




Presentation

7.1 A problem

Many of us have found ourselves with a report that has to be completed by a deadline, with the result (Figure 7.1) that the dining room table, extended to its twelve-guest state, is covered by piles of paper as well as reports, books, clippings and slides; perhaps with more arranged on the floor and on a couple of chairs. There may even be piles on top of piles. Such a presentation of vital information makes a lot of sense: everything relevant is to hand (we hope!) and, moreover, its very visibility acts as a reminder (Bolt, 1984) of what might be relevant at any particular juncture, possibly triggering a situated action (Suchman, 1987). In this environment I can concentrate on creative tasks rather than organization.

Despite the availability of high-resolution displays and powerful workstations I still write most of my reports in this way. Why? Because the display area provided by the typical workstation is far too small to support, visibly, all the sources that are relevant to my composition.

7.2 The presentation problem

I am not alone in the sense of having too much data to fit onto a small screen. A very large and expensive screen, for example, would be needed to display the London Underground map in sufficient detail (Figure 1.1), and it would be difficult or impossible to present, on a normal display, the complete organization

FIGURE 7.1

Support for
report
preparation!



chart of IBM or ICI. Moreover, the recent emergence of small and mobile information and communication devices such as PDAs and wearable displays has additionally identified a pressing need for a solution to the 'too much data, too little display area' problem: the presentation problem. How can it be solved, mindful of the need to support the activity of visualizing the underlying data?

7.2.1 Scrolling

An obvious solution is to scroll the data into and out of the visible area. In other words, to provide a means whereby a long document can be moved past a window until it reaches the required 'page' (Figure 7.2). This mechanism is widely used, but carries with it many penalties. One relates to the 'Where am I?' problem: I'm working on Chapter 2, (it may be Section 2.3, I'm not sure) and I want to remind myself of a figure that is in Chapter 5, it may be in Section 5.3 – or was it 5.6? All I can do is operate the scrolling mechanism and look out for the figure I need, albeit assisted by various cues such as the page number indicated in the scrolling mechanism. With a scrolling mechanism, most of a document is hidden from view. I have the same problem when using a microfilm reader. A similar difficulty applies to my use of the famous London 'A to Z' street directory. I'm driving along a road that goes off the edge of the page, so I desperately need whatever page contains the continuation of that road (and quickly!). Even if I get it, I will typically have trouble locating the same road on the new page. These and other similar problems can be ameliorated by the provision of context. Much of this chapter, in fact, is concerned with deciding how to provide context.

7.2.2 Context map

Some drawbacks of the scrolling technique can be ameliorated by the provision of a map with the location of the detail of interest indicated. Thus (Figure 7.3), a detailed display of part of a map can be placed near a global, but necessarily less detailed map, with the location of the detailed map shown. This may be an improvement on the situation where there is no context at all, but there is now a continuity problem: it is difficult to transfer mentally from the edge of the detailed region to the corresponding point in the global map. You may say, 'Why not bring the two together very simply by using a magnifying glass technique?' A major disadvantage immediately obvious from Figure 7.4 is that the magnifying glass occludes the immediate context of the detailed region.





FIGURE 7.2
Scrolling hides most of a document

The problem of occlusion was addressed by Ware and Lewis (1995) who proposed the DragMag Image Magnifier. The approach adopted (Figure 7.5) is to avoid occlusion by allowing the separation of the (small) region of interest (in the red square) and the resulting magnified image. The region to be magnified, shown with lines connecting it to the magnified image, can itself be panned around the 'base image' independently of the location of the magnified image, which can itself be independently placed. Controls beside the magnified image allow zooming. Nevertheless, the continuity problem still remains.


Another interesting solution to the problem of occlusion is to make the detailed map transparent (except for essential features such as roads and place

FIGURE 7.3

A separate display of context

**FIGURE 7.4**

Conventional magnification masks detail



names), as in Lieberman's (1994, 1997) *Macroscope*, shown in Figure 7.6. After the area to be magnified is identified (by the red rectangle) it is effectively enlarged to the same size as the original map and then made available for *simultaneous* display. Indeed, a manual control can vary the relative prominence of each map, even to the extremes where either the original map or the magnification of the region of interest is visible on its own.






