Chapter 2
Issues

What are the principal issues that need to be addressed by a book about information visualization? To answer this question we examine how and why an information visualization tool is typically employed. It can be employed in many ways for very different purposes. In one situation we may be concerned with a single – possibly printed – image that is intended for viewing by many people: a railway time table, and Minard's map come to mind. In these examples interaction between user and image may be described as 'passive'. In a very different situation a user seeking information or a product may interact intensively with what is seen on a display. The approach adopted in this brief chapter is to select a few representative tasks to see what important issues they identify. Those issues are presented in **boldface**.

But first, a reminder of what information visualization is. It is *not* what a user sees on a computer display. To repeat the earlier definition,

Visualization is the formation of a mental model of something.

This definition must be kept in mind as we identify important issues associated with information visualization.

2.1   A Task

The first task we address is that of finding a car to buy, a specific but representative example of the generic and widespread task of selecting one **object** from among a collection of objects on the basis of its **attributes**. A variant would be the 'whittling down' of a large number of objects to identify a few worthy of further consideration. An associated and vital subtask, that of **gaining insight** into – that is, forming a mental model of – a collection of data, is widely relevant and also an essential component of activities such as data mining and decision support.

Although we shall first address the task of selecting a car to buy using a digital information space we shall note useful parallels with the physical world of car showrooms.

2.2 Fuzzy Goals

Typically, someone looking for a car to buy will not express their needs with precision: "nice looking", "sporty", "around £3,000" are examples of articulated requirements. Other requirements of which the buyer is aware but which may not initially be articulated could well include shape and colour. Not surprisingly then, and as happens with many tasks, it is frequently the case that **problem formulation** at first takes place only in general terms: more precision may be introduced as the task proceeds (Schon 1983). We often do not know what we do not know.

2.3 The Data

Despite a general preference for choosing a car by visiting a showroom, data describing a collection of cars usually originates in the form of a table and is frequently presented to a car buyer in this way (Fig. 2.1). Each row corresponds to a single car (in general, an **object**), while each column is associated with an **attribute** of that object. Not all attributes are of the same type: some are numeric (e.g., price), some are categorical (e.g., make), some are ordinal (e.g., rating) and one may be an image (e.g., the appearance of a car). A photograph of a car can often exert considerable influence on choice.

Make	Price (£)	MPG	Rating	Age (yrs)	Photo
Ford	8000	42	*****	5	
Renault	6500	37	***	3	

Fig. 2.1 Example of the tabular representation of car attributes

2.4   Table Presentation

Notwithstanding the public's widespread familiarity with them, tables are often of limited help especially if there are many rows (cars) and ten or more columns (attributes). In the very unusual situation in which a buyer knows *precisely* what they want the table can certainly be searched, by eye or some automatic search mechanism, until a car satisfying those requirements is found (or, more likely, not found!). Nevertheless, even if a precise match is found, the purchaser may still wonder what else may have been acceptable and, on further reflection, perhaps more desirable. Even if a precise match fails to be found it is unlikely that a line-by-line search of the table will contribute much to the buyer's **mental model** of the data. Nevertheless, a table is not without its attractions. As we shall see later when discussing the Table Lens, an **interactive rearrangement** of table rows according to some attribute (e.g., Price) can be immensely helpful.

2.5   Derived Data

There are alternatives to the table representation. For example, the price of a car can be **represented** by a dot on a scale (Fig. 2.2) and the actual value estimated to within an accuracy acceptable during the early stages of a search. Using this same representation the price of each of a collection of cars can be presented to a buyer in the manner shown in Fig. 2.3a. With regard to Price, what is shown in Fig. 2.3a can be called **raw data**. But if the prospective car buyer is simply trying to get some rough idea – an **overview** – of car prices an alternative representation may be more effective. One (and only one of many) such representation of what we call **derived data** is shown in the Tukey Box Plot of Fig. 2.3b. We see there (1) the *median* price (which simply divides the number of data items into halves – it is *not* the mean),[1] (2) the bottom end of the box representing the 25 percentile (i.e., below which one quarter of the cars are to be found) and the top end representing the 75 percentile, (3) the separate horizontal bars representing the 5 and 95 percentiles and (4) 'outliers' represented as points. What the Tukey Box Plot does not fully show, however, is the distribution of the data points, but this information is easily added by a


Fig. 2.2   One representation of the price of a car

1800   1900   2000

---

[1] The median is much more useful than the mean, because one false data observation can seriously throw off the mean, whereas the median is not affected until half the data observations are false. In other words, the median is a much more robust measure.

