Robert Spence

Sensitivity, residue and scent

Concepts to inform interaction design for the support of information space navigation

Keywords: cognitive processing, human-computer interaction, information spaces, navigation, perception, sensitivity, residue, scent

Navigation through an information space that is often unknown, unstructured and extensive can be challenging, especially when the target is not fully formulated and exploration of the space is deemed to be beneficial. The design of an information space requires the provision of effective navigational cues to indicate to users what information the space contains and how to get to that information. We suggest, and demonstrate by means of examples that the concepts of sensitivity, residue and scent can inform the design of such navigational cues. In order to show the wide application of these concepts we identify their potential to facilitate navigation in three very different contexts. One is the traditional information space in which information is explicitly available to users. A second comprises information spaces in which sensitivity information can be computed, and is especially relevant to interactive design (e.g., engineering or financial) and decision-making. More speculatively, we suggest that the concepts of sensitivity, residue and scent may be of value in socially defined contexts in which navigational cues arise from the collective or individual behaviour of other users.

1. Movement in information spaces

Many types of information space are in daily use. They include the World Wide Web, online shopping catalogues, engineering design databases and virtual reality environments as well as data bases concerned with the investigation of fraud and other criminal activities. They are often very large, unstructured and unfamiliar to the user. Moreover, as reflected in Bertin's (1981) remark that "… a graphic is never an end in itself: it is a moment in the process of decision making", a user may typically undertake from 50 to 100 'movements' in an information space in the course of a working session, whether choosing a new car to buy or planning an investment portfolio.

Any human being moving in an unfamiliar, unstructured and extensive *physical* space, and attempting to reach a goal which might initially not be wholly formulated, would at almost every step be asking – often with some degree of urgency – two vital questions:

"Where can I go from here?"
"How do I get there?"

The same applies to the navigation of an information space, but with the complication that, as Wittenburg (1997) remarked: "The concept of navigation in cyberspace has a completely different physics from navigation in the physical world"

In other words, familiar aids to navigation such as landmarks, regions and boundaries will not always transfer effectively or at all to an information space. Nevertheless, the navigator of an information space will constantly be asking the same two questions:

> "Where can I go from here?"
> "How do I get there?"

In physical space, answers to these questions are often obtained by the interpretation of cues indicating spatial relationships. Thus, the Polynesian seafarer (Figure 1), close to an island which is beyond the horizon and therefore invisible, will nevertheless see the cloud that is positioned above it. That cloud is a cue whose interpretation is useful to the seafaring navigator. We are therefore moved, in the context of an information space with its different physics, to ask what cues – indicating semantic relationships (Dourish & Chalmers, 1994; Dourish, 1999) - can be made available to the information space navigator, and what concepts are available to inform the interaction designer whose task it is to design those cues (Jul & Furnas, 1997). This paper offers some answers.

Figure 1. The cloud formed above an island invisible beyond the horizon provides a navigational cue

2. Sensitivity

Navigation, which is the interactive control of (usually) iterative movement in information space (Cleveland, 1985), necessarily involves two important variables. One is a movement in that space and the other is the interaction required to achieve that movement. These are directly and respectively related to the two questions continually posed by a user: "Where can I go from here?" and "How do I get there?"

To aid the design of cues whose interpretation can answer these questions we propose a definition of *sensitivity*:

> **sensitivity**: a movement in information space and the interaction required to achieve it.

For convenience we express sensitivity S as a 2–tuple

S = SM, SI

where SM denotes a movement in information space and SI the interaction needed to achieve that movement. Two examples, one from a physical space and the other from an information space, will illustrate the concept of sensitivity and its two components.

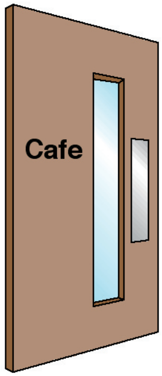

Figure 2. The label 'Café' and the flat plate provide navigational cues by showing where the user can go (the café) and how they can get there (push the door)

Figure 2 shows a door through which a person may need to pass. The actual affordance (Norman, 1988; Gibson, 1979) of the door is that it allows passage by pushing the door. For most people the perceived affordance will be identical with the actual affordance, first because there is no way of pulling the door, and secondly because it is generally understood that the flat metal panel denotes the need to push. That panel is the SI cue: it indicates the interaction required (pushing) to open the door. SI is, in fact, an affordance. The label 'café' indicates the nature of the movement: this SM cue indicates that passage through the door leads to a café.

Figure 3 shows part of a web page, labelled 'Holidays-to-go'. Five areas distinct from their surroundings, together with informative labels, indicate that movement (SM) is possible to other pages concerned with different types of holiday, and that a mouse-click (SI) on the appropriate area will lead to movement to the corresponding page. Thus, the same cues encode both SM and SI. In some cases such cues are easy to interpret, whereas in badly designed cases the interpretation of a cue may be difficult: much depends upon the skill of the interaction designer (Nielsen & Tahir, 2002).

Figure 3. Part of a Web page: each label and surrounding grey area indicate that a mouse click on the area (SI) will cause movement (SM) to another page concerned with a selected type of holiday

3. Residue

While sensitivity, as defined above, is the combination of a single movement in information space and the interaction necessary to achieve it, its representation by a cue or cues would be even more useful if those cues could additionally indicate 'what lies beyond' that single movement. It was for this reason that Furnas (1997) identified the valuable notion of *residue* with the remark that "as well as considering how an information space appears to a user we can usefully think about how the space looks from the perspective of a target which 'wants' a world in which it can be found". To this end we define residue as:

> **residue**: an indication of distant content in the SM encoding

where 'distant' implies content requiring more than one movement to reach it. Here, therefore, is an additional requirement placed upon the interaction designer: the need to design a cue whose interpretation by a user will not only correctly identify the outcome of a *single* movement in information space, but will also suggest content that can be reached *beyond* that single movement. Thus, residue encodes distant content and, in this sense, can be considered to be a generalised form of sensitivity encoding if the definition of sensitivity is extended to multiple sequential movements. Since there is a tendency to regard 'distant content' as that which has not yet been visited it is vital to acknowledge that content already visited may be *revisited*, typically by use of a BACK control. Thus, design for navigation should ensure that consideration is given to the provision of residue to support revisitation as well as visitation.

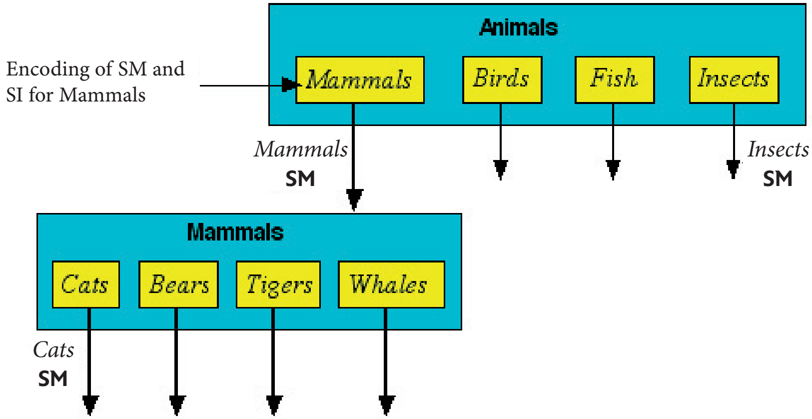

Figure 4. Representation of the top two levels of an hierarchically-structured menu-based system providing information about animals

The notion of residue – as well as sensitivity – can be illustrated by the example of the hierarchical menu system shown in Figure 4. A collection of animals has been arranged in a hierarchy to facilitate navigation towards information about a particular animal. At the top level of the menu system one of four available selectable options is labelled 'Mammals' to indicate that a mouse-click on it will lead to another menu in which options correspond to different kinds of mammal. Thus, the option label 'Mammals' encodes an SM and an SI. But the label 'Mammals' also provides residue for each of the mammals that can be accessed from the second level, as well as animals further down the hierarchy. For example, an Abyssinian cat, located at the level immediately below 'Cats' has, in the label 'Mammals', residue at the top level of the menu.

