

User Centred Design of Hypertext/Hypermedia for Education

Cliff McKnight⁽¹⁾, Andrew Dillon⁽²⁾ and John Richardson⁽³⁾

(1) Department of Information and Library Studies, Loughborough University of Technology, Ashby Road, Loughborough, Leics. LE11 3TU, UK. <c.mcknight@lut.ac.uk>

(2) School of Library and Information Science, Indiana University, Bloomington, Indiana 47405, USA. <adillon@indiana.edu>

(3) HUSAT Research Institute, Loughborough University of Technology, Elms Grove, Loughborough, Leics. LE11 1RG, UK. <j.h.richardson@lut.ac.uk>

This item is not the definitive copy. Please use the following citation when referencing this material: McKnight, C., Dillon, A., and Richardson, J. (1996) User Centered Design of Hypertext and Hypermedia for Education. In: D. Jonassen (ed) Handbook of Research on Educational Communications and Technology. New York: Macmillan, 622-633

Abstract

The chapter begins by describing the fundamental concepts of hypertext and gives a brief overview of the different philosophical perspectives manifest in the key figures of the field. It then considers the role of hypertext in learning, concluding from a review of empirical evaluations that many of the claims for hypertext have failed to be substantiated. It is argued that for a variety of conceptual and methodological reasons, it is extremely difficult to evaluate hypertext experimentally in an educational context. However, rather than simply abandon either hypertext or empirical evaluation, the chapter concludes by arguing for an empirically grounded, user centred approach to the design of hypertext based on a knowledge of the users, their tasks, the information space and the context in which the three interact.

What is hypertext, what is hypermedia?

The prefix 'hyper' usually means 'more than' so we may begin by asking what is it that hypertext has that makes it more than text. The simple answer to this is that as well as text, hypertext has 'links'. The text is usually organised into chunks, units, or 'nodes' as they have come to be known and the links form connections between certain nodes.

There are no rules about how big a node should be or what it should contain. Similarly there are no rules governing what gets linked to what. Hence, there can be many different kinds of hypertext in the same way that there are many different kinds of text.

Furthermore, hypertext allows the concept of 'document' to be extended since, logically, entire documents can be treated as nodes and linked together to form a single hypertext.

The concept of links between units of information has a history almost as old as writing itself. Think of a footnote marker in a text: it links the main text with the footnote text, although in this case the marker is a static link — the reader must make the movement between text and footnote. In hypertext, the links are active — selecting the link moves the reader to the linked text in some way. It is this dynamic aspect of presentation that is the principal difference between text and hypertext — an alternative term for hypertext which never gained the same currency was 'interactive documentation' (Brown, 1986) and a recent paper used the term 'responsive text' (Hillinger, 1990).

If hypertext isn't a new idea, why has it recently become so popular? In fact, hypertext has been an idea waiting for technology to catch up with it. In order to implement active links, it is necessary to use a dynamic display medium such as a computer screen. Hence, in the 1960s in the research laboratories, several groups started using large mainframe computers in order to explore the potential of hypertext. Advances in computer technology have led to the development of the minicomputer and the microcomputer, making considerable computer power available to the individual user — the personal computer with so-called 'user friendly' interfaces. On such computers, 'popular' hypertext becomes feasible.

Hypermedia is presented as a further development of hypertext. As computers have moved from being able to present little more than upper case text to being able to present information in a variety of communication media — sound, graphics, video — so it is possible to link these media together using hypertext techniques, hence the term hypermedia. However, in the same way that a book can contain text, drawings, tables, photographs or even pop-up models, so the distinction between hypertext and hypermedia is somewhat arbitrary. We will use either term to refer to a set of nodes of information which are dynamically linked.

By way of example, consider how information relating to the European Union (EU) could be presented. At the top level the reader could be presented with a map of Europe and

then be able to select various aspects of the EU or any particular member state for more information [see Figure 1]. This may take the form of a more detailed map showing the major cities which, when selected, showed some information about the city. Alternatively, the reader may be offered other information about the country, such as the currency, culture, audio-clips of common words in the language such as *please* and *thank you*, and so forth [see Figure 2]. The city maps could offer sites of interest such as museums, local transport details, entertainment sites such as theatres or even restaurants and pubs; they could include video clips of the major tourist sites, on the beach or at the disco. Now the tourist ‘brochures’ can feature live action and be interactive in a way that the present printed chapter cannot. Hence, the extra- or hyper- text nature of the medium.

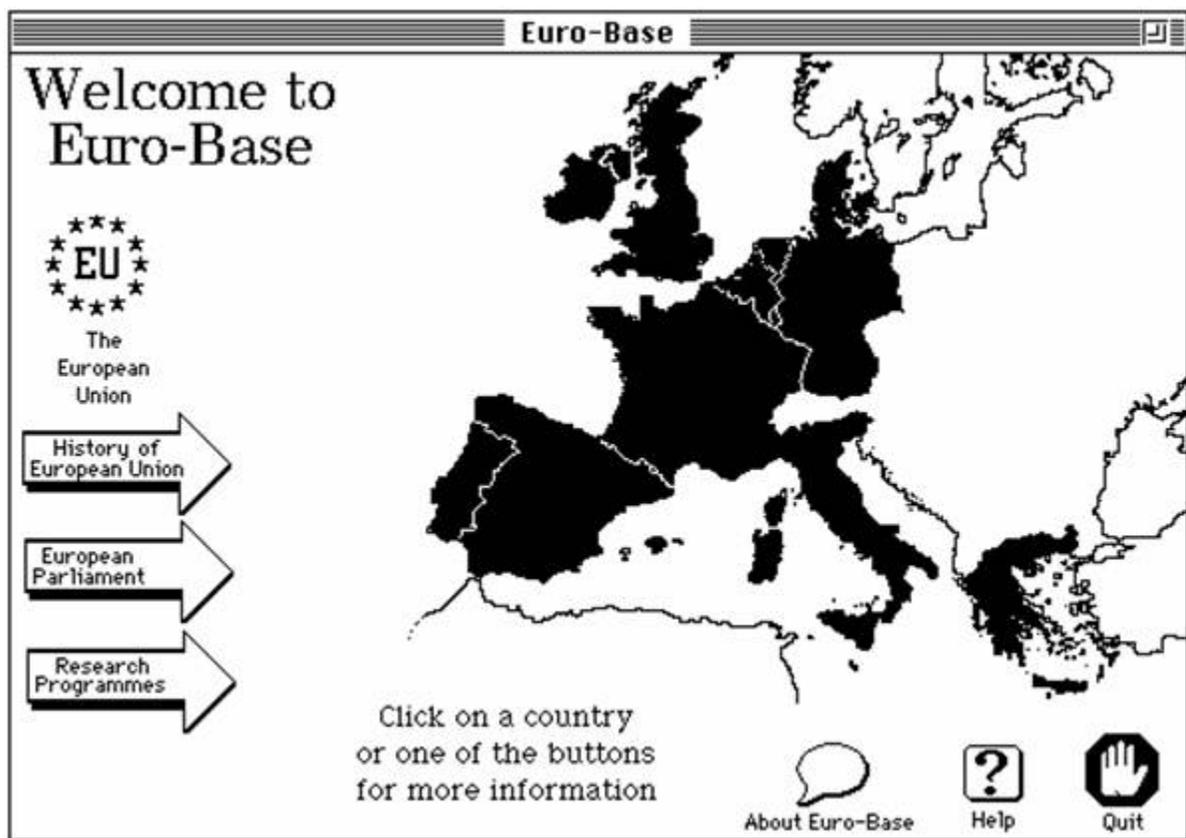


Figure 1: The top level of the Euro-Base hypertext — a wealth of information at the click of a button.

