Chapter 1

How to Graph Badly or What NOT to Do

"The aim of good data graphics is to display data accurately and clearly"
— H. Wainer (1997), pg. 12.

"The greatest value of a picture is when it FORCES us to notice what we never expected to see."
— John Tukey, quoted in (1997), pg. 47.

1.1 Introduction

Like umpiring a sports event, graph-making is best when completely inconspicuous. Therefore, much of the art of visualization is best learned by studying how to make poor, uninformative graphs and then doing the opposite. To put it another way, much of the skill in visualization lies in not making mistakes.

1.2 Chartjunk

"Chartjunk" is Edward Tufte's term for extraneous features that add nothing to the INFORMATION CONTENT of a graph. The worst collections of chartjunk are illustrations in newspapers and non-technical magazines. These figures are drawn by professional artists who are art-school graduates. To them, aesthetic beauty or visual interest is more important than statistics or the shape of a transcendental function. Indeed, many artists had their interests pushed towards art by flunking mathematics and/or science courses! Furthermore, drawing a minimalist graph — just the curves, M'am — is rather boring to an artist who spent years and years in college creating complex drawings and paintings with subtle shadings and color variations.

For all these reasons, commercial artists tend to decorate even the simplest graphs — perhaps *especially* the simplest graphs — with lots of extraneous decoration, i. e., chartjunk. Instead of a simple bar graph in which the height of the bar is proportional to the quantity being graphed, such as the annual oil production of a given country, the artist will draw fancy oil barrels of different sizes whose area is proportional to oil production. The oil barrels will draw the newspaper reader's eye to the graph, but the information is not improved at all by replacing bars by barrels. Indeed, the quantitative usefulness of the graph is actually reduced. Psychological studies have shown that people are much worse at estimating areas than lengths. If one asks a person to compare the relative magnitude of two areas, the

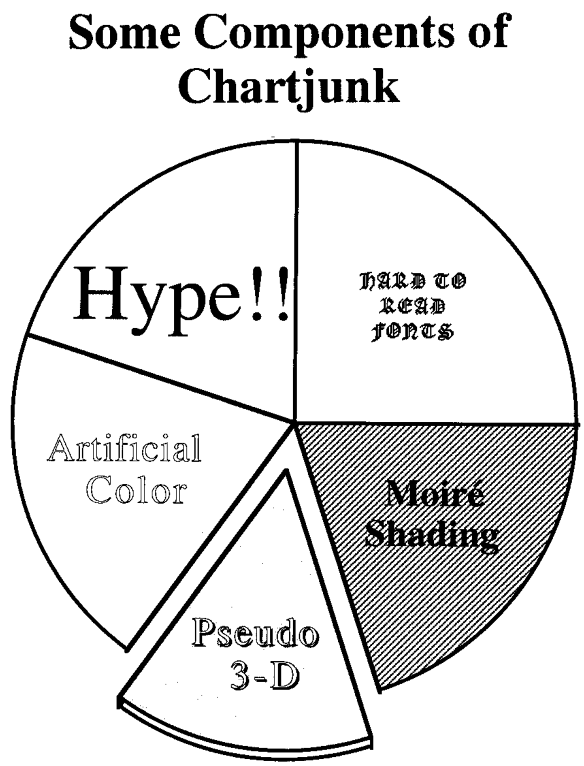

Figure 1.1: Several species of chartjunk. The piece in the upper right quadrant is "Hard to Read Fonts" using so-called "black letter" or "Gothic font" in a small typesize which IS very hard to read.

estimate will almost always underestimate the ratio of areas, and the error worsens as the area ratio increases.

Scientific graphs rarely make the error of using area where length would do. Nor are scientists inclined to add decorations or cute icons to graphs, or make bars into barrels. (In resisting the temptation to add chartjunk, poor drawing ability is actually an advantage!)

Nevertheless, there are some kinds of chartjunk that scientists are prone to including all depicted in Fig. 1.1.

1.2.1 Fonts

It is unlikely that an engineer would label a graph in a Gothic font. However, modern computers offer an enormous range of choices in fonts. This can sometimes tempt one to use five different fonts in a single graph. This may be good if (i) the fonts are all distinctive and are applied to *different* elements of the graph and (ii) look good *together*. Since engineers have no training in font design, the results often are rather garish.[1]

Howevr, the big danger is not ugliness but illegibility. Graphic designers are taught to always think in terms of the end medium, whether printed page or Web site. Engineers don't, and often use type sizes which are too small. A graph may be reduced by 50% in printing; if one doesn't allow for that, a figure that looks fine in the manuscript may be an incomprehensible mess after the publisher has resized the graph, and made all the labels too small to read.

1.2.2 Moiré Shading

Tufte and Wainer, social scientists who have written books on graphics, feel the same way about shading that Scipio Africanus did about Carthage[2] Scientists and engineers are more tolerant because to label complicated regions, there may be no practical alternative, at least in black-and-white graphics. Still, it is important to understand the reasons for Tufte's pique.

The most important is that shading is very attention-getting. If part of a region is marked by white, and the rest by cross-hatching, our eyes will be drawn only to the cross-hatched region. With sufficient will and attention, we can recognize the shape of the other region that is delineated only by a line curve enclosing white. However, this does take will and attention. A hurried reader will remember only the shaded region.

If the shaded region indeed is the most important, then all is well. If the unshaded region is the more interesting, or if both areas are equally important, then we have committed the graphical crime of using an emphatic, attention-getting feature for a part of the graph that doesn't deserve it.

Tufte, Wainer, and illustrators whose training has been primarily in graphic design or commercial art get very upset about graphic emphasis of the wrong thing. Furthermore, the various patterns of shading and cross-hatching, which are often called "Moiré" patterns by analogy with the patterns found on watered silk, are frequently garish and ugly, and visually overwhelm the rest of the graph. Cross-hatched lines often seem to "shimmer" if one looks at them too long. Greyscale shading is free of this problem, but the reproduction of greyscales is rather dicey; a shade that looked good on your local laser printer may look horrible after being copied and reprinted by the publisher.

[1] TEXand LATEXrestrict the user to families of fonts that go well together; for example, Knuth's Computer Roman font, the default in TEX, was designed for visual compatibility with the italic fonts used for equations and mathematical symbols.

[2] Scipio was a Roman politician, flourished circa 170 B. C., who ended every speech in the Roman Senate, even if he was talking about a completely unrelated topic, with *"Carthago delenda est"*, which means "Carthage must be destroyed".

