

# Information Visualization Principles, Techniques, and Software

## Perceptual Properties

Of all the human senses, vision is considered a primary and powerful channel of input. It takes in everything from the world around us and transmits this information to the brain for an instant analysis of what it means. With this intimate connection that vision has with cognition comes the first step toward learning and creating visual designs, which is through understanding how human visual perception works. Visual perception is relative, constantly scanning, constantly adjusting focus, and constantly adapting. A well-known example, which is presented in figure 2.1, demonstrates how written words are perceived in a way that is different from what most people may initially think. According to research in word recognition and how the human brain reads letters, the order of letters in a word does not matter; the only important thing is that the first letter and last letter are in the right place.<sup>1</sup> Our vision relays the information to our brain, which takes care of processing what has been seen. Even when an entire paragraph contains numerous spelling errors, people can still comprehend it without any problems and often without even realizing that there are any errors in the text. Figure 2.1 serves as a nice demonstration of how human vision and cognition interact: the human mind does not read every letter by itself; instead, it processes the word as a whole. This example also illustrates the complex and mysterious nature of human vision. Even more, it creates extensive opportunities for visual designers to take advantage of human visual perception as it pertains to information visualization. Among all the unique characteristics of vision, two perceptual properties are commonly applied in information visualization: pre-attentive processing and the Gestalt Laws.

I cdnuolt blveiee  
That I cluod aulacty uesdnatnrd what I was rdanieg.  
The phaonmneal pweor of the hmuan mnid  
Aoccdrnig to rscheearch at Cmabridge Uinervtisy,  
It deosn't mtttaer in what oredr the ltteers in a wrod are,  
The olny iprmoatnt thing is  
That the first and last ltteer be in the rghit pclae.  
The rset can be a taotl mses and you can still raed it wouthit a porbelm.  
This is bcuseae the huamn mnid deos not raed ervey lteter by  
istlef, but the wrod as a wlohe.  
Amzanig huh?

What the above text says:


I couldn't believe  
That I could actually understand what I was reading.  
The phenomenal power of the human mind.  
According to research at Cambridge University,  
It doesn't matter in what order the letters in a word are,  
the only important thing is  
That the first and last letter be in the right place.  
The rest can be a total mess and you can still read it without a problem.  
This is because the human mind does not read every letter by  
itself, but the word as a whole.  
Amazing, huh?

### Figure 2.1

Word recognition: how our brain reads letters (from Jinxi Caddel, "Word Recognition: How Our Brain Reads Letters," *Learn, Think, Inspire* [blog], February 4, 2010, [www.jinxi-boo.com/blog/2010/2/4/word-recognition-how-our-brain-reads-letters.html](http://www.jinxi-boo.com/blog/2010/2/4/word-recognition-how-our-brain-reads-letters.html)).

## Pre-attentive Processing

Pre-attentive visual processing, which takes place in the sensory memory, is fundamental for creating visual representations. Pre-attentive visual attributes are perceived by viewers almost instantaneously and without the intervention of consciousness. One of the key issues in the study of information visualization is to investigate how the human visual system processes and analyzes an image. An important initial result of this research was that a small number of visual



**Figure 2.2**

Search for a target red circle based on a difference in hue (from Christopher G. Healey, “Perception in Visualization,” North Carolina State University, Department of Computer Science, last modified September 24, 2016, <https://www.csc.ncsu.edu/faculty/healey/PP/index.html>).

properties can be detected rapidly and accurately simply by the “low-level” visual system. These properties were termed pre-attentive since their detection appears to precede focused attention.<sup>2</sup> It may be difficult to fully perceive the value of this insight, but understanding pre-attentive processing is important for visualization design. In particular, by exploiting pre-attentive processing capability, several important design questions can be answered, such as, What can be perceived immediately? Which visual properties are good discriminators? What can mislead viewers? How can one design the information so that it pops out?

Christopher G. Healey from North Carolina State University has created a website to demonstrate what pre-attentive processing means in regard to the experience of information visualization.<sup>3</sup> One of his examples is given in figure 2.2. This figure shows how to search for a target circle based on a difference in colors—the darker circle pops out immediately among all the lighter circles. This is a basic, yet common use of pre-attentive processing. Some other basic tools used in pre-attentive design include length, width, and intensity. There are also more advanced symbols, which include lighting and motion direction.

