

The copyright of this thesis vests in the author. No quotation from it or information derived from it is to be published without full acknowledgement of the source. The thesis is to be used for private study or non-commercial research purposes only.

Published by the University of Cape Town (UCT) in terms of the non-exclusive license granted to UCT by the author.

Hypertext navigation, goal (dis)orientation, and the role of mental models: An empirical investigation.

Jonathan Ipser

Submitted for the degree of
Master of Arts (Research Psychology),
in the Department of Psychology,
Faculty of Humanities
University of Cape Town
September 2002

Supervisor: Frank Bokhorst

Acknowledgements

I would like to thank my supervisor, Frank Bokhorst, for always being available when I needed him, and for helping me think more clearly through otherwise complex issues. Thanks go to Bhavesh Prag for helping me carry out the experimental sessions, and to Floretta Boonzaier, Loraine Townsend, and Ines Meyer for providing moral support and being dedicated coffee-break accomplices. I would also like to thank the people in the Multimedia Education Group at UCT, and in particular Marion Walton and Maria Loophuyt, for their guidance and assistance. Prof Martin Hall, Assoc. Prof. Patrick Harries Dr. Chirevo Kwenda, Assoc. Prof. Andrew Smith, Melissa Steyn and Prof. Alan Morris, provided the web-based material for the Images of Africa course, and let me base my study on the notes they provided, for which I am grateful. I am also indebted to Prof Martin Hall, Assoc. Prof Patrick Harries, and Nick Shepherd, for completing the word-sorting task, thanks to Maria Loophuyt's offer to act on my behalf. And finally, a word of thanks to the PSY205S students who participated in my study, without whom my study would not have seen the light of day in the first place.

Abstract

An exploratory investigation into the capabilities of hypertext as an educational medium was conducted. To this purpose, a model was advanced in an attempt to provide a theoretical foundation for the explication of hypertext navigation and disorientation in terms of relevant user characteristics. 75 2nd year psychology students from the University of Cape Town participated in a study in order to investigate the claim that the motivational theory of goal orientation has the greatest utility in accounting for the exploitation of links within a hypertext, a necessary precondition for the development of mental models which incorporate both the conceptual and structural elements required for effective learning from, and the reduction of disorientation within, this medium. The results obtained, although compromised, provide some support for this argument. The paper concludes with a discussion of the methodological limitations of the study, as well as possible directions for future research.

Contents

Acknowledgements

Abstract

Literature Review	1
Introduction	1
Hypertext as a godsend for educationalists	2
The problem of Hypertext disorientation	7
Concluding statements regarding the relationship of hypertext navigation and learning	15
Theoretical foundations of an integrated theory of hypertext learning . . .	16
Learner orientated factors	16
The impact of text characteristics on learning	28
Hypertext navigation and motivation	41
How motivated does one have to be to learn?	41
The theory of Goal Orientation	42
Conclusion and Hypotheses	47
Method	52
Participants	52
Materials	52
Measures	55
User characteristics	55
Navigation	56
Mental Models	60
Disorientation	65
Learning	66
Research design	66
Procedure	66
Data analysis	70
User characteristics	70
Hypertext characteristics	70
Navigation	71

Mental models	76
Learning	77
Results	79
Characterisation of the data	79
Reliability and validity	79
Exploratory analyses	84
The effect of goal orientation on navigation	94
Mental models and disorientation	96
Structural aspect of mental model	96
Comparison of results with Smith's measure of lostness	100
Conceptual aspect of mental model	100
The influence of mental models on factual knowledge gain	104
The influence of goal orientation on recall	105
Knowledge gain vs. recall	106
Discussion	109
Goal orientation	111
The effect of goal orientation on hypertext navigation	113
The measurement of navigational pathways	114
The issue of mental models	116
Disorientation	118
Learning and Recall	119
Limitations	120
Measure of conceptual knowledge	123
General discussion	126
Future directions for research	129
Conclusion	132
References	134
APPENDIXES	
Appendix A	1

Graphical representation and descriptive statistics for the navigational clusters in the Hierarchy condition	2
Graphical representation and descriptive statistics for the navigational clusters in the Network condition	3
Tests of assumptions for analysis of covariance of affect of independent variables on recall	4
Appendix B	5
Hypertext System diagrams and screenshots	6
Comparison of navigational structure of hypertexts	7
Screenshot of homepage and sample hypertext node	9
Path diagrams for participants in Network condition	11
Appendix C: Questionnaires	14
General - Knowledge Questionnaire	15
Goals Inventory Questionnaire	17
Factual - Knowledge Questionnaire	18
Motivated Strategies for Learning Questionnaire	22
Computer and Internet Skills scale	24
Miscellaneous questions	25
Fact Finding Task	26
Word sorting task	28
Appendix D: Additional Materials	30
Custom LSA data corpus	31
Composition of tasaALL space	32
Author Note	
Colophon	

List of Tables

1	Kintsch's typology of inferences	38
2	Correlations between variables characterising research participants . .	80
3	Correlations amongst navigationally dependent metrics	82
4	Correlations between hypertext efficacy, background knowledge and navigational measures	83
5	Comparison of means and standard deviations for clusters in the Network condition	85
6	Comparison of means and standard deviations for clusters in the Hierarchy condition	86
7	ANOVA results for differences between goal - orientation and associated variables for the two cluster solution	88
8	Motivational variables for 3 k-means cluster analysis	91
9	ANOVA results for differences between goal - orientation and associated variables for three cluster solution	93
10	Results of the effect of goal orientation on the navigational metrics. .	95
11	Results of the effect of goal orientation on navigational metrics in the Network condition	96
12	Descriptive statistics for mental model - dependent variables as a function of hypertext structure and goal orientation.	98
13	ANOVA summary table of specific effects of goal orientation and hypertext structure on mental model variables	99
14	ANOVA summary table of specific effects of Smith's measure of lostness	102
15	Descriptive data for two LSA procedures	102
16	ANOVA summary table for WebLSA	103
17	ANOVA summary table for CustomLSA	103
18	Descriptive statistics for analysis of the recall of factual knowledge . .	105
19	ANCOVA summary table of specific effects for analysis of interaction of goal orientation and hypertext on recall	106

List of Figures

1	Conceptual model of relationships between goal orientation, navigation, mental models and knowledge acquisition.	49
2	Example of a graph and matrix representation of simple network hypertext (grey arrows represent the hyperlinks)	59
3	2 K-means cluster profiles for goal orientation and associated variables	90
4	Profile of 3 k-means clusters on motivational variables	91
5	The interaction of goal orientation and hypertext structure on disorientation and structural hypertext knowledge	97
6	Effect of interaction of goal orientation and hypertext on Smith's measure of lostness	101
7	The effect of goal orientation on recall in the Network and Hierarchical conditions	107
8	Revised model of relationships between goal orientation, navigation, mental models and knowledge acquisition.	109

Literature Review

Introduction

This dissertation represents an attempt to consolidate two aspects of hypertext research which are commonly treated separately. The first area of interest concerns itself with a feature of hypertext which is considered to serve it to good advantage as an educational medium. This is the non-linear nature of the hypertext. Praised both for the greater control it gives the student, as well as the naturalistic quality it imbues to the knowledge stored within it, it has also been identified as its greatest weakness. The majority of hypertext system development has been geared towards correcting what is seen as its most pernicious attribute, that of disorientation, or "getting lost in hyperspace" (Edwards & Hardman, 1989).

Once the nature of the non-linear quality of hypertext has been fully considered, the measures taken to promote its educational aspect and to minimise any attendant hypertext disorientation will be discussed. It is in the attempt to place these measures within a theoretical framework that one of the major themes of this dissertation will emerge. This is that there resides an inherent tension between those qualities of a hypertext most beneficial for learning, and many of the developments inspired by the need to facilitate navigation within the hypertext system. It is instead suggested that navigation through a hypertext system and comprehension of the content found within the system are two sides of the same coin, and navigational aids which do not take learning into consideration, do so at their own peril. The effort in developing these aids would be better invested in conducting research into those characteristics of the hypertext user which may affect their patterns of navigation, and to determine whether the features of a specific hypertext which promote learning vary as a function of those user characteristics. One trait which is likely to be particularly pertinent in this regard is that of motivation

The argument which will be outlined will of necessity contain a speculative element, attempting as it does to integrate research from a number of different areas. The wide-ranging nature of the thesis will unfortunately preclude the possibility of conducting an in-depth analysis of theoretical constructs to which a number of books have already been devoted. I will, however, attempt to provide greater coverage of the select topics on which the argument developed in this

dissertation most greatly depends.

Before proceeding, it is advisable to clarify the terminology to be used in this paper. Although the terms *hypertext* and *hypermedia* are frequently used interchangeably in the literature (Dieberger, 1994), and both contain elements of interactive non-linearity (Rouet & Levonen, 1996), the latter can contain a combination of images, video and audio media with the text presented¹, while the text stands alone as the only changing element within a hypertext. The thesis presented in this paper will focus primarily on hypertext for the following three reasons: (1) One of the largest factors impeding the comparison of the efficacy of hypertext/media instructional programs with one another, and with traditional texts, revolves around the great variety of (often *ad-hoc*) features associated with the systems used in the programs². A focus on the only element most commonly associated with all of these systems, namely the hyperlinked text, will therefore expand the generalisability of the conclusions reached in this paper. (2) In a point related to the first, the widening of the scope of interest through selective attention to the features of hypertext brings the educational implications of not only closed hypermedia instructional systems into the purview of the research agenda, but also the loosely structured chaos of the internet (3) Finally, the emphasis on text makes hypertext amenable to analysis in terms of theories of text comprehension.

Hypertext as a godsend for educationalists

There is a rift between scholars as to the true worth of hypertext as an educational medium. On the one hand hypertext has been widely cited as possessing unparalleled potential within the field of education (Landow, 1990; Strommen, 1995). On the other hand, there has been an increasing tendency over the last decade or so to express doubt about its status as a panacea for all of the world's educational woes. This section will be dedicated to detailing the reasons behind the emergence of these doubts.

In an article entitled *Misleading theoretical assumptions in hypertext/hypermedia research* (1997) S. Tergan questions three widely held beliefs regarding the efficacy of hypertext/hypermedia as a learning tool. Two of these misconceptions are of particular relevance to hypertext. The first, dubbed the

¹As captured in the formula: hypertext + multimedia = hypermedia

²Refer to Tergan (1997a) for a discussion of the claim that these instructional features of hypertext studies are frequently confused with the system's delivery features.

Plausibility Hypothesis (Tergan, 1997b), postulates that the linking of nodes within the text determines a structural and functional equivalence between it and the human mind (Jonassen, 1993)³ and that the provision of those links allow for self-regulated learning.

This belief is commonly acknowledged as owing its ancestry to Vannevar Bush's seminal article, *As we may think* (Bush, 1945), in which a direct comparison is drawn between the associative capability of the human mind, and the interconnected hypertext system (Bartasis & Palumbo, 1995). By implication the content matter presented in a hypertext system should impress itself more readily on the student's mind than it would were it to be accessed in a more conventional paper-based format, existing as it does in what amounts to a pre-processed form. It is this same characteristic associative network structure, in which the clicking with a pointing device, such as a mouse, on one of multiple activated word sequences in a text (otherwise known as "hotwords" or "links") enables the reader to traverse the connection between two knowledge nodes (web-pages), which is presumed to give the student greater control over the sequence of nodes that they visit, and hence over their entire learning experience. The greater learner control afforded the user is regarded as beneficial, as the user is assumed to be the expert on what they need to know (Charney, 1994; Jonassen, 1993), and because of the perception of the ideal student as one who is actively engaged with the material presented to him/her, rather than the passive recipient of knowledge provided by someone else. It is the student's engagement with the text which facilitates their construction of a model of the hypertext and information in it, and it is for this reason that people asserting this belief have been labelled "radical constructivists" (Tergan, 1997b).

Unfortunately, such an optimistic outlook reveals itself as naive when tested empirically. This was apparent in an examination of the meta-reviews conducted of the learner control literature, which found no benefit in the increasing learner control made possible through the medium of computer-assisted instruction (Niemiec, Sikorski, & Walberg, 1996). Moreover, on a more general level, while the results of a meta-review of 46 studies comparing the academic performance of participants in hypermedia learning conditions vs. non-hypermedia

³A point of view finding its most radical expression in the Principia Cybernetica Project's prediction that the common associative qualities of the internet and human mind enable their co-evolution as part of a "Global Brain" (Heylighen, 1995)

conditions was that the hypermedia format was beneficial (Liao, 1999), this conclusion needs to be tempered with the observation that the results of 17 of the 46 studies reported favoured the non-hypermedia forms of instruction (traditional instruction, computer assisted instruction or videotapes) over the hypermedia (computer-based interactive video-discs, computer simulators, or interactive multimedia) condition. In addition, the largest positive effects for hypermedia were found in those studies employing the smallest sample sizes, and which made use of a repeated measures experimental design, both factors which reduce the weight one can attach to the significance of their findings.⁴ Finally, a review of the impact of 92 benchmark measures for usability features of hypertext (Nielsen, 1989) found that these features had little effect. Those features which improved performance on the testing criteria most substantially were more strongly orientated towards differences between subjects (age, motivation and activity level) than the systems they used. This is consistent with other evidence regarding the importance of user characteristics (Balciytene, 1999; Lawless & Kulikowich, 1996, 1998), as well as task differences (Jonassen, 1993) in determining the extent to which people benefit from a hypermedia system. Claims regarding the educational potential of hypertext can, accordingly, not afford to omit consideration of these factors.

Advocates of the radical constructivist approach might still hold out hope that future studies will support their point of view, were it not for a number of theoretical weaknesses which plague such a simplistic interpretation of the benefits of hypertext. First of all, the characterisation of the distinction between traditional text and hypertext as one of linearity is misconceived. So-called "linear" texts, and in particular, works of literature, often surpass the constraints of the medium in which they are set by use of complex relational devices and associations, which resonate in the mind of the reader (Whalley, 1990).⁵ Reference works, such as dictionaries, also possess non-linear features (Panero, 1995). Conversely, the hierarchical structure of hypertexts can often be traversed in a sequence which preserves the linear flow of the content. Should the the reader, moreover, decide to

⁴The interpretation by the author of the reduction in effect size reported with larger sample sizes as indicative that hypermedia is most suitable as an educational tool for small groups, is also somewhat dubious, given the conventional recommendation that large samples be employed as a more reliable indicator of whether a phenomena exists at all.

⁵It is for the same reason that poetry is relatively opaque to text analysis techniques (Kintsch, 1998).

practise the choice given him/her by clicking on the hyperlinks, the supposed freedom s/he enjoys would only extend to the selection of the links provided in a hypertext⁶, with the links themselves being pre-determined by the author of the hypertext (Schroeder & Grabowski, 1995; Calvi & De Bra, 1998). This forms the basis of the argument that hypertexts impose greater constraints on the reading strategies of the user than traditional linear texts (Altun, 2000)⁷. While the advantages of the ability of the reader in the conventional text to move unhindered from one part of the text to the other is often undermined by the provision of a single coherent reading throughout the text, the potential of hypertext in overcoming this restriction would only be realised were the pathway followed within a hypertext optimally suited to each particular reader.

There is good reason, however, to suspect that this is not usually the case. Empirical evidence suggests that students frequently do not employ the most effective learning strategies when given a study task (Jonassen, 1993). There is also a large body of evidence which indicates that it is only under certain special circumstances that people display any competence at all in determining how well they know a particular topic (Lin & Zabrocky, 1998; Glenberg & Epstein, 1987; Maki, Foley, Kajer, Thompson, & G., 1990). The lack of ability to accurately assess whether one understands something, known in the literature as poor *calibration of comprehension*, makes a lie of the claim that hypertext users are automatically the most qualified, on the basis of what they know, in determining what they need to extract from a hypertext.

The possibility that students do not know the status of their knowledge implies that greater freedom of movement granted students is not automatically equivalent to greater control, a fact which has prompted certain researchers to suggest the combination of (the more traditional) instructional and constructional modes of teaching as optimal.⁸ This point acquires extra salience when one considers that the hypertext user is only in a position to make an informed decision as to the pathway s/he should follow in circumstances in which the hyperlink's role and the

⁶Although refer to Rouet (1996) for a different perspective.

⁷This point may be particularly applicable to readers who are both unfamiliar with the navigational strategies at their disposal within a hypertext, as well as the layout of the nodes in a particular hypertext.

⁸cf. Shabo, 1997. This might also account for the finding that hypertext is most effective as an adjunct to traditional teaching methods (Liao, 1999), a role frequently recommended for it (Whalley, 1990).

relationship between nodes which it embodies is made explicit (Nielsen, 2000; Kaplan, 2001). It is for this reason that Otter (2000) recommends that hypertext designers take pains in doing just that. The nature of the links are usually left in implicit form in free-for-all environments such as the World Wide Web, however, and are likely to remain so for the foreseeable future.⁹

The case against the identification of the non-linearity of educational media and learner control is further strengthened by the realisation that a student's choice of the optimal pathway through a hypertext system also presupposes the development of the cognitive equivalent of a map of it. Unfortunately, in the absence of an externally provided map of the hypertext, this requires that the student first explore the entire hypertext, thus rendering void any discussion of the selection of the optimum path. In order to pre-empt the argument that familiarity with hypertext systems would do away with the requirement that students possess knowledge of the structure of a particular hypertext, proponents of this line of reasoning would have to argue that generic knowledge of hypertext structure assists in the prediction of the organisation of a particular hypertext system, on the basis of the content contained within it. This claim would be highly suspect, however, as the lack of a one-to-one mapping between the structure of a hypertext and the semantic relations between concepts within its nodes (Whalley, 1990; Dillon, McKnight, & Richardson, 1993; Tergan, 1997a), as highlighted by the large variation in links constructed by different designers for the same text (Alschuler, as cited in Charney 1994), precludes the possibility of anticipating the structure of the hypertext from the material it covers.¹⁰

It will be argued, however, that even hypertext systems which provide graphical representations of the hypertext structure do not foster learner control. The claim that they do is first brought into doubt by the allusion to evidence indicating that graphical maps of a hypertext system facilitate the speed at which people complete tasks in which they are required to find particular pieces of information within a hypertext, but do not aid them in finding information of

⁹Although look at the World Wide Web Consortium's specification of the emerging XML (Bray, Paoli, Sperberg-McQueen, & Maler, 2000), and in particular XLink (De Rose, Maler, & Orchard, 2001), technology, for the capabilities it offers in annotating links on the web. Weinreich (2001) also provides discussion of the possibilities available for the enhancement of hyperlinks

¹⁰It is important that a map be formed of the entire hypertext, as a strategy of choosing the most strongly associated links from node to node, while optimal on a local level, may emerge as sub-optimal in the context of the entire system. This point will be elaborated on at a later stage.

particular relevance to their task. (Boechler, 2001). This would appear to indicate that faith in the benefit of such a map in highlighting the best path for a particular student is misplaced. Furthermore, despite the fact that mention of the implications of the non-linear nature of hypertext for learning almost inevitably leads to a discussion of navigational aids, there is no good reason to equate the ability to navigate through a hypertext with its utility as an educational instrument, as some have done (Calvi, 1997). In fact, it will become apparent through reviewing literature spanning many areas of academic enquiry, that an increased emphasis on improving navigation might actually limit the educational potential of the medium.

The problem of Hypertext disorientation

The discussion thus far has focused on reasons why the non-linear aspect of hypertext should not be accepted without question as elevating it above traditional texts. The faith in the benefits which have been attributed to hypertext in terms of learner control appears largely without substance. There are, however, good grounds to argue that the case which has been made up till this point is not strong enough, as it commits the oversight of not mentioning one of the most widely cited problems associated with the substitution of a hyperlinked hypertext structure for the linear constraints of the traditional text. This is the danger of hypertext disorientation, otherwise known as becoming "lost in hyperspace".

Hypertext disorientation has been claimed as having the potential to seriously limit the educational benefits of the medium (Nunes & Fowell, 1996). The oft-cited definition of disorientation comprises a set of characteristics of the disorientated hypertext user (Edwards & Hardman, 1989), namely one who does:

- not know where to go next
- knows where to go but not how to get there, and
- does not know where he/she is at any particular time in relation to the entire hypermedia system.

This "definition" of hypertext disorientation leaves sufficient elbow room for a number of divergent perspectives, amongst them that (a) disorientation is not intrinsic to hypertext (Bernstein, 1999), but unless deliberately effected (Bernstein, 1998), is actually indicative of poor writing style, (b) that it only seriously affects

novices (Rouet & Levonen, 1996; Altun, 2000) and (c), that it is actually necessary for effective learning (Mayes, Kibby, & Anderson, 1990).

This broad characterisation of hypertext disorientation allows one to re-evaluate the GVU's WWW usage survey report (GVU.96, 1996), in which it was stated that only 6.5% of World Wide Web users who participated in their survey did not know where they were on the Web. Although this was referred to in the report as the "classical lost in hyperspace problem" (p7) (Pitkow & Kehoe, 1996), this broader conceptualisation includes the observation that 34.5% of Web users were not able to find a page which they knew was out there, that 23.7% were not able to find a page once visited, and that 14.3% of the users reported not being able to visualise where they have been and where they can go, as symptoms of disorientation (Pitkow & Kehoe, 1996).¹¹ The last statistic mentioned is particularly consistent with the idea that hypertext users often lack a cognitive map of the structure of a hypertext, an issue to which we now turn.

A person's movement through hypertext is frequently expressed as "navigation", and it is in this same geographic-spatial idiom that hypertext disorientation has been conceptualised. It should therefore come as no surprise that the remedy to such disorientation has often been cast in the form of the development of a map. This way-finding mechanism can either be provided externally, or assume form in the mind of the hypertext user, in what is known as a cognitive map. Dillon (1993) identifies three different levels of knowledge as applicable to navigation in both geographic and information spaces, and which have been used analogously in discussions of the emergence of a cognitive map. These are, in order of increasing inclusiveness, landmark, route, and survey knowledge.

Landmark knowledge refers, within the context of navigation, to the awareness of objects employed as orientating devices for further exploration. In order to traverse the distance separating one landmark from another, a person needs to have knowledge of the specific path to follow between them, a form of knowledge known as route knowledge. Landmark and route knowledge are both required for the development of survey knowledge, which consists of an awareness of how a particular landmark relates to the entire constellation of landmarks within a particular system.

