




Representation

4.1 Symbolic encoding

In the previous chapter we addressed the problem of visually representing the numerical values associated with an object or an effect, basically by points in a multidimensional space. While there are many circumstances where such representation is appropriate, we need to consider both a wider range of requirements (to be able, for example, to handle ordinal and categorical data) as well as alternative representations of numerical values appropriate to different circumstances. We have already encountered some examples, including Minard's use of line width for numerical data, Playfair's use of slope and Harry Beck's use of colour to differentiate Underground railway lines (categorical data).

A wide range of symbolic representations are available to encode data (Tufte, 1983, 1990, 1997). Unfortunately there is very little in the way of underlying theory to help us exploit this extensive palette of encoding techniques. As Card *et al.* (1999) remark, 'Innovation is still the overriding activity, generating "point solutions" for theoretical workers to encompass within a general powerful theory'. We shall therefore examine a number of techniques, point to the benefits and disadvantages of each, and provide some idea of the task and environment to which each may be suited.

4.2 Size

An example of the use of size to encode data has already been introduced in Chapter 2. In the MINNIE computer-aided design system (Spence and Apperley, 1977) a circle is superimposed upon components of interest (Figure 4.1), the size of each circle indicating the extent of the influence of that component on a circuit property. Such a symbolic encoding of numerical data is often appropriate in the early stages of design when the designer merely wants to gain a *qualitative* impression of the effect of different components: the designer is concerned, not with whether the sensitivity has a value of 5.2 or 5.3, but rather whether it is small, medium or large, and this is precisely the information that is provided by the circles in Figure 4.1. In the same system boxes appear to indicate (Figure 4.2) the magnitude of the voltage at various points within the circuit: again, precise values may at some stage of design be relevant, whereas a pattern of box sizes could immediately alert the designer to an unexpected situation. As already explained in Chapter 2, an additional dimension, that of frequency, can be displayed by means of an animation in which circle or box sizes take on values appropriate to a displayed frequency (Figure 2.20).





FIGURE 4.1

The size of each circle indicates the effect of the associated component on the performance of the electronic circuit

FIGURE 4.2

Black squares indicate the voltage magnitude at points within the circuit



4.3 Length and height

Encoding numerical data by length or height is popular when a qualitative feeling for the data is needed, and can be useful for making comparisons. An example has already been provided by the mosaic plots in the previous two chapters. A different, and more complex example can be found in a system (Colgan *et al.*, 1995) which automatically and quite rapidly changes the values of electrical components in order to improve the design of a circuit. The designer who monitors that process and intervenes when necessary has many representations to interpret within the displayed data (Figure 4.3): apart from component values which are changing quite rapidly – maybe once or twice per second – the designer must observe and interpret a large number of red circles (some of which are seen in Figure 4.3) which together provide a view of the current quality of the design. Under these circumstances a numeric display of component values would be wholly inappropriate, so bar heights are employed, as seen in the lower part of Figure 4.3. The length of each bar indicates the current value chosen for that component by the computer, and the maximum length (indicated by the 'box' containing each bar) represents the limits that the designer has tentatively placed on those components. As automated design proceeds, the bars increase and decrease in height quite rapidly as the computer incrementally improves the circuit. The designer can learn a great deal from observation of the bars: for example, if a bar keeps 'hitting' its upper or lower limit – which will normally have been specified quite tentatively – it is probably wise to extend the limit to see if a marked improvement occurs.





FIGURE 4.3
Part of the 'Cockpit' by means of which a human engineering designer observes the progress of automated design. The red circles indicate the quality of different parts of the design and the top row of bars indicate the current values of the components that are being adjusted automatically

Consideration of the value of bar heights as a means of representation provides an opportunity to discuss two issues: one is the value of *qualitative* representation and interpretation and the other concerns the frequent need for a composite visualization tool which allows a user to switch rapidly between qualitative and quantitative interpretation. As an illustration we select the design (Figure 4.4) of an altimeter for the cockpit of a light aircraft.¹ The vertical bar provides, and is enhanced by familiar colour coding, an easily perceived qualitative indication of aircraft altitude: it is easy to see if one is flying at an acceptable or a dangerous height. Closely located is a quantitative indication of altitude, so that little eye movement is required to change from a qualitative enquiry to a quantitative one. The digital indication is carefully chosen to conform with conventional practice, where altitude is often quoted in hundreds of feet. The last two digits of the digital presentation are, for that reason, chosen to be smaller in size so as to avoid confusion but to provide an unambiguous indication of altitude. Since altitude varies continuously rather than in discrete steps the presentation of the last two digits is also continuous and is so chosen that no ambiguity arises when both those digits are zero. The altimeter design illustrates the fact that, in practice, the designer is faced with combining a variety of representations to handle the tasks that must be supported.

