

Internal Models, their Formation and Interpretation


6.1 The need for internal models

In Chapter 1 we adopted the definition that information visualization is the formation of an internal (mental) model of some data, and suggested that the model is created through sight of some display of that data. The very good reason for seeking an internal model is to have a better understanding of the artefact, scheme or situation to which the data refers, and to be able to interpret that model (Figure 6.1) in some useful way, perhaps to make a decision. Simple examples of such models, the sort of interpretation that may follow and the location in this book where the example is discussed, include:

- The engineering designer who observes the value of some property X of an artefact as the value of some other property Y is manually or automatically

FIGURE 6.1

The interpretation of an internal model



varied, and forms an internal model of the relation between *X* and *Y*. The interpretation may be that there appears to be some trade-off between the two properties (Chapter 9).

A house-buyer who discovers that a readiness to pay more does not make more houses available. The interpretation may be that the discovery of an affordable house needs the relaxation of one or more other limits (Chapter 5).

- Certain crops respond well to a particular treatment so further study to investigate the cost of treatments is desirable (Chapter 2).
- As well as the route I usually take, there is also a Circle Line Underground route direct from South Kensington to Liverpool Street station. The model lacks a comparison of journey times.

Since internal models appear to be so closely associated with information visualization we examine, in this chapter, the nature of those models, their interpretation and how they are formed. In doing so we move some way towards identifying implicit guidelines to support the designer of visualization tools.

Before going into detail it should be remarked that even though billions of words have been written about mental processes, there is even now a lack of consensus among professionals regarding the concepts involved. **Gibson (1997)**, for example, in his book *Theories of Visual Perception* remarks that a more appropriate title might be 'Some theories of visual perception'. It is nevertheless useful, in the limited space of this chapter, to introduce some generally accepted concepts related to internal models and which, together with some illustrative examples, will hopefully provide some basis for productive thought about the design of effective visualization tools.

6.2 Navigation

Our discussion will be based on a framework (Figure 6.2) for the human activity we call navigation, one component of which is the user's internal model which is of such importance to information visualization. We begin by briefly explaining the framework in outline (Spence, 1999) and then examining each component




FIGURE 62
A framework
for navigation

in detail, in all cases using illustrative examples drawn from earlier chapters. The framework will be seen to describe the process of navigation of information space where, initially we broadly define navigation as the creation and interpretation of an internal mental model.

The activity¹ identified in Figure 6.2 as browsing occurs, for example, when a user scans a display to 'see what's there'. They may be looking at a **Spotfire** display of pollutants (Chapter 5) in order to see what pattern emerges. The **modeling** of that pattern results in the internal model (or 'cognitive map'): in the same example it could be the association of pollution with some geographic feature. Interpretation of that model follows; perhaps the likelihood that other pollutants may have the same source. Next, based on that interpretation, the user can formulate a new browsing strategy; that could be a decision to display two particular pollutants together to see, again by browsing, if some correlation exists. The entire framework is that of navigation in the sense that the internal model is being created and explored. The overall intention or goal which requires navigation to occur may be a general desire to gain insight into (that is, to visualize) the origin of pollution within Sweden.

Since the framework of Figure 6.2 currently refers only to internal (mental) processes **we shall** gradually extend it to include essential external features such as the externalization of the data that the user is browsing. First, however, we examine some properties of internal models.

6.3 Internal models

6.3.1 Cognitive maps

We often **refer** to an internal model as a cognitive map. Tversky (1993) points out that '**as** mental constructs available to mental inspection, cognitive maps are presumed to be like real maps available to real inspection'. For example, having worked in London for many years I possess a reasonable if limited internal model of the London Underground railway system. For short journeys from Imperial College (close to South Kensington station) I need not look at a printed map, and only miss my destination if reading a particularly gripping book. But for less familiar journeys I need, in order to enhance my confidence, to glance at the printed map (the so-called externalization of the relevant data) before starting my journey. In so doing **I'm** initially and briefly referring to my internal model to identify the parts that need clarifying but then, by looking at the map, I can extend my internal model sufficiently to be able to undertake the journey. I might also choose, in fact, to create a new but related model which is simpler to remember – at least temporarily – for the particular journey: it might be of the form 'east on the Piccadilly, then east on the red line'. Although much remains to be discovered about memory (Solso, 1998), sufficient is already known with some degree of confidence for a simple account relevant to information visualization to be presented here (Norman, 1983; Waugh and Norman, 1965).

