CHAPTER 2

Rearrangement and Interaction

2.1 'A graphic . . . is a moment in the process of decision making'

Anyone who has seen, and especially used, a highly responsive interactive visualization tool will be struck by two features. First, that a mere rearrangement of how the data is displayed can lead to a surprising degree of additional insight into that data. Second, that the very property of interactivity can considerably enhance that tool's effectiveness, especially if the computer's response follows a user's action virtually immediately, say within a fraction of a second.

To illustrate the concept of rearrangement we begin with a simple example in which ten crops (for example, wheat, rice, beans . . .) have been subjected to seven treatments (for example, insecticide, . . .), and the improvement or otherwise noted. If improvement is coded black and degradation white, the result of the experiments might appear as in Figure 2.1(a), where crops are numbered from 1 to 10 and treatments designated by A to G. Note that we are here concerned with categorical data: the ordering has been imposed merely for convenience of reference.

It is probably fair to say that the likelihood of anyone gaining immediate insight from Figure 2.1(a) is negligible. Now, however, imagine the diagram to be cut into rows as shown in Figure 2.1(b) and for these rows to be rearranged in an order which assigns priority to placing black squares as high as possible, first in column A, then in column B, and so on (Figure 2.1(c)). In other words, rows with a black entry in column A are ordered first, followed by rows with such an entry in column *B*, and so on. In Figure 2.1(c) we see some sort of pattern beginning to emerge.

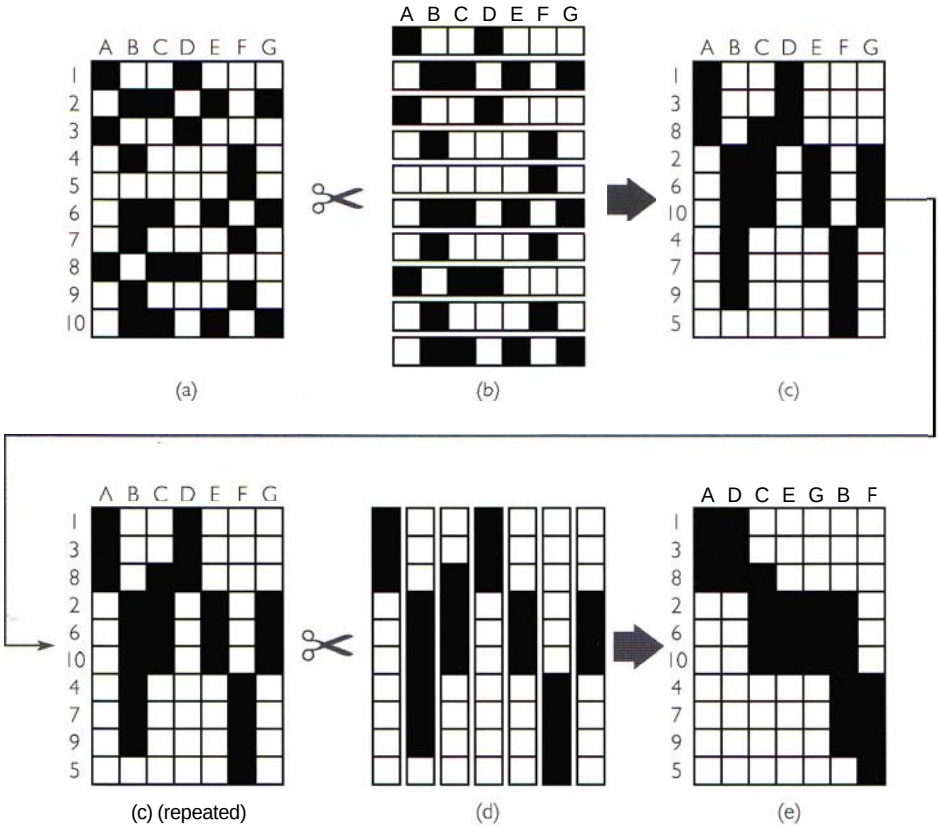

FIGURE 2.1 Kearrangement of data concerning the treatment of crops

If the same process is now repeated, starting with Figure 2.1(c) but with cuts (Figure 2.1(d)) to allow columns to be similarly rearranged, the result has the appearance of Figure 2.1(e). The same data is involved, but its presentation has changed: one can clearly perceive a pattern, showing that certain groups of treatments are appropriate to certain groupings of crops.'

It is at this point that the question 'OK, but what now?' may be asked. The most useful – as well as the most perceptive – answer was provided by Bertin (1981) who commented that:

> *A graphic is no longer 'drawn' once and for all: it is 'constructed' and reconstructed (manipulated) until all the relationships which lie within it have been perceived . . . a graphic is never an end in itself: it is a moment in the process of decision making.*

We have just seen, in Figure 2.1, one such 'reconstruction'. That there will typically be more is confirmed by Cleveland's (1985) comment that:

> *Graphing data needs to be iterative because we often do not know what to expect of the data; a graph can help discover unknown aspects of the data, and once the unknown is known, we frequently find ourselves formulating new questions about the data.*

---

[1] I'm most grateful to Bob Waddington for permission to use the material shown in Figure 2.1.

Thus, following inspection of Figure 2.1(e), a user – and especially a user in possession of domain knowledge (for example, a crop geneticist) – may decide upon further rearrangement or the addition of possibly relevant data. Here is the essence – and certainly quite often the excitement – of interactive visualization: a lively, iterative examination and interpretation of graphically presented data, in which rearrangement is an important part and the outcome is usually not predictable.[2]

2.2 The Table Lens

Our second example is drawn from baseball data, in particular that data contained in the *Baseball Encyclopedia Update 1997*. The data concerns **323** players, each associated with a row in a table; the columns of that table refer to attributes such as salary, position (in the baseball field) and hits. The Table Lens (Rao and Card, 1994) allows each column to be viewed as a histogram (Figure **2.2),** and there is a facility whereby a name can be associated with a particular row of the histograms.

