Int. J. Human-Computer Studies (1999) **51,** 919–945
Article No. ijhc.1999.0265
Available online at http://www.idealibrary.com on IDEAL®

A framework for navigation

ROBERT SPENCE
Department of Electrical & Electronic Engineering, Imperial College of Science, Technology & Medicine, Exhibition Road, London, SW7 2AZ, UK.
email: r.spence@ic.ac.uk

(*Received 16 February 1998 and accepted in revised form 26 April 1999*)

A new schematic framework for navigation is presented which is relevant to physical, abstract and social environments. Navigation is defined as the creation and interpretation of an internal (mental) model, and its component activities are browsing, modelling, interpretation and the formulation of browsing strategy. The design of externalizations and interactions to support these activities, and navigation as a whole, is discussed.

1. The problem

Within the broad field of human–computer interaction, concepts such as navigation, search and browsing have long been employed in the discussion and design of information artifacts. Over the years, these concepts have often been discussed verbally, and with varying interpretations. But the concept of navigation, in particular, is now sufficiently important and wide-reaching as to warrant a more formal and precise definition; a definition, moreover, which allows the concept to be extended beneficially to a much wider range of environments. Also, and especially as the range and complexity of applications has widened, it has been recognized that the issue of interaction design to support navigation is increasingly vital and challenging. It is against this background that we propose, in this paper, a *schematic framework for navigation* applicable to a wide range of *environments* and having particular relevance to *interaction design.*

2. Scope

2.1. PHYSICAL ENVIRONMENTS

For hundreds of years the consideration of navigation and related concepts such as landmarks, beacons and tracks took place in the context of real *physical environments* such as cities, landscapes and continents (Lynch, 1960). "Getting lost" was an ever-present problem, and stimulated the development of numerous navigational aids, some of which exploited spatial memory and other cognitive abilities of the navigator. Overall, there was an intuitive understanding of the meaning of navigation.

2.2. VIRTUAL AND SIMULATED ENVIRONMENTS

During the last decade the problem of "getting lost" has been identified as a severe impediment in the newly emerging *virtual electronic spaces* (Dillon, McKnight &

1071-5819/99/110919 + 27 $30.00/0

Richardson, 1990). Familiar concepts such as landmarks and routes associated with physical navigation were therefore evaluated, and sometimes applied, in these new spaces (Darken & Sibert, 1996*a*, *b*). At the same time, essential differences were acknowledged: some were beneficial, such as the release of the information space designer from physical constraints. Simultaneously, however, the new design freedoms undoubtedly posed challenges for the interaction designer. In place of familiar, constant and pre-existing worlds, the new virtual worlds were unfamiliar and of many different architectures. As Wittenberg (1997) remarked, "The concept of navigation in hyperspace has a completely different physics from navigation in the physical world". While some simulated worlds are based on a *continuous* and stable Euclidean space (Dahlback, 1998), to which familiar spatial concepts might still apply, others such as hypermedia are typically based on node-link representations (as with the web) and permit only *discrete* movement, albeit still allowing the concept of "location".

With the development of virtual worlds came an increasing realization of a distinction between *navigation* and the task of *searching for a known item*: of the fact that users might, sometimes, simply want to *explore* to find out "what's there", without an accompanying need to search for some specific destination or object (Benyon & Hook, 1997). Indeed, our proposed framework is consistent with the fact (Dahlback, 1998) that navigation is concerned with *learning* about a space, whereas searching and other activities *use* that space. Moreover, the characteristics of the new framework are unique to navigation, such that it is clearly distinguishable from problem-solving and human–computer interaction in general. The concept of an internal (mental) model or "cognitive map" will be seen to be crucial to the definition of navigation; something valuable in its own right as well as an essential component of searching and other activities.

2.3. ABSTRACT ENVIRONMENTS

The notion of navigation as *the creation and interpretation of an internal model*, the definition we adopt in this paper, immediately suggests the extension of such a definition, not only beyond physical and simulated Euclidean spaces, but also beyond such environments as hyperlinked spaces. One such extension is to *abstract relationships*.

The creation of a mental model, and its use to achieve some goal, is, for example, a major task for the *engineering designer*. In the creation of a new artifact such as a hi fi, the designer is concerned with the influence of the many parameters under their direct control upon the many performances of the hi fi (Sedra & Smith, 1991). The design of even a small part of that hi fi might involve at least 100 parameters and a similar number of performances and, in addition, many derived quantities such as manufacturing yield and previously unsuspected but vital relationships such as trade-offs. In this context, the notion of *space* ceases to be so attractive, and familiar concepts such as spatial memory and "movement" diminish in value. In fact, the object (e.g. hi fi) of interest may not even be physical: similar abstract systems of relationships are associated, for example, with *financial design* (Black & Scholes, 1973) in which the choice of a portfolio or scheme is the eventual goal, and variables of interest such as profit, risk and market volatility have no physical relevance. Nevertheless, in the sense that a fundamental aspect of the engineering or financial designer's task is that of creating an internal model ("cognitive map") of relevant relationships, the concept of navigation has much to offer.

2.4. SOCIAL ENVIRONMENTS

A general framework of navigation can potentially be applied to the activity of social intercourse (Svensson, 1998). Such intercourse in the real world is well understood, but can now take place in virtual worlds (Dieberger, 1997, 1998). Commonly, a user represented by a cursor will move around some (3D) "information space" and indulge in social intercourse with other "inhabitants" of that world. The proposed framework of navigation appears to be relevant to social navigation, and is treated in some detail elsewhere (Rankin & Spence, 1999). In such navigation, two distinct but inter-related internal models may be involved, one for information and the other for social aspects.

2.5. FOCUS AND VIEWPOINT

While the *range of applicability* of the proposed framework is intentionally very wide, its functionality will be seen to be delineated in such a way as to identify characteristics unique to navigation, and which make navigation quite separate from problem-solving and other human–computer interactions, a requirement usefully enunciated by Dahlback (1998). Indeed, our overall definition of navigation as *the creation and interpretation of an internal mental model* is consistent with it being of value either in its own right—as in the act of exploration—or as a component complementary to some other activity of which searching is just one example.

It is also pertinent to remark that the proposal of a framework was stimulated, not particularly by gaps in current thinking, but principally by challenges that the author has faced, during the creation and design of systems, in trying to understand the issue of navigation. Thus, the proposed framework is grounded in the author's experience, and this paper is written from the viewpoint of an interaction designer. Another, complementary view addressing cognitive issues in some depth would be of considerable value in enhancing and modifying the proposed framework, but is beyond both the scope of this paper and the expertise of its author.

3. Overview

The essence of the proposed framework of navigation is shown in Figure 1. It comprises four cognitive activities, each with its own returned result. Other components complementary to the framework will be added later. Prior to a detailed consideration of each activity, a brief overview will now provide necessary context.

The activity of *browsing* is defined as the *registration* (or elicitation or assessment) *of content*: it answers the question "what's there?", but without integrating the result into some structure or map. There may be no specific target being sought, though there may well be "weights", representing importance, that are associated with certain types of content. One example of browsing is one's rapid glance through the morning newspaper before deciding what to read; another is the rapid scan of an unfamiliar restaurant menu to get some idea of prices, dishes and so on. An important issue, to which we shall return, is the manner in which data should be presented to facilitate the process of browsing.

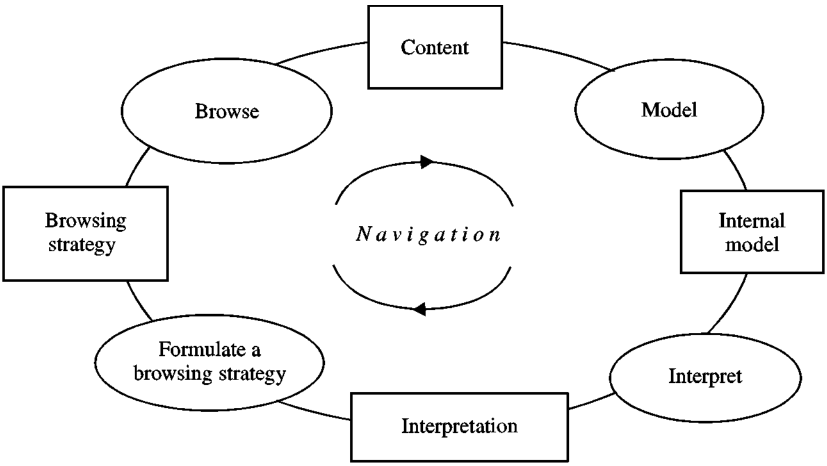

FIGURE 1. The proposed framework for navigation.

The content registered in the browsing activity is available for integration into the internal model, constituting the *formation of an internal model*, or 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".

The internal model, together with an available display of data, can be *interpreted* in a variety of ways. One interpretation could be that no more browsing is needed for the specific task in hand. Or it may be clear that the model is currently inadequate, in which case the manner in which it is judged to be inadequate will influence the subsequent activity, the *formulation of browsing strategy*.

Traversal of this sequence of activities, almost certainly iterative, constitutes the *navigational process*. The control of the various processes is exerted by some higher level cognitve activity, to which brief reference will be made later.

We now proceed to examine the four activities in more detail. Following that, the historical route leading to the new framework will be described and the key aspects of the framework identified.

3.1. BROWSING

Browsing is here defined as the *registration* (alternatively the elicitation or assessment) of *content*, and is very similar to the act of perception (Solso, 1998) in which, for visual browsing, the result of perception is held—albeit momentarily—in sensory storage. Everyday examples include the scanning of a restaurant menu to see what's available, a quick flip through the pages of the morning newspaper to see "what's new" and the riffling of the pages of a new book to determine its structure, layout and composition. Computer-supported examples include the new user of a drawing program viewing a menu to see what operations are possible, and the action taken by a new shift walking into a refinery control room and glancing around to acquaint themselves with the current situation. In all these examples, no search need be involved.