¹ In Figure 6.2 and throughout the chapter, activities are coded blue, their results (which can form the argument of another activity) are coded light purple, while the externalizations (that is, displays) of data are coded yellow.

6.3.2 Cognitive collage

In many situations it is more realistic to speak of a collection of cognitive maps or, as Tversky (1993) calls it, a cognitive collage. Thus, as a result of my journeys on the London Underground, I can recall in detail a number of separate segments, but have difficulty 'joining them up'⁷. Not surprisingly, the existence of a collage can lead to misinterpretations. A familiar illustration of errors that can emerge when accessing two items within a collage relates to the following question:

'Which city is farther west, Reno or Los Angeles?'

A common answer is 'Los Angeles', because that city is associated with California and Reno is associated with Nevada, which is east of California. However, reference to a real map shows (Figure 6.3) that Reno is west of Los Angeles. Association of Reno with Nevada and Los Angeles with California can be thought of as one item in memory, and the relative locations of Nevada and California as another, and the danger lies in the simple combination of these two items.



FIGURE 6.3
A map showing the states of California and Nevada, and the cities of Los Angeles and Reno


6.3.3 Incompleteness of a mental model

The fact that an internal model is not a single homogeneous entity has already been stressed by reference to the concept of a cognitive collage. In addition, the possibility of it being incomplete must also be appreciated. For example, our internal model of the shape of different countries is of little help in recognition if the representation of that country is tilted through an angle, if the relative size of the country is not shown, and if context is removed: try identifying the countries whose shapes appear in Figure 6.4! The problem here is that each country has been seen on so many occasions in an 'upright' position and in proximity to other familiar countries.²


² The countries of Figure 6.4 are, from left to right, British Isles, France, Italy, India and Africa.

FIGURE 6.4


The relative size and orientation of some countries have been altered – identify them!

**6.3.4 Resistance to change**

In the event that I have to consult the London Underground map to enhance my mental model of the part I plan to traverse I would expect to see, pasted to a wall of some station, the familiar map (Figure 6.5). However if, overnight (Figure 6.6), someone had removed all the Underground maps and replaced them with new maps printed upside-down, but with the lettering upside-up, then my browsing and interpretation of the (perfectly valid) map would be far more difficult, simply because of the challenge of simultaneously understanding two maps, one internal and the other external. I would also worry that the 'old map' would never be printed again, thereby requiring me to begin construction of a new internal model and, **simultaneously, suppression** of the old one. This situation may appear farcical, but the electronic equivalent often occurs on a computer screen where externalizations are so easy to modify.

**FIGURE 6-5**

The London Underground map

**FIGURE 6.6**

An inverted London Underground map

6.3.5 Inertial effects

Sometimes an externalization of a given set of data **has to change** in order for some insight to be gained. An example (Robertson et al., 1991) is provided by the so-called Cone-Tree representation of hierarchical data, to be discussed in more detail in Chapter 8. At one instant of time the tree may appear as shown in Figure 6.7, perhaps to allow a user to observe the relation between three people (including close colleagues) in an organization. A later need to bring different data to the fore would require rotation of one or more cones (Figure 6.8). If the change is abrupt (in other words, if Figure 6.7 is immediately replaced by Figure 6.8) then, although the newly required data may easily be visible, one's internal model of the rest of the cone tree may well be destroyed or, at best, not rehearsed.