### Gestalt Laws

Gestalt Laws, which are also known as the Law of Simplicity and the Law of Pragnanz (which entails the entire visual), describe how to arrange visual symbols in a graphical display that is optimized to achieve a better, more effective visualization. These laws focus on how people interpret the world, providing relevant principles in relation to perceptual organization. Gestalt Laws were first proposed by German psychologists in the early 1900s.<sup>4</sup> Their primary purpose was to assist in understanding pattern perception, while also providing clear descriptions of many basic perceptual phenomena. While psychologists call them laws, these principles are more like heuristics, which are mental shortcuts for solving problems.<sup>5</sup>



**Figure 2.3**

Gestalt Laws: the faces-vase drawing (from Edgar Rubin, “Figure and Ground,” in *Visual Perception: Essential Readings*, ed. Steven Yantis [Philadelphia, PA: Psychology Press, 2000], 225–29).

One of the well-known Gestalt Laws is known as “figure and ground,” also called the “faces-vase” drawing, shown in figure 2.3.<sup>6</sup> The foreground of this image is the “figure,” while “ground” is what is behind. This drawing exemplifies one of the key aspects of the figure-ground organization and edge assignment, along with its effect on shape perception. One can notice that in the faces-vase drawing, the perceived shape depends critically on the direction in which the border between the black and white regions is assigned. The perception of figure, as opposed to ground, can be thought of as part of the fundamental perceptual act of identifying objects.

It must be acknowledged that Gestalt Laws play a significant role in information visualization. These laws apply in the creation of graphs in which visual elements are designed so as not to interfere with each other. Specifically, according to Gestalt Laws, the design of visualizations should avoid unexpressive marks, use perceptually effective encodings, and not distract the audience. Creating an effective and expressive visual design with a truthful message is a high priority in information visualization.

### Information Visualization and Cognition

In one of the definitions of *information visualization*, it is stated that “the use of computer-supported,

interactive visual representations of data to amplify cognition” is necessary.<sup>7</sup> This close connection between information visualization and human cognition is emphasized. One of the noted benefits of graphics is to help simplify the search for information that is needed for task completion through the aid of human vision. In particular, in order to crystalize knowledge, data is needed. When data alone is not able to easily provide information, the use of information visualization can help. Visualization aids cognition due to the ways that it increases memory and the processing of resources available through pre-attentive properties, Gestalt Laws, and many other perceptual properties. By simplifying the search for information and enhancing the recognition of patterns, visualization enables perceptual inference operations.

## Tufte’s Principles for Information Visualization

Edward Tufte is the man the *New York Times* called “the da Vinci of data” because of his concisely written and artfully illustrated books on the visual display of data.<sup>8</sup> In addition to being noted for his writings on information design, Tufte is also a pioneer in the field. Through his work, he proposed key principles in designing visualizations, which have offered substantive and important insights into displaying information for maximum effect and ease of comprehension.

One consequence of Tufte’s principles is the ability to achieve graphical excellence through information visualization. According to Tufte, graphical excellence is “the well-designed presentation of interesting data—a matter of substance, of statistics, and of design.”<sup>9</sup> Knowing that graphical excellence consists of complex ideas communicated with clarity (no ambiguity or confusion in graphs), precision (truthful results and distortion-free presentations), and efficiency (a minimal amount of chart “junk”), this is a critical point of concern. Graphical excellence is what will give the viewer the greatest number of ideas in the shortest amount of time and with the least ink possible in the smallest amount of space. Such graphical excellence can be achieved by accomplishing two goals: (1) telling the truth about the data, which is achieved by graphical integrity principles, and (2) visualizing with clarity and precision, which is achieved by design principles.

Graphical integrity principles are concerned with having clear, detailed, and thorough labeling, which can be used to prevent graphical distortion and ambiguity. Tufte also defined the “lie factor” as “a value to describe the relation between the size of effect shown in a graphic and the size of effect shown in the data.”<sup>10</sup> Specifically, the representation of numbers, as physically measured on the surface of the graphic

itself, should be directly proportional to the quantities represented. Another design principle that Tufte introduced was the “data-ink ratio.”<sup>11</sup> Tufte referred to data-ink as “the non-erasable core of a graphic, the non-redundant ink arranged in response to variation in the numbers represented.”<sup>12</sup> By definition, non-data-ink is the ink that does not convey information but is used for scales, labels, and edges. To further expand on this, the data-ink ratio is “the proportion of ink that is used to present actual data compared to the total amount of ink (or pixels) used in the entire display.”<sup>13</sup> The lesson in all of these details and objectives is that good graphics should include only data-ink; all non-data-ink should be deleted whenever possible to avoid drawing attention to distracting, irrelevant elements in the presentation. The goal is to design a display with the maximum possible data-ink ratio without eliminating something that is necessary for effective communication.

