

Visualizing Time

Designing Graphical Representations
for Statistical Data

Statistics and Computing

Series Editors:

J. Chambers

D. Hand

W. Härdle

For further volumes:

<http://www.springer.com/series/3022>

Graham Wills

Visualizing Time

Designing Graphical Representations
for Statistical Data

 Springer

Graham Wills
Hidden Spring Dr. 1128
60540-4112 Naperville, Illinois
USA
graham@spss.com

Series Editors:

J. Chambers
Department of Statistics
Sequoia Hall
390 Serra Mall
Stanford University
Stanford, CA 94305-4065

D. Hand
Department of Mathematics
Imperial College London,
South Kensington Campus
London SW7 2AZ
United Kingdom

W. Härdle
C.A.S.E. Centre for Applied
Statistics and Economics
School of Business and
Economics
Humboldt-Universität zu
Berlin
Unter den Linden 6
10099 Berlin
Germany

ISSN 1431-8784
ISBN 978-0-387-77906-5 e-ISBN 978-0-387-77907-2
DOI 10.1007/978-0-387-77907-2
Springer New York Dordrecht Heidelberg London

Library of Congress Control Number: 2011940977

© Springer Science+Business Media, LLC 2012

All rights reserved. This work may not be translated or copied in whole or in part without the written permission of the publisher (Springer Science+Business Media, LLC, 233 Spring Street, New York, NY 10013, USA), except for brief excerpts in connection with reviews or scholarly analysis. Use in connection with any form of information storage and retrieval, electronic adaptation, computer software, or by similar or dissimilar methodology now known or hereafter developed is forbidden.

The use in this publication of trade names, trademarks, service marks, and similar terms, even if they are not identified as such, is not to be taken as an expression of opinion as to whether or not they are subject to proprietary rights.

Printed on acid-free paper

Springer is part of Springer Science+Business Media (www.springer.com)

Although this book contains tributes to famous men and women who have invented unique and novel visualizations, and to experts who have distilled knowledge and advanced the science of information visualization, this book is dedicated to those people who designed reports, published charts, and created visualizations and were not content to use the defaults but instead took the extra effort to make their work more truthful, more beautiful, and more useful.

Thank you!

Preface

Art or science? Which of these is the right way to think of the field of visualization? This is not an easy question to answer, even for those who have many years of experience in making graphical depictions of data with a view to helping people understand them and take action. When we look at beautiful hand-drawn pictures of data, carefully composed by talented individuals, we are drawn to the artistic side. In some ways those charts are discouraging; their artistic elegance implies that the creation of good visualizations is not an option for most of us. There are books that provide rules and advice on how to draw graphs. Some give general advice, suggesting that such and such is good, but this other is bad. Others give specific advice such as requiring all charts to have a title or all axes to go to zero, but these are often tied to specific visualizations and so are not general enough to qualify as scientific principles. They are valuable for describing existing visualizations, but not general enough to provide guidance for future visualizations. If you are designing something new, advice on a bar chart is not especially helpful.

In this book I want to bridge the gap and not simply give rules and advice but base these on general principles and provide a clear path between them, so that the rules and guidance fall into place naturally, due to knowledge of those principles. In terms of the art/science split, I want to advance the scientific component. There are excellent books describing artistically superb plots; however, my goal is not simply to be descriptive, but to be *prescriptive* – to allow people to start with a goal in mind and design a visualization that fulfills that goal clearly, truthfully, and actionably. Because I have an essentially scientific direction in mind, I will concentrate on reproducibility. A chart that is wonderful for exactly one data set is of little interest. It can be appreciated and enjoyed, but the important question must always be: What can I learn from this graphic that I can apply to other data? With this in mind, the examples in this book have been chosen to be realistic rather than exemplary. I have made a definite attempt not to choose data that make a picture look good, but rather to choose data for which a chart should be applicable. If the result is not perfect, I prefer to present imperfection and explore remedies rather than look for a different data source.

This book is concerned with the graphical representation of *time* data. Time is *special* – it doesn't behave quite like other variables. It has an inherent direction and determines causality. Time can be recorded in many ways: it can be linear or cyclic, categorical or continuous. Even the written format of a piece of time data can be curiously baroque; mixtures of words, numbers, and special symbols make up the time “Monday the 13th of October, 11:45 am.” What other form of data might occur in so obscure a format? All data are recorded at a certain time, and so all data have a time component, even if it has been removed or deemed a priori as uninteresting. This makes time data both unique and universal, so understanding how best to portray them not only is challenging but has wide applicability.