FIGURE 7.5
The DragMag
image magnifier




(a)



(b)



(c)



(d)


FIGURE 7.6
Within a large
area (a), detail
within a small
region (b) is of
interest. The
small region is
magnified (c) to
reveal detail,
and then
superimposed
(d) upon the
original map

7.3 Focus⁺Context

Not surprisingly in view of the importance of context, a great deal of attention has been directed during the last two decades to the solution of the problem created by the need to have context information beneficially co-existing with detail of the focus of attention. One of the most successful concepts, and one which has led to numerous derivative techniques, is the Bifocal Display (Imperial College Television Studio, 1980; Spence and Apperley, 1982).

The principle of the bifocal display is perhaps most easily demonstrated with the physical analogy of a long strip of paper (Figure 7.7) showing a representation of my in-tray: a variety of symbols represent incoming letters, telephone calls, emails and documents I'm preparing. However, the items are so numerous that this information space cannot be viewed through the normal aperture of a display screen: wherever an aperture is positioned, most of the information space – and hence most of the content (and therefore context) – is hidden from view.

FIGURE 7.7
Representation
of an in-tray



The solution is simple (Figure 7.8): imagine the strip of paper to be pulled back across two posts, but *in such a way that all of it is still in view*. One or two items will appear in the center and, if they are documents, will be readable, but the viewer is still aware of the (albeit distorted) presence of all other items. An electronic implementation of the bifocal concept might appear as sketched in Figure 7.9. This is where the skill of the interaction designer is relevant: if the encoding of the distorted items is carefully chosen, then an informative 'overview' of information space is immediately available. For example, colour might be used to denote the sender of an item, 'X-position' to denote time, 'Y-position' to denote type of communication, and shape to represent extent. For example, the thin red rectangle in the left-most distorted region is a message from my boss: it is readily noticed, and can easily be read if the information space is scrolled until the red rectangle has expanded upon entering the central region. Suitably implemented, a bifocal display might appear as in Figure 7.10.




FIGURE 7.8
The principle of the Bifocal Display






FIGURE 7.9
The Bifocal Display as seen by a user

**FIGURE 7.10**

An early implementation of the Bifocal Display

The bifocal concept was invented long before a satisfactory implementation could be achieved, and was implemented as the Perspective Wall (Mackinlay *et al.*, 1991) eleven years later on an advanced workstation (Figure 7.11). The Perspective Wall simulates a three-dimensional effect, but is essentially demonstrating the same beneficial distortion effects as the Bifocal Display. Other terms associated with the concept are the 'fish-eye lens' (see later, p. 120) and the 'distortion technique'.

**FIGURE 7.11**

The Perspective Wall

Source: Inxight Software Inc

As illustrated, the bifocal display is based on distortion in the X-dimension, shown schematically in Figure 7.12. But distortion can additionally be applied in the Y-dimension, as illustrated in Figure 7.13. An example of where this com-




FIGURE 7.12
Distortion in the X-dimension




FIGURE 7.13
Distortion in both X- and Y-dimensions

bined distortion might be useful is shown in Figure 7.14, which is a distorted version of the London Underground map. The advantage over a display such as Figure 7.3 is clear: there is a continuity across the boundaries where the distortion begins, allowing a smooth visual transition between focus and context. Thus, the traveler can easily trace a line out of the detailed region and into the compressed region or, if an appropriate interactive control is provided, may simply move the non-distorted region in order gradually to identify a complete journey. One should not overlook the fact that Harry Beck could be said to have

FIGURE 7.14
The London Underground map distorted in both X- and Y-dimensions

Source: Reproduced by kind permission of London Transport Museum



introduced distortion in his original map of the Underground, namely distortion of the underlying geography. Another example where concurrent X- and Y-distortion can be useful is the diary of Figure 7.15 (Spence and Apperley, 1982): horizontal scrolling anywhere on the screen will flip successive weeks into the central region, which *itself* is distorted so that ample space can be devoted, by vertical scrolling, to any selected day. The same principle was applied to the Document Lens by Robertson and Mackinlay (1993).

Mar	April	May	June	July			Aug	Sept	Oct
				11 Sun	Check slides, notes Family barbeque				
				12 Mon	Fly LA Kathy to airport Model Maker				
				13 Tue					
				14 Wed					
				15 Thur					
				16 Fri	Flight to SFO Tutorial set-up Tutorial United flight Heathrow Pointer Color OHs Jane+Jonh Call Kathy				
				17 Sat	Fly LHR Kathy to collect Chapter 2/see Dave March				

FIGURE 7.15


A diary exploiting both X- and Y-distortion

The control of a distortion-based display can be achieved in many ways. In the original proposed version of the Bifocal Display concept an item in the distorted region was moved into the central focus region by a simple scrolling action (for example, by moving a finger across the display): hand-eye coordination was sufficient to achieve the repositioning without difficulty. An alternative approach in which the required item is identified by mouse-click, whereupon it automatically moves to the focus region, requires more attention to detail.

