1

INTRODUCTION

1.1 THE CONTENTS OF THE BOOK

Chapter 2: Principles of Graph Construction

Figure 1.1 graphs an estimate of average temperature in the Northern Hemisphere following a nuclear war involving 60% of the world's arsenal of nuclear weapons. The data are from a *Science* article, "Nuclear Winter: Global Consequences of Multiple Nuclear Explosions," by Turco, Toon, Ackerman, Pollack, and Sagan [127]. The temperatures are computed from a series of physical models that describe a script for the nuclear war, for the creation of particles, for radiation production, and for convection. Figure 1.1 shows that the predicted temperature drops to about −25 °C and then slowly increases toward the current average ambient temperature in the Northern Hemisphere, which is shown by the dotted line on the graph.

In Figure 1.1 the data region is enclosed by a rectangle, the tick marks are outside of the rectangle, the size of the rectangle is set so that no values of the data are graphed on top of it, and there are tick marks on all four sides of the graph. Principles of graph construction such as these are the topic of Chapter 2. The focus is on the basic elements: tick marks, scales, legends, plotting symbols, reference lines, keys, labels, and markers. These details of graph construction are critical controlling factors whose proper use can greatly increase the information gotten from a graph.

Chapter 3: Graphical Methods

Figure 1.2 is a graphical method called a *dot chart*, which was invented in 1981 to display data in which each value has a label associated with it that we want to show on the graph [28]. The large dots convey the values and the dotted lines enable us to visually connect each value with its label. The dot chart has several different forms depending on the nature of the data and the structure of the labels.

The data in Figure 1.2 are the number of speakers for 21 of the world's languages [138, p. 245]. Only languages spoken by at least 50 million people are shown. The data are graphed on a log base 2 scale, so values double in moving left to right from one tick mark to the next.

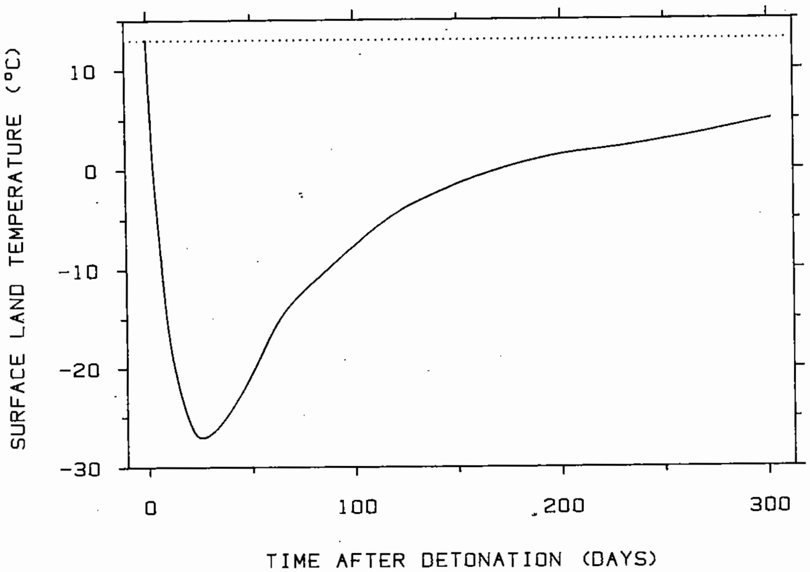

Figure 1.1 PRINCIPLES OF GRAPH CONSTRUCTION. The graph shows model predictions of average temperature in the Northern Hemisphere following a 10,000 megaton nuclear exchange. On the graph, the data region is enclosed by a rectangle, the tick marks are outside of the rectangle, the size of the rectangle is set so that no values of the data are graphed on top of it, and there are tick marks on all four sides of the graph. Chapter 2 is about principles of graph construction such as these.

Figure 1.3 is a graph of ozone against wind speed for 111 days in New York City from May 1 to September 30 of 1973. The graph shows that ozone tends to decrease as wind speed increases due to the increased ventilation of air pollution that higher wind speeds bring. However, because the pattern is embedded in a lot of noise, it is difficult to see more precise aspects of the pattern, for example, whether there is a linear or nonlinear decrease. In Figure 1.4 a smooth curve has been added to the graph of ozone and wind speed. The curve was computed by a method called *robust locally weighted regression*, often abbreviated to *lowess*, that was invented in 1977 [26]. Lowess provides a graphical summary that helps our assessment of the dependence; now we can see that the dependence of ozone on wind speed is nonlinear. One