Clearly there are many ways in which the rows can be ordered. Alphabetic ordering, for example, would enable a particular player to be located with ease. A more interesting reordering however, is shown in Figure 2.2. Here, the rows have first been reordered according to the (ordinal) attribute 'position' and then according to 'number of hits'. It therefore shows the number of hits achieved by each batting position. By examining the 'position' and 'hits'

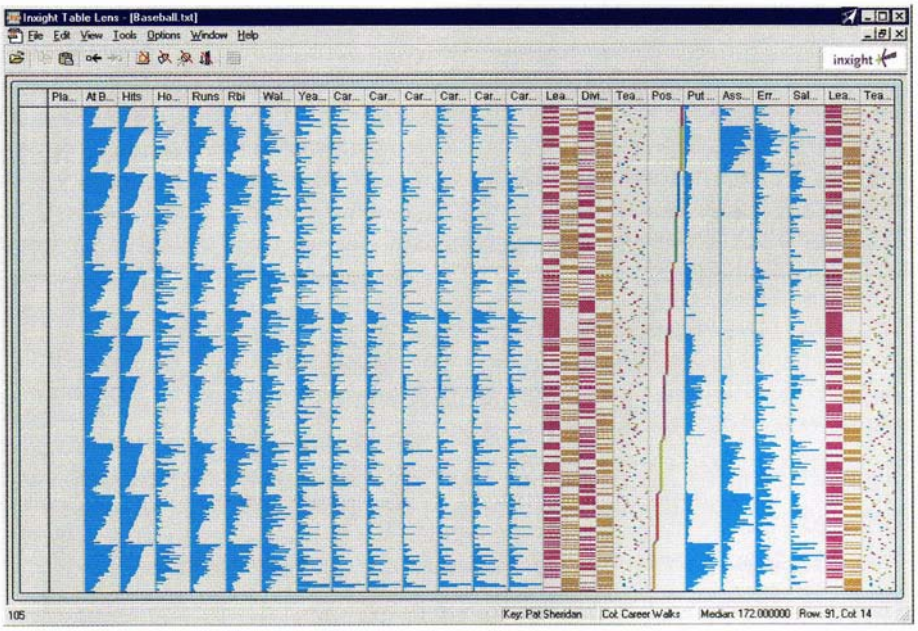

FIGURE 2.2
A Table Lens displaying baseball data

Source: Inxight Software

[2] Rearrangement, of course, is not always beneficial, as when a well-intentioned cleaner 'tidies up' one's desk (Malone, 1983)!

columns one can see that different batsmen in the same position are characterized by a range of batting skills. The Table Lens possesses the advantage, as far as ease of use is concerned, that it is based on a table, a concept with which most people are familiar.

We shall revisit the Table Lens again in Chapter 7 to identify additional beneficial features: in the context of rearrangement, however, an important question to be answered now is how the rows are reordered. The reordering action chosen by Rao and Card (1994) is simple, and illustrated in Figure 2.3(a): a movement of the mouse down the column which is to be ordered. The movement need not be precisely vertical. Another example of how a gesture with the mouse can lead to rearrangement is shown in Figure 2.3(b): a 'tick' with the mouse in the column associated with attribute B followed by a similar tick for attribute A triggers the generation of a new column containing a histogram of the values of *B/A*.

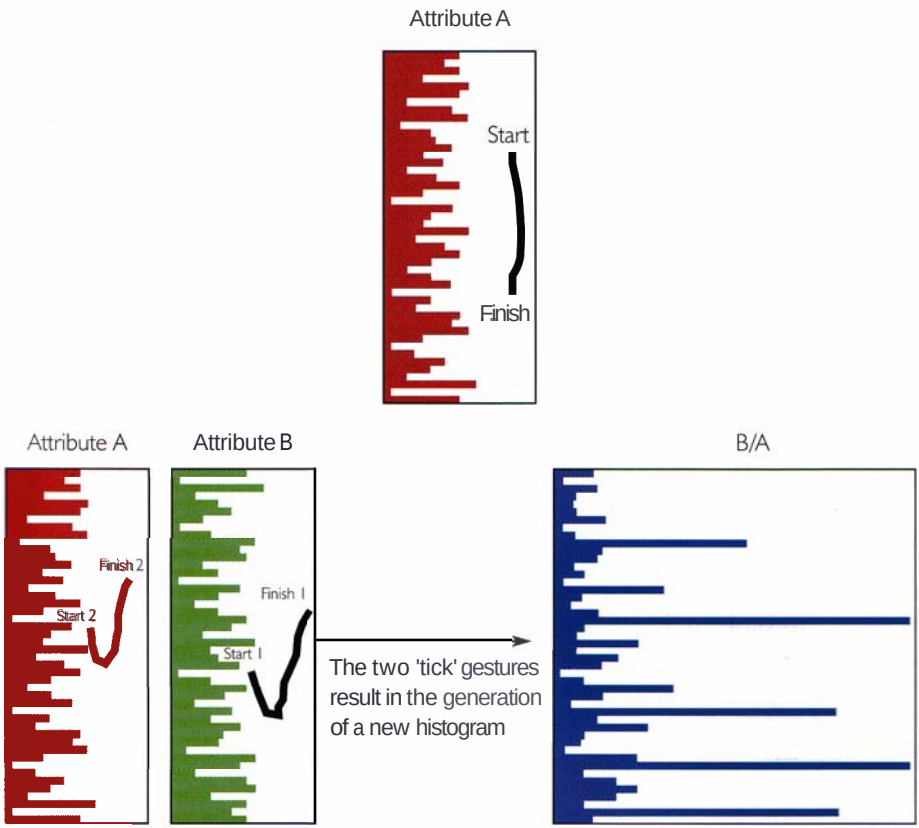

FIGURE 2.3 (a) Reordering of the entries in the Table Lens (b) Calculation of a new column in the Table Lens

Affordances

The example of the Table Lens illustrates an important aspect of rearrangement to which the designer of an information visualization tool must give attention: it is called **affordance.** The concept is perhaps most effectively introduced by a

physical example taken from Norman (1988). You are approaching a door through which you eventually want to pass. The door, and the manner in which it is secured to the wall, permits opening by pushing it from its 'closed' position. We say that the door **affords** (or allows, or is for) opening by pushing. On approaching that door (Figure 2.4) you observe a flat plate fixed to it at waist height on the 'non-hinge' side, and possibly some sticky fingermarks on its otherwise polished surface. You deduce – correctly in this case – that the door is meant to be pushed open: you therefore push on the plate, whereupon the door opens and you pass through. Here, there is a perceived affordance, triggered by the sight of the plate and the fingermarks, that is identical with the actual affordance. Note that the affordance we discuss is neither the door nor the plate: it is a property of the door ('the door affords opening by pushing').

