

Space - the Final Chapter or Why Physical Representations are not Semantic Intentions

Andrew Dillon, Cliff McKnight and John Richardson

From: McKnight, C., Dillon, A. and Richardson, J. (eds) *Hypertext: A Psychological Perspective*. Chichester: Ellis Horwood, 1993.

Introduction

The term 'hypertext' evokes many images (e.g., nodes and links, semantic webs, non-linear access and so forth) but perhaps one of the most common is that of users struggling to find their way around a complex information space. As a result, navigation has become a subject of great interest to many researchers in the field. In this chapter we will discuss navigation through hypertext in terms of its relevance as a concept as much as its presence as an issue and try to draw lessons for design and research from the psychological work that has been carried out on navigation in physical space. We will attempt to show that while relevant to hypertext, discussion of navigation is prone to difficulty when researchers and designers misapply arguments and evidence from the physical domain to the semantic domain.

The extent of navigation difficulties in hypertext

Although there is a striking consensus that navigation is a difficulty for users of hypertext and frequent reference is made to "getting lost in hyperspace" (e.g., Conklin., 1987), empirical demonstrations of the problem are less than clear cut. As Bernstein.; (1991) put it:

while the so-called 'navigation problem' has come to dominate hypertext research, evidence for its existence and nature is distressingly thin. (p. 365)

However, much of the difficulty results from the traditional methodological problems in reading research of adequately measuring the reading process as opposed to the reading outcome (Dillon., 1992). In the absence of a suitable measure of process activities, researchers tend to infer from outcomes so, for example, the impact of information structure on reading speed, accuracy or comprehension is often explained in terms of process difficulties such as navigation (e.g., Monk.; Walsh.; and Dix.;, 1988).

If by the expression 'navigation difficulty' we accept the Elm.; and Woods.; (1985) definition of users not knowing how the information is organised, how to find the information they seek or even if that information is available then we can appreciate how such an approach may be justified. With paper documents there tend to be at least some standards in terms of organisation. With books, for example, contents pages are usually at the front, indices at the back and both offer some information on where items are located

in the body of the text. Concepts of relative position in the text such as 'before' and 'after' have tangible physical correlates. No such correlation holds with hypertext. If users perform slower or less accurately with a hypertext than with paper and claim to have had problems finding material then a conclusion of navigation difficulty may appear justified.

There is some direct empirical evidence in the literature to support the view that navigation in hypertext can be a problem. Edwards.; and Hardman.; (1989), for example, describe a study which required subjects to search through a specially designed hypertext. In total, half the subjects reported feeling lost at some stage (this proportion is deduced from the data reported). Such feelings were mainly due to "not knowing where to go next" or "not knowing where they were in relation to the overall structure of the document" (descriptors provided by the authors). Unfortunately, without direct comparison of ratings from subjects reading a paper equivalent we cannot be sure such proportions are solely due to using hypertext. However it is unlikely that many readers of paper texts do not know where they are in relation to the rest of the text.

McKnight.;, Dillon.; and Richardson.; (1990) proposed time spent browsing contents and indices as a suitable metric for navigation problems in a study of users retrieving information from a document on winemaking presented as hypertext, word processor file or paper. Their results also suggested that navigation is a difficulty for hypertext users, since users of both the word processor and paper versions spent significantly less time in these sections.

Indirect evidence comes from the numerous studies which have indicated that users have difficulties with a hypertext (e.g., Gordon.; *et al.* 1988; Monk ,Walsh and Dix.;, 1988). Hammond. and Allinson. (1989) speak for many when they say:

Experience with using hypertext systems has revealed a number of problems for users... First, users get lost... Second, users may find it difficult to gain an overview of the material... Third, even if users know specific information is present they may have difficulty finding it. (p. 294)

There are a few dissenting voices. Landow.; (1990, 1991), for example, describes the navigation problem in hypertext as a fallacy or pseudo-problem and claims that even discussing navigation can lead to false assumptions about hypertext usage. Brown.(1988) argues that:

although getting lost is often claimed to be a great problem, the evidence is largely circumstantial and conflicting. In some smallish applications it is not a major problem at all. (p. 2)

This quote is telling in several ways. The evidence for navigational difficulties is often circumstantial or at least inferential, as noted above. The applications in which Brown.; claims it is not a problem at all are, to use his word, "smallish" and this raises a crucial issue with respect to hypertext. When we speak of documents being so small that a reader

cannot 'get lost' in them or so large that navigation aids are required to use them effectively, the implication is that information occupies 'space' through which readers 'travel' or 'move'. Hammond.; and Allinson.; (1987) talk of the "travel metaphor" as a way of moving through a hypertext. Canter.;, Rivers.; and Storrs.; (1985) speak of "routes through" a database. Even the dissenters believe that the reader or user navigates through the document, the only disagreement being the extent to which getting lost is a regular and/or serious occurrence.

Psychological models of navigation

Navigation as an activity in itself is rarely seen as central to psychology and this is reflected in the subject matter and coverage of most textbooks. As Gould.; (1973) put it:

psychologists, in their concern with 'perception', have barely touched upon the investigation of mental images of geographic space, for many of their efforts have concentrated upon the physics and physiology of the senses, often within highly controlled laboratory conditions. (p. 185)

While it might be tempting to dismiss those comments as the product of their time, it is only justifiable to do so if we can point out significant developments in psychology's understanding of navigation in the last twenty years. Yet these have not emerged. However, although the activity of navigation may have been given short shrift by the discipline it is not hard to see how aspects relevant to the study of navigation are dealt with in work on spatial imagery, orientation, distance judgment, environmental perception and so forth. Furthermore, as Neisser.; (1976) remarks, geographers, planners and other professionals have always been interested in navigation and if, as Downs.; and Stea.; (1973) claim, "each writer ultimately turns to psychology for the answer" the activity cannot be said to have been overlooked (even if the answers have not been forthcoming).

The disparities of approaches to navigation render it difficult to draw together a cohesive theory but general agreements do exist. Tolman's (1948) paper on cognitive maps is frequently cited as seminal as it postulates the existence of a cognitive map, internalised in the human mind which is the analogue to the physical layout of the environment a view that is taken as axiomatic by most researchers despite its inherent presumptions about human cognition. According to Tolman, information impinging on the brain is:

worked over and elaborated... into a tentative cognitive-like map of the environment indicating routes and paths and environmental relationships... (p. 192)

Recent experimental work takes the notion of some form of mental representation of the environment for granted, concerning itself more with how such maps are formed and manipulated. We have proposed elsewhere (Dillon.;, McKnight.; and Richardson.;, 1990; McKnight.;, Dillon.; and Richardson.;, 1991) that navigation in general can be

conceptualised in psychological terms as involving four levels of representation: schemata, landmarks, routes and surveys and that several of these levels are of direct relevance to hypertext design. In the following section we review this approach which draws heavily on schema theory and consider how alternative views on navigation might shed further light on the concept.