To overcome this potential disadvantage, the inventors of the Cone-Tree arranged for the transition to proceed smoothly over a period of about one




FIGURE 6.7
A Cone-Tree representation of hierarchical data: a company organization chart




FIGURE 6.8
A new presentation of the Cone-Tree of Figure 6.7, corresponding to a new focus

second. In this way the cone rotations are sufficiently slow that the user does not lose awareness of the spatial form of the model and have to construct a new one: the user is being shown how the two states relate to each other. In the video presentation of the Cone-Tree concept it is pointed out that '... the perceptual phenomenon of object constancy enables the user to track substructure relations without thinking about it'. As we shall see in Chapter 8, the user is additionally helped by sight of the 'shadow' cast by the cones on a base plane, which in the illustrations of Figures 6.7 and 6.8, would be vertical and to the right of the tree.

6.3.6 Model robustness

When we say we are familiar with something we are in effect saying that we have a well-established and robust internal model of that thing. Such a property can

be exploited. One example occurred in Chapter 4 where magnification was the encoding mechanism: one's approximate feeling for the size of various countries and their relative spatial layout is such that magnification can provide a reasonable qualitative indication of some property associated with each country.


6.4 Model formation

We have already seen, in earlier chapters, the many circumstances in which an internal model can be formed, and the various activities that can assist that formation; the selection of a new price limit is an example (Chapter 5). It is worth subjecting these activities to a more detailed study in order to draw more generic conclusions.

At the most primitive level in information visualization, the user typically views a collection of items on a display screen. As the user's eye moves over the display, those items in foveal vision (Anstis, 1967; Bouma, 1970) are registered with more precision than those in peripheral vision: there is a visual acuity gradient moving out from the fovea. We shall refer to this activity as browsing, defined as the registration (or elicitation or assessment) of content. As an example, the tourist visiting an unfamiliar town and consulting a newly acquired map may almost immediately notice a green area, a blue area, an area of narrow winding streets, a factory area and a railway line, though initially without forming an internal model. Similarly the diner visiting a new restaurant will scan the menu to see what dishes are available and how the menu is laid out. The term 'scan-browse' has been used by Carmel et al. (1992) to describe this activity.

The content acquired by browsing is, probably almost immediately, integrated (Figure 6.9) in some way to begin forming a mental map, a process referred to as 'review-browse' by Carmel et al. (1992) and as **modeling** in this book. Thus, the tourist will note that the narrow streets are close to the harbor, and that the railway line passes close to their hotel. Similarly the diner will have quickly formed an impression of the standard of cuisine, the range of choice, the general price range and the likelihood of finding something attractive to their palate.

FIGURE 6.9
Browsing of
externalized
data leads to an
Internal model



6.4.1 Externalization

The interaction designer must, of course, arrange for relevant data to be externalized so that it can be browsed (Figure 6.9). While such externalization is normally in graphic form on a computer display, it can also be visual, aural, tactile or olfactory. In the physical world it could be all four, for example, during a

medical examination when the physician views a diseased limb, manipulates it, listens to it and is sensitive to any **odour**, all in order to form an internal model of the underlying problem. For the moment we concentrate on the more common graphic externalization while acknowledging that other sensory mechanisms might be able to provide additional or alternative beneficial capacity.

6.4.2 Perception and cognition

A familiar model provided by cognitive scientists, and relevant to the process of browsing leading to the formation of an internal model is shown in Figure 6.10. The process termed 'perception' ends with a representation in **memory**, but as yet no meaning. The representation is then subjected to the process of cognition (Humphreys and Bruce, 1989) in which a model representing some understanding is established in **memory**. The processes of perception and cognition are, in fact, extremely complex and not fully understood. We shall therefore work with the navigation framework of Figure 6.2 and use the verb 'model' to describe that part of the cognition process that takes the result of browsing and forms it into an internal model or an addition to an existing model. We already **have**, of course, many internal models, and some of these may be relevant to the task currently being undertaken, so that what is labeled 'model' in Figure 6.9 should be placed in the context of other internal models (Figure 6.11) and a return arrow added to indicate that the formation of an internal model may well require reference to an already existing one.