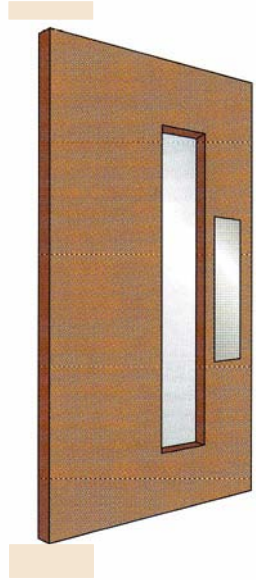

FIGURE 2.4 The door is perceived, correctly, to afford opening by pushing

You now have to proceed through a second door, and observe (Figure 2.5) that it has a handle, again at about waist height. Your interpretation is that you are meant to pull the handle towards you, thereby opening the door. Though a reasonable interpretation, it happens to be incorrect – the door does not move. You try pushing it instead, again to no avail. After some exploration you then discover that the door opens by sliding (it affords opening by sliding), and that the correct action is to 'turn' the handle and apply a horizontal force. In this situation the perceived affordance (that the door affords opening by pulling) is not the same as the actual affordance (the door affords opening by sliding). Clearly, designers of doors, and particularly their mechanisms and associated devices such as handles, will try to ensure that perceived and actual affordances are identical.[3] It is important to note from these simple examples that experience and convention influence one's perception of affordance.

[3] Those who are required quite frequently to stay in hotels may well wonder whether designers of shower controls are familiar with the difference between perceived and actual affordances.

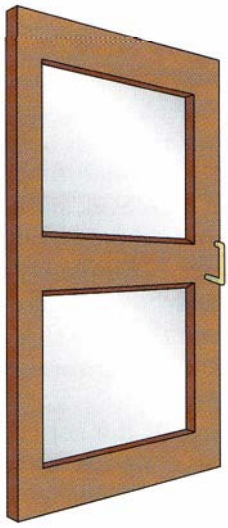

FIGURE 2.5 The door is perceived, incorrectly, to afford opening by pulling

In designing interactive displays to facilitate a beneficial rearrangement of the visual presentation of data, the same design objective is present: the designer should ensure, as far as possible, that perceived and actual affordances are identical. But that is not always easy to achieve. As an illustration **we** choose (Figure 2.6) the familiar scroll controls associated with a word-processing application. A 'mouse-down' on the control **A** (that is, anywhere within the rectangle containing the triangular arrow) will, for example, cause the contents of the **display** to scroll downwards until 'mouse-up' occurs. Thus, the *actual* affordance of scrolling must be deduced by the user. In this case the sensitive control icon **A** is chosen by the interaction designer with the expectation that the perceived affordance of scrolling will be identical to the actual affordance of scrolling. Nevertheless, the interaction designer will be aware that an understanding of the icon's function depends upon many factors including convention and how suggestive the icon is.

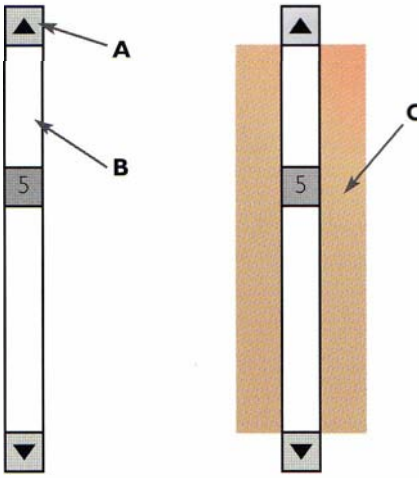

FIGURE 2.6 Scroll controls familiar from a word-processing application

In some cases a perceived affordance is difficult to formulate, sometimes because the relevant control may be far from obvious. An example is provided by area **B**, the rectangular area extending from the upper edge of the page indicator to the lower edge of the 'scroll-down' control. At a first glance it is not obvious if a mouse-click on this area will achieve anything at all. However, through exploration or training it is found that a mouse-click in **B** causes flipping to the extent of one page.[4] Moreover, prolonged mouse-down in **B** causes page flipping to continue at a convenient rate until terminated by a mouse-up.

Other controls may not even be completely externalized: to ease the user's task and avoid demanding precisely vertical movement of the cursor within the relatively thin scroll-bar, it is often the case that the mouse-down drag can deviate considerably from the scroll-bar (within area C) and yet be interpreted as being on the scroll-bar (Figure 2.6).

With the Table Lens (Figure 2.2) the actual affordance is far from obvious and must be learned. Nevertheless, once learned, considerable insight into the underlying table of data is possible, as is fluency of use.

Hair and *eyes:* the **Mosaic Display**

The two illustrative rearrangements presented above – the crop treatments and the Table Lens – were intentionally simple, even though the Table Lens offers other useful functions which are discussed in Chapter 7. We now examine another example of rearrangement called the Mosaic Display (Hartigan and Kleiner, 1981, 1984; Unwin et *al.*, 1996; Unwin, 1999).

Once upon a time, as all interesting stories seem to begin, a lecturer in statistics (Snee, 1974) recorded, for each of his 592 students, the colour of the student's eyes and hair. The result is shown in Table 2.1, from which little insight can immediately be gained. Instead, we consider (Figure 2.7) a diagram

TABLE 2.1 Data concerning the hair and eye colour of 592 students

Eye colour	*Hair colour*				*Total*
	Black	*Brown*	*Red*	*Blond*	
Brown	68	119	26	7	220
Blue	20	84	17	94	215
Hazel	15	54	14	10	93
Green	5	29	14	16	64
Total	108	286	71	127	592

[4] Except for the last line, to provide continuity.

Eye colour	Black (108)	Brown (286)	Red (71)	Blond (127)
Green (64)	*11.7* **5**actual	*30.9* **29**actual	*7.7* **14**actual	*13.7* **16**actual
Hazel (93)	*17.0* **15**actual	*44.9* **54**actual	*11.2* **14**actual	*20.0* **10**actual
Blue (215)	*39.2* **20**actual	*103.9* **84**actual	*25.8* **1**Tactual	*46.1* **94**actual
Brown (220)	*40.1* **68**actual	*106.3* **119**actual	*26.4* **26**actual	*47.2* **7**actual

Hair colour

FIGURE 2.7 Diagram showing all combinations of eye and hair colouring

to be constructed so that each rectangle, of which there are 16 corresponding to all combinations of eye and hair colour, has a height proportional to the number of students having a given eye colour, and a width proportional to the number of students having a given hair colour. Now suppose – incorrectly, as we shall see – that a person's eye color is independent of hair colour: the area of each of the 16 rectangles in Figure 2.7 (and indicated by the small *italic* inscribed values) would then be a measure of the probability of occurrence of all combinations of eye and hair colour. In fact, these values are those probabilities multiplied by the total number of students, so that the total is 592, the number of students.