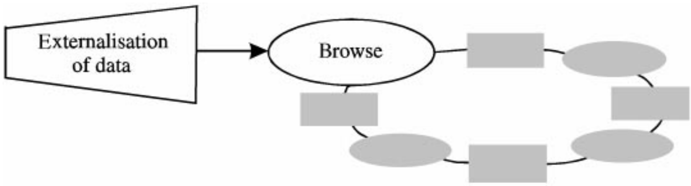

FIGURE 2. Data are externalized for browsing.

For browsing to take place, some externalization (display) of the relevant raw data must be available (Figure 2). The selection, encoding and presentation of that data must be specifically designed, not only to support browsing, but additionally other activities (see later) within the navigational process that require a view of data. As Apperley and others have identified (Jul and Furnas, 1997), inherent data structure must be transformed into an imposed structure the better to lead to the formation of an internal model.

3.1.1. Characteristics

Browsing has, of course, already been studied by others. Carmel et al (1992), for example, identified three activities termed "scan-browse", "review-browse" and "search-browse". Their "scan-browse" ("*scanning for interesting information*") is very close if not identical to the browsing activity of the proposed framework.

The "content" which is registered during browsing may take a variety of forms. In contrast to the simple visual registration of static data encountered in a newspaper or menu, many highly interactive "What if?" scenarios common to engineering and financial design constitute a form of exploratory ("ping-pong") browsing with the ultimate objective of gaining insight into (i.e. a cognitive map of) a property "hidden" within some mathematical relation. In such cases the manual adjustment ("ping") of a cursor may be a component of the browsing activity, being the driver of the display of which the change ("pong") is registered by the user. The concept of browsing can also be generalized to the registration of *function* in which, for example, the user of a new package, uncertain about the meaning of displayed icons, carries out an interactive exploration to determine their functionality: no *search* for a particular function may be involved.

Browsing is commonly undertaken via visual perception, but aural, tactile and olfactory browsing—or indeed a combination of all these as in some medical examinations—may be appropriate.

3.1.2. Weighted browsing

Some aspects of browsing, such as the way it is undertaken by novices and experts, can be accounted for by the concept of *weighted browsing*: Darken (1997) uses the term "primed". For example, a user may, consciously or unconsciously, have selected the nature of the content to be registered, perhaps based on anticipated use or value. Thus, while a novice may have little idea how to "weight" their perception of a conventional map, an experienced "orienteer" may have a well-defined set of weights. As browsing proceeds and an internal model becomes highly developed, weights may become very well defined ("now we need to see what museums there are in this area"). By contrast, opportunistic browsing may begin with a total absence of consciously selected weights.

3.2. MODELLING

As browsing proceeds, the user is able to construct an internal model of the browsed data. The visitor to a new holiday destination will acquire, by browsing a map, some idea not only of the *existence* of an old-walled town, a park, a tramway and a harbour, but also of the relative spatial locations of, and the routes between, these entities, initiating the formation of an internal model. In a similar way, the engineering designer, having observed the variation of two artifact properties as a parameter is varied, might model the trade-off between these two properties, an item of information vital to the process of design. Browsing and modelling will often proceed virtually concurrently.

The activity of modelling seems to be close to the action which Carmel, Crawford and Chen (1992) call "review-browse" and which *integrates* the registered content. As navigation proceeds, the internal model will be extended and may become a *collage* (Tversky, 1993) as, in the case of the engineering designer, exploration generates a collection of internal models. Indeed, some will already be in existence as a result of years of experience, and may be referred to in order to support the extension of the model (Figure 3). There could well be a hierarchy of models as well as models that are disjoint, as in the Geospace system (Lokuge & Ishizaki, 1995).

3.2.1. Externalization

Modelling is profoundly affected by the manner in which the raw data is externalized, and by the affordances by means of which the externalized data can be rearranged to enhance the formation of an internal model. As Apperley and others (Jul & Furnas, 1997) have identified, there are three levels of structure involved; the inherent structure of the data is transformed into an imposed (and externalised) structure, which in turn impacts upon the user's cognitive map (Figure 4). The transformation involves not only the selection, encoding and presentation of raw data but the manner in which all these processes can be influenced through interaction. It is the imposed structure that is browsed to create an internal model which then supports navigation. In the sense that the externalization is designed to enhance the user's internal model of some data, the

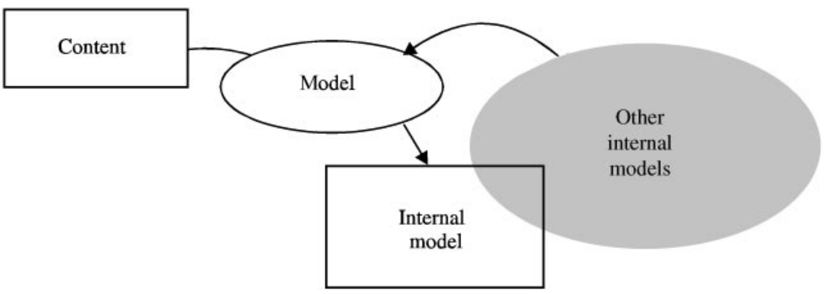

FIGURE 3. Existing internal models may be referred to when integrating newly registered content.

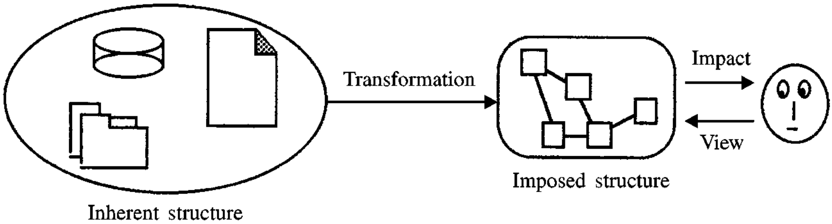

FIGURE 4. Transformation of inherent data structure into imposed structure with a view to influencing the formation of an internal model.

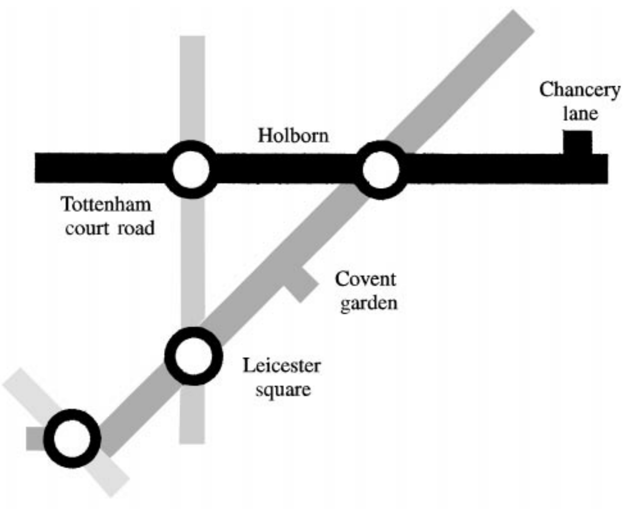

FIGURE 5. Sketch of part of the London Underground Map.

mechanisms of browsing and modelling are directly relevant to the process of Information Visualization, the literature of which (Card, Mackinlay & Scheiderman, 1999) is a fertile source of concepts and techniques.

Depending upon the purpose of navigation, the appropriate externalization can take on a wide range of forms. At one extreme is the well-known non-interactive (printed) example of Harry Beck's London Underground map (Figure 5) which, in view of its topological basis, is far more effective in supporting the creation of a useful internal model than were previous, geographically-based externalizations (Garland, 1994). At another extreme are exceedingly complex visualization tools (e.g. Ahlberg and Wistrand, 1995) which employ a range of sophisticated dynamic techniques to facilitate the many sub-activities involved, for example, in demographic studies.

3.2.2. Characteristics

While we know very little about the nature of internal models it is reasonably safe to assume that, in order to form an internal model of data that is essentially topological (e.g.

Beck's Underground map), the externalization provided for browsing should be designed to make the topological nature of the data clear. Equally, for data that is essentially continuous, it is likely that an analog externalization will be more effective. Nielsen (1990) has identified two concepts to which he assigned the names "*context-in-the-small*" and "*context-in-the-large*", and which appear to have some relation respectively to analog and topological data.

"Context-in-the-small" refers to the difficulty of establishing position within a continuum. Examples of how such context can be perceived via an externalization include (Figure 6) the familiar scroll-bar indicating where the visible page lies within a document, and the count-down clock (Figure 7) of the MINNIE computer-aided electronic circuit design system (Spence and Apperley, 1977) which shows how much time remains before the result of a requested calculation will be displayed.

"Context in the large" refers to location in a discrete space, and is essentially concerned with connectedness (topology). Here it is likely that an essentially topological display will enhance the creation of an internal model. An example can be taken from a hypertext system, where the diagram of Figure 8 might be used to indicate where the user is located in the system. Externalized topological models have, in fact, served as a basis, not only for controlling, recording and displaying the trajectory of past selections in a menu

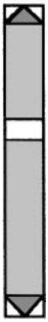

FIGURE 6. Scroll bar.

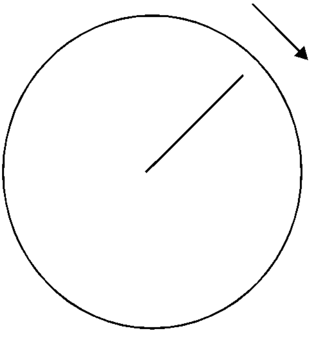

FIGURE 7. Count-down clock.