The portrayal of time data is ubiquitous. Any newspaper will feature time-based plots; any company report will show historical data as charts. Even the gas bill for my home invites me to compare a time series of the cost of heating my home against one of average monthly temperature. Because of this generality, I have written this book to cover a range of different users. A visualization expert designing tools for displaying time will find it valuable, but so also should a financier assembling a report in a spreadsheet or a medical researcher trying to display gene sequences using a commercial statistical package. You have data, you have a goal in mind. Now all you need are the tools to graph the data and so achieve the goal. Read on!

Graham Wills

Acknowledgements

The only way to know the effort needed to write a book is to do so yourself, and only authors know the debt of gratitude they owe to others. Warm thanks are due to many people, broadly classified as shown in the diagram below. Any errors and mistakes within the book are entirely my own.




Fig. 1 A modified Venn diagram showing where acknowledgement is due; it shows the major sources but does not include everyone who has helped or influenced my thinking or who has taught me or argued with me over the years. The sum total of those contributions would be large; as this book will indicate, it is often small contributions that make or break any major endeavor

A Note on the Figures

One of the advantages of working for a company, rather than in an academic position, is that you get a different perspective on how visualizations are used. That turns out to be “every way you might ever think of, and then some.” Once your algorithm for generating histogram bin widths has been used by tens of millions of people, on hundreds of thousands of data sets, in over a hundred countries, and any time it didn’t work out for them you got a note to that effect, you start to appreciate the value of robustness not just as a statistical quality but as a *visualization* quality. Now, whenever I see a cool new technique being presented or see an algorithm described that works on one form of data, I immediately think: Will that work in general, or is it too fragile?

Not only is this a pervasive theme in the book, but it has also motivated the design of the figures and, in particular, the lack of postproduction editing. It has been very tempting to take some figures in the book and import them into a graphics editing environment and do a little subtle rearrangement or improvement. Nothing dramatic – just the sort of thing that magazines do to their cover pictures; smooth over imperfections, thin down the bulges, make something already attractive even more so. Several of my reviewers indeed recommended I do exactly that.

Instead, I have chosen to leave the outputs unedited. The charts I have created (all those in the main chapters that have not been attributed to others) are straight from production. I have used the *VizML* language (an XML specification language) to create chart specifications, and the output in the book is exactly what came out of the package. *VizML* is a basic language used in virtually all output of SPSS (now part of IBM) and is available and usable by anyone who owns the major SPSS products. In fact, all the *VizML* figures in this book were generated automatically by a set of Python libraries I wrote that encapsulated common actions for editing *VizML* specifications. As an example, Fig. 3.6 on page 72 was generated by the following Python fragment:

```
Movies = Datafile('Movies.csv')
thin = 'element{stroke-width:0.25px} visualization{margin:5mm}'
T.Histogram.make(x=Movies.Sales, name="MovieSalesA") \
    .remove(Axis,1).addStyle(thin).outputSize("4in", "3in")
```

```
T.LogDistributions.make(X=Movies.Sales, name="MovieSalesB") \  
    .remove(Axis,1).addStyle(thin).outputSize("4in", "3in")
```

I used a standard histogram template (`T.histogram`) for the first part of the figure and a template I designed myself for the second one. For each chart I killed the y axis, added some styles, and set the output size. To generate the figures for the book, I kick off a Python script, wait 10 minutes, and have all my figures.

That is my goal: not to present a set of graphics that are as good as any designer could produce, but instead to produce graphics that a data professional could create to solve a task. Visualization is a tool, and I want a tool that *works*.

Contents

1	History	1
1.1	The Importance of Time	1
1.2	Ancient Visualizations of Time	2
1.2.1	Summary	5
1.3	Playfair	6
1.3.1	Summary	8
1.4	Napoleon’s March	9
1.4.1	A Fortunate Correlation	11
1.4.2	Summary	14
1.5	Comic Books	15
1.5.1	Summary	18
1.6	Further Exploration	20
2	Framework	21
2.1	How to Speak Visualization	21
2.2	Elements	23
2.2.1	Point	24
2.2.2	Line	25
2.2.3	Area	26
2.2.4	Interval	29
2.2.5	Schema	33
2.2.6	Multiple Elements	33
2.3	Statistics	35
2.3.1	Local Smooths	36
2.3.2	Complex Statistics	39
2.4	Aesthetics	41
2.4.1	Categorical and Continuous Aesthetics	44
2.4.2	Combining Aesthetics	46
2.5	Coordinates and Faceting	49
2.5.1	Coordinates	50
2.5.2	Faceting	55