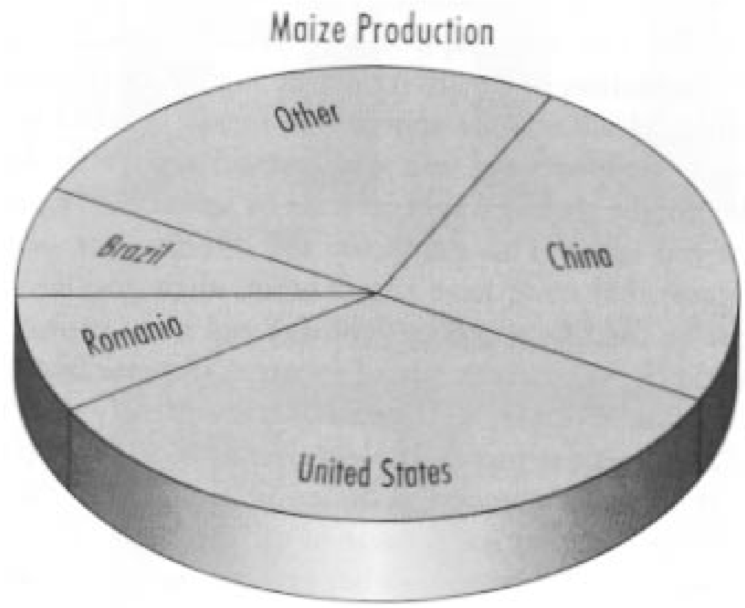

Figure 1.2: The lip that has been added to make this pie chart three-dimensional gives a misleading impression. The eye perceives both the front lip and the segment on top of the "pie" as a single entity, making the share of the United States seem larger than the others.

Another point to remember is that black-and-white can be stored as 1 bit-per-pixel. Grayscale drawings require 8 bits-per-pixel. It is rather silly to amplify storage needs just to make the drawing more confusing!

1.2.3 Pseudo 3-D

Plots that project into the third dimension, such as surface mesh or "net" plots, can be very useful. However, newspapers and magazines are often victimized by the artist's creativity. A simple bar chart may be drawn with elaborate three-dimensional columns in perspective. The difficulties of accurately perceiving size in three-dimensional space — a column that appears small may look small only because it is located at the back of the projected space — that the graph becomes almost useless for gaining a visual feel for the numbers.

Fig. 1.2 shows an ordinary pie chart which has been elaborated, quite needlessly, into a three-dimensional cylinder. The height of the cylinder magnifies the perceived area of the pie segment which is nearest the viewer. In this case, a lot of extra care and trouble in drawing the graph has made it better art but poorer science.

As computer graphics software becomes more ubiquitous and capable, scientists and engineers can easily succumb to the same temptation.

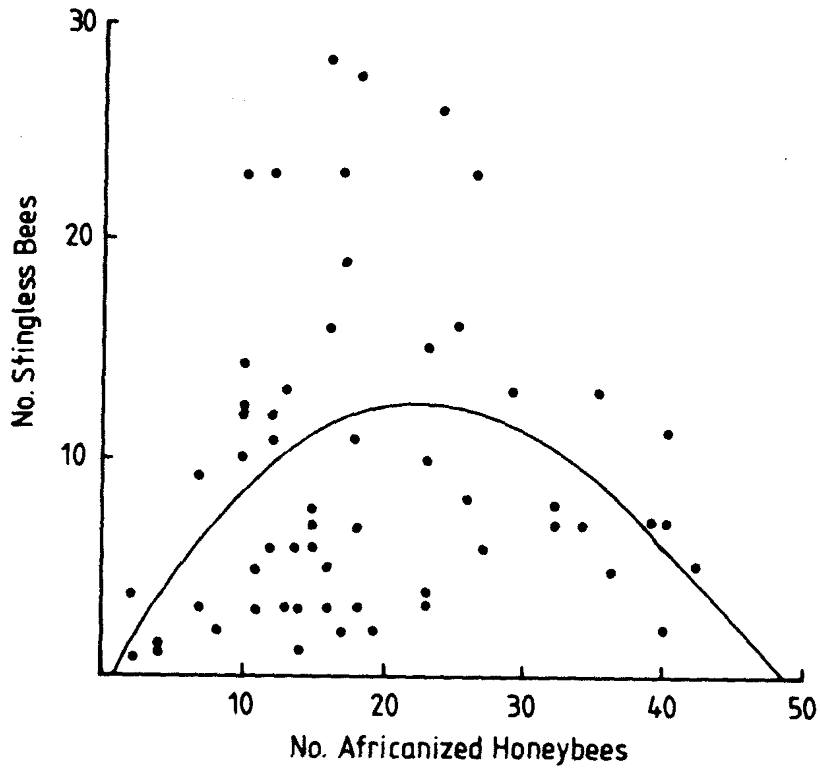

Figure 1.3: The relative abundance of the Africanized (horizontal axis) and stingless bees (vertical) as measured on flowering *Melochia villosa*. The solid line is a quadratic polynomial which gave the best fit to the points. (Taken from a note by by David W. Roubik, 1978.)

1.2.4 Artificial Color

Color is a powerful tool in scientific visualization, but precisely because it is powerful, it is easily misused. It has been my experience that when a research group purchases a color printer, the result is a flood of illustrations which are very colorful but don't actual convey much information. Furthermore, color is very expensive in storage (a color image requires at least eight times as much storage as a line graph (black-and-white) of the same resolution). It is very expensive to print, and so most journals require the author to pay the cost of color reproduction, often $1000/page or more.

So, the moral is: Use color. But use it wisely, which means when you really need it to make clear what would otherwise be confusing.

1.2.5 Hype, I: Overinterpretation

One big danger for scientists and engineers is over-interpreting the data. Sometimes, experiments just don't work. The experimental noise may overwhelm the signal. The honest reaction is to admit this. The pressures to keep publishing so that one can keep grant funding, get tenure, etc., make it difficult to do this.

Fig. 1.3 shows the results of an experiment in biology: measurements of the numbers of bees of two different species of honeybee on a single species of flower. The purpose of the experiment was to understand the relationship between the two species. The Africanized bees are ferocious, aggressive stingers, very dangerous to humans. The stingless bees are

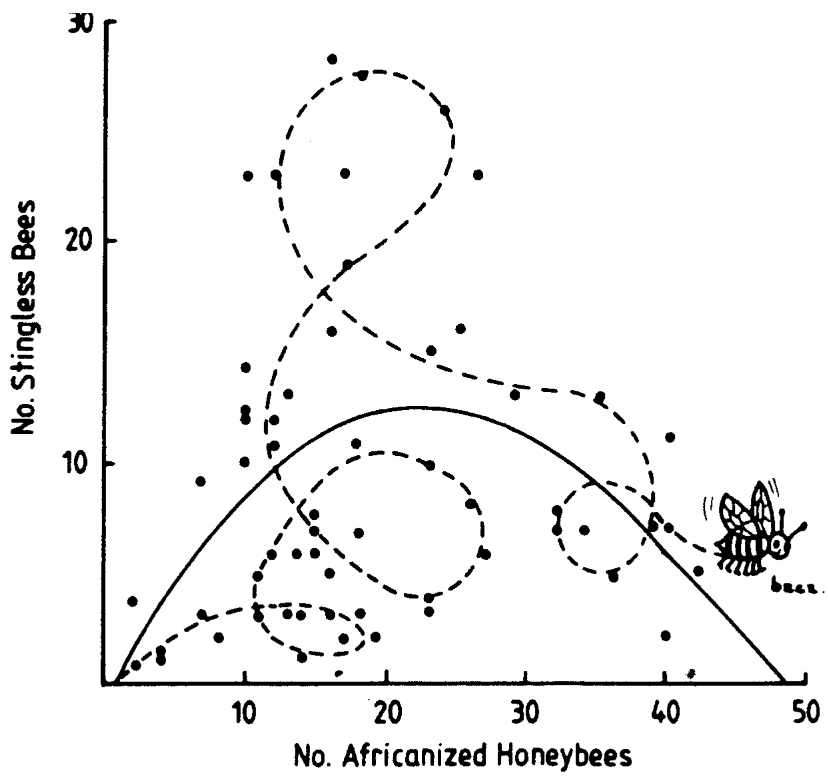

Figure 1.4: Satirical redrawing of the figure from Roubik(1978), published by Hazen (1978).