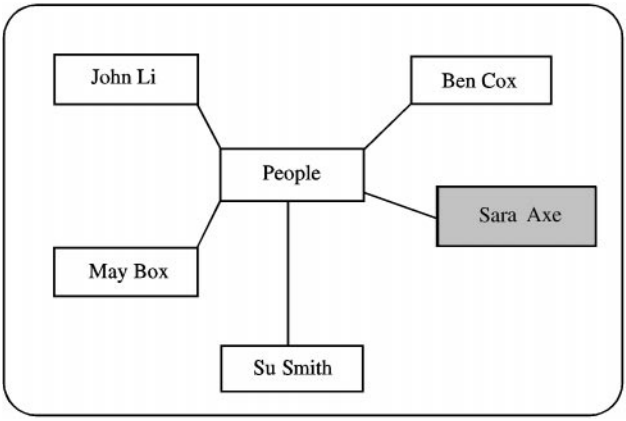

FIGURE 8. Location of a user in a hypertext system.

system, but also to enhance the user's internal model of that menu system (Field & Apperley, 1990).

3.2.3. Dynamic modelling

The dynamic nature of internal models must be recognized. An internal model may decay due to lack of use; be refreshed, for example, by regular viewing of an externalization; or be extended. Indeed frequent use may not only extend the internal model but link together previously separate models. In a complex task such as engineering design it would, in fact, be extremely beneficial for the designer's understanding (i.e. internal model) of a variety of relationships to be coherent rather than disjoint.

3.2.4. Model hierarchy

That internal models can be both disjoint and hierarchical in nature can be illustrated by many examples. Part of the author's internal geographical model relates to the streets around his home, the commuter railway network and the streets around his College. Within each of these areas the internal model is more detailed, allowing the author, for example, to locate his office without undue cognitive exertion.

3.3. INTERPRETATION

Decisions as to how the process of navigation should proceed are based upon interpretations, not only of the internal model but also of externalized data (Figure 9). Once an interpretation is available, a decision can then be made as to how and whether further browsing should proceed, or indeed whether the original task should be modified.

There is a relatively straight-forward form of interpretation which exploits *sensitivity information*, and which we consider first. Two examples will suffice at this stage.

The first (Figure 10) is a special presentation (Spence & Apperley, 1982; Mackinlay, Robertson & Card, 1991) of data-space in which parts are "bent backwards" at an angle so that everything fits into the display area.

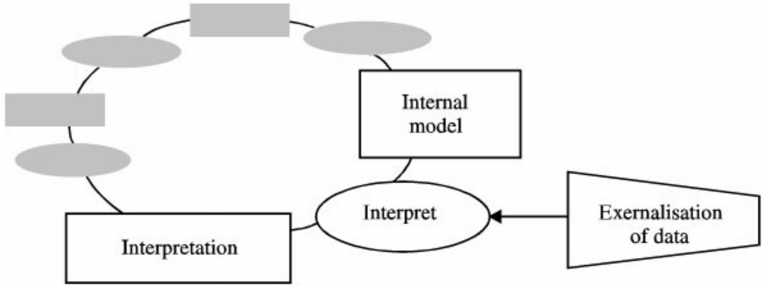

FIGURE 9. Interpretation is based on externalized data as well as the internal model.

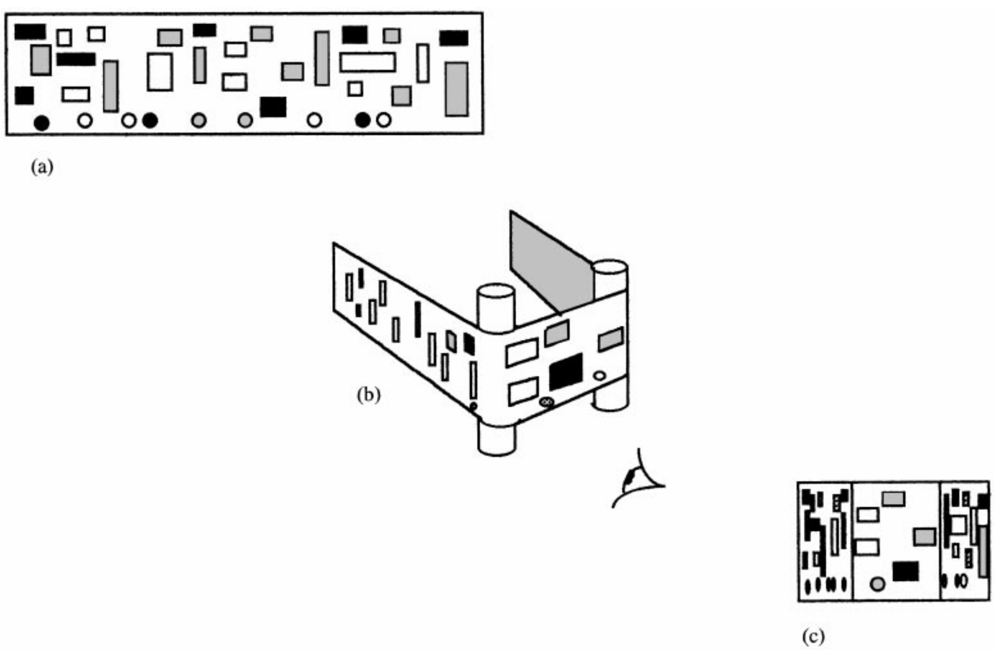

FIGURE 10. Principle of the bifocal display: (a) information space, much larger than the screen, (b) the same space "wrapped" around two uprights and viewed as shown, and (c) the view seen by a user, and scrollable as suggested in (b).

Since items in the peripheral areas would be unreadable because they are "squashed", their text and similar small detail is automatically removed; nevertheless, the colour of each item is still discernable. Thus, a red rectangle towards the left edge of the information space denotes a letter from the user's boss, an item that can be scrolled into the central region, there to expand horizontally and become readable. Such a display offers a clear indication of *sensitivity* in the sense that the user now knows the approximate distance of the boss's letter from the central region, and hence the distance it has to be moved to make it readable.

A second example, involving the aural presentation of information, was provided by the Media Room (Bolt, 1979). Here, a user panning across an image of the MIT campus on a screen measuring 11′ by 8′ becomes 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". The sound provides the sensitivity information.

Although the concept of sensitivity can be useful, and may even be widely applicable, it will in many cases be a gross oversimplification of the manner in which data are 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 externalised and internalized data; and may well be influenced by existing domain knowledge possessed by the user.

3.4. FORMULATION OF A BROWSING STRATEGY

It is unfortunate that definitions of browsing 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 certainly one browsing strategy, but only one among many. When a reader picks up their copy of *The Times*, they know where the weather report is, on which page the obituaries are to be found, where the Comment is and where, on a Saturday morning, the humorous Letters to the Editor are positioned. Their *planned* but largely unconscious browsing is brisk, is essentially saying "what's there?" and is followed by a decision about what to read first. Nevertheless, though that browsing is planned, that plan may suddenly be abandoned, and the user's activity become *opportunistic*, upon noticing the obituary of their old teacher (Suchman, 1987).

A comparable situation often occurs in an abstract environment. In the course of an engineering design, for example, and while executing a planned exploration, the designer might notice an improvement of some valuable property as another property remains constant. Realizing the significance of that effect, the previous plan might be abandoned immediately and attention focused opportunistically on the newly discovered effect. Thus, (Figure 11) the externalization of data can play an important part in the formulation of a browsing strategy.

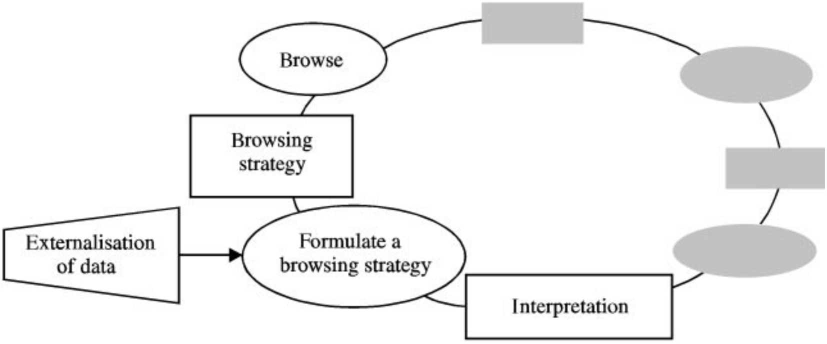

FIGURE 11. Formulation of a browsing strategy draws upon the interpretation and the externalization.

3.4.1. Cognitive determinant

There are two determinants of the formation of a browsing strategy (Tweedie, 1995). One is *cognitive*, based either upon the interpretation that has been made or as a result of a new idea, but not directly influenced by what is displayed. As an illustration, we select the task of searching for a house to buy. Either through access to a database and/or by driving through a selected locality, an internal model will gradually be developed. The interpretation of that model may be that navigation should be terminated because a small group of houses worthy of more detailed consideration has now 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 (Figure 12). The house-hunter may, on the other hand, suddenly decide to adopt a new course of action to see if it proves helpful, that of randomly sampling an estate agent's brochure: a *cognitively* initiated (i.e. not influenced by the display) opportunistic strategy.

3.4.2. Perceptual determinant

The other determinant is perceptual. In other words, the browsing strategy that is formulated is influenced by what the user sees displayed on the screen. For example, sight of an ordered set of coloured "blobs", each representing a house selection previously thought worthy of at least temporary record, might prompt the user to decide that the next browsing strategy will be to re-examine one of those records: a *planned* action stimulated *perceptually*. On the other hand, sight of a large and inexpensive house in a previously rejected area may lead to a *perceptually* initiated opportunistic action, also called a situated action (Suchman, 1987).