Schema theory as an explanatory framework

Schema theory provides a convenient explanatory framework for the general knowledge humans seem to possess of activities, objects, events and environments. Regardless of its actual truth or explanatory power over alternative views of cognition such as prototype or feature-count theories (see e.g., Anderson., 1980), we can convincingly postulate the existence of some form of general knowledge of the world that aids humans in navigation tasks.

It seems obvious, for example, that we must possess schemata of the physical environment we find ourselves in if we are not to be overwhelmed by every new place we encounter. Presumably acquired from exposure to the world around us, schemata can be conceptualised as affording a basic orienting frame of reference to the individual. Thus, we soon acquire schemata of towns and cities so that we know what to expect when we find ourselves in one: busy roads, numerous buildings, shopping areas, people, etc. According to Downs.; and Stea.; (1977) such frames of reference exist at all levels of scale from looking at the world in terms of east and west or First and Third Worlds, to national distinctions between north and south, urban and rural and so on down to local entities like buildings and neighbourhoods.

In employing schema theory as our explanatory framework it is worth making a distinction between what Brewer.; (1987) terms "global" and "instantiated" schemata. The global schema is the basic or raw knowledge structure. Highly general, it does not reflect the specific details of any object or event (or whatever knowledge type is involved). The instantiated schema however is the product of adding specific details to a global schema and thereby reducing its generality. An example will make this clearer. In orienting ourselves in a new environment we call on one or more global schemata (e.g., the schema for city or office building). As we proceed to relate specific details of our new environment to this schema we can be said to develop an instantiated schema which is no longer general but is not sufficiently complete to be a model or map of the particular environment in which we find ourselves.

Global schemata remain to be used again as necessary. Instantiated schemata presumably develop in detail until they cease to be accurately described as schematic or are discarded when they serve no further purpose. In the above example, if we leave this environment after a short visit we are likely to discard the instantiated schema we formed, but if we stay or regularly return we are likely to build on this until we have a model of the environment. This is a very simple sketch and leaves many questions unanswered, such as what details are required to make a global schema become an instantiated one or if we always recall a specific detail about a place we visited but nothing else does this memory

constitute an instantiated schema? We will not attempt answers to these questions here but recognise them as typical theoretical issues for schema theory to address.

While schemata are effective orienting guides, in themselves they are limited. In particular they fail to reflect specific instances of any one environment and provide no knowledge of what exists outside of our field of vision. As such, they provide the basic knowledge needed to interact with an environment but must be supplanted by other representations if we are to plan routes, avoid becoming lost or identify short cuts activities in which humans seem frequently to engage.

Levels of schema instantiation: landmarks, routes and surveys

The second level of representation (but first stage of instantiation for the schema) proposed for navigation is knowledge of landmarks, a term used to describe any features of the environment which are relatively stable and conspicuous. Thus we recognise our position in terms relative to these landmarks, e.g., our destination is near building X or if we see statue Y then we must be near the railway station and so forth. This knowledge provides us with the skeletal framework on which we build our cognitive map.

The third level of representation is route knowledge which is characterised by the ability to navigate from point A to point B, using whatever landmark knowledge we have acquired to make decisions about when to turn left or right. With such knowledge we can provide others with effective route guidance, e.g., "Turn left at the traffic lights and continue on that road until you see the large church on your left and take the next right there..." and so forth. Though possessing route knowledge a person may still not really know much about her environment. A route might be non-optimal or even totally wasteful.

The fourth level of representation is survey (or map) knowledge. This allows us to give directions or plan journeys along routes we have not directly traveled as well as describe relative locations of landmarks within an environment. It allows us to know the general direction of places, e.g., "westward" or "over there" rather than "left of the main road" or "to the right of the church". In other words it is based on a world frame of reference rather than an egocentric one.

Navigational knowledge development

By navigational knowledge development we do not mean the acquisition of spatial knowledge and world views as they are discussed in developmental psychology but rather how adults become familiar with their environments. However there exists an interesting literature on children's spatial knowledge development that has parallels with this work (see e.g., Hart.A.; and Moore.T.;, 1973), some of which will be alluded to in this discussion.

Current thinking is dominated by the view that the last three levels of representation: landmark, route and survey knowledge are points on a continuum rather than discrete

forms. The assumption is that each successive stage represents a developmental advance towards an increasingly accurate or sophisticated world view. Certainly this is an intuitively appealing account of our own experiences when coming to terms with a new environment or comparing our knowledge of one place with another and has obvious parallels with the psychological literature which often assumes invariant stages in cognitive development but it might not be so straightforward.

Obviously landmark knowledge on its own is of little use for complex navigation and both route and survey knowledge emerges from it as a means of coping with the complexity of the environment. However it does not necessarily follow that given two landmarks the next stage of knowledge development is acquiring the route between them or that once enough route knowledge is acquired it is replaced by or can be formed into survey knowledge. Experimental investigations have demonstrated that each form of representation is optimally suited for different kinds of tasks (see e.g., Thorndyke.; and Hayes-Roth.;, 1982; Wetherell.;, 1979). Route knowledge is cognitively simpler than survey knowledge but suffers the drawback of being virtually useless once a wrong step is taken (Wickens.;, 1984). Route knowledge, because of its predominantly verbal form, might suit individuals with higher verbal than spatial abilities, while the opposite might be the case for survey knowledge. Furthermore, age differences in navigational knowledge have been demonstrated which show older males relying more on landmark rather than route knowledge (Francescato.; and Mebane.;, 1973), a complete reversal of what we would expect if the invariant stage development model was true.

Thus while the knowledge forms outlined here are best seen as points on a continuum and a general trend to move from landmark to survey knowledge via route knowledge may exist, task dependencies and cognitive ability factors mediate such developments and suggest that an invariant stage model may not be the best conceptualisation of the findings.