FIGURE 6.10
The processes of perception and cognition




FIGURE 6.11
A model relevant to a particular task may be formed by reference to an existing model and may then form part of it

6.4.3 Weighted browsing

For the purpose of describing the browsing and modeling process in a way that might help our understanding of information visualization, we introduce the concept of **weighted** browsing. In a typical browsing situation a user – for example the tourist or diner referred to above – will not assess all the content, partly because there is so much of it and partly because **much might** be suspected to be of little or no interest. Thus, when handed the menu in an unfamiliar Italian restaurant I initially put low weighting on (in other words, I have little interest in) the soups and pizzas, I put average weighting on pasta dishes and a very high weighting on scallopi marsala and chocolate profiteroles. The internal model I


first create by such 'selective attention' will thereby contain little or no information about pizzas, while the presence of scallopini and profiteroles will certainly be highlighted.

6.4.4 Externalization for browsing

For a user to browse and gradually build an internal model there must obviously be something there to browse! The question therefore arises as to how data should be presented on the display screen to ensure effective browsing (Figure 6.12). We begin with two contrasting illustrations regarding externalization design.

FIGURE 6.12


Externalization of the data that is to be browsed



The display of Figure 6.13, making sparse use of display space, offers an encyclopaedia user a choice of Art, Science or Religion, the selection of one of these options taking the user to the corresponding subordinate menu. These three options can certainly be browsed, but only allow a very simple model to be created. In fact, their main purpose may legitimately have been to allow the user to reach a particular target word as conveniently and quickly as possible (Norman, 1991) and with no intention that the user should build an internal model, of something as large as an encyclopaedia, to facilitate later use.

FIGURE 6.13

A menu system which does not support extensive browsing, possibly intentionally so



By contrast, the opening scene (Figure 6.14) of the CDi title Richard Scarry's *Busiest Neighborhood Disc Ever* makes full use of the available screen space. Moreover, in the opening scene, what is displayed changes automatically

and smoothly, in a 'panning' mode. As a result, after a period of about ten seconds, the user has been able, without performing any control action at all, to browse an area about five times the size of the screen. As a consequence, that user has been enabled to build a reasonable internal model of the content sufficient to decide what next to explore in some detail.





FIGURE 6.14
A presentation offering considerable opportunity for browsing and the formation of an internal model

Source: Reproduced by permission of Philips Media

6.4.5 Influence as content

The definition of browsing is not limited to the viewing of a fixed and passively displayed object. A particularly valuable form of browsing, encountered in the dynamic queries examples in Chapter 5, takes the form of registering cause and effect – a sort of 'ping-pong' browsing often associated with exploration stimulated by a 'what if?' question (Figure 6.15). Thus, in the house-seeking example, movement of an upper price limit ('ping'), together with the resulting effect ('pong') constitutes a sample of some relation, and therefore is a form of content which is browsed and may then be incorporated into the internal model.

FIGURE 6.1!**'Ping-pong'**
browsing

6.4.6 Short-term memory


One feature of the human cognitive system that constrains activities such as **modeling** is short-term memory (STM). Although controversy still surrounds this concept it is known that STM has a very limited storage and processing capacity, but the possibility of a trade-off between the two. That STM is not only severely limited, but requires **rehearsal** to maintain its content, was demonstrated in a classic experiment by **Peterson and Peterson (1959)**. In their experiment subjects were read a three-letter cluster and asked to recall it after varying periods of time during which subjects had to count backwards by threes from a three-digit number presented immediately after the three-letter cluster:

Experimenter: CHJ 506
Subject: 506, 503, 500, 497 . . . (until time to repeat the three-letter cluster).