¹¹It should be borne in mind that these statistics are probably conservative estimates of disorientation amongst the population of web-users, due to the self-selection of respondents who have the expertise to successfully complete the online questionnaires (GVU.96, 1996).

Although originally applied within a geographic context, these forms of knowledge are useful in the analysis of navigation within a hypertext, especially with regards to their involvement when navigation breaks down. Their definition is consistent with those points which have been characterised as the constitutive elements of disorientation. For example, knowing where to go but not how to get there points to a deficit in route knowledge, while not knowing where one is at any particular time in relation to the entire hypermedia system implicates a lack of survey knowledge.

Disorientation has been operationally defined in a number of studies as a ratio of the number of nodes required to reach a target node to the total number of nodes visited (Chou & Lin, 1998; Dias & Sousa, 1997). Although hypertext disorientation, as defined thus, probably results from a lack of both route and survey knowledge, the evidence suggests that subjective feelings of lostness are more sensitive to inadequate levels of the latter than the former. For instance, Otter and Johnson's (2000) finding that the size of a hypertext is cited by students as one of the biggest causes of disorientation, is explicable when one considers that the adequacy of a person's survey knowledge of a hypertext would be strongly influenced by its size.¹² Route knowledge is no substitute for the awareness of one's location relative to the entire system, in circumstances in which a hyperlink displaces one to a node lying outside a familiar pathway. The observation in a particular study (Modjeska & Marsh, 1997) that users of websites tend to prefer a base node as a starting position for further navigation rather than orienting themselves on the basis of landmark nodes, coupled with the finding that the study's participants were not able to provide accurate drawings of site structure, suggests that hypertext users will only have the confidence to make effective use of landmark knowledge once they have some idea of where the landmarks lie in relation to the rest of the site.^{13 14}

Although Dillon cites one study in particular which he regards as suggesting the possibility of the development of survey knowledge amongst hypertext users (p184), there are several factors mitigating against it. Readers of traditional texts have the advantage of being able to employ narrative schemas, or stereotyped

¹²Refer to Ahuja (2001) for an argument in favour of using subjective reports of lostness to the inference of disorientation from the actions of the hypertext user.

¹³In the absence of survey knowledge distinctive landmarks may masquerade as base nodes, despite the fact that they do not lead to any useful pages (Panero, 1995)

¹⁴While it is true that the World-Wide-Web is too large by many orders of magnitude for users to develop a conception of its entire structure through exploration (Boechler, 2001), it is still sensible to speak of the development of survey knowledge of particular websites

perceptions of narrative structure associated with particular text genres, in facilitating their comprehension of a text (Charney, 1994). The structure encoded in such schemas is not available for hypertext, which does not comply with any standard format (Boechler, 2001; Foltz, 1996). In addition, the burden of having to apportion limited cognitive resources to the construction of a model of the site structure (particularly in the role that those models might play in self-regulated learning (Winne, 1995)), and simultaneously having to learn the information stored in the hypertext, is liable to induce a state of cognitive overload (Albers, 1997; Jonassen, 1993) amongst hypertext users.

When these considerations are taken into account, it becomes readily apparent why a number of the technological add-ons developed in order to counter hypertext disorientation, including global and local maps, externalise the system structure¹⁵. In addition to portraying the structure of a hypertext, or portions thereof, maps in the form of graphical browsers frequently perform the function of facilitating movement within the system. The finding that hypertext users are often themselves not able to replicate the structure of the hypertext they have read (Panero, 1995; Modjeska & Marsh, 1997), appears to support the wisdom of providing these structural supports. The observation that the inclusion of a site map leads to less confusion in a hypertext than the use of hyperlinks alone (Beasley & Waugh, 1995), is consistent with these findings.

There is, however, sufficient contradictory evidence to suggest that it would be unwise to adopt an uncritical attitude towards the use of graphical maps. Dias and Sousa (1997) interpreted the findings of a study they conducted as indicating that a map of a non-hierarchically structured or network hypertext was not effective in helping locate information relevant to the completion of a task. Evidence has also been cited that while hierarchically structured navigational aids are more effective for navigational purposes than alphabetical indices (Rouet & Levonen, 1996), this advantage is only specific to hierarchically structured hypertexts (Beasley & Waugh, 1995; Simpson & McKnight, 1990), with other forms of representation, such as spider-maps, inducing greater degrees of disorientation in these systems (cf. Beasley (1995) for references to additional studies which support this finding). While researchers who have made these discoveries generally conclude that the form of

¹⁵Other navigational aids include methods such as the use of "bread-crumbs" or history lists to indicate nodes which have already been visited.

representation used should reflect the structure of the hypertext system, one can anticipate that this may be problematic with hypertexts which are designed in a more free-form and less rigidly hierarchical fashion. Finally, there is empirical evidence to suggest that the mere provision of structural maps is not sufficient to induce learning (Jonassen, 1993).

Given the conflicting evidence regarding the benefits of using graphical browsers and maps, how much weight should be attached to claims such as those made by Calvi (1997), that the "user's spatial ability, that is, the user's capability at orienting him/herself in a physical setting, has been identified as one of the main factors effecting comprehension in a hyper-textual environment"? (p306) The observation that graphical browsers only represent the exploitation of one area of the spatial spectrum in representing the layout of a hypertext, albeit a popular one, suggests that the credibility of this statement might still be rescued through the examination of other forms of spatial representation.

The perception of hypertext as spatial in nature finds its most explicit expression in attempts to achieve various degrees of conversion between the virtual space in which its information is stored, and an Euclidean, 3 dimensional space. These attempts have ranged from the superimposition of different regions of a hypertext onto a 2 dimensional plane, so as to render what is known as a 2 1/2 dimensional space (Shum, 1990), to the representation of a textual virtual environment as a city (Dieberger, 1994, 1995, 1996; Dieberger & Frank, 1998), to the full immersion of the hyperspace user in a virtual reality environment (Modjeska, 1998; Frecon & Smith, 1998). The benefits cited in representing information in these ways include taking advantage of the innate spatial sense people are endowed with, overcoming the limitations of screen size plaguing the 2 dimensional display of large hypertexts (Chen & Czerwinski, 1998; Chen, 1998), and representing spatial relations using metaphors of real world spaces, such as the city, through which people are expert in navigating.

Unfortunately, these alternative representations come with their own set of potential problems. For instance, the discrepancy between the ability of students to traverse, with one click, the distance between a node and another spatially portrayed as lying far from it, referred to as "tunnelling through space" (Dieberger 1994, p74) has been known to disorientate the hypertext user. On one level this can be ascribed to the fact that each link in a hypertext is semantically over-determined.

The convention of representing increasing differences in the semantic nature of the content matter contained in each node, through increasing the representation of the distance between the two nodes, might yield a representation which is not consistent with their semantic distance, as determined by the reader. The background knowledge that a student brings with him/her to the hypertext will affect the associations they draw between linked nodes, and determine to a large extent the degree of match between the semantic and the spatial layout of the hypertext.

This dilemma can also be read, however, as indicating limitations in the usefulness of space as a metaphor of hypertext.¹⁶ As Boechler (2001) has indicated, the characteristics of space which have been used as criteria in estimating the similarity of two physical objects do not appear to translate easily to hypertext. Of the three criteria of minimality, symmetry and triangular inequality (Tversky, 1977)¹⁷, symmetry, or the notion that “the distance between two points is the same no matter the order” (Boechler 2001, p31), seems to be particularly inapplicable in an environment where the semantic “distance” between concepts within two nodes is in a state of continual fluctuation. The reader’s perception of the similarity of two connected nodes in a hypertext can only be understood as dynamic in nature, given that it is a function of the reader’s very navigation through that hypertext, and the subsequent information that s/he acquires.

The capability of these distance axioms in capturing the nature of hypertext representations in which the connections between nodes are made explicit, while failing to apply to that which is being represented, indicates a mismatch which might seriously effect the utility of such representations in terms of both hypertext navigation and learning. These observations, in combination with that of semantic over-determinism, indicates that a mismatch with the representation of a hypertext can occur on both the level of the semantics of the hypertext contents, as well as the hypertext’s structure. It is these discrepancies which are more likely than not to lead to disorientation. Put another way, users of hypertexts in which there is a large discrepancy between the hypertext (content or structure), and the representation thereof, would be predicted as being more likely to experience feelings of disorientation. A case in point might be in hypertext systems which employ

¹⁶Refer to Landow (1990) for one of the earlier expressions of doubt as to the utility of space as the metaphorical framework for hypertext.

¹⁷See Aldenderfer (1984) for a discussion of these criteria in the context of the definition of a “metric”, and their applicability to statistical measures of similarity.

graphical maps, in which the presentation of the relationships between nodes are inflexible, due to the over-specification of their nature — through making them visible as lines, for instance — and where the material conveyed is sufficiently complex conceptually as to warrant the conception of the relations between knowledge nodes as multiple.

Ironically, this criticism of the use of spatial representations in counter-acting hypertext disorientation suggests a reason to continue subscribing to the metaphor of hypertext as a “hyperspace”. Forms of representation which provide an analog representation of the spatial relations found between elements of the physical environment, can, paradoxically, provide more information about the nature of what is depicted when less information is contained within the depiction (Schonotz & Bannert, in press). This has implications not only for the efficient use of screen “real estate” in the representation of the hypertext, but is also important with respect to the cognitive representation of the hypertext’s structural and imbedded conceptual relations. The limited capacity of a person’s cognitive resources renders it imperative that cognitive models of the hypertext are as information efficient as possible.

One of the most compelling justifications for the spatial representation of hypertext is therefore its potential economy as an information storage mechanism. In order for this potential to be realised, however, the representational format must possess certain qualities. The nature of these attributes becomes most apparent when examining the typical rationale which might plausibly be provided for the spatial representation of hypertext. One such argument would be that diagrammatic and Euclidean representations of hyperspace take advantage of the ability of images to render implicit conceptual relationships between segments of text explicit, facilitating the display of the relative positions of the nodes in the system, and subsequently fostering the user’s survey knowledge of the hypertext. The graphical presentation of the hypertext, in the form of a site - map, for instance, side-steps the possibility of the user suffering from cognitive overload, thereby freeing the resources required for effective learning and navigation.

While there is evidence to support the validity of this rationale with respect to simple systems, and students with low levels of background knowledge appear to benefit the most from the provision of graphical maps (Balcytiene, 1999), presumably for many of the reasons given, the tendency of such explicit

representation to inflict damage on the user's conceptualisation of the relations between concepts contained in connected nodes, as is more likely to be the case in conceptually complex and less easily categorisable material, would be expected to result in greater degrees of disorientation. This could be overcome, however, through resorting to the portrayal of relations between nodes in a more abstract and less explicit fashion.

Instead of portraying relations between nodes statically by means of connecting lines, one could indicate differences in a manner akin to that employed in the document management system, VIKI2 (Marshall, Shipman, & Coombs, 1994; Marshall & Shipman, 1997; Shipman, Hsieh, Maloor, & Moore, 2001; Shipman, Marshall, & LeMere, 1999). This "spatial hypertext" system allows the user to signify differences between node types through assigning different shapes or colours to different nodes, for instance, while conceptual relatedness between the content within nodes can be conveyed in an analog fashion by the proximity with which the nodes are placed in relation to one another. A similar convergence between the user's conceptualisation of a system and the manner in which the system is portrayed can be achieved in environments in which there is an abundance of sensory stimuli, as found in virtual reality representations, and which allow the user to assign multiple and idiosyncratic meanings to those aspects of the environment which the users latch onto in forming their own representations of the connection between knowledge nodes¹⁸. For example, Gamberini and Bussolon (2001) describe a dual graphical browser/VR system in which readers are free to assign significance to arbitrary features of that environment, such as the colour of a wall. The reader might then form an association between the colour of the wall representing one node and the wall colour of a related node.

The information storage efficiency of spatial representations can accordingly only be called on as support for their use, to the extent that they can store the information that hypertext users assign to them. The rendering of implicit relations explicit, while apparently providing more information, does harm to the student's ability to integrate their own background knowledge with what is being learnt. Therefore, the representation of lines between nodes, apart from being cumbersome

¹⁸Although such a flooding of stimuli tends to be expensive in terms of bandwidth. One should also take note of the argument put forward by Noel and Hunter (2000) that synthetic environments do not need to capture physical reality down to the smallest detail in order to achieve task - adequate "psychological reality".

when the system reaches a particular size (Dillon et al., 1993), actually reduces the information content of the system. This is particularly the case if one accepts the notion that the ability to recall information is dependent on the integration of incoming and background information (Kintsch, 1998).

The description of implicit images as both information rich and poor is similar to that given of mental models by Schonotz and Bannert (in press). While the 'Plausibility hypothesis' central hypothesis regarding the effects of similarity between hypertext and the human mind was given short shrift, it does not seem unreasonable to claim that a representation of hypertext which mimics the features of one's model of it is more likely to be successful as a navigational aid than one with opposing features. The "plausibility" of this hypothesis would appear strengthened by the degree to which an implicit structure accommodates differences in user's prior knowledge, by, in effect, leaving the user to fill in the gaps. This line of argument will be elaborated upon in the closer examination of the nature of mental models, and the inference generation mechanism by which those gaps are filled.

Concluding statements regarding the relationship of hypertext navigation and learning

The factors that have been considered make it apparent that instead of combating the phenomenon of becoming lost in hyperspace, navigational aids can actually function as significant causes of hypertext disorientation. This is true of systems in which the external representation of the hypertext structure does not mesh with the user's conceptualisation of relations between nodes. A navigational representation will only achieve its potential if it is flexible enough to allow the user to make the connections needed in order to capture the dynamically fluctuating nature of their representation of complex material.

Although the preceding discussion did not make a strong distinction between the navigation of a hypertext, and learning the material contained within it, when the distinction needs to be made, it has to be conceded that once user input is acknowledged as vital in effectively navigating complex hypertext systems, the distinction falls away. Effective navigation implies effective learning, an assertion which resonates with Spence's (1999) portrayal of navigation as the development of the cognitive map of a hyperspace. This realisation indicates that the problems identified with the development of a student's mental map of a hypertext might

actually be more symptomatic of the naivety of the belief in the automatic transfer of the map provided to the cognitive domain of the user, than of any weaknesses in the argument that the development of a cognitive map facilitates navigation.

Actually, as will become apparent with a closer examination of a construct which has utility in the characterisation of hypertext education, mental models are better suited in overcoming the non-alignment of hypertext content and structure than the cognitive map's exclusive focus on structural relations between nodes. A more reasonable claim might be that navigational aids possess the potential to bridge the chasm between comprehension and navigation, on the proviso that they simulate and stimulate features of the mental model of the hypertext which enable this convergence to occur.

The remainder of the literature review will be devoted to the explication of three theoretical frameworks, with the purpose of further elucidating the nature of the relationship between hypertext navigation, disorientation, learning, and the cognitive models these phenomena rely on. These theories are, in order of discussion, Constructivism, the theory of Mental Models, and the Construction-Integration hypothesis. It will be contended that the active participation of the student is a vital component in the consolidation of these theoretical frameworks, upon which the realisation of the educational implications of hypertext learning depends. It will further be argued that the motivational theory of Goal Orientation provides the mechanism through which the student can effect this consolidation, in the interests of combining navigation and comprehension.

Theoretical foundations of an integrated theory of hypertext learning

The discussion of the literature which follows will be partitioned into two sections. The first deals with those areas which place a heavier emphasis on the role of the student in an educational setting, while the second section concerns itself with the means by which qualities of text in general are hypothesised as interacting with the student whilst learning.

Learner orientated factors

The Constructivist Theory.

Constructivism is the theory of education most strongly allied to hypertext

research. Constructivism (Eklund, 1995; Tergan, 1997a), or variants thereof, such as the Experiential (Nunes & Fowell, 1996) or Discovery Learning school of thought (Bartasis & Palumbo, 1995) has formed the basis of many optimistic approaches to the effects of using hypertext on learning. These theories all regard the student as an active participant in the educational process (Jonassen & Grabinger, 1990; Jonassen, 1993). They argue that the student's participation in processing course material fosters an idiosyncratic learning experience, which in turn manifests itself as an individualistic representation of knowledge in a particular domain. The unexamined notion of constructivism most commonly applied to hypertext, coined as "radical constructivism" (Tergan, 1997a), expounds two beliefs regarding the efficacy of hypertext as an educational medium. They are the, by now quite familiar, propositions that the medium's non-linear nature encourages student participation through the freedom of choice given the student in choosing which particular link s/he wishes to follow, and that the model of knowledge structure embedded within the hypertext will stimulate the construction of the user's own internalised model of how concepts in the hypertext are related to one another.

The vacuous theoretical arguments and the lack of empirical backing for the first assertion has already been exposed. It has also become clear that it is in response to the dangers of the student's greater freedom of exploration that hypermedia designers have employed the tactic of externalising the structure of their systems in the form of maps and graphical browsers. These navigational aids have received the lion's share of attention, as they make the putative connection between the knowledge model in the hypertext and the user's own mental model most explicit. It is indeed the constructivists contention that these externalised aspects of the computer interface will somehow subliminate into the student's mental model (Calvi, 1997). This provides the rationale for the design of the hypertext on the basis of an expert's cognitive mapping of the knowledge domain covered (Eklund, 1995; Jonassen & Grabinger, 1990; Jonassen, 1993) According to this viewpoint, the externalisation of that expert knowledge structure should ease the way for its absorption by the student (Reed, Ayersman, & Liu, 1996)

The problematic nature of this line of reasoning becomes apparent with the recognition of the conceptual leap involved in the replacement of the relatively weak claim that the progression through a structure which appears to bear certain

similarities to the putative associational qualities of thought ¹⁹ might benefit learning, to the argument that the imposition of an experts conceptual model on the structure will foster the adoption of his/her particular pre-wired associational complex. Objections to this sleight of mind take a variety of forms. One obvious criticism concerns the conflation of the process of learning with its end product (Charney, 1994). The eagerness with which researchers are prepared to assume that the results of educational studies which extend over an entire course term are applicable to a field of research in which longitudinal studies are still a scarcity, betrays the failure to appreciate the time required for the process of learning (Bartasis & Palumbo, 1995). Instead, learning is portrayed as the more or less immediate transferral of the expert's knowledge structures to the student's mind. This, as noted by Lehtinen et al. (cited by Tergan, 1997a), betrays a static, correspondence perception of knowledge which is anathema to the constructivist ethos.²⁰ It also fails to address issues identified as important for learning, such as the effect of the sequence and timing in which nodes are encountered (Charney, 1994).

I wish to argue that a portrayal of learning as a haphazard, gradual process which is punctuated by what, from the expert's point of view, might be regarded as many wrong turns and dead ends is actually more sympathetic to the constructive conception of education. In other words, the claim is that the divergence between the student pathway and the sequence advised by the expert, as dictated by the model imposed on the hypertext, is not necessarily negative, but may actually benefit learning. Seen in this light, one can consider hypertext as providing the technological analogue of the model of learning characterised by Rumelhart and Norman (1978),²¹ and later embroidered upon by Jonassen (1991,1993). According to their theory, learning can be divided into three different stages, on the basis of how new information to which the student is exposed interacts with his/her background knowledge. The first stage, Accretion, refers to the process in which the incoming information merely builds on the student's pre-existing knowledge structures. When the newly acquired information does not slot comfortably into those structures, however, a process known as Tuning takes place, in which the

¹⁹A perception of the mind which has gained even greater currency with the introduction of connectionist models of learning.

²⁰This objection is only slightly less valid with regards to the fixed modelling of the hypertext on the domain knowledge of a particular user group (Otter & Johnson, 2000)

²¹Which, in turn, owes a heavy debt to the Piagetian theory of schemas

resident knowledge structures are modified in order to accommodate the new information. Finally, in circumstances in which the discrepancy between what is known and what is being learnt is too great, a critical mass of new and discordant information will eventually be attained which requires the development of entirely new knowledge structures. This has been labelled the Restructuring phase.

The facility with which a hypertext allows the user to select links which lead to weakly related or discordant nodes of information nominates it as an implementation of this theory of learning. Seeing learning as essentially a process of tension-reduction, hypertext can be characterised as providing the mechanism by means of which the tension is created. This is consistent with the finding that refutational expository text facilitates student's conceptual change (Qian & Alvermann, 1995). Regarding discrepancy between nodes as a function of both the hypertext structure and the students background knowledge allows one to suggest the contribution of both towards achieving an optimal level of non-coherence between nodes. Hypertext content which is too coherent might not be stimulating enough to foster learning (McNamara, Kintsch, Songer, & Kintsch, 1996), while few students would be sufficiently motivated to attempt to learn material which they find completely unintelligible.²² The latter circumstance might also precipitate the creation of new knowledge structures when the newly acquired knowledge can be integrated with what is already known. This would be hardly be optimal, particularly if those pre-existing structures provide the memory cues for recall of in-coming information (Kintsch, 1998).

The scenario sketched with regard to the interaction of prior and newly acquired knowledge would have one predict that students with little background knowledge of a domain might obtain sufficient stimulation through the requirement to create new knowledge structures in a highly coherent hypertext structure, while more advanced students would be expected to benefit more from a hypertext structure with a degree of non-coherence built into it. This is supported by the fact that while those with little relevant knowledge generally appear to benefit the least from hypertext (cf. Lawless 1996, 1998), in one study in which they actually learnt

²²Moreover, it should be noted that there are a number of constraints which provides an upper limit to the amount of non-coherence between nodes which might still benefit learning. Factors such as the degree of interest shown in the material by the student, their self-efficacy, as well as their IQ, would be predicted as being influential. While the role of self-efficacy will be considered, it is the central role played by the student's motivation which shall receive primary focus in a later section

more, the hypertext was hierarchically structured (Balcytiene, 1999). With reference to the previous discussion, this would also appear to indicate that one of the reasons why explicit graphical browsers might only be adequate for simple, hierarchically organised hypertexts, is that those would be the systems to which introductory students would be exposed. A hypothesis which explicitly theorises the transition from introductory hypertexts to those more suited for advanced students is known as Cognitive Flexibility Theory.