Other issues

While this approach might be seen as a useful condensation of current thinking it is by no means indisputable and leaves several issues unresolved. We have made little reference to knowledge of how to behave in certain environments which must be tied at some level to our navigational schemata. Roads must be crossed in certain ways, e.g., at pedestrian crossings or when there is no traffic, or you must pay if you want to use public transport. It may be slightly artificial (not to mention dangerous) to separate such knowledge from our schemata for places. In this sense the frame of reference is close to the concept of script (Schank.; and Abelson.;, 1976).

Furthermore, as Shum.; (1990) noted when describing the likely content of cognitive maps, most human factors work has addressed only the descriptive attributes such maps possess and said little or nothing about their evaluative attributes such as "is this place worth visiting?" We could extend this to say that such value judgments are only part of what seems to be missing from the standard treatment of cognitive maps. We could consider affective issues to be associated with places such that when we enter an

environment or even think about entering it, associated emotional responses might come into play. Too little is known about these aspects of human response to environments and their impact on navigation.

Cohen.; (1989) proposes that three variables are important to consider in understanding memory for places: scale, complexity and familiarity. Her argument is that essentially, variation in these factors alters the nature of the navigation task being performed. If this is the case, the nature of the representation employed by the human navigator may well alter and simple classifications such as "routes are represented procedurally and maps are represented propositionally" (see e.g., Bartram.; and Smith.;, 1984) are unlikely to tell the full story. Cohen.; argues that mental models (as used by Johnson-Laird.;, 1983) or analogues could also provide representational formats for navigational information, to which we could add representational formats such as episodic memory if we want to account for some other possible phenomena.

Our view is that the precise nature of the representation is less important to workers in the field of interactive technology than the insights any theory or model of navigation provides. To this extent we propose that a model based on schema theory and including landmarks, routes and surveys as instantiations of basic knowledge is of some utility in considering the design of electronic information spaces. In the following sections we develop this argument.

Navigation in the paper domain

Documents as structures

If we are seriously to consider hypertext documents as navigable spaces there is no logical reason for denying paper texts the same status. There is evidence to suggest that from a reader's perspective at least, such a view is meaningful. For example, studies of regular readers' perceptions of journal articles (Dillon.;, Richardson.; and McKnight.;, 1989) and software manuals (Dillon.;, 1991a) have shown that such readers conceptualise documents as possessing a prototypical form or structure that aids location of material, suggesting at least that the idea of navigating through such structures is a valid notion.

Unfortunately, the term 'structure' is used in at least three distinct ways by different researchers and writers in this field. Conklin.; (1987) talks of structure being imposed on what is browsed by the reader, i.e. the reader builds a structure to gain knowledge from the document. Trigg.H.; and Suchman.A.; (1989) refer to structure as a representation of convention, i.e. it occurs in a text form according to the expected rules a writer follows during document production. Hammond.; and Allinson.; (1989) offer a third perspective, that of the structure as a conveyer of context. For them, there is a naturally occurring structure to any subject matter that holds together the 'raw data' of that domain.

In reality, there is a common theme to all these uses. They are not distinct concepts sharing the same name but different aspects or manifestations of the same concept. The

main role of structure seems to differ according to the perspective from which it is being discussed: the writer's or the reader's, and the particular part of the reading/writing task being considered. Thus the structure of a document can be a convention to both the writer, so that he conforms to expectations of format, and to the reader, so he knows what to expect. It can be a conveyer of context mainly to the reader so he can infer from, and elaborate on, the information provided, but it might be employed by a skilled writer with the intention of provoking a particular response in the reader. Finally, it can be a means of mentally representing the contents to both the reader so he grasps the organisation of the text and to the author so that he can appropriately order his delivery.

It can be seen from the comments of subjects (as readers) in the journal and manual usage studies cited above that structure is a concept for which the meanings described seem to apply with varying degrees of relevance. Certainly the notion of structure as convention seems to be perceived by readers of journal articles, while the idea of structure supporting contextual inference seems pertinent to users of software manuals. Beyond these manifestations, research in the domain of linguistics and discourse comprehension lends strong support to the concept of structure as a basic component in the reader's mental representation of a text.

The theory of discourse comprehension proposed by van Dijk.A.; and Kintsch.; (1983) places great emphasis on text structure. According to this theory, readers acquire (through experience) schemata, which van Dijk.A.; and Kintsch.; term 'superstructures', that facilitate comprehension of material by allowing readers to predict the likely ordering and grouping of constituent elements of a body of text. To quote van Dijk.A.; (1980):

a superstructure is the schematic form that organises the global meaning of a text. We assume that such a superstructure consists of functional categories... [and]... rules that specify which category may follow or combine with what other categories. (p. 108)

In addition to categories and functional rules, van Dijk.A.; adds that a superstructure must be socioculturally accepted, learned, used and commented upon by most adult members of a speech community.

They have applied this theory to several text types. For example, with respect to newspaper articles they describe a schema consisting of headlines and leads (which together provide a summary), major event categories each of which is placed within a context (actual or historical), and consequences. Depending on the type of newspaper (e.g., weekly as opposed to daily, tabloid as opposed to quality, etc.) one might expect elaborated commentaries and evaluations. Experiments by Kintsch.; and Yarborough.; (1982) showed that articles written in a way that adhered to this schema resulted in better grasp of the main ideas and subject matter (as assessed by written question answering) than ones which were re-organised to make them less schema conforming.

The van Dijk.A.; and Kintsch.; theory has been the subject of criticism from some cognitive scientists. Johnson-Laird.; (1983), for example, takes exception to the idea of any propositional analysis providing the reader with both the basic meaning of the words

in the text and the significance of its full contents. For him, at least two types of representational format are required to do this. He provides evidence from studies of people's recall of text passages that it is not enough to read a text correctly (i.e. perform an accurate propositional analysis) to appreciate the significance of that material. He proposes what he terms mental models as a further level of representation that facilitates such understanding. Subsequent work by Garnham.; (1987) lends further support to the insufficiency-of-propositions argument in comprehension of text.