- 2.6 Additional Features: Guides, Interactivity, Styles 56
 - 2.6.1 Guides 56
 - 2.6.2 Interactivity 58
 - 2.6.3 Styles 58
- 2.7 Summary 61
- 2.8 Further Exploration 62
- 3 Designing Visualizations 63**
 - 3.1 Guiding Principles 63
 - 3.1.1 The GQM Methodology 64
 - 3.2 Goals 65
 - 3.2.1 Presenting What Is Important 66
 - 3.2.2 Seeing General Patterns 66
 - 3.2.3 Spotting Unusual Features 68
 - 3.3 Questions 69
 - 3.3.1 One Variable: Unusual Values 70
 - 3.3.2 One Variable: Showing Distribution 71
 - 3.3.3 Two Variables: Showing Relationships and
Unusual Values 73
 - 3.3.4 Multiple Variables: Conditional Relationships,
Groups, and Unusual Relationships 78
 - 3.3.5 Multiple Variables: Showing Models 80
 - 3.4 Mappings 85
 - 3.5 Systems of Visualizations 86
 - 3.5.1 Narrative Structure 88
 - 3.5.2 Consistency 89
 - 3.5.3 Stereotypes 89
 - 3.6 Top-Down Versus Bottom-Up 89
 - 3.7 Summary 91
 - 3.8 Further Exploration 93
- 4 Types of Data 95**
 - 4.1 Four-Minute Mile, Day of the Week, Bottom of the Ninth 95
 - 4.1.1 Scales of Measurement 96
 - 4.1.2 Form Follows Function 97
 - 4.2 Events and Intervals 99
 - 4.3 Regular and Irregular Data 101
 - 4.4 Date and Time Formats 102
 - 4.5 Summary 103
 - 4.6 Further Exploration 104
- 5 Time as a Coordinate 105**
 - 5.1 Put It on the Horizontal Axis 105
 - 5.2 Event Occurrences 108
 - 5.2.1 Many Events 114

- 5.3 Regular Categorical Sequences 116
 - 5.3.1 Patterns in Sequences 118
- 5.4 Summary 121
- 5.5 Further Exploration 121
- 6 Coordinate Systems, Transformations, Faceting, and Axes** 123
 - 6.1 Time Series 123
 - 6.1.1 Aspect Ratio 124
 - 6.2 Coordinate Transformations 127
 - 6.3 Axes 129
 - 6.3.1 Drawing Time Axes 132
 - 6.3.2 Formatting Time Ticks 135
 - 6.4 Faceting 136
 - 6.4.1 Faceting by Time 138
 - 6.4.2 Faceting Complexity 140
 - 6.4.3 Time Within a Faceting 144
 - 6.4.4 Faceting When Data Are Not Categorical 148
 - 6.5 Summary 150
 - 6.6 Further Exploration 150
- 7 Aesthetics** 151
 - 7.1 Time as a Main Aesthetic 151
 - 7.1.1 Representing Counts 152
 - 7.1.2 Summarizing and Splitting Aesthetics 154
 - 7.2 Specific Aesthetics 156
 - 7.2.1 Coloring by Time 156
 - 7.2.2 Sizing by Time 157
 - 7.2.3 Shaping by Time 158
 - 7.2.4 Other Aesthetics and Time 161
 - 7.3 Time as a Secondary Aesthetic 161
 - 7.4 Summary 165
 - 7.5 Further Exploration 166
- 8 Transformations** 169
 - 8.1 Distortions of Time 169
 - 8.2 Time as Frequency 172
 - 8.3 Converting Between Categorical and Continuous 176
 - 8.3.1 From Categories to Continuous 176
 - 8.3.2 From Continuous to Categories 178
 - 8.4 Summary 179
 - 8.5 Further Exploration 180
- 9 Interactivity** 181
 - 9.1 A Framework for Interactivity 181
 - 9.1.1 Display Pipeline 182
 - 9.2 Modifying Parameters 184
 - 9.2.1 Modifying Element Parameters 184