The non-data-ink concept also leads to the discussions of avoiding “chart junk.” Chart junk is all the visual elements in charts and graphs that are not necessary for the viewer to comprehend the information represented on the graph or that distract the viewer from this information.<sup>14</sup> Examples of unnecessary elements that are chart junk include heavy or dark grid lines, unnecessary text, or inappropriately complex font faces. In addition, ornamented chart axes and display frames, pictures or icons within data graphs, and ornamental shading and unnecessary dimensions are in danger of being considered chart junk.

Interestingly, there has been research on people’s acceptance of this “minimalist” approach to visualizing information.<sup>15</sup> Research that disagrees with what Tufte has written suggests that there are circumstances in which chart junk is considered useful and effective. These information visualization scholars have argued that elaboration is not all bad and that visual embellishments may have other benefits. In particular, when graphs are used with the purpose of persuasion or presentation, they should be designed with the specific objective of aiding the memorability of the presented data. Through evaluation of existing research, it has been supported that a data graphic must engage the reader’s interest, something that can be achieved through the use of graphic imagery.<sup>16</sup>

There is another widely adopted design principle proposed by Tufte that addresses “data density” as being beneficial for a good graphic. Data density is defined as “a ratio of the number of entries in data array to the area of data graphic.”<sup>17</sup> Intuitively, there are three ways to increase data density: including small multiples, including sparklines, and simply maximizing the amount of data shown. The approach that is referred to as “small multiple” figures addresses a collection of miniature illustrations arrayed as a single figure, which is designed to be perceived as one.

When talking about the sparklines approach to create a graphic of small multiples, Tufte described them as being “data-intense, design-simple, and word-sized graphics.”<sup>18</sup> The third approach, that of maximizing the amount of data shown, simply refers to showing as much data as possible.


## Information Visualization Techniques

There is a wide range of techniques to be utilized when it comes to information visualization. These techniques range from simple charts, such as bar and line graphs, to more advanced techniques, such as heatmaps and scatter plot matrixes. Some of the techniques are created for a specific purpose, thereby making it hard to categorize them. In this section, the primary focus will be on introducing several commonly used information techniques.

Let’s start by defining *multivariate data*, as it is a common term and concept used in many of these techniques. *Multivariate data* refers to data of more than three dimensions. In other words, there are more than three variables per case in the data set.

The first approach to visualizing multivariate data is through multiple views.<sup>19</sup> With this technique, if one cannot present the dimensional data all in one graph, one can try to display them in multiple graphs, with each of the graphs conveying a certain message. Figure 2.4 provides an example of the multiple views technique.<sup>20</sup> In this example, a matrix of data is presented on the left of figure 2.4, with five variables in each row and five variables in each column. The intersection presents the value corresponding to the specific row and column. This matrix is a good choice for numerical representation, but it is also difficult to retrieve any pattern or relationship of the data from the matrix at first glance. The logical progression of thought would be to determine a more intuitive way to give each variable its own display, as shown on the right of figure 2.4. In this visual display, each variable, 1 to 5, is given a graph, with variables A to E on the x-axis and their corresponding values on the y-axis. The visual display demonstrates how a multivariate data representation can be partitioned into five sub-graphs, with each sub-graph mapping the data on a two-dimensional plane. While this visualization technique is easy to implement and intuitive, an immediate drawback exists: as the number of variables increases, the number of sub-graphs needed for each variable also increases. This increase creates a risk of having too many sub-graphs presented at one time, which increases the noise level, making information comprehension and processing challenging.

A more sophisticated approach to visualizing multivariate data is through the use of a scatter plot