The inscribed numbers, however, do not correspond to the data in Table **2.1,** because it transpires that eye and hair colour are not independent. The actual numbers, taken from Table 2.1, are shown bold and labeled 'actual' in Figure 2.7, and are seen to be different and in many cases very different, from the independent' numbers. The areas of the 16 rectangles are therefore misleading. A simple modification of Figure 2.7 is all that is required to provide a far more informative presentation: in the absence of independence the so-called Mosaic Display of Figure 2.8, which is similar to a divided bar chart, can provide more insight (Hartigan and Kleiner, 1981, 1984; Friendly, 1994). Here, the width of

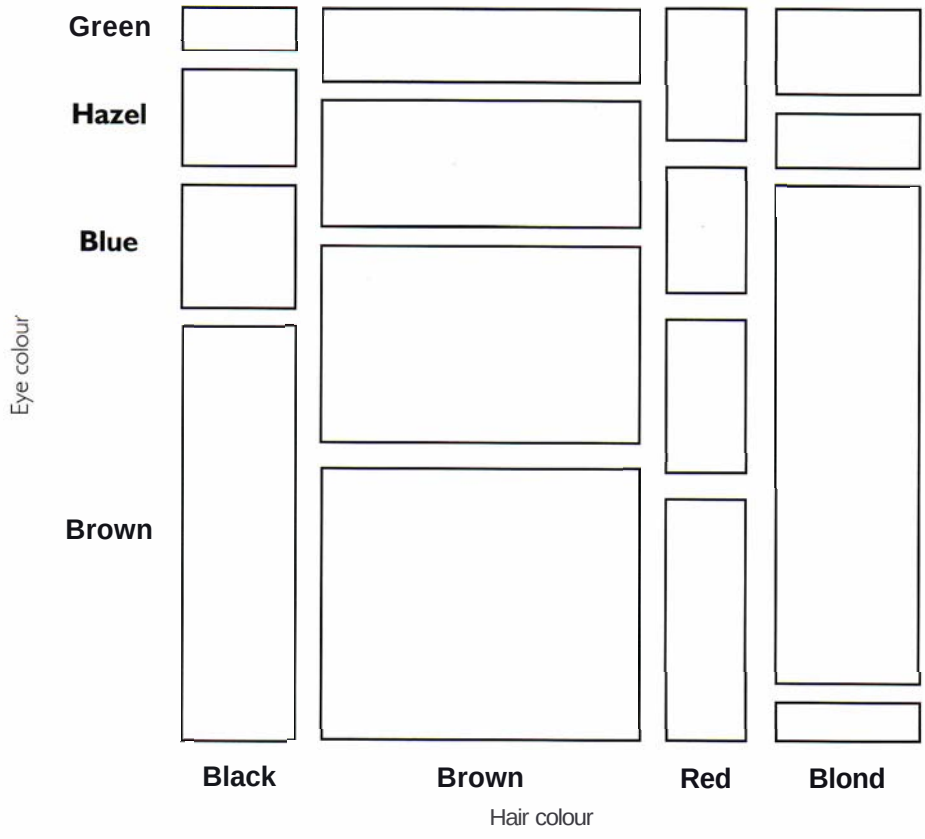

FIGURE 2.8 A Mosaic Display, with areas proportional to the actual numbers possessing various combinations of eye and hair colour

each 'tile' is still proportional to the relative occurrence of the four hair colours, but now the height is governed by the *actual* number of students with the appropriate eye colour. Unlike Figure 2.7, where equal heights indicate an assumed independence, the mosaic display of Figure 2.8 reflects the actual dependence and provides more insight. Certain tiles stand out: there are, for example, more blue-eyed blonds and brown-eyed black-haired people than would occur under independence. With the exception of the first and last rows, the rows are not aligned, which makes it harder to make comparisons within eye-colour groups. This drawback can be remedied by generating a new mosaic plot with the eye colours determining the widths and hair colours determining the heights.

Further rearrangement can be helpful. Figure 2.9 recognizes the deviation from independence between eye and hair colour. Shades of red and green indicate the extent of the deviation from independence, and the ordering of the eye and hair colours (which, of course, are categorical and not ordinal variables) has been chosen to ensure that deviations from independence have an 'opposite-corner' pattern. Extensions of the mosaic display are treated in Chapter 3: the underlying concept is first introduced here to provide yet another example of the value of rearrangement.

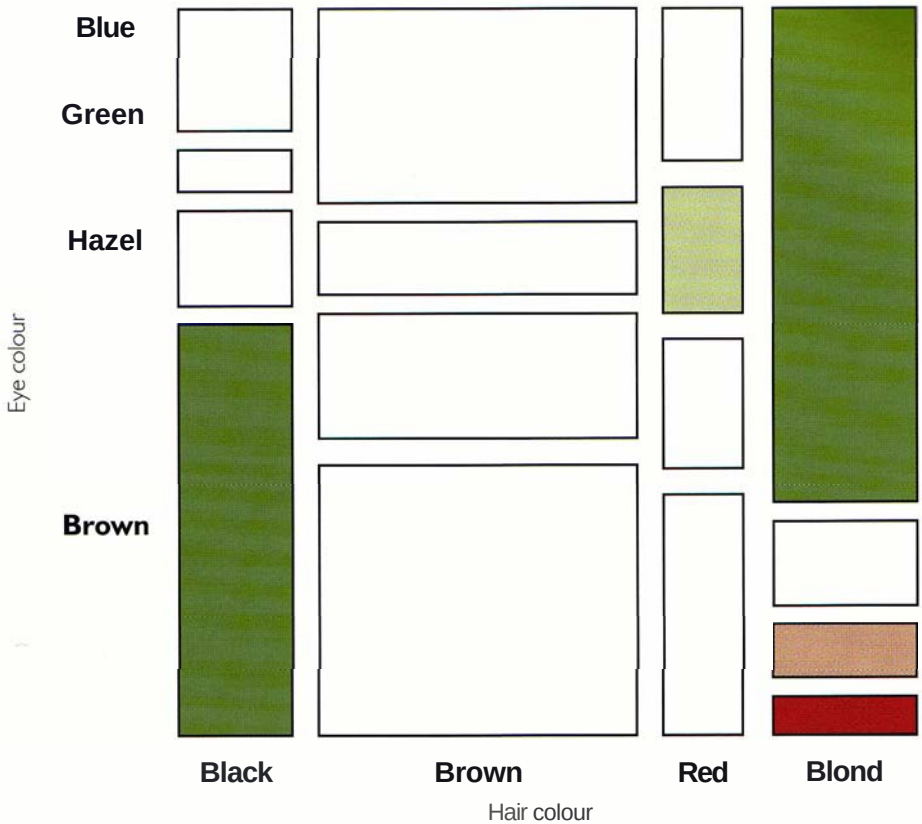

FIGURE 2.9 In the Mosaic Display, colour indicates the deviation from independence

Network data

A dramatic benefit arising from the rearrangement of data can also be associated with 'network' or 'connectivity' data. Suppose that the telephone calls made and received by 13 people (labeled 'A' to 'M') during a specific period are recorded and presented in a table (Table 2.2). Even a careful study of Table 2.2 will typically fail to disclose any interesting pattern or features of interest. Even when represented in the form of a node-link graph (Figure 2.10) it is still not easy to discern a pattern in the same data.