- 9.2.2 Modifying Aesthetic Parameters 186
- 9.2.3 Modifying Coordinate Parameters 188
- 9.2.4 Modifying Statistic Parameters 191
- 9.2.5 Modifying Scale Parameters 194
- 9.2.6 Modifying Facet Parameters 195
- 9.2.7 Modifying Transform Parameters 196
- 9.3 Interacting via the Data 196
 - 9.3.1 Brushing and Linking 198
 - 9.3.2 Drill-down 204
 - 9.3.3 Summary 205
- 9.4 Further Exploration 206
- 10 Topics In Time** 207
 - 10.1 Large Data Sets 207
 - 10.1.1 Aggregation 208
 - 10.1.2 Augmenting Traditional Displays 213
 - 10.2 Time Lines and Linked Events 217
 - 10.2.1 Linked Events 218
 - 10.2.2 Timelines 220
 - 10.3 Summary 224
 - 10.4 Further Exploration 224
- 11 Gallery of Figures** 227
 - 11.1 Chart Complexity 227
 - 11.1.1 Complexity Study Procedure 228
 - 11.1.2 Initial Analysis 230
 - 11.1.3 Model Fitting 233
 - 11.1.4 Discussion 234
 - 11.1.5 Application to the Figures in This Book 235
 - 11.2 The Gallery 236
- References** 247
- Index** 253

Chapter 1

History

HISTORY, n. An account, mostly false, of events, mostly unimportant, which are brought about by rulers, mostly knaves, and soldiers, mostly fools.

—Ambrose Bierce, The Devil’s Dictionary (1881–1886)

1.1 The Importance of Time

Measurement and recording are part of the scientific approach, no less for time than for any other form of data. (Chapter 4 discusses details of how we define data and ways to describe time in detail.) The history of how those measurements have been made – even the words that have been used for units of time – makes for a fascinating study. For example, Brian Hayes describes the astronomical clock of the Strasbourg Cathedral in his essay “Clock of Ages” (reprinted in [53]). This is a mechanical clock built over 160 years ago that measures in addition to the usual values:

- Sidereal time (measured by the Earth’s rotation),
- Local solar time (where noon is when the sun is highest),
- Local lunar time,
- A counter for the years,
- The current date, including leap year calculations, and
- The dates of movable church feasts, including Easter (a complex calculation, only standardized in 1582 by Luigi Lilio, as described in [81]).

This entire system works with gears, from fast-turning ones to a small gear that turns only once every 2500 years. The Strasbourg clock is a tribute to the importance of being able to calculate and measure time in multiple ways.

Although vital to the smooth running of civilization, calendars and the measurement of time have not always been standardized. Holford-Strevens [56] introduces and describes the major different calendar systems that have been employed across

the ages; the major division into lunar and solar calendars and how various systems attempted to resolve the differences between them. He gives the histories of the Babylonian, Egyptian, Jewish, Roman, Julian, and Gregorian systems and their evolutions. Today, virtually all countries use either the Gregorian or Revised Julian calendars, which are not due to differ until 2800,¹ so dates under modern notation can safely be assumed to be comparable.

The *accurate* measurement of time was another historical goal with important applications. In the eighteenth century, accurate measurement of time was necessary to be able to calculate a ship's longitude with precision. This was considered important enough that million-dollar prizes (in modern-day dollars) were awarded and ship owners were prepared to pay amounts up to a quarter of their ship's value in order to buy a highly accurate clock. How well did they manage? In 1761 John Harrison's H4 clock made the trip by sea from England to Jamaica and lost only five seconds on the voyage [103]. Pretty impressive.

Although the history of measurement and recording of time is a fascinating study, in this book we are concerned primarily with the display of time. In line with our aim of making *informative* visualizations, this chapter will take a look at a set of historical visualizations with the goal of learning what made them useful and what lessons we can apply to our visualizations.

1.2 Ancient Visualizations of Time

Figures 1.1 and 1.2 show a pair of petroglyphs (also known as “rock carvings” or “rock engravings”) that were created probably a couple of thousand years ago near the area of the Grand Canyon in Arizona, USA. They are not particularly notable or exceptionable specimens; in fact, they portray a fairly common pattern of markings that experts believe may represent early calendars (or, at the least, records of past time periods). The earliest forms of data recording consist of notches or scratches representing counts of various qualities, and so the representation of time as a sequence of counts of units past a known event is a natural first step in recording time.