3.4.3. Affordances

The browsing strategy that is formulated must be one which is supported by the affordances available for interaction by the user. The detailed affordances most appropriate to the task will, of course, vary considerably. A Bifocal Display (Figure 10) for

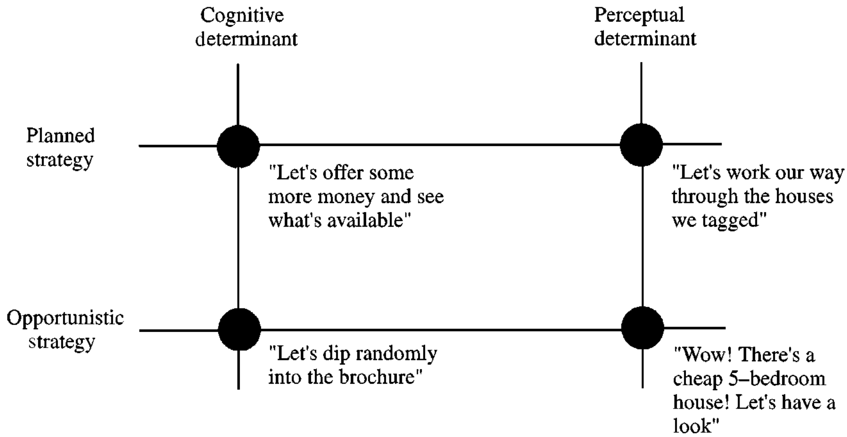

FIGURE 12. Planned and opportunistic browsing strategies may be triggered cognitively or perceptually.

professional office use may, for example, be supported by only one affordance in the form of a scrolling mechanism, and that mechanism may simply be a touch screen sensitive only to the "horizontal" component of finger movement, together with a tapping mechanism to magnify any document in the central region. By contrast, a visualization tool intended for specialized use by an engineering designer (Tweedie, Spence, Dawkes & Su, 1995) or demographic investigator (Ahlberg & Wistrand, 1995) will almost certainly offer a large number of affordances in view of the user's intention both to make serious use of the tool and the wide variety of questions likely to be posed during the act of exploration. Another example of a visualization tool offering a wide range of affordances is shown later in Figure 26.

3.5. NAVIGATION

In proposing a framework of navigation it has been our objective to delineate the activities it embodies and, concurrently, to maximize the range of environments to which it is relevant. Though a well-defined activity, navigation is nevertheless an important component of many human–computer interactions including search and problem-solving. Three examples will suffice.

The house-hunter's search for a house may first engage the navigational process at the stage of formulating a browsing strategy, for example upon being handed a map or offered the use of a car to view houses within a locality. The navigational process may later be terminated at the interpretation stage when it is concluded that a sufficiently good internal model has been formed for some recorded houses to be suitable for more detailed examination.

The engineering designer may, at some stage of design, decide to explore relationships between the properties of a hi fi, possibly to see if any unexpected relationships exist, and may decide to do so using a known affordance to select increasing ranges of a selected property (Tweedie *et al*., 1995; Spence, Tweedie, Dawkes & Su, 1995; Spence, 1999): thus, the formulation of a browsing strategy is again the point at which the navigational process is initiated. Termination of the process may be triggered, again at the interpretation stage, when it is seen that no unexpected behaviour is present.

The internal model itself may be a starting point. A Londoner wishing to undertake an Underground journey might immediately inspect their internal model of that transportation network, and proceed to interpretation. Again, termination of navigation could occur when a mental model of the particular journey "East on the red line, change to the blue line" has been adopted as appropriate.

The three tasks discussed above all involve other cognitive and perceptual processes such as goal formation and overall strategy decision: there is no intention to incorporate these processes within the navigational framework.

3.5.1. Metaphors

We have intentionally not discussed metaphors in the context of navigation. Much of the existing literature on navigation refers to, and sometimes successfully exploits, the geographical metaphor with its landmarks and routes and maps (Lynch, 1960). Nevertheless, and particularly in the context of abstract environments such as engineering and financial design, we suggest that a readiness to consider alternative metaphors would be beneficial. Dahlback (1998) has pointed to the danger of attempting artificially to extend familiar spatial metaphors to environments for which they are inappropriate, or porting successful design solutions from one kind of space to another.

4. Navigational frameworks

A comparison of the proposed framework with alternative definitions of navigation can be achieved via an historical review.

Essentially before the advent of electronic worlds, Downs and Stea (1973) described navigation as (1) orienting oneself in the environment, (2) choosing the correct route, (3) monitoring this route and (4) recognizing the destination has been reached. Such a characterization, presumably intended to be iterative, could have some relevance in physical environments, but suffers from the assumption of a "searching goal" with a unique "correct" route, thereby excluding the act of exploration. No explicit reference is made to an internal model or the externalization of data. The absence of memory is also characteristic of the Ahlberg and Truve (1995) model of the information visualization process which, as has been noted, has an important part to play in navigation.

Wickens' (1984) characterization, again firmly based in the Euclidean space metaphor, explicitly involves a *cognitive map* ("survey knowledge") formed by combining knowledge derived from the identification of *landmarks* and *routes*. Broadly speaking, there is some correspondence between the identification of landmarks and routes and our definition of *browsing*. Nevertheless, there is no loop involved explicitly making navigation an iterative activity. The place of data externalization is not explicit.

More recently, at a seminal workshop on Navigation in Electronic Worlds (Jul & Furnas, 1997), Darken (1997) proposed a schematic model of the navigation process (Figure 13). It explicitly represented a task being performed, and included *Strategy*, *Movement* and *Progress Evaluation*. The involvement of an internal model was implied by reference to "*a priori* spatial knowledge". There is some similarity between the actions of *Strategy* in Darken's model and *Strategy formulation* in the proposed framework, between *Movement* and *Browsing*, and between Darken's *Progress Evaluation* and our *Interpretation*. Darken's model does not introduce any externalization of data. Iteration is present, but the essential formulation of the internal model is not explicit.

At the same workshop Spence (1997) proposed his initial model of the navigation process (Figure 14), now superseded by the framework proposed in this paper. In Figure 14, the importance of an internal model was recognized, and "Gradient Perception" refers to the sensitivity-based interpretation discussed earlier. At the workshop, however, it was modified by the participants to contain explicit representations of the Goal, and further modified to the form shown in Figure 15.

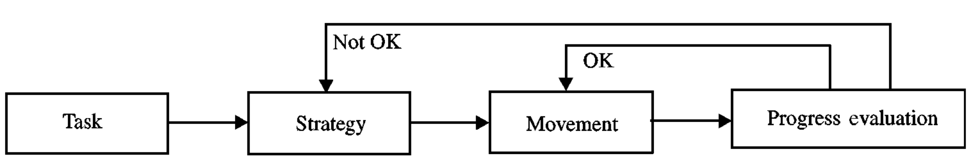

FIGURE 13. Darken's model of navigation.

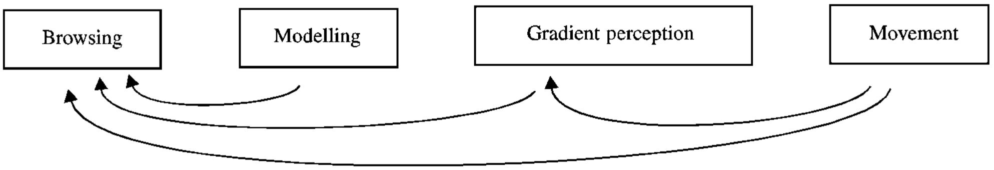

FIGURE 14. Spence's (1997) earlier proposal of a navigational framework.

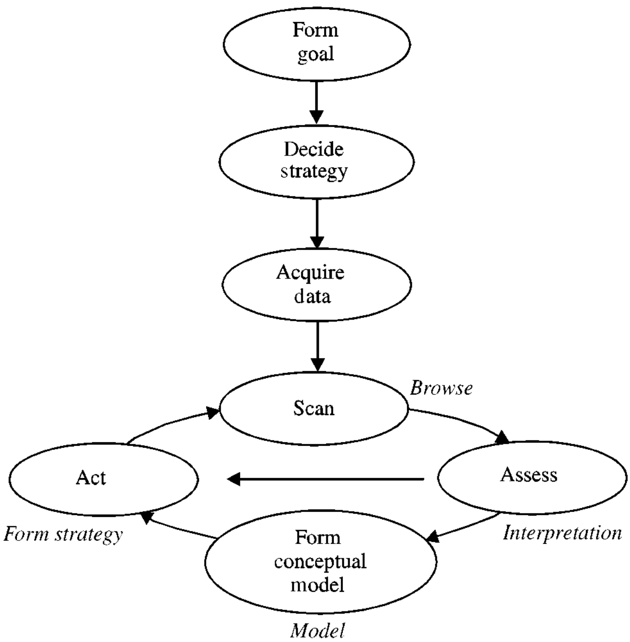

FIGURE 15. Modification to Spence's (1997) original proposal (italic annotation added in this paper).

Although the lower-half contains activities which can loosely be compared with the four activities of the new framework (and are so annotated in italics), two appear to have been interchanged, and it is unclear where the browsing ("scanning") strategy is being formulated.