Fig. 2.4 A Tukey box plot and a violin plot for the eruption length of the geyser Old Faithful
(*From* Hintze and Nelson (1998) *reprinted by permission of the American Statistical Association)*

representation known as a Violin Plot (Hintze and Nelson 1998), as illustrated by
data relevant to the eruption length of Old Faithful, the geyser in Yellowstone
Park (Fig. 2.4).

Fig. 2.5 A histogram of car prices

Fig. 2.6 A bargram representation of the data represented by the histogram of Fig. 2.5

Another familiar representation of derived data is the histogram (Fig. 2.5): again, an **overview** of prices is facilitated. One advantage of these representations of derived data is that rapid judgments can be made (e.g., "cars are cheap or expensive – nothing in between" or "there are too few cars – let's not bother"). In other words, a **qualitative awareness** of some aspect of data can be achieved and, as we shall see later in the book, sometimes within about 100 ms.

2.6   Alternative Representations

The representations of Price data shown above are only three of many that are possible. A useful variant of the histogram (Fig. 2.5) is achieved by 'pushing over' its columns and then abutting them as shown in Fig. 2.6, a **representation** known as a **bargram**. But why should a bargram be of any interest, especially as it omits the information that no cars are available in the £14k–£16k range? As well as offering an overview of car prices it provides a convenient framework for adding more information, as shown in Fig. 2.7. Here, individual cars are represented by small circles (icons). The advantage of this is that we can **encode**, quite compactly in the circles, additional information about each car. For example (Fig. 2.8) the appearance of a car can be clearly associated with its price by the use of colour. The car can be selected for identification by **interaction**, either with its image or its icon. Clicking on an icon to change its colour and result in the framing of the corresponding image

Fig. 2.7 Icons positioned above price ranges represent individual cars


Fig. 2.8 Colour coding can be used to associate a car with its appearance. Either the image or the icon could be selected interactively

is called **brushing**, and is an extremely powerful technique (Cleveland et al. 1988). Since it will often be the case that a purchaser notices an attractive car and wants some idea of its price, **brushing** in the opposite direction can provide an immediate answer. The price could, of course, be displayed in a label just below its image, but there is no reason why Price – just one of maybe ten important car attributes – should be chosen for that label. Similarly, clicking on a price range bar could cause brushing into the corresponding car images as we shall soon see in Fig. 2.11.

2.7   Adding a Representation

A buyer will usually have in their mind the attributes of their **ideal** car. Although that car may not – and probably will not – be part of the available collection, it can be useful to **represent** it amongst the available ones (Fig. 2.9) to permit comparison with what is available. Similarly if, in the course of exploring available cars, one is considered to be a 'possible' and worth remembering, it can be **tagged** as shown in Fig. 2.9.

Fig. 2.9 The positioning of an ideal (and possibly non-existent) car (*red triangle*) and the tagging of an existing car that may be worth remembering

2.8   Multiple Attributes

Price is not usually the sole criterion on which the selection of a car is based. If the miles-per-gallon (MPG) performance is also of interest another bargram can be positioned above the Price bargram (Fig. 2.10). If object icons representing cars are also positioned above this new bargram then, because they represent the same cars, the two bargrams will be of the same length in terms of numbers of cars. But because, in general, there will be no correlation between the positions of the two sets of icons, a particular car will not necessarily be represented at the same position on each bargram. Indeed, Fig. 2.10 shows that one particular car has been selected (perhaps through interaction with its image) and that selection has been brushed into the two rows of icons: reverse brushing would have the identical result.