The experimental results, shown in Figure 6.16, show that recall is seriously eroded in the absence of rehearsal.

6.4.7 The magic number seven plus or minus two

How extensive is STM? Short term memory is generally accepted to be able to hold around seven 'units of information', hence the title of George Miller's (1956) paper, 'The magical number seven, plus or minus two: Some limits on our capacity for processing information'. Thus, if asked to remember and then

**FIGURE 6.16**

Recall as a function of recall interval, when rehearsal was prevented

recall a string of letters, it would be difficult to do so without significant error if more than seven letters were involved. However, if the letters happened to be the following sequence:

I B M B B C C I A F B I S O S

then it is likely that more than seven letters would be recalled, but only because the subject would have been able to 'chunk' each group of three in view of their familiarity. Indeed, it is widely accepted that a subject would only be able to remember around seven 'chunks' of information.


How is STM relevant to information visualization? The basis for an illustration has already been provided in Chapter 5 by both the Dynamic Queries tool and the Attribute Explorer. As the use of either of these tools proceeds, new images containing potentially useful information are examined and, at the same time, new questions formulated. Thus, in the course of searching for a house to buy, not all presentations will have been remembered.³ Since some intermediate results may later be thought to have been significant ('maybe we could go for that larger house – remember? – and take in paying guests'), but might be difficult to recall in detail, it could be useful to provide some means – such as a simple 'tagging' mechanism – whereby potentially interesting intermediate situations could be recorded for later retrieval and re-examination.

The activity of electronic circuit design provides another example. In Chapter 2 it was shown how data concerning the effect, on the performance of an electronic circuit, of small changes in individual component values could be encoded by the size of circles superimposed on the corresponding component symbols (Figure 6.17). Moreover, animation was employed to show the manner in which these circles vary in size as the frequency varies over the range from bass to treble and back again. What we know about STM suggests that, if there were three important and broadly distinguishable regions of frequency (for example, bass, mid and treble) it might just be possible to observe and recall the nature of size changes for two or three components.

Of possibly equal importance, neither will the questions which ultimately led to the formation of these presentations.

FIGURE 6.17

With the animated circles it may be possible to recall, in outline, the size changes occurring at bass, mid and treble frequencies for around three components




6.5 Model interpretation

The execution of any task involving information visualization will be motivated by the user's intention and influenced by many factors. One of these is the user's internal model. Another is the visible externalization of some data. A decision as to how – as well as whether – to proceed will depend upon an interpretation of these sources of information (Figure 6.18). Once an interpretation is available a decision can be made as to how and whether further browsing should proceed, or whether the original overall intention should be modified. We begin with an example from Chapter 5.

FIGURE 6.18

Interpretation of the internal model and externalized data



The house-seeker may well have visited a number of houses and have built an internal model of certain features: the presence of nice looking schools near Pearson's Park, the worrying occurrence of graffiti near the otherwise attractive lakeside cottages and the steep rise in house prices as one moves towards the center of Thurstonbury. The house-seeker may well have jotted down some of these observations as a memory aid. They may also be viewing an externalization of some data – a map, perhaps, or a histogram, or some images of houses that have been tagged as being affordable and potentially satisfactory. On the basis of all this information, partly internal and partly external, the user's interpretation may be that no further exploration is necessary, since internal and external models are sufficiently detailed that a visit to an estate agent is now appropriate to obtain more detailed information regarding the candidate houses. On the other hand an interpretation may indicate the likelihood of real bargains along the Scottswood Road, in which case an appropriate display to permit

browsing in that area must be obtained. In that sense the term 'browsing' includes 'locomotion' to a new area. That locomotion may be physical (as in a car) or virtual, the latter possibly involving only an eye movement.

6.5.1 Interpretation of externalized data

We begin by selecting three illustrations of how externalized data can be interpreted. The first (Figure 6.19) is a particular presentation ('bifocal' or 'fisheye' – see Chapter 7) of data in which parts are 'bent backwards' so as to allow everything to fit into the display area.