The proposed framework reported in this paper, however, emerged from a comparison of Figure 14 with Figure 16. The latter is the well-known "Design-Execute-Analyse-Predict" (DEAP) model (Box & Draper, 1987; Deming, 1996; Su, Nelder, Wolbert & Spence, 1996) familiar to statistical modellers, and is an essential activity followed in engineering design such as the creation of artifacts to exhibit high quality. Knowledge of the DEAP model was instrumental in leading to the proposed navigational framework. The direct analogy between the DEAP model and the proposed framework is identified by placing the new framework (Figure 17) in juxtaposition with the DEAP model. The direct analogy, in fact, offers the possibility that concepts and techniques from the field of statistical modelling may be worth examining to establish their relevance to the activity

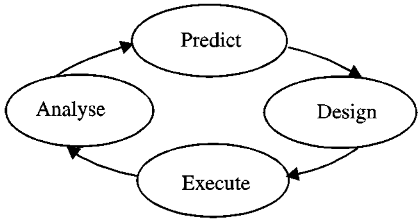

FIGURE 16. The Design-Execute-Analyse-Predict cycle associated with statistical modelling and relevant to the engineering design of high-quality artifacts.

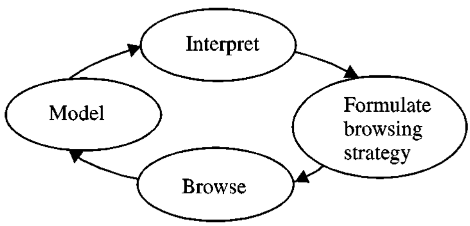

FIGURE 17. The proposed framework.

of navigation in real and virtual worlds. For example, Malik, Su and Nelder (1998) studied the "sampling strategy" (similar to browsing strategy, and represented by the activity of design in Figure 16) that would yield the highest information content.

In summary, it is worth reviewing, briefly, the significant characteristics of the proposed framework: (1) it is schematic rather than verbal; (2) it does not attempt to model the entire process of problem-solving, and human–computer interaction in general, but merely the unique activity of navigation which is called upon by external cognitive decision-making processes; (3) it involves the creation and interpretation of an internal model as an essential component; (4) it is concerned with learning about a space rather than necessarily using it, and thereby clarifies, for example, the distinction between browsing and searching; and (5) it allows specific consideration of the interaction design appropriate to the four activities that together constitute navigation. Such consideration is the topic of the following section.

5. Support by externalization and interaction

In keeping with a focus on interaction design, we now examine each of the activities within the navigational framework to see how they can be supported by the externalization of data and by interaction with that data. The four activities will be considered in turn.

5.1. SUPPORT FOR BROWSING

There are many ways in which browsing can be facilitated and many ways in which it can be impeded. An excellent example of the facilitation of browsing is the opening scene of the CDi title *Richard Scarry's Busiest Neighborhood Disc Ever* (Figure 18). It not only makes full use of screen space to present a pictorial representation of content, but very simply extends its coverage by having the "camera" pan across the imaginary town in a few seconds. In this way, the viewing child very quickly registers a great deal of content.

The ability to browse images effectively and rapidly is also exploited in the Elliot and Davenport (1994) "Video Streamer" browser (Figure 19). Fast movement of the cursor along the side of the "video block" shows many full frames, each for a short time. The same effect could, of course, be achieved by moving a slider along a separate scale, but the

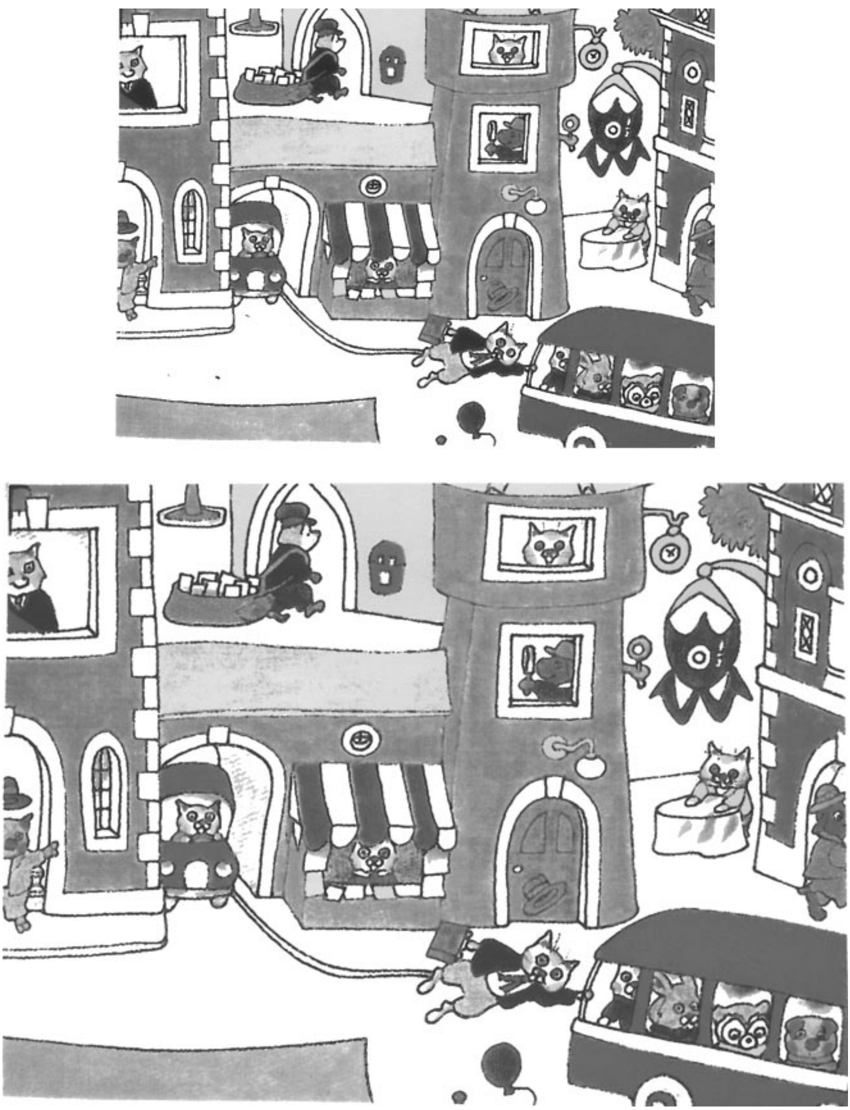

FIGURE 18. A CDi opening scene.

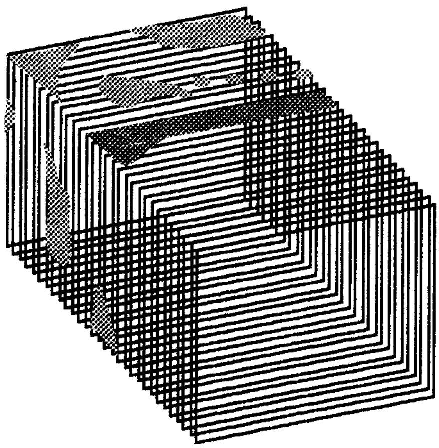

FIGURE 19. The Video Streamer.

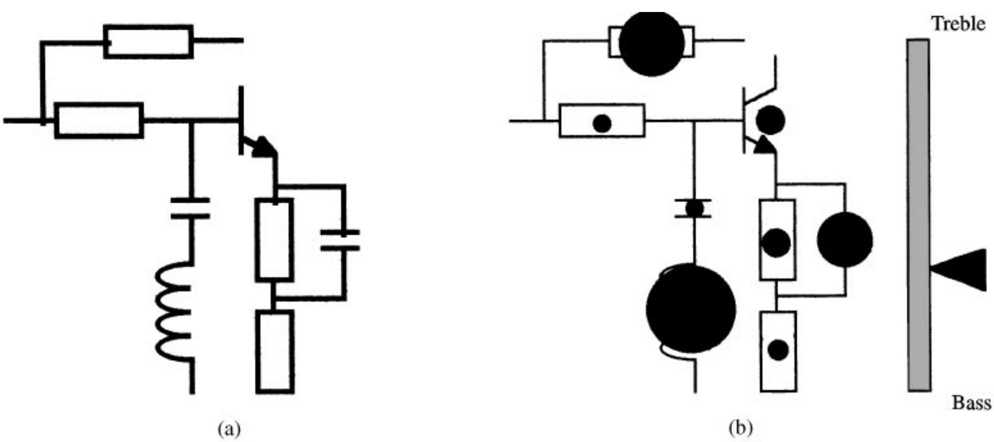

FIGURE 20. (a) Part of a hi fi electronic circuit. (b) Circles indicate the extent to which the hi fi's performance depends upon a component's value, and change in size as the frequency varies continuously between bass and treble.

very shape and composition of the block additionally provides assistance for the modelling which is supported by the patterns visible along the top of the block. A similar approach to image browsing, identifying an important space–time trade-off relevant to interaction design for this activity, has been demonstrated by Lam and Spence (1997) and Wittenberg, Ali-Ahmad, Lahiberte and Lanning (1998).

Animation can be of considerable assistance in browsing. Consider, for example, an electronic circuit such as a hi fi, part of which is illustrated in Figure 20(a). Each component influences the performance of the circuit, and it is immensely valuable for the designer to be aware (at least qualitatively and initially) of the extent of that influence. In the MINNIE system, supporting the interactive graphic design of electronic circuits (Spence & Apperley, 1977), in which the drawn circuit is displayed on a screen,

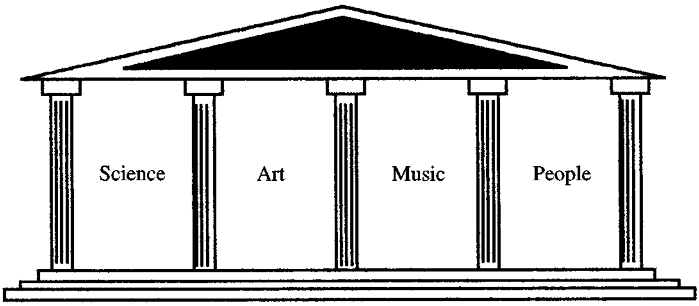

FIGURE 21. Little opporunity to browse, but maybe intentionally so.

a circle is superimposed on each component [Figure 20(b)], circle size indicating the influence of that component on circuit performance. The circles are then changed in size as the frequency of operation of the hi fi is varied linearly up and down between bass and treble limits indicated on the scale. The designer's observation (browsing) of which circles are large or small at which locations in the circuit and at which parts of the frequency scale can usefully enrich their internal model of circuit performance (Spence & Drew, 1971).