The differences between Johnson-Laird. and van Dijk.A.; are mainly a reflection of the differences between the psychologist's and the linguist's views of how people comprehend discourse. From the perspective of the human factors practitioner it is not clear that either theory of representation format is likely to lead to distinct (i.e. unique) predictions about electronic text. Both propose that some form of structural representation occurs it is just the underlying cognitive form of this representation that is debated. The similarity of their views from the human factors perspective is conveyed in this quote from Johnson-Laird.; where he states that mental models:

appear to be equally plausible candidates for representing the large-scale structure of discourse the skeletal framework of events that corresponds to the 'plot of the narrative', the 'argument' of a non-fiction work and so on. Kintsch.; and van Dijk.A.'s proposal that there are macro rules for constructing high-level representations could apply *mutatis mutandis* to mental models. (p. 381)

In other words, the issue is not if, or even how, readers acquire a structural representation of texts they read (these are accepted as givens) but what form such structures take: propositions or mental models? This is not an issue of direct concern to the designer of hypertexts, what is of importance is the provision and support of document structures that aid accurate structural representations (of whatever form) in the reader's mind.

Applying global schemata of documents

At a more global level Dillon.; (1991b) tested readers' ability to impose structure on randomly presented paragraphs and sentences of text. In the first of two experiments, subjects were given a selection of paragraphs from academic journal articles and asked to organise them into one article as fast as they could. To avoid referential continuity, every second paragraph was removed. In one condition headings were provided, in the other they were absent. The results indicated that readers had little difficulty piecing the article together into gross categories of Introduction, Method, Results and Discussion (over 80% accuracy at this level) but had difficulties distinguishing the precise order at the within-section level. When provided with headings, subjects formed the same major categories but were less accurate in placing second level headings in the correct section. This suggests that experienced journal readers are capable of distinguishing isolated paragraphs of text according to their likely location within a complete article with respect to the major categories. Interestingly this could be done without resorting to reading every word or attempting to understand the subject matter of the paper.

In the second study, subjects read a selection of paragraphs from two articles on both paper and screen and had to place each one in the general section to which they thought it belonged (Introduction, Method, Results or Discussion). Again subjects showed a high degree of accuracy (over 80%) with the only advantage to paper being speed (subjects were significantly faster at the 5% level in the paper condition) which is probably explicable in terms of image quality variables and manipulation differences. Taken together, these results suggest that readers do have a model of the typical journal article that allows them to gauge accurately where certain information is located. This model does not seem to be affected by presentation medium.

In this way the schema (superstructure) constitutes a set of expectancies about a text's usual contents and how they are grouped and positioned relative to each other. Obviously, in advance of actually reading the text we cannot have much insight into anything more specific than this, but the generality of organisation within the multitude of texts we read in everyday life affords stability and orientation in what is otherwise a complex information environment.

Cognitive maps of document spaces: instantiating the global schema

If picking up a new book can be compared to a stranger entering a new town (i.e. we know what each is like on the basis of previous experience and have expectancies of what we will find) how do we proceed to develop our map of the information space?

To use the analogy of the navigation in physical space we would expect that generic structures such as indices, contents, chapter headings and summaries can be seen as landmarks that provide readers with information on where they are in a text, just as signposts, buildings and street names aid navigation in physical environments. Thus when initially reading a text we might notice that there are numerous figures and diagrams in certain sections, none in others, or that a very important point or detail is raised in a section containing a table of numerical values. In fact, readers often claim to experience such a sense of knowing where an item of information occurred in the body of the text even if they cannot recall that item precisely and there is some empirical evidence to suggest that this is in fact the case.

Rothkopf, Z.; (1971) carried out an experiment to test whether such occurrences had a basis in reality rather than resulting from popular myth supported by chance success. He asked people to read a 12-page extract from a book with the intention of answering questions on content afterwards. What subjects did not realise was that they would be asked to recall the location of information in the text in terms of its occurrence both within the page (divided into eighths) and the complete text (divided into quarters). The results showed that incidental memory for locations within any page and within the text as a whole were more accurate than chance, i.e. people could remember location information even though they were not asked to. There was also a positive correlation between location of information at the within-page level and accuracy of question answering.

There have been several follow-up studies by Rothkopf, Z.; and by other investigators into this phenomenon. Zechmeister.; and McKillip.; (1972) had subjects read eight pages of text typed into blocks with four blocks per page. Subjects were asked to read the text before being tested on it. The test consisted of fill-in-the-blank questions, confidence ratings on their answers and location of the answer on the page. Again, an effect for knowledge of location was observed which was correlated to accuracy of answers, suggesting that memory for location and for content are independent attributes of memory that can be linked for mnemonic purposes. Interestingly no interaction of memory for location and confidence in answer was found. Further work by Zechmeister.; *et al.* .;.;(1975) and by Lovelace and Southall (1983) confirm the view that memory for spatial location within in body of text is reliable even if it is generally limited. Simpson.; (1990) has replicated this for electronic documents.

In the paper domain at least, the analogy with navigation in a physical environment is of limited applicability beyond the level of landmark knowledge. Given the fact that the information space is instantly accessible to the reader (i.e. she can open a text at any point) the necessity for route knowledge, for example, is lessened (if not eliminated). To get from point A to point B in a text is not dependent on taking the correct course in the same way that it is in a physical three-dimensional environment. The reader can jump ahead (or back), guess, use the index or contents or just page serially through. Readers rarely rely on just one route or get confused if they have to start from a different point in the text to go to the desired location, as would be the case if route knowledge was a formal stage in their development of navigational knowledge for texts. Once you know the page number of an item you can get there as you like. Making an error is not as costly as it is in the physical world either in terms of time or effort. Furthermore, few texts are used in such a way as to require that level or type of knowledge.

One notable exception to this might be the knowledge involved in navigating texts such as software manuals or encyclopedias which can consist of highly structured information chunks that are inter-referenced. If for example a procedure for performing a task references another part of the text it is conceivable that a reader may only be able to locate the referenced material by finding the section that references it first (perhaps because the index is poor or she cannot remember what it is called). In this instance one could interpret the navigation knowledge as being a form of route knowledge. However, such knowledge is presumably rare except where it is specifically designed into a document as a means of aiding navigation along a trouble-shooting path.

A similar case can be made with respect to survey knowledge. While it seems likely that readers experienced with a certain text can mentally envisage where information is in the body of the text, what cross-references are relevant to their purpose and so forth, we must be careful that we are still talking of navigation and not changing the level of discourse to how the argument is developed in the text or the ordering in which points are made. Without doubt, such knowledge exists, but often it is not purely navigational knowledge but an instantiation of several schemata such as domain knowledge of the subject matter, interpretation of the author's argument, and a sense of how this knowledge is organised that come into play now. This is not to say that readers cannot possess survey type

knowledge of a text's contents, rather it is to highlight the limitations of directly mapping concepts from one domain to another on the basis of terminology alone. Just because we use the term navigation in both situations does not mean that they are identical activities with similar patterns of development. The simple differences in applying findings from a three-dimensional world (with visual, olfactory, auditory and powerful tactile stimuli) to a two-dimensional text (with visual and limited tactile stimuli only) and the varying purposes to which such knowledge is put in either domain are bound to have a limiting effect.