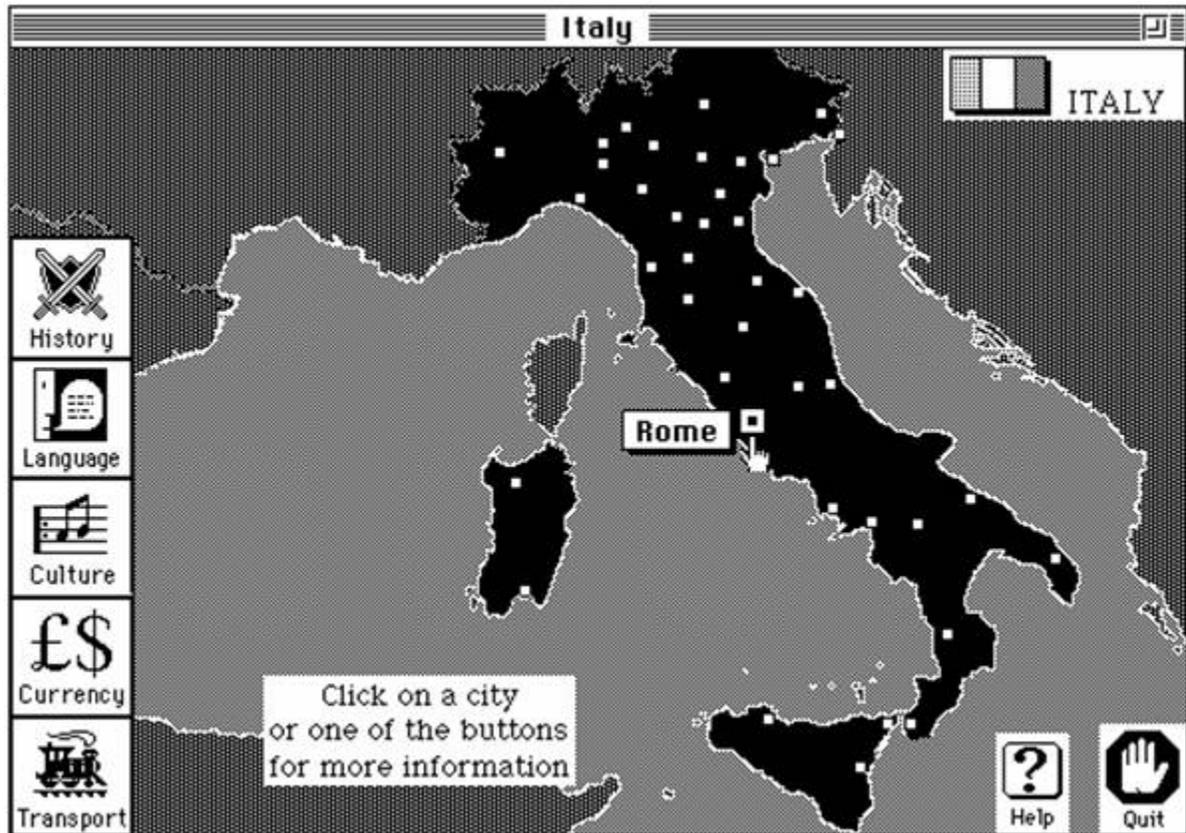


Figure 2: See and hear the sights and sounds of Europe.

The Genesis of Hypertext

We do not intend this section to be a general overview of everything which has gone before. The reader who is interested in such a description would do well to read Jeff Conklin's excellent review (Conklin, 1987) followed by Jakob Nielsen's popular book on the subject (Nielsen, 1995). Rather, we would like to present some of the historical highlights as being representative of three different conceptual stances towards information technology and we represent each of these by an influential figure in the field. These stances are:

- technology designed to work in a manner akin to human cognition — the view of Vannevar Bush;
- technology as an augmenter of the human intellect — the personal philosophy of Doug Engelbart;

- technology as a flexible and usable access mechanism to a complete world of inter-related information — the dream of Ted Nelson.

The article most often cited as the modern birthplace of hypertext is Vannevar Bush's "As we may think" (Bush, 1945). Bush was appointed the first director of the Office of Scientific Research and Development by President Roosevelt in 1941. He saw clearly the problems associated with the ever-increasing volumes of information that were the product of World War II research and technology initiatives:

"There is a growing mountain of research. But there is increased evidence that we are being bogged down today as specialization extends. The investigator is staggered by the findings and conclusions of thousands of other workers — conclusions which he cannot find time to grasp, much less to remember, as they appear. Yet specialization becomes increasingly necessary for progress, and the effort to bridge between disciplines is correspondingly superficial."

It surprises many people to learn that this statement was made 50 years ago — the information explosion is not such a modern phenomenon after all! To cope with this plethora of information, Bush conceived the *memex*, a device

"...in which an individual stores his books, records, and communications, and which is mechanized so that it may be consulted with exceeding speed and flexibility."

More than a simple repository, the memex was based on the principle of association, the notion that all facts, concepts and ideas are linked in the mind and that any one knowledge chunk can act as a stimulus or trigger to remember another.

In devising a supporting technology of this kind, Bush sought to create a means of information storage and retrieval that was intuitive to the user by virtue of its similarity to the workings of the mind. The appeal to intuitive operation predates many of the arguments that have subsequently emerged in the design of information technology. For Bush, it was sufficient that a principle such as association — at that time a popular view of human cognition — would satisfy the organisational requirements of information users, an approach to design that is also partly mirrored in current user-centered design thinking. However, given the volume of information a person might need to handle and what we now know about the vagaries of human association (e.g., Howe, 1980), it

appears unlikely that such a principle alone could act as a sufficient design constraint on a hypermedia system.

Objections aside, it is interesting that Bush viewed the need for system design in such a user centred manner. Unfortunately, the technology of the day was not as developed as the ideas — Bush conceived of the memex as being microfilm based. Hence, his ideas lay dormant for at least two decades, waiting for the technology to catch up. However, it is still commonplace to hear hypertext spoken of as a ‘natural’ information system and even a viable model of the human mind ***CLIFF WE DO NEED A REF to SUCH COMMON SPEECH

Whereas Bush (and others) sought to devise systems based on cognitive equivalence, Doug Engelbart has been developing his conception of hypertext since the early 1960s by placing his emphasis on augmenting or amplifying of human intellect (cf. Engelbart , 1963), as reflected in the naming of his system as *Augment*. Engelbart’s first implementation was NLS (oN Line System) which was meant as an environment to serve the working needs of his Augmented Human Intellect Research Center at Stanford Research Institute. This was a computer-based environment containing all the documents, memos, notes, reports and so forth in addition to supporting planning, debugging and communication. As such, NLS can be seen as one of the earliest attempts to provide a hypertext environment in which computer-supported collaborative work (CSCW) could take place.