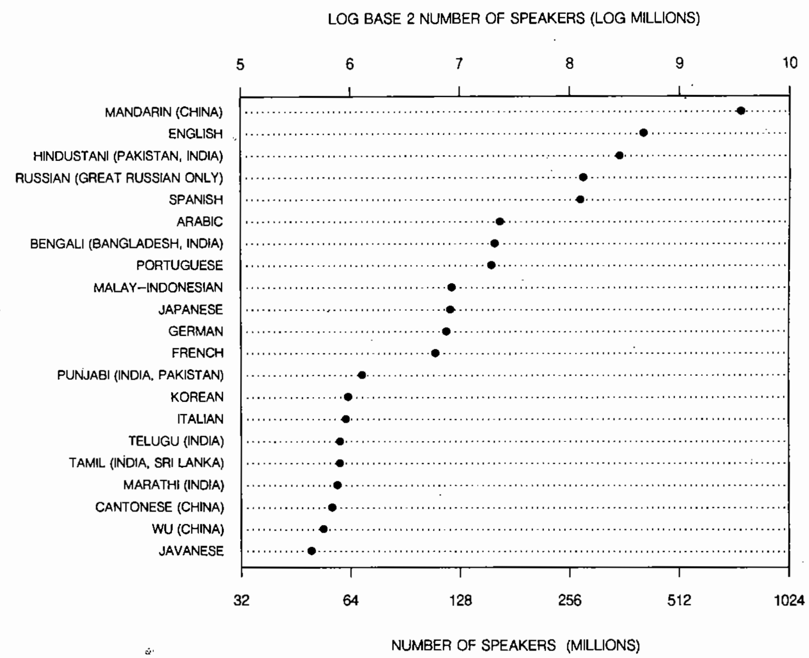

Figure 1.2 GRAPHICAL METHODS. The figure shows a graphical method called a dot chart, which can be used to show data where each value has a label. The data are the number of speakers for the world's 21 most spoken languages. The data are graphed on a log base 2 scale, so values double in moving left to right from one tick mark to the next.

important property of lowess is that it is quite flexible and can do a good job of following a very wide variety of patterns.

Chapter 3 is about graphical methods such as the dot chart, lowess, and graphing on a log base 2 scale. Some of the graphs are methods by virtue of the design of the visual vehicle used to convey the data; the dot chart is an example. Other methods use the standard Cartesian graph as the visual vehicle, but are methods by virtue of the quantitative information that is shown on the graph; graphing a lowess curve is an example of such a method.

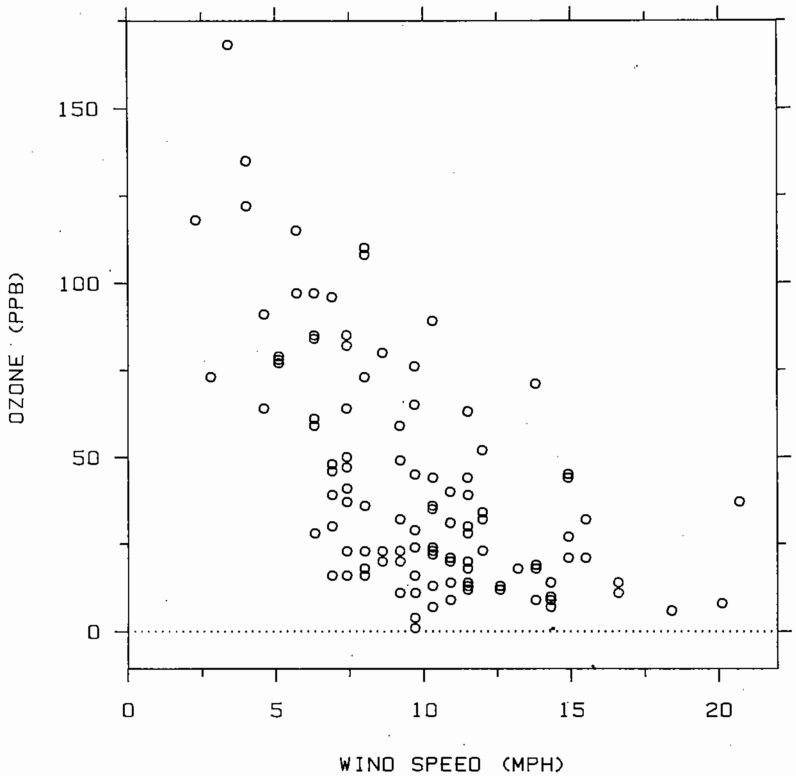

Figure 1.3 GRAPHICAL METHODS. An air pollutant, ozone, is graphed against wind speed. From the graph we can see ozone tends to decrease as wind speed increases, but judging whether the pattern is linear or nonlinear is difficult.

Chapter 4: Graphical Perception

When a graph is constructed, quantitative and categorical information is *encoded*, chiefly through position, size, symbols, and color. When a person looks at a graph, the information is visually *decoded* by the person's visual system. A graphical method is successful only if the decoding process is effective. No matter how clever and how technologically impressive the encoding, it is a failure if the decoding