The ability of the label 'Mammals' to effectively provide residue as well as encode sensitivity is a result of the hierarchical structure adopted for the Animals information space. By contrast, in an unstructured information space such as the Web, the addition of an Abyssinian cat on one web page will not be associated with a residue on many – if any – other pages. Even with an hierarchical menu structure, of course, label design may not be straightforward: for example, a user's incomplete knowledge might lead to the erroneous selection of 'Fish' in the search for information about whales.

4. Scent

So far we have discussed the need to design cues that encode both sensitivity and distant content in such a way as to enhance the likelihood of correct interpretation. The need to interpret cues arises from the fact that this activity must be followed by *evaluation* if anything other than a random choice is to be made from available movements. Thus, the user must assess the *benefit* of each available movement (SM), in effect asking, not

> "Where can I go from here?"

but rather

> "Where can I *most beneficially* go from here?"

It is for this reason that the concept of scent was introduced (Chi et al., 2001).[1] Our use of the term 'scent' should not be confused with its aromatic connotation. Scent can be defined as:

> **scent**: the perceived benefit associated with a movement in information space, evaluated following interpretation of one or more cues.

The term 'scent' arises from the extension of foraging theory, originally developed in a biological context (Stephens & Krebs, 1986), to the search for information (Pirolli & Card, 1999). Implicit in the term 'benefit' is a consideration of the 'cost' of the movement in information space.

It is essential to recognise that the evaluation of scent involves higher-order cognitive processes. It makes reference, for example, to the user's current (though usually ever-changing) internal model of the information space, as well as the strategy – again ever-changing – being adopted to carry out a task. The user may for example be exploring to enhance a mental model of information space and/or to refine a goal or, alternatively, they may be moving as directly as possible towards a temporary or final target. Quite often it is some combination of the two. It is therefore useful to be reminded of Dahlback's (1988) comment that navigation must be defined as something quite separate from problem-solving and other complex human-computer interactions. Thus, while sensitivity and residue are primarily concerned with *local* movement, scent provides the essential link to higher-order cognitive processes – such as mental modelling and strategy formulation – of considerable complexity. While our understanding of these higher order processes is certainly sufficient for the interaction designer to be aware of their profound importance, it is insufficient to provide anything other than the most general guidelines for that designer. Nevertheless, ongoing investigations (e.g., Chi et al., 2001) attempt to employ the concept of scent to develop mathematical models relating user goals, interaction behaviour and visible cues, with potential application to Web design. The relation between sensitivity, residue and scent is illustrated in Figure 5.[2]

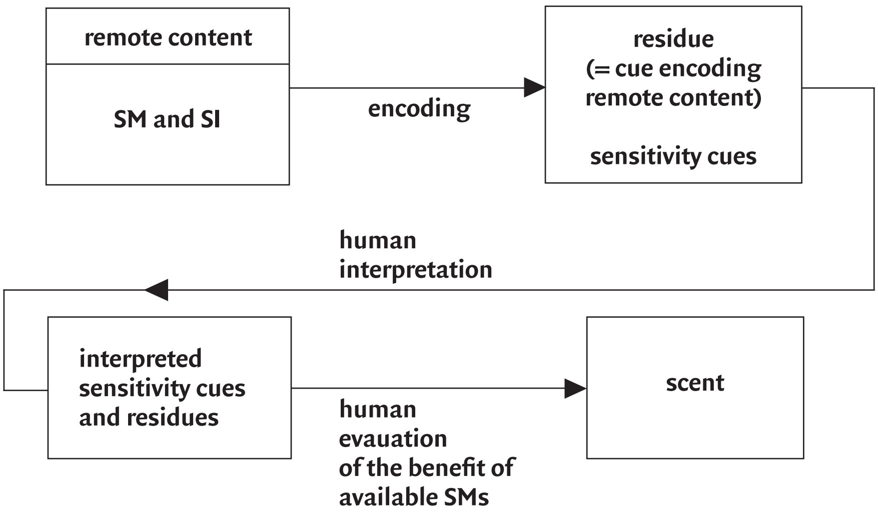

Figure 5. The relation between sensitivity, residue and scent

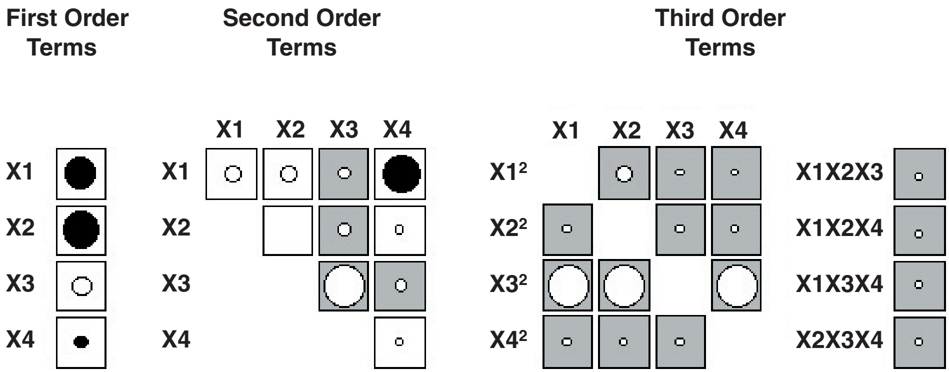

Figure 6. The interface of the Model Maker. Inscribed circles indicate the benefit of including (white circle) or excluding (black circle) the corresponding term within the polynomial being fitted to measured values

Despite the inherent cognitive complexity of scent evaluation, some simple encodings can nevertheless be very effective. A good illustration of encoding supportive of scent evaluation is provided by the Model Maker (Smith et al., 2001), a tool that supports the fitting of a polynomial to measured points: it is specifically developed for users inexperienced in statistics. In the Model Maker interface (Figure 6) every possible polynomial term is represented by a small box clearly encoding SM and SI: a mouse click on a box includes the corresponding term in the polynomial if it is not already included, and removes it if it is. However, within each box is a circle whose size indicates the degree to which the inclusion or exclusion of the corresponding polynomial term is beneficial, thereby providing considerable support to scent evaluation. Figure 7 provides another example in which the size of the selectable menu options might usefully indicate the extent of the data that each option will reveal. Even such a simple example, however, may not be straightforward, since there is often a tendency for a user to make a decision regarding a movement before all the available SM cues have been interpreted (Nielsen, 2000) and the benefit of each movement evaluated.

5. Experimental evidence

Navigation in information spaces has received experimental attention from many investigators. Here we risk the danger inherent in selection and summarise just three studies that are particularly relevant to our discussion of navigation, and especially to the concept of residue.