¹ I'm grateful to Neil Mothew (a pilot) for this example.

FIGURE 4.4
 DesignTor an altimeter which provides both qualitative and quantitative indications of altitude
 Source: Neil Motthew, 1999



4.4 Magnification

It is possible to employ a human being's well-developed memory of the **conventional** world atlas to use magnification as an encoding scheme for geographic data. Thus, if the number of bicycles per head of population in New Zealand is ten times that in Australia, then this fact can be encoded in a map (Figure 4.5) in which New Zealand is magnified relative to Australia to reflect the ratio of bicycle ownership. This method of encoding has been extensively and effectively employed in *The New State of the World Atlas* (Smith, 1999). In Figure 4.6, con-

FIGURE 4.5
 The use of magnification to illustrate relative bicycle ownership in two countries



cerned with population, the small size of Canada and Australia is immediately noticeable, as are the comparatively large sizes of India and Japan. Though effective in providing an immediate qualitative impression, use of this encoding mechanism relies heavily upon a good memory for sizes.




FIGURE 4.6

Area magnification is used to indicate national populations


Source: Smith (1999)

4.5 Faces

Professor **Herman** Chernoff, a statistician at Stanford University, observed that human beings are very sensitive to a wide range of facial expressions and appearance, and suggested that facial features such as the size of the eyes, the height of eyebrows above the eyes and the shape of the mouth are quite numerous and could, in a cartoon face, take on a sufficiently large number of 'values' to offer a useful encoding mechanism (Figure 4.7). He therefore applied this technique (Chernoff, 1973) to the study of geological samples, each characterized by 18 attributes (for example, salt content, water content), and found that the display of so-called Chernoff faces facilitated the identification of interesting groups of samples. If this use of computer-controlled cartoon faces should appear frivolous, it should be remarked that accountants, not perhaps typically known for frivolous behavior, have also explored the use of Chernoff faces to display accountancy data (Stock and Watson, 1984). Studies by De Soete (1986) have established the relative value of the various facial features. Following Chernoff's proposal it was recognized that half a face carries the same information as a full symmetrical face, stimulating studies of asymmetrical faces (Flury and Riedwyl, 1981).

FIGURE 4.7

Chernoff faces, in which facial characteristics encode the values of variables



4.6 Multidimensional icons

A Chernoff face is an example of a multidimensional icon: in the example investigated by Chernoff himself the dimensionality was 18. It is not difficult to propose other multidimensional icons appropriate to different tasks and different domains. One example specifically designed to explore the effectiveness of multidimensional icons addressed the problem of selecting, from a number of alternatives, a house satisfying certain requirements (Spence and Parr, 1991). Sample icons are shown in Figure 4.8: colour encodes a price band (red is over £400,000, orange between £300,00 and £400,00, yellow between £200,00 and £300,00 and white between £100,00 and £200,00) while shape encodes a categorical variable (house, houseboat, flat, and cottage). The number of bedrooms was indicated by the number of windows, and these were coloured black or white to denote, respectively, a bad or good state of repair. Garden size was indicated by size, the presence of a garage and central heating each by a symbol and journey time by a clock segment. Overall, eight dimensions are represented.




FIGURE 4.8
Multidimensional
icons encode
eight
characteristics
of a dwelling

house
£400,00
garage
central heating
four bedrooms
good repair
large garden
Victoria 15 mins

Flat
£300,00
no garage
central heating
two bedrooms
poor repair
small garden
Victoria 20 mins

houseboat
£200.00
no garage
no central heating
three bedrooms
good repair
no garden
Victoria 15 mins

FIGURE 4.9
Textual
presentation of
the same data
encoded in
Figure 4.8

Controlled experiments were undertaken to investigate any advantage such icons might offer over a simple textual description of the form shown in Figure 4.9. Typically, a subject was given the following task:

You can spend up to £200,000 on accommodation. Locate the best you can with regard to the number of bedrooms and the size of garden, but it must have central heating.

Overall, it was found that the time taken to identify the appropriate house using icons was about half what it was using a textual description. It was suggested that one significant advantage of the icons was their semantic relation to the task, a relation not present when using faces, for example, to represent geological samples.