quite harmless to people. One expect that the aggressive African bees would chase the other species away, or sting them to death. The actual result of the measurements is a nearly-random cloud. Nevertheless, the experimenter went ahead and fitted a parabola through the data. The mathematical form of the parabola is given in the original caption, but the average value of the difference between the parabola and the data points is not omitted from both caption and abstract.

Robert M. Hazen, a geophysicist who has written half a dozen popular (and very readable) nonfiction books about science, was moved by what he describes as the "rather fanciful curve fitting of Roubik" to propose "an alternative interpretation".

Dr. Roubik replied, "I think Dr. Hazen was right in being sceptical, but I do not think it would justify disregarding the study or my conclusions. I thought that his letter to the editor was hilarious, but some of my colleagues did not. It seems to me that biologists are often obliged to take a different view of quantitative data from that of physical scientists. They have more or less set rules, while we must often try to discover nature's meaning. And there is a lot of slop in nature."

Well, yes, but it was still idiotic to fit a parabola. The danger of fitting a curve to data that is all over the place is that the tidiness of the curve suggests that we actually understand the underlying phenomena, that we have a useful theory for it. If one looks at the center of the graph (20 Africanized bees on a flower), one sees that the observed number of stingless bees ranges from a low of 2 to a high of about 25. What predictive value has a theory so crude?

It is likely that the ratio of stingless bees to Africanized bees is really controlled by factors different from the only factor considered in the graph: the number of bees of the other species.

One can't be too hard on Dr. Roubik. He had the good grace to describe Hazen's satire as "hilarious". Graphical crimes as bad or worse are very common, alas, and most of the time there is no Hazen to inject a note of common sense.

1.2.6 Hype, II: Graphical Carpet Bombing

The ease of making graphs by computer has lead many students to the following (false!) attitudes:

- I graph, therefore I think.
- I graph, therefore I work.
- I graph, therefore I progress in understanding.

Graphs can be very useful to the goals of thinking and understanding; indeed, the whole purpose of the course is to improve the connection. It is very easy, however, to become so caught up in the mechanics of running cases and printing graphs that one forgets what it all *means*.

The result is that many Ph. D. theses consist largely of an endless progression of graphs, interrupted by brief bits of text that do little except repeat what's in the captions. To make the thesis even more boring, the graphs are usually all of the same type: there are contour plot theses, and line graph theses, but rarely are different types of graphs mixed together in the same thesis.

Complicated phenomenon unfortunately necessitate using a lot of intricate graphs. However, it is easy to fall into the trap of lobbing graphs at the reader, one after another, until the reader is hopelessly lost and bored. Students are especially prone to this because there is no page limit for a thesis. Older scientists are less vulnerable because during the review process, they will be yelled at by referees and editors, and thereby encouraged to shorten their journal articles, and be more selective in choosing graphs for inclusion.

We may dub this a kind of graphical "carpet bombing" after the military practice of dumping large numbers of bombs to saturate an area when individual targets can't be identified. In the Pacific Gulf War, for example, carpet bombing was done by three-plane elements of B-52s, all dropping their loads of 84 500-pound bombs per plane simultaneously over an area of perhaps a square kilometer or so. Many scientists do the same thing only with almost-identical illustrations replacing the bombs.

If the goal is to daze and confuse, carpet bombing is quite effective. The Allies took tens of thousands of prisoners who had been protected by trenches from the bomb fragments, but who nevertheless were left too dazed, deafened and disoriented to fight.

If the goal is to enlighten rather than to concusse into submission, carpet bombing is *not* a good idea. When a phenomenon is complicated and intrinsically requires a lot of graphs, more visualization skill is needed, too.

There are several remedies when one needs to present a lot of complicated information graphically including:

- Varying the graph type; mix contour plots, surface plots, line graphs and so on.
- In each graph, highlight the key figures in the caption and on the graphic itself.
- Triage[3] the analysis: spend a lot of words (and graphs) on the important points, but few words and perhaps no graphs on the unimportant ones.
- Combine many closely-related graphs into a single multi-panel graph.

Fig. 1.5 shows a rather extreme example of a multi-panel graph: 384 separate line graphs combined into one illustration. There are some obvious difficulties with combining so much information into a single figure. (The journal, which uses 8.5 inch by 11 inch pages, tried to ameliorate the legibility problem by turning the figure sideways and then filling the entire page with this single graph.) Nevertheless, it is a lot better than printing each of these 384 panels on a separate page — although I have seen technical reports of similar ocean data that did just that!

[3] "Triage" is a medical term for a practice which is now standard in most American emergency rooms, but was originally developed by French military surgeons during World War I. Overwhelmed by casualties, the French field hospitals split the wounded into three groups ("tri") by preliminary examinations: (i) Those whose wounds were not serious and would recover without treatment (ii) those whose wounds were fatal even with treatment and (iii) those who had serious wounds but might survive if treated promptly. The triage system has saved uncounted numbers of lives. Scientists need to set priorities with graphs, too.