FIGURE 6.19 Principle of the bifocal display
(a) Information space, much larger than the display
(b) The same space 'wrapped around' two uprights and viewed as shown
(c) The view seen by a user, and scrollable as suggested in (b)

Since items in the peripheral areas would be unreadable because they are 'squashed', the text and similar detail is automatically removed;⁴ nevertheless, the colour of each item is discernable. Thus, a red rectangle towards the left edge of the information space denotes a letter from my boss, an item that can be scrolled into the central region to expand horizontally and become readable. Here my interpretation is that there is an urgent letter presently located in iconic form towards the left of the display. Thus, an important property of the display is that it provides a measure of *sensitivity*: I can easily form some idea as to the extent to which I must scroll the information space to achieve my goal of reading the letter from my boss.

A second example is provided by the black houses in Figure 5.14: they provide some measure of the direction as well as the distance one has to move a particular limit in order to acquire the corresponding data. They provide a measure of sensitivity, or a 'gradient' with respect to houses that fail only one limit. A similar measure of sensitivity is provided by the tight coupling concept (Ahlberg, 1996) illustrated in Figure 5.6.


A third example, involving the aural presentation of information, was provided by the Media Room (Bolt, 1979). Here (Figure 6.20), a user panning across an image of the MIT campus on a screen measuring 11 feet by 8 feet becomes

⁴ See the discussion of semantic zooming in Chapter 7, Section 7.8.1

FIGURE 6.20

The Media Room at MIT

Source: MIT




aware, on hearing the (stereo) roar of a football crowd, that there is a large and well-packed football stadium 'somewhere off the left edge of the screen'. Again the act of interpretation could be classified as one of 'gradient perception' or an awareness of sensitivity.


Although the concepts of 'gradient perception' and sensitivity can be useful, and may even be widely applicable, they will in many cases be an oversimplification of the manner in which data is interpreted. The topic is complicated by the fact that interpretations may be formed consciously or unconsciously; may be based on part or the whole of the externalized and internalized data; and may well be influenced by existing domain knowledge possessed by the user. Also, as we shall see in the next section, the interpretation could well be influenced by knowledge of the available browsing mechanisms that can be harnessed to exploit the interpretation.

6.5.2 Model adequacy


One aspect of interpretation that must always be borne in mind is the *adequacy* of a model. This property is perhaps best illustrated by analogy with the process of fitting a mathematical model to some available quantitative data. For example, if the data about how some property *X* varies as some parameter *Y* is changed is contained within five data points (Figure 6.21), it could be dangerous to base any interpretation on a smooth curve (Figure 6.22) fitted to the data, since the *actual* variation of *X* with *Y* might be as shown in Figure 6.23. It is in situations like this that domain knowledge and experience can play an important part and, if appro-

**FIGURE 6.21**

Data points

**FIGURE 6.22**

A possible model?

**FIGURE 6.23**

The actual relation

priate, help identify the fact that a model is impoverished in some respect. If a model is impoverished in any way it may be difficult or impossible to perceive any valid gradient or sensitivity indication or make any useful interpretation. Such a situation might obtain, for example, after examining only four or five houses in the house-seeking scenario of Chapter 5, so that questions such as 'can we increase the number of bedrooms without increasing the price?' simply cannot be answered with any useful degree of confidence. In this case the interpretation would be to the effect that further content should be browsed.

6.6 The formulation of browsing strategy

Until now we have said little about browsing strategies and how they are formulated. It is unfortunate that definitions of browsing offered in the literature sometimes imply a rather unstructured and random activity, perhaps with serendipitous intent, and even with a hint that it might not have been consciously planned. That is one possible browsing strategy, but only one among many. When I pick up my copy of *The Times* newspaper I know where the weather report is, on which page the obituaries are to be found, where the 'Comment' is and where, on a Saturday, the humorous 'Letters to the Editor' are positioned, and I habitually register their content in that order. My planned but unconscious browsing is brisk, is essentially asking 'what's there?' and is followed by a decision about what I want first to read in detail.