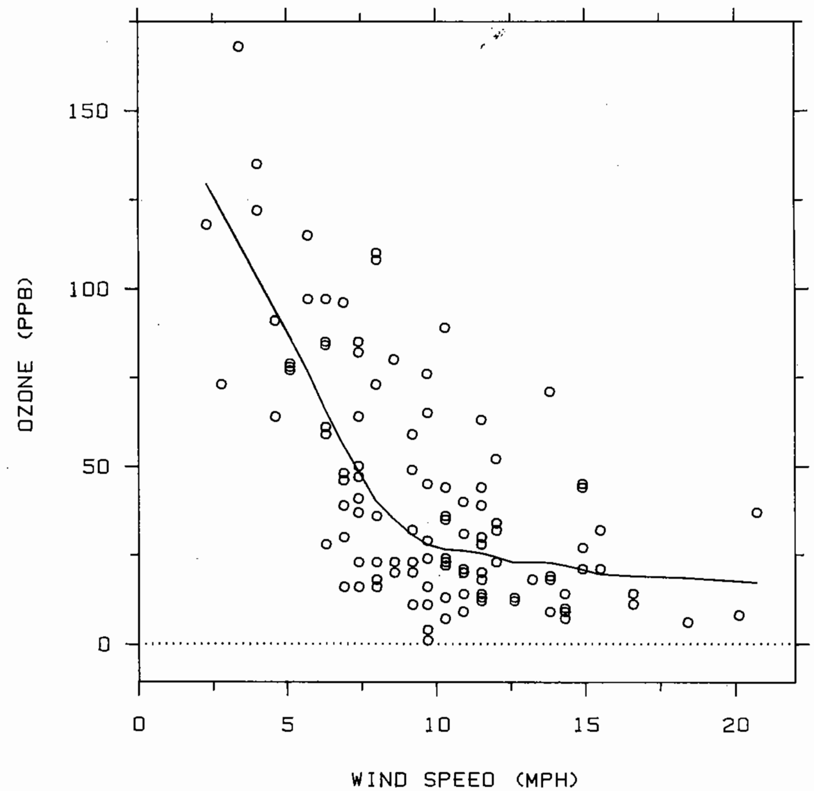

Figure 1.4 GRAPHICAL METHODS. A method of smoothing data called lowess was used to compute a curve summarizing the dependence of ozone on wind speed. With the curve superposed, we can now see that the dependence of ozone on wind speed is nonlinear. Chapter 3 is about graphical methods such as lowess, dot charts, and graphing on a log base 2 scale.

process is a failure. Informed decisions about how to encode data can be achieved only through an understanding of the visual decoding process, which is called *graphical perception*.

Consider the top panel of Figure 1.5 which graphs the values of imports and exports between England and the East Indies. The data were first shown in 1786 on a graph of William Playfair [108] that will

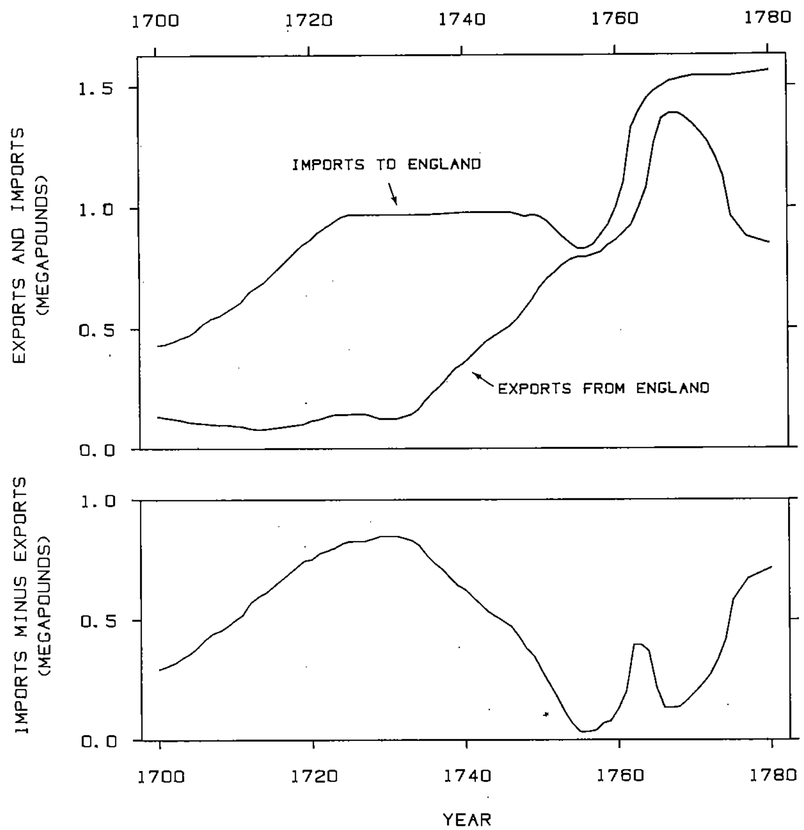

Figure 1.5 GRAPHICAL PERCEPTION. The top panel is a graph of exports and imports between the East Indies and England. The data are from a graph published by William Playfair in 1786. It is difficult to visually decode imports minus exports, which are encoded by the vertical distances between the curves. Imports minus exports are graphed directly in the bottom panel, and now we can see that their behavior just after 1760 is quite different from what we visually decode in the top panel. Chapter 4 deals with issues of graphical perception such as this.

be discussed in Chapter 4. To visually decode the import data we can make judgments of *positions* along the vertical scale; the same is true of exports. Another important set of quantitative values on this graph is the amounts by which imports exceed exports; to decode these values we must judge the vertical *distances* between the two curves.

There is a problem with the top panel of Figure 1.5. It is exceedingly difficult for our visual system to judge vertical distances between two curves when there is a large change in the slopes; we tend to judge minimum distances, which lie along perpendiculars to the tangents of the curves. For example, from the top panel of Figure 1.5 the visual impression is that imports minus exports do not change by much during the period just after 1760 when both series are rapidly increasing. This visual impression is quite incorrect. Imports minus exports are graphed directly in the bottom panel of Figure 1.5 so that the values can be decoded visually by judgments of position along a common scale, and now we can see there is a rapid rise and fall just after 1760.