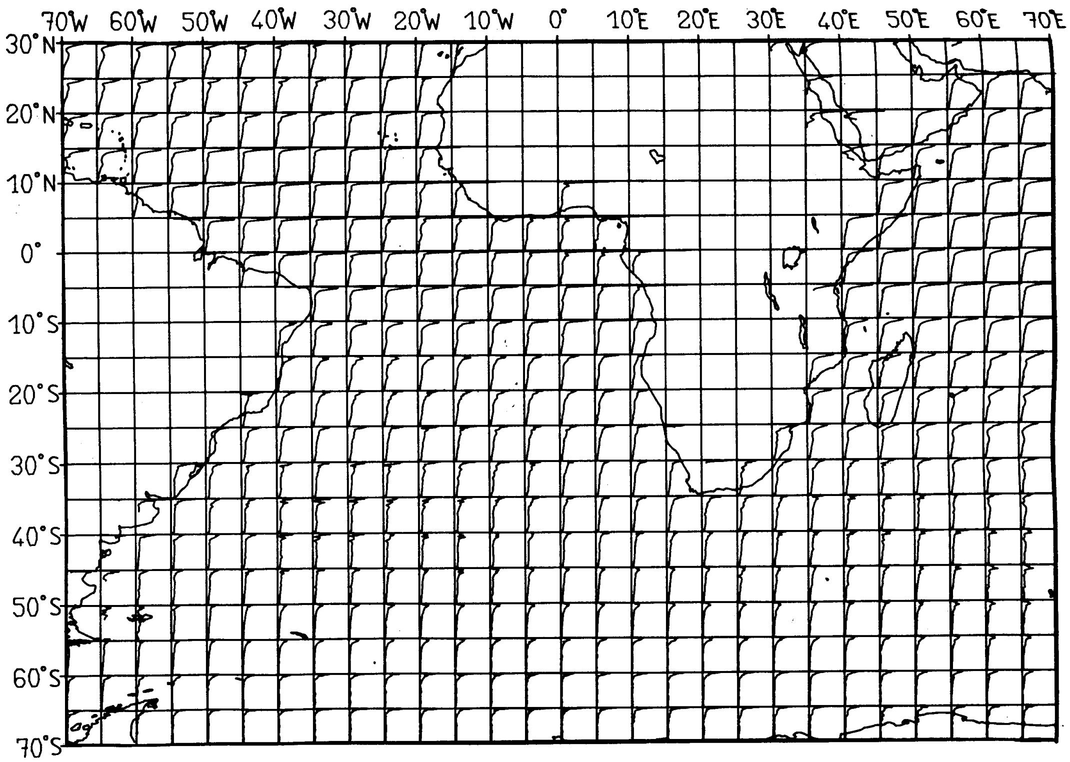

Figure 1.5: A graph containing 384 separate line graphs (excluding the land areas). Each one-dimensional plot shows the mean Brunt-Vaisala frequency within a given square. From Houry, Dombromsky, de Mey and Minster(1987).

1.3 Wainer's Rules for Bad Graphs

1. Show as little data as possible [minimize the data density]
2. Hide what data you do show [minimize the ratio of data/ink]
3. Show the data inaccurately [ignore the visual metaphor and randomize the connection between graphical elements and the numbers]
4. Use length as the visual metaphor when the area of two-dimensional icons is what is actually perceived
5. Graph data out of context [sparse captions and vague text]
6. Obfuscation #1: Change scales in mid-axis
7. Obfuscation #2: Emphasize the trivial [ignore the important]
8. Obfuscation #3: Jiggle the baseline [use different axis ranges for two graphs which will be printed side-by-side and need to be compared]
9. Obfuscation #4: Alabama first! [Order the data by some criterion, such as alphabetical order, which is irrelevant to all of the interesting patterns in the data]
10. Obfuscation #5: Label: (a) illegibly (b) incompletely (c) incorrectly (d) ambiguously
11. Obfuscation #6: More is murkier: (a) more decimal places and (b) more dimensions
12. If it has been done well in the past, think of a new way to do it [New graph types are sometimes needed, but they require a lot of concentration from the reader, and should be used sparingly in GOOD graphics]

In the next few sections, we will illustrate some of these principles.

1.4 High Data Density

A theme that is reiterated again and again by Tufte, Wainer and other authors is that high-density illustrations are good: it is possible to pack a tremendous amount of information in a single picture if it is designed carefully. Fig. 1.6 is an example.

Conversely, graphing badly is to display graphs with little information. Occasionally, a low-density graph is justified if used to emphasize information or concepts that are very important. A low-density graph is always justified if it is sufficient to illustrate a concept.

Sometimes, though a bar graph that presents only four or five numbers is a just a bad alternative to a table.

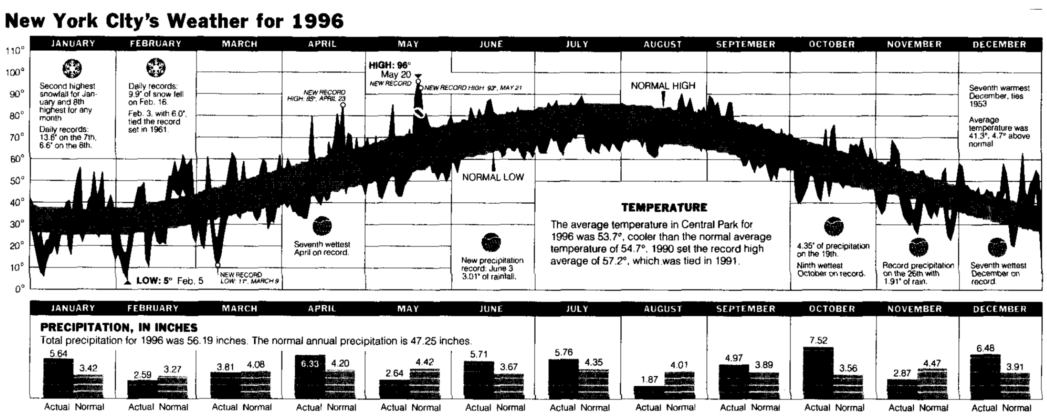

[h]

Figure 1.6: A good example. Although this graph is crammed with a great diversity of information, it is quite readable. One can learn the high and low temperatures for the year, the annual trends, unusual weather on individual days and visually perceive the fluctuations of temperature with time. Taken from annual weather summary of the *New York Times* newspaper.

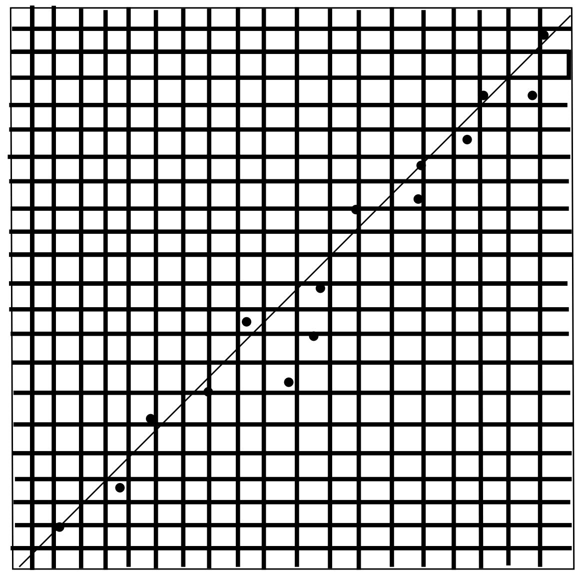

Figure 1.7: A bad example of "Data Hiding" : the grid is so heavy that it almost totally obscures the theoretical line and the data points. Inspired by an example published in a scientific journal and later analyzed by Tufte and Wainer.

1.5 Data-Hiding

An excellent strategy for making a really bad graph is to add so many extraneous elements to the graph that the data and/or information content is obscured. Fig. 1.7 is similar to one in an article in a scientific journal, subsequently analyzed (in very unfriendly fashion) by Tufte and Wainer. The grid is so heavy that the data points are almost invisible. But any child can make a grid; the data points are the heart of the illustration!