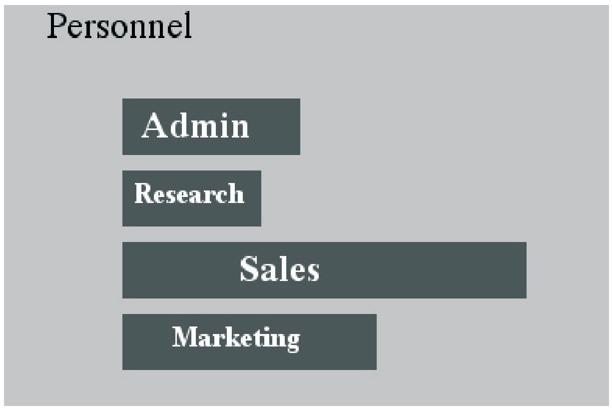

Figure 7. Encoding to support the evaluation of scent

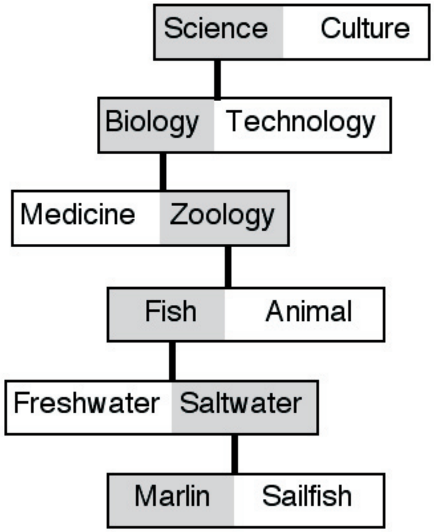

Figure 8. That part of a 2^6 menu to be traversed in a successful search for the target word 'Marlin'

5.1 *Menu structure*

Snowberry et al. (1983) studied the navigation of hierarchically-structured menu-based systems. Subjects were asked to search for a target word within an hierarchically structured database of 64 common English words. Since the number 64 is equal to 8^2, 4^3 and 2^6 it was possible to compare shallow and broad menu structures with narrow and deep ones. That part of the 2^6 menu required to reach the target word "Marlin" directly (i.e., without retreats) is shown in Figure 8.

Each subject was shown a target word and then asked to make successive menu selections, without any backtracking, to arrive at the target word. Figure 9 shows, for the 2^6 menu structure, the number of errors (i.e., incorrect selections) made at different levels of the menu structure. Not surprisingly, most errors occurred at the first and second encountered levels, emphasising the difficulty of providing residue at a distance from a target. In a separate experiment investigating the relative merits of, on the one hand, broad and shallow structures and, on the other, narrow and deep ones, analysis of the results shown in Figure 10 revealed that percentage error

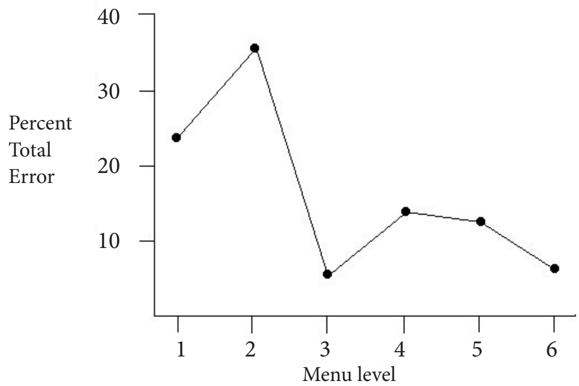

Figure 9. Errors made at different levels of a narrow and deep six-level menu in the search for a target at the lowest level

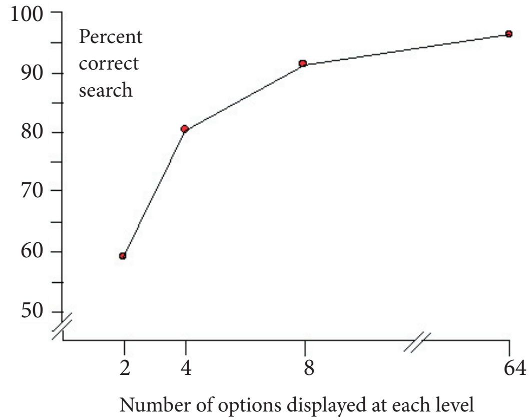

Figure 10. Percentage correct search as a function of menu structure

was significantly ($p<0.0001$) affected by menu structure. There is, here, a clear suggestion that broad and shallow menu structures offer the possibility of a good residue of a target word at the top level(s). Hierarchical structures other than those investigated by Snowberry et al. are of course possible: some have been studied by Norman & Chin (1988).

5.2 *Help fields*

The same investigators (Snowberry et al., 1985) also directed their attention to what they called *help fields*, of which an example is shown in Figure 11. Below each selectable option in the menu is displayed an unselectable sample subset of options that would appear next if that option had been selected. The intention is that sight of these sample labels will clarify the meaning of their superordinate options and, thereby, enhance the confidence of interpretation. The question being posed, in effect, was whether the help field, together with the category labels, constituted useful residue.

It was found, using the 2^6 structure, that the presence of the upcoming help field led to only 8 to 10% error as opposed to about 22 to 28% error when no help field was displayed, a result that was found to be statistically very significant ($p<0.001$). It seems, therefore, that the upcoming help field was effective in providing residue at each level.

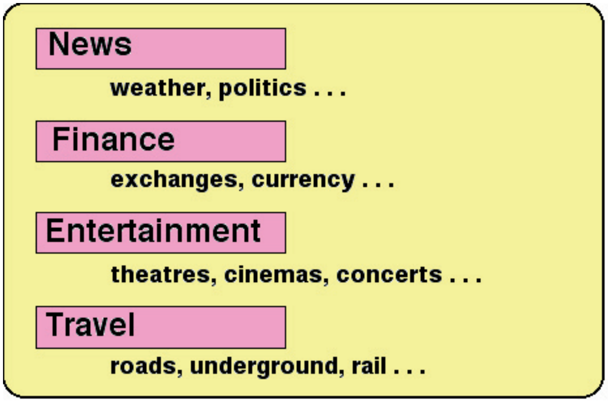

Figure 11. Example of the provision of an 'upcoming' help field, where samples from the next lower level help to enhance confidence in the interpretation of the menu options

5.3 *Complex tasks within extensive databases*

Field and Apperley (1990) removed many of the constraints inherent in the studies of Snowberry and her colleagues by addressing the more realistic task of navigating a reasonably large and non-uniform menu structure to solve a complex task capable of more than one acceptable solution. Added realism was acknowledged by permitting exploration, including retreat, specifically forbidden in the Snowberry experiments. Their information space was a 'videotex' database describing various aspects of the fictitious city of Carlton, and the task given to subjects required the planning of an evening's entertainment with a friend involving travel, a cinema visit, a meal and a return journey. Their study is particularly valuable in the context of our discussions because it identifies a broader aspect of navigation within which the concepts of sensitivity, residue and scent must be positioned.

Field and Apperley compared the two menu systems illustrated in Figure 12. One, called Standard Menu, allowed normal selection from a menu but also the selection, via short typed command, of either the 'home' page or the immediately previous page. The other, called Selective Retreat, not only displayed a 'trace' of previously selected options at each level, but allowed retreat to any one of them via a typed number.