Cognitive Flexibility theory (Spiro, Coulson, Feltovich, & Anderson, 1989; Spiro, Feltovich, Jacobson, & Coulson, 1996) provides an alternative to the blatantly untenable nature of the radical constructivist position towards hypertext, advocating as it does a *laissez faire* attitude towards a student's freedom to move through a hypertext, while in the same breath condemning the students knowledge structures to mould themselves to that of the hypertext model provided. CFT manages to integrate constructivist principles with user qualities, and is commonly cited as one of the only applications of a theoretical framework to the design of non-linear systems for educational purposes (Rouet & Levonen, 1996).

CFT attempts to counter the creation of simplistic knowledge representations, as fostered by introductory courses within ill-structured knowledge domains, through exposing advanced student's to concepts within multiple contextual settings. The revisitation of concepts from many different directions, in what is metaphorically referred to as a "criss-crossed landscape" (Spiro et al., 1996), reduces the likelihood of the development of inappropriate generalisations, in domains such as medicine, where real-world "cases" do not align in perfect unitary correspondence with the domain's conceptual boundaries. While the requirements stipulated by CFT provide the specification for custom-made hypermedia systems, its central tenet, that concepts must be situated within multiple frameworks, is consistent with the access provided by networked hypertext to the same node (and concepts therein) from a number of other nodes, as well as acknowledging the dynamic nature of the knowledge representations formed by the student. Although Spiro and colleagues do not adhere to the notion of the construction of general knowledge representations, and actually regard the formation of such a representation in a negative light, in their focus on the transfer of knowledge to novel situations, they concede the validity of the construction of case - specific truths.

CFT's negative characterisation of the construction of general, global models of

all the content in their instructional systems is a direct contradiction of the traditional belief amongst hypertext researchers that the extent to which hypertexts facilitate the construction of such models is an indicator of their educational efficacy. It could be argued, however, that this perspective might still apply to ill-structured knowledge domains, to the extent that a premium is placed on the retrieval of information, as a person's ability to trace their way from one concept to another has been portrayed as the intellectual equivalent of moving from one part of a text to another (Britton & Gulgoz, 1991). If this is the case though, it represents the point of departure between navigation and learning, where the developers of a system pay tribute to hypertext's history as an information retrieval system (Jonassen, 1993; Tergan, 1997a), through using that system purely for information retrieval purposes. Furthermore, the disorientation which it has been conjectured results from mismatches between representation of a hypertext and the relations between nodes in the hypertext would still be expected to reduce the utility of a global conception of the hypertext for purposes of navigation.

In order to evaluate the exact nature of the contribution of mental models towards hypertext navigation and learning, a thorough examination of the technical notion of what constitutes a mental model needs to be undertaken. This will reveal that the fast and loose way in which hypertext researchers have made use of the concept does not disqualify its application to hypertext learning, but that its more accurate usage suggests possibilities for its general applicability to all forms of hypertext learning.

The Case for Mental Models.

Theories of knowledge representation fall into two primary categories, corresponding to the symbolic or propositional schools of thought (Graesser, Gernsbacher, & Goldman, 1997). Symbolic theories of mind regard knowledge as being represented in some form by objects, a perspective which contrasts with the perception of thought as composed of propositional statements. While there is on-going debate concerning which is the more appropriate for different types of cognition,²³ knowledge gleaned from hypertexts has consistently been described in

²³Within the debate surrounding the form of knowledge representation most applicable to reasoning, the main distinction is between rule-based and mental models theories of reasoning. Refer to Bonati (1994) and O'Brien, Braine and Yang (1994) for arguments claiming inadequacy of the mental model based approach to reasoning

terms of mental models (Panero, 1995). One of the pioneers in the development of the theory of mental models is Johnson-Laird (Johnson-Laird, 1983, 1990, 1996; Johnson-Laird, Schaeken, & Byrne, 1992; Johnson-Laird, Girotto, & Legrenzi, 1998). He cites Kenneth Craik's 1943 formulation of a model as a "relation-structure" which mimics the processes it represents (Johnson-Laird, 1983, p11). Mental models provide a common currency for knowledge representations extracted from text, from perception, and from one's background knowledge (Garnham, 1997).

It is instructive to compare mental models with other mental structures which are often postulated as underlying learning. Schema and Frames, for instance, are stereotypical devices borne of the student's previous experience, which by their very nature as general knowledge structures do not emerge from any particular text. Instead they serve primarily as orientating mechanisms (Otter & Johnson, 2000; Dillon et al., 1993), slotting particular instances into the variables from which they are composed (Rumelhart & Norman, 1978). Pre-existing knowledge structures which are imposed from above, they are their most useful in processing familiar situations.²⁴ Mental models, on the other hand, emerge from a particular learning experience and substitute both structural integrity as well as tokens which are true to the number and character of the objects they represent, for the individual idiosyncratic variable-bound structures of schemas. They are therefore better suited for the acquisition of novel information.

Although mental model theory has primarily been applied to syllogistic reasoning, it is capable of providing an account of text comprehension and spatial reasoning (Johnson-Laird, 1996), thereby suggesting the possibility of capturing the nature of the relationship between the spatial representation of a hypertext and the model of its content. Its promise seems even more manifest when one considers one of the central principles constraining Johnson-Laird's notion of mental model, known as the principle of structural identity. This states that "the structures of mental models are identical to the structures of the states of affairs, whether perceived or conceived, that the models represent." (Johnson-Laird, 1983, p419). There are therefore grounds to claim that models of the hypertext, as realised through maps and graphical browsers, are visual analogues of mental models to the

²⁴Although Brewer (1987) takes issue with the distinction between schemas and mental models as generic and specific, respectively, he concedes that the former are generally pre-compiled. The identification of which schema to apply to a particular situation therefore depends on the recognition of conditions onto which a schema has been mapped in the past

extent that they are true to the structure of the hypertext . The distinction between mental models and hypertext models becomes even more blurred with the description of mental images as corresponding to "views of models" (Johnson-Laird, 1983, p157). The relationship between mental images and models corresponds with that between the hypertext map image and the mental model of the hypertext.

Mental images form only a subset of mental models, however (Johnson-Laird et al., 1998). There is good reason to suspect that the properties of images might mitigate against their ability to facilitate the formation of mental models of the hypertext. Firstly, there is the possibility, mentioned on page 12, of the lack of a match between the structure of a hypertext, the representation of the hypertext ²⁵, and the semantic relations within its content (Dillon et al., 1993). This possibility tends towards certainty if one considers the role of background knowledge, and how the strength of the relationship between the nodes in a hypertext for a particular user varies as a function of the interaction between what they are learning and what they already know. One defining aspect of images, however, is their economy of expression, capturing as they do the information content of a multitude of propositions (Kintsch, 1998). They do this by making the implicit relations between structures explicit, and thereby not allowing for any room for variation in mental structures.

Mental models differ from images in that many of the relationships within them are left in implicit form (Schwamb, 1990; Johnson-Laird et al., 1992). Although this quality has been used to account for lapses in reasoning (Johnson-Laird et al., 1992), it does have the beneficial effect of necessitating the use of inferences. Inferences, and in particular those which are not made automatically, enable, and indeed require, the contribution of a person's individualistic background knowledge in order to make a specific text coherent (Kintsch, 1993). Hypertext maps which provide the student with the freedom to make inferences would therefore go some way towards binding the hypertext's semantic content, its structure, and the knowledge of the student. They would also assist the student in navigating the hypertext in the absence of an explicit map provided, through their construction of a model of the hypertext system. Not only would recall be facilitated by the

²⁵If the hypertext structure is not organised according to the same principles as the content, then the representational image employed is guaranteed to misrepresent either the structure or the semantic content of the hypertext

construction of a mental model with maplike qualities, through enhancing both the navigation of the hypertext system and the mirror of that system as provided by their own knowledge structure (Britton & Gulgoz, 1991), but the integration of the student's background knowledge with knowledge of the hypertext would meet the criteria specified by many as conditions for comprehension and conceptual learning (Rumelhart & Norman, 1978; Kintsch, 1998). This scenario provides the vision of the development of navigational aids which stimulate the formation of internal models, thereby uniting navigation with comprehension.

Fortunately, there are already systems in place which provide users with the leeway to commit their own background knowledge in the quest for structure. For example, the document management system, VIKI (Marshall et al., 1994; Marshall & Shipman, 1997; Shipman et al., 1999, 2001), evolved from the observation in a study (Marshall & Rogers, 1992) that people tended to make far greater use of the strategy of manipulating icons into groups which were identified as part of the same group on the basis of their proximity to one another, than of linking documents to one another by means of explicit links. This allowed users to forego the a priori creation of arbitrary and unwieldy categories in which to place the documents they encountered, and instead create categories to best fit the nature of the documents they found, as they found them. This can be seen as analogous to the distinction between having to relate an expert's model to the content of a hypertext, and creating your own model through exploration of the system.²⁶

Another approach is that taken by Persson (1998), who suggests the use of narrative as a means of easing a users navigation through a hypertext. Narrative would presumably provide the organisational sequence which people could take advantage of in retracing their way through a hypertext system (Landow, 1990). The example Persson provides is that of an animated figure which pops up randomly and tells stories relating to the content of the node the user is currently occupying. Of particular interest, however, is the point he makes about the propensity of individuals to construct narrative from even the most information-poor stimuli, such as geometric figures. Such abstract figures would appear to have particular promise as navigational aids which are tailor-made for

²⁶ While such systems facilitate the building of one own's spatial map of relations, they require the introduction of aids, such as links, to make the meaning of the relations implied by the implicit arrangement of objects accessible to the novice. This constitutes further evidence that loosely structured navigational aids might be better suited for more advanced students.

each individual, as one would imagine that the user, in creating a story, might combine both background knowledge and node information in compensating for the information poverty of the figures themselves.

The argument thus far is of the less is more variety. Both images and propositions provide more information about a specific content area than do mental models (Glasgow & Malton, 1999). To the extent that learning involves the integration of old knowledge with new, those knowledge representations which are less complete and therefore require a greater knowledge investment on the part of the student are predicted to facilitate learning (Shipman & Marshall, 2000). In other words, by conveying less information, mental models *enable* a greater degree of knowledge. This conceptualisation of the benefits of mental models is in agreement with Frank Shipman and Catherine Marshall's (1999b) discussion of the difficulties people have in creating and employing strongly formalised knowledge structures, including those found in hypermedia, one of the reasons for which is that they do not effectively capture the evolving nature of the individual's conceptualisation of the knowledge found in the systems they focus on. There is another reason for suspecting that mental models might be of relevance to hypertext, however. This is borne of the notion that mental models consist of procedures for their own creation (Payne, 1993). One wide - ranging implication of this reconceptualisation is that it calls into question the conventional separation of declarative from procedural knowledge.

In an influential article, Brewer (1987) makes the distinction between causal and mental models. The former differs from the latter in that procedures of use are (implicitly) imbedded within the model (Kleer & Brown, 1983; Young, 1983), with the focus on models of application most commonly relegating causal models to discussions of the use of technological equipment. Determining how to operate a particular item of machinery corresponds to the "running" of its causal model. The fit between the model and reality will be determined by the extent to which the model enables one to make proper use of the equipment.²⁷

²⁷The level of detail and the specific components which form part of an individual's model of a technological item are largely determined by the role it plays. For instance, calculator designers would require a far more detailed understanding of the internal workings of the machine than people who make use of the functions determined by the calculator's design (Young, 1983; Jonassen, 1993). Their respective mental models would accordingly require qualities which facilitate the effective realisation of the differing roles occupied by the model, and would exact a high cognitive price were they to assume qualities of the other.

This conceptualisation of mental models has recently been turned on its head by the claim that it is the operationalisation of procedures for constructing these models, and not so much the running of the models themselves, which can be used to determine their accuracy (Payne, 1993).²⁸ It is specifically the overlap between the memory for propositions expressing the order of the procedural steps used in constructing two models, which determines whether these models describe the same situation. This characterisation of mental models emerged through the reinterpretation of the findings of Mani and Johnson-Laird's (1982) classic study, which demonstrated that people are more likely to confuse passages which are inferred from a study text with the text, when that text describes indeterminate spatial layouts, or layouts which can take more than one form, than when it describes only one conceivable arrangement of objects. This finding has traditionally been explained as resulting from the development of models which incorporate these inferences and therefore can not be used to distinguish between texts with or without them. By demonstrating that this effect disappears when the degree of overlap is taken into account, support is provided to the notion that considerations of mental models need to acknowledge the importance of memory of procedure as well as structure.

The distinction between causal and other mental model types is therefore not that manifest, at least with respect to their degree of implicit procedural content. This is especially clear in the application of mental model theory to temporal sequences (Johnson-Laird, 1996; Glasgow & Malton, 1999). Where they do differ, however, is in the test of their "truth value". While causal models can be tested for accuracy against the real-world consequences of holding such models, mental models of conceptual structures are often only testable against other mental models. If the process of constructing a model becomes the means by which it is compared with other knowledge, then is it any wonder that providing the hypertext user with a pre-fabricated navigational structure does not facilitate its integration with background knowledge? The comparison between the information stored in a person's long term memory and the information s/he acquires through sensory input and reasoning is surely a requirement for the consolidation of the novel information with his/her background knowledge base. Hypertext navigation can therefore truly

²⁸A point echoed by Johnson-Laird (1996) in the context of his discussion of the possibility that mental diagrams might be conducive to the formation of inferences

be seen as the pre-requisite for the development of an accurate mental model of a hypertext.²⁹ Moving through the hypertext allows the student to compare what s/he encounters with his/her mental model of the hypertext's structure, and make adjustments to his/her mental model on the basis of mismatches revealed. The provision of an externally provided and overly specified representation of the hypertext has the risk associated with it of lulling the student into a false sense of confidence regarding their access to information about the hypertext structure, and may prevent the formation of a mental model which is actively re-interpreted in the light of what one knows and what one is in the process of finding out.

A final heroic attempt to dismiss the importance for learning, of the comparison between what one knows and new information one is exposed to, might allude to claims that the validity of a mental model is solely a product of its construction. This argument has been made in connection with diagrams (Lindsay, 1988), where adherence to a fixed set of rules of assembly guarantees that the information the diagram contains is valid. The same can not, however, be said for models of conceptual information, where there are no generally accepted rules of construction, and where the individual is the final arbitrator in deciding whether to accept or reject a model. Mental models have the advantage over cognitive maps of the hypertext structure in that their qualities allow them to function as a framework for the co-construction of representations of knowledge of both hypertext structure, as well as the conceptual relations between knowledge units in the text itself. They thereby provide for the alignment of these two forms of knowledge in the service of facilitating both navigation through the hypertext, and learning of the hypertext's content on a factual and conceptual level. It is this amalgamation of structural hypertext and conceptual knowledge elements which shall be referred to in a subsequent discussion of the integrative aspect of mental models.

To summarise, the investigation of mental model literature has provided us with an independent route to the same conclusion reached when discussing hypertext disorientation. This is that movement through a hypertext is a strong determinant of the cognitive model one has of it. In the case of disorientation, this conclusion presented itself as the logical endpoint to an argument based on the

²⁹This is consistent with the discovery of what is known as the enactment effect, where people who perform actions have greater recall than those who must remember descriptions of those actions (Hammond, 1993).

recognition that the ability of a user to effectively navigate a hypertext depends on the participation of the user in constructing a model of the system. A closer examination of the nature of mental models reveals the dissolution of the distinction between the procedures involved in constructing a mental model and the final form the model takes, and the revelation that the traditional dichotomy between factual and procedural knowledge might not be legitimate. This leads one inexorably to the same conclusion, namely, that navigation of a hypertext is knowledge of its content. Both lines of argument are moreover sympathetic to the possibility that overly explicit expert maps of the hypertext structure do more harm than good (Whalley, 1990). They tend to interfere with the integration of old knowledge with new, as they short-circuit attempts to develop inferences on the part of the user, and thereby reduce the probability that the student acquire anything more than a superficial understanding of the hypertext content. It is these inferences, and the subsequent coherence of the hypertext, which provide the pathways for navigation and information retrieval,³⁰ by enabling the integration of knowledge of the hypertext structure, the material the hypertext contains, and an individual's background knowledge.

While the discussion thus far has helped identify mental models as the form of knowledge most likely to be beneficial to the student making use of a hypertext and provided insight into the constructivist processes underlying the development of that knowledge, the issue of the qualities of the hypertext most likely to promote such forms of learning has not yet been addressed. Given the decision to focus exclusively on the role of the text in hypertext learning, it is necessary to examine how the characteristics of text might encourage the integration of old knowledge with new. It is with this in mind that closer attention will now be paid to what the text comprehension literature might have to offer in this regard.

The impact of text characteristics on learning

The Construction-Integration Hypothesis.

The emphasis in this paper has been on the detrimental effect of the provision of a complete specification of the relationships between hypertext nodes on learning (Whalley, 1990). The processing of information and its subsequent integration with what one knows is not merely a product of the presence of gaps within a hypertext

³⁰In both the actual hypertext and one's model of it.

model, however, but in addition also arises from discontinuities within the text itself.

Hypertexts which possess discontinuities between and within nodes can be characterised as globally and locally non-coherent, respectively. Although the everyday usages of the term "coherence" more closely reflects its Oxford English Dictionary (1989) definition as "Consistency in reasoning, or relating, so that one part of the discourse does not destroy or contradict the rest", in the text comprehension literature it has a much more specific meaning. Global coherence within traditional texts refers to the cohesion provided to the text by titles, section headings and thematically-bound sentences, or more succinctly put, as the "gist of the text". It is this which is most aptly captured by the under-determined nature of mental models. Local coherence, on the other hand, operates on the higher resolution of units of text. For example, the random adjunction of a pair of sentences are most likely to result in a poor degree of local coherence.

Hypertext is unusual in that it frequently does not discriminate on a qualitative basis between text within nodes — providing local coherence — and hyperlink text, or those sections of the hypertext playing a structural role — and hence providing global coherence.³¹ An understanding of the processes involved in text comprehension is therefore vital to fully appreciate the contribution of both textual and structural discontinuities on learning. As one of the most prominent text comprehension theorists, Walter Kintsch's Construction-Integration Hypothesis is helpful in gaining insight into those processes. The following discussion of the CI hypothesis will draw primarily on Kintsch's (1998) comprehensive explication of his theory in the book entitled "Comprehension: a paradigm of cognition".

According to the CI hypothesis, the fundamental processing unit in text is that of the proposition (Kintsch, 1988). Propositions can be decomposed into predicates, arguments, and modifiers, loosely corresponding to verbs, nouns, and adjectives, respectively. An atomic proposition such as *Mary gave Fred the book* has *give* as its predicate³², and *Mary*, *book*, and *Fred* as its arguments. Atomic propositions can be combined to form compound propositions, and written using any of a number of

³¹Although web browsers such as Internet Explorer and Netscape Navigator distinguish hyperlinks from the surrounding text by colouring them, these links still form part of the text. See Weinreich et. al. (2001) for a further discussion of the forms that hyperlinks have taken.

³²Note that when forming propositions, considerations of tense are left out. In addition, the decision as to exactly which elements to include in the propositionalisation of a text is left entirely up to the researcher and his/her research agenda.

notations. For example, the preceding proposition might be combined with *The book was old*, to yield *GIVE[agent:MARY,object:OLD[BOOK],goal[Fred]]*, where *old* is a modifier. When the atomic propositions in a compound proposition are subordinated to a core meaning, they form a complex proposition, which preserves much of the structural integrity of what is commonly known as a sentence.

Although Kintsch pleads lack of concern over the exact nature of a person's knowledge representations, it is the theoretical applicability of the process of propositionalisation to both the text and one's background knowledge, which has given rise to CI theory's recent manifestation as a general model of comprehension. The same process involved in decomposing a text into its constituent propositions (forming what is known as a *text base*), can be performed on one's background general and domain-specific knowledge as well. Where this knowledge takes the form of images (and presumably other spatial representations as well), the description of the imagery forms the text which is converted. A person is regarded as having understood a text, when the representation of the text arrived at merges with their propositionalised knowledge base, to form a *Situation model*. The similarity between the situation and mental model constructs has been commented on (Kintsch, 1988; Payne, 1993).

It is therefore apparent that discussions of hypertext learning in terms of propositions and mental models need not be mutually exclusive. Both are bottom-up generative constructs which are sensitive to specific details in a text. In the case of CI, the consolidation of new knowledge with old first involves the recurrent time-dependent elimination of all of the contextually irrelevant associations activated through the initial exposure to a word, until one interpretation is ultimately converged upon (referred to as the *Construction* phase). It is only after the first 100 ms of exposure to a word that all concepts associated above a certain threshold value are activated in working memory. After approximately 300 to 350 ms the strength of association between the activated concepts and the meaning of the word as determined by its context — commonly accepted as not extending beyond two or three propositions in advance of the word — determines which semantic value is the most contextually appropriate, and hence will rise to consciousness. The generation of thematic inferences which provide coherence to the text and assist in its consolidation with background knowledge, known as the *Integration* phase, only takes place after approximately one second

has passed.

One of the major benefits of the CI hypothesis is that it provides a computationally tractable means of analysing the coherence of texts. Although text coherence is a function of both the text's microstructure, or relations between consecutive propositions, and its macrostructure, or larger units of text in which the proposition is imbedded (Graesser et al., 1997), the emphasis on argument overlap between propositions as an index of text coherence provides for its relatively straightforward calculation. The external validity of defining coherence in this way, and its applicability to learning, has been demonstrated by a study carried out by Britton and Gulgoz (1997), in which texts that were revised in order to overcome gaps in coherence, as identified by a program based on Kintsch's notion of coherence, resulted in greater recall of the material. Larger correlations between the ratings of strength of relation between words provided by the participants in the study and the authors of the original texts, were also obtained.

The comprehension of a text is not merely a linear function of the repetition of words in adjacent sentences, however. In a point analogous to that made previously in discussing the relationship between a student's background knowledge and the appropriateness of an expert-based map of a hypertext, the comprehension of a concept's semantics is unlikely to emerge through the creation of a text which does not stimulate the processing required to combine the information contained in the text base with what the student already knows. In cases in which the student possesses little background knowledge, there is probably little harm done in the provision of such a text, and the improved memory for the text base (relative to the situation model) as fostered by a highly coherent text (Moravcsik & Kintsch, 1993; McNamara et al., 1996) might actually provide the factual basis required for more advanced study. When students do possess grounding in a particular knowledge domain, however, failure to engage with that knowledge in attempting to make sense of a text will merely result in a shallow understanding of how any new material relates to what one knows.