Fig. 2.10 By the use of multiple bargrams a number of car attributes can be represented, and brushing used to link different attributes for a selected car

2.9   Attribute Specification

Usually, the prospective car buyer will **explore** the available cars by specifying acceptable ranges of attributes, achieving this by interacting with – perhaps by clicking on – the appropriate ranges in the bargrams, whereupon the identified range is highlighted to provide confirmatory feedback. In Fig. 2.11 the buyer has expressed an interest in cars costing between £12,000 and £14,000 and having an MPG performance of around 40. Colour-coded icons and images then represent the two cars that satisfy both selected attribute ranges. Figure 2.11 illustrates the important action of filtering out items considered, at least temporarily, to be irrelevant. It would, of course, be impossible to know which car icon corresponds to which car image, so the technique of **brushing** is again employed. Identification (by mouseover, for example) of *any* of the three representations of a single car will highlight the other two.

Fig. 2.11 The interactive selection of a Price range and an MPG range, resulting in the colour coding of icons and the presentation of images representing the two cars that satisfy both of the specified ranges

In the course of a search for an acceptable car the buyer will undoubtedly **explore** – by selection and deselection – the effect of different range selections. As this process proceeds the buyer will hopefully build up some understanding – i.e., create a **mental model** – of what is available. Eventually the buyer may decide to concentrate on only one Price range and only one MPG range, in which case a facility to remove from view – i.e., **filter out** – all other (distracting) ranges (as in Fig. 2.11) may reduce the **cognitive effort** involved in a search.

2.10   Display Space Limitations

As many as ten attributes of a car may be examined before a final decision is made as to which one to buy. Unfortunately there may be no room on a display to contain so many bargrams, as well as the image collection, in the form and detail shown in Fig. 2.8. It is of course always possible to implement a scrolling function, but the disadvantages of hiding a majority of bargrams at any one time will be familiar to anyone who has encountered situations in which only a fraction of available data can be seen.

An alternative and advantageous **presentation** of those bargrams is shown in Fig. 2.12. The vertical size of all but a few bargrams is reduced, allowing them *all* to appear on the display, but necessarily requiring range values to be suppressed and car icons to be superimposed on the bargrams. However, if the buyer is exploring Price and MPG, any significant change on the other bargrams would immediately be apparent and could be investigated. When necessary, a stretching action brings other bargrams into full view so that range values can be discerned.
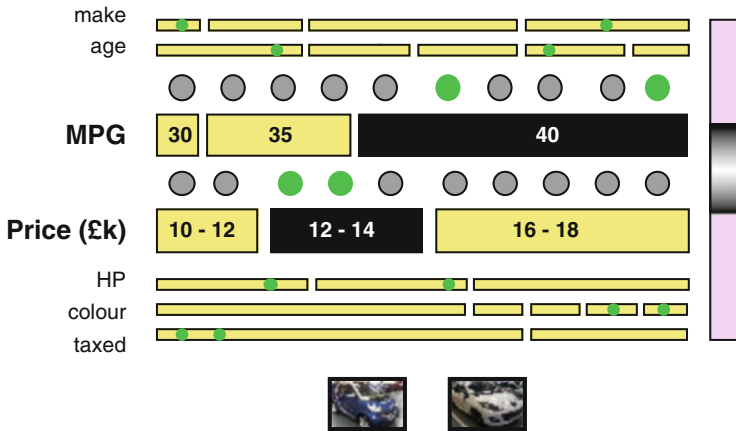
Fig. 2.12 If display space is too limited to accommodate all attribute bargrams, many can be reduced in size vertically, albeit with a loss of the range values but with the advantage of a better overview

2.11   Mental Models

It has been stressed that **visualization** means 'the formation of a **mental model**[2] of something'. A mental model has been described as an internal representation of external reality. We don't know very much about mental models, but we can be sure about some of their features. For example, we do not usually form a single model about some data, but rather (Tversky 1993) a collection ('collage') of sub-models associated with each other either tentatively, confidently or perhaps not at all. And the collage varies with time: it may be enhanced as we accumulate new insight but also pruned as we conclude that some content is no longer relevant (perhaps because of a decision to concentrate on inexpensive cars). Enhancement of a mental model will result from both overview and detailed information. The model will not always be quantitative: a qualitative overview may sit alongside quantitative detail ("the sporty looking car for less than £12,000").