Early historical documents often refer to dates in a similar manner, for example:

- “In the four hundred and eightieth year after the Israelites had come out of Egypt, in the fourth year of Solomon's reign over Israel, in the month of Ziv, the second

¹2800 is a leap year in the Gregorian calendar, but not in the Revised Julian calendar. Under the latter system, leap years for centuries only occur when the year divided by 900 has remainder of 200 or 600. This system is more accurate than the Gregorian calendar in the (very) long run, requiring less adjustment to match the actual year length.



Fig. 1.1 Petroglyphs near the Grand Canyon. These cave carvings date from as far back as 500 B.C. and were found at a site believed to be used as a temporary hunting camp. Although this particular image looks somewhat like the skeleton of a fish, there are many similar diagrams that look more tabular. These images are often thought to represent a “hunting calendar”



Fig. 1.2 This figure shows another “calendarlike” petroglyph found near the Grand Canyon. Beside it is an iconic representation of what appears to be a volcano. If so, a natural interpretation is that these boxlike divisions indicate periods of time since a significant event



Fig. 1.3 Stonehenge may be the most famous ancient monument in Britain. For over 850 years, researchers have searched for explanations of its purpose and why and by whom it was built. Photograph ©Mark Radford 2009, used with permission [89]

month, he began to build the temple” – an account of the building of Solomon’s Temple, taken from the Book of Kings.²

- The traditional Japanese calendar consists of eras based on the reign of the emperor. This calendar is still used in several official contexts – the Japanese patent office only changed from this dating system in 2002, or “Heisei 14” as that year is called in the traditional system.
- The terms “D-Day” and “H-Hour” are used in military operations to designate the time an operation begins. Times are measured relative to this event so that “H-3” means 3 hours before the operation commences and “D+7” a week after its commencement. This convention has been used since at least 1918.

From a small, common example of prehistoric creation we leap to one of the most famous sites in the world: a collection of large stones, carefully arranged in a set of concentric rings – Stonehenge. Evidence makes clear that it was built in successive stages over a period of several hundred years, and considerable effort was put into its creation, so there must have been a strong reason (or reasons) to have it built.

²The phrase “in the *N*th year” of someone’s reign gives historians some trouble as even within a single source it may be used in different ways. If a king starts ruling in a given year, sometimes that is described as “the first year,” but other recorders may use that term only for the first *full* year, using a phrase like “the year when so and so became king” for the partial year of ascension.

The stones have been placed with a precision of up to a few centimeters, leading to several theories for this concern for accuracy. One theory, dominant for the last 200 years, is that the site is a giant astronomical observatory or calendar, designed to allow its users to observe or measure seasons, lunar cycles, and similar time-based occurrences.

Despite the lack of strong evidence to back this theory up,³ the theory clearly has been compelling. It is both a cautionary tale in trying to guess the motives of other people and a statement about the importance that we ourselves give to time. In essence, we think that if a prehistoric people were to go to such great lengths to make such a structure, it must be for a very important reason – and what is more important than time?

1.2.1 Summary

The major lesson we can learn from these ancient representations is that *time* is important. Many representations of time exist in the ancient world, and even when we are unsure about an artifact's importance, we can see that some features are time related. Since our goal is to make useful visualizations, it is important to remember that these visualizations were created for a purpose. We should do the same, creating visualizations that serve a purpose. Chapter 3 describes a design methodology that accomplishes this.

These depictions of time show how time is often *divided* into discrete units; calendars divide time into days, months, and years – although we speak of time flowing, in practice we often divide it up into chunks. In Chap. 8 we address this technique and use “chunked time” throughout the book in an informal way. Relatedly, time can be recorded as an ordinal variable – a set of ordered categories – that corresponds to the way calendars are laid out. Chapter 4 deals with the various ways in which we can consider time data. One such way is to consider *durations*, which can be defined as the time elapsed since fixed events.

In Chap. 6 we discuss *axes* – guides that show time dimensions and allow us to identify what time a region of the chart represents. The main axis of Stonehenge is a literal axis being used to represent the time dimension. For these examples, measurement might be the most important feature, but for all time visualizations we must be able to answer the question: When did this interesting feature happen? Guides have a particular significance for time data, as measurements are often made relative to known events. In designing visualizations, adding known reference times is highly advisable.

³Johnson [64] gives strong arguments that the only real calendric detail is the alignment of the major axis along the line of the summer–winter solstices and that an alternative explanation should be sought for other features.