Amazingly, a textbook author made an even worse graph by reprinting the earlier figure *without* the data points (Fig. 1.8)! If the curve were complicated and tracked the data closely, this decision could perhaps be defended — but probably not; why not include the points? Here, however, the plot is sheer lunancy because the fitted curve is just a diagonal straight line. Never use a graph where words would be just as clear.

The right way to present the information is to include both the scattered points and the fitted curve but to omit the grid as shown in Fig. 1.9.

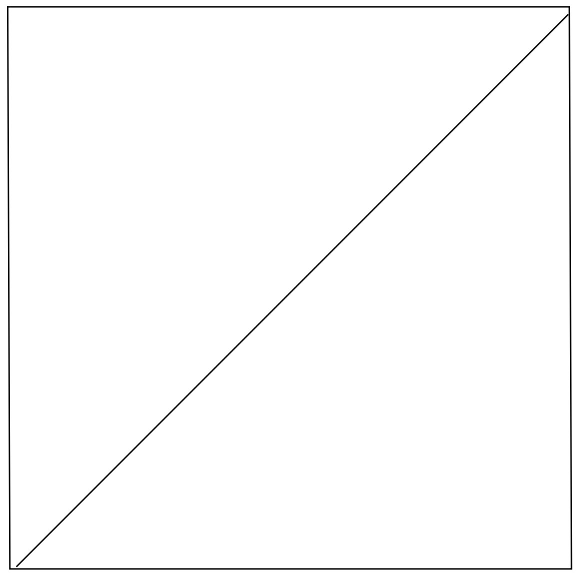

Figure 1.8: A later author reprinted the figure in a textbook, but omitted all the data points! The straight line, which is the result of a theoretical model, is a diagonal line with a slope of exactly forty-five degrees. This graph really conveys no information at all — a classic bad example.

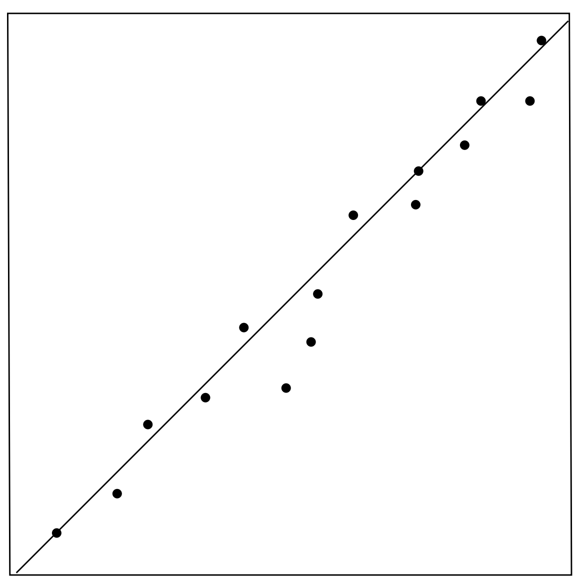

Figure 1.9: E. Tufte's redrawing of the figure: eliminating the grid enormously improves clarity. Nothing remains but the essentials: the scatter points, the straight line which is fitted to them, and the axes.

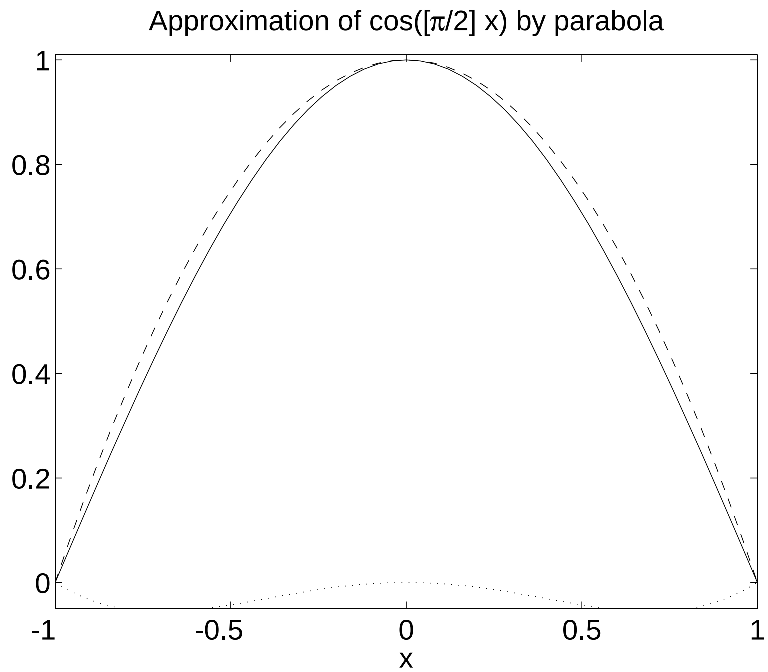

Figure 1.10: Comparison of the cosine function (solid) with a parabola $u_{parab} = -x^2$ that approximates it (dashed) and the difference between them (dotted). The difference is so small compared to the cosine, however, that it is almost invisible.

1.5.1 Data-Hiding by Graphing Disparate Quantities on the Same Scale

A common problem is to depict quantities which are all significant even though their magnitudes are very different. It is then very easy for the curve of the small to be so dwarfed by the graph of the large that no useful information about the small variable can be gleaned from the graph except its order of magnitude. Fig. 1.10 depicts such a case. The difference between the cosine function and the approximating parabola is important because it tells us the accuracy of the approximation. Unfortunately, the error is so small that its graph, when plotted on the same scale for the cosine, is almost a horizontal line. Is the maximum difference 6%? 2%? Hard to tell from the graph.

The best strategy is to split the graph into multiple panels and graph the big and the small separately, each with its own scale, as in Fig. 1.11. It is now possible to draw a couple of conclusions from the graph: (i) the maximum error is about 6% of the maximum of the function which is being approximated and (ii) the error is highly oscillatory and uniform over the interval, unlike the error in the power series approximation $u_{power} \approx 1 - (\pi^2/8)x^2$, which rises sharply away from the origin.