2.12   Near Misses

It is quite common, after finding a car whose attributes are judged to be satisfactory, to wonder "what would happen if we had another £400 to spend? Perhaps a very much better car would be available" An answer to this question could, of course, be obtained by selecting the next higher price range and observing the result. But such

---

[2] Also referred to as a cognitive map (Tversky 1993).

Make:

	Chevrolet	Daewoo	Dod	Ford	Honda	Hyundai	Kia	Ma	Mit	Nis	Por	Sat	Sub	Suz	Toyota	Vol

ConsumerGuide Rating

| Be | Recommended | | Bu | | | | No Rating | | | | | | | | |

Base Price
Dollars

| 9.2 | 9.5 | 9,700 | 10,000 | 10, | 10, | 11, | 11, | 11,800 | 11, | 12,000 | 12, | 12, | 13, | 13, | 14, | 14, | 15, | 18, |

Loaded Price
Dollars

| 10, | 10,200 | 10, | 11, | 12 | 12 | 12 | 13, | 15, | 15, | 15, | 15, | 15, | 16,000 | 16, | 16, | 17, | 19, | 20, | 21, | 24, |

City Fuel Efficiency
Miles/Gallon

| 22 | 23 | 24 | 26 | 27 | 28 | 29 | 31 | 33 | 34 | 35 | 39 | 41 | 42 | 61 |

Highway Fuel Efficiency
Miles/Gallon

| 29 | 31 | 33 | 34 | 35 | 36 | 37 | 38 | 39 | 40 | 41 | 43 | 47 | 49 | 70 |

Vehicles matching your feature selections. (Color a vehicle in all rows by selecting a color over the picture.)

Dodge/Plymouth Neon   Ford Escort/ZX2   Mazda Protege   Toyota Corolla
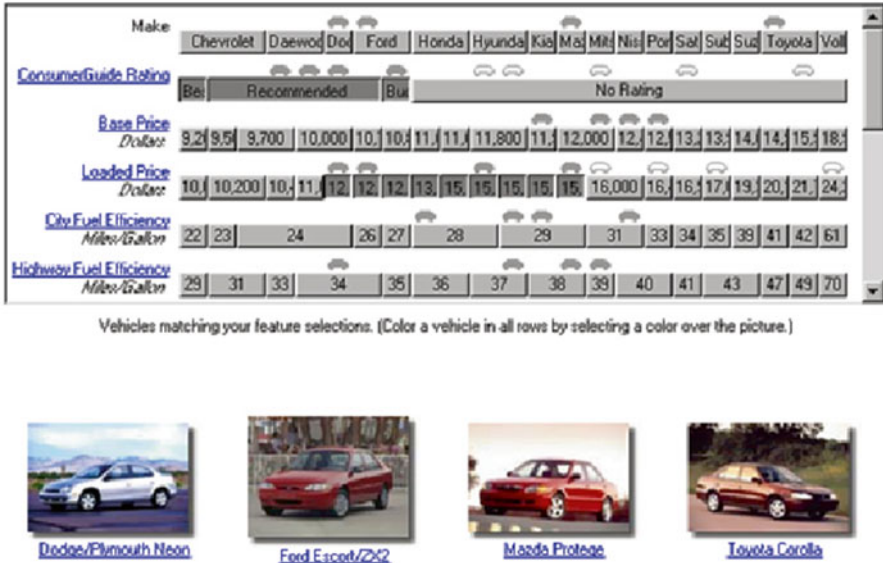
Fig. 2.13 A view of the EZChooser, showing sensitivity information in the form of outline cars (*Courtesy of Kent Wittenburg*)

a step would then have to be reversed in order to ask a similar but different question about another attribute: "what would happen if we weren't concerned about MPG?" A much more effective way to answer such "what if?" questions is to encode 'near miss' features in an icon. In the (real) example (Wittenburg et al. 2001) of the EZChooser (Fig. 2.13) the icons above the bargrams take two forms: a 'filled' car shape and an 'outline' car shape. The former indicates a car that satisfies all the specifications, whereas an outline car indicates that the bargram range below it must be selected for that car to satisfy the new set of needs. What this encoding is doing is providing **sensitivity information**[3] which can be extremely useful in supporting a user's **exploration** and **search**.