Examples abound where, at a first glance, an opportunity for browsing appears to have been wasted. The presentation of Figure 21, for example, reminiscent of many early multi-media encyclopaedias, appears to restrict browsing to only four items. Nevertheless, their main purpose may legitimately have been to allow the user to reach a particular word as conveniently and quickly as possible, and with no intention that the user should build an internal model (of something as large as an encyclopaedia) to facilitate later use.

The above examples all identify an important guideline for the interaction designer, which is to be in possession of a deep understanding of the overall task which draws, at one or more points, on the activity of navigation.

5.2. SUPPORT FOR MODELLING

Perhaps the most obviously effective support for modelling is the provision of appropriate externalization. In other words, to ensure that the content which is browsed is well matched to the type of internal model most relevant to the task involved. An example is the London Underground map (Figure 5), skilfully designed through its use of selection (of relevant data), encoding (colour, and symbols for intermediate and interchange stations) and presentation (essentially piecewise-linear) to be easily committed to, and recalled from, internal memory.

Effective internal modelling can also be enhanced via the phenomenon of "tight coupling", a term descriptive of the spatial and temporal proximity of the user's input and the system's output in situations wherein browsing is of the "ping-pong" variety. Indeed, examples are known where it is judged beneficial to wait for a lengthy calculation if it is otherwise impossible to then carry out a "what if?" exploration in which response times are less than about one-fifth of a second.

In the special case of menu systems, the activity of modelling can be eased considerably by the provision of an externalization, not only of current state, but additionally a trace of the trajectory leading to the current state. In an experiment specifically designed to test the internal models of subjects, Field and Apperley (1990) concluded that, following previous encounters with the external representation of the trace, a user '*has a much better contextual picture of his current position within a database, and navigation ceases to be a short-term memory exercise*". The task of modelling has been recognized (Hendley, Drew, Wood & Beale, 1995) as particularly acute in large anarchic systems such as the World Wide Web.

5.2.1. Maintenance of the internal model

An internal model must not only be created but be maintained. There are many situations in which a sudden change in an externalization is unwittingly destructive of a well-formed internal model. This situation is illustrated in Figure 22. Figure 22(a) shows a sequence of three states that might characterize some successive stages of a World Wide Web trajectory, or some engineering documention search. What the reader of this paper sees is precisely what the user does not see: by the time the user is exclusively viewing frame **, they have either forgotten frames P and Δ or, perhaps needlessly and unsuccessfully, have expended considerable effort in incorporating all three frames into an internal model that needs, in the absence of an external representation, to be continually rehearsed. The alternative, of partially overlapped frames [Figure 22(b)] is a simple example of how an external representation can significantly reduce cognitive effort by helping to support an internal model. A similar example is provided by Lieberman's (1994) Macroscope which, superimposed on a conventional map, displays (transparent) detail of a selected and magnified locality.

5.2.2. Necessary change

If change has to occur it is immensely helpful, as far as minimizing the cognitive load associated with the maintenance of a good internal model is concerned, if the external representation can change smoothly. With the Cone-Tree representation (Figure 23) of a tree structure (Robertson, Mackinlay & Card, 1991) the rotation of the cones needed to bring a requested part of the tree to the fore is intentionally not instantaneous. A smooth change over a period of about 1 s, and imaginatively accompanied by the casting of

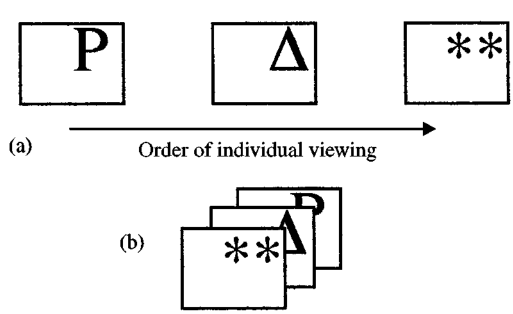

FIGURE 22. Representation of the context of a currently viewed frame.

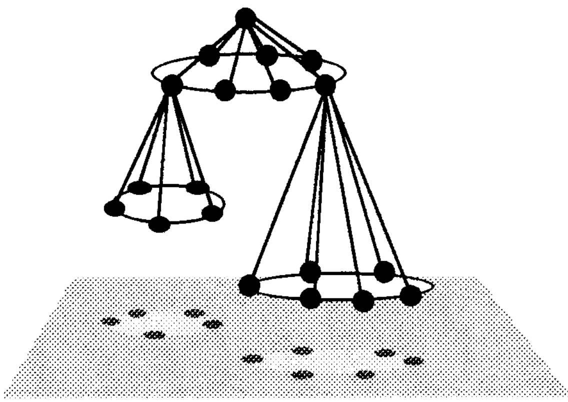

FIGURE 23. The Cone Tree presentation of hierarchically related data.

shadows on a base-plane, constitutes "*.. animation [which] allows the perceptual system to track rotations. The perceptual phenomenon of object constancy enables the user to track substructure relations without thinking about it*". Similarly, in the Lamping, Rao and Pirolli (1995) Hyperbolic Display, topology is again constant and the changes are controlled by (analog) direct manipulation. Similar remarks apply to the Colby and Scholl (1991) "Z-thru" mapping technique which illustrates the additional "stabilizing" effect of a background map.

5.2.3. Model simplification

In many situations, such as when a technician works to rectify a fault within a very large system, a full externalization can be distracting: a simplified externalization containing only immediately relevant context (or with irrelevant matter "greyed out") allows the user to focus on essentials. Nevertheless, spatial layout should be maintained so that restoration of the complete externalization does not cause needless discontinuity. Here, Furnas' (1986) "Fisheye" technique provides a solution with its definition of a "Degree of Interest", a technique used to considerable advantage in the Netmap investigative tool (Davidson, 1993) and in Mitta's (1990) engineering drawing scheme. The technique of "greying-out" in place of removal is familiar from many menu-based schemes.

5.3. SUPPORT FOR INTERPRETATION

Interpretation plays an important part in the process of navigation, so that techniques for its facilitation require explicit study and innovation. Happily, the published literature (Card *et al*., 1999; Spence, 2000) already offers many ideas and results.

Often, techniques to support interpretation hinge on the provision of what has been termed "sensitivity information" and is referred to in the literature as "scent" or "navigational residue" (Furnas, 1997*a*; Pirolli, 1997). These terms refer to the visibility, in the current location, of a "trace" of a remote "target". Relatively simple examples of sensitivity information have already been illustrated in the context of the Media Room and the Bifocal Display. Nevertheless, the challenge faced by the interaction designer in facilitating interpretation is considerable, and increases in severity as the richness, content and extent of virtual worlds increases while the visible frame (the display) remains essentially fixed in size.

Menu systems provide examples of the interpretation of "topological" sensitivity. Menu "look ahead", for example, was investigated by Snowberry, Parkinson and Sisson (1985) and found to have a significant beneficial effect. Later results by Field and Apperley (1990), in which a trajectory of previous selections is displayed and made sensitive to permit selective retreat, show the value of sensitivity display. Encouragingly, the process of placing sensitivity perception on an analytic basis has been given impetus by Furnas' (1997*b*) study, relating the provision of navigational residue to the problem of scale, and has led to strategies for design.

5.3.1. Multidimensional sensitivity

The ability of an engineering designer to perceive sensitivity in multidimensional design space can be *extremely* beneficial. In the Influence Explorer (Tweedie *et al*., 1995; Spence *et al*., 1995), histograms indicate (Figure 24), by colour (shown as grey-scale in the figure), how many customer requirements on performance have been violated, and therefore provide sensitivity information. For example, a design lying just outside a customer limit and known, by its colour, to have failed only one limit will, by definition, become acceptable if that limit is relaxed. A very useful indication of zero sensitivity is provided by the tight coupling indicator (Figure 25) of "ineffective range of limit" employed in the Spotfire application (Ahlberg & Wistrand, 1995) .

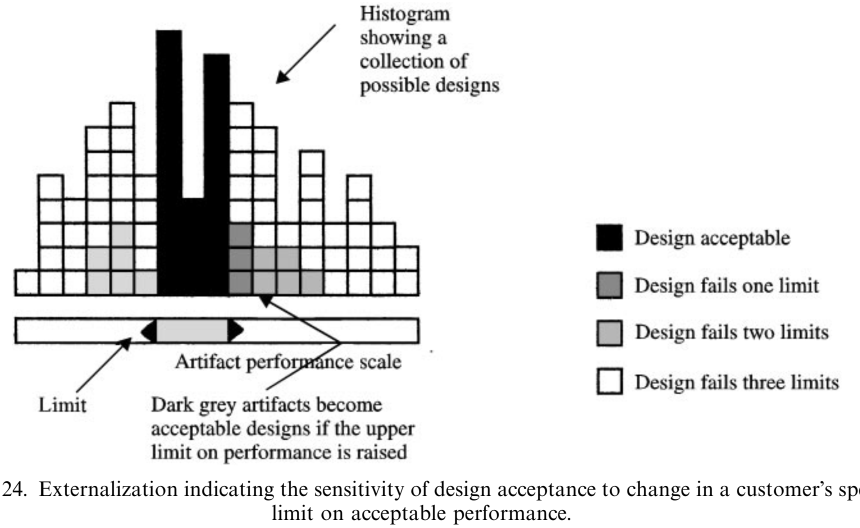

FIGURE 24. Externalization indicating the sensitivity of design acceptance to change in a customer's specified limit on acceptable performance.