The result of this experiment was a rich collection of results and conclusions deserving more detailed study than can be offered here. There were, for example, significant differences in the number of pages accessed in the two conditions (means for Standard Menu and Selective Retreat were 50.3 and 39.7 respectively). If 'efficiency' is defined as 'the minimum number of frames

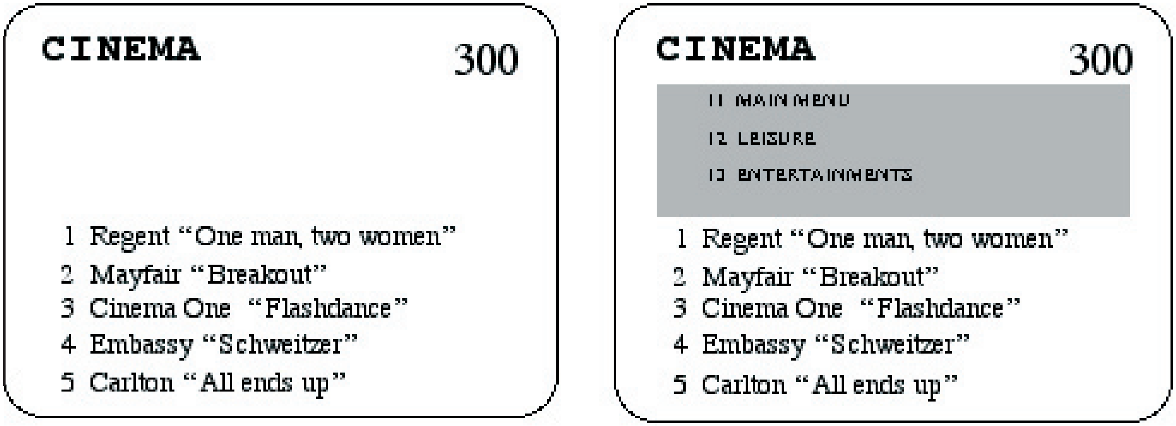

Figure 12. The two menu systems compared by Field and Apperley. The Standard Menu system on the left allows selection by typed numeral. A user types *0@ to retreat to the main menu, and *@ to retreat to the immediately previous menu. The Selective retreat system on the right allows selection, either to move forwards or backwards, by typed numeral

to target divided by the actual number of frames to target', the Selective Retreat group were 63.3% efficient and the Standard Menu group 51.3% efficient.

In the context of our discussion perhaps the most relevant conclusion drawn by Field and Apperley is that subjects using the Selective Retreat structure gained a better understanding of the complexities of the database – that is, they seemed to acquire a better contextual map (i.e., internal model) of the city of Carlton. As Dahlback (1988) has reminded us, navigation is concerned with *learning* about an information space as well as using it. Selective Retreat is, in fact, a very common movement in information spaces. For example, about 60% of all movements within the World Wide Web are *revisitations*, to pages already – and normally very recently – visited (Tauscher & Greenberg, 1997). There is, in fact, on average, a 39% chance that the next URL visited will be found within a set containing the six previously visited pages. As Tauscher and Greenberg remark, "the success of [the] Back [control] is in line with our observation that extreme recency is a good predictor of what page will be revisited". Thus, in the context of interaction design to support navigation, an important conclusion is that the design of sensitivity cues and residue should fully acknowledge the likelihood of revisitations.

6. Sensitivity types

Our discussion of sensitivity has, so far, intentionally been restricted to familiar information spaces in which a user moves from one view of available data to another, making use of fixed navigational cues created by an interaction designer. The concept of sensitivity has, however, much broader application and potential through its extension to computable data in addition to pre-existing data, and through its extension to navigational cues which are created over time by a community of users. These powerful extensions immediately broaden the range of application of sensitivity and related concepts to such activities as engineering design and collaborative work and play.

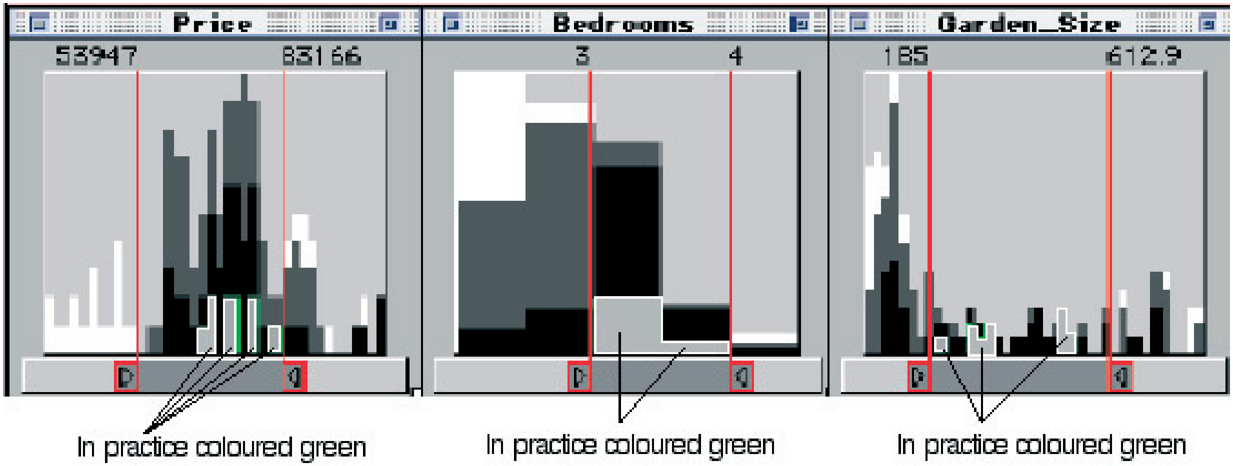

Figure 13. The Attribute Explorer indicates in green (artificially shown here as grey with white outline in the monochrome figure) those houses which satisfy all limits placed on three attributes. Colour coding indicates the extent to which other houses fail those limits

In this section we first consider information spaces for which relevant data is immediately available, and illustrate some issues involved in the design of cues to support navigation. We then, in Section 6.2, extend the concept of sensitivity to situations in which it can be computed. Finally, in Section 6.3 we discuss cues appropriate to social navigation in which the sensitivity data is *facilitated* by an interaction designer, but *collectively created* by users.

6.1 *Available sensitivity data*

There are many information spaces for which sensitivity data is immediately available, and all that is needed is to consider what cues are appropriate to the data and the use that will be made of that data.

A first example is provided by an interface – the Attribute Explorer – designed to support the selection of one object from among many on the basis of its attribute values. The illustrative example (Figure 13) addresses the problem of selecting a house to purchase.

In the Attribute Explorer (Tweedie et al., 1994; Spence & Tweedie, 1998; Albinsson & Morin, 2002) the data for each attribute is represented in the form of a histogram, each house contributing one small rectangle to that histogram. The concurrent display of histograms associated with a number of attributes allows limits to be defined on all of them, as shown in Figure 13. Those houses satisfying all attribute limits are coloured green (artificially shown as grey with white outline in the monochrome figure, see the online edition of this article for a full-colour presentation) on all the histograms, whereas those failing only one attribute limit are coloured black. Grey encodes houses that fail two limits, and so on. A mouse click on any rectangle or group of rectangles could, for example, lead the prospective purchaser to more detail, for example a photograph of a house. Use of such a tool often includes exploration, especially when requirements are far from being formulated precisely.