It might be that rather than route and survey knowledge, a reader develops a more elaborated analogue model of the text based on the skeletal framework of landmark knowledge outlined earlier. Thus, as familiarity with the text grows, the reader becomes more familiar with the various landmarks in the text and their inter-relationships. In effect the reader builds a representation of the text similar to the survey knowledge of physical environments without any intermediary route knowledge but in a form that is directly representative of the text rather than physical domain. This is an interesting empirical question and one that is far from being answered by current knowledge of the process of reading.

Navigation in electronic space

Schemata and models

The concept of a schema for an electronic information space is less clear-cut than for physical environments or paper documents. As we have stated elsewhere (see e.g., McKnight.; Dillon.; and Richardson.;, 1991) computing technology's short history is one of the reasons for this but it is also the case that the medium's underlying structures do not have equivalent transparency. With paper, once the basic *modus operandi* of reading is acquired (e.g., page-turning, footnote identification, index usage and so forth) it retains utility for other texts produced by other publishers, other authors and for other domains. With computers, manipulation of information can differ from application to application within the same computer, from computer to computer and from this year's to last year's model. Thus using electronic information is often likely to involve the employment of schemata for systems in general (i.e. how to operate them) in a way that is not essential for paper-based information.

The qualitative differences between the schemata for paper and electronic documents can easily be appreciated by considering what you can tell about either at first glance. We have outlined the information available to paper text users in the section on paper schemata above. When we open a hypertext document, however, we do not have the same amount of information available to us. We are likely to be faced with a welcoming screen which might give us a rough idea of the contents (i.e. subject matter) and information about the authors/developers of the document but little else. It is two-dimensional, gives no indication of size, quality of contents, age (unless explicitly stated) or how frequently it has been used (i.e. there is no dust or signs of wear and tear on it

such as grubby finger-marks or underlines and scribbled comments). At the electronic document level, there is usually no way of telling even the relative size without performing some 'query operation'. Such a query operation will usually return a size in kilobytes and will therefore convey little meaning to the average reader.

Performing the hypertext equivalent of opening up the text or turning the page offers no assurance that expectations will be met since many hypertext documents offer unique structures (intentionally or otherwise). At their current stage of development it is likely that users/readers familiar with hypertext will have a schema that includes such attributes as linked nodes of information, non-serial structures, and perhaps, potential navigational difficulties! The manipulation facilities and access mechanisms available in hypertext will probably occupy a more prominent role in their schema for hypertext documents than they will for readers' schemata of paper texts since they differ from application to application. As yet, empirical evidence for such schemata is lacking.

The fact that hypertext offers authors the chance to create numerous structures out of the same information is a further source of difficulty for users or readers. Since schemata are generic abstractions representing typicality in entities or events, the increased variance of hypertext implies that any similarities that are perceived must be at a higher level or must be more numerous than the schemata that exist for paper texts.

It seems therefore that users' schemata of hypertext environments are likely to be 'informationally leaner' than those for paper documents. This is attributable to the recent emergence of electronic documents and comparative lack of experience interacting with them as opposed to paper texts for even the most dedicated users. The current lack of standards in the electronic domain compared to the rather traditional structures of many paper documents is a further problem for schema development.

Acquiring a cognitive map of the electronic space

The roots of navigation problems in electronic space lie in the literature on users interacting with non-hypertext databases and documents as well as with menu-driven interfaces where it has been repeatedly shown that when users make an incorrect selection at a deep level they tend to return to the start rather than the menu at which they erred and that the actual-to-minimum ratio for screens of information accessed in a successful search is 2:1 (Lee.; *et al.*, .;.;;1984).

Research by Snowberry.,; Parkinson.; and Sisson.; (1985) indicates that the main source of difficulty in menu navigation is the relatively weak associations users have between category descriptors at the highest level of menu and the desired information at the lower. This is a fault of design where little attempt is made to identify the user's conceptualisation of the information space. Significantly enough, Lee.; *et al.*, .;.;;(1984) discovered considerable variation among experts in terms of what they believe constitutes a 'good' or well-organised menu.

In terms of the model of navigational knowledge described above we should not be surprised by such findings. They seem to be classic manifestations of behaviour based on limited knowledge. For example, returning to the start upon making an error at a deep level in the menu suggests the absence of survey-type knowledge and a strong reliance on landmarks (e.g., the start screen) to guide navigation. It also lends support to the argument about route knowledge that it becomes useless once a wrong turn is made. Making 'journeys' twice as long as necessary is a further example of the type of behaviour expected from people lacking a mental map of an environment and relying on landmark and route knowledge only to find their way.

Jones.P.; and Dumais.T.; (1986) empirically tested spatial memory over symbolic memory for application in the electronic domain, citing the work of Rothkopf.Z.; and others as indicators that such memory might be important. In a series of three experiments subjects simulated filing and retrieval operations using name, location or a combination of both stimuli as cues. Like the preceding work on texts they found that memory for location is above chance but modest compared to memory for names and concluded that it may be of limited utility for object reference in the electronic domain.

Therefore, we know that navigational difficulties exist where users need to make decisions about location in an electronic information space. There seems to be some evidence that the first stage of knowledge about navigation is of the landmark variety and that the organising principles on which the information structure is built are important. We now turn to the more specific evidence for hypertext.

Acquiring a cognitive map of a hypertext document

McKnight.; Dillon.; and Richardson.; (1990) found that subjects reading hypertext spent significantly greater proportions of time in the index/contents sections of documents than readers using paper or word processor files, indicating a style of interaction based on jumping into parts of the text and returning to base for further guidance a seemingly sub-optimal style for hypertext and concluded from this that effective navigation was difficult for non-experienced users of a hypertext document.

Once more this is a classic example of using landmarks in the information space as guidance. Subjects in the linear conditions (paper and word processor versions) seemed much happier to browse through the document to find information, highlighting their confidence and familiarity with the structure presented to them. Similar support for the notion of landmarks as a first level of navigational knowledge development are provided by several of studies which have required subjects to draw or form maps of the information space after exposure to it (e.g., Simpson.; and McKnight.;, 1990). Typically, subjects can group certain sections together but often have no idea where other parts go or what they are connected to.