Engelbart’s pragmatic position is less concerned with supporting via mimicking the information processing tendencies of the human (a la Bush) than with extending cognition, although to some extent such a position is equivalently determined by an image of the mind which one is trying to enhance or augment. In conceptualising the technology in this way, Engelbart espoused the notion of hypertext as some form of cognitive artefact, extending the capabilities of the human and offering the potential to attain performance levels in information tasks which would be difficult or impossible to achieve without hypertext. This is not dissimilar to McLuhan’s ideas of technologies as extensions of human faculties, for example, wheels as an extension of legs (McLuhan, 1964), ideas which had their own philosophical roots in the ‘personal knowledge’ work of Michael Polanyi (1957).

The actual term *hypertext* is attributed to Ted Nelson who called his dream system Xanadu — a ‘docuverse’ in which the entire literature of the world is linked, a “universal instantaneous hypertext publishing network” (Nelson, 1988). In Xanadu nothing ever needs to be written twice; a document is built up of original, or *native*, bytes and bytes which are *transclusions*, a term which implies the transfer and inclusion of part of one document into another. However, an important aspect of Xanadu is that the transclusion is virtual, with each document containing links to the original document rather than copies of its parts.

While Nelson’s view may appear to be the most ambitious, in many ways his advocacy of this form of system is the most realistic since it rests less on contemporary assumptions drawn from the psychological theories of the day and more on the requirement for a technology to provide rapid, easy access to the world of information. At an abstract level, the World Wide Web developed by Tim Berners-Lee and colleagues at CERN (Berners-Lee *et al.*, 1994) comes close to realising Nelson’s dream since it allows links between documents anywhere in the world (or at least anywhere with a connection to the Internet)⁽¹⁾. Indeed, this may be the route through which some of the ideas behind Xanadu will now be given substance. When Autodesk Inc., the CAD company, purchased Xanadu in 1988 there was hope that an actual product would be developed. However, seven years and reputedly \$5 million later, Autodesk has abandoned the project and the Xanadu name is back with Nelson who, at the Hypertext ’93 conference, was talking about ‘Xanadu Light’ based on existing tools such as Gopher and Telnet.

The three views represented here are not mutually exclusive; it is possible to advocate a hypertext system which provides ready access to all information and therefore allows users to perform new tasks. Indeed, there is a fine line between these idealised positions and it is not possible, or indeed very useful, to describe any particular contemporary system (or system designer’s viewpoint) in terms of one or any of them. However, the fact that different views can proliferate illustrates the point that *hypertext* is not a unitary concept, not a single thing which can be defined any more precisely than in terms of nodes and links. It is for this reason that hypertext software packages with completely different ‘look and feel’ (cf. Guide and HyperCard) can be produced and still claim to embody the concept of hypertext.

The corollary to this is that we should not expect any particular hypertext system to be ideal in all task situations. It is not surprising that Conklin’s (op. cit.) historical survey of

hypertext systems groups them in a largely task-based way. He uses four categories: macro literary systems, which center on the integration and ready accessibility of large volumes of information; problem exploration systems, which are designed to allow the interactive manipulation of information; systems for structured reading/browsing/reference; and finally, systems which might have been applied to a specific application but whose real purpose in construction had been the experimental investigation of hypertext technology itself.

As we shall attempt to demonstrate here, the pedagogic implications of hypertext rest on a mixture of these three underlying philosophies. However, a major weakness of the field has been the inadequate theoretical explication of the nature of learning which hypertext systems might support. In particular, an over-simplistic acceptance of the naturalistic argument advocated by Bush remains, even if the models of the mind subscribed to by contemporary learning theorists (e.g., Hammond 1989) are more sophisticated than the principle of association current in Bush's time. Consequently, the field lacks a dominant theoretical perspective that can draw on significant empirical support and much of the literature in the field is rhetorical in style, reflecting the pre-paradigmatic nature of the field.

Hypertext and Learning

For over 40 years, technological 'solutions' have been offered to the teaching profession in order to improve their effectiveness, ranging from 'programmed text' and 'teaching machines' through to the modern fascination with computers. Areas such as Computer Based Learning (CBL) or Computer Aided Instruction (CAI) aim to provide some of the functions of the teacher. The technologist's dream, of course, was the provision of a workstation for every learner so that they may proceed at their own pace and, to some extent, the dream remains today. However, between the earliest teaching machines and the latest hypermedia environments, there has been a radical shift in prevailing pedagogy, from the repetitive reinforcement schedules of the behaviourists through the cognitivist movement and latterly the constructivists. Each movement has sought to make the technology their own and to make a case for the use of hypertext in their own terms.

It is not our intention to expound these arguments here. Suffice it to say that few theorists of any philosophical persuasion within the learning discipline have been able to specify a distinctive learning environment within the technology that uniquely supports their

position. Indeed, critics of the old behaviorist teaching machines characterised that technology as betraying a crude, mechanistic model of learning where teaching is reduced simply to optimising the presentation of material — a criticism that is not lessened in any way by recent improvements in technology *per se*.

Within education, hypertext has been seen by some as a valuable new constructivist tool for supporting teachers and students (Cunningham, Duffy and Knuth, 1993) whereas others have seen it as simply a fancy new ‘jug’ (Whalley, 1993). There is merit in both positions, but for us the crucial point is that the argument is currently showing no sign of resolution since neither side’s claims have been subjected to sufficient empirical testing. Instead, educational technologists have embarked on seemingly interminable theorising of technological empowerment without commensurate empirical correction.

The perceived advantages of hypertext as an educational medium are usually ascribed to its non-linear property. This is often contrasted with the assumed linearity of traditional text, for example:

“In contrast [to hypertext] most standard text documents are constructed to be read *linearly*, from beginning page to ending page.” (van den Berg and Watt, 1991, p.118)

We have discussed the myth of linear text at length elsewhere (McKnight, Dillon and Richardson, 1991, Dillon 1994). Suffice it to say that most texts are *not* constructed to be read from beginning to end. This becomes blatantly obvious after the briefest of observations of people using “standard” text documents — e.g., journal articles, course text-books, encyclopædiae and newspapers. How are you reading the book you are holding? Did you reach this page after reading all the preceding pages, or did you browse the contents, flick through and glance at some chapters before deciding which to read now? Even within this chapter you may have glanced at the references first (have we referred to your work?), skipped sections (maybe you are not interested in the historical roots of hypertext) and so forth. Apart from the novel, very few texts are constructed to be read linearly in their entirety. As we have pointed out elsewhere, even the novel may be used in many non-linear ways in, say, a literature class in tertiary education (Dillon and McKnight, 1990).