However, if the data is processed in a very simple way to identify unconnected subgraphs, and represented by a node-link diagram (Figure 2.11), we begin to see interesting features. While their interpretation cannot at this stage be deduced it is clear that there are three distinct groups of telephone users between which no telephonic communication took place during the time period of the recording. If the duration and/or frequency of each call was additionally recorded, then the encoding of this information, for example by line width, could provide additional insight (Figure 2.12(a)); so also could colour-coding to represent the proportion of outgoing and incoming calls (Figure 2.12(b)).

TABLE 2.2 Telephone communication between 13 people

Originator	*Receiver*
A	H
C	L
I	M
B	E
F	H
G	I
I	B
B	M
K	B
G	B
K	E
C	J
D	C
F	A
J	L

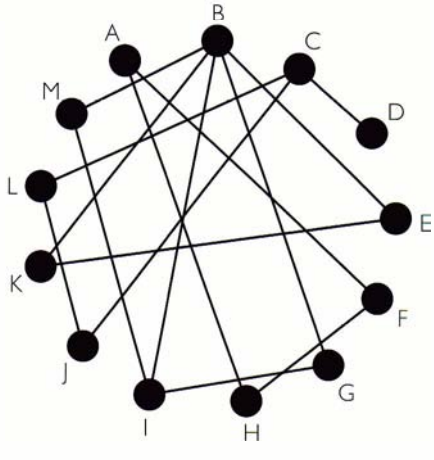

FIGURE 2.10 A node-link graph representing telephone communication between people

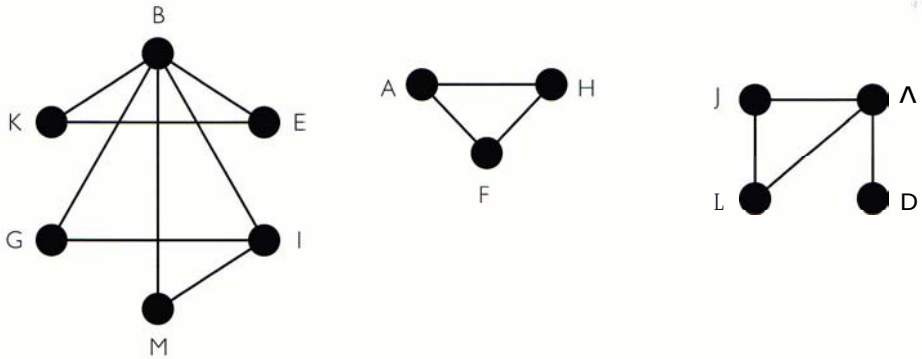

FIGURE 2.11 A rearrangement of the graph of Figure 2.10, identifying an interesting feature

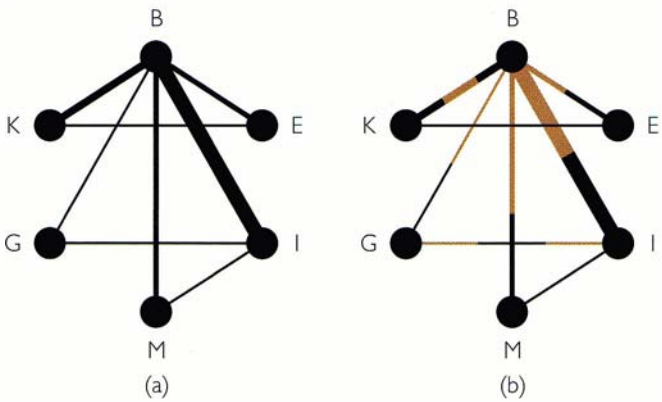

FIGURE 2.12 Methods of encoding parameters characterizing a link

Selective manipulation

The display of Figure 2.13 shows the geographic location of a supply distribution network for a relief effort in a large-scale crisis, as presented by the Selective Display and Manipulation (SDM) system (Chuah *et al.*, 1995). Supply centres are represented by cylinders, main routes between them by dark lines on the 'floor plane'[7] and shelters where supplies are needed by rectangular bars. Heights indicate material quantities. Figure 2.13 might be regarded as a very 'busy' and cluttered display, attracting the comment that perhaps too much data is being presented concurrently, with the attendant danger of occlusion. On the other hand, it could be argued that one's interpretation of (say) the green items depends at times on the layout of the purple items, and that the display of only one type of item would lead to an impoverished display. A useful technique under these circumstances (Figure 2.14) is that of 'raising' the reference level for one class of item so that it can be examined in a clutter-free manner, though still within context. Shrinking the width of irrelevant bars is also possible, and can be helpful. A typical drawback of a three-dimensional presentation – the difficulty of comparing patterns, widths and heights that are at different 'distances' from the user – is overcome by allowing selected bars associated with a user-drawn reference line to be projected onto a two-dimensional presentation that allows comparisons to be made (Figure 2.15).

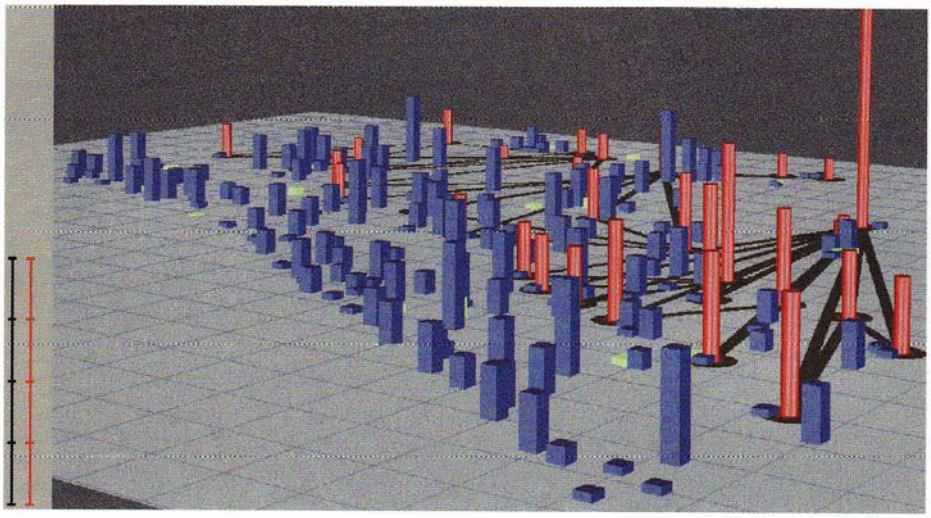

FIGURE 2.13 Display of the geographic location of a supply distribution network for a relief effort in a large-scale crisis