Unfortunately it is difficult to chart the development of navigational knowledge beyond this point. Detailed studies of users interacting with hypertext systems beyond single experimental tasks and gaining mastery over a hypertext document are thin on the

ground. Edwards.; and Hardman.; (1989) claim that they found evidence for the development of survey-type navigational knowledge in users exposed to a strictly hierarchical database of 50 screens for a single experimental session lasting, on average, less than 20 minutes. Unfortunately the data are not reported in sufficient detail to critically assess such a claim but it is possible that given the document's highly organised structure, comparatively small size and the familiarity of the subject area (leisure facilities in Edinburgh), such knowledge might have been observed. Obviously this is an area that needs further empirical work.

While it is clear that empirical work on hypertext is limited, numerous designers and researchers have considered the navigation issues in less experimental ways, often without concerning themselves with the development of mental representations of the information space. In the following section we discuss the design of suitable maps, browsers and landmarks for users to aid navigation.

Providing navigational information: browsers, maps and structural cues

A graphical browser is a stylised representation of the structure of the database aimed at providing the user with an easy-to-understand map of what information is located where. According to Conklin.; (1987) graphical browsers are a feature of a "somewhat idealized hypertext system", recognising that not all existing systems utilise browsers but suggesting that they are desirable. The idea behind a browser is that the document can be represented graphically in terms of the nodes of information and the links between them, and in some instances, that selecting a node in the browser would cause its information to be displayed.

It is not difficult to see why this might be useful. Like a map of a physical environment, it shows the user what the overall information space is like, how it is linked together and consequently offers a means of moving from one information node to another. Indeed Monk.;, Walsh.; and Dix.; (1988) have shown that even a static, non-interactive graphical representation is useful. However, for richly interconnected material or documents of a reasonable size and complexity, it is not possible to include everything in a single browser without the problem of presenting 'visual spaghetti' to the user. In such cases it is necessary to represent the structure in terms of levels of browsers, and at this point there is a danger that the user gets lost in the navigational support system!

Some simple variations in the form of maps or browsers have been investigated empirically. In a non-hypertext environment Billingsley.; (1982) had subjects select information from a database aided by an alphabetical list of selection numbers, a map of the database structure or no aid. The map proved superior, the no aid group performing worst.

In the hypertext domain a number of studies by Simpson (1989) have experimentally manipulated several variables related to structural cues and position indicators. She had subjects perform a series of tasks on articles about houseplants and herbs. In one experiment she found that a hierarchical contents list was superior to an alphabetic index

and concluded that users are able to use cues from the structural representation to form maps of the document. In a second study she reported that users provided with a graphical contents list showing the relationship between various parts of the text performed better than users who only had access to a textual list. Making the contents lists interactive (i.e. selectable by pointing) also increased navigational efficiency.

Manipulating 'last card seen' markers produced mixed results. It might be expected that such a cue would be advantageous to all users but Simpson reported that this cue seemed of benefit only during initial familiarisation periods and for users of non-interactive contents lists. Further experiments revealed that giving users a record of the items they had seen aided navigation, much as would be expected from the literature on physical navigation which assumes that knowledge of current position is built on knowledge of how you arrived there (Canter., 1984). In general, Simpson found that as accuracy of performance increased so did subjects' ability to construct accurate post-task maps of the information space using cards.

Such work is important to designers of hypertext systems. It represents a useful series of investigations into how 'contents pages' for hypertext documents should be designed. Admittedly, it concerned limited tasks in a small information space but such studies are building blocks for a fuller understanding of the important issues in designing hypertext systems. As always, more research needs to be done.

Several writers have suggested novel navigational tools for use in hypertext. However, such tools are rarely, if ever, evaluated but thrown into the designer's tool-box for future possible use. For example, Utting.; and Yankelovich.; (1989) provide a good review of various systems and a description of the rationale behind the "Web View" navigation aid but all is done in terms of a hypothetical scenario and there is no serious attempt at evaluation. Gloor.A.; (1991) offers us the "Cybermap", a form of hierarchical overview system constructed using automatic indexing and clustering techniques. The Cybermap has some potentially useful features such as dynamic links based on user profile and reading history. However, the algorithm used to generate the higher levels of the map (termed "hyperdrawers") imposes the restriction that all hyperdrawers should contain approximately the same number of nodes which seems extremely unrealistic and needs empirical justification before it can seriously be advocated as a design target.

Lai.; and Manber.; (1991) offer a means of "flying" through hypertext, a technique analogous to flipping the pages of a book. They stress that this is intended as an additional tool rather than as a replacement for other navigation or manipulation techniques. However, like the Cybermap, until flying is subjected to user evaluation its real usefulness cannot be known. Methods of navigation in a speech—only application have even been proposed (e.g., Arons.;, 1991) but such papers are probably most useful for the questions they raise rather than the answers they provide.

Navigating the semantic space

One aspect of the whole navigation issue that often appears overlooked in the hypertext literature is that of the semantic space of a text or electronic document. In other words, to what extent does a user or reader need to find his way about the argument that an author creates as opposed to, or distinct from, navigating through the structure of the information?

It is probably impossible to untangle these aspects completely. We noted earlier, in the section on readers' memory for spatial location on pages, that there was a correlation between memory for location and comprehension. This is attributed to the fact that they are independent aspects of memory which are capable of being linked for mnemonic purposes. In other words, memories may consist of a constellation of attributes in which the recall of any one attribute is facilitated by the recall of others.

One complication is that while we can easily compare how near ideas are in terms of location in a structure, i.e. how many links exist between two nodes or the number of selections/button presses need to be made to access node Z from node A, we cannot offer a similar measure of semantic distance. The extent to which two ideas are related may seem intuitively easy to assess but is unlikely to have a broadly agreed quantifiable metric (see e.g., the work of Osgood.E., Suci.J.; and Tannenbaum.H., 1957 and Kelly.A., 1955).

Ultimately, we believe the idea of directly navigating semantic space has to be spurious. Semantic space is an abstract psycholinguistic concept which cannot be directly observed, only represented by way of alternative instantiations. By definition, semantic space is n-dimensional and practically unbounded. In order to visualise the semantic space it needs to be given physical representation and in so doing, it becomes at most three-dimensional (though more often two-dimensional) and physically bounded. In this form it is easy to see how concepts such as navigation appear relevant and thus we may talk of moving through semantic space in a manner equivalent to navigating physical environments.