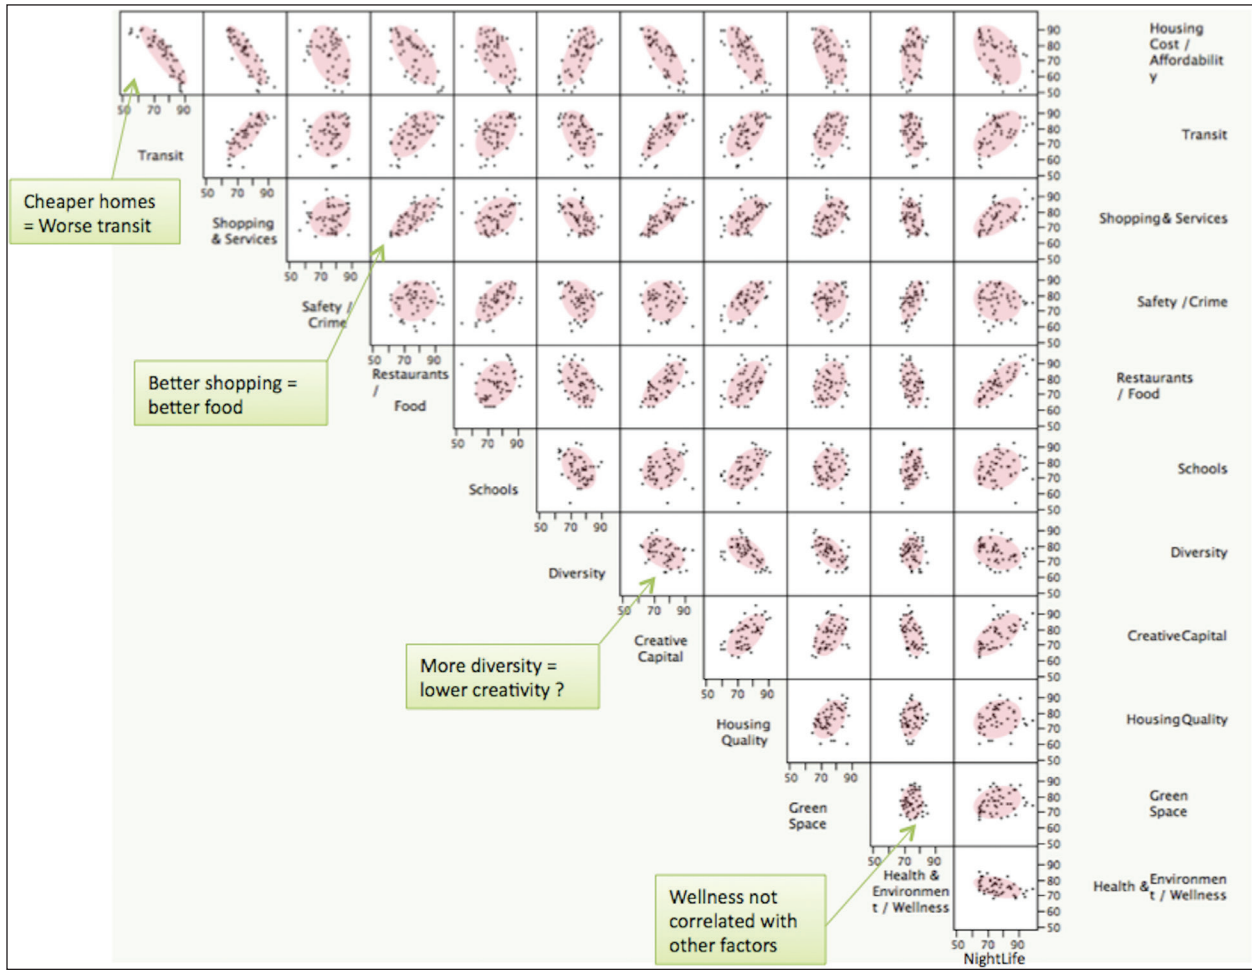


**Figure 2.4**

A multiple-views technique example (from John Stasko, “Multivariate Visual Representations 1” [lecture, Georgia Institute of Technology, Atlanta, GA, September 14, 2016]).

matrix. The scatter plot matrix is one of the graphical tools beloved by visualization designers. A scatter plot renders a visual display to capture the correlation between a pair of variables. Given a set of  $n$  variables, there are  $n$ -choose-2 pairs of variables, and thus the same numbers of scatterplots.<sup>21</sup> These scatter plots can be organized into a matrix, making it easy to represent and observe all pairwise correlations in one place. Figure 2.5 shows an example of a scatter plot matrix visualization based on Nate Silver’s feature article about the most livable neighborhoods in New York.<sup>22</sup> In this article, Silver’s ranking formula considers twelve factors, such as housing affordability, green space, transit, and night life, and calculates an overall score using the individual scores from the factors based on their chosen weights. For example, housing affordability is weighted 25 percent. Scores are then converted to ranks. The intent of this visualization by Kaiser Fung is to look at whether the individual factors are correlated. For example, what is the correlation between housing price and housing quality in a neighborhood? Are more diverse neighborhoods also more creative? What about better schools? Several observations can be easily drawn from this visualization. For example, in the top left corner, the slant shows that the more affordable the houses are, the worse the transit is; the better the shopping is, the better the dining is. An interesting observation, though only a moderate correlation, is that more diversity seems to lead to lower creative capital.