Hypertext has certainly become a popular term to be bandied about in education, not surprisingly since it is relatively new, technically impressive and, until the novelty wears off, often fun to use. A brief search of the literature from 1989 to 1993 using the search

terms 'hypertext' OR 'hypermedia' AND 'learning' will yield over 100 references. The issues covered and the disciplines and journals represented are many: using multimedia with Navajo children to alleviate problems of cultural learning style, in *Reading and Writing Quarterly*; monitoring hypertext users, in *Interacting With Computers*; hypertext in cognate-language learning, in the *Journal of Computer Assisted Learning*; opportunities for hypertext in interactive learning, in the *Journal of the American Society for Information Science*.⁽²⁾

Unfortunately for present purposes, very few of these report on the systematic evaluation of hypertext in an educational setting — by systematic evaluation we mean the empirical testing of real users interacting with the artefact in an ecologically valid task. A recent extension of this work by Gabbard and Dillon (1995) shows that the situation has not changed since then. The vast majority of the literature published on this topic either says how wonderful hypertext is or might be, or describes some hypertext-based courseware in detail without evaluating its effectiveness. Some even equate Apple's HyperCard with hypertext, assuming that if it is built in the former it must be an example of the latter (e.g., Horton, Boone and Lovitt, 1990). Although these authors do perform an evaluation, their results provide evidence for computer aided instruction rather than hypertext.

One of the earliest attempts to evaluate hypertext's potential in a learning environment was that reported by Beeman *et al.* (1987). Their paper reports on attempts to use the Intermedia hypertext system in two of Brown University's existing courses — an English literature course and a plant cell biology course. The evaluation of the effect of introducing hypertext was far from easy. Each of the courses was closely observed by a team of social scientists, once prior to the introduction of the hypertext and once when the hypertext materials were in use; instructors and students were interviewed several times throughout the evaluation; a group of students was asked to keep diaries of their activities during the time the courses were taught; and the use of a specially set up computer laboratory by both students and instructors was monitored.

At first sight the effects of introducing hypertext seem to have been positive. Beeman *et al.* report a small positive correlation ($r = 0.29$ at 0.05 level of confidence) between high Intermedia use and high grades. However, they also report an unexpected finding which suggests that improvements may not have been attributable to the introduction of hypertext *per se*, but rather to factors related to its introduction. Because the Intermedia workstations were not ready in time, the professor in charge of the English course was

forced to teach the course without using the system but having already prepared the Intermedia material. The result of this was that he changed the way he taught the course and subsequently felt that students grasped 'pluralistic' reasoning styles better than in previous years. Furthermore, the students were more satisfied with the course than in previous years. This suggests that the need to rethink the course design may have been the major contributor to the improvement in grades. A professor who has taught the same course for years may not be as 'inspiring' as he used to be, but interest may well be rekindled and communicated to the students by having to redesign the course for a new medium.

A further difficulty in making any strong statements about the apparent improvements in grades is also raised by Beeman. By themselves, the studies offer no evidence that the style of thinking fostered in these two courses transfers to other courses. As Beeman points out, students are generally good at adapting to teachers because they are interested in doing well. Hence their results may indicate a course-specific adaptation rather than a genuine change in thinking style. However, it would be extremely difficult to test such an hypothesis since it would involve the comparative evaluation of students across many courses.

The Beeman studies are an excellent illustration of the difficulties involved in assessing the effect of introducing not only hypertext but any new teaching technique or technology into an educational context. The Hawthorne effect, so-called because of the industrial context in which it was first systematically observed, is just as likely to appear in an educational setting. Furthermore, Beeman *et al.*'s most interesting conclusion was that the significant learning effects observed through the use of Intermedia were more pronounced for the people involved in producing the materials rather than students using the system, apparently substantiating the adage that 'the best way to learn something is to teach it'!

Despite the topic of learning having formed a major section of psychology for many years, we still know very little about the cognitive processes of learning. As psychologists we recognise this as a parlous state of affairs since many people recognise the ability to learn as a key attribute of human behaviour. From the outset, psychology has sought to define, explain and predict learning, from associative principles to laws of effect, from principles of reinforcement to cognitive skill acquisition, from mental model acquisition to constructivism. However, as Norman (1980) states boldly, "it has all come

to nought.” Learning, once the backbone of psychology, is now rarely found as a major component of psychology degrees but resides in specialised sub-disciplines such as education. While this might be interpreted as a position of strength engendered by the subject’s inherent importance, such a position is hard to justify in the absence of similar developments for other major psychological issues such as thinking, memory, perception, social interaction and so forth. In effect, 100 years of studying learning has provided little by way of systematic knowledge for ensuring desirable learning outcomes. Hence, evaluating the interactive technology to support this process is by no means straightforward.

Hammond (1989) claims that we seem to be good at providing some appropriate environments for learning:

“... in *teaching* a child to talk the parent merely needs to give appropriate stimulation at appropriate times; details of intermediate states of knowledge and the processes of acquisition can safely be left to the child and to the research psychologist.” (p.172)

This comment is telling in that it suggests that ‘merely’ providing the appropriate stimulation at the right time might have the desired results (a rather Skinnerian perspective for such a cognitivist to take, though with obvious Chomskian or Pinkerian resonance!) with the learners themselves doing the rest. If this really is the best we can do, then psychology’s role is less to worry about the nature of learning than to concentrate on the provision of suitable environments...and hypertext may be one such environment.

Hammond and Allinson (1989) suggest that hypertext can provide the basis for an exploratory learning system but that by itself it is insufficient, needing to be supplemented by more directed guidance and access mechanisms. In order to investigate this suggestion, they conducted an experiment in which all subjects used the same material held in a hypertext form, but with differing guidance and access facilities available. The baseline group had a basic hypertext with no additional facilities, while other groups had either a map or index or guided tours available, and a final group had all three facilities (map, index, tours) available. Half of the subjects were given a series of questions to answer while accessing the material (a directed task) while the other half were instructed to make use of the material to prepare for a subsequent multiple-choice test (an exploratory task).

Perhaps surprisingly, Hammond and Allinson report no reliable differences between task conditions for the three groups which had a single additional facility, although in all three groups the facilities were used to a substantial extent. However, in the group having all three facilities available there was a significant task-by-facility interaction. Those subjects performing the exploratory task made little use of the index but significant use of the tours, while those performing the directed task made little use of the tours and far more use of the index. Thus, Hammond and Allinson argue that after only 20 minutes subjects were able to employ the facilities in a task-directed manner.

The additional facilities also allowed more accurate overviews of the available material and resulted in a higher rate of exposure to new rather than repeated information. However, there were no significant differences in task performance between groups. Hammond and Allinson attribute this lack of difference to the fact that neither of the tasks required any strategic organisation of the material and they therefore caution against extrapolating such results to situations other than simple rote learning of relatively unstructured material. Indeed, even the subjective judgements of the subjects — that the system was easy to use, getting lost was not a major problem, and the system was rated as “better than a book” — should be viewed in the light of the fact that the hypertext used was very small (consisting of only 39 information screens) and subjects only used the system for a maximum of 20 minutes. Although it discusses ‘learning support environments’ and is clearly aimed at an educational context, Hammond and Allinson’s work provides a contrast to the Beeman study in that its strength is its controlled, experimental nature, whereas the strength of Beeman’s work is its applied, ‘real world’ nature. Both types of study have a role to play in the attempt to discover the effects of hypertext in education.