Source: *Chuah (1995), Chuah et al. © 1995 Association for Computing Machinery, Inc. Reprinted by permission*

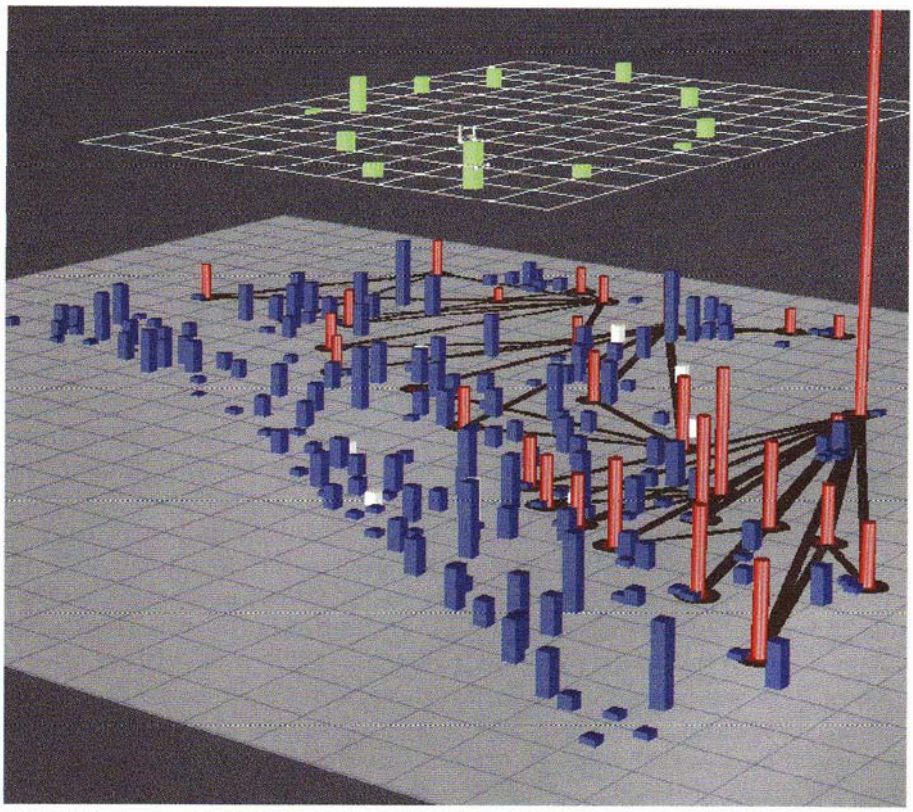

FIGURE 2.14 Raising a reference for one class of item allows that class to be examined more easily

Source: *Chuah (1995), Chuah et al. © 1995 Association for Computing Machinery Inc. Reprinted by permission*

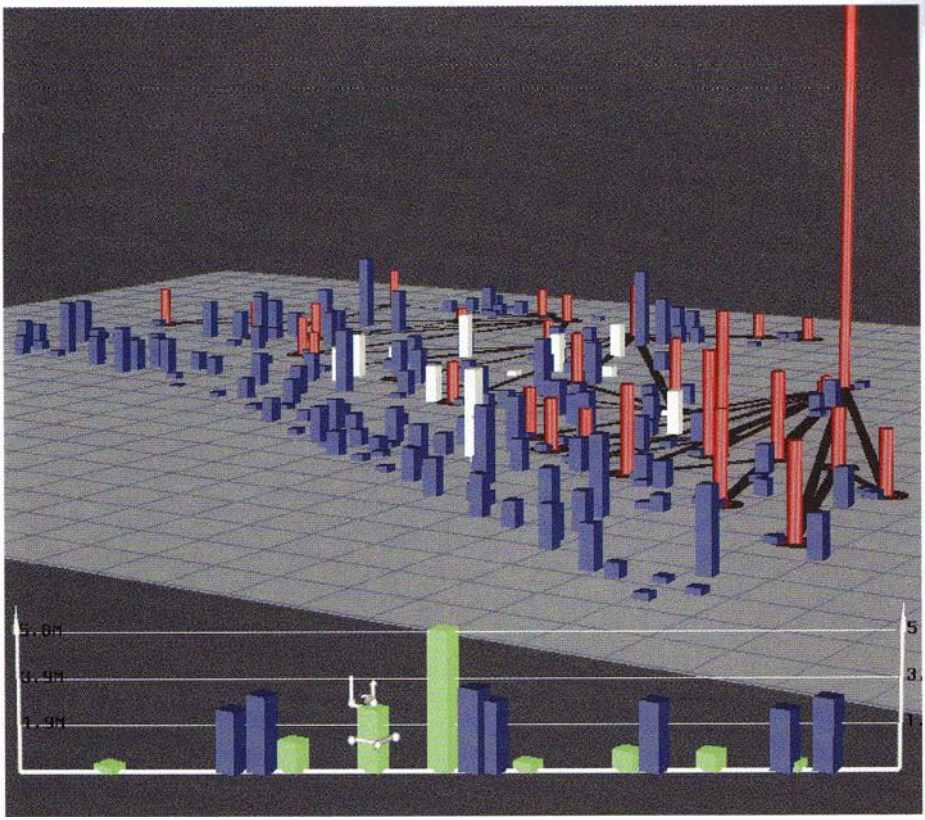

FIGURE 2.1 A reference line enables comparisons to be made

Source: *Chuah (1995). Chuah et al. © 1995 Association for Computing Machinery, Inc Reprinted by permission,*

Algorithms

Many visualization techniques are only made possible by algorithms of considerable complexity which operate upon available data. We choose for illustration an example that relates to a user's search among a collection of documents to identify one (or a group) which is particularly appropriate to that user's interest.[5] Obviously the user does not want to have to read through a large number of documents – and especially those which are of little or no interest – to find the most interesting one; rather, they might be happy instead to supply either some keywords or an existing document known to characterize their interest to some degree and then, via an effective visualization tool, either discover a suitable document or be guided towards a revision of their query. Thus, as with other examples in this chapter, the same data is available but is selected, and its presentation rearranged, to suit the task being undertaken.