Finally, the use of a polar coordinate system is seen in many ancient depictions of time. The sun travels around the world, the world rotates around the sun, and so many of our systems for measuring time are cyclical in nature; sundials must use a polar system for laying out the hours, mechanical clocks have dials, and so on. Not only is this needed for mechanical reasons, but using a circular coordinate system allows us to visualize data that may have a cyclical nature. Many phenomena have obvious dependencies on natural cycles (weather patterns, for example), and so a final lesson to learn is that we should not restrict ourselves to portraying time only as a linear dimension.

1.3 Playfair

William Playfair has been credited with the invention of the time series chart and in 1786 published a book, *Commercial and Political Atlas* ([87], republished as [123]), that introduced a number of new visualizations to the public, including many charts showing time data. Playfair was a strong character and is the focus of the first part of Wainer’s book [122] on graphic revelations. As an example of what Playfair was like, I offer up the title of Chap. 3 of this book: “William Playfair: A Daring Worthless Fellow.” His worth in the realm of statistical graphics, however, is unquestioned.

Playfair made a number of strong statements in his introduction in favor of graphical representations, of which the following is representative:

Figures and letters may express with accuracy, but they can never represent either number or space. A map of the river Thames, or of a large town, expressed in figures, would give but a very imperfect notion of either, though they might be perfectly exact in every dimension.

In an affair of such consequence, as the actual trade of a country, it is of much importance to render our conceptions as clear, distinct, and easily acquired, as possible . . . A man who has carefully investigated a printed table finds, when done, that he has only a very faint and partial idea of what he has read.

Figure 1.4 is a reproduction of one of Playfair’s figures. The Annenberg Rare Book and Manuscript Library at the University of Pennsylvania provided the original and assisted in recovering a high-quality reproduction of the original for [123], which has been reproduced here. The figure is intended, as stated by Playfair, to show the balance of trade clearly and succinctly. Playfair’s charts mark a major milestone in the history of data visualization, and for time visualization in particular they can be regarded as the first quantitative charts of time.

Playfair’s charts show several elements in the same graphic. Here we have lines for both imports and exports and the area element that is defined by their difference. Further, the lines are given unique colors to distinguish them, and the areas are colored depending on the sign of the difference in values. These are examples of the use of aesthetics, which take a basic chart and add information without changing




Fig. 1.4 This figure shows two graphs: the balance of trade between England and Denmark/Norway and the balance of trade between England and Sweden. In both graphs there are two lines, one for imports and one for exports (from England's point of view). The area between the lines has been colored to indicate the direction of the balance. The original of this figure is in [87], reprinted in [123]

the structure⁴ and are described in detail in Chap. 7. Playfair uses multiple elements in other charts, including a bar and line chart with different vertical dimensions. He also shows stacked areas and uses patterns and colors for aesthetics on bar charts as well as areas. And it would be remiss not to mention that he also invented the bar chart – the first depiction of data in a purely abstract coordinate system without either dimension representing space or time.

A second innovation is the use of an *area* to represent the trade amounts. As we will see in Chap. 2, an area is a good choice when the values being represented on the vertical dimension are additive. We can add money, so an area makes sense. Further, the size of the slice of area between the two vertical lines represents the sum of all exports in the time period between those lines, so we can make comparisons of total area.

The charts are also *aligned*; the time values are aligned vertically between the two figures and their vertical axes span the same range, making it possible to do area comparisons accurately. This technique, here applied to two charts, can be applied to many charts using a technique termed here *faceting*, elsewhere called “paneling,” “small multiples,” and “trellis.” In Fig. 1.4 we can directly compare the

⁴In this book the word *aesthetic* is defined as a mapping from data to graphic attributes of a chart element, such as color, size, pattern, opacity, etc.

areas between the two charts and also compare the areas “in favor” and “against” in the upper chart and make correct judgments based on them.⁵ Faceting is studied in detail in Chap. 6.

As if this were not enough, Playfair also gets the details for his *axes* right. Gridlines are clear, but not prominent, with major gridlines more heavily weighted than minor ones. His labeling of the tick marks drop unneeded precision and are placed at appropriate tick locations. In Chap. 6 some of these details are discussed, and a brief look at the default output of most charting packages will indicate that the creation of a good axis is not an easy task.