In addition to student motivation (to be discussed), two factors have been identified as impinging on the extent to which people process a text. These are the epistemological beliefs held by the individual, and the degree to which the text must be supplemented by additional information in order to compensate for deficits within it. The misconception amongst students, particularly rife in the physical

sciences (Hofer, 2000), of knowledge as simple and consisting of absolute facts is recognised as particularly detrimental in this regard (Schommer, 1993, 1994). The perception of knowledge as a collection of facts encourages rote learning strategies, with the student tending to interpret his/her memorisation of the facts as comprehension of the text (Schommer, 1993). Calibration of comprehension studies have, moreover, revealed the processing of the text as important even when not taking the learners epistemology into account (Lin & Zabucky, 1998). This is evident in individuals' improved ability to estimate their understanding of a text from which sequences of letters have been excised (Maki et al., 1990).

The relationship between text difficulty and the cognitive processing of the text can be extended to text coherence as well. In a study conducted by McNamara, et. al. (1996) participants with low levels of domain knowledge obtained the best results on both text-based, as well as situational inference and word - sorting measures of knowledge, when the text was coherent, while high-knowledge subjects demonstrated improved performance on the latter measures of conceptual understanding when the text was less coherent. The authors concluded that the benefits resulting from the encouragement within reading tasks of active processing of the material is most likely to manifest itself in changes in the situation model. They were able to elaborate on the pattern revealed in the study conducted by Moravcsik and Kintsch (1993), by supplementing the finding that superficial and deep processing benefited factual and conceptual knowledge gains, respectively, with evidence of the interaction of background knowledge and the different forms of knowledge.

McNamara, et. al. demonstrated that individuals learn less on a conceptual level when a text is either too coherent or not coherent enough, relative to the background knowledge they possess. While the relationship between the effect of coherence of a text and how much factual and conceptual knowledge a person extracts from it can be portrayed as linear when the individual has little background knowledge, the linear relationship becomes curvi-linear with regard to the relationship between conceptual learning and increasing domain knowledge. A person who has little background knowledge would be expected to experience a degradation of performance in a test of conceptual understanding with the introduction of any degree of incoherence, while someone who possesses a higher degree of knowledge within a particular domain may display a deterioration of

his/her performance unless the coherence of the text has been adjusted to match their level of expertise. That this argument is applicable to hypertext as well is evident in the contradiction of the finding by Lawless and Kulikowich (1996) that those hypertext users with the highest levels of domain knowledge were those who appeared the most determined to obtain knowledge from a hypertext dealing with psychology, by the results obtained in a subsequent study they conducted (1998), in which this group of participants revealed themselves to be the most apathetic in their exploration of the system.

The argument that text which is either too easy or too difficult will not be processed sufficiently to ensure its incorporation within an individuals situational model accords well with the prediction made by Johnson-Laird (1983), that knowledge of information contained within a text is likely to remain in propositional form, rather than form part of a mental model, when the text is ambiguous. The practical equivalence of mental and situational models suggests that the conclusion that the text most likely to induce conceptual change in a student is determined on an individualistic basis as a function of both the argument overlap within a text and the domain knowledge possessed by a person, is equally true of both constructs.

Latent Semantic Analysis.

One disadvantage of CI theory is that it has not been possible to develop an algorithm to automatise the propositionalisation of texts. Consequently, in order to determine how strongly words in a passage are related semantically, an essential pre-requisite in determining the ease with which a text is comprehended, it has been necessary to undertake the labour-intensive task of hand-coding the propositions, thus imposing severe limits ³³ on the size of the texts analysed. One means of overcoming this limitation has been through the application of a computational technique known as Latent Semantic Analysis.

In the words of Landauer, Foltz, and Laham "Latent Semantic Analysis (LSA) is a theory and method for extracting and representing the contextual-usage meaning of words by statistical computations applied to a large corpus of text" (1998, p259). Endorsed by Kintsch as a valuable statistical extension to the CI hypothesis (1998), it provides a derived measure of the semantic relatedness of words in the text upon which it is trained. It achieves this through transformations

³³typically under 1000 words (Foltz, 1996)

applied to a matrix consisting of the words in the text along one dimension, and the passages, or "documents" in which they are found on the other.³⁴ The particular transformations in question are based on a class of matrix operations referred to as Singular Value Decomposition. SVD transformations, of which factor analysis is a special case, are based on the well known theorem that any matrix can be decomposed into three constituent matrices (Landauer, Foltz, & Laham, 1998). Two of these matrices represent the values in the rows (by convention, the [T]erms) and columns (the [D]ocuments) of the original,³⁵ respectively, while the third is a diagonal matrix (S) composed of constants known as Singular Values. These convey the scale of the factors in T and D. When these matrices are multiplied with one another, their product is that of the original matrix. Each cell in this reconstructed matrix contains a measure of the co-occurrence of the term and "document" represented by the respective row and column which intersect at that point. Although a paragraph is often selected as a meaningful unit of text for the "document", if the researcher decides upon using sentences, then the matrix constitutes a lookup table for the number of words which overlap between sentences, and hence is compatible with the standard criteria of coherence.

One of the attributes of LSA which most endears it to its advocates, however, is its ability to capture the degree to which words in a text are associated semantically, above and beyond the measure provided by word overlap. This is achieved through removing a pre-specified number of the smallest of the singular values in S, prior to multiplication with the other two matrices. The deviation of the resulting matrix from the original can be interpreted as arising from the extraction of the influence of the spurious and incidental juxtaposition of terms within the original body of text, leaving the values in the matrix cells as more sensitive indicators of general contextual associations. When the appropriate number of dimensions, which are equivalent to the number of singular values used, are retained, the (cosine) values in the cells provide a more accurate measure of the relationship between terms. Moreover, these similarity values are also sensitive to the back-propagated influence of the constellation of terms which form the environment of the words used for comparison. For instance, although the term

³⁴Readers are referred to Deerwester, et.al (1990) for a more technical description of the computations involved in LSA.

³⁵The columns and rows are identical in the case of factor analysis.

"combine-harvester" might never be used in conjunction with "broccoli", LSA might nevertheless indicate that they are related, through their common association with the term "farming". The reduction of the number of dimensions from those used in the original matrix,³⁶ to somewhere in the range of 100 to 300, has resulted in a threefold increase in the accuracy with which LSA can predict the strength of association between words as judged by human subjects (Landauer et al., 1998)

The validity of the values produced by LSA as measures of semantic association has been bolstered by the results of empirical studies. Despite the fact that LSA reduces a body of text to a "bag of words" (Landauer et al., 1998, p21), its neglect of syntax and word order appears to be justified (Landauer, Laham, Rehder, & Schreiner, 1997) by its remarkable success in simulating the performance of students in synonym - matching tasks (Foltz & Wells, 1999; Landauer et al., 1998), and in predicting the results obtained on comprehension and multiple-choice tasks (Foltz, Kintsch, & Landauer, 1998). It has also been successfully incorporated into a number of commercial computer programs, such as *Summary Street*, which evaluates and fosters the ability of students to summarise what they have read (Kintsch, Steinhart, & Stahl, 2000), and the Intelligent Essay Assessor which automatically grades student essays (Knowledge Analysis Technologies, 2002). Unlike hypertext, its information retrieval origins do not seem to have jeopardised its effective application within the domain of education.

LSA's strengths in relation to text coherence are graphically illustrated by Foltz's (1998) re-analysis of material used in the Britton and Gulgoz (1997) and McNamara et.al. (1996) studies. The LSA measurement of the coherence of the sentences in the Britton experiment, averaged across the texts used, was not only strongly correlated with their coherence as indexed by word-overlap, but also with the three comprehension tests administered. In order to test LSA's ability to predict the comprehension of texts whose coherence was less determined by the shared terms, it was used to analyse the study material employed in the McNamara et.al. research. These texts differed from those used in the Britton study in that coherence breaks within them were not remedied in accordance with recommendations made by software based on the CI propositionalisation procedure. The high correlation between the LSA measure of coherence and the subject's performance on the comprehension tasks in the McNamara et. al. study confirms its

³⁶Of either the number of rows or columns, depending on whichever is smaller.

effectiveness in capturing higher-order associations between terms.

LSA's relevance to the argument presented in this paper extends beyond its use as a method of determining the coherence of a text, however. LSA has been characterised as both a method and a model of learning. Its sensitivity to the higher-order proximity of terms in semantic space, as determined by the co-occurrence of the contexts which define them, has been proposed as providing the solution to Plato's paradox, or the question of how children learn so much information without being directly exposed to it (Landauer & Dumais, 1997). This expression of LSA as a model of learning, and not merely a method, is also evident in the association of the material LSA is trained on,³⁷ with the particular knowledge domains from which the texts have been extracted, as well as the amount of reading material an average student up to a particular Grade might have read.³⁸ LSA therefore possesses the potential to dissolve the distinction made earlier between the text and the reader's background knowledge as sources of coherence, by equating the semantic associations between terms used in the text ³⁹ with their relative positions in the student's semantic space. This facilitates the matching of texts of varying degrees of difficulty with groups of children (Wolfe et al., 1998), in a manner which is sensitive to the curvilinear relationship between text coherence and background knowledge.

Establishing coherence through inference generation.

Any discussion of the importance of text coherence can hardly afford to neglect the means by which it is established. As mentioned in the description of the Construction Integration hypothesis, inferences are central to the consolidation of what one knows with what one is learning. This is reflected in the fact that when a person encounters a coherence break in text, the mechanism which is commonly espoused as resolving it is that of the generation of a bridging inference (Kintsch, 1988).

Inferences have been partitioned into those necessary for establishing continuity in the text, and those which are extra-textual, and include elaborative inferences

³⁷see the Latent Semantic Analysis @ CU Boulder web-site for examples of such training data-sets

³⁸One of the main reasons put forward to account for the occasions when LSA fails to deliver on its promises is that the training material was not truly representative of the amount of text to which a person at a particular level of expertise would have been exposed (Wolfe, Kintsch, & Landauer, 1998)

³⁹the average of the cosines between terms in a sentence standing for the coherence of each sentence

(Graesser & Zwaan, 1995). Examples of the former include anaphors, or referential inferences assumed by pronoun-noun relations, and depend on knowledge of grammar and the world in general. Although a well written text contains cues for the decoding of anaphoric references and other coherence - generating inferences, without the need to employ background knowledge, such knowledge becomes crucial when gaps in coherence are introduced. Given the high probability that these gaps will be a quality of hypertexts which employ hyperlinks, this point is particularly pertinent to academic courses which present their material through the medium of hypertext. When the hypertext material takes the form of expository, rather than narrative texts, as is typical of "courseware" hypertexts, it becomes necessary to supplement the general knowledge used in comprehending the narrative (Graesser & Clark, 1985; Graesser & Zwaan, 1995) with domain-specific knowledge.⁴⁰ The role of background domain knowledge in inference generation accounts for the finding in a study conducted by Kintsch and Franzke (1995) that although participants made the same number of inferences regardless of how much they knew about a topic, the inferences made by those with higher levels of background knowledge were significantly more accurate. The strong association between inferences and background knowledge underlines the possibility of being able to explain inferences in terms of underlying knowledge representations, a necessary step if one acknowledges that there is a bi-directional interplay between the two (Graesser & Clark, 1985).

Kintsch (1993, 1998) has devised a classificatory scheme for inferences which might be useful in this regard. Discriminating between inferences on the basis of whether or not they are automatically made, and when they are made, whether information is generated or merely retrieved, he has separated inferences into four different types (reproduced in Table 1). The first class of inferences, retrieval inferences, make direct use of a student's existing knowledge in patching up non-coherent text. In the case of automatic retrieval, the activated knowledge is so strongly associated in memory with the sentences that a person is often not aware of making the inference, giving rise to the inability to distinguish between read and inferred sentences referred to earlier. For example, when presented with the sentences "The man was involved in a car accident. He was taken to the hospital",

⁴⁰Although Singer 1997 provides evidence that inferences in expository text are not limited to this type of knowledge.

Table 1: Kintsch's typology of inferences

Retrieval	
A _a	Bridging inferences, associative elaborations
B _b	Transitive inferences in a familiar domain
Generation	
C _a	Searching for bridging knowledge
D _b	Logical Inferences

^aInferences which are the result of controlled processes.

^b Inferences which are the result of automatic processes.

most people would automatically infer that the man was injured, despite the fact that this was not explicitly mentioned in the text.

The other three types of inference correspond to the distinction made earlier between schemas, mental models, and images. Retrieval inferences which are not made automatically, but instead rely on stereotypical frameworks to provide the missing information, employ generalised background knowledge as encoded in schemas and scripts.⁴¹ Inferences which are characterised by knowledge generation align themselves with the bottom-up situation-specific forms of knowledge representation discussed in the section on mental models. The automatic generation of knowledge is facilitated by the construction of mental images, which convey new knowledge to the student through making implicit relations between conceptual structures explicit. The final form of inference, and the only form which Kintsch believes warrants the label, requires the effortful and deliberate use of strategies of knowledge construction for their realisation. A solid candidate for the type of knowledge representation which would aid this last class of inferences is that of the mental model.

The notion that stimulating the increased processing of a text facilitates comprehension of it is consistent with the effort required in the generation of the last class of inferences. It therefore provides an additional line of reasoning to the

⁴¹Scripts are here defined as schemata for action sequences (Graesser et al., 1997)

conclusion that mental models stand as the most appropriate form of knowledge representation in learning from a hypertext. This is fortunate, as there is no clear-cut way of relating the functional roles according to which inferences are often partitioned in the literature, to Kintsch's categories. For example, Graesser and his colleagues' (Graesser & Clark, 1985; Graesser & Zwaan, 1995; Graesser et al., 1997) constructivist claim that students routinely attempt to create global coherence through making inferences regarding the super-ordinate goals of characters, the causal antecedents of events and states, and global themes in a text, does not commit itself to any position regarding which of these inferences are made automatically. McKoon and Ratcliff's opposing stance, as exemplified by their minimal inference hypothesis (McKoon & Ratcliff, 1992, 1995), expounds that under ordinary reading conditions, in which no special reading goals or instructions are involved, the only inferences which are made are those required for local coherence, or which are based on readily available background knowledge.⁴² These inferences are made automatically, as emphasised by the claim that speeded item recognition tasks are sensitive to such inference generation (McKoon & Ratcliff, 1995).

The minimalist hypothesis runs into difficulties in conditions in which local coherence in a text is not established by the reader, as well as by the recognition that even when it is, the inferences responsible often depend on constructive processes (Garnham, 1992).⁴³ Moreover, according to Kintsch (1998), automatic inferences do not in fact qualify as inferences at all. In addition to doubts as to the automatic inference's ontological status, there is some question of its ecological validity. While Graesser et al.'s (1994) description of the behaviour of readers as supporting the "search after meaning" principle does not ring true under all circumstance, the characterisation of normal reading tasks as goal and strategy-free seems even less plausible. The allusion by McKoon and Ratcliff to what are essentially task differences (McKoon & Ratcliff, 1995), as determinants of whether people adopt special strategies and goals in reading a text, allows them to overlook a potentially less variant source of goal driven behaviour, that of the reader

⁴²Ratcliff and McKoon do not regard global non-coherence by itself as being sufficient to stimulate inferencing. However, see Albrecht and O'Brein (1993) and McNamara et al. (1996) for experimental evidence which indicates that the need to establish global coherence is an important source of inference generation

⁴³Garnham (1992) argues that these processes of construction may be triggered by linguistic signals, which may account for the linguistic marker's facilitation of the recall of text (Sanders & Noordman, 2000)

him/herself. One user characteristic which is particularly pertinent in this regard is that of motivation.

The theory of Goal Orientation specifically addresses the impact of motivation on goal directed behaviour, and is therefore particularly pertinent as far as the discussion of the role of motivation on inference generation is concerned. Goal Orientation theory provides a means of discriminating between those individuals who tend to try and avoid non-coherent texts, by which are included the majority of hypertexts, from those who regard the presence of such texts as an invitation to generate coherence, and hence, comprehension. Before the implications of this theory are fully spelt out, however, it is necessary to clarify exactly why motivation is such an important factor in determining whether students exploit the educational potential inherent within hypertext.

*Hypertext navigation and motivation**How motivated does one have to be to learn?*

Hypertext has been portrayed as a structure which only effectively imparts knowledge through the degree to which it facilitates the convergence of the user's mental model of the structure with his/her model of the conceptual relations embedded in it. The literature on text coherence predicts that there exists an optimum level of text difficulty for each student, as a function of how much the student knows. Given the haphazard nature of the links found in "naturalistic" hypertexts such as those on the WWW, it follows that users are likely to encounter nodes which are from their point of view non-coherently linked, even if solely on a probabilistic basis. It is only through the continued exploration of the system that the student can hope to acquire sufficient information to resolve these coherence gaps. Seen in this light the shallow browsing strategies elicited by the World-Wide-Web and which are regarded as excluding it as an optimal educational medium (Whalley, 1990) ⁴⁴ might actually promote a person's ability to establish coherence in the system.

What is needed is a theory of motivation which has implications for the extent to which people do make use of the explorative facilities offered by hypertext. Not only does this conceptual construct need to explain the persistence of exploratory efforts required for hypertext learning in an educational environment, in which the absence of guidance from an instructor can result in disorientation, but also the continued participation in online distance education courses (Miltiadou, 1999, 2000), in connection with which it has been cited that "attrition rates... tend to be 40%-50% higher than... in traditional face-to-face classrooms" (Miltiadou, 2000, ¶2). Given these demands, and the assumption that the hypertext format automatically induces learning, it is perhaps not surprising just how little attention has been paid to motivation as a factor in effecting how much is learnt from hypertext.

⁴⁴A failure to engage with the material is also evident in the constant speed with which people navigate through web-sites (Thury, 1998).

The theory of Goal Orientation

Goal orientation theory (Dweck, 1986; Licht & Dweck, 1984; Elliot & Dweck, 1988) offers a theoretical framework within which the relationship of hypertext and user characteristics can be explicated. By helping one predict the degree to which disorientation is likely to incite feelings of anxiety within a particular individual, and whether a person will respond to this anxiety through persevering in their efforts, or by giving up, it provides a theoretical basis according to which one can determine which one of these two strategies people adopt in order to minimise the effect of the gaps in coherence identified as the root cause of hypertext disorientation.

The determination of the nature of a person's response to non-coherence in a hypertext is facilitated by the distinction made between performance and learning orientated individuals. While both performance and learning orientation are separate dimensions, studies which focus on this construct generally classify a person as learning or performance orientated when they obtain a score above the median on the performance or learning orientation sub-scales, respectively (Meece & Holt, 1993). Unless otherwise stated, this is the operational definition which shall be employed when references are made in this paper to goal orientation.⁴⁵

People who are performance orientated have been portrayed as being, on a fundamental level, entity theorists, who perceive ability as a fixed individual characteristic (Dweck, Chiu, & Hong, 1995; Braten & Olaussen, 1998). These individuals tend to interpret difficulty in learning material as a failure of ability, which translates to greater anxiety in conditions in which their performance is evaluated, such as those found in academic tests (Roedel, Schraw, & Plake, 1994). Their equation of performance with ability accounts for their pre-occupation with doing well in tests, and the rapidity with which people who fall in this group allow failure at mastering a text to dissuade them from further attempts. It explains their tendency, moreover, to avoid material which they feel threatens to expose them as lacking in ability (Braten & Olaussen, 1998; Garner & Alexander, 1989).

The adoption, among those individuals orientated towards learning, of an incremental theory of ability (Dweck et al., 1995), results in the formation of an

⁴⁵It is the emphasis on academic studies which renders the distinction between learning and performance orientation useful to the present paper, rather than the equivalent distinction between task or ego-orientated individuals, commonly used in analysing the motivation underlying sports achievement (Duda & Nicholls, 1992).

entirely different constellation of characteristics. Their recognition of performance as reflecting on abilities which are within their power to change allows them to regard difficult material as challenging, rather than threatening, and therefore to persevere in their attempts to master such material⁴⁶ (Dweck, 1986; Licht & Dweck, 1984; Elliot & Dweck, 1988). The signs of effort which they attach positive connotations to are seen as an indication of lack of ability amongst performance orientated individuals. Learning orientation has, in addition, been associated with a more sophisticated epistemological stance towards knowledge, which is likely to result in a greater degree of learning on a conceptual level than the rote learning procedures associated with less developed epistemologies (unpublished study, as cited in Hofer & Pintrich, 1997).

The pivotal contrast between performance and learning orientated students with respect to the hypothesis presented thus far, is that whereas the former avoid material they find difficult for fear of possible difficulties experienced reflecting poorly on their abilities (Ames & Archer, 1988), learning orientated individuals tend to actively seek out challenging material, and tend to regard failure as an indication that they did not invest a sufficient amount of effort in the task (Meece, 1994). This observation suggests that the distinction between these two groups of individuals might be particularly relevant in a medium such as hypertext, in which an emphasis is placed on learner control, and in which disorientation can be viewed as that aspect of hypertext which makes learning from it a challenge. Their respective responses towards hypertext disorientation might both be viewed as adaptive, in that whereas performance orientated individuals avoid material which increases the likelihood of their becoming lost in hyperspace, learning orientated students attempt to limit disorientation through rendering that material coherent.

Their adoption of opposing strategies in their approach to hypertext material, however, is likely to result in a concomitant amplification of the differences between the navigational behaviour of performance and learning orientated students. This is a knock-on effect of the self-reinforcing nature of disorientation, as the failure of the former group to persevere in their encounters with difficult material leads to a self-fulfilling cycle of failure in their attempts to master it, thereby further dissuading navigational exploration. The pattern of helplessness which results bears many similarities to the phenomenon of "learned helplessness" observed in people with

⁴⁶Hence, their alternative description as "mastery - orientated"

low perceived self-efficacy (Schunk, 1991), and suggests that a well-developed sense of self-efficacy might counter the effects of performance orientation as an impediment to learning. Although a high sense of self-efficacy has been observed to possess this moderating influence (Braten & Olaussen, 1998; Kaplan, 2001), self-efficacious performance orientated individuals are still disadvantaged relative to their learning orientated peers.