Stanton and Stammers (1990) suggest that the reasons why a non-linear environment might be superior are that it (a) allows for different levels of prior knowledge, (b) encourages exploration, (c) enables subjects to see a sub-task as part of the whole task, and (d) allows subjects to adapt material to their own learning style. In their experiments (1989, 1990), one set of subjects was given the freedom to access a set of training modules in any order, while another set of subjects was presented with the modules in a fixed order. They reported that performance was significantly improved when subjects trained in the non-linear condition. Although such comparisons may provide valid experimental designs, extrapolating the results to realistic learning situations is difficult, particularly in higher education where students are rarely forced to access material in a

rigid, predetermined order. Hence, the results may reflect the advantage not so much of non-linear environments but rather of giving the learner some degree of control over the learning environment — a return to the more straightforward notion of providing accessible material and letting the learners ‘get on with it’ themselves.

It could be argued that a hypertext environment does provide for greater learner control and therefore possesses advantages over traditional paper-based learning materials. However, this admits of two equally plausible interpretations: greater control over the user’s access to the hypertext’s contents by way of the links provided by the author/designer, or greater control by the user because they are free to follow the pathways of their choice — an option that is allegedly more difficult with printed text. The second option seems more optimistic and attractive but experience does not give much encouragement. Given the option of following their own path through a hypertext courseware or taking a path suggested by the tutor, how many are likely to follow their own inclinations? Furthermore, how many tutors would put as much energy into generating paths through the hypertext which support a stance antithetical to their own? Students using hypertext courseware will tend to follow the paths provided by the course tutor or hypertext author. If either of these possibilities is true then a hypertext courseware may prove more constraining than the books it replaces, which can be opened at any page.

In principle, the whole of a book or journal volume is available to the reader simply by turning page after page, whereas in hypertext the learner is at the mercy of the author, reliant on him having provided suitable links. Even if learners are given the facility to add their own links, they must have seen the nodes at both ends of the link in order to make the judgement that a link is desirable. This makes the process of adding links seem a little more ‘hit-and-miss’ than it is usually described.

Although the notion of control is an important one in education, it is far from clear that hypertext provides the learner with more control than traditional media. While Duchastel (1988) states that computers promote interaction through a manipulative style of learning where the student reacts to the information presented, the fact that the learner is using a mouse to select items and move through the information space does not make the process any more ‘active’ than consulting an index, turning the pages of a book, underlining passages and writing notes in the margin. In the main part, interactivity in education comes, as it has done for over 2,000 years, from verbal discourse.

If current hypertext systems appear to provide greater opportunities for learner control and best support exploratory styles of learning, this may in part explain the excessive zeal among their proponents for the importance of these aspects of learning. Romiszowski (1990) provides a welcome degree of caution:

“There is, however, some doubt as to whether all these process oriented aspects of hypertext systems are necessarily ‘a good thing’ in all manner of learning situations. The research on learner control of the learning process is, to say the least, mixed. There is much evidence to suggest that learners, when free to select their own strategies, do not always select wisely.” (p. 322)

This echoes earlier suggestions that the majority of students are not able to set learning objectives for themselves and study autonomously (Bunderson, 1974; O’Shea and Self, 1983). Many undergraduate programmes start with study skills courses.

A heavy emphasis on exploratory learning for younger students may soon, if not already, be seen as yet another progressive enthusiasm that flowered in the 1960s. Today it is more likely to be seen as just one useful tool in the teacher’s armoury — perhaps a stimulus and motivator to be used at regular intervals between sessions of more structured classwork.

At the other end of the student age spectrum, access to a wide ranging and richly interconnected hypermedia database may be of real interest to, say, the fledgling humanities PhD candidate who is trying to identify common factors which have influenced a diverse range of human activities. The term fledgling is used intentionally since who else will incorporate primary documents which are accessed so very rarely (unless thesis requirements are changed to dictate that all referenced material is incorporated into the hypertext ‘canon’!) Undergraduates will probably still feel so overwhelmed by the extent of their workload that following associative flights of fancy, rather than sticking to the prescribed texts of the reading list, will seem at best a luxury if not sheer folly. As Whalley (1990) points out:

“The hypertext reader might flit between the trees with greater ease and yet still not perceive the shape of the wood any better than before.”

The scientific principle of parsimony, of adopting the most simple explanation, is worth applying to learning and the new media. ‘Pluralistic reasoning’ to some is ‘confused

thinking' to others, and given the far from certain results of much cognitive psychological work on mental representations and knowledge construction, as well as the on-going revisionism in the educational field, we might best consider hypertext as an information accessing medium and the learner as a seeker of information before positing elaborate notions of thinking style which prove difficult to validate empirically.

Adopting such a perspective allows us to shift our concern from theorising about the mental activity of learning to designing an information environment which can support task performance. We will return to this point later, but before outlining an approach that could help improve the design of hypertexts we will review some of the studies which have actually attempted, with varying degrees of success, to evaluate hypertext systems in an educational context.

In a study by Gordon and Lewis (1992), 80 subjects read a tutorial about a state-of-the-art hi-fi video cassette recorder in either a linear format or one of two hypertext formats (unconstrained network or constrained structure). They were then asked to summarise the tutorial, answer a series of questions and solve two problems using the hi-fi VCR to perform particular tasks. For the factual questions, subjects using the linear form of the information scored significantly higher than either of the hypertext groups which did not differ from each other. For the problem solving tasks, subjects using the constrained hypertext performed equivalently to the linear subject, with the free network subjects performing significantly worse. The authors conclude that: 'if it is critical that students learn the details of material in a document, the instructor cannot rely on the student to hypertext [sic] through the information and acquire it; linear formats should probably be retained.' (p. 307) It is a pity these authors did not include a condition in which subjects used the traditional paper manual.

A study by Higgins and Boone (1990) compared the effectiveness of a hypertext study guide used either in combination with or instead of lectures. Forty ninth-grade students (mean age 14.6 years) enrolled in a course in Washington State History took part. Of the 40, there were 10 with learning disabilities, 15 remedial students and 15 regular students. They were randomly allocated to the three conditions of lecture alone, lecture plus study guide, and study guide alone. Their results led them to conclude that 'The hypertext computer study guide is as effective an instructional medium for students with learning disabilities, remedial students, and regular education students as a well-prepared lecture, as measured by recall and retention of information both factual and inferential.' (p. 539)

They drew identical conclusions for the combination of study guide and lecture and suggested that the study guide could provide some students with the practice necessary to increase quiz performance. Strangely, the authors appear not to have considered subjective preference. If there was not much to choose between them in terms of outcome, then the more attractive system has a major advantage in terms of student motivation, especially with less able students.

Despite our criticism above of van den Berg and Watt's (1991) view of text linearity, the study reported in their paper is worthy of consideration. In a design similar to that of Higgins and Boone (1990) they compared the effectiveness of a hypertext document containing material on introductory statistics and hypothesis testing. In a 'competitive' condition a random sample of 28 students used the hypertext during a six-week period when they did not attend lectures, while the remaining 81 students continued to attend lectures. In a 'supplementary' condition, 30 randomly selected students were given the hypertext to use at their discretion while the remaining 72 were not allowed access to it. In the 'replacement' condition, the hypertext served as the sole source of instruction and there were no accompanying lectures. These conditions were run over three consecutive semesters and subjects were senior level Communications Sciences majors.