The distortion principle inherent in the Bifocal Display concept has also been applied to tables, to allow fine detail to be examined. The Table Lens (Rao and Card, 1994) shown in Figure 7.16 will be familiar from Chapter 2 where the baseball data was seen to be more easily interpreted by interactive rearrangement of the rows. Now, with the Table Lens, detail about individual players can be examined in the context of all the data.

FIGURE 7.16
The Table Lens,
exploiting the
distortion
principle

Source, *Inxight
Software, Inc*



The principle of the Bifocal Display will be familiar to those who frequent poster shops, and especially to those who actually purchase illustrations with titles such as 'A New Yorker's View of the World'. Such a view acknowledges the fact that, to a New Yorker, the postbox, the subway entrance and the corner shop are of far greater day-to-day concern than the entire landmass of Australia, which is therefore given token display as a mound somewhere out to sea. Actually, this presentation involves not only distortion but also suppression, a technique examined later in this chapter.

The focus+context issue continues to stimulate new solutions (Holmquist, 1997) and applications (Yamaashi *et al.*, 1993) and, in fact, is liable to increase in importance with the proliferation of hand-held information and communication devices, where display area is at a premium. Figure 7.17 shows Holmquist's (1997) flip zoom technique implemented on a simulated palmtop display to function as a web browser (Bjork *et al.*, 1999b).

7.4 Suppression

7.4.1 Furnas' fisheye concept

The concept of suppression is perhaps best introduced by an example in which an aircraft maintenance technician is examining the diagram of a very complex mechanism, perhaps in order to rectify a fault (Mitta, 1990). It would not be unusual if the diagram contained far more detail than is relevant to the task at hand, in which case the technician's concentration on the problem could probably be enhanced if irrelevant detail were removed. Thus, the engineering drawing of Figure 7.18 (which is very simple as engineering drawings go!) might be more useful if temporarily simplified to the form shown in Figure 7.19 if some




FIGURE 7.17
The flip zoom
technique
embodied in a
palmtop display
Source: Lars
Holmquist




FIGURE 7.18
An engineering drawing
Source: Mitta (1990)




FIGURE 7.19
The engineering drawing simplified
Source: Mitta (1990)

system property having to do with the relation between the two screws was of prime concern. The problem is, how do we define relevance, and how do we arrange for irrelevant detail to be suppressed? There are, of course, many different aspects to the relation between two screws, such as their physical occupation of space, the torque transmitted and their shear forces.

A powerful technique for organizing the suppression of irrelevant data was suggested by Furnas (1981, 1986) and called the *fish-eye technique*, a term since used more frequently to describe the distortion techniques introduced earlier. To quote Furnas, the fundamental motivation for the strategy he proposed ‘... is to provide a balance of local detail and global context’. At first sight the motivation is identical with that which led to the bifocal display concept. The essence of Furnas’ proposal, however, is that it is often helpful to display data only if its perceived value exceeds some threshold set by the user.

The concept was originally illustrated by reference to a tree-structured hierarchical database in which the ‘degree of interest’ (or importance or priority or worth) of any node decreases by one the further one moves in the tree away from a node designated as the focus. Thus, if the root node is the focus, node weights are assigned as shown in Figure 7.20(a), decreasing by one with each link down the tree. The same principle applies if, say, one of the lowest nodes (which is still a root node of another tree) is designated as the focus of interest: the assignment of degrees of interest is then as shown in Figure 7.20(b). If, as with the two screws shown in Figure 7.18 and 7.19, the situation is such that both the nodes selected as foci in Figures 7.20(a) and 7.20(b) are concurrently selected as foci, the nodal degrees of interest are obtained by addition, as shown in Figure 7.20(c). If the threshold for display is now set at -4 , only the data associated with those nodes shown in black in Figure 7.20(c) is displayed.