2.13   Time

The simple example of car purchase has identified many of the issues discussed in this book. Nevertheless, to identify other important facets of information visualization we now focus a little more on **the human user** and especially the processes of **perception** and **cognition** (Fig. 2.14).

Users don't just want to gain insight into data – they usually want to do so as quickly as possible. Fortunately, computation can be extremely fast, and this leads to many opportunities, some familiar from the real world. For example, imagine that you

---

[3] The concept of sensitivity has very wide relevance, and will be a recurring theme of this book.

Fig. 2.14 By riffling the pages at a fast rate (e.g., ten per second) an illustration can quickly be found

want to show a friend a particularly nice painting in a book. Rather than look up the index you might instead **search** for the item by rapidly riffling the pages – say at a rate of about 10 per second – in which case you would probably find that painting very rapidly. This is because you are able to recognise the painting in around 100 ms through a process known as **preattentive vision** (Ware 2012). Equally, you may riffle through a car magazine to see what attractive cars are featured. You often employ the same ability to handle rapid change in the images on a display when you fast-forward and rewind recorded programmes on your television. This remarkable property of the **human visual system** has many applications. Imagine a search for a Christmas present. One approach, of course, is to wander through a department store and see what's available. But that takes time – why not stand still and arrange for the department store to move past you? Impossible in the physical world, but not in the digital one: Fig. 2.15 illustrates the **presentation** of products by 'floating' them towards you and allowing them to disappear 'over your shoulders', much as the signs on motorways 'move' past you. The ability to reverse and stop such a presentation facilitates **browsing** and offers many advantages over the familiar interface in which you are interrogated about what you are looking for before you are allowed to see any products.

Fig. 2.15 Representations of products 'float' towards you as if you were moving quite quickly through a department store. If you see something of interest you can slow the movement down, reverse it or stop it to interactively gain more detail

Time is also a key feature of exploration when trying to learn how things are related. Suppose you need a financial loan. A typical website might appear as in Fig. 2.16, implicitly making the assumption that you (a) know how much money you need, and (b) know the period over which you will repay the loan. Since the relation between Term, Amount (of loan) and Monthly Payment is known to the Lender, why not provide the user with a **fluid interactive interface** such as the one in Fig. 2.17? It allows you to vary *any* of the three important quantities and see its effect on the others. If that is possible the user can not only explore, but – most effectively – explore *dynamically*. As they smoothly move the sliders up and down and see the resulting effect they are building up a **mental model of** – in other words they are **visualizing** – the relations between variables that are of interest to them.

Fig. 2.16 An interface that shows the monthly payments required for a specified loan amount over a given term
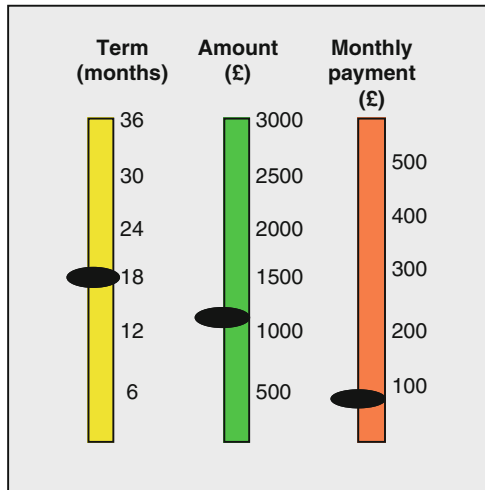
Enter Term (months)

Enter Amount (£)

Your monthly payment is

Fig. 2.17 An interface that allows a user to dynamically explore the relation between loan amount, term and monthly payments