An underlying problem with searches for interesting documents is the huge amount of data required to characterize the word content of a single document. Words within a document will typically be examined to determine the frequency

[5] A more detailed discussion of this task is the subject of Chapter 10.

of their occurrence, often resulting in a collection (a 'vector') of thousands of numbers, each denoting the frequency of occurrence of a particular word. Comparison of that vector with keywords supplied by a user will enable the relevance of that document to be estimated. Fortunately, not only are algorithms available for carrying out such characterizations and comparisons, but others – also of considerable complexity – are available which can represent the documents in two-dimensional space on a conventional display. Frequently, this representation is so arranged that 'similar' documents are represented by points which are close together, whereas those with little in common are positioned far apart, and for these points to be displayed (Chalmers, 1993) in a 'landscape' presentation (Figure 2.16) which permits interactive interrogation. Chalmers et *al.* (1996) have proposed two features of a landscape presentation to aid search (Figure 2.17). One is to colour the area bounding those documents which are of interest, and the other, shown at the top of the display in the form of miniatures, is to display previous results and associated keywords; a visible trace of recent landscapes can be valuable in suggesting new or modified queries. In many visualization tools a visible record of earlier steps has often been found to be useful when a step-by-step approach to some goal is involved.

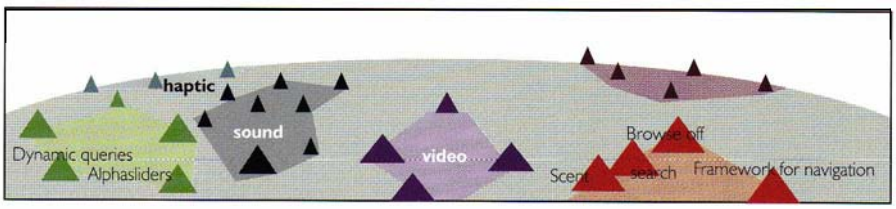

FIGURE 2.16 A landscape presentation of data about documents
Source: *After Matthew Chalmers*

The nature of rearrangement

Rearrangement can occur in different ways. In the Table Lens and other examples a new presentation appeared following a discrete command of some sort. But there is another class of useful rearrangement in which the display of data is dynamically rearranged. Two modes are of interest. In one, a smooth manually controlled rearrangement enables cause and effect to be displayed simultaneously, often to good effect. In the other mode, rearrangement is automated, leaving the user free to concentrate upon interpretation. We briefly examine these two modes of dynamic rearrangement.

2.8.1 Dynamic manual rearrangement

Insight into the relation between two or more quantities is often sought by the manual variation of one variable and concurrent observation of the consequent variation of one or more remaining variables. An example is provided by the activity of engineering design.

The design of a physical object, whether it be a radio or a kettle, is not a straightforward process but, as Chapter 9 will demonstrate in detail, is one that can benefit considerably from information visualization. Figure 2.18 shows part

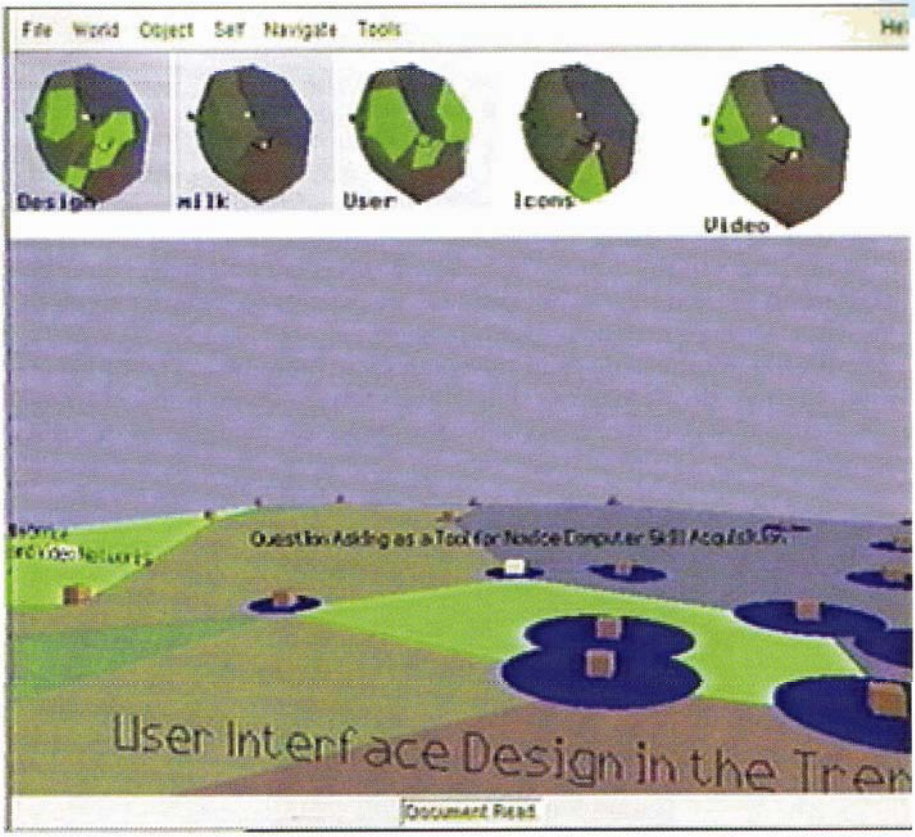

FIGURE 2.17 A visible trace of earlier steps in the examination of colour-coded landscapes

Source: *Matthew Chalmers*

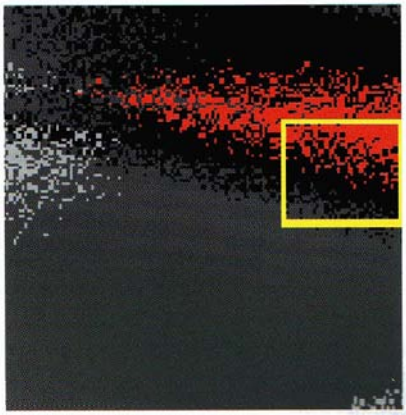

FIGURE 2.18 The red area of this display contains all acceptable designs of an artefact and the yellow box contains all the manufactured samples

of a display that can be of value to a designer. A point **within the display** area represents a design, since its projection onto the horizontal **axis** is one dimension of the object being designed (the base area of a kettle, for example), and its projection on the vertical axis is another – perhaps the thickness of the kettle. The red part of the display contains points (that is, designs) that satisfy the customer: one of the boundaries of the red region might correspond to the cost of manufacturing the kettle. But because a mass-produced object such as a kettle

cannot be manufactured with infinite precision, the many points representing the many mass-produced kettles will lie anywhere within the yellow box. There is therefore much to be gained (especially 100 per cent manufacturing yield) if the yellow box can be made to lie completely within the red region.