4.7 Spatiality

Alistair Cooke, the author for more than 40 years of the BBC radio program *Letterfrom America*, had the problem of storing his vast collection of books about the American states. Many people would have chosen an alphabetical ordering. Instead, Cooke (1973) relied on his excellent knowledge of the geographical location of the American states and used that 'mental map' as the basis for distributing his books around the bookshelves. Thus, books about Florida were to be found on the lowest shelf at the right, whereas those about California were at waist height at the left-hand end. It is not recorded where books on Hawaii were stored – perhaps on a table to the left of the bookshelves!

Cooke's solution exploits the human being's excellent spatial memory, a characteristic that can be exploited in many ways. Malone (1983), for example, remarked upon how an office worker creates piles of documents on a desk to signify that they need attention, and how different piles have different significance. The spatial characteristics and visual appearance of the piles constitute the externalization of attributes: the user's internal model is continually reinforced by the action of placing a document on a pile and creating new piles.

The relevance of spatial memory to visualization is clear. Since visualization is the creation and enhancement of an internal model, it is useful to know that spatiality as a basis appears to lead to a robust internal model, and therefore that phenomena having no inherent location could perhaps be assigned an artificial location – consistently used throughout – in order to enhance the likelihood of a robust internal model.


4.8 Patterns

Patterns can provide a valuable approach to encoding, especially when combined with other encoding techniques. The patterns shown in Figure 4.10(a) and (b) appear in a visualization tool adopted by a large company to investigate and help diagnose problems occurring in the manufacture of mass-produced domestic equipment.² The pattern is formed in three dimensional space in which the 'horizontal' plane has axes associated with month of production (MOP) and months in service (MIS). Plotted vertically is the failure rate. In Figure 4.10(a) the pattern indicates an epidemical failure. Other characteristic failure patterns that can occur, and their interpretation, are shown in Figure 4.10(b). Area 2, for example, occurs exactly twelve months after manufacture, corresponding to the end of the guarantee period; area 7 corresponds to an epidemical failure associated with one particular month's production, and area 4 refers to latent defects. By examining such patterns the investigator can draw conclusions about the production process as it proceeds.

Another field in which human pattern recognition can be useful is that of social intercourse. Although social interaction via the Internet is becoming common through the use of online 'chat environments', present-day embodiments leave much to be desired. The associated display (Figure 4.11) suffers from many constraints with regard to the social cues that can be made explicit. Thus, anyone not messaging (but possibly 'listening') does not appear in the summary; a participant may therefore feel compelled to post a message simply to establish a presence; there is a cost in having repeatedly to read the names of participants; and the conversational dynamics are not explicit. Many of these drawbacks are removed, and other benefits accrue, from a recently invented graphical interface (Figure 4.12) called Chat Circles (Viegas and Donath, 1999) which allows dialogs to be visualized. Here, where time proceeds downwards, coloured circles represent a participant (contributory or otherwise), the colour being used as a discriminant rather than as an identifier,³ the latter function being supplied by a label containing the participant's name. The size of the circle expands, and is immediately noticed, as a message is posted; the message appears within the circle and the colour of the circle

² I'm very grateful to Harald Noordweijk for drawing my attention to this example and supplying the figure from his thesis.

³ The use of colour to enable a user to discriminate between participants poses some difficulty. The colour red, for example, is a sign of danger in the West but is interpreted as good luck in China.

**FIGURE 4.10**

(a) Failure rate of a manufactured artefact is plotted against month of production and month in service
 (b) Characteristic patterns that can occur, and their interpretation

Source: Harold
 Noordhoek





FIGURE 4.1 |

Typical display
of a 'chat'
environment

```
<lebeau> where can i find wvdial
<The-Vox> lebeau: freshmeat.net
<le_beau> ok thanx!
<The-Vox> np
***le_beau has left #linuxhelp
<W0w08> hello
<DA9L> is there any external win modems
<D-side> Yo
*** knowbodyhas joined #linuxhelp
<{Q-Ball}> The-Vox: thanks
```

FIGURE 4.12

Chat Circles
allow the
visualization of
an Internet
discussion



brightens. After the message has been posted the circle gradually decreases in size. As the inventors point out, 'The resulting graphical cadence becomes an important articulator of the flow of conversation . . .'. Other parameters of the display encode additional valuable information: proximity of circles is related, for example, to whether or not a person can hear a particular conversation.