Interestingly, although there were no significant differences in the objective performances of the groups, the subjective acceptance of hypertext was highest in the supplementary condition and lowest in the competitive condition. The authors suggest that subjects in the competitive condition may have been influenced by the contrast between being left to their own devices compared with the guidance they perceived the control group receiving. They conclude that the most suitable use for hypertext might be as a replacement for in-person instruction where teaching is not available.

Jonassen, one of the most prolific of authors in the field of hypertext and learning, recently reported on a series of three studies attempting to evaluate hypertext (Jonassen, 1993). In terms of structural knowledge acquisition, Jonassen was forced to conclude that his results called into question 'the ability of learners to engage in meaningful learning rather than information retrieval from hypertext, especially in the context of a learning environment' (p 165). Far from being the 'natural' learning environment which somehow reflects semantic memory, Jonassen suggests that 'a fair evaluation of learning from hypertext can only come from hypertext-literate learners who have developed a useful set of strategies for navigating and integrating information from hypertext' (p. 165).

Marchionini and Shneiderman (1988) have suggested that hypertext is more suited to browsing than directed retrieval tasks. Following from this suggestion, Jones (1989) hypothesised that more incidental learning would occur in a browsing task than in a task requiring the use of an index. The argument advanced by Jones was that the links in a hypertext node represent an embedded menu and that the context provided by the node should encourage the connection of the ideas at either end of the link. In other words, the learner's semantic net is more likely to be elaborated or more learning is likely to occur.

Two groups of subjects were used in Jones's experiment. Both groups used the same hypertext database, but one group was shown how to browse through the information using the links and were explicitly instructed not to use the index, while the other group were instructed in the use of the index and were not informed about the active nature of the highlighted words on screen (which were described to them as 'clues to other index entries'). Subjects were given five questions to answer from the database, but afterwards were given 10 questions to measure incidental learning.

Although Jones's argument has intuitive appeal, her experiment failed to support her hypothesis. No significant differences were observed in terms of performance on the incidental learning questions. It is possible that the nature of the questions given to the subjects to answer from the database did not encourage incidental learning. This is certainly suggested by the low overall success level of subjects on the incidental learning test — the *highest* mean number correct for any group was 1.56. Even in the five target questions, taking all groups together, no question was answered correctly by more than half of the subjects. This suggests that the task was not particularly sensitive to the effect of the experimental manipulations and hence we can do little more than agree with Jones that "much more research is needed."

In total these studies illustrate the problems that befall most evaluations of hypertext in education: difficulties in controlled experimentation, difficulties in finding ecologically valid tasks, difficulties in describing process and difficulties of defining — let alone measuring — the outcomes of learning. Marchionini (1990) sums it up:

"The essential problem of evaluating highly interactive systems is in measuring both the quality of the interaction as well as the product of learning. Evaluations of hypermedia-based learning must address the process of learning and the outcomes of learning."

(p.20.6)

If there has been inadequate evaluation of the hypertext systems that have been implemented, does this prevent the design of a new generation of improved hypertext systems? We answer this question both 'yes' and 'no'! There can be no possibility of progress without reliable and convincing evaluation but there are certain established procedures to improve the design of the initial system. These are the views underlying user centred system design and the view of hypertexts and their users as one form of socio-technical system.

User Centred Design

User-centred design has emerged as the prevalent technological design philosophy in recent years in response to an increased awareness that interactive computer based systems often failed to achieve the goals of their designers, especially in relation to user requirements and consequently user satisfaction. The conceptual theory has been recast in more operational terms as enhancing a system or product's 'usability' with measurable objectives such as speed of learning, reduction in error, task efficiency, effectiveness in and reduction in need of support requirements.

To understand the usability issues underlying hypertext it is helpful to conceptualise the system in terms of four basic factors: the users, their tasks, the information space in which the task is being performed and the environment or context in which all these interact.

The users

Users vary tremendously in terms of the skills, habits, intentions and numerous other attributes that they bring to the computer when interacting with hypertext. Information technology's rapid development over the last decade and its penetration into almost every sphere of human activity has raised new challenges in terms of design assumptions. It can no longer be assumed that users will be trained computer scientists or IT professionals who have experience of a wide range of systems or applications. User interfaces, particularly in educational and public access applications, must be designed with the assumption that some users may well bring no previous computing knowledge with them to the interface. Since contemporary thinking rightly stresses that technology should be designed with the users' needs in mind, an essential first stage of user centred design methods is the analysis of users' skills and requirements.

The tasks

The tasks that can be performed with documents are also extremely variable. People read texts for pleasure, to solve problems, to stimulate ideas or to monitor developments for example. Since such interactions also vary widely in terms of time, effort and skill involved, then when we consider the development of a new information presentation medium we must also determine the nature of the tasks it is intended to support if we are to avoid catch-all phrases or claims such as “hypermedia is better than paper”.

The information space

These first two terms are relatively self-explanatory but ‘information space’ is a more vague term by which is meant the document, database, texts and so forth on which the user works. An information space could therefore be as small as a single text fragment or as large as an online database, national archive or university library. It is the presence of a boundary rather than the size or number of items which defines the concept. The information spaces currently in existence are numerous (just think of all the types of documents that are available in the paper domain) but these are likely to be overrun with new forms of information space in the hypermedia world when graphics, sound, video and immersive graphics eventually appear. Information type can be shown to interact with both users and tasks i.e., people utilise information differently depending on its type, their experience and what they want from it (Dillon *et al.*, 1989).

The centrality of context

Of course all three basic factors come together to provide some context — the scenario in which certain users interact with particular information to perform specific tasks. In embracing user-centred principles for design it is essential that the contextual variables are clearly specified. A doctoral student and a solicitor may both be searching an online database for references to relevant material but one may be looking for *any* relevant item while the other may be looking for a specific supporting case. To assume that this is the same task, or that these users are equivalent and could be supported by the same system would be simplistic and lead to the sort of confused reasoning that posits ‘hypertext is better (or worse!) than paper’. Such statements betray a naïve acceptance of commonality of purpose and ability on the part of all potential users and is almost certain to lead to lack of usability in systems design.

Usability and hypertext

Jonassen (1990) recognises the central connection between hypertext design and user requirements:

“The most significant problem in creating hypermedia is deciding how to structure the information. The answer to this question depends, in part, upon how the hypermedia will be used.” (p.12)

Wright (1990) also supports a strongly user centred approach to hypertext design for educational applications. In response to the rhetorical question as to ‘what really matters?’ in hypertext design she gives an unequivocal answer:

“The answer will nearly always depend upon the task that the learner is engaged in. This task will determine the functionality required by the learner...” (p.171)

For some learning tasks, minimal requirements might seem obvious: CD quality sound for poetry and language students; high resolution screens for students of the visual arts and large format screens for designers. Yet even these would hardly ensure learning and at best could only be considered the starting point for a user centred design. However, for many learning contexts the functional requirements are not even that specifiable, particularly when interface and information structures are at issue. The only solution is a fine grained task analysis that can not only determine the functionality but the particular instantiation that is most appropriate for the range of users. As Wright points out, there are often many ways in which a function can be provided — just think of the number of ways in which a database can be searched.

The microcomputer has been justifiably celebrated as a real general purpose tool. Ironically, the hypertext designer may have to constrain the computer in order to empower the user. While the computer can be programmed to undertake highly complex transformations on abstract data structures, the majority of users are more inclined to think in terms of simple transformations on more ‘concrete’ or familiar data structures.