Chapter 4 is about issues of graphical perception such as this. A *paradigm* for graphical perception is presented. ("Paradigm" is used here in the sense of Thomas S. Kuhn to mean a framework that organizes information [84].) Elementary graphical-perception tasks that people perform in visually decoding quantitative information from graphs are identified. Then, using both the theory of visual perception and experiments in graphical perception, the tasks are ordered based on how accurately people perform them. Also, the roles of detection and distance in graphical perception are investigated. The paradigm has an important application: data should be encoded on graphs so that the visual decoding involves tasks as high in the ordering as possible. This is illustrated by many examples. One result is that new methods are developed and some of the most-used graphical forms are set aside.

1.2 THE POWER OF GRAPHICAL DATA DISPLAY

The premise of this book is that infusing the new knowledge about graphical data display into science and technology will lead to a deeper understanding of the data that arise in scientific studies. Graphs are exceptionally powerful tools for data analysis. The reason is nicely encapsulated in a sentence from a 1982 letter written to me by W. Edwards Deming: "Graphical methods can retain the information in the data." Numerical data analytic procedures — such as means, standard deviations, correlation coefficients, and t-tests — are essentially data reduction techniques. Graphical methods complement such numerical techniques. Graphical methods tend to show data sets as a whole, allowing us to summarize the general behavior and to study detail. This leads to much more thorough data analyses.

One reason why graphical displays can retain the information in the data is that a large amount of quantitative information can be displayed *and* absorbed. This is illustrated in Figure 1.6. Panel 1 (the top panel) is a graph of monthly average atmospheric carbon dioxide concentrations measured at the Mauna Loa Observatory in Hawaii [75]. The panel shows two striking phenomena. One is the persistent long-term rise in CO_2 concentrations due to the burning of fossil fuels. This rise, if continued unabated, will produce the famous greenhouse effect: global temperatures will rise, the polar ice caps will melt, the coastal areas of the continents will be put under water, and the climates of different regions of the earth will change radically [57, 85].

The second phenomenon is the yearly rise and fall of the CO_2 concentrations. This is due largely to vegetation in the Northern Hemisphere. When the foliage grows in the spring, plant tissue absorbs CO_2 from the atmosphere, and atmospheric concentrations decline. When the foliage decreases at the end of the summer, CO_2 returns to the atmosphere, and the atmospheric concentrations increase.

We can get substantial insight into the variation in the CO_2 data by a combination of numerical and graphical procedures. Panels 2 and 3 in Figure 1.6 show numerical descriptions of the long-term trend in the concentrations and of the seasonal oscillations. These trend and seasonal components were computed by a complicated algorithm called SABL [29]. Panel 4 of Figure 1.6 is the variation in the CO_2 that is neither seasonal nor trend; this remainder is just the CO_2 data minus the trend component and minus the seasonal component. On the vertical scales of the four panels of Figure 1.6 the number of units per cm varies. The bars on the right help to show the relative scaling by portraying changes of the same magnitude on the four panels.

Panel 1 of Figure 1.6 allows us to see the overall behavior of the CO_2 data; the bottom three panels allow us to see more detailed behavior. The trend panel shows that the rate of the CO_2 increase is increasing since the slope of the trend curve increases through time; the global CO_2 increase is worsening.

The seasonal panel shows that the seasonal oscillations are getting slightly bigger. For a long time it was thought that these seasonal oscillations were stable and not changing through time, but then around 1980 three groups — one at CSIRO in Australia [106]; a second at Scripps Institution of Oceanography in California [6]; and a third at