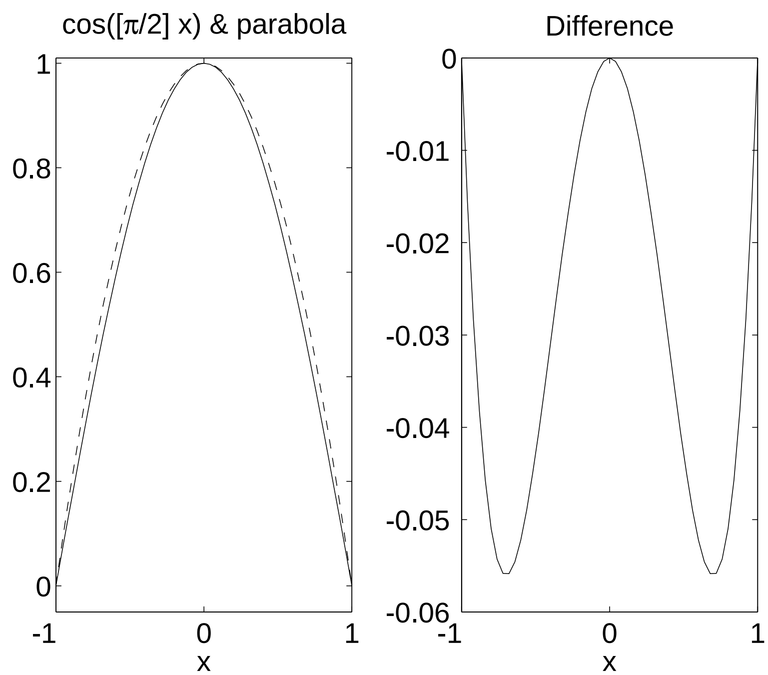

Figure 1.11: Same as the previous figure except that the plot has been split in two so that the error (right panel) can be plotted on its own scale.

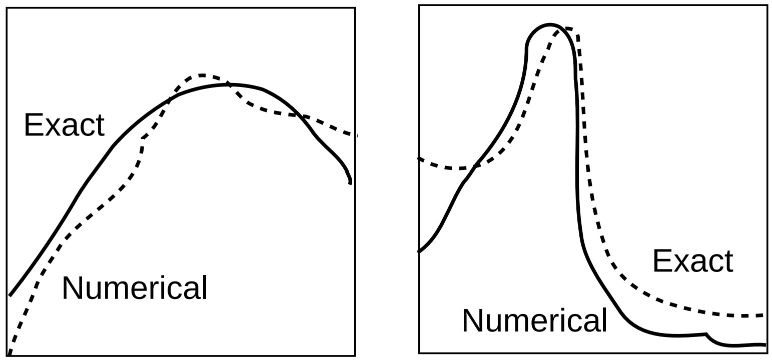

Figure 1.12:

1.6 Inconsistent Visual Metaphor

"Visual metaphor" is a broad term for the relationship between specific graphical elements and the data. For example, in a sequence of graphs that compare the exact solution with the corresponding numerical solution, one visual metaphor might be to associate the solid curve with the exact solution and the dashed line with the numerical calculation. Fig. 1.12 shows a graph that apparently employs this metaphor except that in the right graph, the metaphor has been ACCIDENTALLY REVERSED so that the exact solution is dashed.

Inconsistent use of a visual metaphor is rather common. When testing a numerical algorithm against a set of problems with known solutions, it may take a week or two to run all the test cases. It is terribly easy to forget on Wednesday that the exact solution was a solid curve on Monday.

1.7 Context-Free Data

The "context" of a graph is the engineering or physics background. It is difficult to describe "context" in purely graphical terms because the lack of adequate context is more a technical deficiency than a visual failure. A few general remarks are possible.

First, a graph is a failure if the text and caption fail to give the reader enough information so that the graph is comprehensible. A good graph-with-caption will label all the elements of the graph and specify the key parameters of the numerical calculation or experiment that generated the graph. However, it is not sufficient to clearly label a curve "Supercalifragilousness" and state in the caption that the "humdinger" was set at 360 "klingons" if these terms are unfamiliar to the reader and the text fails to explain. A physicist or engineer may be reasonably expected to know the meaning of "acceleration" and "mass", but it is not reasonable to expect all readers to know the atomic number of praseodymium or the decay rate of the chief isotope of meitnerium.

Second, much of science is about comparisons. Indeed, it has been argued that the primary function of graphs is to facilitate comparisons. A graph will fail through lack of context if its curves fail to make all the important comparisons.

For example, suppose one has developed a new algorithm for numerical quadrature. A graph that demonstrates that the error decreases with an increasing number of quadrature points is encouraging. But a single curve is limited by its lack of context. How does the algorithm compare with older methods like Romberg integration and Gauss-Legendre quadrature? How does the execution time vary with the number of quadrature points, both in absolute terms and also relative to competing methods?

The most important characteristic of a good graph is that it show enough curves — and the article as a whole contain enough information — so that these kinds of questions can be answered. A graph showing error versus the number of points N is meaningless. A graph showing three curves for three different algorithms may make you immortal.

An analogy may be helpful. Each year, over two thousand home runs are hit in major league baseball. Most are insignificant. A home run by Mark McGwire in September of 1998 made him one of the immortals of baseball, and guaranteed his future election to the Baseball Hall of Fame in Cooperstown, New York. The context of his 62d home run of the season was: no one in the hundred year history of professional baseball had ever previously hit more than 61 in a single season. The home run by itself was nothing. (His team lost the game!) The context was everything.

1.8 Area Instead of Length as a Visual Metaphor

Newspaper artists love dearly to turn even the simplest graph into a Work-of-Art. This is often a really bad idea.

Fig. 1.13(top) has several flaws. First, beer sales is represented by various sizes of beer barrels. The use of an icon can be helpful; it is almost impossible to forget the topic of the graph — beer sales — while staring at icons of barrels plus the name of one of the world's largest brewers (Schlitz) in bold letters on the graph. However, the surface area and volume of the barrels both grow wildly out of proportion to the associated numbers. The largest barrel has at least ten times the volume of the smallest, but sales actually grew by less than 30%.

Even if the artist had been more careful, the use of area or volume to represent quantity is questionable. The problem is that it is difficult for the eye to estimate areas and volumes accurately. Psychological studies have shown that estimates of volume increase more slowly in the mind than on the graph. The reason seems to be that unconsciously, we simultaneously estimate the length, area and volume of three-dimensional objects. If we see two barrels that differ by 30% in volume, the mind also registers that their surface area differs by just 20% and the width and height by only 10%. When test subjects are asked to estimate the change in volume, these three different numbers seem to be averaged in the reader's mind to arrive at an answer of 20% when the volumes differ by 30%.

The figure also has unnecessary greyscale shading (in the main panel) and solid black shading (with "Schlitz" incised in white within) in the inset. The barrels themselves also distract the eye from the numbers and the axis labels, which are almost invisible in comparison.

The lower graph is Wainer's minimalist rendering of the same data. It is much easier to see what is really happening. First, overall sales are up (good!) but not by a huge amount. Second, Schlitz' market share is going *down* for the last couple of years. Schlitz management cannot take credit for the increase in sales; beer has simply become more popular, and Schlitz' competitors have been increasing their sales even faster. No wonder Schlitz wanted to hide the truth in rapidly-expanding beer barrels and garish shading!