A specific type of visualization that most consider to be highly interesting, perhaps even fascinating, is called Chernoff Faces. Herman Chernoff is an applied mathematician, statistician, physicist, and educator who invented Chernoff Faces to display multivariate data using the shape of a human face. In Chernoff’s paper from the *Journal of the American Statistical Association* in 1973, he proposed that simplified, cartoon-like face shapes were able to represent a number




**Figure 2.5** A scatter plot matrix technique example (from Kaiser Fung, “The Scatter-Plot Matrix: A Great Tool,” *Junk Charts* (blog), June 17, 2010, [http://junkcharts.typepad.com/junk\\_charts/2010/06/the-scatterplot-matrix-a-great-tool.html](http://junkcharts.typepad.com/junk_charts/2010/06/the-scatterplot-matrix-a-great-tool.html)).

of variables in a data set.<sup>23</sup> By mapping numbers to head eccentricity, eyebrow shape, eye size, eye eccentricity, pupil size, nose size, nose width, mouth curvature, mouth width, and mouth openness, a visual understanding could be achieved. The Chernoff Faces help viewers more quickly and precisely detect patterns, groupings, and correlations. The rationale was built upon the fact that the human mind has the innate ability to recognize small differences in facial characteristics and to assimilate many facial characteristics at once. One well-known example of using Chernoff Faces has been used to display life in Los Angeles in 1970, as shown in figure 2.6. This depiction was created by Eugene Turner, a geography professor from California State University.<sup>24</sup> In this map, the four variables to describe life in Los Angeles include affluence, unemployment rate, urban stresses, and the percentage of white population. These variables are presented by facial elements, including face shape, mouth curvature, eyebrow slope, and face color, respectively.

The well-known map accurately and aesthetically captured the living conditions in the Los Angeles area. People attribute the success of the award-winning map to the effective, symbolic use of Chernoff Faces. Turner described this map as “probably one of the most interesting maps I’ve created because the expressions evoke an emotional association with the data.”<sup>25</sup>

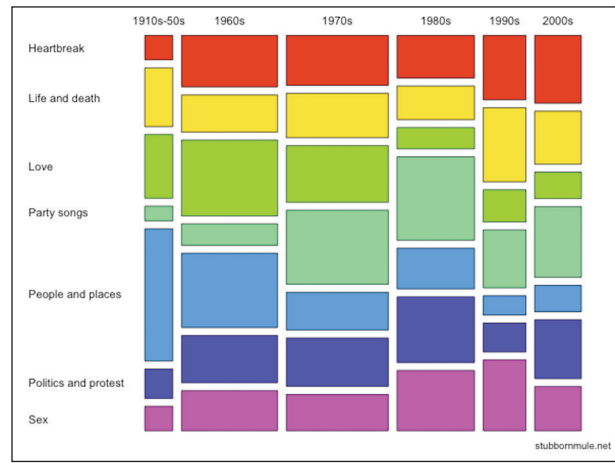
When a data set is comprised of purely categorical data, a popular visualization tool called a mosaic plot is used. A mosaic plot allows analysts to examine the relationship among two or more categorical variables. The mosaic plot starts as a square with length one, which is then divided horizontally based on the proportions of the probabilities associated with the first categorical variable. Then each horizontal bar is further split vertically based on the proportions of the conditional probabilities of the second categorical variable. Additional splits can be made using a third and fourth variable, and so on.<sup>26</sup> Figure 2.7 shows a simple example of a mosaic plot that shows



**Figure 2.6**  
A Chernoff Faces visualization for life in Los Angeles in 1970 (from Eugene Turner, "Life in Los Angeles 1970," California State University Northridge, accessed October 4, 2016, [www.csun.edu/~hfgcg005/eturner/images/Maps/lifeinla.gif](http://www.csun.edu/~hfgcg005/eturner/images/Maps/lifeinla.gif)).

cross-sectional distribution through time of different musical themes in the *Guardian*'s list of "1000 songs to hear before you die."<sup>27</sup>