A related tool supports the visualization of an archive of a conversation. In the display of Figure 4.13, in which time again proceeds downwards, a **colour-coded** vertical line is associated with each conversational participant, and horizontal bars represent utterances. In this simple representation – a sort of conversational landscape – we see that person C is a 'lurker', constantly listening but not contributing, whereas contributor A is dominant. Logical zooming allows the actual conversion to be monitored.




FIGURE 4.13
An archive of a conversation carried out by Chat Circles

4.9 Software engineering

The design, development and evaluation of software is, like any other branch of engineering, a task of great complexity in which success is enhanced by the ability of the designer to develop relevant models of that software's behavior under various conditions and at different stages of development. It is for this reason that tools are available that can provide statistics regarding, for each line of code, such aspects as authorship, date of last change, purpose of last change (bug fix or new functionality), how it is reached, how often it is executed, and so on. Nevertheless, the sheer volume of code (perhaps thousands or millions of lines) is such that it is not easy to gain insight or obtain a global view of the software from these statistics. Fortunately, the techniques of information visualization can help immensely, especially when integrated within a well-designed tool. **Seesoft** (Eick *et al.*, 1992) is one such tool.

Seesoft embodies three information visualization techniques which are illustrated in Figure 4.14. This figure shows the representation of a directory comprising twenty source code files containing a total of 9365 lines of code. Files are represented as columns, and individual lines of code by thin rows. The colour of each line shows that line's age: the scale used is the rainbow scale with the newest lines in red and the oldest in blue. Rows are just large enough so that block comments, functions and control structures (for example, *if* and *case*) are visible merely by their indentation.

FIGURE 4.4

The representation of a directory containing 20 source code files


Source: © 1997
IEEE



4.10 Sound

In one sense it is unfortunate that the term 'visualization' contains the term 'visual', since an aural presentation of data – on its own or combined with visual presentation – has much to offer. Classical concert-goers and rock band enthusiasts alike are capable of identifying not only the different musical instruments, of which there may be twenty or more, but also the different notes, chords, tempo, and rhythm that are being played, so the potential for encoding data in sound appears considerable. Aural presentation also possesses some features that are not present in visual presentations. As Card *et al.* (1999) remark, 'The capacity to pick up outliers, akin to the missed note in a symphony performance, may be greater in sonic environments'. As they additionally point out, auditory information is omnidirectional and can easily be heard when not being attended to, and they illustrate this property by describing the playing, through a loudspeaker, of the sound of laboratory results, allowing staff to move around freely. Indeed, when using an automatic teller machine, we are now used to hearing a 'beep' to indicate a change of state of the machine, a very simple but extremely effective way of encoding data in sound and thereby allowing the state of the machine to be visualized.

Figure 4.15 refers to the use of audio encoding in the study of a brain tumour. What is seen in that figure is the result of an MRI scan of a brain, with a tumour clearly visible (Grinstein, 1999). However, the clinical investigator can also point to a particular part of the brain scan and deduce, from the continuous sound that is triggered, some aspect of that tumour.

**FIGURE 4.1B**

Visual presentation of a brain tumour whose details can additionally be visualized by sound triggered by the location of a cursor

Source: Grinstein (1999)

4.1.1 Virtual worlds

One is aware, when viewing a Web page, that thousands of other people may concurrently be viewing the same page; nevertheless, one has no way of **knowing**. A more **striking** realization, however, is that social interaction with them is impossible: there is no opportunity to ask them about the content being viewed; there is no opportunity to meet someone with similar interests. To overcome these severe constraints, research is currently **underway** into electronic imaginary worlds ('virtual worlds'), rendered in three dimensions, with shared user access over a network. In such a world people can explore, socialise, acquire information, teach, play, and generally do most things possible in the real world (Rankin *et al.*, 1998).

The concept of virtual worlds identifies many visualization issues. One is that of representing the user within the virtual world. One approach, of course, is to employ a representation resembling, as closely as possible, the actual appearance of the user. But for many reasons, such as privacy, a symbolic representation may be preferred. One example is called the **StarCursor** (Rankin *et al.*, 1998) and has the general appearance illustrated in Figure 4.16.


The anthropomorphic **StarCursor** is characterized graphically by rudimentary head, body, limbs, eye, and aura: all of these are multimedia channels for personal disclosures, communicative signals and actions. A person's interests filter what they see about other cursors: personal data about medical history, hobbies and current mood can be made available at the discretion of the 'owner' of a cursor. Within the huge range of research questions raised by such a facility, however, we concentrate – perhaps cursorily in view of the many issues involved – on some parameters of the **personal** cursor. The body, for example, provides an opportunity for 'clothing' to reflect mood (as in real life), body shape can indicate attitude, and the aura (both visible and audible) can indicate interest or mood. One might hesitate to approach the group shown in Figure 4.17, but feel quite uninhibited about making contact with the group depicted in Figure 4.18.