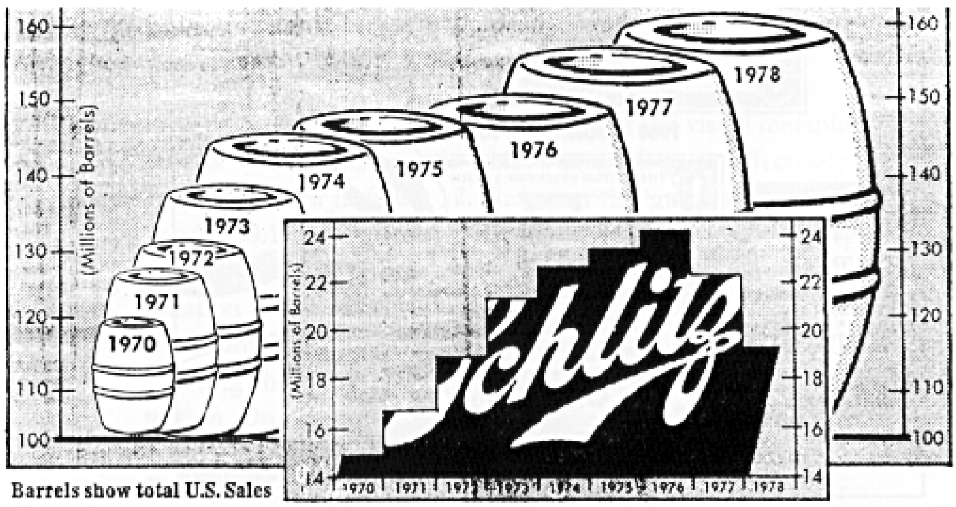

(Millions of Barrels)
1970
1971
1972
1973
1974
1975
1976
1977
1978
Barrels show total U.S. Sales

U.S Beer Sales Grew Steadily Throughout the 1970s
Millions of Barrels
Year
Schlitz's Share of the Market
Percent
Year

Figure 1.13:

1.9 Label Woes

Obfuscation #5 is to use labels that are illegible, incomplete, incorrect, and ambiguous. These would seem to be such obvious boo-boos that no one past elementary school would make them, but they are actually rather common.

Illegible labels usually result from the following causes:

1. Too small type size.

2. Poor placement.

3. Too few labels.

One must be careful about choosing a sufficiently large type size because of reproduction during publishing. When the journal is typeset by the publisher, figures are almost always reduced in size from the manuscript. When a figure is published in an Internet journal or by being included on your own Web page, the resolution is reduced to the 72 dpi of a standard monitor — perhaps lower if the figure has to be converted from its original format (Postscript, say) to a format popular on the Web, such as JPEG. Thus, an illustration which looks good as 300 dpi laserprinter output may display illegible labels on the screen.

Label placement can be a difficulty, too. A contouring routine offers automatic placement of numbers to label each contour line. However, to find sufficient room to place the numbers without overlapping other contour lines, the labels may be scattered widely — the label for 1.0 near the top, the label for 0.75 on the far left, that for 0.5 on the far right, etc. This can be very confusing for the reader. Manual placement is sometimes a great improvement. However, manual label placement is time-consuming at best. At worst, one may be faced with clustering of lines which makes effective labeling between the lines impossible. Sometimes, what one really needs to do is to erase parts of each line to insert the labels in the erasure regions, but most computer graphics programs do not allow this. Sometimes, a drawing program will permit such modifications.

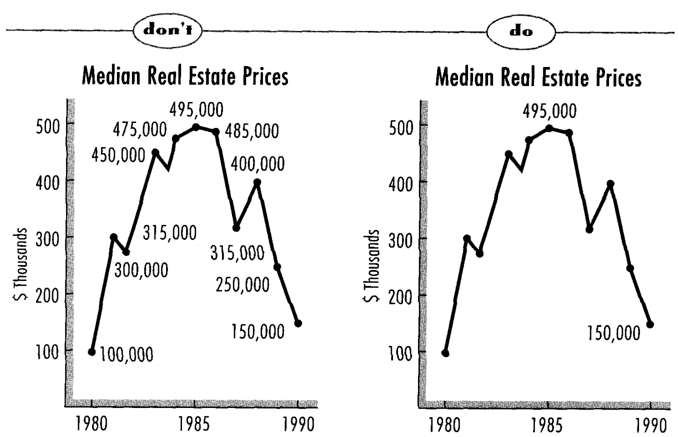

Figure 1.14: The left panel is over-labeled; so many numbers are written out that they distract the eye from the plotted curve. Does one really need to know all the local minima or maxima, or is the overall shape of the curve (zig-zag and irregular) the important thing? The right curve is better because the labelling is more restrained. Labelling one or two key numbers, here the global maxima and minima, may be useful if these are important numbers that the author wants badly to convey to the reader. Often, however, information such as maxima and minima is given in the caption instead.

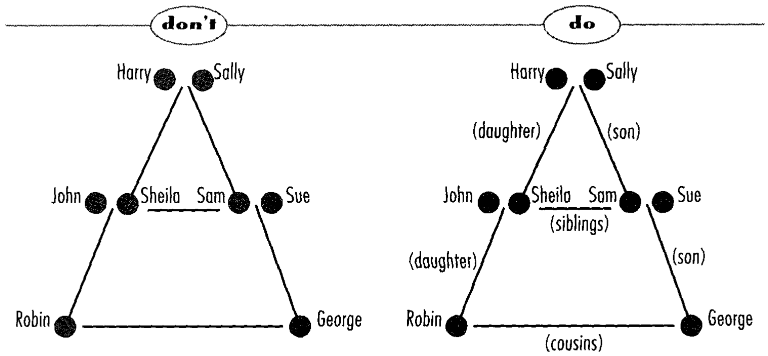

Every important piece of information, whether entity or relationship, should be identified in the display.

Figure 1.15: These two figures, taken from Kosslyn, are identical except that the graph on the right has a few additional labels which greatly improve clarity.

Choosing the proper number of labels can be difficult. If one adds labels with too much enthusiasm, the plotted curves may disappear under a blizzard of text. Fig. 1.14 is over-labeled on the left because the labels merely specify the numbers which are already indicated by the plotted points and the axis label. The reader is likely to spend so much time reading the labels that the curve, which shows the rather unpredictable ups and downs of the real-estate market, may be lost in the perception. Labelling a couple of key values, such as the absolute minimum and maximum as done on the right, may be useful because (i) a couple of labels does not distract from the curve very much and (ii) if the reader wants to know the maximum and minimum, it is easier to read labels than to estimate numbers from the curve and axis labels.

If the graph has too few labels, the information content can be seriously degraded. Fig. 1.15 is an example of inadequate labeling. The "Don't" graph already has eight labels; every black dot (representing a person) is already labelled. Nevertheless, the clarity of the graph can be significantly improved by adding another six labels. Without the additions, the reader has to strain to decode some of the relationships between people.

1.10 Emphasize the Unimportant

Emphasizing the essential elements of a graph can be tricky. Fig. 1.16 shows two versions of the same bar graph. The use of differential shading — white for one bar, cross-hatching for another, solid black for the third — has inadvertently emphasized the black bar and deemphasized the unshaded bar. It is often a good idea to use different shadings to distinguish different graphical elements from one another. Here, however, the differential shading is inappropriate because all three elements are bars, and the eye can distinguish them just fine without shading. The "Do" version shades all three bars equally so as to let the numbers themselves do the talking.