A visualization technique similar to the mosaic plot is a treemap. Treemaps are ideal for situations in which large amounts of tree-structured (or hierarchically structured) data need to be visualized. The space of a treemap is first split into rectangles, which are sized and ordered based on a quantitative variable. The levels in the hierarchy of a treemap are captured with the visualization of nested rectangles.<sup>28</sup> For example, a rectangle representing a country may contain several rectangles representing states in that country. Each rectangle representing a state may in turn contain rectangles representing cities in these states. A number of different algorithms can be used to determine how the rectangles in a treemap should be sized and ordered, taking into account the general rule of treemapping. In general, the rectangles in the treemap range in size: the largest rectangle will be



**Figure 2.7**  
A mosaic plot technique example (from Sean Carmody, "Mosaic Plot Showing Cross-Sectional Distribution through Time of Different Musical Themes in the *Guardian*'s List of '1000 Songs to Hear Before You Die,'" *Wikipedia*, s.v. "Mosaic plot," last modified July 24, 2009, [https://en.wikipedia.org/wiki/Mosaic\\_plot#/media/File:Mosaic-big.png](https://en.wikipedia.org/wiki/Mosaic_plot#/media/File:Mosaic-big.png)).

placed in the top left corner, and the smallest rectangle will be placed in the bottom right corner, with everything else in between placed in the middle part of the treemap. Hierarchies are presented when the rectangles are nested.

Another widely adopted and useful visualization technique is that of *parallel coordinates*. Parallel coordinates are one of the most famous visualization techniques and among the most common subjects of academic papers on visualization.<sup>29</sup> While one's initial impression of parallel coordinates may be that they are confusing, once they are understood it is easy to see why they are a very powerful tool for understanding multidimensional numerical data sets. How parallel coordinates work can be best explained through an example. Figure 2.8 shows an example of parallel coordinates in which the relationships of car variables are captured, including mileage per gallon (MPG), number of cylinders, horsepower, weight, and the year they were introduced.<sup>30</sup> In this visualization, each of the car variables is mapped onto a vertical axis. As a result, each data value falls within some place on the line and is scaled to lie between the minimum at the bottom and the maximum at the top. And for each case in the row, data points are connected and form the polylines. Useful insights can be gained from this visualization. For example, the cylinders' axis is worth noticing because it has only a few different values. And since the number of cylinders can only be an integer, and there are only eight values in this example, all of the lines will go through a small number of points. In the space between MPG and cylinders, one can tell that eight-cylinder cars generally have lower mileage than six- and four-cylinder ones.



**Figure 2.8**  
Parallel coordinates visualization of car variables (from Robert Kosara, “Parallel Coordinates,” EagerEyes, May 13, 2010, <https://eagereyes.org/techniques/parallel-coordinates>).

The “look” of the lines—that is, how they cross and how they fold—can tell more than just the data. For example, heavy crossings of lines are indicators of an inverse relationship; for instance, the more cylinders, the lower the mileage. A similar correlation can also be found between cylinders and horsepower: the more cylinders, the more horsepower. There are also some crossing lines that show that more cylinders do not always mean more power; however, the general trend is clearly there. Between horsepower and weight, the situation is similar: more horsepower means heavier cars in general, but there is some spread in the values. One can also see that there is a single exception of a high-horsepower eight-cylinder car that is very light. Finally, the lines between weight and year heavily cross, which implies that cars have gotten lighter over the years.

## Software

Information visualization is employed in many cases to answer a question, communicate information, support decisions, and increase efficiency. Good information visualization software is important since it

improves the audience’s understanding of the information being presented. The audience has a higher likelihood of retaining the information, compared to information they receive without visualization. The data visualization engages the audience, and this makes the presentation more interesting. By relaying data to the intended audience in the shortest time possible, presenters also benefit. They do not need to use excessive energy explaining the data to support their message. The use of information visualization software also assists in promoting the credibility and trustworthiness of the data being presented to the audience.

This section provides a list of popular software tools with the objective of encouraging readers to research these tools and give them a try. Through these efforts, readers can connect with the software that best suits their specific needs and preferences.