FIGURE 4.16


A Starcursor representing a human being in a virtual world

Source: StarCursor
Design © Philips
Electronics N V
1998 Reproduced
with permission

The anthropomorphic cursor

**FIGURE 4.17**


A group of inhospitable Starcursors



4.12 Colour

As already demonstrated, numerical data can be encoded by colour. Use of colour should be undertaken with caution, however, and not with the hysterical abandon of a child discovering its first paint box. There is, for example, a minimum size of line or point whose colour can easily be discriminated, and a

FIGURE 4.18
A group of
welcoming
StarCursors



consequent need to exercise care in the construction of graphs. Colour scales to represent magnitude should also be chosen carefully. A popular choice is the spectrum, with red denoting large elements and blue the small elements. Another was mentioned in Section 4.6 in connection with encoding the price ranges of houses for sale. Here, the scale red–orange–yellow–white was employed to denote decreasing values, and it is interesting to note that, during the associated experiments, none of the hundred or so subjects asked to be reminded about this scheme. An alternative familiar from geographic atlases is the green, light green, light brown, dark brown, gray, and white sequence. Grayscale shading (Figure 4.19) can also be used to advantage, though there is quite a low limit to the number of shades that can satisfactorily be discriminated. There are many other guidelines to the use of colour (Jackson *et al.*, 1994). A good illustration of the use of both colour coding and animation is provided by the GVis ('geographical visualization') system (MacEachren *et al.*, 1998) as shown in Figures 4.20 and 4.21. It also illustrates a parallel with the electronic circuit example discussed in Section 4.2 in that spatial layout is meaningful: the investigator using GVis is aware of the layout of American states just as the circuit designer knows where each component is located as well as its anticipated effect on circuit performance. The display of Figure 4.20 is concerned with georeferenced health statistics: it uses the colour ordering dark purple, medium purple, light purple, grey, light green, medium green, dark green




FIGURE 4.19
Encoding by
grayscale
shading

FIGURE 4.20


Illustration of the use of colour in the display of georeferenced health statistics

Source: MacEachren

**FIGURE 4.21**

A bivariate cross map, illustrating the benefits of interactive rearrangement

Source: MacEachren



to depict white male, age-adjusted, heart disease mortality rates by health service area (HSA) over the period 1979-1981. The colour scheme puts emphasis on the two extremes and on the middle of the range. At the lower edge of the display, familiar video controls facilitate automated animation of an adjustable speed along a 'time line' though direct manipulation is also supported.

The benefit of interactive rearrangement, discussed in Chapter 2, is illustrated in the bivariate cross map of Figure 4.21: it depicts white male, age-adjusted, heart disease mortality rates in the period 1979-81 crossed with per capita income. Both of these variables are focused on the 80th percentile. Blue represents HSAs with a mortality rate in the top 20 per cent, and light shades depict income in the lower 80 per cent with a few dark blues depicting the anomalous high-income-high death-rate HSAs. In this way light blue indicates HSAs with a high rate of heart disease and low income. Characteristics of expert use of the GVis system are discussed by MacEachren *et al.* (1998).

4.13 Comment

Exposure to the many ways of encoding data visually and aurally prompts the obvious question 'Which is best?' In one sense the answer is 'it depends', since the task for which the user seeks to form a mental model can take many different forms, and the domain for which the data is encoded can also vary widely. In another sense the answer is 'we don't know'. Various attempts have been made to identify, separately for numeric, categorical and ordinal data, the comparative benefits of a variety of encoding techniques such as length, size, angle and colour. Cleveland and McGill (1984), for example, presented working hypotheses about this issue, as did Mackinlay (1986) with a view to automating the design of visualization tools. Nevertheless, the fact remains that the issue is still unresolved. Contributory factors to this lack of knowledge include the fact that encoding techniques are rarely used in isolation, that the task itself influences the best choice of encoding, that a real visualization tool is normally quite complex and must satisfy a number of constraints, and that early assessments of encoding techniques concentrated on passive and some active visualization tools rather than on the more recent dynamic tools. It is for this reason that the approach of this and the previous chapter has been to present a rich palette of techniques whose application in a wide range of tools is illustrated in the remaining chapters.