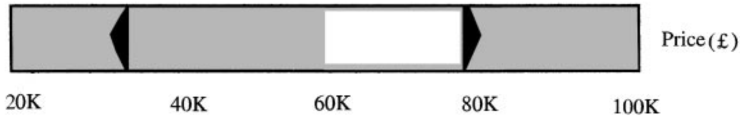

FIGURE 25. A limit slider additionally indicating that the upper limit to Price can be lowered substantially, from about £80 to about £60K without affecting the selection of houses that obey all limits.

5.4. SUPPORT FOR THE FORMULATION OF A BROWSING STRATEGY

Many browsing strategies exist, and it is the aim of the interaction designer, aware of the task(s) that may be carried out, to decide which browsing strategies are appropriate to those task(s) and should be supported by affordances. Such a decision and the subsequent interaction design is a skilled activity, with few guidelines. To illustrate the influence of affordances on the formation of browsing strategy, we select a detailed example based on the Attribute Explorer (Spence & Tweedie, 1998) discussed earlier, for which a possible externalization is shown in Figure 26; here the task supported is that of locating a house which satisfies the gradually formulated requirements of househunter.

The numerous and varied affordances available to the user are indicated by "A" in Figure 26. Together they define the set of interactions that can take place, and hence candidates from which a browsing strategy can at any moment be selected. It follows that each affordance should be clearly defined, and that actual and perceived affordances should be identical: here again, the skill of the interaction designer is paramount.

Also shown in Figure 26 are various visual items, of which some are affordances, which might suggest the use of particular browsing strategies. For example, (top left) the appearance of attribute labels in a menu might remind the user of untested explorations, and the icons (top) representing recorded selections, as well as being affordances, might suggest new planned or opportunistic browsing strategies.

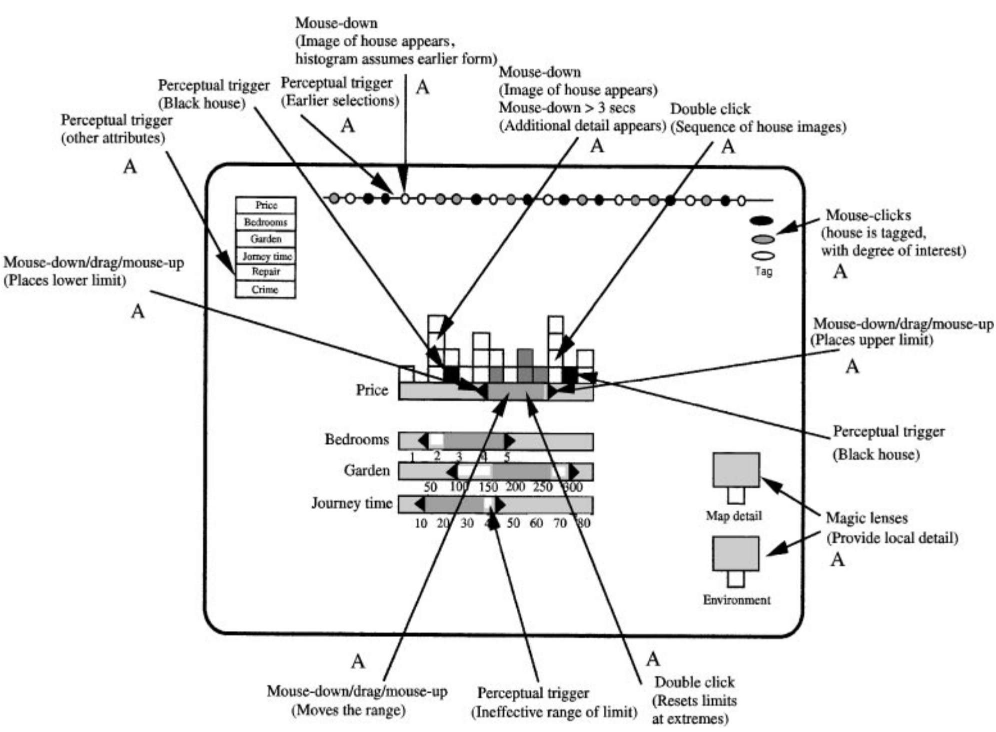

FIGURE 26. A possible design of an Attribute Explorer, showing a large number of affordances and perceptual triggers for browsing activity.

5.5. SUPPORT FOR NAVIGATION

Many of the interaction techniques used for illustration actually support more than just one of the four activities comprising navigation. For example, the displayed menu trajectory investigated by Field and Apperley (1990) not only indicates topological sensitivity but was shown, by experiment, to enhance the formation of an internal model. The displayed menu also supported selective retreat through the menu system, and could therefore influence the selection of a browsing strategy. We see, therefore, that careful judgement is required regarding the most effective use of screen space. The challenge to the interaction designer was identified by Apperley et al (in Jul & Furnas, 1997), who remarked that "navigation is strongly influenced by the extent to which the information space has been moderated; that is, the extent to which its structure has been coordinated and controlled". For an "in-house" world, structure can be "built in" by the designer. In contrast, in a non-moderated world such as the Web, structure tends to be anarchic. In designing to support navigation, therefore, the nature of the task will exert considerable influence over the most appropriate interaction, leaving interaction design for navigation as a skill, albeit one which is able to draw upon a rapidly expanding pallette of techniques (Card *et al*., 1999). An additional challenge is posed by individual differences in navigational behaviour (Hook, Dahlback & Sjolinder, 1996). As with all design for human–computer interaction, a deep understanding of the task to be supported is paramount.

5.5.1. Navigation as a component of search

It has already been emphasized how navigation, as defined, is a self-contained activity but often accessed in the course of tasks such as searching. In view of the current debate concerning the difference between browsing and searching we offer a brief example here for clarification.

In the search for a house to buy, the house-hunter will typically begin with a *vaguely* expressed set of requirements: Cost? *About* £70K. Number of bedrooms? Two, *but* it would be nice to have three. Location? A *nice* neighbourhood, *perhaps* close to Granny who could babysit for us. In view of the vagueness of many requirements, and in the absence of knowledge about the area where a house is sought, there is a clear need to *explore* to create an internal model. Weighted browsing, with weights consciously or unconsciously placed upon such attributes as a green environment, proximity to schools, distance from an industrial area, house appearance and absence of grafitti, and perhaps supported by browsing using interactive histograms of the form shown in Figure 26, will gradually create an internal model whose interpretation may, at some point, show that two or three houses may be worth further study in greater detail. But this process is not only triggered by the need to search, but continuously provides information, in the form of the internal model, which may well alter the search parameters or even the goal itself ("why don't we just rent?"). Thus, *outside* the navigational framework, the activities of problem reformulation, result evaluation and searching strategy will take place, all essential components of searching activity but not involved in navigation. In the limit, where a satisfactory internal model already exists, a search may not require any browsing of an externalization at all.

It follows, of course, that the interaction designer must be aware that externalizations of data must support, not only navigation, but those additional activities which, from time to time, call upon and sense navigational behaviour.

6. Conclusion

The new framework of navigation is relevant to a wide range of physical and abstract (including social) environments, and constitutes a well-defined and self-contained module of activity available as a component in a wide range of human–computer interaction. Examination of mechanisms to support the four constituent activities suggests that the framework is a useful basis for organized thought about the activity of navigation and the design of externalizations to support it.

In view of the rich nature of navigational activity and the extensive range of interaction techniques available, it would be an oversimplification to select a small number of guidelines: rather, the guidelines are many and some are implicit in this paper.

As with many other existing tools and concepts for which no final evaluation is available, the value of the proposed framework will only be established through use or disuse.

I am most grateful to Ken Chakhawata of First Derivative Associates for encourging the work which led to the thoughts presented here and for his insightful comments, to Dr Lisa Tweedie for critical observations and to Paul Rankin for the opportunity to collaborate on examining the relevance of the framework to social navigation in virtual worlds. It is a pleasure also to thank Professor John Nelder FRS and Dr Zahid Malik whose collaboration in research into statistical modelling was responsible for suggesting the final form of the framework, and Dr Oscar de Bruijn and Dr Kosthas Stathis for critical comment. Much benefit was derived from attendance at the Workshop on Navigation in Electronic Worlds associated with the CHI'97 conference, and I gratefully acknowledge the stimulation provided by the participants in that enjoyable activity.

References

AHLBERG, C. & TRUVE, S. (1995) Exploring Terra Incognita in the design space of query devices. *Proceedings of EHCI'95*, pp. 49–67 Amsterdam; North-Holland.

AHLBERG, C. & WISTRAND, E. (1995) IVEE: an information visualisation and exploration environment. *Proceedings of InfoViz'95*, pp. 66–73.

BLACK, F. & SCHOLES, M. (1973) The pricing of options and corporate liabilities. *Journal of Political Economy*, **81,** 637–654.

BENYON, D. & HOOK, K. (1997) Navigation in information spaces: supporting the individual. *Proceedings of INTERACT'97*, pp. 39–45. New York: Chapman & Hall.

BOLT, R. (1979) *Spatial Data Management*. Cambridge, MA: Massachusetts Institute of Technology.

BOX, G. E. P. & DRAPER, N. R. (1987) *Empirical Model-Building and Response Surfaces*. New York: Wiley.