We do not expect to see any interactivity in hand-drawn charts from the 1700s, but Playfair does lay the groundwork for several key interactive techniques. Consider Fig. 1.5 (plate #1 in [87]). This is an *overview chart*, summarizing the trade between England and all its trading partners. One interesting point to note is that the chart shows a focus on the years from 1760 onward. The lines for both imports and exports are less smooth, leading us to believe that there are more data available in this period, and Playfair has given us minor gridlines only in this time period – a static technique that focuses attention and helps us concentrate on this part of the chart. In Chap. 9 this technique is expanded to a fully interactive context; the use of *distortion* techniques that selectively magnify part of the axis is demonstrated. Playfair paves the way for this very recent interactive technique with the selective use of guides to focus our attention on this area.

1.3.1 Summary

Playfair establishes the basic structure of time series charts with these graphics. Most important is the basic structure – a *two-dimensional* coordinate system with time running horizontally and a quantitative *y* dimension. This basic structure, *time* × *quantity*, forms the standard framework for a time series chart and is used extensively throughout this book. His use of multiple elements, aligned axes, and axis drawing techniques is also an important design development, which we will explore in the course of the book.

A final point to make is on the presentation of the charts themselves. Playfair structures the book by giving us an overview chart as the first figure and then moving on to show details of individual countries’ trade in subsequent chapters. This is another important interaction technique – start with an overview of the data, then filter by a relevant variable (in this case, the country being traded with), and show details on those drill-down charts only on request. In this static medium, the details

⁵Up to a point. Wainer and Spense indicate in [123] that Playfair may not have been as accurate with his depictions of the data as he could have been. The method is correct, but the execution seems to be a little haphazard.




Fig. 1.5 This area graph shows the overall balance of trade between England and all countries with which England traded over the period 1700 through 1800. Most notable in the graph is the single dip “into the red” that occurs between 1780 and 1782. At this time England was heavily involved in fighting the Revolutionary War in North America, as well as a couple of conflicts in India; the first Anglo-Maratha War, fought between the British East India Company and the Maratha Empire in India; and the second Anglo-Mysore War, started opportunistically by Great Britain on the grounds that since they were already at war with France (who joined the Revolutionary War on the US side), they might as well have a go at kicking them out of India, which was a far more important possession in terms of economic value. Spain joined the war (as an ally of France) and the Dutch also joined in, with Britain declaring war on them in 1780 to start the Fourth Anglo-Dutch War, which worked out very well for Britain, as they took a number of Dutch colonies. This graph shows the effect of all these wars on the balance of trade; the next important step would be to drill down and see how trade with each of the opposing countries was affected. The original of this figure is in [87], reprinted in [123]

are the body of text in the chapter. A famous basis for visualization systems, the “Visual Information-Seeking Mantra” of Shneiderman [98], states:

Overview first, then zoom and filter, and finally, details on demand.

Playfair, 210 years earlier, used this basic organization principle for his exploration of trade balances, and this approach is described further in Chap. 3.

1.4 Napoleon’s March

In this section we will look at a chart that is widely regarded as a classic of information design, Minard’s map of *Napoleon’s march on Russia*, reproduced as Fig. 1.6. Charles Joseph Minard was an engineer who created many innovative

- *Geography*: The main chart is a map and depicts a selection of rivers, cities, and battles. The path showing the advance and retreat is located using map coordinates.
- *Path*: The path of the army is drawn directly on the map and is color coded by direction: gold heading into Russia, black retreating out.
- *Count*: The number of soldiers is represented by the width of the path, from 480,000 at the beginning to 10,000 at the end.
- *Temperature*: For the retreat only, the air temperature is given at selected points along the journey, represented by a line chart at the bottom, with thin lines linking the two charts.
- *Time*: Time runs right to left as the army retreats, and the line chart at the bottom gives dates at several locations.

The geographic information is essentially a guide in this display – a background that gives context and helps us understand the main effects – for example, showing that when rivers are crossed in cold weather, people die. The rest of the data can be thought of as a single table with each row being a location along the march, and the variables being *latitude*, *longitude*, *count*, *temperature*, and *time*.

This figure has a lot to live up to. Tufte [111] declared that it “may well be the best statistical graphic ever drawn,” and it has received similar accolades by other authors. In this book our goal is not so much to analyze charts for their unique beauty and individual utility, but to understand how useful the chart is and what general lessons we can learn from it. The original was hand-created, but numerous versions of it have been created, some in an attempt to remain true to the original, some more radical reworkings. Several are detailed in [45] and collected on the Web at [44]. The version in Fig. 1.7 was created using the *VizML* visualization system used throughout this book. Only the cities have been retained from the geography, and more prominence has been given to the temperature chart, since our goal is to emphasize the time component more than Minard originally did.