The relevance of goal orientation to hypertext - based learning becomes particularly apparent with consideration of the findings of an experiment carried out by Licht and Dweck (1984). They exposed fifth grade children classified as either helpless or mastery orientated, on the basis of the degree to which they attributed failure to lack of effort, to conditions in which study text was either preceded by passages of text which were readily comprehensible, or in which sentences were constructed using a syntax designed to make them difficult to understand. They found that there was an interaction between learning condition and attributional style in determining whether the students were able to answer multiple choice questions based on the study material, with helpless children being unable to master the material despite being given a number of opportunities to do so. The deliberate lack of continuity between the initial passages and the study text strengthens the argument that it was the effect of induced anxiety on the approach taken to subsequent learning material which was the salient factor.

On the basis of the evidence discussed, it is predicted that a person's goal orientation should be a strong indicator of the degree to which they will exploit hyperlinks in the text. This is due to the fact that the hyperlinks present the greater likelihood of non-coherence than those navigational aids which lead one along the sequence determined by the order in which the material was originally written. In effect, what is being stated is that, given the self-reinforcing nature of the state of helplessness induced by hypertext disorientation amongst performance orientated individuals, their failure to exploit a mode of navigation which theories such as CFT predict are necessary for the acquisition of additional forms of conceptual knowledge, places them in the position of those students with low background knowledge, in which the relationship between their degree of conceptual understanding and text coherence is linear. Although this relationship is positive, the upper ceiling of their conceptual knowledge gain is quickly attained, and the only possibility by means of which they can surpass this limit is through the

exploitation of the hypertext's links. It is this more exploratory learning strategy which it is predicted will be adopted by those students who are learning orientated.

Accordingly, it seems plausible to suggest, on the assumption that text coherence falls on a continuum,⁴⁷ that the acquisition of the additional conceptual knowledge resulting from the strategy of exploratory hypertext navigation will lead to a feedback loop, in which text which was previously inaccessible becomes increasingly more accessible with the generation of the inferences that that additional knowledge facilitates. The ideal would be realised through the generation of a state of global coherence within a hypertext, in which the tension introduced through non-coherence has been reduced to a minimum.⁴⁸ This can be contrasted with the patches of local coherence which might result from the navigational behaviour of performance orientated individuals, which, while locally coherent, on a global level would result in greater overall tension, through the magnification of inconsistencies between the sections from which a hypertext is composed.

The scenario painted above can be expressed in terms of mental models in the following manner. The greater the exposure of students to a variety of means of conceptualising relations between knowledge nodes (which may or may not correspond to nodes within a hypertext), as determined by their motivational patterns, the more abstract their model of the hypertext becomes,⁴⁹ and the more inadequate the representation of that hypertext by means of the rigid relations specified by graphical browsers or site maps. This flexibility opens up the possibility of determining a close match between conceptual relations imbedded within the hypertext content, which should exist in abundance within a rich conceptual structure, and the structural layout of the hypertext. The presence of an appropriate representation of the hypertext provides the student with the leeway to overcome the disparity the structure of a hypertext and its content shall inevitably

⁴⁷Which it can only do with the admission that text coherence is a function of background knowledge.

⁴⁸The logic described here is similar to that which guides the use of the spring model underlying the pathfinder's algorithm, used by Chen (1998a) and Chen and Cherwinzky (1998b) for the creation of Virtual Reality displays of the relationships between nodes. The spring model ensures that the sum tension between the nodes is minimised with regards to the value of the dimension represented in an analog fashion by the distance which separates them. The dimension employed in the Chen studies was that of the strength of the semantic relationship between content within the nodes, as calculated by LSA.

⁴⁹Which is in agreement with Johnson-Laird's (1990) portrayal of the mental models of experts as being more abstract than those of novices.

suffer from, given that the degree to which they correspond is primarily a function of the highly variable amount of background knowledge which students bring to the study session.

This line of reasoning implies that while students who do not make use of hyperlinks will develop mental models which provide a level of detail sufficient to facilitate their discovery of factual information, this model will actually hinder their ability to conceptualise how the information in different parts of the system relate to one another. In lieu of the hypertext user's own set of conceptualisations concerning the nature of the relationships between concepts in the hypertext, s/he adopts the over-specified relationships as conveyed in the representations which are commonly provided of the system. Consequently, not only do they have a very poorly developed conceptual understanding of the material, but they are also particularly prone to the form of disorientation identified as arising from the use of inappropriate representations of hypertext structure. On a more general level, these readers are likely to be especially vulnerable to the discrepancy between hypertext structure, representation, and semantic relations which are characteristic of ill-structured hypertext systems. Their possession of an overly detailed cognitive map will tend to exacerbate the degree to which they suffer from disorientation, primarily due to its interference with the development of global survey knowledge of the system.

The claim by the pioneer of mental model theory that text which is ambiguous or describes scenes which are indeterminate is likely to remain in a propositional form (Johnson-Laird, 1983) resonates with the suggestion that learning only occurs once the integration of incoming and background knowledge has taken place, as it is only then that an adequate situation model can be constructed. The characterisation of route knowledge as possibly textual (Dillon et al., 1993), therefore provides further reason to argue that disorientation is likely to occur when the global model which underlies survey knowledge is absent, as the propositional representation of knowledge within the hypertext would only be expected to facilitate navigation between nodes which are adjacent to one another in the hypertext. However, even this expectation might be optimistic, given the role that it is contended mental models play in recall, and without which the ability to navigate through the system would be seriously jeopardised.

Conclusion and Hypotheses

Broad coverage has been provided of the contribution of user characteristics, qualities of the text, and the nature of both user-generated representations and those provided as part of the system, on learning from a hypertext. It has been suggested that the danger of disorientation which has frequently been identified as curtailing the educational potential of the medium is actually a pre-requisite for the development of the knowledge representations which underlie conceptual understanding. This is implied through the claim that disorientation in hypertext can be viewed as resulting from coherence gaps introduced through the navigational possibilities inherent within hypertext, as well as the discrepancies between the hypertext structure and the conceptual relations embedded within its content. This in turn opens up the possibility that inference generation can counter the feelings of disorientation experienced when navigating a hypertext. The degree of non-coherence present, and consequently, the inferences which are required to bridge them, depend on the amount of background knowledge possessed by the hypertext user.

The definition of background knowledge as consisting of both knowledge the student brought to the study session, as well as that obtained during the process of browsing the hypertext, introduces a dynamic aspect to coherence generation. When considered in the context of the characterisation of learning as the integration of newly acquired information with a person's background knowledge, it suggests that the student will only learn effectively from a hypertext when their representation of it is sufficiently abstract as not to interfere with shifts in how a student conceptualises the material. This type of representation is consistent with the description of a mental model, a knowledge form which has been characterised as a rich source of implicit information, and which, additionally, can be seen as containing the rules for its own construction.

Seen in this light, mental model theory promises to be a particularly fertile source of ideas for hypertext research, as it makes the connection between structure and procedure explicit. It accordingly supports the theory which has been developed, namely that the comprehension of the conceptual relations within a hypertext facilitates its navigation, and *vica versa*. The abundance of potential relations between conceptual knowledge units provides a rich smorgasbord of possibilities for the integration of that information within the constantly changing

landscape of what one already knows, and in addition, facilitates one's navigation of complex material through providing the possibility of a closer fit between the hypertext structure and conceptual relations within it. The over-determined nature of the relationships between concepts in an ill-structured domain corresponds to the infinity of possibilities which can be contained by mental models, due to their under - specified character (Garnham, 1992).

Given the crucial role attributed to the generation of inferences in learning, as well as the relationship between exposure to conditions which are likely to induce disorientation and the possibility of making these inferences, it is predicted that an important factor which will determine how much is learnt from a hypertext is the extent to which people are prepared to navigate through non-coherent text in the first place. Goal orientation theory has been identified as useful in providing a theoretical basis according to which one can determine who are likely to hazard the risks of disorientation. Learning orientated individuals are more likely to interpret material which induces feelings of disorientation as challenging, and therefore persevere in their efforts to establish coherence. It is predicted that they will also make more extensive use of hyperlinks, for the same reason. Performance orientated students, on the other hand, will become increasingly wary of behaving in a manner which they feel will reflect poorly on their ability. The failure of this latter group to engage with the material will be exacerbated by their reliance on overly-specified graphical browsers or site maps, which while serving as a crutch in attempting to cope with feelings of disorientation, is of such a nature as to exaggerate them. The differences between these two groups in terms of navigational behaviour and disorientation will persist, even when variations in self-efficacy are controlled for.


The theoretical framework which has been developed is represented diagrammatically in Figure 1. Goal orientation is depicted as the primary determining factor in the way in which people navigate a hypertext. Whether or not an individual can be categorised as performance or learning orientated determines how actively they explore the system, even when their degree of background knowledge of the knowledge domain contained therein, and their sense of control in using the hypertext features, are controlled for.⁵⁰ This will in turn influence the extent to which they are familiar with the structure of the hypertext, and the

⁵⁰Hence the labelling of the background knowledge and self - efficacy variables as covariates of goal orientation.

conceptual relationships they form between the knowledge units contained within each of the nodes. These have been characterised as the constituent elements of integrative mental models of the hypertext system, and it is the nature of these models which feed back into the navigational experience. Mental models are theorised as influencing hypertext navigation directly, as well as via the effect of disorientation on subsequent exploration of the system. While this latter effect will be one of the curtailment of exploratory behaviour amongst performance orientated individuals, learning orientated participants are expected to respond to disorientation through pursuing further exploration of the system. Finally, the characterisation of mental models as integrative suggests that those participants with more developed mental models will experience greater degrees of conceptual and factual knowledge change as a result of their use of the hypertext system, but that this advantage will only be evident in hypertexts which employ hyperlinks.

The model outlined above is rephrased below in terms of the major themes into which the literature review has been partitioned.

Figure 1. Conceptual model of relationships between goal orientation, navigation, mental models and knowledge acquisition.



Navigation: Goal orientation is expected to influence the number of hyperlinks made use of. The paths followed by learning orientated participants is predicted as being characterised by a greater reliance on the use of hyperlinks than those taken by performance orientated individuals. The effects of goal orientation on hypertext navigation will be exerted both directly, as well as through the mediation of disorientation. Performance orientated individuals

are predicted to respond to disorientation through restricting their passage through the hypertext to structural links, while learning orientated participants will regard any adversity they may encounter in the form of disorientation as a representing a challenge, and shall accordingly use the challenging aspects of the hypertext, namely, the hyperlinks, to an even greater extent. Various metrics will be applied in quantifying the degree to which goal orientation effects hyperlink usage, providing us with a profile of the differences that do occur.

Mental Models: The more well - developed a mental model, the more effective it should be in amalgamating the structural aspects of the hypertext system with the conceptual aspects of the material contained within it. It is predicted that learning oriented individuals will demonstrate significantly more well - developed mental models than performance oriented subjects.

The extent to which this model facilitates the process of conceptual integration, as well as the degree to which it results from it, should be evident in a superior performance on tests of both structural and conceptual knowledge. The integrative nature of mental models will be most apparent when comparing the conceptual and structural knowledge of students using a hypertext with hyperlinks, and the performance of those who explore a hypertext without these navigational devices, as it is the presence of the non-coherence introduced by the hyperlinks which fosters the consolidation of these knowledge forms. Latent semantic analysis will be used in providing a measure of the semantic relatedness of concepts found in the hypertext, independently of their actual locations within it. The structural accuracy of participants' mental models will emerge in their performance on a fact-finding task.

Disorientation: Participants who possess a well developed mental model, as defined above, would be expected to demonstrate fewer signs of disorientation. It is predicted that the converse will be true for students who fare poorly on both the structural and conceptual mental model measures.

Learning: The ability of students to learn from a text has been described as a function of the extent to which they are able to integrate what they know with

what they are in the process of learning. The degree to which they are capable of achieving this once again depends on the effort made in establishing coherence in the text, in combination with the amount of knowledge they already have at their disposal. Given the dynamic process of growth which it has been postulated is characteristic of the levels of background knowledge possessed by those people who continually expose themselves to new information, it is argued that goal orientation is ultimately a more important factor than a person's initial levels of background knowledge in determining how much a person will learn. The ability to recall factual information from an ill-structured hypertext will depend on the mental model which has been developed of it, as it has been argued that recall is determined by the extent to which an adequate situation model of the text has been constructed. Given that a situation/mental model is composed of both a structural and conceptual knowledge component, these forms of knowledge are expected to be positively correlated with one another, and with the disorientation which results from a poorly realized mental model.

The reader will have noted that there has been no mention of the examination of the effects of different types of externally- provided representations of the hypertext system on the student's mental model development. While this may seem a glaring oversight, given the emphasis in the literature review on the influence of implicit versus explicit spatial representations of hypertext systems, it was decided, on the basis of time and resource constraints, to test a circumscribed version of the theoretical framework as presented. By keeping the structure of the hypertext systems used across comparison groups constant, and omitting the impact of the influence of an external representation through not providing one, it has been possible to focus exclusively on the influence of the employment of hyperlinks on mental model formation. The exploitation of hyperlinks forms the mediating connection between goal orientation and the mental representation of the hypertext as a product of conceptual-structural integration. This manner of implementing the research hypotheses thus effectively tests the most important components of the theory discussed.

Method

It was decided to partition the experimental procedure into two parts. This allowed the researcher to be expansive in the decision as to which scales and questionnaires to use. It also ensured that a degree of redundancy was built into the data, which assisted in cross-validating the questionnaires used, and in providing greater insight into the precise nature of the constructs being examined. The latter point is particularly pertinent given the exploratory nature of much of this research study. The wealth of information provided by all of the scales used, in combination with the non-intrusive navigational data collected during the experimental sessions, has led to the adoption of a liberal approach in the selection of a variety of statistical techniques for data analysis.

Participants

Ninety two students participated in the experiment for course credit. They were enrolled for a second year Psychology course at the University of Cape Town, entitled "Language, Learning, and Cognition". The research programme was officially registered as the tutorial module for the course, as well as forming a substantial part of its "Learning" component. Nineteen of the students were male, and 73 female.

Materials

The hypertext content consisted of an abridged version of the web - based material from the 1999 season of the Images of Africa course, coordinated by the Centre for African Studies at the University of Cape Town. A course offered to Humanities and Social Science first year postgraduate students, it comprises different perspectives of the representation of Africa, its people and its history, as informed by archaeological, historical, and anthropological perspectives. It focuses in particular on how the self-serving nature of these representations led to the colonial subjugation of the African people.

Two versions of the hypertext were created, each consisting of 63 nodes. While both hypertexts contained identical structural links, the distinguishing presence of hyperlinks within one of the hypertexts employed allowed for testing of the prediction that there will be an interaction effect observed between hypertext coherence, which will decrease as a result of the introduction of hyperlinks, and goal

orientation. The two versions of the hypertext which were provided will be referred to as the Hierarchical⁵¹ and Network conditions, due to the distinguishing presence of hyperlinks lending greater inter-connectivity to the hypertext employed in the latter condition.

Each version of the hypertext was organised in a hierarchical structure, with the top level of structural links leading from an introductory "home" page to one of four pages corresponding to the major modules into which the course had been divided.⁵² The modules were entitled as follows: (1) Africa as Paradise, (2) Race, (3) God and (4) African cities. These sections were in turn subdivided into a number of additional subtopics, such as that of the "Origins of Racism" node within the Race module (Appendix B.1). Each subtopic was comprised of a series of sequential nodes, which ranged from 2 to 8 nodes in length. In addition to the navigational aids offered by the web browser used by the students to gain access to the hypertext, the hypertext included structural links at the top of each page to the home page and each of the module base nodes, as well as buttons at the bottom of each page which allowed the student to move forward or backward in the sequence in which s/he found him/herself. Each node consisted on average of two paragraphs of text.

The additional hyperlinks within the Network condition were embedded within the text, allowing students to bypass the sequential mode of navigation by jumping to nodes located within other regions of the hypertext (Appendix B.2). The section of the destination node deemed most relevant to the text embedded within the hyperlink was rendered bold on following that link (Appendix B.3). If a participant subsequently accessed the same destination page via a different hyperlink, only the appropriate text would be emphasised. This departure from standard hypertext design protocol was regarded as justified by the degree to which it facilitated the location of the relevant text by the student.⁵³ Although it might be argued that a

⁵¹With an uppercase H, so as to distinguish it from the discussion of hierarchically structured hypertexts

⁵²Refer to the figures in Appendix B.2 for a diagrammatic representation of the hypertext structure, and the respective navigational possibilities offered by each version.

⁵³The highlighting of the target text was deemed necessary given the constraints introduced by the original material upon which the hypertexts were based. The effort to conserve the meaning of the content made it difficult to introduce substantial changes to the phrasing and structure of the content so as to render it more compatible with the hyperlinks created. In addition, the standard means of bringing the relevant section of text to the students attention by moving it to the top of the screen was not effective in this case, as the fact that many of the pages did not extend beyond the end of the computer monitor meant that the inclusion of identification markers within the hyperlink

major cause of disorientation is exactly the disconcerting nature of the jump from a link to text which it is not immediately possible to establish as relevant, it is proposed that it is only in the most poorly designed hypertext that the difficulty in identifying those segments of text which are linked is a greater source of non-coherence than the lack of fit between the text units themselves.

All of the questionnaires and both hypertext versions of the study material were delivered from a Pentium I Linux (Mandrake ver. 7.0.) server, running the Apache web-server software. These materials were presented to the students on Microsoft Windows 98, Intel Pentium Pro networked workstations, with 14 inch colour monitors. The problems posed in keeping a log of the web-pages revisited by the participants, due to the default caching behaviour of the Internet Explorer browser, was overcome by means of the use of the "Probe" script, down-loaded off the World Wide Web (Woude, 2002). Each of the questionnaires was produced with the aid of the Perseus Survey Solutions software (Perseus Development Corporation, 2000). The actual experiment took place in two separate venues. The first session occurred in the left wing of the South-one computer laboratory, which contained twenty computers. The second session took place in a room in the Humanities computer lab, which was slightly larger with 30 workstations. Both laboratories are located on the Upper campus of the University of Cape Town.

The hypertext was based on the material used in the Images of Africa course for a number of reasons. Firstly, the division of the course material into separate components, each of which was contributed to by different lecturers, meant that it was accurate to regard the hypertext content as ill-structured. This ensured that there was a multitude of different ways in which the content of the hypertext could be mapped to its structure. The modular nature of the course, moreover, led to variations in the difficulty of the text across sections. For example, the Paradise component was written in a more accessible fashion, and contained material, in the opinion of this researcher, which was easier to understand on a conceptual level than the other sections. This increased the chances that at least certain parts of the material could be understood by all of the second year students participating in this study, while providing more difficult terrain as a proper test of student motivation. In addition, the observation made within the report of the initial evaluation conducted of the course that "Images of Africa seemed to lack a coherent course

target tags did not result in this movement.

design, aim, and purpose" (SAIDE, 1999) suggests that the combined material possessed a degree of non-coherence which suited it well to this research. The fact that the material used was extracted from an actual university program also helps in contributing towards the external validity of the results obtained.

Measures

User characteristics

Goal Orientation.

The Goal Inventory scale (Roedel et al., 1994), a 5 point Likert scale consisting of 25 items, will provide a measure of goal orientation (Appendix C.2). Five items on the scale are sensitive to performance orientation ($\alpha = .75$), and 12 items to learning orientation ($\alpha = .80$), with the rest serving as filler items. Due to the centrality of the concept of goal orientation to this study, the Test anxiety (five items, $\alpha = .80$), Regulation of effort (four items, $\alpha = .69$), and Intrinsic (four items, $\alpha = .74$) and Extrinsic (four items, $\alpha = .62$) motivation sub-scales (Pintrich, Smith, Garcia, & McKeachie, 1991; Pintrich, Smith, Garcia, & McKeachie, 1993) of the Motivated Strategies Learning Questionnaire (MSLQ) were included in the research design. Given the theoretical relevance to goal orientation of the dimensions rated using these sub-scales, the total of 17, 7 point likert items from which they are composed will prove useful in determining the reliability of the goal orientation measurements. Roedel et.al. (1994) verified the existence of a significant negative and positive correlation of performance and learning orientation, respectively, with the results of the Samson's (1984) "Reactions to Test" scale, a measure of test anxiety. Learning and performance orientated participants were, in addition, originally conceived as more intrinsically and extrinsically motivated, respectively (Dweck, 1986). The former group are also expected to be characterised by their willingness to expend more effort in tasks of comprehension than the latter.

Self-efficacy.

The subject's sense of computer self - efficacy was measured by means of a minimally modified version of the 5 point computer-user self - efficacy Likert-type scale, developed by Murphy, Coover, and Owen (1989). The questions dealing with mainframe computer use were omitted, due to the present day dominance of the personal computer, thus leaving a scale comprised of 21 items. To this were

appended 8 questions of the same format, taken from Eastin and LaRoses' Internet Self-Efficacy scale (Eastin & LaRose, 2000).

Background knowledge.

Background knowledge was determined through the responses given to a 10 item multiple choice questionnaire (Appendix C.1). The items were partitioned according to specificity, such that the first five tested for more general knowledge of African colonial history, while the latter half consisted of questions which could be answered by information found in the hypertext. The participants could select the correct answer to each question from 4 candidate responses, after which they had to indicate on a 5 point Likert-type scale the extent to which they were confident that they had answered the question correctly. Due to the fact that the hypertext covered material taken from a large number of knowledge domains, it was decided that the ten questions asked might not provide a true indication of the amount of background knowledge possessed by the participants, especially given that a quarter of those questions could be answered correctly by chance alone. The possibility that the confidence with which participants answered certain of the items might be determined as accurately predicting the correctness of the response, argued for their inclusion. Their potential in facilitating the identification of those items which are most likely to reflect background knowledge would assist in boosting the reliability of the scale.

Navigation

The choice of methods used in determining the navigational pathway followed by individuals through a hypertext can only be properly appreciated when one is aware of the methodologies used in the past. Attempts to gain insight into how people use a hypertext system through analysing the navigational pathway they follow is a well - established approach in hypertext research (Edwards & Hardman, 1989; Lawless & Kulikowich, 1996, 1998; Otter & Johnson, 2000; Bouras, Konidaris, Konidari, & Sevasti, 2001). It is most commonly accomplished through filtering of the web server log file, which keeps a record of the nodes visited, the identity of the computer the participant used in accessing those nodes [through listing its Internet Protocol (IP) number], as well as the duration spent in each of the nodes.