Term (months)	Amount (£)	Monthly payment (£)
36	3000	
		500
30	2500	
		400
24	2000	
		300
18	1500	
12	1000	200
		100
6	500	

2.14   Archives

'**Archival time**' is fundamental to a great deal of data. News stories accumulate over time and provide a rich source of archival material for both historians and the general public to interpret and understand. Similarly, measurements of physiological data must be **represented** and **presented** in such a way as to facilitate interpretation. There are many ways of representing such data: an attractive one is called ThemeRiver (Havre et al. 2002) and is illustrated in Fig. 2.18. Here, coloured 'currents' represent changing content in a document collection.


Fig. 2.18 The ThemeRiver representation of time-based data (©*2002 IEEE Reprinted*, *with permission*, *from* Havre et al. 2002)

2.15   Unwanted Information

I'm driving north from Huddersfield to Halifax and my navigator is using the map shown in Fig. 2.19. We've successfully driven north through Huddersfield and simply wish to get to Halifax without undue difficulty. One problem is that the map contains far more information than we need. We do not, for example, plan to visit
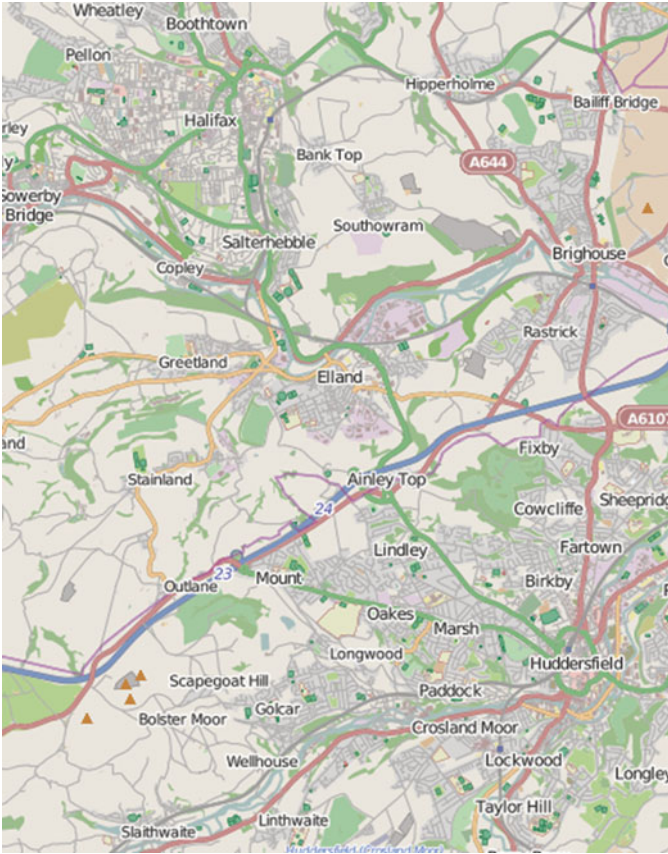
Fig. 2.19 A conventional map of the region between Huddersfield and Halifax (© *OpenStreetMap contributors*)

the undoubtedly charming village of Woodhouse. In fact, it would help enormously if a new map was available that took account of my **degree of interest** in the various features contained within the conventional map of Fig. 2.19. Such a map is shown in Fig. 2.20. All the driver needs to know is that, once through Huddersfield, he must travel along the A629 and then the A6025 to get to Halifax. Purely for reassurance a landmark (the motorway over which the driver will pass) is shown. There are many other situations where documentation is complex and can be simplified by the removal of items of little or no interest.