- **D3.js:** The full name of D3.js is “Data Driven Documents.” The software combines HTML, CSS, and SVG to render charts and diagrams. D3.js is well known for its power and flexibility to create very specific, creative visualizations. It is also packed with features, interactive, and extremely

beautiful. D3.js is free and open-source, which makes it a great tool for amateur learners.<sup>31</sup>

- **FusionCharts:** FusionCharts is well known for its exhaustive collection of charts (90+ chart types) and maps (over 1,000 maps). FusionCharts supports major data formats, such as JSON and XML. Charts can also be exported into a variety of formats, such as PNG, JPEG, PDF, or SVG.<sup>32</sup>
- **Google Charts:** Google Charts uses HTML5 and SVG to create charts that are portable and compatible across different browsers and platforms. It offers a wide range of chart options, such as bar, pie, line, map, and gauge charts. It is also well known for its flexibility and ease of use.<sup>33</sup>
- **Dygraphs:** Dygraphs is an open-source charting library based on JavaScript. It is well known for its ability to handle large-scale data sets. To accommodate “Big Data,” the software is scalable, flexible, and highly customizable. It also has an active support community, which is a great resource for visualization learners.<sup>34</sup>
- **Tableau:** Tableau is one of the most popular, commonly adopted visualization tools. It supports a wide range of charts, maps, graphs, and other visual designs. It also offers strong support for use for academic purposes.<sup>35</sup>
- **Infogram:** Infogram is a software tool that provides the capability to create both charts and infographics. It has a user-friendly interface and well-designed basic charts.<sup>36</sup>
- **Plotly:** Plotly provides a web-based interface for data analysis and visualization. It supports a wide collection of chart types and provides social sharing features. The visualization creation process is intuitive, from loading the data to customizing a visual display.<sup>37</sup>
- **IBM Watson Analytics:** IBM Watson Analytics offers the benefits of advanced analytics without the complexity. It’s a smart data discovery service available on the cloud, which guides data exploration and automates predictive analytics, while also enabling effortless dashboard and infographic creation.<sup>38</sup>

## Conclusion

Through the rapid changes of information visualization that have taken place over the decades, the ability to create compelling, visually stimulating content that can be represented with a variety of techniques is greater than it ever has been. Through dedicated efforts, researchers, scholars, and practitioners have developed visualization theories, techniques, software tools, and applications that can be used for a great variety of purposes and can be directed toward

target audiences in a compelling way. In order to take maximum advantage of these discoveries and techniques, it is important for visualization learners to be open-minded, constantly seeking new knowledge and skills while keeping up with trends and changes.

## Notes

1. Graham Rawlinson, “The Significance of Letter Position in Word Recognition” (PhD thesis, Nottingham University, 1976); Jinxi Caddel, “Word Recognition: How Our Brain Reads Letters,” *Learn, Think, Inspire* (blog), February 4, 2010, [www.jinxiboo.com/blog/2010/2/4/word-recognition-how-our-brain-reads-letters.html](http://www.jinxiboo.com/blog/2010/2/4/word-recognition-how-our-brain-reads-letters.html).
2. Colin Ware, *Information Visualization: Perception for Design* (Burlington, MA: Morgan Kaufmann, 2012).
3. Christopher G. Healey, “Perception in Visualization,” North Carolina State University, Department of Computer Science, last modified September 24, 2016, accessed September 26, 2016, <https://www.csc.ncsu.edu/faculty/healey/PP/index.html>.
4. Rudolf Arnheim, “The Gestalt Theory of Expression,” *Psychological Review* 56, no. 3 (1949): 156–71; Wolfgang Kohler, *Gestalt Psychology: The Definitive Statement of the Gestalt Theory* (New York: Liveright, 1970).
5. Robert J. Sternberg, *Cognitive Psychology* (Fort Worth, TX: Harcourt Brace College Publishers, 1999).
6. Daniel L. Schacter, Daniel T. Gilbert, and Daniel M. Wegner. *Psychology* (New York: Worth Publishers, 2010); Edgar Rubin, “Figure and Ground,” in *Visual Perception: Essential Readings*, ed. Steven Yantis (Philadelphia, PA: Psychology Press, 2000), 225–29.
7. Stuart K. Card, Jock Mackinlay, and Ben Shneiderman, *Readings in Information Visualization: Using Vision to Think* (San Francisco: Morgan Kaufmann, 1999), 7.
8. Edward Tufte homepage, accessed September 26, 2016, <https://www.edwardtufte.com/tufte>.
9. Edward Tufte, *The Visual Display of Quantitative Information* (Cheshire, CT: Graphics Press, 2001), 51.
10. Tufte, *Visual Display of Quantitative Information*; Edward Tufte, *Visual Explanations: Images and Quantities, Evidence and Narrative* (Cheshire, CT: Graphics Press, 1997), 57.
11. Tufte, *Visual Display of Quantitative Information*, 93.
12. *Ibid.*, 93.
13. *Ibid.*
14. *Ibid.*, 107.
15. Scott Bateman, Regan L. Mandryk, Carl Gutwin, Aaron Genest, David McDine, and Christopher Brooks, “Useful Junk? The Effects of Visual Embellishment on Comprehension and Memorability of Charts” (paper presented at ACM Computer Human Interaction Conference, Atlanta, GA, April 10–15, 2010).
16. *Ibid.*
17. Tufte, *Visual Display of Quantitative Information*, 161.
18. *Ibid.*, 171.
19. Michelle Q. Wang Baldonado, Allison Woodruff, and Allan Kuchinsky, “Guidelines for Using Multiple Views in Information Visualization” (paper presented