The log file is invaluable in aiding attempts to associate groups of hypertext users with particular patterns identified within the navigational data. Lawless and

Kulikowich (1996, 1998) for instance, were able to distinguish between three groups, which they labelled "knowledge seekers", "feature explorers", and "apathetic hypertext users". They were able to do so on the basis of differences observed in the number of times users visited thematically organised cards, the time spent in the nodes, the number of times they deviated from those themes, and the number of times they used the navigational resources supplied. The analytical technique they employed was cluster analysis, which, by setting limits on the variation which can be tolerated between with-in group members, assists in delineating groups from one another.

The hypothesis that learning orientated individuals will display different navigational profiles from those who are performance orientated can be tested through the use of a technique similar to that employed by Lawless and Kulikowich. In this instance, though, the analysis will be of a more exploratory nature, with the variables which have been identified as potentially relevant with respect to goal orientation being ⁵⁴

- Average time spent in individual nodes
- Number of distinct nodes visited
- Total number of nodes revisited
- Total number of times the base node is returned to.

While findings concerning the way in which the variables listed above relate to the scores obtained on the learning and performance orientation subscales could be useful in assisting in the interpretation of these navigational patterns, their a theoretical nature might hamper attempts to determine their implications for hypertext systems based on other design principles. Claims concerning the relationship between goal orientation and hypertext learning would only have comparative validity were they based on quantifiable characteristics of a hypertext which have been identified as affecting how well suited it is for purposes of

⁵⁴ Although it was originally felt that the average coherence between the text in the links followed and the text to which they lead should be included, the proposal that the differences between the navigational patterns of the two orientational groups would be amplified suggests that any of the links embedded in the text would quickly lead to the cessation of link use amongst performance orientated participants. While this would tend to limit the value of comparing link coherence, its effects would be reflected in the number of hyperlinks used, and the total number of nodes which are visited


education.⁵⁵

Fortunately, the derivation of techniques of analysis originally applied in the arena of Social Network theory shows promise in relating both the structural features of a hypertext with its potential as an educational medium (Botafogo et al., 1992; Astleitner & Leutner, 1996), as well as the "structure" imposed by the navigational pathway followed by the user with how much that potential is realised (McEneaney, 1999, 2001). This is apparent when one considers that the matrix format which has been used in representing the relationships between agents in any closed social network, such as those between the employees of a corporation, can also be used in the representation of the interconnected nodes of a hypertext. For example, the ability to move from one particular node in a hypertext to all others (which might be discussed in Social Network theory in terms of the influence of one employee over his colleagues) can be read off the horizontal axis or row in which the cell representing the node is found in the matrix, while the ability to reach that node from the others (or the extent to which the employee's colleagues can influence him) is evident from the values found in the y axis, or column in which the cell is found. This bi-directional relationship can be represented at its simplest in binary format, where the digit "1" represents the presence of a connection between two nodes, while "0" stands for its absence. The equivalence of these two forms of representation is evident in the portrayal in Figure 2 of the same hypertext in both diagrammatic and matrix format.

The matrix representation of hypertext structure facilitates the calculation of the average ease with which a student can move from one node in the hypertext to another. This is achieved through the division of the total number of connected nodes (summing the marginal row or column totals) by the maximum possible number of connections, thereby producing a statistic referred to as density. The freedom which the hypertext user has in moving from one node to any other node in the system can thus be determined as greater the closer the density is to 1. When the binary notation is used instead to represent the connections actually made use

⁵⁵Two qualities which have potential in this regard are the degree to which nodes within the hypertext are inter-linked, as well as the possibilities available to the participants to deviate from a sequential pattern of movement within the hypertext. Known respectively as compactness and stratum (Botafogo, Rivlin, & Schneiderman, 1992), the importance of the former characteristic should be self-evident, as it is directly related to the ubiquity of hyperlinks within a system, while stratum reflects the degree to which students are exposed to less coherent links.

Figure 2. Example of a graph and matrix representation of simple network hypertext (grey arrows represent the hyperlinks)



of by participants, then density provides a concise means of summarising the extent to which they have exploited the navigational routes available to them.

Another advantage of the matrix format is that it can be manipulated algebraically, yielding products such as that of the distance matrix, in which the value found in a cell indicates the number of steps required to reach a particular node from all others, or vica versa (Botafogo et al., 1992; Astleitner & Leutner, 1996; McEneaney, 1999). This technique of characterising the structural aspects of a hypertext have been extended by McEneaney (1999, 2001), in a manner akin to the density conversion discussed. By substituting the number of times a user makes the passage between two nodes for the presence of said passage, he has managed to derive a path matrix from the distance matrix. This provides a means of capturing user navigation in a form which is amenable to analysis in terms of *path compactness* and *path stratum*, both with potential educational salience as analogues of structural metrics (cf. footnote 55).

Path compactness is based on a formula which measures the complexity of a user's path, as an index of the connectivity it is characterised by, while path stratum provides a truer reflection of the linearity of the path followed. Path compactness compares the connectivity of the path matrix as a ratio of the maximum connectivity possible in the network. While it is similar to the density metric, the path compactness formula includes a measure of the total number of

paths possible given the number of nodes that the user has accessed.⁵⁶ Concerns surrounding the validity of interpreting this metric as a measure of the navigational possibilities which are actually available to the participant led to the inclusion of the density matrix, which compares the number of links exploited relative to the number which were available to the participant to exploit.

Path stratum, on the other hand, provides a normalised measure of the ratio between vertical and horizontal movement within a hypertext. When these dimensions are interpreted in terms of their meaning with relation to the path distance matrix, they indicate the ease with which someone retracing the steps taken by a hypertext user would be able to return to the node from which that person started. Within the context of this study, this measure reflects the degree to which a particular individual restricts his/her movement to the sequence dictated by the static navigational aids within the hypertext. Conversely, in the Network condition, it would indicate the extent to which a person has taken advantage of the hyperlinks, as consistent use of hyperlinks would reduce the number of nodes one is required to traverse in returning to a position one started from.

These metrics have been developed as quantitative complements to McEneaney's more intuitively accessible visual presentation of the user path, and help to imbue conclusions about differences between navigational pathways with greater comparative validity.

Mental Models

Structural aspects.

The requirement that subjects locate facts within the hypertext allows one to supplement navigational qualities observed during the reading of a hypertext with a measure of the degree to which the path followed was indicative of a poor conception of the structure of the hypertext. Participants with poorly formed mental models are likely to experience difficulty in locating factual knowledge within the hypertext.

The ability with which participants can locate particular factual information has been calculated as the ratio of nodes visited relative to the number it was necessary to pass through in order to attain the information (Dias & Sousa, 1997; Otter & Johnson, 2000; Boechler, 2001), and the ratio of the total number of nodes

⁵⁶Although refer to the final section of this paper for a discussion as to whether it is more appropriate to use the total number of nodes in the hypertext for normalisation purposes.

accessed to the number of different nodes. Smith's measure of lostness (Smith, 1996), for instance, encapsulates these metrics in the following formula:

$$L = (N/S - 1)^2 + (R/N - 1)^2$$

where

N equals the number of different nodes visited during the search,

S the total number of nodes visited, and

R the minimum number of nodes it was necessary to visit in completing the task

The value of L is predicted to diverge from zero in proportion to the degree of disorientation experienced. There are, however, reasons to suspect that this formula is not entirely appropriate as a measure of the hypertext user's structural knowledge of the system, particularly with regards to goal orientation.

Although Smith (1996) uses N/S and R/N as measures of both lostness and efficiency, she recognises that while the latter formula might be particularly susceptible to the effects of knowledge of where specific information is located in the system, and hence might legitimately be regarded as an indicator of structural knowledge, the ratio of total versus distinct nodes visited could conceivably be more influenced by the confidence of the user. Moreover, recognition should be given of the fact that the formula embodying this ratio might more faithfully reflect the efficacy of the mental model at integrating hypertext structure and the embedded conceptual relations of its content, rather than structural knowledge alone. Were the student not able to achieve this fusion of structure and conceptual content, it would be expected to result in his/her inability to remember the content which they had come across in the information search, in a manner reminiscent of Kintch's (1998) portrayal of the interaction of a person's situation model with recall,⁵⁷ and lead to their return to paths which had already been investigated. In addition, the definition of a learning - orientated individual as one who is less anxious and exerts more effort in his/her approach to a task suggests that these characteristics might result in over-extended paths to the factual targets, and that this would minimise

⁵⁷This suggestion is also consistent with Dillon's (1993) speculation on evidence that the memory of content and the location of that content might be linked mnemonically. This supports Kintsch's proposal if it is conceded that memory of location would require survey knowledge of the hypertext.

evidence of an improved structural mental model. This is particularly the case if Smith's formula is exclusively applied to those search attempts which were successful.

In light of these considerations, the decision was made to adjust this formula so as to more effectively measure those qualities of interest in this study. Structural knowledge of the hypertext will be gauged through the use of the following formula:

$$SK = (R/N \times R_{tot}/N_{tot}) \times C/Fn$$

where

- SK = structural knowledge of the hypertext system,
- R = the sum of the minimum number of nodes it was necessary to visit to complete those tasks on which success was obtained,
- N = the number of distinct nodes visited during the search for those tasks which were completed successfully,
- R_{tot} = the sum of the minimum number of nodes it was necessary to visit to complete all of the tasks,
- N_{tot} = the sum of the non-replacing distinct nodes in the path of each student, regardless of whether the fact was found or not,
- C = the number of questions correctly responded to, and
- Fn = the total number of facts required

SK will be used as the measure of structural knowledge of the hypertext, with the value of SK falling between 0 and 1. The closer the value to 1, the greater the participant's knowledge of the hypertext structure. The ratio of distinct nodes versus required nodes for those facts which were successfully obtained was multiplied by these variables for the entire fact-finding session, in order to compensate for the possibility that those successful attempts at answering questions might not accurately reflect the state of structural knowledge possessed by the students. This would be predicted as being especially probable in circumstances in which there is a large differences in the degree to which particular facts are associated with specific regions of the hypertext. A dissociation in performance would be expected to emerge amongst those students who only possess locally

specific mental models of the hypertext, and would result in unrealistically high values of SK if this weighing function were not employed.

Despite the weighing of R/N by the performance throughout the fact-finding exercise, the possibility still exists that the structural knowledge statistic arrived at would be distorted. This would be anticipated to be most probable amongst students who are either highly anxious or who have very low levels of motivation, and only attempt to find the information for those questions which they feel confident they know the answer. It would seem likely that they would answer a few of the questions in a very efficient manner, and not to make any serious efforts to answer the others. These students would be distinguished from those who take more effort to answer questions they are unsure of by both the high SK value they obtain, as well as the low number of questions they answer correctly. The multiplication of the weighed SK value by a number which is proportional to the questions which were answered successfully is one way of redressing the imbalance introduced by what is essentially the distinction between performance and learning orientation.⁵⁸

Conceptual aspects.

The efficacy of the study session in effecting changes in the subject's conceptual model of how words are related to one another was measured by means of a word-sorting task. Students were instructed to sort the words into different groups, on the basis of which words they felt "belonged together". The relatedness of the words in the clusters formed was arrived at through the construction of a matrix containing the values for the strength of association, as derived from Latent Semantic Analysis, between all 40 words used. This open-ended method was used in preference to the more conventional a priori recognition of word similarity by means of the assortment of words into a number of fixed categories, as it was felt this method would be more sensitive to variations in how people group concepts. Given the abstract nature of many of the concepts in the word list, and the constructionist's recognition of the equivalent validity of a number of different knowledge structures, it was felt the open-ended methodology would be more appropriate.

Two separate word-similarity matrices were produced, one for each of two

⁵⁸The scores of the high SK answer group are a manifestation of the partially adaptive qualities of the learning strategies adopted by the performance orientated group, as hinted at in the literature review.

semantic spaces. The first matrix was derived through the use of the service offered by the Latent Semantic Analysis CU Boulder web-site (Laham, 2002). Their provision of the *tasaAll* semantic space was taken advantage of in providing one measure of word relatedness. The corpus of text which this space had been trained on was estimated as equivalent to the amount of reading material a person initiating their first year of education at a tertiary institution has been exposed to.⁵⁹ Material included for training was determined on the basis of readability levels. The number of documents and terms included in the corpus was 37 651 and 92 409, respectively, with 419 factors being used in computing the similarity scores. The entire corpus consists of approximately 11 million words (Landauer et al., 1998).

The second, or "custom", semantic space⁶⁰ was based on all of the text used in the reading task, the majority of articles which were referenced on the Images of African website, as well as a number of web resources.⁶¹ SVD was performed on the material using the "General Text Parser" software for Linux (Berry & Dumais, 2001), developed by S. Howard, H. Tang, M. Berry and D. Martin. This yielded three matrices of 10349 documents (paragraphs), by 3012 words (excluding a list of 413 common words), with 289 singular values,⁶² which were recombined by means of a script⁶³ to deliver the cosine values for the similarity of the 40 words of interest. The cosine values ranged between 0 and 1, with the higher values indicating a greater degree of relatedness between the words. Three "experts" were approached to classify the words into groups, and their performance was used to establish the validity of the LSA similarity readings.

⁵⁹Refer to the CU@ Boulder web-site for additional information regarding how this data-set was compiled. Landauer and Dumais (1997), in addition, provide a general explication of the procedure involved in estimations of this nature.

⁶⁰Hereafter referred to as the CustomLSA data corpus.

⁶¹Refer to Appendix D for the partitioning of the documents used in constructing the *tasaAll* space according to subject area, as well as for a partial listing of the sources used for the second semantic space. Interested parties may contact the author for the full listing.

⁶²Great emphasis is placed in the literature on the importance of the number of factors or dimensions chosen when reconstructing the matrix. It is generally agreed that a number between 100 and 300 is effective (Foltz & Wells, 1999), with larger numbers presumably being more appropriate for bodies of text which encompass a broader range of knowledge domains, and hence possess words with more varied meanings. Although I ran an SVD procedure with both 128 and 289 factors, I ultimately settled on the larger number in consideration of the varied topics covered in the Images of Africa course.

⁶³Unless otherwise indicated, each of the operations described hereafter was performed by means of scripts written by the author in the Perl (version 5.005_03) programming language.

Disorientation

The establishment of hypertext disorientation is problematic, since while it has traditionally been evaluated in terms of the number of times the subjects jump back to the start node, revisit the same nodes, and backtrack to previous nodes along the route they have taken, the possibility has also been conceded that these behaviours may actually constitute beneficial strategies as far as learning is concerned (Rouet & Levonen, 1996). For instance, the revisiting of nodes in a hypertext, albeit within different contexts, is a pivotal part of a successful learning strategy in the eyes of CFT theorists.⁶⁴

An alternative measure of disorientation emerged in the discussion of Smith's lostness metric, which yielded the ratio of distinct nodes versus total nodes in the path of a user searching for facts within particular nodes. This ratio revealed itself as having potential as a measure of the integration of structural with conceptual knowledge, and therefore is well suited as an indicator of disorientation as defined in this paper. Moreover, the choice of this ratio as a measure of disorientation is supported by its similarity to measures employed by other researchers (Stanton, Correia, & Dias, 2000).

The disorientation statistic can, in addition, be supplemented by the total number of facts which were ultimately located by each participant (Smith, 1996). Although it could be claimed that this latter metric is more indicative of a poor structural conception of the hypertext than the degree of disorientation suffered, the failure to include it in the formula for disorientation would result in the same problems discussed in relation to structural knowledge. Their omission, and the subsequent subtraction of the Lostness value from one, would ensure that someone who received a high SK score for the reasons described would also obtain a Lostness value of close to zero. It is for the purpose of controlling for that possibility that the weighing functions are included.

$$Lostness = 1 - (N/S \times N_{tot}/S_{tot}) \times C/Fn$$

where

⁶⁴While one could argue that this represents another reason why the ratio of total versus distinct nodes is not an appropriate measure of structural knowledge, as discussed in connection with the fact-finding task, this point is probably less relevant when someone is more focused on locating facts than learning them

$Lostness$ =	the degree to which the student is lost,
N =	the number of different nodes visited during the search,
S =	the total number of nodes visited,
S_{tot} =	the sum of the total number of nodes visited,
N_{tot} =	the sum of the non-replacing distinct nodes in the path of each student, regardless of whether the fact was found or not,
C =	the number of questions correctly responded to, and
Fn =	the total number of facts required

Learning

Changes in factual knowledge were ascertained through an 18 point multiple choice questionnaire (Appendix C.3). Each of the questions was formulated on the basis of information found in the hypertext. The degree to which participants can recall information is predicted as depending on the amount of material content to which they have been exposed, how much background knowledge they possess, as well as the facilitation of recall provided by a well developed conceptual model of the content.

Research design

A pre-post experimental design was employed. The experiment was carried out over two sessions, due to concerns that an attempt to include all of the materials in a single session might not only be too demanding, but also render the post-tests susceptible to carry-over effects from the tests completed prior to the reading task.

Procedure

The experiment was conducted in two stages. The same grouping of students who participated in the first session, as determined by the dates for which they entered their names in tutorial lists, was used for the second stage. The 8 participants who attended the first group of the first session were treated as a pilot group, and their data for both sessions was omitted from the analyses. All of the experimental sessions took place at two o'clock in the afternoon. All of the questionnaires and the instructions for their completion were presented on computer. After each student had been placed in front of a computer during the

first session, and once the instructor had finished introducing himself,⁶⁵ the participants proceeded to complete the exercises presented them. The questionnaires were presented in the following order: (1) Background knowledge questionnaire, (2) Goal Orientation scale, (3) Factual knowledge test, (4) MSLQ scale, (5) Computer and Internet self-efficacy scale, and (6) the word sorting task. All of the questions posed in each the questionnaires had to be answered before the student was able to proceed to the next questionnaire.

The reading task was presented to the participants in the second session, after all the students had successfully attended the first tutorial meeting.⁶⁶ Of the five groups which took part in the second session (the first of which consisted of the pilot study participants), the second and third group fell in the Network hypertext condition, while the last two groups were provided with the hierarchical hypertext structure. This session was divided into three stages. In order to prevent the students from being able to backtrack from a later stage to a previous one, as well as to ensure that each of the students began a new section at the same time, a shortcut icon was placed on the screen for each of the sections. The participants were instructed to close each browser session before clicking on the icon to initiate a new one, and were only able to proceed past the first page of instruction for each stage by entering a user-name and password into a dialog box. This information was provided to them by the instructor once everyone was in a position to commence with the next stage.

Prior to the second stage the participants were given 5 minutes in which to familiarise themselves with the navigational features of the respective systems to which they were to be exposed, by browsing through a hypertext containing the instructions for the subsequent reading tasks, each of which were structured according to the navigational conventions employed within the hypertexts. The instructions for the Network condition additionally suggested that the participants use the hyperlinks, and attempt to ascertain how the texts connected by the hyperlinks were related. These instructions were included in order to encourage processing of the text — of which the hyperlinks are considered to form a part — in

⁶⁵The researcher undertook the role of instructor in all of the studies sessions, barring one meeting in the first stage, when he took ill. He was, in addition, assisted by a graduate student in the administration of the first meeting of the second session.

⁶⁶Barring four students who were only able to participate in the first session after the others, and whose data were not included in the analyses.

light of evidence that students will not do so without being provided with a task (Jonassen, 1993). While these instructions might be regarded as interfering with the effects of goal orientation, the expectation that second year students would process text which was originally intended for a post-graduate course, especially when they were informed that they were not to be evaluated on it,⁶⁷ was seen as placing too great a demand on the motivation of the student.

The training phase was followed by the main reading exercise, in which the participants were given a minimum of an hour in which to read through the course material. They were also informed that once they had completed the reading, they would be asked questions relating to what they had read. This was revealed to them so as to control for the effects that the possible speculation by certain of the participants regarding testing might have on the data generated.⁶⁸ The participants were instructed that were they to finish the text before the hour was completed, they should read through it again, as they may have missed something.

Once the hour allocated to the main reading task was over, and the students had completed reading the material in the hypertext, they were required to complete another series of questionnaires. Before they began, however, they had to read a web-page warning them that a number of the questionnaires that they were about to be presented with were the same as those completed in the first session. The first questionnaire consisted of two questions (Appendix C.6), the first of which provided them with a choice of five estimates of the hypertext's size, in page numbers, from which they were to select one. This item was included due to the identification in the literature that the inability to determine the size of a web-site is a major cause of disorientation (Otter & Johnson, 2000). The second question could only be answered with information contained within a reading provided to the students prior to the session. The article was entitled "Cognitive flexibility, Constructivism, and hypertext: Random access instruction for advanced knowledge acquisition in ill-structured domains", and was co-authored by Spiro, Feltovich, Jacobson, and Coulson (1996).

Although there was a great deal of concern that the provision of a reading dealing with topics which are so central to the present research programme would

⁶⁷Refer to the comment on page 70 with regards to this.

⁶⁸ The students were constantly reminded, however, that they would not be evaluated on the basis of their performance.

bias the results obtained, it was ultimately decided by the course convenor and I, that, given the status of the programme as a major part of the "Learning" component of the course, it was essential that we provide the students with material which heightened the relevance of the research programme to the content of the course. Unfortunately, time restrictions necessitated that the reading be given to the students prior to the second stage of the experiment. The question taken from the reading was therefore added in order to provide a gauge of who had actually read the article, and control for any bias that may have been introduced by this factor.

Upon submitting their responses to the first two questions asked, the participants were once again presented with first the factual, and then the conceptual knowledge assessment tasks employed in the first tutorial meeting. The multiple-choice and word-sorting tasks used were identical to those provided in the first session. The factual knowledge task was presented in both the first and second session, despite containing questions derived from the hypertext content, as the performance during the first session was to be used as a baseline against which to detect improvements in knowledge due to the reading. It was on the basis of the same rationale that the word-sorting task was repeated.