Fig. 2.20 A map reflecting a
driver's degree of interest in
the journey between
Huddersfield and Halifax

2.16 Movement

Anyone who has **searched** the internet for information will be aware that they are
'moving', step-by-step, through **discrete information space** from one web page to
another, partly to **explore** that space (and form a mental model of it) and partly to
progress towards the page that contains the required information. Such movement is
often made difficult by poor interface design: either the **interaction** required to
move from one web page to another is not obvious or the available **destinations**
(web pages) are themselves not obvious. Figure 2.21 illustrates the confusion that

Fig. 2.21 An example
chosen to illustrate the
difficulties of navigating
discrete information spaces
(see text)
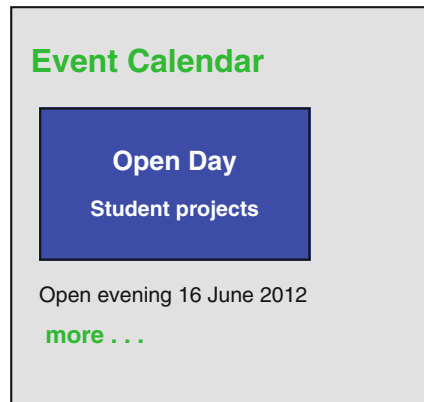
can arise from a poorly designed web page, confusion that impedes the formation of a **mental model** of the immediate region of information space. The small part of a College's web page appears to emphasise the Open Day that I want to attend in order to see this year's student projects. So I click somewhere on the large blue area that advertises the Open Day. Nothing happens. So I try to click on the text giving the date. Again, nothing happens. I then notice the word 'more', with dots to indicate it might lead to something, so I click on it. Success! – I get transferred to a page with illustrations of the exhibits and details of how to register. But wait! – I noticed that the sensitive 'more . . . ' is in a particular shade of green. The same shade of green was used for 'Event Calendar', so perhaps I can see a list of events, one of which is the Open Day. So I click on Event Calendar. Nothing happens.

There are many lessons to be learned from just this small part of a home page (Nielsen and Tahir 2002), and all are associated with the **navigation** of discrete information spaces. In particular, it is essential to make it obvious which items are **sensitive** to **interaction**.

2.17   Conclusion

We have identified a number of important concepts in the field of information visualization, and in general we have noted that there are many applications in which data is first **represented** (e.g., by bargrams – Fig. 2.13) and then **presented** (e.g., by a simplified map – Fig. 2.20). An interpretation of that presentation of represented data by a **human user** engaged in a variety of tasks can then usefully lead to **interaction** (e.g., to navigate to a new web page). That sequence is illustrated by the Reference Model of Fig. 2.22.

Let us examine the Reference Model in more detail. On the left we have raw **data**, the content into which a user wishes to gain insight. Usually that data is vast in size, but purely for illustration we take an example (a) of a 5×5 numerical array. (We also saw that a view of **derived data** can sometimes be preferable, but for the moment we implicitly include that within the box marked 'data'). There are many ways in which numerical data can be **represented**: in Fig. 2.22 (b) we have chosen to represent each number by the size of a circle to convey, more easily, the location of large numbers. However, it may well be that, while keeping all the circles in view, the user wishes to examine, in more detail, some aspect of the value in the centre, in which case (c) the centre of the display is 'stretched' to provide more room for that detail. This is only one of many ways in which represented data can be **presented** for viewing and comprehension. Additionally, as pointed out, the **human user** may well wish at some point to **interact** to change either the data that is selected for representation, the manner in which it is represented or the way in which it is presented. At the right of Fig. 2.22 two examples are illustrated. In the uppermost example (d) colour coding is employed to represent some feature of each of the numerical entries shown below the Data box, while below it (e) is a sketch suggesting an interactively driven animation showing, for example, how a patient's blood pressure changes with time.
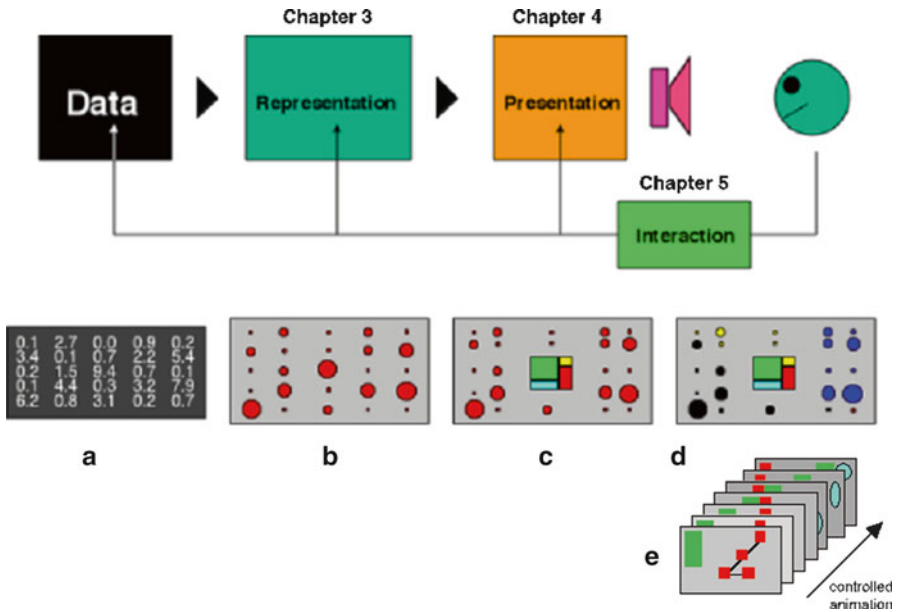
Fig. 2.22 A model of a system supporting information visualization (often called the Reference Model) and very simple examples to illustrate the function of each component