In the Attribute Explorer we see that the interaction designer has not chosen a separate cue to indicate movement (SM) to each house within a huge collection. Rather, in view of the importance of initially receiving guidance about possibly beneficial changes to attribute limits, the houses failing only one limit have been aggregated and coded black: thus, many SMs are identically encoded, but conveniently grouped. A major advantage of the black encoding is that even when limits are so stringent that no green houses are visible, a black house just outside a limit provides a valuable identification of a limit that might beneficially be relaxed. Thus, there is considerable opportunity to evaluate some scent, in the sense that the 'sacrifice' entailed in relaxing a limit to turn a black house into a green house is represented by the distance of the black house from the corresponding limit.

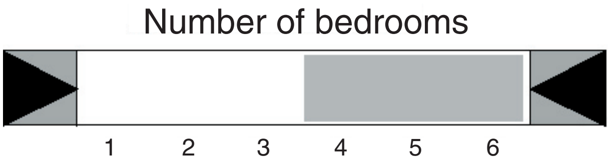

Figure 14. In a limit positioning tool, colour coding indicates that object selection will be unaffected while the lower limit stays within the white region. When a limit moves into the grey region selection will be affected

Another example in which aggregate sensitivity is usefully and very simply encoded is provided by the limit positioning mechanism shown in Figure 14. It was originally employed in the Dynamic Queries interface (see below) and is currently used in the Spotfire™ visualization tool (Spotfire). The white area of the slider shows limit positions which will have no effect whatsoever on the information displayed, whereas movement of a limit within the grey area will cause movement in information space. Such an encoding can be immensely valuable in what have been termed "What would happen if?" (or simply "What if?") situations: the encoding

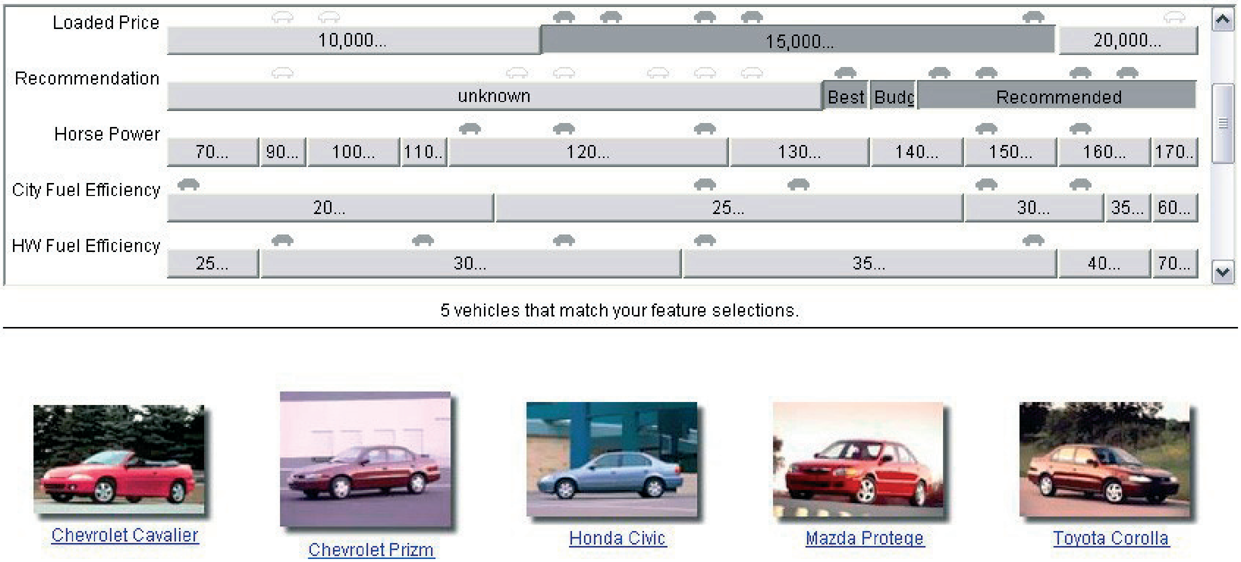

Figure 15. The EZChooser interface, here used to facilitate the purchase of a car. The user has specified some attribute ranges. Icons above the ranges indicate cars that satisfy the user's requirements, and also offer guidance by encoding sensitivity (Figure courtesy of Kent Wittenburg)

performs the useful function of indicating that, for some "what-if?" questions, the answer is "Nothing".

The benefit to be gained from encoded sensitivity is also well illustrated by the EZChooser interface (Wittenburg et al., 2001), illustrated in Figure 15 in the context of online car sales. The interface exploits the concept of bargrams (Apperley et al., 2001). A bargram comprises a number of attribute ranges (e.g., price), above which icons represent individual cars. User requirements are indicated by the interactive selection of attribute ranges, whereupon cars satisfying all requirements are indicated by shaded icons above the selected ranges. Brushing allows the position of a given car to be highlighted on all bargrams. Sensitivity is encoded by outline icons which indicate a car that would satisfy all requirements if the attribute range underneath it were to be selected. An evaluation of scent is supported in the sense that a previous – and possibly tentative – selection can be indicated (for example by a red, green or blue object) and available for comparison with a new potential selection. Such a feature can usefully be extended if a user is allowed to position their 'ideal' (even though possibly unattainable) object on each attribute bargram, again allowing comparison of any potential selection with the ideal (Apperley et al., 1999)

The value of sensitivity information can be emphasised by an example in which the relevant data is immediately available *but not exploited*. The Dynamic Queries interface (Williamson & Shneiderman, 1992; Ahlberg, 1996) is illustrated by the Homefinding example in Figure 16. It is designed to facilitate the task of finding a house to buy. Within a geographical area shown on a map, it identifies by dots those houses which satisfy limits set by a user on the values of certain attributes. A major disadvantage, however, especially when there are no houses that satisfy the limits, is that no sensitivity information is provided – no guidance is given as to the variations in limit values that might lead to an acceptable house being shown on the map. As a consequence, the formation of a mental model of the collection of available houses can only be achieved, and with considerable difficulty, by the adjustment of each limit in turn, a time-consuming and tedious procedure if, as is typical, there may be ten or more such attribute limits. As well as an

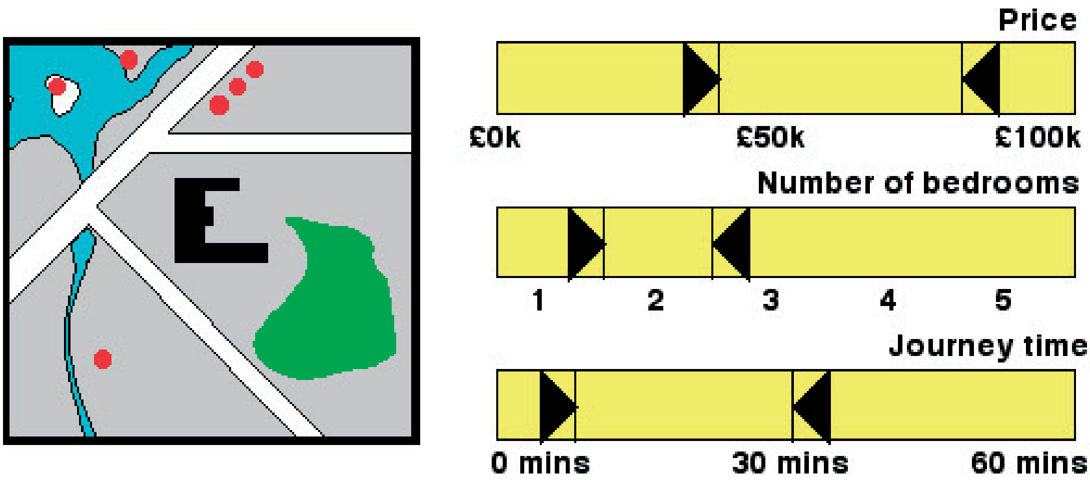

Figure 16. The Dynamic Queries interface. Limits placed on house attributes by a user lead to the display of houses satisfying those limits on the map