Once both of these tasks had been completed, the students proceeded to a page containing the instructions for the fact-finding task. A paper copy of these instructions was handed to the students at the beginning of the third stage of this session, should they have forgotten during the course of the task what exactly was expected of them.⁶⁹ The page following the instructions was partitioned into two frames. The first question that the students were required to answer was to be found in the top frame, while an exact replica of the hypertext which they had previously been exposed to was contained within the lower frame. The students were instructed to locate the information required to answer the questions in the hypertext, even if they felt they already knew the answer, for confirmation purposes. Once the participants had entered the information they found, and submitted it, they could proceed to the next question. The system was designed so that each time the participant proceeded to a new question, the hypertext was reset to the home page. This made it possible to compare the number of nodes visited by the student with the number it was necessary to pass through from the base node in

⁶⁹Refer to Appendix C.7 for these instructions, an example of the format in which the questions were presented, and a list of the actual questions asked.

order to reach the relevant information.⁷⁰ Once the participants had proceeded to the next question of the nine originally posed,⁷¹ it was not possible for them to return to any of the questions previously responded to.

Before proceeding with the exercises in both sessions, the instructor took pains to emphasise that the students were not going to be evaluated on their performance in the tasks assigned them. This was regarded as necessary so as to counteract the degree to which a competitive setting would enhance the qualities associated with performance orientation (as demonstrated by Ames 1988), and attenuate the participant's desire to learn for its own sake (Meece, Blumenfeld, & Hoyle, 1988).

Data analysis

All of the statistical analyses will be carried out by means of the Statistica (5.5) software package (StatSoft, 2000).

User characteristics

The reliability of the Goal Orientation scale will be determined by means of correlating the scores obtained on the performance and learning orientation dimensions with the sub-scales of the MSLQ. Pearson product-moment correlations will additionally be computed between these scores and levels of background knowledge and self-efficacy. A k-means cluster analysis procedure, as described below, will be performed in order to provide confirmation of the validity of the conceptualisation of goal orientation as discriminating between groups of individuals.

Hypertext characteristics

The stratum and compactness values for the hypertexts used in both the Network and Hierarchical conditions will be reported. This will provide one with an indication of the extent to which these measures are sensitive to differences between the hypertexts, and accordingly, how likely they are to detect variations found between the individual pathways taken by the participants.

⁷⁰A similar design was employed by Otter (2000), for that same reason.

⁷¹The number of questions was subsequently reduced to seven, for reasons which will be discussed in the Results section

Navigation

A tree-clustering procedure will be employed in detecting clusters of orientation-consistent navigational behaviours. Cluster analysis can be cast as a mathematical technique of classification, whereby variations in data are consolidated into groups or clusters in accordance with the results of an algorithm which attempts to reduce inter-group differences. A popular class of such algorithms, referred to as hierarchical agglomerative procedures, initiate the clustering process by regarding each data point as distinct, and through application of a linkage rule, merging those points which are most similar into clusters. One of the most commonly used of these procedures in behavioural science research is known as Ward's method, which operates on the same principle as the ANOVA statistical technique, by forming clusters with the purpose of reducing variation within groups, and increasing variation between them (StatSoft, 2000).

Hierarchical agglomeration can be compared with iterative partitioning algorithms, which select initial values, known as centroids, and then attempt to find data points which are close to these centroids (Aldenderfer & Blashfield, 1984). This procedure is repeated with the newly computed centroids until a stable state is reached in which no data points switch clusters. While the iterative nature of these rules overcome the major shortcoming of agglomerative procedures, which is their sensitivity to the initial clusters that are formed, they are computationally intensive, and require the use of heuristic rules to reduce the partitions into which data can be grouped. K-means cluster analysis, a popular variant of this class of algorithm, additionally requires the specification of the number of clusters searched for, prior to iteration through the data. The k-means method is thus more appropriately suited to circumstances in which the researcher is aware of the approximate number of groupings which best characterise the data (StatSoft, 2000). Agglomerative procedures, on the other hand, are better designed for exploratory research, in which the poor reliability of the technique is compensated for by the extent to which it facilitates one's ability to interpret otherwise complex data.

It is with these points in mind that a hierarchical cluster analysis, employing Ward's method of linkage, will be used to identify groups within the navigational data which may be associated with goal orientation. Attempts will be made to determine the reliability of the results of the analyses, by comparing them with the clusters formed using another, more conservative linkage rule, known as complete

linkage. A subsequent analysis will adopt an inverse approach to the identification of a relationship between goal orientation and hypertext navigation, by employing a k-means cluster analysis procedure in examining the nature of motivation amongst the participants in this study. While this latter analysis will be undertaken primarily due to suggestions that academic motivation might not be characterised in its entirety by the two dimensions of learning and performance orientation (Meece et al., 1988; Meece & Holt, 1993), as well as the finding, as conceded by the authors of a validation study of the goal orientation scale, that a third factor is occasionally discovered in the item analyses of responses to the scale (Roedel et al., 1994), by reclassifying the participants according to goal orientation, it should provide clarity as to the exact nature of the relationship between this construct and hypertext navigation.

Density.

The structural density of the hypertexts employed in the Hierarchical and Network conditions was computed by means of dividing the number of potential links between the 63 nodes of the hypertext by the number of actual connections which were present in each structure. The derivation of these statistics has already been described, and was performed using the Excel program of the Microsoft Office 97 package.

In order to determine whether it was true that learning-orientated participants exhibited a denser path than performance orientated subjects, it was necessary to compute a density statistic for each individual participant. This was due to the fact that the process of collapsing the navigational pathways for all members within each group led to the differences between the groups becoming obscured, as a result of the overlap caused by within-group variance. The procedure used for calculating individual density statistics involved the assignment, via a script, of a unique identifying number to each hypertext node in the web log. These numbers were then extracted from the log, and the list of nodes traversed by each participant was placed in a separate file by the author. A density matrix of 63 columns by 63 rows was subsequently generated for each participant. The ratio of the total number of links (which included both structural and hyperlinks in the Network condition) which could be traversed within the hypertexts to those which were actually employed was subsequently determined in the manner previously described.

In addition to the determination of a density statistic which gauged the total

number of links used, a variant (henceforth referred to as link density) which reflected the use of hyperlinks in the Network condition was also calculated. This necessitated the removal of the digits ("1") indicating common links in both the manually created density matrices of the hierarchical and network hypertexts. This left a density matrix for the hyperlinks in isolation. The hyperlink usage for each of the students was then determined by the degree of commonality between the hyperlink matrix and the student matrix.⁷² The grand total of the marginal row totals was calculated for each participant, as this indicated the number of links which were shared between these matrices. The division of this total by the total number of hyperlinks in the hypertext yielded a figure which provided a measure of the density or degree to which a particular student took advantage of the hyperlinks in the network hypertext condition.

Path Compactness and Stratum.

The formulae used for calculating path compactness and stratum are presented below.⁷³ All the matrix operations used in calculating these metrics were carried out by means of UCINET v.4 for Dos (Borgatti, Everett, & Freeman, 1992).

Path Compactness

$$P_{Cp} = \frac{P_{Max} - \sum_i \sum_j PC_{ij}}{P_{Max} - P_{Min}}$$

where

⁷²Achieved through the addition of the digits representing the presence/absence of hyperlinks in the system, and numbers indicating whether the hyperlinks were used. Thus, in each cell, the presence of the digit "0" stood for the absence of a link, "1" for an unused link, and "2" for a link that had been made use of.

⁷³Formulas taken from McEneaney, 1999

- P_{Cp} = Path compactness value
- PC = converted distance matrix of path (conversion referring to the entry of the number of nodes in user's path in cells which are completely disconnected)
- P_{Max} = converted distance value for completely disconnected network
- P_{Min} = converted distance value for completely connected network

P_{Max} and P_{Min} are calculated by means of the following equations

$$P_{Max} = K(n^2 - n)$$

$$P_{Min} = (n^2 - n)$$

where n equals the order of the path matrix and K the number of nodes in the original path matrix.

Path Stratum

$$P_{St} = \frac{\text{path absolute prestige}}{LAP}$$

The path absolute prestige refers to the absolute (unsigned) indicator of the difference between the number of steps required to reach all the other nodes in the hypertext from a particular node, and the number of nodes which need be traversed in order to reach that node.⁷⁴ Each of the prestige values for the entire hypertext is subsequently summed.

LAP stands for linear absolute prestige, and is calculated in the following way

⁷⁴In social network terms, someone has more prestige the fewer the number of intermediaries which have to be negotiated in order to reach a desired person, and the larger the number that others have to pass through in reaching oneself (Botafogo et al., 1992).

$$LAP = \begin{cases} \frac{n^3}{4} & : \text{ if } n \text{ is even.} \\ \frac{n^3-n}{4} & : \text{ if } n \text{ is odd.} \end{cases}$$

The inclusion of number of nodes in the formulae for both C_p and in calculating the linear absolute prestige for P_{St} serves a normalising function, and allows for the comparison of these metrics across hypertexts of different sizes. McEneaney (2001) recommends that one use the number of nodes accessed by the hypertext navigator when comparing the compactness and stratum of individuals' paths in the same hypertext. The substitution in the path distance matrix of the number of nodes accessed for the constant which represents all of the nodes which are not reachable from one another (known as the conversion constant), as well as the use of this value for normalisation purposes, facilitates the ease with which one can compare the navigation between participants.

Another factor which needs to be considered, as it has been shown to distort the stratum values produced for different hypertext systems (Botafogo et al., 1992), is the presence of reference nodes, or nodes which are reachable from a large number of other nodes in the system. The possible influence of this factor will be accommodated through the calculation of a separate path stratum index for a distance matrix derived from a matrix which excludes links leading to and from the home page, and the four module base nodes. These nodes are selected as reference nodes, due to their accessibility from all others, as is demonstrated by the fact that in both the Network and Hierarchy hypertexts the total number of steps required to access them is more than 1.5 standard deviations below the mean.

A one-way analysis of variance procedure will be employed in determining whether the navigational metrics of path compactness, stratum, density and link density vary significantly on the basis of goal orientation. An additional method of representing the qualities which these metrics capture is in a visual format, by means of path diagrams. These graphical illustrations have been used to good effect by McEneaney (1999), and are used here to display the combined navigational pathways followed by performance and learning oriented students. The diagrams were rendered using the Graphviz program (AT&T Lab-Research, 2002), and are presented for each of the groups in the Network condition only. Differences in the extent to which hyperlinks are exploited should be evident in the linearity of the graphs produced for the different groups.

Mental models

In order to determine the conceptual change which took place as a result of the reading task, both of the word-sorting lists created by each student were transcribed, with each word being assigned a number from 1 to 40. The strength of the associations between the words within the categories formed was determined by using an algorithm which extracted the cosine value between each of the combinations of the words found in that category from one of two LSA matrices,⁷⁵ and then dividing the total of these values with the number of comparisons made, in order to calculate the average coherence of the words in each category. All possible combinations were used, in order to overcome any word-order effects. The sum of the category coherence values obtained for each of the word lists was subsequently averaged across the number of categories formed, and the difference between the average cosine value obtained for all the categories between the first and second word sorting task provided a measure of conceptual change undergone by each participant.

This procedure was performed using the matrix derived from the TasaAll database available on the World-Wide-Web, as well as the CustomLSA database. The GTP software used to create this custom database was configured to use the log entropy weighing algorithm, which determines the significance of a word's occurrence within a particular context as a function of the number of times the words occurs both globally and within that particular context (Foltz & Wells, 1999). If a term appears very frequently within an entire corpus of text, as well as within a particular passage, log entropy weighing would reduce the strength of association which would ordinarily have been derived between this term and others in the passage, were that term's presence in the passage more distinctive.

In order to estimate the reliability of the LSA measures of word similarity, advantage was taken of the offer made by a member of staff in UCT's Multimedia Education Group, to liaise on the researcher's behalf with three people who were/are intimately involved in the operation of the Images of Africa course. The 40 words were presented on rectangular cue cards to Patrick Harries and Martin Hall, who lectured and provided the material for the Race and African cities components of the 1999 course, respectively, and Nick Shepherd, who was the course co-ordinator

⁷⁵excluding the cosine of 1 obtained from comparing a word to itself

at the time. They were instructed to sort the cards into groups of similar words. The confirmation of the prediction that the experts' word similarity scores are significantly higher than the scores for the categories formed by the participants in the first tutorial session would stand as verification of the reliability of this procedure.

The prediction that the structural and conceptual knowledge of the hypertext depend to a large extent on whether a person is learning or performance orientated will be tested by means of a full-factorial MANCOVA. Changes in word-sorting performance, and the SK metric will be the dependent variables, while goal orientation and hypertext condition will function as independent variables. In the interests of parsimony, and due to the similarity of the SK and Lostness formulas, the Lostness measure will also be included as a dependent variable in this particular analysis.

The appropriateness of dis-assembling the formula for Smith's measure of lostness into its constituent components will be verified by means of comparing the results obtained in the previously described analysis with those obtained from a factorial analysis of variance of Smith's measure across goal orientation and hypertext conditions.

Learning

The dependence of the ability to learn factual information from a hypertext is expected to reveal itself in a positive correlation between factual knowledge gain and both the structural and conceptual aspects of the mental model. This will be apparent in the Network condition, where the greatest potential occurs for the integration of these two forms of knowledge. A correlational analysis will therefore be employed to determine whether this relationship exists, as well as whether the effects of a poorly developed mental model on disorientation will be evident in the form of a negative correlation between factual knowledge gain and disorientation. The possibility that the possession of a greater degree of background knowledge might mask the influence of goal orientation and hypertext condition on recall, by reducing the amount of additional information a person is likely to be able to recall following the study session,⁷⁶ will be controlled for by conducting an ANCOVA of

⁷⁶Individuals who obtained higher pre-reading knowledge scores will in effect have less of a margin for the demonstration of improved recall. Although the failure of any of the participants to correctly

the factual knowledge recalled following the reading, and using the pre-reading background knowledge levels of information within the hypertext as a covariate. Should these background knowledge levels not influence the final degree of information recalled, this procedure will have the effect of removing the interference this variable introduces into a knowledge change metric, and thereby increase the chances of determining whether recall levels are influenced by goal orientation and hypertext condition.

answer all of the questions might be regarded as evidence that the danger of a ceiling effect has not been realised, such an effect might still play a role to the extent that it becomes increasingly difficult for students to improve on their performance.

Results

The results section of this paper will be divided into two components: the first attempting to more fully characterise and validate the constructs upon which the theory depends, while the latter component will examine the interaction between these constructs in their contribution to a coherent model.

Before proceeding, a caveat is in order regarding the sample size on which the findings to be presented are based. The full sample, excluding the pilot group, and those who did not complete the study in exactly the same conditions as the others, consisted of 75 participants. Due to the mechanical failure of the server while delivering the fact-finding task to one of the groups of students, and the subsequent loss of this information as a result of the premature termination of the session, the sample size for comparisons involving the fact-finding data dropped to 52.⁷⁷ This number decreases in size even further when reporting comparisons involving the LSA data, as a number of the participants did not return the print-outs of their completed word-sorting task to the instructor, resulting in the loss of data for 10 additional participants (N=42). The sample size discrepancy will be accommodated by reporting the sample sizes for the separate analyses conducted.

Characterisation of the data

Reliability and validity

The relative novelty of many of the measures used in this dissertation, at least insofar as their application in this domain is concerned, necessitates a demonstration of both their reliability, and their validity. The former is achieved through providing statistics indicating the internal consistency of the scales used, while the external validity of the measures employed can be determined through the agreement of the correlations obtained between the partially redundant measures included in this study with those which would be predicted on the basis of findings reported in the literature.⁷⁸

⁷⁷It was for this reason, that, in order to decrease the load on the server when students were engaged in the fact-finding task, the number of facts they were required to find was reduced from nine to seven.

⁷⁸The figures reported in this section have been rounded off to two decimal points, with the exception of analyses in which the magnitude of the values reported warrants greater precision. In these cases three decimal points have been used.

Participant variables.

Table 2: Correlations between variables characterising research participants

	1	2	3	4	5	6	7	8	9
1. Performance Orientation	–								
2. Learning Orientation	.07	–							
3. Extrinsic Motivation	.3*	-.11	–						
4. Intrinsic Motivation	-.07	.41**	.05	–					
5. Anxiety	.64**	.22	.34**	.08	–				
6. Effort	.02	.63**	-.34**	.2	.16	–			
7. Hypertext efficacy	.01	.36**	-.11	.28*	.16	.1	–		
8. Background knowledge	.11	-.04	.08	.05	-.16	-.14	.02	–	
9. Gender	.05	.23	-.01	.13	-.00	.07	.21	.25*	–

* $p < .05$ ** $p < .01$

A Cronbach alpha of .76 and .71 was obtained for the learning and performance motivation subscales of the goal orientation scale, respectively. As expected, and as observed in Table 2, the scores on the learning and performance orientation sub-scales did not correlate significantly with one another. Evidence of the validity of the motivational groups identified by the Goal Orientation scale is provided by the high correlation between performance orientation and both the extrinsic motivation ($r = .3, p < .05$), and anxiety ($r = .64, p < .01$) subscales of the MSLQ, while learning orientation correlates significantly with the intrinsic motivation ($r = .41, p < .01$) and effort ($r = .63, p < .01$) subscales. This is consistent with the nature of the relationships one would predict from the literature. With the possible exception of the Intrinsic motivation subscale of the MSLQ (4 items at $\alpha = .58$), the other subscales demonstrate relatively high levels of internal consistency (Anxiety [5 items, $\alpha = .67$], Extrinsic motivation [4 items, $\alpha = .77$], Effort [4 items, $\alpha = .68$]), especially considering the low number of items within each subscale.

The Computer and Internet self efficacy scales both proved to be reliable at $\alpha = .91$ and $\alpha = .84$, respectively. These scales were, in addition, very strongly correlated with one another ($r = .56, p < .01$), which, given the prominent involvement of computers in the use of the internet, lends the scores obtained from them greater validity. For purposes of ease of interpretation, and given the strength

of the correlation between these two measures, it was decided to sum the results of the scales. Henceforth, unless otherwise indicated, all references to hypertext self-efficacy will refer to this combined measure.

Gender was included in the analysis due to observations that it has been a determinant of levels of goal orientation in the past (Dweck, 1986; Kaplan, 2001). No significant correlation was detected between gender and either learning or performance orientation, although gender was a nominally significant predictor of self-efficacy levels ($r = .22, p < .07$). The positive correlation between gender and background knowledge is difficult to interpret, given doubts as to the reliability of the scale used for detecting the latter.⁷⁹

Hypertext variables.

Both hypertext systems consisted of 63 nodes and 419 structural links. The hypertext systems in the Network condition, contained an additional 199 hyperlinks. The influence of the hyperlinks was evident in the density score of 0.158 for the hypertexts in the Network condition, versus 0.107 in the Hierarchical hypertexts. The stratum values obtained for the Network and Hierarchical hypertexts were 0.028 and 0.08, respectively, with compactness values of 0.68 and 0.37.⁸⁰ These compactness values are consistent with the expectation that the hypertext in the Network condition will be both more greatly interconnected as a result of the presence of the hyperlinks, than is the case for the system employed in the Hierarchical condition. It is this greater degree of linkage between nodes which accounts for the lower stratum value of the Network hypertext, as it ensures a larger number of access routes to, and exit points from, the different nodes. The finding that these values are consistent with predictions provides one with some confidence in the reliability with which they will detect differences between the paths taken by participants.

⁷⁹As reported in the discussion of the reliability of the knowledge tests employed, on page 84

⁸⁰The compactness values were obtained through substitution of the maximum compactness score for the hypertext in the Network condition into the formula for calculating the compactness of the hypertext without hyperlinks. This was undertaken due to the finding of near equivalence for the scores (at 0.68), for the hypertexts when their respective maximum values were used. This can be considered acceptable practise, however, as the maximum chosen is largely a matter of convention, fulfilling as it does a normalisation function.

Table 3: Correlations amongst navigationally dependent metrics

	SK	Lostness	Density	Compactness	Stratum
SK	—				
Lostness	-.52** _a	—			
Density	.12 _a	-.03 _a	—		
Compactness	-.08 _a	-.15 _a	.15 _b	—	
Stratum	.19 _a	-.21 _a	.08 _b	-.13 _b	—

Note. The normalising variables and conversion constants used for the path compactness and stratum metrics are based on number of nodes in the user's path.

^a $n = 46$. ^b $n = 75$.

$p < .05$ ** $p < .01$

Navigational variables.

The correlations between all of the measures in this study which are based on the pathway followed in hypertext navigation are presented in Table 3. It is apparent from the table that the SK and Lostness scores are inversely proportional to one another. The magnitude of the negative correlation observed is partially attributable to a degree of redundancy between the formulas used, as was illustrated through the failure to find significance when the unweighted measures were compared ($r = -.06; p > .1$). Although this might cast doubt on the validity of the measures, the emergence of a positive relationship between them when SK is not weighed provides reason to conclude that the inclusion of each of these elements was justified. This pattern of results is consistent with a strategy employed by hypertext users of restricting their searches to those questions which they were confident of answering.

While the correlation between path stratum and compactness is in the expected direction, the relationship is not statistically significant. This may be largely the result of the relatively small sample employed, as the effect size of the relationship has been demonstrated to be sensitive to variations in sample size (McEneaney, 2001). Moreover, when the total number of hypertext nodes is substituted for normalisation purposes, and as the conversion constant in the distance matrix, a significantly positive correlation ($r = .49, p < .01$) emerges between these two variables. The same pattern comes to the fore with regards to the relationship

Table 4: Correlations between hypertext efficacy, background knowledge and navigational measures

	Density	Link density	P. compactness	P. stratum
Hypertext efficacy	.195	.058	-.047	.196
Background knowledge	.060	.143	-.024	.028

between density and path compactness, where the non-significant relationship as reported in Table 3 achieves significance ($r = .42, p < .01$) when the formula for compactness includes the hypertext node count. The latter manifestation of this pattern most probably reflects the influence of variations in the number of nodes which are visited by the participants, as these only reveal themselves in relation to the size of the hypertext in its entirety. Support for this proposition is provided by the significant correlation evident between the number of unique nodes the subjects encountered, and both their hypertext compactness ($r = .48, p < .01$) and density ($r = .59, p < .01$) scores. This bolsters the contention made earlier that density and path compactness are to be interpreted in a similar fashion.

A t-test for dependent measures revealed no significant difference ($df = 74, t = .83, p > .1$) between the path stratum measure which was based on all of the nodes traversed by the participants ($Mean = .47, SD = .47$), and that from which the influence of the reference nodes was excluded ($Mean = .42, SD = .3$). The former path stratum metric will accordingly be included in subsequent analyses.