Another example is Fig. 1.17. The theme of the graph is to illustrate the driver's controls. Perhaps the best way to do the graph would be show the driver's compartment only, especially for an audience of engineers. However, the illustrator wanted to place the controls is context, as might be appropriate for a non-technical readership, by showing the whole car. If the publication depicts the controls for several different types of vehicles, such as trucks, minivans, etc., then providing this graphical context would be especially useful. The vehicle type could be specified in the caption, but the less information the reader needs from the

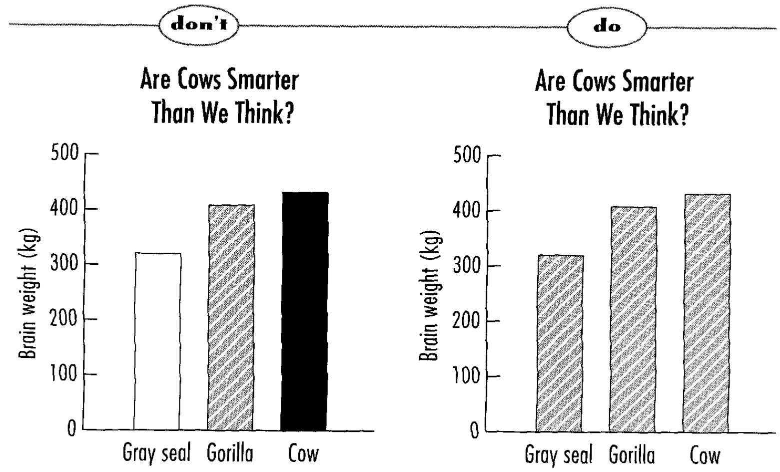

The progression from gray seal to cow looks inappropriately sharper when additional ink makes each bar increasingly salient. If the background were dark, or if the display were on a computer screen, adding more lightness would have the same effect.

Figure 1.16: The solid black color of the rightmost bar in the "Don't" figure makes it practically jump off the page; we have emphasized this datum through choice of color whether we intended this or not. The "Do" figure shades all three bars equally so that the bars are distinguished only by their length, that is, by their numerical value.

caption, the more rapidly she will comprehend the graph.

The left version of the figure, however, is a bad way to provide context because the entire car is depicted in exactly the same style so that nothing is emphasized. The "Do" figure is superior because the detailed line drawing illustrates *only* the driver's compartment.

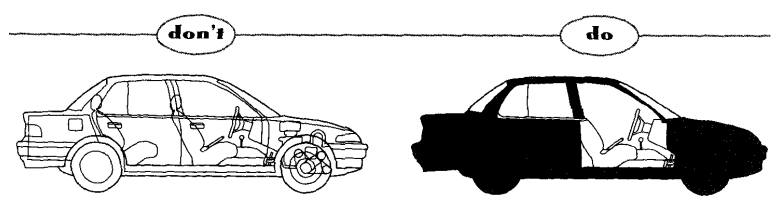

A diagram intended to show the driver's controls is easier to read if the irrelevant components are eliminated and the relevant ones highlighted.

Figure 1.17: The goal of this pair of figures is to depict the driver's controls. The left figure is bad because it shows the rest of the car in the same detail and visual appearance as the important stuff: the driver's controls. The right graph is better because the extraneous details have been shaded out; the eye is drawn to the driver's compartment. At the same time, the shaded areas give a context; the controls are those of an automobile.

1.11 Unnecessary Graphic Novelty

One aim of this course is to encourage imaginative graphics when conventional line plots or contour plots are inadequate. However, unfamiliar species of graphs should only be used from desperate necessity. The reason is that unfamiliar graph types make heavy demands on the reader, who must not only master the data of the figure, but also learn the organization and content of a new species of graph.

Every graphics maven has a weakness for certain graph species. Howard Wainer, whose book is otherwise a masterpiece of visualization technique and common sense, has an ordinate fondness for the so-called Nightingale rose. This is a kind of polar plot which is so named because it was popularized by Florence Nightingale in her successful efforts to dramatize British losses to disease during the Crimean War, and the need to establish military hospitals, nursing corps, and improved sanitation.

Fig. 1.18 compares two versions of the same data. Wainer greatly prefers the top graph because it is a Nightingale rose. However, I find this incomprehensible because I lack Wainer's familiarity with rose charts. The ordinary pie chart is much more successful with me in presenting the information: Most Francophones (native speakers of French) can speak English as a second language, but only a minority of the English-speaking population of Canada can speak French.

The moral of this story is that in choosing graphical species, ask: What is familiar to my audience? What will they find comprehensible?

For example, meteorologists are extremely familiar with contour plots because the usual weather map is this type of graph. One may use contour graphs with reckless abandon in *The Journal of the Atmospheric Sciences*. In contrast, biologists use contour plots much less frequently. Therefore, one should employ a contour plot in article intended for a biology journal only when it is really necessary.

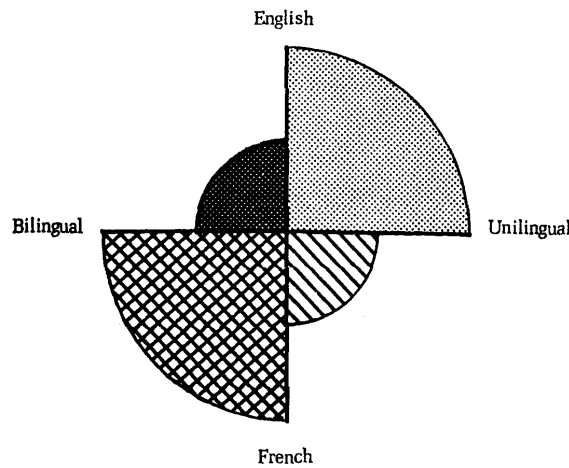

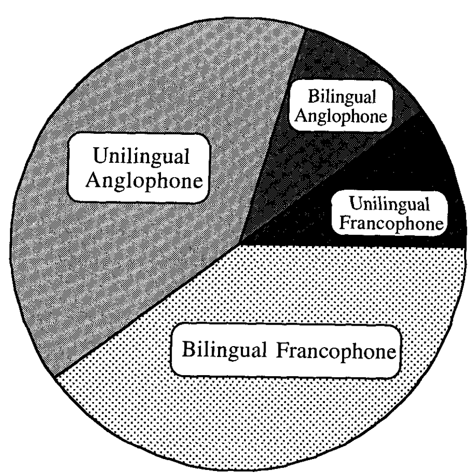

Figure 1.18: Two plots of the same data which show linguistic ability among Canadians as correlated with their primary language. (Both French and English are official languages in Canada.) The top graph is a Nightingale rose chart, prepared by Stephen Fienberg in 1975. Wainer greatly prefers this to the ordinary pie chart (bottom graph) which he prepared as a bad example.