The designer can easily move the box (in this case vertically) so that more designs lie within the red region. However, by exploration, the designer has discovered that the picture can be changed to that of Figure 2.19, to provide a larger red area, by choosing a new value for a third dimension of the object – perhaps the height of the kettle. The yellow box has not only been better placed with regard to the red region, but can now be positioned within a much larger red region, ensuring a high manufacturing yield.

As with the other examples in this chapter, no new data is being generated or acquired: here the designer is simply exercising a design freedom by selecting, for view, a different subset of the data, in this case such that more yield can be achieved.

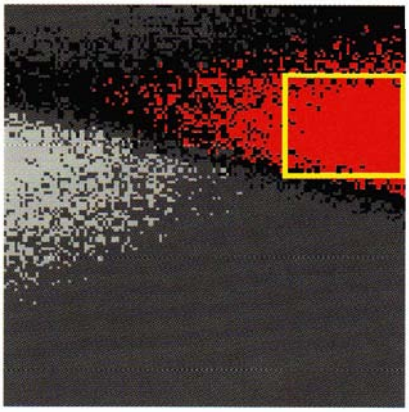

FIGURE 2.19 Alteration of the design leads to a greater area (red) in which acceptable designs can lie enabling the yellow box to be positioned to lead to a greater manufacturing yield

2.8.2 Dynamic automated rearrangement

Figure 2.20 is a display which shows the designer of an electronic circuit the extent to which each component within it affects the circuit's overall performance. Following the designer's sketch of the circuit diagram on a computer display (Figure 2.20(a)), the area of the circle superimposed[6] on a component's symbol (Figure 2.20(b)) provides an indication of the change in circuit performance that would result from a 1 per cent change in component value. This effect is termed 'sensitivity' and the circles are known as 'sensitivity circles' (Spence and Drew, 1971; Spence and Apperley, 1977). Such an indication of sensitivity can be immensely valuable to the designer.

However, if the circuit is a hifi, the designer is interested in the sensitivity to each component over a range of frequencies reaching from bass to treble, so that a static display, as in Figure 2.20(b), is of limited value. A simple solution to

[6] In the actual presentation the shape of the symbol can be discerned through the circle, thereby removing any doubt in the designer's mind about the nature of the component to which a circle refers.

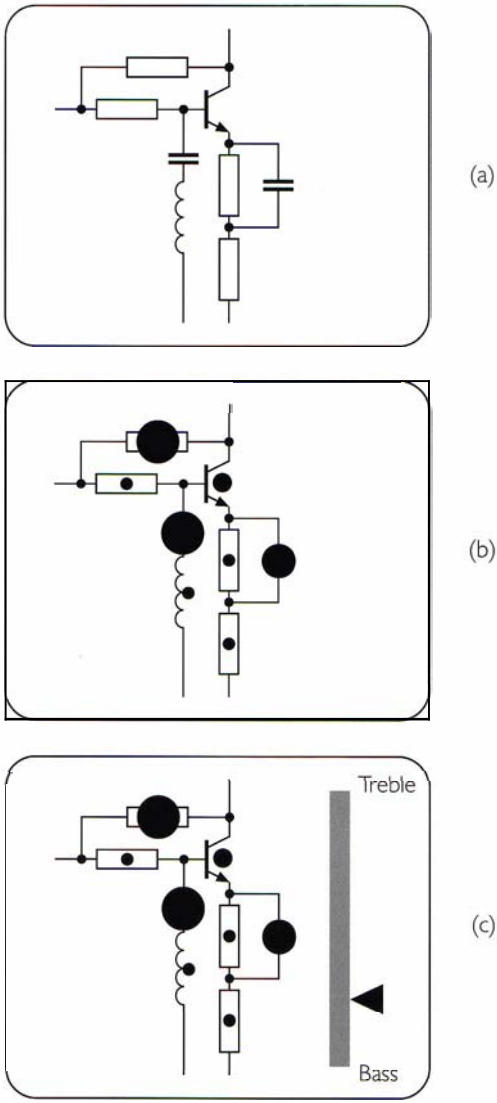

FIGURE 2.20
(a) Part of an electronic circuit diagram drawn *on* a computer display
(b) Circles provide a qualitative indication of the sensitivity of the circuit's performance to a change in each component
(c) A frequency indicator allows animation to display sensitivity as a function of frequency

the accommodation of an additional variable (frequency) is to employ animation: to change the size of the circles as the frequency is automatically varied over a range of interest (Figure 2.20(c)). As the frequency indicator on the right moves smoothly up and down the frequency scale between bass and treble limits, the circle sizes vary accordingly. A designer's observation of which circles are large or small at which parts of the frequency range can immensely enhance their understanding of the circuit. The scanning rate – in other words, the speed of animation – requires careful choice: if circle sizes change too rapidly the designer may fail to comprehend the meaning of the changing circles: too slow and the designer has difficulty in integrating what is seen on the display into a useful mental model. An animation period of about eight seconds was found by experience to be optimum.

For a particular electronic circuit the visualization tool shown in Figure 2.20 displayed a sudden, substantial but temporary increase in the size of two circles. This unexpected result was noticed with some surprise by the designer, but quickly recognized as an indication of 'resonance'. This phenomenon is perhaps more familiar from the example of the destructive oscillation of the Tacoma Narrows bridge, and is equally undesirable in an electronic amplifier: a hifi emitting a constant whistle does not lead to customer satisfaction! Here is an example of an animated rearrangement of data leading to discovery. Fortunately for the designer, the frequency on the bass-to-treble scale could also be adjusted manually, and the precise frequency of resonance read directly from a numerical display beside the frequency scale.

2.9 Conclusion

We have seen, from many examples, how interactive rearrangement of the way in which data is presented can provide an opportunity for additional insight into that data. We have examined categorical data (crops and treatments, eye and hair colour), numerical data (baseball statistics), topological data (telephone connections), symbolic representations (sensitivity circles and complex geographical operations) and textual data (landscapes), merely as representative examples rather than to establish an exhaustive categorization. Many other examples exist, a particularly impressive one having to do with the 'O'-ring disaster in the fated Challenger shuttle (Tufte, 1997). In the chapters which follow, extensive use is made of this potential, and many other examples will be encountered.