 shows the result, which uses lines, colors, and labels to enable comparisons across voters and states in both dimensions (social and economic).

Figure 9 scales the ladder of abstraction in several ways. The economic and social scores represent aggregated responses to multiple survey questions while the income subgroups are determined by the income question alone. These scores create an abstract space of ideology, in which individuals can be placed. Instead of individuals, we plot income-state subgroups, which is another abstraction. States are categorized as Republican-leaning, battleground, or Democratic-leaning.

Once Figure 9 has been understood, we can go one more level of abstraction to Figure 10, which displays the same comparisons but separately for people who attend church regularly, occasionally, or not at all. We did not want to clutter the display with standard errors for all these averages, but uncertainty can be deduced from apparently random variation in the plots.

Figure 10 represents the greatest level of abstraction in this series of graphs. Presented on its own, it would take some effort to understand well, which is why we find it helpful to lay out the path in several steps, starting from raw survey data on income and voting and gradually adding more information. In designing the graphs, we have used certain strategies. Specific cases are frequently presented first, then generalized individuals, and finally subgroups that are abstractly defined by models. In statistics, we are so used to talking about averages, without recognizing that




Figure 10: Continuation of Figure 9, further characterizing people by religious attendance. This plot is complicated, but it is possible to read it by building up from earlier, less-abstract graphs.

they represent an abstract construct.

The effectiveness of graphs such as Figure 10 will depend on the problem being studied and on the data being plotted. Here it helped that the lines for the three groups of states were clearly separated and that the paths from poor to middle-income to rich voters had generally similar trajectories. The cleaner the data, the less burden is placed on statistical modeling. As with any applied problem, this is not the end of the story. There are many other ways of looking at data on income and voting, and the patterns we showed above have changed in the years since those graphs were published (see, for example, Feller et al., 2012, and Gelman and Azari, 2017).

## 4. Discussion

There is a long history of abstraction in visual representation, going far beyond statistical graphics. Consider, for example, the phrase “the map is not the territory” (Korzybski, 1933) and the satirical vignette of Borges (1946) pointing out the uselessness of a perfectly faithful 1:1 scale map. Figure 11 shows an amusing example discussed by Wainer and Friendly (2021) of an early infographic that cleverly conveys information directly through a non-abstract image. Literature in cartography (Fabrikant, 2004) and information visualization (Cui et al., 2006, Viola and Isenberg, 2018) have considered the cognitive and perceptual aspects of abstraction in visual communication. But when it comes to displaying mathematical or statistical relationships, we typically take for granted the abstractions of scatterplots, time series graphs, bar charts, and so on.

Our contribution in this paper is to demonstrate the general strategy of presenting an increasingly more abstract world using a series of graphs starting with simple plots of special cases. This approach should be increasingly useful to statisticians and graphic designers in a world of big data and information visualization. Statistical graphics can be beautiful, informative, and sophisticated, but they do not always speak for themselves. It is well known that a graph can be much improved with a clear title and caption: as the saying goes, a picture plus a thousand words is worth more than two pictures or two thousand words. Here we have shown how it can also be helpful to build up to an informative graphic by explicitly showing intermediate graphs at lower levels of abstraction.

Throughout the two examples, we apply abstraction in various ways when moving up the ladder.




Figure 11: Graph from Panckoucke (1783) illustrating the portions of trees that could be used to construct different parts of a ship. This is an example of a non-statistical visual representation at a low level of abstraction.

Strategy	Basketball example	Voting example
Define an abstract space or transform it	Figures 4 and 5	Figure 8
Use modeled or optimized values	Figure 4	Figures 7 and 8
Use statistical averages		Figure 8
Expand coverage from single year to multiple years		Figure 8
Expand sampling from single to multiple	Figures 2 and 5	
Define subgroups or select representative instances		Figures 7 and 8
Allow a fixed parameter to vary	Figures 4 and 5	

Table 1: Some strategies of abstraction and where they are used in the two examples of this paper. Each strategy represents a way that a statistical analysis or summary can be generalized.