at Advanced Visual Interfaces, Palermo, Italy, May 24–26, 2000).

20. John Stasko, “Multivariate Visual Representations 1” (lecture, Georgia Institute of Technology, Atlanta, GA, September 14, 2016).
21. Kaiser Fung, “The Scatter-Plot Matrix: A Great Tool,” *Junk Charts* (blog), June 17, 2010, [http://junkcharts.typepad.com/junk\\_charts/2010/06/the-scatterplot-matrix-a-great-tool.html](http://junkcharts.typepad.com/junk_charts/2010/06/the-scatterplot-matrix-a-great-tool.html).
22. Ibid.; Nate Silver, “The Most Livable Neighborhoods in New York,” *New York* magazine, April 11, 2010, <http://nymag.com/realestate/neighborhoods/2010/65374>.
23. Herman Chernoff, “The Use of Faces to Represent Points in K-Dimensional Space Graphically,” *Journal of the American Statistical Association* 68, no. 342 (1973): 361–68.
24. Eugene Turner, “Life in Los Angeles 1970,” California State University Northridge, accessed October 4, 2016, [www.csun.edu/~hfgeg005/eturner/images/Maps/lifeinla.gif](http://www.csun.edu/~hfgeg005/eturner/images/Maps/lifeinla.gif).
25. As cited in J.G. Spinelli and Y. Zhou, “Mapping Quality of Life with Chernoff Faces” (proceedings of Twenty-Fourth ESRI International User Conference, 2004), <http://www.dis.uniroma1.it/~santucci/InformationVisualization/Slides/References/MappingQualityLifeChernoffFaces.pdf>.
26. Sean Carmody, “Mosaic Plot Showing Cross-Sectional Distribution through Time of Different Musical Themes in the Guardian’s List of ‘1000 Songs to Hear Before You Die’,” *Wikipedia*, s.v. “mosaic plot,” last modified July 24, 2009, [https://en.wikipedia.org/wiki/Mosaic\\_plot#/media/File:Mosaic-big.png](https://en.wikipedia.org/wiki/Mosaic_plot#/media/File:Mosaic-big.png); Michael Friendly, “Extending Mosaic Displays: Marginal, Partial, and Conditional Views of Categorical Data” (paper presented at Workshop on Data Visualization in Statistics, Madison, NJ, July 6–10, 1998).
27. Carmody, “Mosaic Plot.”
28. Ben Shneiderman and Catherine Plaisant, “Treemaps for Space-Constrained Visualization of Hierarchies,” University of Maryland, Department of Computer Science, last modified September 2014, accessed September 26, 2016, [www.cs.umd.edu/hcil/treemap-history](http://www.cs.umd.edu/hcil/treemap-history).
29. Alfred Inselberg, “The Plane with Parallel Coordinates,” *Visual Computer* 1, no. 2 (1985): 69–91.
30. Robert Kosara, “Parallel Coordinates,” EagerEyes, May 13, 2010, <https://eagereyes.org/techniques/parallel-coordinates>.
31. Mike Bostock, Data-Driven Documents homepage, accessed October 24, 2016, <https://d3js.org>.
32. FusionCharts homepage, accessed October 24, 2016, <http://www.fusioncharts.com>.
33. Google Charts homepage, accessed October 24, 2016, <https://developers.google.com/chart>.
34. Dygraphs homepage, accessed October 24, 2016, <http://dygraphs.com>.
35. Tableau homepage, accessed October 24, 2016, [www.tableau.com](http://www.tableau.com).
36. Infogram homepage, accessed October 24, 2016, <https://infogr.am>.
37. Plotly homepage, accessed October 24, 2016, <https://plot.ly>.
38. “IBM Watson Analytics,” accessed October 24, 2016, <https://www.ibm.com/analytics/watson-analytics/us-en>.