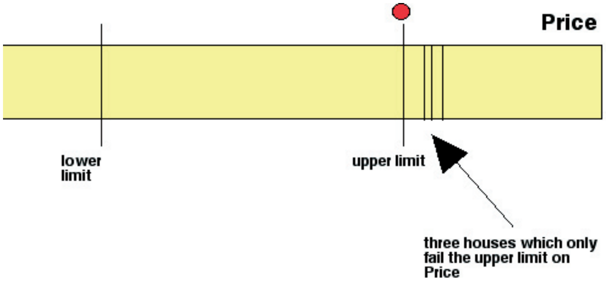

Figure 17. A possible modification to the Dynamic Queries interface. Houses violating only one limit are identified, so that sensitivity is explicit rather than having to be discovered by manual movement of the limits

absence of residue and cues encoding sensitivity, there is no way in which scent can be evaluated other than through the gradual development of a mental model of the data base of houses. As with the Attribute Explorer, a great deal of useful 'non-directed' exploration may precede a more directed movement to a desirable house.

As it happens, the 'black house' sensitivity cues could easily be introduced to a Dynamic Queries interface, as illustrated in Figure 17. To enhance the evaluation of scent some sign (a red circle in this example) could usefully indicate the number of houses lying just outside a limit.

7.2 *Computable sensitivity*

Many situations exist in which sensitivities are not available *a priori*, but their value to the navigational process is so high that consideration should be given to their calculation. Three examples will be given.

Andrienko & Andrienko (2003) acknowledged the value of computed sensitivity and incorporated it in navigational cues in their design of a geographical information system used, in one example, to decide upon the distribution of limited funds to counties within the state of Idaho to help them attract health care professionals. For each county, ten numerical health-related attributes were taken into account, ranging from average fertility rate for the county to the use of medical insurance. A freedom available to the user is the allocation of weights, indicative of perceived importance, to each of the ten attributes, under the constraint that the weights should sum to unity. Following the necessarily subjective allocation of weights, an algorithm ranks the counties in an order of preference indicative of the need for funds.

The question naturally arises "If I had chosen slightly different weights, would the ranking of the counties have been radically changed?" To provide an answer, rankings were recomputed for a range of weights around the values originally chosen, and the result encoded in a histogram for each county as shown in Figure 18. The four bars of a histogram indicate the minimum, mean, maximum and median of the computed ranking, and collectively provide, for each county, some impression of the sensitivity of its ranking to the subjective choice of weights.

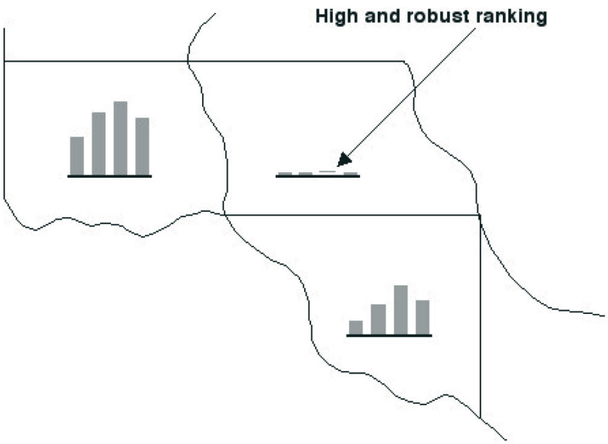

Figure 18. On a conventional map, histograms indicate the computed robustness of the allocation of health funding to changes in subjectively chosen weights

A second example relates to engineering design. If the designer of an artefact is trying to achieve a particular performance F of that artefact, it is immensely helpful to know how F will change in response to small changes in each of the parameters (p) whose value can be chosen by that designer. An incremental version of the sensitivity we have already defined is, in fact, embodied in the partial differential $\delta F/\delta p$; this quantity is the ratio, as changes become infinitesimally small, of a change (or 'movement') in F and the change in parameter p (the 'interaction') which causes it. The relation between SM and a change in F, and between SI and a parameter change, becomes apparent.

It so happens that, for certain electronic circuits, the partial differential $\delta F/\delta p$ can be computed at negligible cost once the circuit has been simulated (Director & Rohrer, 1969); the latter procedure is in any case commonly and frequently carried out during the iterative design process. There are many ways in which this sensitivity information can be presented. A very effective approach (Spence & Drew, 1971) is to map this information directly onto the circuit diagram and, where qualitative performance is of interest (as it very often is), to use size – for example the radius of a circle (Figure 19) – to encode the numerical value of $\delta F/\delta p$ for the parameter on which the circle is superimposed. Qualitative encoding has the additional advantage that the presentation can effectively be animated to show how sensitivity varies with change in another variable (for example as an amplified sound moves along the scale from bass to treble).

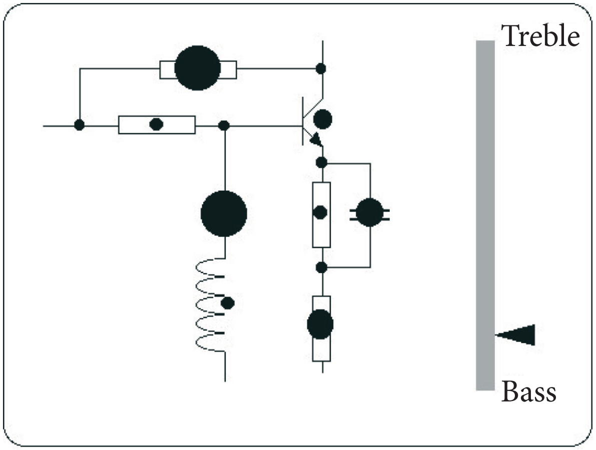

Figure 19. Circles indicate, to an electronic circuit designer, the qualitative effect of variation in the corresponding component on some overall circuit property

A third example is concerned with electromagnetic devices, familiar to us from the ubiquitous electric motor and the huge electromagnets that lift scrap metal. Like electronic circuits they have to be designed to perform a specific function, and their design is complicated both by the nonlinear relationships that characterise the materials involved (e.g., iron) and by the complexity of their shape. Fortunately the partial derivatives of some overall 'worth' of the device with respect to those of its features a designer can choose – in other words, the sensitivity – can be computed at low cost (Cowan & Lowther, 2003). The computed sensitivities can often be displayed to advantage as a two-dimensional map to provide navigational cues for the designer. Thus, to support the improvement of an electromagnet, the shaded regions of Figure 20 indicate, to the designer, where a change in the material would have the largest effect. Conversely, they also indicate where a manufacturing error, or 'tolerance' would be most damaging when trying to maintain consistency of performance over a large number of devices.