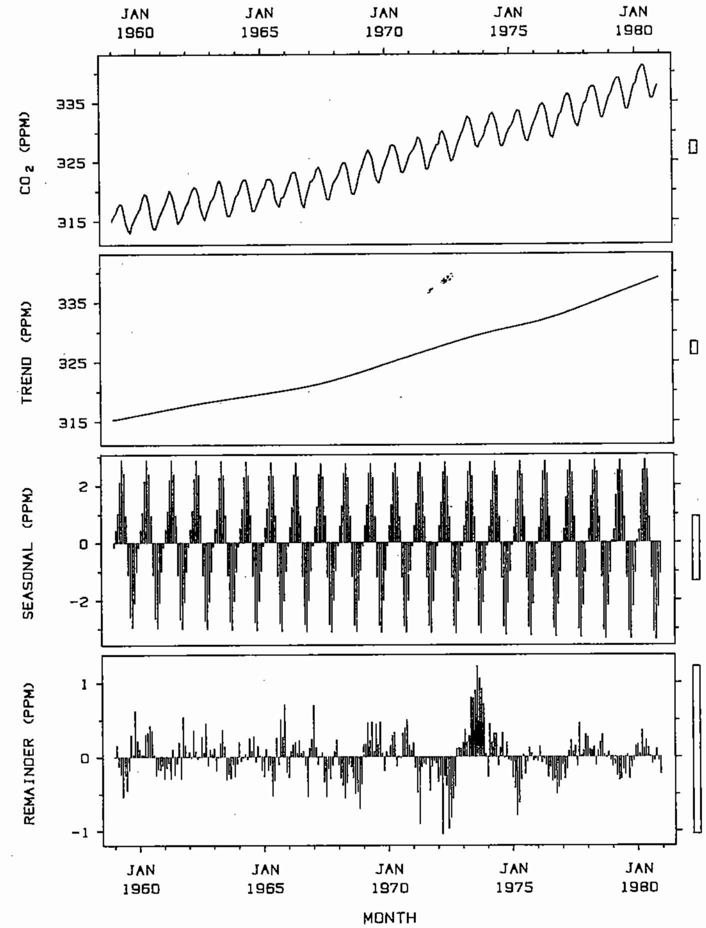

Figure 1.6 THE POWER OF GRAPHICAL DATA DISPLAY. Panel 1 (the top panel) shows monthly average CO_2 concentrations from Mauna Loa, Hawaii. Panel 2 shows a numerical description of the long-term trend in the concentrations, caused by the burning of fossil fuels. Panel 3 shows a numerical description of the seasonal oscillations, which are caused by the increase and decrease of foliage on the earth during the year. Panel 4 displays the CO_2 concentrations minus the trend component and minus the seasonal oscillations. The bars on the right portray changes of the same magnitude on the four panels. A graph like this enabled one group to discover that the amplitudes of the CO_2 seasonal fluctuations are increasing. This visual display shows 2112 numbers. No vehicle other than a graph is capable of conveying so much quantitative information so readily.

AT&T Bell Laboratories in New Jersey [30] — independently discovered the small, but persistent change in the Mauna Loa seasonal oscillations. No one yet has a good understanding of what is causing the change, but a number of scientists are working to determine if it is due to a slow change in the seasonal rise and fall of foliage on the earth or some other mechanism. This small change could well be the harbinger of an important change in the way the earth is working.

Panel 4 of Figure 1.6 shows the effect of another global phenomenon. The values of the remainder show slow oscillations of several years in length; this is revealed by stretches in which the remainder is predominantly above or below zero. These changes in the CO_2 concentrations are correlated with changes in the Southern Oscillation index, which is a measurement of the difference in atmospheric pressure between Easter Island in the South Pacific and Darwin, Australia [5]. Changes in the index are also associated with changes in climate. For example, when the index drops sharply, the trade winds are reduced and the temperature of the equatorial Pacific increases. This warming, which has important consequences for South America, often occurs at Christmas time and is called El Niño — the child [77].

Figure 1.6 conveys a large amount of information about the CO_2 concentrations. We have been able to summarize overall behavior and to see very detailed information. It may come as a surprise just how much quantitative information is shown; there are 1104 data points on this graph and each data point specifies a concentration and a time; thus 2208 numbers are displayed. No vehicle other than a graph could convey so much quantitative information so readily.

1.3 THE CHALLENGE OF GRAPHICAL DATA DISPLAY

Graphical data display is surprisingly difficult. Even the most simple matters can easily go wrong. This will be illustrated by two examples where seemingly straightforward graphical tasks ran into trouble.

Aerosol Concentrations

Figure 1.7 is a graphical method called a *percentile comparison graph* which will be discussed in detail in Chapter 3; the figure shows the graph as it originally appeared in 1974 in a *Science* report written by T. E. Graedel, Beat Kleiner, Jack Warner, and me [31]. (As with almost

all of the reproduced graphs in this book, the size of the graph is the same as that of the source.) The display compares Sunday and workday concentrations of aerosols, or particles in the air. First, the graph has a construction error: the 0.0 label on the horizontal scale should be 0.6. Unfortunately, the error makes it appear that the left corner is the origin; many readers probably wondered why the line $y = x$, which is drawn on the graph, does not go through the origin. A second problem is that the scales on the graph are poorly chosen; comparison of the Sunday and workday values would have been enhanced by making the horizontal and vertical scales the same. (Scale issues such as these are discussed in Chapter 2.) Finally, because in 1974 many of the principles of graphical perception that are discussed in Chapter 4 had not yet been formulated, it did not occur to us then that it is not easy to compare the vertical distances of the points from the line $y = x$; the solution to this problem is a graphical method called the *Tukey sum-difference graph,* which will be discussed in Chapters 3 and 4.

Brain Masses and Body Masses of Animal Species

Figure 1.8 is a graph from Carl Sagan's intriguing book, *The Dragons of Eden* [113]. The graph shows the brain masses and body masses, both on a log scale, of a collection of animal species. We can see that log brain mass and log body mass are correlated, but this was not the main reason for making the graph.

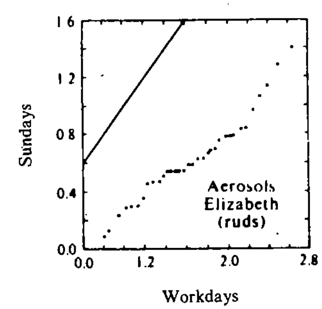

Figure 1.7 THE CHALLENGE OF GRAPHICAL DATA DISPLAY. This graph, made in 1974, compares Sunday and workday concentrations of aerosols. The line shown is $y = x$. The graph has problems. There is a construction error: the 0.0 label on the horizontal scale is wrong and should be 0.6. The horizontal and vertical scales should be the same but are not. Furthermore, it is hard to judge the deviations of the points from the line $y = x$.
Figure republished from [31]. Copyright 1974 by the AAAS.