1.4.1 A Fortunate Correlation

In statistical analysis, if we are trying to predict a target variable, and two of the variables we are using as predictors are highly correlated, then this can often cause problems. In statistical graphics, however, it may turn out to be a good thing. I would argue that Minard's chart is a great chart because it is based on a highly successful and simple time-based chart that has a fortunate correlation with geography.

Refer back to Playfair's figures on balances of trade (Figs. 1.4 on page 7 and 1.5 on page 9). They are time series charts that display values of a variable using an area, where that variable is a measure of the size of something – a display very similar to the paths in Fig. 1.6. Minard's chart works because it shows a time series of army size, with the *time* variable replaced by the highly correlated *longitude* variable. That

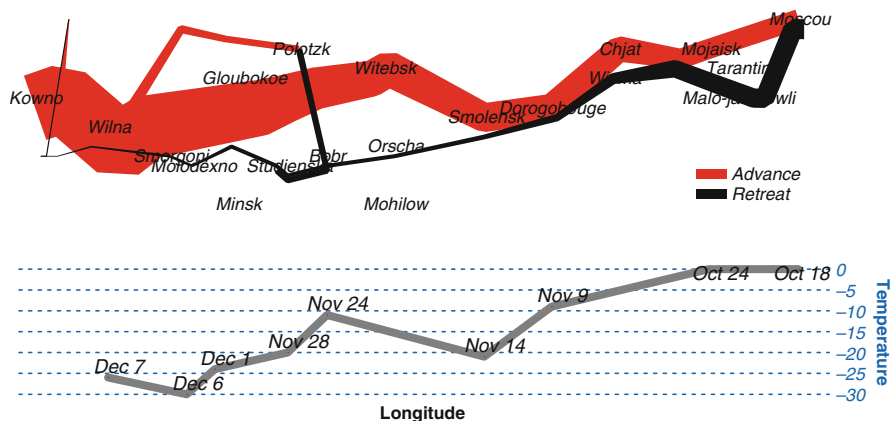


Fig. 1.7 Napoleon’s March on Russia. This figure is a reworking of Fig. 1.6, simplified to allow it to fit into a smaller space and remain legible. It also has allocated equal space to both the map and the temperature graphs, making them equal parts of the story

allows the time series of army size to be naturally integrated into the map display without needing any explicit linking. This has a few consequences:

- One potentially interesting factor – how fast the army advanced – is not available to us. The temperature graph that is linked to the main chart helps us infer the pace of the retreat by linking locations to a line chart with time annotations. We would need another linked view for the advance.
- The area of the path and the rate of diminution show the army size per mile of ground covered, which is a somewhat unusual and less informative quantity than would have been shown using undistorted time for the horizontal axis. In that case we would have seen the rate of change per day. However, it could be argued that since the goal was to get to Moscow at all, rather than in a given time period, base units of distance traveled make just as much sense.
- Since for the complete march the longitude does not correlate well with time, Minard had to split the chart in two, one path for the advance and one for the retreat. Each is dealt with separately.
- For the retreat, which is the most recognizable “time series” chart on the bottom, time runs right to left, which is disorienting as even in cultures where text might read in other directions, time is almost always shown running left to right. A time series shown for the advance would be oriented more naturally.

Figures 1.8 and 1.9 show the result of directly plotting army size and air temperature. In the first figure we use the (reversed) longitude for the horizontal dimension, and in the second figure we show the time. Not only can the drop in




Fig. 1.8 Napoleon's retreat. An area element shows the size of the main component of the army (survivors) and a line shows the temperature. These have different y dimensions and are color coded to match the corresponding axis. The x dimension indicates the position of the army as a longitude. The x dimension has been reversed so that time runs *left to right*




Fig. 1.9 Napoleon's retreat. An area element shows the size of the main component of the army (survivors) and a line shows the temperature. These have different y dimensions and are color coded to match the corresponding axis. The x dimension indicates the date

troop numbers and temperature be seen, but by comparing the two figures we see that there is not much distortion in the way that time and distance are conflated in Minard's chart.

At an overview level, at least for the main component of the army in the retreat direction, the two charts can be thought of as the same. These additional views give us confidence in Minard's technique, so that we can make assumptions about the horizontal dimension as a time dimension and know we are not being misled.