2.18   Remainder of the Book

The model of Fig. 2.22 defines the theme of the following chapters. Chapter 3 addresses the topic of Representation, Chap. 4 that of Presentation and Chap. 5 the topic of Interaction. What is *not* considered in a separate chapter are the issues associated with the capabilities and limitations of the **human user**, for one simple reason: those issues are intimately associated with how we represent, present and interact with data in order to form a mental model of that data – in other words, to **visualize** it. Therefore, aspects of human perception and cognition are introduced as and when appropriate.

Following the three main chapters identified above we conclude with a chapter on design, for one simple reason. An introductory course in information visualization is enhanced immensely by carrying out a design, preferably in collaboration with two or three other students. Therefore, Chap. 6 (**Design**) offers advice to project groups undertaking a short (e.g., 5 day) commission of the sort that might be provided by industry. And an appendix contains brief descriptions of the short **videos** that illustrate many of the concepts and techniques introduced in this book.

Exercises

Exercise 2.1

What are the drawbacks and advantages of the bargram representation of data?

Exercise 2.2

A car is described by a number of attributes: price, appearance, make, recommendations, horse-power, year of manufacture and age. Say which of these attributes are numerical, ordinal or categorical.

Exercise 2.3

Based on your experience of buying a mobile phone, washing machine, bicycle, car or other multi-attribute object, express your view as to the meaning of the term 'overview' and 'detail' and give some examples.

Exercise 2.4

For which of the attributes mentioned in Exercise 2.2 would it be possible to order the objects above the corresponding bar of a bargram?

References

Cleveland S, Becker RA, Weil G (1988) The use of brushing and rotation for data analysis. In: Proceedings IASC world conference on computational statistics and data analysis, International Statistical Institute, Voorburg, pp 114–147

Havre S, Hetzler E, Whitney P, Nowell L (2002) ThemeRiver: visualizing thematic changes in large document collections. IEEE Trans Vis Comput Graph 8(1):9–20

Hintze JL, Nelson RD (1998) Violin plots: a box plot-density trace synergism. Am Stat 52(2):181–184

Nielsen J, Tahir M (2002) Homepage usability. New Riders, Indianapolis

Schon DA (1983) The reflective practitioner: how professionals think in action. Basic Books, New York

Tversky B (1993) Cognitive maps, cognitive collages and spatial mental models, in 'Spatial information theory – a theoretical basis for GIS'. In: Proceedings of the European conference COSIT'93. Lecture notes on computer science 716. Springer, Berlin/Heidelberg, pp 14–24

Ware C (2012) Information visualization: perception for design, 3rd edn. Morgan Kaufman, Amsterdam

Wittenburg K, Lanning T, Heinrichs M, Stanton M (2001) Parallel bargrams for consumer-based information exploration and choice. In: ACM, proceedings of UIST'01, pp 51–60