In effect we cannot navigate semantic space, at least not in the way we navigate physical environments, we can only navigate the physical instantiations that we develop of the semantic space. In this case, it is meaningful to utilise the concept of navigation in the manner outlined above (with all the inherent limitations) but we must be clear that here we are not navigating through, or on the basis of, semantics. Rather, we are imposing a physical structure on the semantics and expecting people to employ cues for distance, size and form in a manner based on physics that can have an effect on exposure to issues at the level of semantics.

Thus, any presentation medium offers its own physical limitations on the representations available and authors can impose further on these to provide an instantiation of semantic space in the form of an article or document. In navigating around the document the reader is exposed to the medium-constrained physical manifestation of the author's semantic space. The reader can never demonstrate absolute grasp of or conformity to the author's semantics. At best, a reader can replicate or demonstrate skill in constructing similar

physical representations of that semantic space and it is this ability that is regularly tested in education, training and learning scenarios.

Perhaps to all intents and purposes such a distinction is unnecessarily fine, but it is important to avoid falling into the trap of viewing the structure of arguments and ideas as equivalent to their physical manifestation in the presentation medium. This is too simplistic. There is a relationship, but it is not so straightforward; and blurring the distinction leads to the false impression that navigation can be meaningfully talked about in semantic terms or that simply offering trails and backtracking facilities helps readers grasp the contents of an argument and thereby automatically improves reader comprehension levels. These are empirical issues and we emphasise again, it is important that such issues are investigated experimentally.

Conclusion

Navigation through hypertext documents is an issue worthy of attention from researchers and designers. The psychology of navigation, based as it is on studies of the human information processor interacting with a physical environment, offers some insights into how we acquire relevant information and develop knowledge of our environments, but does not provide a complete set of findings that are directly transferable to the electronic domain.

What are needed are well-controlled experiments which examine the best means of supporting navigation through large and complex information spaces. The field is replete with claims for wonderful interface styles and metaphors without the commensurate evaluation of such claims which is the hallmark of a true user-centred design process.

The issue is clouded by the confusion of terms such as information space and semantic space and the mistaken assumption that a document's physical instantiation through a presentation medium is equivalent to the semantic space an author intended to convey. Clearer definition of terms and an appreciation of the role of the medium in the message are not new recommendations, but then hypertext is nothing if not a new vehicle that raises old issues.

References

Anderson, J. R. (1980) *Cognitive Psychology and its Implications*. San Francisco: W. H. Freeman.

Arons, B. (1991) Hyperspeech: navigating in speech-only hypermedia. In *Hypertext '91*, Proceedings of the Third ACM Conference on Hypertext. New York: The Association for Computing Machinery. 133—146.

- Bartram, D. and Smith, P. (1984) Everyday memory for everyday places. In J. Harris and P. Morris (eds.) *Everyday Memory, Actions and Absentmindedness*. London: Academic Press.
- Bernstein, M. (1991) Position statement for Panel on Structure, Navigation and Hypertext: The Status of the Navigation Problem. In *Hypertext '91*, Proceedings of the Third ACM Conference on Hypertext. New York: The Association for Computing Machinery. 365—366.
- Billingsley, P. (1982) Navigation through hierarchical menu structures: does it help to have a map? *Proceedings of the Human Factors Society 26th Annual Meeting*. Santa Monica, CA: Human Factors Society. 103—107.
- Brewer, W. (1987) Schemas versus mental models in human memory. In I. P. Morris (ed.) *Modelling Cognition*. Chichester: John Wiley and Sons. 187—197.
- Brown, P. (1988) Hypertext the way forward. In J. C. van Vliet (ed.) *Document Manipulation and Typography*. Cambridge: Cambridge University Press. 183—191.
- Canter, D. (1984) Wayfinding and signposting: penance or prosthesis? In R. Easterby and H. Zuraga (eds.) *Information Design*. Chichester: Wiley and Sons. 245—264.
- Canter, D., Rivers, R. and Storrs, G. (1985) Characterising user navigation through complex data structures. *Behaviour and Information Technology*, 4(2), 93—102.
- Cohen, G. (1989) *Memory in the Real World*. London: Lawrence Erlbaum Associates.
- Conklin, J. (1987) Hypertext: an introduction and survey. *Computer*, September, 17—41.
- van Dijk, T. A. (1980) *Macrostructures*. Hillsdale, NJ: Lawrence Erlbaum Associates.
- van Dijk, T. A. and Kintsch, W. (1983) *Strategies of Discourse Comprehension*. London: Academic Press.
- Dillon, A. (1991a) Requirements analysis for hypertext applications: the why, what and how approach. *Applied Ergonomics*, 22(4), 458—462.
- Dillon A. (1991b) Readers' models of text structures: the case of academic articles. *International Journal of Man-Machine Studies*, 35, 913—925.

- Dillon, A. (1992) Reading from paper versus screens: a critical review of the empirical literature. *Ergonomics*: 3rd Special Issue on Cognitive Ergonomics, 35(10), 1297—1326.
- Dillon, A., McKnight, C. and Richardson, J. (1990) Navigation in hypertext: a critical review of the concept. In D. Diaper, D. Gilmore, G. Cockton and B. Shackel (eds.) *Human-Computer Interaction: INTERACT '90*. Amsterdam: Elsevier. 587—592.
- Dillon, A., Richardson, J. and McKnight, C. (1989) The human factors of journal usage and the design of electronic text. *Interacting with Computers*, 1(2), 183—189.
- Downs, R. M. and Stea, D. (1973)(eds.) *Image and Environment: Cognitive Mapping and Spatial Behaviour*. London: Edward Arnold.
- Downs, R. M. and Stea, D. (1977) *Maps in Minds: Reflections on Cognitive Mapping*. New York: Harper and Row.
- Edwards, D. and Hardman, L. (1989) "Lost in hyperspace": cognitive mapping and navigation in a hypertext environment. In R. McAleese (ed.) *Hypertext: Theory into Practice*. Oxford: Intellect. 105—125.
- Elm, W. and Woods, D. (1985) Getting lost: a case study in interface design. *Proceedings of the Human Factors Society 29th Annual Meeting*. Santa Monica, CA: Human Factors Society. 927—931.
- Francescato, D. and Mebane, W. (1973) How citizens view two great cities: Milan and Rome. In R. M. Downs and D. Stea (eds.) *Image and Environment: Cognitive Mapping and Spatial Behaviour*. London: Edward Arnold. 131—147.
- Garnham, A. (1987) *Mental Models as Representations of Text and Discourse*. Chichester: Ellis Horwood.
- Gloor, P.A. (1991) CYBERMAP yet another way of navigating in hyperspace. In *Hypertext '91*, Proceedings of the Third ACM Conference on Hypertext. New York: The Association for Computing Machinery. 107—121.
- Gordon, S., Gustavel, J., Moore, J. and Hankey, J. (1988) The effects of hypertext on reader knowledge representation. *Proceedings of the Human Factors Society 32nd Annual Meeting*. Santa Monica, CA: Human Factors Society. 296—300.
- Gould, P. (1973) On mental maps. In R. M. Downs and D. Stea (eds.) *Image and Environment: Cognitive Mapping and Spatial Behaviour*. London: Edward Arnold. 182—220.