It may seem that we are arguing for a different hypertext instantiation for every learning context based on a new and detailed task analysis. Without a backward look at existing user models of data structures this would probably be a recipe for chaos. Evidence for

this suggestion comes directly from the evolution of our most successful information technology — the printed book.

We have had nearly 500 years experience of using printed text books and they not only support a wide range of applications but users also have such a strong mental model of their generic structure and organisation that they can successfully adopt an equally wide range of usage strategies. What we understand as the book's standard structure, both physical and organisational, evolved over time and readers' models also developed and accommodated these changes. However, the enduring success of the book as an artefact is largely due to a faithful adherence to common user expectations — change took place but only very slowly.

While it is clear that hypertext can support activities impossible or very difficult to perform with paper, we must be sure that we introduce such designs in order to improve our support of, or to enable, valid learning tasks. It is not sufficient that we can browse a million pages on our desktop, or link 100 articles together for rapid retrieval at the click of a mouse button, — such capabilities are only important in terms of their utility to human learners. Yet there are few signs that most learning scenarios require such support, and little knowledge on how we might best provide it in terms of usability, even if it were required.

Beyond media differences — information structures and knowledge representation

It is clear that readers form mental representations of a paper document's structure in terms of spatial location (Lovelace and Southall, 1983) and overall organisation. Dillon (1991) for example has shown that experienced academic journal readers can predict the location within an article of isolated selections of material from it with a high degree of accuracy, even when they cannot read the target material in detail. Such representations or models are derived from years of exposure to the information type or can be formed in the case of spatial recall from a quick scan of the material. Such models or superstructural representations (van Dijk and Kintsch, 1983) are useful in terms of predicting where information will be found or even what type of information is available. Consideration of existing models is vital in the design of new versions so as to avoid designing against the intended users' views of how the information should be organised (see e.g. Dillon, 1994).

The qualitative differences between the models readers possess for paper and electronic documents can easily be appreciated by considering what you can tell about either at first

glance. A paper text is extremely informative. 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.

Hypertext may appear to be a completely new presentation format and therefore free to establish new user models based on the radically new technology. Unfortunately for the creative system designer, the current generation of potential users approach the new technology with expectations that are grounded in print. This is not surprising when the primary content of most hypermedia systems is, and may well continue to be, text. It may be electronic rather than printed, but that can be a minor difference to the user.

Some researchers have accepted this legacy and constructed hypertext systems which rest heavily on not only the conceptual structure but the physical format as well. Benest's (1990) hypertext system employs a realistic on-screen representation of an open book with pages that can be 'turned'. The system was not designed, or developed, with the benefit of empirical studies but is intuitively appealing and an impressive achievement in its time.

In contrast, the SuperBook developed at Bellcore and described comprehensively by Landauer *et al.* (1993) has benefited from exhaustive user studies. SuperBook incorporates a number of familiar print conventions but does not attempt to re-invent the printed book on screen. Instead, development effort has been directed to enhancing the application's intelligent features in a manner that users can employ effectively. In comparative trials SuperBook has proved superior to printed texts in terms of speed and accuracy for search type tasks — but only after repeated tests and redesigns on the basis of those tests.

What is important to note is that experienced readers have acquired expectations of how information spaces are organised and we need to be aware of this in our designs. It is precisely because such models of information are often ignored that we read so much of the navigation problem for users of hypertext (see Dillon, McKnight and Richardson, 1993).

Conclusions

If there is little reliable evidence to support the claims that hypertext systems can really support alternate and superior modes of learning (Landauer (1995) suggests only nine scientifically satisfactory studies have compared paper with hypertext), and we have few effective means of measuring the process of learning anyway, what does that leave us? Three compatible options are available: a reduction in expectations, a switch from process to outcome, and a concentration on an evolutionary approach to development, building on user models rather than trying to make a step-function change — Mao's thousand mile march started with the first step!

A reduction in expectations is probably now due, given the amount of hype which has accompanied the popularisation of hypermedia. As any marketing executive will tell you, a certain amount of hype is necessary in order to stand out against the background noise. However, now that hypermedia has gained a certain amount of acceptance we can afford to be a little more realistic in our expectations.

A switch from process to outcome might seem reasonable, given our arguments to the effect that psychology does not have a good understanding of the process. Unfortunately, it is no simple matter to measure 'learning' as an outcome. Notwithstanding the fact that the education system lumbers on, those involved in the system continue to search for ways which truly reflect meaningful changes in learners as a result of their experiences with the system. It is somewhat ironic that the current emphasis within the British tertiary education system is on measuring the quality of teaching, yet we do not have a philosophically defensible measure of learning against which to judge the teaching function.

Our own preference, as the foregoing has hopefully argued convincingly in favour of, is for an evolutionary approach to system development. However, we take a proactive view of evolution rather than see it as a simple environmental winnowing. We believe that, rather than generating many systems and relying on the 'survival of the fittest', it is possible to design 'the fittest' for any context by utilising user centred, task based design grounded in an empirical methodology. While we have certain sympathies with the humanistic emphases of the constructivist movement, we find the general scientific approach more fruitful. When the constructivists have achieved as much as the scientists,

we may be persuaded otherwise. After all, the essence of science is a mind open to data. For now, the debate continues.

Footnotes

(1) At the detailed level there are still numerous important differences between Xanadu and the World Wide Web, the former including consideration of, for example, copyright payment, storage and usage charging, addressing to the character level, typed links and of course transclusions.

(2) Note that we are not concerned here with information retrieval from hypertext, a topic large enough to merit a separate chapter. The reader interested in exploring this topic is referred to Smeaton (1991) as a convenient starting point.

Acknowledgements

When we started writing this chapter, we all shared an office in Loughborough University's HUSAT Research Institute — happy days! During its writing, Andrew moved to Indiana and Cliff moved to another department still within Loughborough University but some miles away from John. We are grateful to David Jonassen for his patience while we came to grips with the realities of computer supported collaborative working. We are also grateful to John Leggett for his comments on an earlier draft of this chapter.

References

Beeman, W.O., Anderson, K.T., Bader, G., Larkin, J., McClard, A.P., McQuilian, P. and Shields, M. (1987) Hypertext and pluralism: from lineal to non-lineal thinking. *Proceedings of Hypertext '87*. University of North Carolina, Chapel Hill. pp. 67–88.

Benest, I.D. (1990) A hypertext system with controlled hype. In R. McAleese and C. Green (Eds.) *Hypertext: State of the Art*. Oxford: Intellect. pp. 52–63.

van den Berg, S. and Watt, J.H. (1991) Effects of educational setting on student responses to structured hypertext. *Journal of Computer-Based Instruction*, 18(4), 118–124.

Berners-Lee, T., Cailliau, R., Luotonen, A., Nielsen, H.F. and Secret, A. (1994) The World-Wide Web. *Communications of the ACM*, 37(8), 76–82.

Brown, P.J. (1986) Interactive documentation. *Software — Practice and Experience*, 16(3), 291–299.

Bunderson, C.V. (1974) The design and production of learner-controlled software for the TICCIT system: a progress report. *International Journal of Man-Machine Studies*, 6, 479–492.