FIGURE 7.2

Establishing a degree of interest

(a) Node weights when the designated root node is the focus

(b) Node weights when a leaf node is the focus

(c) Two nodes are foci concurrently



Furnas generalized the concept to handle multiple foci and provided a number of examples ranging from diaries to software. In Mitta’s work, shown earlier in Figures 7.18 and 7.19, Furnas’ concept was generalized to non-hierarchical information structures, as was the meaning of each arc in the linear graph representing the information structure.

7.4.2 Z-thru mapping

The Z-thru mapping technique invented by Colby and Scholl (1991) also exploits selective display. Here (Figure 7.21) a ‘back-plane’ – in the example a map of the Boston area – can be thought of as potentially being covered by a number of information ‘layers’, each associated with some feature of the Boston area. Features could include crime statistics, air corridors, weather, obstructions over 2000 feet and traffic density. The user is able to discern information of interest by controlling the focus and transparency of each ‘layer’. Here, rather than a degree of interest being assigned within a hierarchical structure, there is a menu of options which can be mixed according to a user’s need.




FIGURE 7.21
The Z-thru
mapping
technique


*Source: Colby and
Shall, Frames,
February 1992*

7.4.3 Mixed techniques


The presentation techniques described above can often be combined to good effect. One hypothetical example is the 'Really Useful Map' shown in Figure 7.22(a). It addresses the problem of traveling from a hotel in an unfamiliar city to a relative's house in another unfamiliar city. The presentation employs the distortion concept, using 'rubber sheet' distortion (Kadmon and Shlomi, 1978) as illustrated in Figure 7.22(b), to provide necessary fine detail to support navigation within the two cities. It also employs suppression to acknowledge that,

FIGURE 7.22

(a) A Really Useful Map employing distortion and suppression
 (b) Rubber sheet distortion



(a)



(b)

once the vehicle is on the motorway, the driver is best served by only the occasional landmark, the sight of which can instil confidence that the vehicle is on the correct route. In this way the map of Figure 7.22(a) combines both distortion and suppression with encoding (Spence, 1993).

To place the concepts discussed above in historical perspective it has to be pointed out that, once the concepts of distortion, suppression and encoding are understood, it is not difficult to identify such techniques in old maps. Thus, in the mid-thirteenth century the Benedictine monk Matthew Parish provided such an example in his map of the journey between London and Jerusalem (Barber and Board, 1993). Currently, attempts are being made to find a formal framework for focus+context tools (Bjork *et al.*, 1999a).

7.5 Magic lenses

Reflecting the message of Chapter 2, we can say that the process of presentation is that of arranging and rearranging the view that we have of data. So far, in all the examples in this chapter, we can observe that a visualization tool has been constructed and that the data that is subsequently rearranged has been brought to the tool. But why not bring the tool to the data? To illustrate this approach, and decide whether it offers benefits (Stone *et al.*, 1994), we first consider two examples drawn from creative writing and house-seeking.

In creative writing one often needs clarification of a word, either because someone else has used that word or because one has used the word oneself without being absolutely sure about its definition. An appropriate Magic Lens

(Figure 7.23) would first be moved from its resting position on the display to lie roughly over the sentence in which the word appears; the word itself would then be identified by means of a cursor; whereupon the lens would **resize** itself to provide sufficient room for the definition, and the definition then appears. Adequate discrimination must be achieved between the text and the definition, possibly by colour encoding supported by suitable positioning. In this way the tool (generically termed a magic lens) is brought to the data, the latter represented in this example by familiar text.




FIGURE 7.23

A Magic Lens employed to clarify the meaning of a word

Why is the magic lens concept a good idea? Most obvious is the attractive metaphor – that of bringing a magnifying glass to something in order to discern more detail. Furthermore, limiting the display of detail to a small region removes the danger of clutter over the whole display, and is less expensive computationally than if the function of the lens were to be applied globally. Indeed, leaving the area outside the lens unmodified provides useful context.

The concept of a magic lens is, in fact, quite general, since it is concerned with accessing the underlying data structures of an application and reformatting the data to generate a modified view. This property, could, for example, be usefully applied to the activity of house-seeking: a different lens, placed over part of a map, could show the images of a number of houses, all selected on the basis of limits set on price, number of bedrooms and other parameters: thus, a magic lens can be parameterized. If space is limited the house images could automatically and continuously move through the viewing area. The lens could alternatively be adjusted to identify the locations of schools and hotels.

A useful property that can be exhibited by magic lenses is that, if they overlap, their separate effects can be combined in appropriate order, thereby offering the possibility of visual macros.