- Hammond, N. and Allinson, L. (1987) The travel metaphor as design principle and training aid for navigating around complex systems. In D. Diaper and R. Winder (eds.) *People and Computers III*. Cambridge: Cambridge University Press. 75—90.
- Hammond, N. and Allinson, L. (1989) Extending hypertext for learning: an investigation of access and guidance tools. In A. Sutcliffe and L. Macaulay (eds.) *People and Computers V*. Cambridge: Cambridge University Press. 293—304.
- Hart, R. A. and Moore, G. T. (1973) The development of spatial cognition: a review. In R. M. Downs and D. Stea (eds.) *Image and Environment: Cognitive Mapping and Spatial Behaviour*. London: Edward Arnold. 246—288.
- Johnson-Laird, P. (1983) *Mental Models*. Cambridge: Cambridge University Press.
- Jones, W. P. and Dumais, S. T. (1986) The spatial metaphor for user interfaces: experimental tests of reference by location versus name. *ACM Transactions on Office Information Systems*, 4(1), 42—63.
- Kelly, G. A. (1955) *The Psychology of Personal Constructs*. New York: Norton.
- Kintsch, W. and Yarborough, J. (1982) The role of rhetorical structure in text comprehension. *Journal of Educational Psychology*, 74, 828—834.
- Lai, P. and Manber, U. (1991) Flying through hypertext. In *Hypertext '91*, Proceedings of the Third ACM Conference on Hypertext. New York: The Association for Computing Machinery. 123—132.
- Landow, G. (1990) Popular fallacies about hypertext. In D. Jonassen and H. Mandl (eds.) *Designing Hypermedia for Learning*. Berlin: Springer-Verlag. 39—60.
- Landow, G. (1991) Position statement for Panel on Structure, Navigation and Hypertext: The Status of the Navigation Problem. In *Hypertext '91*, Proceedings of the Third ACM Conference on Hypertext. New York: The Association for Computing Machinery. 364.
- Lee, E., Whalen, T., McEwen, S. and Latrémouille, S. (1984) Optimising the design of menu pages for information retrieval. *Ergonomics*, 27(10), 1051—1069.
- Lovelace, E. A. and Southall, S. D. (1983) Memory for words in prose and their locations on the page. *Memory and Cognition*, 11(5), 429—434.

- McKnight, C., Dillon, A. and Richardson, J. (1990) A comparison of linear and hypertext formats in information retrieval. In R. McAleese and C. Green (eds.) *Hypertext: State of the Art*. Oxford: Intellect, 10—19.
- McKnight, C., Dillon, A. and Richardson, J. (1991) *Hypertext in Context*. Cambridge: Cambridge University Press.
- Monk, A., Walsh, P. and Dix, A. (1988) A comparison of hypertext, scrolling, and folding as mechanisms for program browsing. In D. Jones and R. Winder (eds.) *People and Computers IV*. Cambridge: Cambridge University Press, 421-436.
- Neisser, U. (1976) *Cognition and Reality*. London: Freeman.
- Osgood, C. E., Suci, G. J. and Tannenbaum, P. H. (1957) *The Measurement of Meaning*. Illinois: University of Illinois Press.
- Rothkopf, E. Z. (1971) Incidental memory for location of information in text. *Journal of Verbal Learning and Verbal Behavior*, 10, 608—613.
- Schank, R. and Abelson, R. (1976) *Scripts, Plans, Goals and Understanding*. Hillsdale, NJ: Lawrence Erlbaum Associates.
- Shum, S. (1990) Real and virtual spaces: mapping from spatial cognition to hypertext. *Hypermedia*, 2(2), 133—158.
- Simpson, A. (1989) Navigation in hypertext: design issues. Paper presented at *OnLine '89 Conference*, London, December.
- Simpson, A. (1990) Towards the design of an electronic journal. Unpublished PhD Thesis, Loughborough University of Technology.
- Simpson, A. and McKnight, C. (1990) Navigation in hypertext: structural cues and mental maps. In R. McAleese and C. Green (eds.) *Hypertext: State of the Art*. Oxford: Intellect. 73—83.
- Snowberry, K., Parkinson, S. and Sisson, N. (1985) Effects of help fields on navigating through hierarchical menu structures. *International Journal of Man-Machine Studies*, 22, 479—491.
- Thorndyke, P. and Hayes-Roth, B. (1982) Differences in spatial knowledge acquired from maps and navigation. *Cognitive Psychology*, 14, 560—589.
- Tolman, E. C. (1948) Cognitive maps in rats and men. *Psychological Review*, 55, 189—208.

Trigg, R. H. and Suchman, L. A. (1989) Collaborative writing in Notecards. In R. McAleese (ed.) *Hypertext: Theory into Practice*. Oxford: Intellect. 45—61.

Utting, K. and Yankelovich, N. (1989) Context and orientation in hypermedia networks. *ACM Transactions on Information Systems*, 7(1), 58—84.

Wetherell, A. (1979) Short-term memory for verbal and graphic route information. *Proceedings of the Human Factors Society 23rd Annual Meeting*. Santa Monica CA: Human Factors Society.

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

Zechmeister, E. and McKillip, J. (1972) Recall of place on a page. *Journal of Educational Psychology*, 63, 446—453.

Zechmeister, E., McKillip, J., Pasko, S. and Bepalec, D. (1975) Visual memory for place on the page. *Journal of General Psychology*, 92, 43—52.

Andrew Dillon is Professor and Dean of GSLIS at the University of Texas-Austin. He can be reached by email at: adillon@gslis.utexas.edu