Since there are many occasions in engineering and other types of design when "What if?" questions are formulated but not explored in view of the expected prohibitive computational cost, it should be noted that efficient methods of sensitivity calculation are available for nonlinear as well as linear systems, and for dynamic as well as static ones (Brayton & Spence, 1980). As well as artefact design, the human activity of designing a model also benefits immensely from sensitivity data, as exemplified by the simulation of metabolic networks (Qeli et al., 2004).

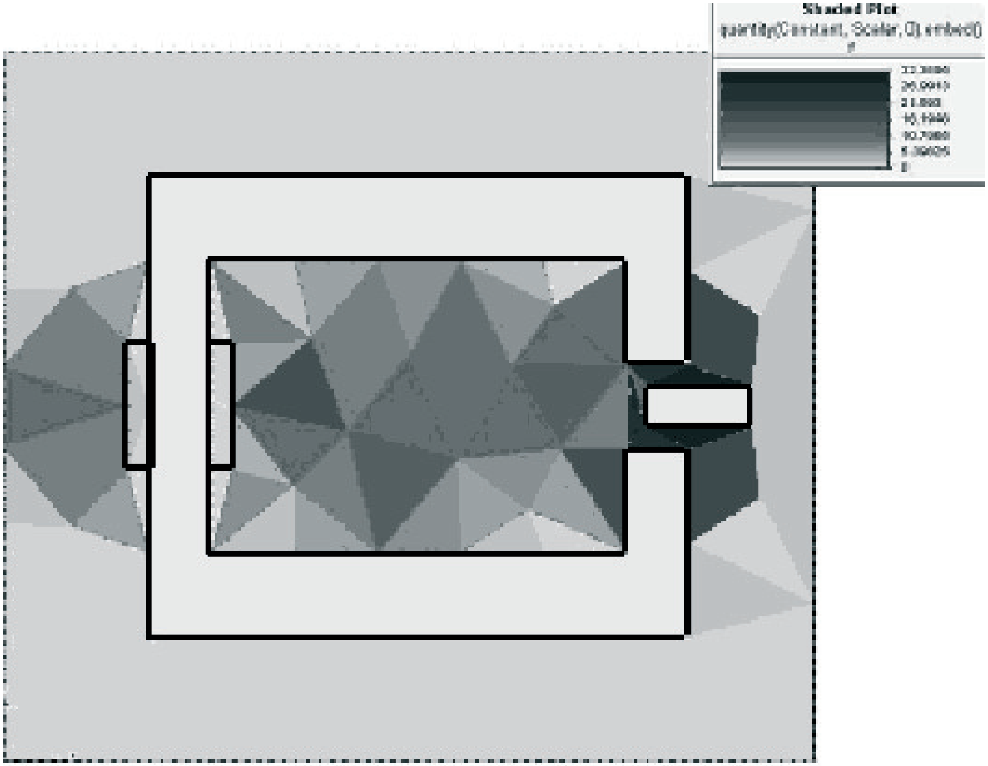

Figure 20. For the designer of an electromagnet, the shaded areas indicate the magnitude of the effect, on a major property of the electromagnet, of changing the magnetic material

6.3 *Socially defined sensitivity*

Navigation of an information space can often be enhanced by sensitivity data created by other users, a situation giving rise to the term *social navigation* (Munro et al., 1999). One form of social navigation is provided by 'recommender systems' (Konstan & Riedel, 2003) in which, for example, a user may be informed that "people who bought this book also bought …". Here, the interaction designer does not design cues directly but rather facilitates their creation by others. Cues may be the result of aggregation and be generated by a community of users, as with a typical recommender system or, perhaps via an annotation on a document, by a single person. A number of systems supportive of social navigation have been designed and evaluated (e.g., Hill et al., 1995).

Simple examples can easily conceal the inherent richness and attendant complexity of social navigation. Harper (1999), for example, in his lucid description of the activities of a 'desk officer' working for the International Monetary Fund, remarks that 'technologically mediated information needs to indicate the socially organised provenance of that information, … indicating .. which institution produced it; who owns that institution; what is the known perspective of that institution (political view, etc.), and so on". His remarks are echoed and developed by Dourish (1999) who points out that social navigation is a more general phenomenon than current practice would suggest, and notes the relevance of anthropological and sociological expertise. It is no surprise, therefore, that Hook et al. (2003) remark that "Social navigation does not have a single underlying theoretical framework.", and that "We do not yet understand how to design for social navigation in various different domains". Thus, while the concepts of sensitivity, residue and scent discussed in this paper would at first sight appear potentially relevant to social navigation, their true value in that context may take some time to emerge.

7. Conclusions

Our discussion aims to inform interaction design which is undertaken with a view to easing the navigation of an information space. To that end we have attempted to separate concepts directly related to navigation from complex higher-order cognitive processes, whether the latter invoke navigation to support a general task or are needed to identify a beneficial movement in information space. The concepts we have identified, and which are closely related to navigation, are sensitivity, residue and scent. We have also identified three classes of sensitivity and residue data to illustrate the considerable potential that exists to enhance the navigational process, particularly in applications not traditionally regarded as involving information spaces.

The limited number of examples in this paper are specifically chosen to illustrate the concepts of sensitivity, residue and scent. By contrast there is a vast literature – comprising informed opinion, concepts and a rich collection of techniques – available for application by the interaction designer concerned with both semantic and social navigation (see, for example, Nielsen (2000), Nielsen & Tahir (2002), Tognazzini (1992), Wildbur & Burke (1998), Tufte (1983, 1990, 1997) and Furnas (1997)). This literature is, in fact, so vast that any attempt to summarise it concisely here would run the risk of superficiality. Rather, it is anticipated that the practical value of that literature will be enhanced through its interpretation in the light of the concepts of sensitivity, residue and scent.

Acknowledgements

Discussions with Oscar de Bruijn, David Lowther and Mark Apperley have been instrumental in helping to form the views presented in this paper. Support from the EPSRC (Grant GR/N20881) is gratefully acknowledged. Comments from two reviewers were particularly helpful.

Notes

1. Scent has been defined by Chi (2002) and Pirolli (2002) as "The imperfect, subjective, perception of the value, cost or access path of information sources obtained from proximal cues such as Web links, or icons representing the content sources".

2. Card et al. (2001) state that "Such intermediate information has been referred to as "residue" by Furnas (1997). In keeping with foraging terminology, we have called this *scent*". We disagree with this suggested equivalence, as our definitions and Figure 5 make clear. Scent is a *perceived benefit*, based on a user's interpretation of sensitivity cues and residue as well as higher-order cognitive processes such as current strategy (e.g., exploring to form a mental model or moving as directly as possible to a target). Residue, in contrast, is an encoding (designed by the interaction designer and not by a user) of distant content.

References

Ahlberg, C. (1996). *Dynamic queries*. PhD dissertation, Chalmers University of Technology, Sweden.

Albinsson, P-A., & Morin, M. (2002). Visual exploration of communication in command and control. *IEEE, Proc. 6th Intl. Conference on Information Visualisation*, pp. 141–146.

Andrienko, N., & Andrienko, G. (2003). Informed spatial decisions through coordinated views. *Information Visualization, 2*(4), 270–285.

Apperley, M.D., Spence, R., & Gutwin, C. (1999). *The neighbourhood explorer*. Working paper 99/15, Department of Computer Science, University of Waikato: New Zealand.