CARD, S. K., MACKINLAY, J. D. & SHNEIDERMAN, B. (1999) *Readings in Information Visualisation*. San Francisco: Morgan Kaufman.

CARMEL, E., CRAWFORD, S. & CHEN, H. (1992) Browsing in Hypertext: a cognitive study. *IEEE Transactions on Systems, Man and Cybernetics*, **22,** 865–884.

COLBY, G. & SCHOLL, L. (1991) Transparency and blur as selective cues for complex visual information. *Proceedings*, vol. 1460, pp. 114–124. International Society for Optical Engineering.

DAHLBACK, L. (1998) On spaces and navigation in and out of the computer. In *Exploring Navigation: Towards a Framework for Design and Evaluation of Navigation in Electronic Spaces*. Technical Report SICS T98.01. Swedish Institute of Computer Science.

DARKEN, R. P. (1997) A simple model of navigation in large-scale virtual spaces. Position paper for Workshop on *Navigation in Electronic Worlds* (see Jul & Furnas below).

DARKEN, R. P. & SIBERT, J. L. (1996*a*) Navigating in large virtual worlds. *International Journal of Human Computer Interaction*, **8,** 49–72.

DARKEN, R. P. & SIBERT, J. L. (1996*b*) Wayfinding strategies and behaviors in large virtual worlds. *Proceedings of CHI'96*, pp. 142–149. New York: ACM.

DAVIDSON, C. (1993) What your database hides away. *New Scientist*, **9,** 28–31.

DEMING, W. E. (1986) *Out of the Crisis*: *Quality*, *Productivity and Competitive Position*. New York: McGraw-Hill.

DIEBERGER, A. (1997) Supporting social navigation on the World Wide Web. *International Journal of Human Computer Studies*, **46,** 805–825.

DIEBERGER, A. (1998) Social Connotation of spatial metaphors and their influence on (direct) social navigation. HOOK, MUNRO & BENYON Eds., *Workshop on Personalised and Social Navigation in Information Space* Swedish Institute of Computer Science Technical Report SICS T98.02, pp. 5–14.

DILLON, A., MCKNIGHT, C. & RICHARDSON, J. (1990) Navigation in Hypertext: a critical review of the concept. *Proceedings INTERACT'90*, pp. 587–592. Amsterdam: Elsevier.

DOWNS, R. & STEA, D. (1973) *Image and Environment*: *Cognitive Mapping and Spatial Behaviour*. London, Edward Arnold.

ELLIOT, E. & DAVENPORT, G. (1994) Video Streamer. *CHI'94 Companion Proceedings*, pp. 65–66. New York: ACM.

FIELD, G. E. & APPERLEY, M. D. (1990) Context and selective retreat in hierarchical menu structures. *Behaviour and Information Technology*, **9,** 133–146.

FURNAS, G. W. (1986) Generalised fisheye views. *Proceedings of CHI'86*, pp. 16–23. New York: ACM.

FURNAS, G. W. (1997*a*) Space-scale diagrams: understanding multi-scale interfaces. *Proceedings of CHI'95*, pp. 234–241. New York: ACM.

FURNAS, G. W. (1997*b*) Effective view navigation. *Proceedings CHI'97*, pp. 367–374. New York: ACM.

GARLAND, K. (1994) *Mr. Beck's Underground Map*: *A history*. Harrow Weald, Capital Transport Publishing.

HENDLEY, R. J., DREW, N. S., WOOD, A. M. & BEALE, R. (1995) Narcissus: visualising information. *Proceedings of InfoViz'95*, pp. 90–96. New York: IEEE.

HOOK, K., DAHLBACK, N. & SJOLINDER, M. (1996) Individual differences and navigation in hypermedia. In (GREEN, CANAS & WARREN, Eds.) *Proceedings of ECCE-8 Cognition and the Worksystem*, pp. 71–76.

JUL, S. & FURNAS, G. W. (1997) Navigation in electronic worlds. A CHI'97 Workshop. New York: ACM. SIGCHI Bulletin, **29,** 44 –49.

LAM, K. & SPENCE, R. (1997) Image browsing: a space-time trade-off. *Proceedings of INTERACT'97*, pp. 611–612.

LAMPING, J., RAO, R. & PIROLLI, P. (1995) A Focus + Context Technique based on hyperbolic geometry for visualising large hierarchies. *Proceedings of CHI'95*, pp. 401–408. New York: ACM.

LIEBERMAN, H. (1994) Powers of Ten Thousand: navigating in large information spaces. *Proceedings of the Conference on User Interface Software and Technology*, pp. 15–16. New York: ACM.

LOKUGE, I. & ISHIZAKI, S. (1995) Geospace: an interactive visualisation system for exploring complex information spaces. *Proceedings of CHI'95*, pp. 409–414. New York: ACM.

LYNCH, K. (1960) *The Image of the City*. Cambridge, MA: MIT Press.

MACKINLAY, J. D., ROBERTSON, J. J. & CARD, S. K. (1991) The perspective wall: detail and context smoothly integrated. *Proceedings of CHI'91*, pp. 173–179. New York: ACM.

MALIK, Z., SU, H. & NELDER, J. (1998) Informative experimental design for electronic circuits. *Quality and Reliability Engineering International*, **14,** 177–186.

MITTA, D. (1990) A Fisheye presentation strategy: aircraft maintenance data. *Proceedings of INTERACT'90*, pp. 875–878.

NIELSEN, J. (1990) The art of navigating through hypertext. *Communications of the ACM*, **33,** 296–310.

PIROLLI, P. (1997) Computational models of Information Scent-following in a very large Browsable text collection. *Proceedings of CHI'97*, pp. 3–10. New York: ACM.

RANKIN, P. & SPENCE, R. (1999) A contrast between information navigation and social navigation in virtual worlds. In (MUNRO, HOOK & BENYON Eds.), *Social Navigation of Information Space* Berlin: Springer-Verlag (in press).

ROBERTSON, G. G., MACKINLAY, J. D. & CARD, S. K. (1991) Cone Trees: animated 3D visualisations of hierarchical information. *Proceedings of CHI'91*, pp. 180–194. New York: ACM.

SEDRA, A. S. & SMITH, K. C. (1991) *Microelectronic Circuits* (3rd edn), New York: Holt, Rinehart and Winston.

SNOWBERRY, K., PARKINSON, R. & SISSON, N. (1985) Effects of help fields on navigating through hierarchical menu structure. *International Journal of Man-Machine Studies*, **22,** 479–491.

SOLSO, R. L. (1998) *Cognitive Psychology* (5th edn). Boston: Allyn and Bacon.

SPENCE, R. (1997) Towards a framework for organised thought about navigation. Position paper for Workshop on *Navigation in Electronic Worlds* (see Jul & Furnas).

SPENCE, R. (1999) The facilitation of insight for lanalog design. *IEEE Transactions on Circuits & Systems* II, **46,** 5, pp. 540–548.

SPENCE, R. (2000) *Information Visualisation*, London. Reading, MA: Addison-Wesley (in press).

SPENCE, R. & APPERLEY, M. D. (1977) The interactive man-computer dialogue in computer-aided circuit design. *IEEE Transactions on Circuits and Systems*, **CAS-24,** 49–61.

SPENCE, R. & APPERLEY, M. D. (1982) Data base navigation: an office environment for the professional. *Behaviour and Information Technology*, **1,** 43–54.

SPENCE, R. & DREW, A. J. (1971) Graphical exploration in electrical circuit and model design. *Proceedings*, *2nd Man-Computer Comunication Seminar*, Ottawa, NRC, pp. 61–70.

SPENCE, R. & TWEEDIE, L. (1998) The Attribute Explorer: information synthesis via exploration. *Interacting with Computers*, **11,** 137–146.

SPENCE, R., TWEEDIE, L., DAWKES, H. & SU, H. (1995) Visualisation for functional design. *Proceedings of InfoViz'95*, pp. 4–10. New York: IEEE.

SU, H., NELDER, J., WOLBERT, P. & SPENCE, R. (1996) Application of generalised linear models to the design improvement of an engineering artifact. *Quality and Reliability Engineering International*, **12,** 101–112.

SUCHMAN, L. (1987) *Plans and Situated Actions*. Cambridge: Cambridge University Press.

SVENSSON, M. (1998) Social navigation. In *Exploring Navigation*: *Towards a Framework for Design and Evaluation of Navigation in Electronic Spaces*. Technical Report SICS T98.01, Swedish Institute of Computer Sceince.

TVERSKY, B. (1993) Cognitive maps, cognitive collages and spatial mental models. In *Proceedings of European Conference COSIT'93*, *Spatial Information Theory — A Theoretical Basis for GIS*, Lecture Notes in Computer Science, pp. 14–24. Berlin: Springer-Verlag.

TWEEDIE, L. (1995) Interactive visualisation artefacts: how can abstractions inform design? *Proceedings of HCI'95*, pp. 247–265.

TWEEDIE, L., SPENCE, R., DAWKES, H. & SU, H. (1995) The influence explorer, *Conference Companion Procedings*, *CHI'95*, pp. 129–130. New York: ACM.

WICKENS (1984) *Engineering Psychology and Human Performance*, Columbus: Charles Merrill.

WITTENBERG (1997) Navigation and search - what's the difference? Position paper for Workshop on *Navigation in Electronic Worlds* (see Jul & Furnas).

WITTENBERG, K., ALI-AHMAD, W., LALIBERTE, D. & LANNING, T. (1998) Rapid-fire image previews for information navigation. *Proceedings of Advanced Visual Interfaces AVI'98*, pp. 76–82. New York: ACM.

Paper accepted for publication by Associate Editor Dr A. Dillon