7.6 What's there?

There are many situations in which browsing, as defined in Chapter 6, plays a particularly important part. One arises when a person returns home after a tiring day, sits in their favourite armchair and attempts to see if something interesting is offered by any of a very large number of television channels. How does one design a supportive display to satisfy that need, a need characterized by the question 'What's there?' A similar need relates to a user who has either forgotten what was put in a given folder, or has been given some folders and quickly wants to assess the nature of their content.


One answer is to allow the user to 'riffle' through content, just as one is able to pick up a potentially interesting book, riffle through its pages in a very short time (for example, three seconds) but sufficient to gain some idea of what the book contains. Two applications of the riffling technique are shown in Figures 7.24 and 7.25.

FIGURE 7.24

A video-on-demand system supporting the browsing of video posters



As illustrated in Figure 7.24, Lam and Spence (1997) reported a video-on-demand presentation to support the browsing of video posters. A large number of posters are arranged in three rows on a wall which allows some posters to be viewed 'full face' in the center, but also be 'pulled out' from the side walls, an arrangement very similar to the bifocal display described earlier in this chapter. Movement of the cursor along the edge of the side walls causes each poster to 'pop out' and be fully displayed for as long as the cursor remains at the position of that poster. Anecdotal evidence suggests that display rates of around five per second were satisfactory for the recognition of one or more posters of potential

**FIGURE 7.25**

The presentation of goods and services to potential customers

Source: AV/98


interest. In one sense the example illustrates a trade-off between space and time, since simultaneous full view of all the posters might require display space larger than the screen. Using a similar principle Wittenburg *et al.* (1998) have described an application (Figure 7.25) designed to present goods and services for purchase to potential customers. As time proceeds, images appear in sequence on the display.

A third application of the riffling technique (de Bruijn and Spence, 1999, 2000) is shown in Figure 7.26. A user has forgotten what is in the folder labeled 'new'. A mouse-click in the semicircular area triggers a rapid serial visual presentation (RSVP) of images of the content, the trajectory suggesting that the images are emerging from one side of the folder and then returning via the other. At a rate as high as about ten per second or more a user is able to gain some idea about content and perhaps recognize an image that is being sought. The black segment will completely fill the semicircle when all images have been displayed, thereby providing some idea of the time it will take to view the entire set of images. A second click will stop the riffling. Other controls (for example,

FIGURE 7.26

Rapid Serial
Visual
Presentation
(RSVP) of
images within a
folder

Source: GTE
Laboratories



increment/decrement by one) can be added according to the application. Rather than provide each folder explicitly with a mechanism to control the riffling, the riffling function might best be provided by a magic lens.

7.7 Effective view

It is sometimes the case that a constraint is imposed upon the presentation of data by some unrelated requirement, for example by the need of a user to rearrange the presentation in a manner reflecting 'movement' in some information space. As an illustrative example we consider the ordered list of items (fruit) shown in Figure 7.27(a), and assume that the user (who is aware of the alphabetic ordering) will wish to search for a particular item. Typically the list will be so long, or the display area so small, that the user will only be able to view part of the list: in our simple illustrative example the viewing window contains five items (Figure 7.27(b)). 'Movement' through the list towards a desired item is achieved by 'clicking' on an item currently within the viewing window. The number of choices available to the user is called the 'out degree' or 'view size'.

Unfortunately, with a small view size and a long list, the user will be faced with a large number of successive selections: thus, if the word 'quince' is sought, clicking on the bottom line within the window ('kale' in the example) causes that line to appear in the center of the window. Available jumps within the window are shown in Figure 7.27(c). Furnas (1997) therefore proposed a measure of the 'diameter' of the graph in terms of how many interactions (for example, clicks) are involved in moving from one end of the list to the other. For the scheme shown in Figure 7.27 the diameter is proportional to N , the number of items in the list. In the design of such a presentation one would aim to keep both the view size and the diameter reasonably small. For very long lists (for example, a town's telephone directory) the number of interactions would be prohibitive, which is why a scroll bar is so useful.




FIGURE 7.27
An effective
view of a list,
and movement
through it

A different view of items within the list can significantly increase the diameter, thereby reducing the number of interactions required to reach a given item. The scheme shown in Figure 7.28 resembles the bifocal display to some extent because, from a central item in the viewing window, the user can both see and click on items that are at a distance of 1, 2, 4, 8 and so on away in the list: it shows the 'local' part of the list in some detail, and further regions in successively less detail. With this view, movement from one item to another by clicking on a visible item is a process that looks much like binary search (Figure 7.28): indeed, the diameter turns out to be of the order of $\log N$, which is much less than N .