What Sagan wanted to describe was an intelligence scale that has been investigated extensively by Harry J. Jerison [68]. Sagan writes that this measure of intelligence is "the *ratio* of the mass of the brain to the total mass of the organism." Later he adds, referring the reader to the graph, "of all the organisms shown, the beast with the largest brain mass for its body weight is a creature called *Homo sapiens*. Next in such a ranking are dolphins."

The first problem is that Sagan has made a mistake in describing the intelligence measure; it is not the ratio of brain to body mass but rather is (brain mass)/(body mass)$^{2/3}$. If we study a group of related species, such as all mammals, brain mass tends to increase as a function of body mass. The general pattern of the data is reasonably well described by the equation

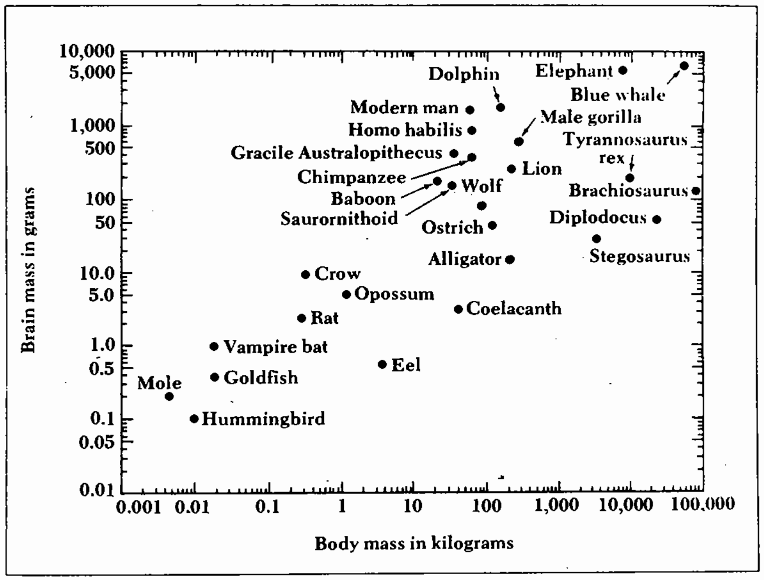

Figure 1.8 THE CHALLENGE OF GRAPHICAL DATA DISPLAY. This graph shows brain and body masses of animal species. The intent was for viewers to judge an intelligence measure; this requires comparing values of $y - 2/3\,x$ for the graphed points, which is difficult to do.

Figure republished from *The Dragons of Eden: Speculations on the Evolution of Human Intelligence*, by Carl Sagan, p. 39. Copyright © 1977 by Carl Sagan. Reprinted by permission of Random House, Inc.

$$\text{brain mass} = c\ (\text{body mass})^{2/3}\ .$$

Since the densities of different species do not vary radically, we may think of the masses as being surrogate measures for volume, and volume to the 2/3 power behaves like a surface area. Thus the empirical relationship says that brain mass depends on the surface area of the body; Stephen Jay Gould conjectures that this is so because body surfaces serve as end points for so many nerve channels [53, pp. 182-183]. Now suppose a given species has a greater brain mass than other species with the same body mass; what this means is that (brain mass)/(body mass)$^{2/3}$ is greater. We might expect that the big-brained species would be more intelligent since it has an excess of brain capacity given its body surface. This idea leads to measuring intelligence by (brain mass)/(body mass)$^{2/3}$.

Let us now return to Figure 1.8 and consider the graphical problem, which is a serious one. How do we judge the intelligence measure from the graph? Suppose two species have the same intelligence measure; then both have the same value of

$$\frac{(\text{brain mass})}{(\text{body mass})^{2/3}} = r\ .$$

Thus

$$\log(\text{brain mass}) = 2/3 \log(\text{body mass}) + \log(r)$$

for both species. This means that in Figure 1.8, the two equally intelligent species lie on a line with slope 2/3. Suppose one species has a greater value of r than another; then the smarter one lies on a line with slope 2/3 that is to the northwest of the line on which the less intelligent one lies. In other words, to judge the intelligence measure from Figure 1.8 we must mentally superpose a set of parallel lines with slope 2/3. (If we attempt to judge Sagan's mistaken ratios, we must superpose lines with slope 1.) This mental-visual task is simply too hard.

Figure 1.8 can be greatly improved, at least for the purpose of showing the intelligence measure, by graphing the measure directly on a log scale, as is done in the dot chart of Figure 1.9. Now we can see strikingly many things not so apparent from Figure 1.8. Happily, modern man is at the top. Dolphins are next; interestingly, they are ahead of our ancestor *homo habilis*. We can also see that this intelligence measure should be regarded as a rough one since it suggests that a goldfish is smarter than a wolf.

It should be emphasized that for some purposes, Figure 1.8 is a useful graph. For example, it shows the values of the brain and body masses and gives us information about their relationship. The point is that it does a poor job of showing the intelligence measure.