In each step, the prior graphs prepare the readers to understand the additional abstraction in the next graph, similar to how Hayakawa defines a more abstract word in terms of less abstract words. These strategies are summarized in Table 1. When preparing a series of graphical displays, it could also be helpful to include example questions that might be asked at each level of abstraction, such as “What does this graph show directly?”, “What comparisons are being facilitated by this display,” “How does this generalize across different subsets of the data?”, “What assumptions are implicit in the model?”, and, often most importantly, “How can this plot be misinterpreted and what can be done to avoid such misreadings?”

Modern treatments of data visualization or the grammar of graphics consider graphs not as static drawn objects but as overlays of information, so that a useful plot is built by adding elements to a general structure (Wilkinson, 2005, Wickham, 2016). The ladder of abstraction discussed in the present paper represents a different sort of forward progression: in addition to adding information to a graph, we are using increasing levels of generality. Our ideas follow the spirit of the grammar of graphics, however, in considering statistical graphics as a form of communication of quantitative information to the reader, with the specific graph being just an instantiation of this goal. As emphasized by Tufte (1983) and Cleveland (1985), choices in data display can be understood with respect to the specific goals of communication. In present paper, we emphasize the challenge that useful visualizations can be highly abstract representations of the information being presented, hence our proposal of stepping readers through a series of graphs. Here we have demonstrated this idea in two different applications of statistical graphics: continuous curves and discrete data.

Once presented, the ladder of abstraction seems natural, but we have rarely seen it in published work, including our own. Our usual approach in graphical communication has been to develop the graphs we want to display and then explain them in words. The examples here suggest that a more explicit unfolding could be useful, by analogy to the paradigm in statistical workflow of fitting increasingly complicated approximations, with each understood in comparison to what came before, an idea we have found helpful for modeling as well as computation (Gelman et al., 2020).

In general, the need for abstraction in statistical graphics arises from the difficulty of representing complicated data structures and functional forms in two linear dimensions. The particular sorts of abstractions needed will depend on the form of the observed data and of latent dimensions of interest, and the larger goals of the problem (finding an optimal shooting angle in our basketball example and exploring voting patterns as a function of income, religious attendance, and state in our political science example). As statisticians we can start from these goals and, from there, figure out useful graphical displays, and then when presenting such displays we can recognize their abstract nature and present them by building up from simpler forms.

## 5. Recommendations

We have not offered any one-size-fits-all strategy to construct such a series of plots, but we think the two examples in this paper demonstrate how it can be done in two very different situations. We have deliberately chosen examples that differ in their general structure (mathematical functions vs. statistical data), subject area (physics vs. political science), and goals (optimization vs. exploration). Even with all these differences, both problems benefit from starting at the core with directly-interpretable graphs and then gradually adding layers of complexity, as indicated in Table 1. We leave it to future researchers to further formalize these ideas and implement them into general-use software.

The present article has three general messages. First, some level of abstraction exists in all but the simplest graphs, and we should be aware of this when using graphs to communicate, especially to readers who are not familiar with data graphics. Second, graphical abstraction is not

just a difficulty; it also is a useful tool that allows us to understand patterns at a greater level of sophistication, and it can be valuable to see this by adding one level of abstraction at a time rather than trying to jump to a single plot that shows everything. Finally, these principles apply not just to graphs of data but also to graphs of mathematical functions such as fitted models.

With this in mind, we offer the following recommendations to practitioners. The starting point is awareness that even very familiar statistical displays such as scatterplots, histograms, and time series plots are highly abstract. The next step is to offer support for the reader by explaining the visual mapping in words and pictures, for example by identifying some individual points on a graph, labeling lines directly (as in Figure 3), and including descriptive captions. There is no reason to expect a graph to speak for itself without explanation. Also remember that it can be a mistake to try to cram too much information on a single plot. It can be helpful to follow the recommendation of Tufte (1990) and display small multiples—grids of graphs that reinforce their structures through repetition, as in several of the figures in this paper. The final step is to construct a series of displays, adding one level of abstraction at a time and connecting each graph to the ones that came before.