Nevertheless, though my browsing is planned, that plan may suddenly be abandoned, and my activity become opportunistic, upon noticing the obituary of my old teacher. But other, entirely different strategies are possible, perhaps using a different source of data; one might, for example, read the 'Summary' provided on the back page of *The Times* instead. The question is, how is my strategy, both global and detailed, determined, and how can the externalization of data best be organized to help decide upon the browsing action (Figure 6.24)? And for interactive (as opposed to passive) displays, how can the interaction best be designed to support the formulation of a browsing strategy?

6.6.1 Cognitive determinant

There are two determinants of the formulation of browsing strategy (Tweedie, 1995). One is cognitive, based either upon the interpretation I have made or as a result of a new idea (Figure 6.25), but not directly influenced by what is

FIGURE 6.24


The formulation of a browsing strategy, partly influenced by the externalization of data



displayed. The interpretation may be to the effect that navigation should be terminated because, for example, a small group of houses worthy of more detailed consideration has been identified. Alternatively it may be to the effect that the original intention is not yet satisfied and that a particular exploration – say, the variation of an upper limit on price – is now needed. Such a conscious and cognitively planned browsing strategy could then be formulated. I may, on the other hand, suddenly decide to adopt a new course of action to see if it proves helpful, that of checking an estate agent's brochure: a cognitively initiated opportunistic strategy.

FIGURE 6.25

Cognitive and perceptual determinants of planned and opportunistic strategies



6.6.2 Perceptual determinant

The other determinant is perceptual. In other words, the strategy I formulate is influenced directly by what I see displayed on the screen. Consider, for example, the screen display of Figure 6.26, one not very different from the Attribute Explorer encountered in Chapter 5. The key to various user actions that are now possible lies in the controls, associated with various affordances, that are provided in the display, and which can be used to rearrange that display as discussed in Chapter 2. These controls are briefly described by annotations outside the screen area in Figure 6.26.

Other features of the display which can influence my actions include, for example, the black houses, the general distributions of the histograms and the 'ineffective range of limit' (Figure 5.6) indicators. Yet another feature of the display is associated with houses thought to be sufficiently interesting at the time they were first encountered to be recorded for later reconsideration. Sight of those cir-




FIGURE 6.26

An application of the Attribute Explorer, illustrating a rich collection of controls and perceptual triggers as well as data

cles might prompt me to decide that my next browsing strategy will be to re-examine one of those records: a planned action stimulated perceptually. Similarly I may notice that other tools are available that I have not used before: an example is the Magic Lens (described later in Chapter 7) which, like a magnifying glass, can be moved over a component of the histogram to reveal aspects of the underlying data that do not appear within the Attribute Explorer, aspects such as the appearance of an individual house or a map containing locations of interest. On the other hand, sight of the display may lead to a perceptually initiated opportunistic action – also called a situated action (Suchman, 1987) – such as a shift of focus to an area near the factory where a very cheap but large house has been noticed. As pointed out before, the new browsing strategy can include a shift of location.

Whether prompted cognitively or perceptually, the execution of the selected browsing strategy may lead either to the enhancement of a currently considered internal model or the initiation of another internal model, part of a collage (Tversky, 1993).

6.7 Comment

Based on the above discussion, navigation might more fully be defined as **cognitively directed movement** in information (or physical) space based on the interpretation of a mental **model and/or** externalized data. Although one aim of the navigation framework is to provide a basis for organized thought about information visualization it **does** not provide precise rules for the designer of a visualization tool to follow. The design of such a tool – like the design of any worthwhile artefact – is an exceedingly complex (and little understood) process. Nevertheless the framework can provide a useful reference during the **course** of visualization tool design, and is frequently referred to in later chapters.