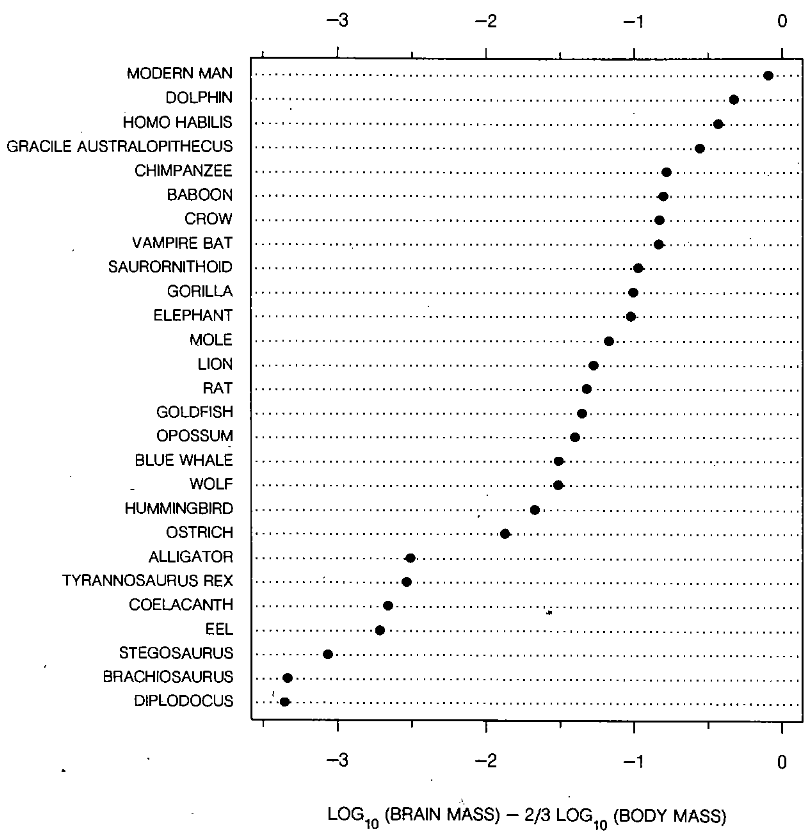

Figure 1.9 DOT CHART. The intelligence measure, log (brain mass) − 2/3 log (body mass), is shown directly by a dot chart. (Both masses are expressed in grams for this computation.) The values of the measure can be judged far more readily than in Figure 1.8. For example, we can see modern man is at the top, even ahead of our very clever fellow mammals, the dolphins.

1.4 SOURCES AND GOALS

Principles of Graph Construction

In 1980 I began a study of graphs in scientific publications. Many people were working hard to develop graphical methods for data analysis, and it seemed reasonable to suppose that buried in the geophysics literature or in the electrical engineering literature or in the literature of many other subjects were clever ideas for displaying data. Indeed, some good ideas were uncovered, but they were a few bright lights standing out in what was mostly a dark picture. In the main, instead of inventiveness, there were errors, poorly explained graphs, graphs where the data could not be seen, graphs where different elements could not be visually disentangled, graphs where the method of display was poorly selected, and graphs that seemed to beg for more or different quantitative information to be shown [79].

In one study, I read the articles and reports of the 1980 Volume 207 of *Science*; there were 249 articles and reports and 67% of them had graphs. I analyzed the 377 graphs, and recorded types, problems, purposes, unconventional practice, possible methods of improvement, and a number of other variables [27]. 30% of the graphs in the volume had at least one of four types of specific problems:

(1) *Explanation* (15.4%) — Something on the graph was not explained.

(2) *Discrimination* (10.1%) — Items on the graph, such as different symbol types, could not be easily distinguished due to the design or size of the graph.

(3) *Construction* (6.4%) — A mistake was made in the construction of the graph such as tick marks incorrectly spaced, mislabeling, items omitted, and wrong scales.

(4) *Degraded Image* (6.4%) — Some aspect of the graph was missing or partially missing due to poor reproduction.

If the only problems uncovered in these studies were those just described, the response could be a few simple guidelines that would eliminate them. But there were deeper problems. First, in many cases the basic graphical form showing the data was poorly chosen. Second, and even more fundamentally, the quantitative information shown on many graphs was poorly chosen. The response to the problems of construction, both superficial and deep, is the principles of graph construction of Chapter 2.

In developing the principles of Chapter 2, I attempted to focus on the basics, to avoid being arbitrary, and to eliminate any principle that was just a matter of style or personal preference. It is a continual challenge in developing principles for graphs not to degenerate into simply expressing personal preferences. William Strunk Jr., prophet of generations of writers and co-author of *The Elements of Style* [120], knew well the tension between freedom and rules. E. B. White writes [120, p. xv]: "Style rules of this sort are, of course, somewhat a matter of individual preference, and even the established rules of grammar are open to challenge. Professor Strunk, although one of the most inflexible and choosy of men, was quick to acknowledge the fallacy of inflexibility and the danger of doctrine." I have tried hard to avoid inflexibility and doctrine in Chapter 2.