Apperley, M.D., Spence, R., & Wittenburg, K. (2001). Selecting one from many: The development of a scalable visualization tool. *Proceedings HCC'01*, pp. 366–372.

Bertin, J. (1981). *Graphics and graphic information processing*. Berlin: Walter de Gruyter. (Translated from Bertin, J. (1977), *La graphique et le traitement graphique de l'information*. Paris: Flammarion).

Brayton, R.K., & Spence, R. (1980). *Sensitivity and optimization*. Amsterdam: Elsevier.

Card, S.K., Pirolli, P., van der Wege, M., Morrison, J.B., Reeder, R.W., Schraedley, P.K., & Boshart, J. (2001). Information scent as a driver of web behaviour graphs: Results of a protocol analysis method for web usability. *ACM, Proceedings CHI'2001*, pp. 498–505.

Chi, E.H., Pirolli, P., Chen, K., & Pitkow, J. (2001). Using information scent to model user information needs and actions on the web. *ACM, Proceedings CHI'01,* pp. 490–497.

Chi, E. (2002). Private communication.

Cleveland, W.S. (1985). *The elements of graphing data.* Monterey, CA: Wadsworth Advanced Books and Software.

Cowan, A., & Lowther, D. (2003). The application of sensitivity analysis in the design process for electromagnetic devices. *Proc. of the 14th Conference on Electromagnetic Field Computation, COMPUMAG, Saratoga, Saratoga Springs, N.Y., USA, June 2003, P16002.*

Dahlback, N. (1998). On spaces and navigation in and out of the computer. In N. Dahlback (Ed.). Exploring navigation: Towards a framework for design and evaluation of navigation in electronic spaces. *SICS Technical Report T98:01*.

Director, S.W., & Rohrer, R. (1969). The generalized adjoint network and network sensitivities. *IEEE Trans. Circuit Theory, CT-16,* 3, pp. 318–323.

Dourish, P., & Chalmers, M. (1994). Running out of space: Models of information navigation. *Proceedings HCI'94*.

Dourish, P. (1999). Where the footprints lead: Tracking down other roles for social navigation. In Munro et al. (Eds.), *Social navigation of information space*. London: Springer.

Field, G., & Apperley, M.D. (1990). Context and selective retreat in hierarchical menu structures. *Behaviour and Information Technology, 9*(2), 133–146.

Furnas, G.W. (1997). Effective view navigation. *ACM, Proceedings CHI'97*, pp. 367–374.

Gibson, J. (1979). *The ecological approach to visual perception.* Boston, MA: Houghton Mifflin.

Harper, R.H.R. (1999). Information that counts: A sociological view of information navigation. In Munro et al. (Eds.), *Social navigation of information space*. London: Springer.

Hill, W., Stead, L., Rosenstein, M., & Furnas, G. (1995). Recommending and evaluating choices in a virtual community of users. *ACM, Proceedings CHI'95*, pp. 194–201.

Hook, K., Benyon, D., & Munro, A. (2003). Editors' Introduction: Footprints in the snow. In Hook et al. (Eds.), *Designing information spaces: The social navigation approach*, pp. 1–13. London, Springer.

Jul, S., & Furnas, G.W. (1997). Navigation in electronic worlds: A CHI Workshop, New York. *ACM SIGCHI Bulletin, 29*, pp. 44–49.

Konstan, J.A., & Riedl, J. (2003). Collaborative filtering: Supporting social navigation in large crowded infospaces. In Hook et al. (Eds.). *Designing information spaces: The social navigation approach*, pp. 43–82. London: Springer.

Munro, A.J., Hook, K., & Benyon, D. (1999). *Social navigation of information space*. London: Springer.

Nielsen, J. (2000). Is navigation useful? Jakob Nielsen's Alertbox, January 9, 2000. http://www.useit.com/alertbox/20000109.html. (last accessed 6 August 2004).

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

Norman, D. (1988). *The psychology of everyday things*. New York: Basic Books.

Norman, K.L., & Chin, J.P. (1988). The effect of tree structure on search in a hierarchical menu selection system. *Behaviour and Information Technology, 7* (1), 51–65.

Pirolli, P. (2002). Private communication.
Pirolli, P., & Card, S. K. (1999), Information foraging. *Psychological Review, 106*(4), 643–675.
Qeli, E., Wiechert, W., & Freisleben, B. (2004). Visualizing time-varying matrices using multidimensional scaling and reorderable matrices. *IEEE, Proceedings Information Visualization IV04,* pp. 561–567.
Smith, A.J., Malik, Z., Nelder, J., & Spence, R. (2001). A visual interface for model fitting. *Quality and Reliability Engineering International, 17*, 85–91.
Snowberry, K., Parkinson, S.R., & Sisson, N. (1983). Computer display menus. *Ergonomics, 26* (7), 699–712.
Snowberry, K., Parkinson, S.R., & Sisson, N. (1985). Effects of help fields on navigating through hierarchical menu structure. *International Journal of Man-Machine Studies, 22*, 479–491.
Spence, R., & Drew, A.J. (1971). Graphical exploration in electrical circuit design and modelling. *NRC, Ottawa, Proc. Man-Computer Communication Seminar*.
Spence, R., & Tweedie, L.S. (1998). The attribute explorer: Information synthesis via exploration. *Interacting with Computers, 11*, 137–146.
Spotfire: www.spotfire.com [last accessed 7 August 2004].
Stephens, D.W., & Krebs, J.R. (1986). *Foraging theory*. Princeton, NJ: Princeton University Press.
Tauscher, L., & Greenberg, S. (1997). Revisitation patterns in World Wide Web navigation. *ACM, Proc. CHI'97*, pp. 399–406.
Tognazzini, B. (1992). *Tog on interface*. Reading, MA: Addison-Wesley.
Tufte, E.R. (1983). *The visual display of quantitative information.* Cheshire, Conn.: Graphics Press.
Tufte, E.R. (1990). *Envisioning Information*. Cheshire, Conn.: Graphics Press.
Tufte, E.R. (1997). *Visual explanations*. Cheshire, Conn.: Graphics Press.
Tweedie, L.A., Spence, R., Williams, D., & Bhogal, R. (1994). The attribute explorer. *ACM, Companion Proceedings CHI 1994.*
Wildbur, P., & Burke, M. (1998). *Information graphics*. London: Thames & Hudson.
Williamson, C., & Shneiderman, B. (1992). The dynamic homefinder: Evaluating queries in a real estate information exploration systems. *ACM, SIGIR'92*.
Wittenburg, K., Lanning, T., Heinrichs, M., & Stanton, M. (2001). Parallel bargrams for consumer-based information and choice. *Proceedings UIST'01*.
Wittenburg, K. (1997). Navigation and search) – what's the difference? Position paper. *Workshop on Navigation in Electronic Worlds* (see Jul, & Furnas, 1997).

ABOUT THE AUTHOR

Bob Spence is Professor Emeritus of Information Engineering, and Senior Research Investigator, at Imperial College London. His research interests include engineering design, for which he was awarded a higher doctorate, and human-computer interaction, for which he received a second higher doctorate. Current research interests include information visualization and information space navigation. Professor Spence is a Fellow of the Royal Academy of Engineering and a Fellow of the IEEE.
Contact: r.spence@imperial.ac.uc