It can take some thought and work to carry the reader along with us the ladder of abstraction, as demonstrated in the two examples above. But we think this effort has value, not just for communicating to others but also for exploring data and models for ourselves. Working through a series of increasingly abstract graphs, and recognizing the abstraction being added at each step, can help us make more sophisticated plots and make more interesting comparisons. We were only able to reach displays such as Figures 5 and 10 by carefully extending them from simpler plots. Any given real-world data or modeling project will probably not look like either of the examples in this article, but we think it will still benefit from a series of visualizations, adding and explaining one aspect of variation at a time.

## References

- Borges, J. L. (1946). Del rigor en la ciencia. *Los Anales de Buenos Aires* **1** (3).
- Cleveland, W. S. (1985). *The Elements of Graphing Data*. Monterey, Calif.: Wadsworth.
- Cui Q., Ward M., Rundensteiner E., and Yang J. (2006). Measuring data abstraction quality in multiresolution visualizations. *IEEE Transactions on Visualization and Computer Graphics* **12** (5), 709–716.
- Fabrikant, S. I. (2004). Abstraction and scale in spatialisation. In *Proceedings of the National Cartographic Conference*, 35–43. New Zealand Cartographic Society.
- Feller, A., Gelman, A., and Shor, B. (2012). Red state / blue state divisions in the 2012 presidential election. *The Forum* **10**, 127–141.
- Gelman, A., and Azari, J. (2017). 19 things we learned from the 2016 election (with discussion). *Statistics and Public Policy* **4** (1), 1–10.
- Gelman, A., Park, D., Shor, B., and Cortina, J. (2009). *Red State, Blue State, Rich State, Poor State: Why Americans Vote the Way They Do*, second edition. Princeton University Press.
- Gelman, A., Vehtari, A., Simpson, D., Margossian, C. C., Carpenter, B., Yao, Y., Bürkner, P. C., Kennedy, L., Gabry, J., and Modrák, M. (2020). Bayesian workflow. <https://arxiv.org/abs/2011.01808>
- Guthrie, A. (1972). Modern semantics can help accounting. *Journal of Accountancy* **133**, 56–63.
- Hayakawa, S. I. (1939). *Language in Action*. New York: Harcourt, Brace.
- Hurst, J. W., and Walker, H. K., eds. (1972). *The Problem-Oriented System*. New York: Medcom.
- Inselberg, A. (1985). The plane with parallel coordinates. *The Visual Computer* **1**, 69–91.

- Jahnke, M., Krisp, J. M., and Kumke, H. (2011). How many 3D city models are there?—A typological try. *Cartographic Journal* **48**, 124–130.
- Korzybski, A. (1933). *Science and Sanity*, 747–761. New York: International Non-Aristotelian Library Publishing Company.
- Leviton, L. C. (2015). Evaluation practice and theory: Up and down the ladder of abstraction. *American Journal of Evaluation* **36**, 238–242.
- Metzgar, C. (1978). Hayakawa’s cow & safety. *Professional Safety* **23** (8), 44–48.
- Moffett, J. (1968). Teaching the universe of discourse. Boston: Houghton Mifflin. Available at <https://wac.colostate.edu/books/landmarks/moffett/universe/>.
- Panckoucke, C. J. (1783). *Encyclopédie Méthodique Marine*. Paris: Chez Panckoucke.
- Tufte, E. R. (1983). *The Visual Display of Quantitative Information*. Cheshire, Conn.: Graphics Press.
- Tufte, E. R. (1990). *Envisioning Information*. Cheshire, Conn.: Graphics Press.
- Viola, I., and Isenberg, T. (2018). Pondering the concept of abstraction in (illustrative) visualization. *IEEE Transactions on Visualization and Computer Graphics* **24** (9), 2573–2588.
- Wainer, H., and Friendly, M. (2021). Displaying information: From woolly mammoths to the Great Migration. In *International Encyclopedia of Education*, 4th edition, ed. R. Tierney, F. Rizvi, K. Ercikan, and G. Smith. Philadelphia: Elsevier.
- Wickham, H. (2016). *ggplot2: Elegant Graphics for Data Analysis*. New York: Springer-Verlag. [ggplot2.tidyverse.org](https://ggplot2.tidyverse.org)
- Wilkinson, L. (2005). *The Grammar of Graphics*, second edition. New York: Springer-Verlag.