Graphical Methods

In the 1960s John Tukey, a renowned statistical scientist and Renaissance man of science, turned his attention to graphical data analysis [126]. Tukey invented a multitude of graphical methods and employed graphs heavily in his book *Exploratory Data Analysis* [125], demonstrating clearly the important role graphs can play in data analysis. This, and the computer graphics revolution, spawned a graphics movement in the field of statistical science, and interest in developing new graphical methods grew rapidly in the 1970s and 1980s [21, 123].

Chapter 3 of this book contains graphical methods that arose in this recent research movement in statistical science, methods from other areas of science and technology, and new methods. The methods selected for discussion in the chapter are useful for all of science and technology and have wide scope in terms of the types of data to which they can be applied. Many specialized methods, useful only in specific fields or only for specialized types of data, are not included. For example, there is a vast methodology for making statistical maps [14, 112] — showing how data vary as a function of geographical location — that is not treated here. Missing also are a number of graphical methods that serve as diagnostic tools for specialized numerical statistical methods [21].

Graphical Perception

In 1981 Robert McGill and I began a series of experiments to probe basic, elementary aspects of graphical perception. The experimentation, together with reasoning from the theory of visual perception, led to the formulation of initial paradigms in a paper in the *Journal of the American Statistical Association* [33] and in an article in *Science* [35]. The material in Chapter 4 draws heavily on these two sources.

Despite the importance of the visual decoding process in graphical data display, graphical perception received very little formal, scientific study in the past. Many have studied the process informally, but informal study is not good enough. Without controlled experiments and measurements there can be no science. Informal study, however, has its value. Intuition flowing from experience is a powerful tool in all areas of science, including graphical perception. We can profitably study graphical perception just by making a graph and looking at it, provided the look is genuinely critical. Certain aspects of the paradigm in Chapter 4 have been derived by researchers in graphical methods — for example, John Tukey [125], Edward R. Tufte [123], Jacques Bertin [14], and Karl G. Karsten [74] — using just such a process of making a graph and studying it.

But intuition and one-subject experiments where researchers study their own graphs can take us just so far. Different researchers will be led to different opinions, some issues are too subtle to submit to just looking, and some phenomena are different from what they seem once you have measured them. To understand graphical perception we need objective numerical measures of people's accuracy in performing graphical-perception tasks, just as measurements are needed in other areas of science. Such a process is behind the paradigm of Chapter 4.

Much of the small amount of experimentation in graphical perception that has been carried out in the past [82, 83] has not led very far because the focus has tended to be the direct comparison of two different types of graphs rather than the probing of basic, elementary aspects. When we visually decode data from a graph, a very complex set of perceptual and cognitive tasks are carried out. Thus, if the basic experimental units are different types of graphs, there is too much complexity and variation to make much progress. In the paradigm of Chapter 4, the complex tasks are broken up into simpler, elementary tasks that then become the focus of the experimentation and theory. Thus the paradigm is an attempt to identify the elementary particles of graphical perception and to describe their interactions and properties.

The more general topic of visual perception has been studied, of course, in great depth. Theories of vision, such as the textons of Bela Julesz [72] and the computational theory of David Marr [94], and the results of experiments in visual perception [8] are important for understanding graphical perception, but the general studies are by no means sufficient for a good understanding of the more specialized topic.

In the past, lack of attention to issues of graphical perception has resulted in the use of data displays that convey quantitative information poorly and in graphical inventions that do not work. Here is one example. In the graphics movement that began in the 1960s in statistical science, much energy was devoted to inventing methods for displaying measurements of three or more variables. An example of such data is daily averages of seven variables — temperature, humidity, barometric pressure, rainfall, solar radiation, wind speed, and wind direction — at one site for 100 days; the data consist of 100 points in a seven-dimensional space. There were many inventions: Chernoff faces [24], Anderson metroglyphs [3], Cleveland-Kleiner weathervane plots [18], Diaconis-Friedman M and N plots [45], Tukey-Tukey dodecahedral views [124], Kleiner-Hartigan trees [78], Andrews curves [4], Tufte rug plots [123, pp. 135-136], and the scatterplot matrix, which is described in Section 6 of Chapter 3. All of the methods in the list, with the exception of the scatterplot matrix, failed in the sense that they almost never showed anyone anything about data that could not be seen more easily by other means. Peter Huber writes [61, p. 674]: "The mere multiplicity of the attempts to deal with more than three continuous dimensions by encoding additional variables into glyphs, Chernoff faces, stars, Kleiner-Hartigan trees, and so on indicates that each of them has met only with rather limited success."

Why did so many methods in the domain of multidimensional data display fail? The answer is that not enough attention was paid to graphical perception. Inventors generated ideas for encoding multidimensional data and did not worry about whether it was easy or hard to visually decode the quantitative information using the methods. Consider Chernoff faces. The values of one point in the space (e.g., the seven values of the meteorological variables mentioned above for one day) are shown by one face. Each variable is encoded by an aspect of the faces (e.g., nose length encodes temperature, the curvature of the mouth encodes humidity, and so forth). The encoding is enormously clever, but the method is of very limited usefulness. Visually decoding the quantitative information is just too difficult.

Chapter 4 is radical insofar as it calls upon us to approach graphs with a new concept: In using graphs and in inventing new graphical methods we should make explicit, conscious use of principles of graphical perception to guide what is used and what is invented.