Bush, V. (1945) As we may think. *Atlantic Monthly*, 176/1, July, 101–108.

Conklin, J. (1987) Hypertext: an introduction and survey. *IEEE Computer*, September, 17–41.

Cunningham, D.J., Duffy, T.M. and Knuth, R.A. (1993) The Textbook of the Future. In C. McKnight, A. Dillon and J. Richardson (Eds.) *Hypertext: a Psychological Perspective*. Chichester: Ellis Horwood. pp. 19–49.

van Dijk, T.A. and Kintsch, W. (1983) *Strategies of Discourse Comprehension*. New York: Academic Press.

Dillon, A. (1991) Readers' models of text structures: the case of academic articles. *International Journal of Man-Machine Studies*, 35, 913–925.

Dillon, A. (1994) *Designing Usable Electronic Text: Ergonomic Aspects of Human Information Usage*. Bristol PA: Taylor and Francis.

Dillon, A. and McKnight, C. (1990) Towards a classification of text types: a repertory grid approach. *International Journal of Man-Machine Studies*, 33, 623–636.

Dillon, A., McKnight, C. and Richardson, J. (1993) Space — the final chapter *or* why physical representations are not semantic intentions. In C. McKnight, A. Dillon and J. Richardson (Eds.) *Hypertext: a Psychological Perspective*. Chichester: Ellis Horwood. pp. 169–191.

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.

Duchastel, P. (1988) Display and interaction features of instructional texts and computers. *British Journal of Educational Technology*, 19(1), 58–65.

Engelbart, D.C. (1963) A conceptual framework for the augmentation of man's intellect. In P.W. Howerton and D.C. Weeks (Eds.) *Vistas in Information Handling*, Volume 1. London: Cleaver-Hume Press. pp. 1–29.

Gabbard, R. and Dillon, A. (1995) How does hypermedia affect learning? - an examination of the empirical literature. Paper presented at *Hypermedia'95*, Indiana University, Bloomington, Oct. 95.

Gordon, S. and Lewis, V. (1992) Enhancing hypertext documents to support learning from text. *Technical Communication*, second quarter, 305–308.

Greeno, J.G., Carlton, T.J., DaPolito, F. and Polson, P.G. (1978) *Associative Learning: A Cognitive Analysis*. Engelwood Cliffs, NJ: Prentice-Hall.

Hammond, N. (1989) Hypermedia and learning: who guides whom? In H. Maurer (Ed.) *Computer Assisted Learning*. Berlin: Springer-Verlag. pp. 167–181.

Hammond, N. and Allinson, L. (1989) 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.

Higgins, K. and Boone, R. (1990) Hypertext computer study guides and the social studies achievement of students with learning disabilities, remedial students, and regular education students. *Journal of Learning Disabilities*, 23(9), 529–540.

Hillinger, M (1990) Responsive text: a training environment for literacy and job skills. In *Proceedings of the Human Factors Society 34th Annual Meeting*. Santa Monica, CA: The Human Factors Society. 1422–1425.

Horton, S.V., Boone, R.A. and Lovitt, T.C. (1990) Teaching social studies to learning disabled high school students: effects of a hypertext study guide. *British Journal of Educational Technology*, 21(2), 118–131.

Howe, M.J.A. (1980) *The Psychology of Human Learning*. New York: Harper and Row.

Jonassen, D.H. and Grabowski, S. (1990) Problems and issues in designing hypertext/hypermedia for learning. In D.H. Jonassen and H. Mandl (Eds.) *Designing Hypermedia for Learning*. Heidelberg: Springer-Verlag. pp. 3–25.

Jonassen, D.H. (1993) Effects of semantically structured hypertext knowledge bases on users' knowledge structures. In C. McKnight, A. Dillon and J. Richardson (Eds.) *Hypertext: a Psychological Perspective*. Chichester: Ellis Horwood. pp. 153–168.

Jones, T. (1989) Incidental learning during information retrieval: a hypertext experiment. In H. Maurer (Ed.) *Computer Assisted Learning*. Berlin: Springer-Verlag. pp. 235–251.

Landauer, T. (1995) *The Trouble with Computers: Usefulness, Usability and Productivity*. Cambridge MA: MIT Press.

Landauer, T., Egan, D., Remde, J., Lesk, M., Lochbaum, C. and Ketchum, D. (1993) Enhancing the usability of text through computer delivery and formative evaluation: the SuperBook project. In C. McKnight, A. Dillon and J. Richardson (Eds.) *Hypertext: a Psychological Perspective*. Chichester: Ellis Horwood. pp. 71–136.

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.

Marchionini, G. (1990) Evaluating hypermedia-based learning. In D.H. Jonassen and H. Mandl (Eds.) *Designing Hypermedia for Learning*. Heidelberg: Springer-Verlag. pp. 355–373.

Marchionini, G. and Shneiderman, B. (1988) Finding facts versus browsing knowledge in hypertext systems. *IEEE Computer*, January, 70–80.

McKnight, C., Dillon, A. and Richardson, J. (1991) *Hypertext In Context*. Cambridge: Cambridge University Press.

McLuhan, M. (1964) *Understanding Media*. London: Routledge and Kegan Paul.

Nelson, T.H. (1988) Managing immense storage. *Byte*, January, 225–238.

Nielsen, J. (1995) *Multimedia and Hypertext: The Internet and Beyond*. Boston: Academic Press.

Norman, D.A. (1980) Twelve issues for cognitive science. *Cognitive Science*, 4, 1–33.

O'Shea, T. and Self, J. (1983) *Learning and Teaching with Computers*. Brighton: Harvester Press.

Polanyi, M. (1957) *Personal Knowledge*. Cambridge: Cambridge University Press.

Romiszowski, A.J. (1990) The hypertext/hypermedia solution — but what exactly is the problem? In D.H. Jonassen and H. Mandl (Eds.) *Designing Hypermedia for Learning*. Heidelberg: Springer-Verlag. pp. 321–354.

Smeaton, A.F. (1991) Retrieving information from hypertext: issues and problems. *European Journal of Information Systems*, 1, 239–247.

Stanton, N.A. and Stammers, R.B. (1989) A comparison of structured and unstructured navigation through a computer based training package for a simulated industrial task. Paper presented to the Symposium on Computer Assisted Learning — CAL 89, University of Surrey.

Stanton, N.A. and Stammers, R.B. (1990) Learning styles in a non-linear training environment. In R. McAleese and C. Green (Eds.) *Hypertext: State of the Art*. Oxford: Intellect. pp. 114–120.

Whalley, P. (1990) Models of hypertext structure and learning. In D.H. Jonassen and H. Mandl (Eds.) *Designing Hypermedia for Learning*. Heidelberg: Springer-Verlag. pp. 61–67.

Whalley, P. (1993) An alternative rhetoric for hypertext. In C. McKnight, A. Dillon and J. Richardson (Eds.) *Hypertext: a Psychological Perspective*. Chichester: Ellis Horwood. pp. 7–17.

Wright, P. (1990) Hypertexts as an interface for learners: some human factors issues. In D.H. Jonassen and H. Mandl (Eds.) *Designing Hypermedia for Learning*. Heidelberg: Springer-Verlag. pp. 169–184.