FIGURE 7.28
A rearrange-
ment of the
effective view
to facilitate
movement
through the list

Furnas presented other approaches, of which we shall examine one. If the original list is transformed into two dimensions (Figure 7.29) the view size can be small but show neighbours in both columns and rows. Now the diameter is of

FIGURE 7.29


Transformation
of the list into
two dimensions



the order of the square root of N . Taking this rearrangement one step further by rearranging (Figure 7.30) the list in three dimensions (much like a telephone directory), and assuming that the view from one page to the next can be supported (by alphabetic tags, for example) then the diameter falls to the order of the cube root of N .

FIGURE 7.30

Transformation
of the list into
three
dimensions



7.8 Zoom and pan

Television watchers and cinemagoers – in other words, virtually everyone – knows what panning and zooming are. Panning (Figure 7.31) is the smooth movement of a viewing frame over a two-dimensional image of greater size. Zooming (Figure 7.32) is the increasing magnification of a decreasing fraction (or vice versa) of a two-dimensional image under the constraint of a viewing frame of constant size. The choice of viewing location afforded by panning, and the ability, offered by zooming, to choose between local detail and global overview, are easily appreciated.

A more demanding requirement can sometimes arise. Suppose a user has completed a local examination of an object **A** within an image (Figure 7.33) and now wishes to similarly examine an object **B** within the same image but currently not




FIGURE 7.31
Panning is the smooth movement of a viewing frame over a 2D image




FIGURE 7.32
Zooming is the increasing magnification of a decreasing fraction of an image (or vice versa)




FIGURE 7.33
Transfer of the focus of attention from item **A** to item **B**

in view. A conventional pan would, at first sight, satisfy this requirement. However, recalling the discussion of internal models in Chapter 6, we would identify the benefit of first zooming out until both foci of interest (**A** and **B**) are in view, and then carrying out a short pan in such a way that a subsequent pure zoom will provide local detail of **B**. In this way the user's internal model of the global area within which **A** and **B** are located is either generated or refreshed. This and other processes concerned with scale can be explained by reference to the concept of space-scale diagrams (Furnas and Bederson, 1995) illustrated in Figure 7.34.





FIGURE 7.34
A space-scale diagram pertinent to Figure 7.33

A space–scale diagram (Figure 7.34) contains a number of copies of a rectangular two-dimensional image arranged vertically along a magnification scale, and stacked to form an inverted pyramid. If image magnification is proportional to vertical distance, it follows that a single point **A** (the black circle in the yellow area) in the image identifies a straight 'ray' passing through the origin. A similar ray tracks the point B, a black square within the green triangle. In the same space the user's viewing frame is represented by a square of constant size, irrespective of where it is placed in the diagram. Thus, horizontal movement of the viewing frame represents panning, while vertical movement describes zooming. For the process described above – a zoom out from **A** until B is within the frame, followed by a short pan and then zooming into B – the corresponding trajectory (W, X, Y, Z) of the viewing frame is shown in Figure 7.35. A single direct pan from W to Z can be disadvantageous, as already remarked, as far as the user's internal model is concerned: arguably, such a trajectory may also be computationally more expensive (Furnas and Bederson, 1995).

FIGURE 7.35

Trajectories in a space–scale diagram



7.8.1 Semantic zoom

Taking the zooming concept further leads to the notion of semantic zoom. With a conventional geometric zoom all objects change only their size: with semantic zoom they can additionally change shape or, indeed, their very presence in the display. As an illustration we select the Spatial Data Management System (Herot et al., 1981; Herot, 1980) which was installed in the *S.S. Carl Vinson*, and its application to naval operations. In certain circumstances a ship's captain needs to be aware of a collection of available ships (Figure 7.36(a)). In a different situation certain details of a particular ship may be of interest, in which case a




FIGURE 6.36
Semantic zoom

geometric zoom combined with a semantic zoom will provide overall details of that ship (Figure 7.36(b)). Further semantic zoom could display additional information, and perhaps replace the ship icon with a photographic rendition (Figure 7.36(c)). Semantic zoom can also be represented on a space–scale diagram (Furnas and Bederson, 1995) and has been incorporated in the Pad++ Zooming Graphical Interface (Bederson and Hollan, 1994).