Table 4 reveals that neither levels of hypertext efficacy or background knowledge are significantly correlated with the navigational variables. Their failure to do so indicates that they do not influence navigation as predicted by the model presented in Figure 1, and they will consequently be omitted from further consideration.

Knowledge tests.

The Cronbach alpha for the confidence with which participants felt they had responded to the background knowledge items correctly was acceptable, at .79 ($N = 74$). However, the same was not true for the responses given to the items themselves ($\alpha = .05$), and indicates that the scale might not be useful in detecting

differences in background knowledge. The failure to observe a significant correlation between correct responses and confidence in total, as well as between responses to the individual items and the confidence ratings, rules out the possibility of using confidence rating as a means of selecting those questions which might be more reflective of the participants background knowledge.

The mean word similarity ratings for the grouping of words by the three experts when using the WebLSA (.145) and CustomLSA (.031) databases falls within the range of those obtained by the students ($N = 62$), both prior to the reading task (Mean = .174 and .031, respectively), and following it (Mean = .167 and .031, respectively). This suggests that the LSA procedures, as employed in this study, do not provide a reliable measure of word similarity. The failure to find significant differences between how experts and novices group terms might arise from a number of different factors, including the lack of sensitivity of the LSA procedures. Other possible underlying causes of the negative result obtained will be discussed in the next section.

Exploratory analyses

Navigational data.

Separate cluster analyses were performed on the navigational data for each of the hypertext groups. This was due to differences anticipated in the navigation of these structures, and the observation that students spent an average of nine minutes⁸¹ longer in the hypertext in the Hierarchical condition, than in the system employed in the Network condition. The variables used as input to the analysis were the average time participants spent in individual hypertext nodes, the number of distinct nodes they visited, the total number of nodes they revisited, and the total number of times the base node was returned to. The similarity measure used as input for the cluster analysis was a correlation matrix for each participant in the network group. This data format was selected in preference to the raw data, as it is the relationship between the variables which is of interest, and as this counters the susceptibility of the Ward's method linkage algorithm to reflect variations in magnitude over that of shape (Aldenderfer & Blashfield, 1984).

The cluster analysis of the network data yielded two major clusters, the

⁸¹This statistic excludes the influence of two outliers, whose data were also excluded from the subsequent cluster analysis.

stability of which was confirmed by the partial replication of these groups with the administration of the complete linkages algorithm.⁸² This is a relatively strong indicator of reliability, given the conservative nature of this algorithm, and the small size of the sample ($N = 37$) used. The descriptive statistics for the relevant user and navigational variables of the group members are provided in Table 5.

Table 5: Comparison of means and standard deviations for clusters in the Network condition

Variable	Cluster No.1 _a		Cluster No.2 _b	
	Mean	SD	Mean	SD
Performance Orientation	14.63	3.79	14.06	4.40
Learning Orientation	28.37	6.23	27.89	6.03
Extrinsic Motivation	19.37	6.74	21.39	6.22
Intrinsic Motivation	18.47	4.142	17.94	3
Anxiety	18.74	5.94	19.44	4.30
Effort	18.79	4.58	16.89	3.98
Hypertext Efficacy	103.9	13.73	105.28	18.45
Background Knowledge	4.9	1.29	5.00	1.28
Distinct Nodes	45.42	9.55	55.06	5.16
Avg. time per node	48.14	15.19	23.25	7.71
Node repeats	23.47	12.01	86.94	47.15
Base repeats	14.21	8.52	22.67	12.36

^a $n = 19$. ^b $n = 18$.

Arriving at a cluster solution for the hierarchy data was more problematic, with the two cluster analytic procedures yielding groups with substantially different memberships. In addition, the similarity measures within the correlation matrix had to be replaced with raw scores, as the former data type produced a large number of small, similar and difficult to interpret clusters when either linkage rule was employed.

The solution arrived at was once again composed of two major groupings. The data for these clusters are presented in tabular format in Table 6.

It is apparent from an examination of Table 5 and 6 that the four navigational

⁸²Only two participants swapped groups between the two procedures

Table 6: Comparison of means and standard deviations for clusters in the Hierarchy condition

Variable	Cluster No.1 ^a		Cluster No.2 ^b	
	Mean	SD	Mean	SD
Performance Orientation	10.71	3.61	11.07	3.17
Learning Orientation	30.19	4.82	31.27	6.38
Extrinsic Motivation	21.81	3.87	18.93	5.59
Intrinsic Motivation	18.00	2.37	19.27	4.17
Anxiety	17.38	5.37	19.27	5.95
Effort	17.00	3.49	17.13	5.93
Hypertext Efficacy	92.62	21.04	102.87	16.62
Background Knowledge	4.81	1.33	4.67	1.54
Distinct Nodes	53.1	9.57	43.00	12.25
Avg. time per node	44.55	5.83	71.05	9.19
Node repeats	19.05	12.07	8.20	6.04
Base repeats	15.14	7.96	10.47	2.17

^a*n* = 21. ^b*n* = 15.

variables which were fed as input to the cluster analyses do not discriminate clearly between groups on the basis of their goal orientation for either hypertexts in the Network or Hierarchical conditions. While the observation that those participants who were extrinsically motivated in the former conditions were prepared to expend less effort is consistent with the literature, the finding that the extrinsically motivated group in the Hierarchical hypertext were less anxious is not. It was nevertheless decided, for the sake of completeness, to report the statistics for the four largest clusters within the network and hierarchy conditions. The graphical presentation of the cluster solutions arrived at, as well as their descriptive statistics, can be found in Appendix A.

Goal Orientation.

K-means cluster analysis was subsequently applied to the goal orientation data, in order to determine whether it is possible to translate the two dimensions which goal orientation has traditionally been characterised as consisting of, into a number

of distinct groups of participants,⁸³ and in a manner consistent with the characteristics which have been associated with these dimensions. The determination of the number of groups required would also be of theoretical relevance, due to contention as to whether the learning and performance orientation groups which have been arrived at through splitting scale - derived scores on the median exhaust the number of motivational groups to be found.

The heuristic employed in the selection of the initial centroids was that of sorting the list of distances between data points, and selecting points at regular intervals to stand as cluster centres. This yielded a structure which was more consistent with the literature and hence easier to interpret than choosing the most clearly differentiated data points as centroids.

⁸³Cluster analysis has an advantage over median - splitting techniques in that it can recover distinct groups of scores which reflect associations between the dimensions otherwise lost when applying a procedure which assumes that the dimensions are homogenous (Meece & Holt, 1993)

Table 7: ANOVA results for differences between goal - orientation and associated variables for the two cluster solution

Source	<i>df</i>	<i>F</i>	η^2	<i>p</i>
Performance orientation	1	0.69	.00	.41
Error	73	(13.93)		
Learning orientation	1	103.68**	.57	.00
Error	73	(13.68)		
Extrinsic motivation	1	2.24	.03	.14
Error	73	(33.34)		
Intrinsic Motivation	1	5.57*	.07	.02
Error	73	(10.73)		
Anxiety	1	5.66*	.07	.02
Error	73	(26.67)		
Effort	1	65.54**	.47	.00
Error	73	(11.1)		

Note. Values enclosed in parentheses represent mean square errors

* $p < .05$. ** $p < .01$.

The two - group cluster analysis discriminated between clusters which were significantly different from one another with regards to learning orientation, intrinsic motivation, anxiety and effort. This is evident in the ANOVA results for the inter - group comparisons between the variables analysed, as presented in table 7.⁸⁴ Of the four significant differences observed, the most extreme were those on the learning orientation and effort subscales, where the clusters formed account for 57% and 47%, respectively, of the total variance of the scores obtained for these variables.⁸⁵ In acknowledgement of the particularly distinctive role of learning orientation, and due to the emphasis in this study on motivation, group of individuals which Figure 3 reveals as scoring high on learning orientation will be labelled as Learning Oriented (LO), while those with a lower profile on this dimension are to be classified as Non-Learning Oriented (NLO).

While participants in the LO group indicated that they would be prepared to expend greater effort in a task, as predicted, they obtained higher anxiety scores⁸⁶ as well, which is more strongly associated in the literature with performance orientation. This association reveals itself when three clusters are selected, as is evident in figure 4, and in which there is a clear dissociation between a learning and performance orientated group.

The claim of differentiation between the groups on the basis of goal orientation is supported by the emergence of a statistically significant difference between the performance orientation score for the different groups, when comparing the tabulated results of the three cluster one-way ANOVA (Table 9, on page 93), to those found in Table 7.⁸⁷ These differences are also apparent when comparing the profiles of the different clusters on the motivational variables for the two (Figure 3)


⁸⁴The significance results are presented in the form of multiple one-way ANOVAs for each individual comparison of means, as this is how Statistica presents them from within its cluster analysis module

⁸⁵While η^2 is recognised as a biased estimate of the magnitude of the experimental effect, it is used here due to the intuitive ease with which it can be interpreted.

⁸⁶A mean of 20.14 vs. 17.3

⁸⁷The interpretation of these results should be made with caution, however, given that as post-hoc statistical tests, they are particularly prone to Type I errors. Their inclusion in this report is only considered justifiable due to the exploratory nature of this component of the study. They do offer the advantage, however, of overcoming the fact that the absolute differences between the groups in terms of their means for the individual variables are partly a function of scale-specific differences, as the scores used were not standardised. For instance, the relatively large difference between learning orientation mean scores compared to that observed for performance orientation can be attributed to a certain degree to the fact that there are 7 additional items in the learning orientation subscale.

Figure 3. 2 K-means cluster profiles for goal orientation and associated variables



and three group (Figure 4, on page 91) scenarios.


The descriptive statistics for motivational variables in the three group solution is presented in table 8. The poor scoring of the third cluster on all of the variables, apart from intrinsic and extrinsic motivation, is consistent with the motivational group identified by Meece and colleagues (1988, 1993) as "work - avoidant".

Table 8: Motivational variables for 3 k-means cluster analysis

	Cluster 1 ^a		Cluster 2 ^b		Cluster 3 ^c	
	Mean	SD	Mean	SD	Mean	SD
Performance Orientation	9.85	4.00	12.62	2.83	7.83	2.81
Learning Orientation	37.25	3.78	30.81	4.08	25.67	3.91
Extrinsic Motivation	15.55	4.77	23.38	4.52	19.11	5.50
Intrinsic Motivation	20.00	3.81	17.87	2.64	17.56	3.78
Anxiety	18.30	5.10	21.65	3.02	12.78	4.33
Effort	22.20	3.30	17.35	3.05	13.06	3.42

^a $n = 37$. ^b $n = 20$. ^c $n = 18$.

Figure 4. Profile of 3 k-means clusters on motivational variables



An inspection of how the participants in the two cluster solution align themselves within three groups provides insight into the most important distinctions between the participants with respect to motivation. An examination of cluster membership reveals that a neat re-alignment occurs when moving from the 3 to 2 cluster groupings, with all of the high learning orientation students falling in the LO group, and all of the work avoidant students finding membership with the NLO group. The additional cases found in the LO and NLO groupings are provided by the performance orientation cluster in the 3 cluster solution. This suggests that the characteristic which most clearly distinguished the participants was that of where they were located on the dimension of learning orientation, and not performance orientation. It is with this in mind that all subsequent analyses in which student motivation will be considered will make the division between learning and non-learning orientation groups. The focus on two, as opposed to a larger number of groups has the additional advantage of swelling the sample size for each of the respective groups, an important consideration when the total sample size is so small.⁸⁸

⁸⁸The identification of two discrete groups through cluster analysis has the advantage of avoiding the exponential increase in group numbers resulting from the partitioning of dimensional attributes.

Table 9: ANOVA results for differences between goal - orientation and associated variables for three cluster solution

Source	<i>df</i>	<i>F</i>	η^2	<i>p</i>
Performance Orientation	2	14.79**	.17	.00
Error	72	(10.11)		
Learning Orientation	2	41.06**	.53	.00
Error	72	(15.69)		
Extrinsic motivation	2	17.71**	.33	.00
Error	72	(23.35)		
Intrinsic motivation	2	3.49	.09	.04
Error	72	(10.68)		
Anxiety	2	30.14	.46	.00
Error	72	(15.86)		
Effort	2	38.75	.52	.00
Error	72	(10.29)		

Note. Values enclosed in parentheses represent mean square errors

* $p < .05$. ** $p < .01$.

The effect of goal orientation on navigation

The intention of using the ANOVA statistical procedure in testing the prediction that the motivational group participants belonged to would significantly influence how they navigated through the hypertext, was rendered problematic by the failure of the data to conform to the assumptions upon which the reliability of the ANOVA's results is based. These are that the distributions of sample data being compared are normal in shape, and that they do not vary significantly in terms of variance. In order to accommodate the unequal sample sizes from which the groups being compared were composed, the Brown-Forsyth test of homogeneity of variances was employed. It was found that the variances across the groups were heterogeneous for the link density ($F = 6.07, p < .05$) and path stratum ($F = 3.96, p < .05$) scores. In addition, the distributions for path stratum and compactness were strongly positively skewed. Although Howell (1997) recommends the application of a logarithmic transformation to data in order to simultaneously homogenise the variances and normalise the shape of positively skewed distributions with standard deviations which are roughly proportional to their means (criteria which were met by the navigational data to be compared), the logarithmic transformation of the raw data did not remedy the failure of the data to satisfy the assumptions.

It was due to concerns regarding the appropriateness of employing statistics which depend on these assumptions that the non-parametric equivalent of the one-way ANOVA was employed. The Kruskal-Wallis one-way analysis of variance avoids the dependence on the normal distribution through compartmentalisation of continuous data, by assigning each data point a rank. It is the extent of the divergence of the sum of the ranks for the smallest comparison group from that which would be predicted, which is used to determine whether the groups being compared originate from the same population.

Although the results reported in Table 10 indicate that path stratum and compactness did not vary as a result of goal orientation when no consideration is given of hypertext condition, the results of the Kruskal-Wallis ANOVA, as presented in Table 11 can be interpreted as supporting the hypothesis that individuals who are learning orientated make more extensive use of the navigational possibilities which distinguish hyperlinked hypertexts from hypertexts which do not feature this navigational device. This is evident in the fact that there was a significant difference in the number of hyperlinks which were used in the Network condition. Moreover,

Table 10: Results of the effect of goal orientation on the navigational metrics.

Variable	Sample Size	Sum of Ranks	Probability
Density			
LO	35	1423.5	.32
NLO	40	1426.5	
Compactness			
LO	35	1361	.74
NLO	40	1489	
Stratum			
LO	35	1354	.80
NLO	40	1496	

Note. The results reported in this table are for comparisons collapsed across both hypertext conditions.

the direction of this difference was consistent with expectations, with the sum of ranks for the LO group being substantially the larger. This is graphically portrayed by means of path diagrams in Appendix B.4, where it is evident that the navigational route followed by the LO group in the network condition is more heavily characterised by the use of hyperlinks.⁸⁹ The relatively low power of non-parametric statistical analyses provides further reason for confidence in the outcome of this test. No differences were observed for comparisons involving either the path compactness or stratum measure.

⁸⁹The connections between nodes in the diagrams represent links which were employed on at least two separate occasions. This not only simplified the diagram, but also removed paths which are more likely to be the result of accident, or were otherwise spurious. The modules within which each node is found is indicated by the shade of the node, while the number assigned to each node reveals where they would be positioned were a participant only to follow the sequence dictated by the static links embedded in each page.

Table 11: Results of the effect of goal orientation on navigational metrics in the Network condition

Link Density				
LO	20	444.5	.048	
NLO	17	258.5		
Density				
LO	20	449.5	.034	
NLO	17	253.5		

Mental models and disorientation

It was initially the intention to compare within one analysis the two orientational groups which have been identified on both the measure of structural and conceptual knowledge, as well as on the Lostness metric. However, the evidence that the LSA measures are not reliable, in combination with the reduction in sample size when these measures are included,⁹⁰ was deemed good reason to analyse the word-sorting data separately.

Structural aspect of mental model

In order to determine whether there was a significant difference in the structural knowledge gained and in the degree of disorientation suffered by participants as a function of the orientational groups that they were members of, and of the hypertext structure they read from, a 2 way full-factorial MANCOVA was conducted. This procedure has the advantage of parsimony with regards to its ability to test differences on multiple dependent variables, and thereby reduce the possibility of committing a Type I error. Hypertext condition and goal orientation were the independent variables and SK and Lostness the dependent variables.

The suitability of the MANCOVA procedure was verified through non-significance on the Box M Multivariate Test for Homogeneity of Variances/Covariances ($\chi^2 = 26.51, p > .1$), and by the fact that the assumption of normality was upheld.

⁹⁰See comment on page 79 regarding sample size discrepancies

The descriptive statistics for the four groups, when compared across hypertext structure and goal orientation, are presented in Table 12, and appear to indicate that knowledge of hypertext structure within the respective orientational groups is influenced by the nature of that structure. The influence is in opposing directions, however, with the addition of hyperlinks to the hypertext benefiting the LO group, while impacting negatively on the structural knowledge of those who fall in the NLO group. This pattern is not duplicated as far as Lostness is concerned, where only the NLO participants appear (adversely) effected by the presence of hyperlinks.

Figure 5. The interaction of goal orientation and hypertext structure on disorientation and structural hypertext knowledge




Table 12: Descriptive statistics for mental model - dependent variables as a function of hypertext structure and goal orientation.

Group	Variable	<i>SD</i>	Means
Hierarchy Hypertext			
NLO _a	SK	0.05	0.10
	Lostness	0.23	0.58
LO _b	SK	0.06	0.09
	Lostness	0.18	0.50
Network Hypertext			
NLO _c	SK	0.04	0.06
	Lostness	0.20	0.72
LO _d	SK	0.07	0.15
	Lostness	0.12	0.47

Note. NLO = Non Learning Oriented students. LO = Learning Oriented students.

^a*n* = 21. ^b*n* = 13. ^c*n* = 6. ^d*n* = 6.

Table 13: ANOVA summary table of specific effects of goal orientation and hypertext structure on mental model variables

Source	<i>df</i>	<i>F</i>	η^2	<i>p</i>
Goal Orientation				
SK	1	4.660*	.091	.037
Lostness	1	5.947*	.123	.019
Hypertext				
SK	1	0.443	.009	.509
Lostness	1	0.714	.015	.403
Interaction				
SK	1	8.19**	.159	.007
Lostness	1	1.78	.037	.189
Error				
SK	42	(0.003)		
Lostness	42	(0.04)		

Note. Values enclosed in parentheses represent mean square errors.

* $p < .05$. ** $p < .01$.

The interaction of hypertext structure and goal orientation on the mental model variables is illustrated in Figure 5. It is apparent from the graph that goal orientation exerts its largest effect within the Network condition. This effect is in the expected direction, with high learners demonstrating an increase in structural knowledge, and a concomitant decrease in disorientation, relative to the low learners. As suggested by the results reported in Table 13, the main effect of goal orientation is significant (Wilks' $\lambda = .85$, Rao R (2, 39) = 3.49, $p < .039$), while hypertext structure by itself does not appear to influence these variables. The interaction between hypertext structure and goal orientation, on the other hand, is significant (Wilks' $\lambda = .84$, Rao R (2, 39) = 3.99, $p < .026$), providing a basis for the claim that goal orientation determines the adequacy of one's mental model of a hypertext, as a function of the presence of hyperlinks.

Comparison of results with Smith's measure of lostness


Given the failure to reject the assumption of the homogeneity of variances for Smith's lostness measure ($F = 2.21$, $p > .1$), as well as the evidence of no substantial departures of the variable's distributions from normality, it was possible to conduct a univariate factorial analysis of variance of the hypothesis that the effect of hypertext condition on Smith's measure of lostness will depend on goal orientation.

The finding of that the LO and NLO groups were not significantly different across hypertext conditions on Smith's measure of lostness is reflected in the ANOVA summary table (Table 14), as well as graphically in Figure 6. This suggests that the dis-assembly of Smith's (1996) construct into separate structural and conceptual/structural components might provide greater sensitivity to the statistical analysis conducted, to detect differences in participant's mental models of hypertext systems.

Conceptual aspect of mental model

Two separate uni-variate factorial ANOVAs were conducted to test the claim that the variation in conceptual knowledge held by participants across hypertext structures is attributable to differences in goal orientation. The distributions of the scores for the two word-sorting measures revealed no unacceptably large departures from normality, while the results of Levene's test of the homogeneity across sub-groups (WebLSA: $F = 0.12$, $p < .1$; CustomLSA: $F = 1.08$, $p > .1$)

Figure 6. Effect of interaction of goal orientation and hypertext on Smith's measure of lostness



indicated that the groups did not differ substantially in terms of within - group variance.

The means in table 15 reveal the change in conceptual knowledge as measured by differences in performance on the word-sorting task to be very small.⁹¹ All of the differences have a mean of less than 1 percent, with a number of the figures lying in the negative direction. This, in combination with the observation that the standard deviations for all of the groups, across both data-bases, are larger than the means, suggests that the LSA measures of conceptual change are highly unreliable.

A two-way analysis of variance was subsequently conducted with the similarity scores for each semantic space tested separately. A non-significant result for Levene's test of the homogeneity of variances, in combination with no substantial

⁹¹These figures, when multiplied by one hundred, can be interpreted as percentage change in conceptual knowledge.

Table 14: ANOVA summary table of specific effects of Smith's measure of lostness

Source	<i>df</i>	<i>F</i>	η^2	<i>p</i>
Goal Orientation	1	0.13	.00	.72
Hypertext	1	0.044	.00	.84
Interaction	1	3.802	.09	.058
Error	40	(0.046)		

Note. Values enclosed in parentheses represent mean square errors

* $p < .05$. ** $p < .01$.

departures from normality, indicated that this test was appropriate.

Table 15: Descriptive data for two LSA procedures

Group	WebLSA		CustomLSA	
	Mean	<i>SD</i>	Mean	<i>SD</i>
Hierarchy Hypertext				
NLO _a	-0.007	0.029	-0.000	0.008
LO _b	-0.011	0.037	0.001	0.007
Network Hypertext				
NLO _c	0.001	0.033	0.005	0.009
LO _d	-0.011	0.03	-0.002	0.006

Note. NLO = Non Learning Oriented students. LO = Learning Oriented students.

^a $n = 19$. ^b $n = 11$. ^c $n = 15$. ^d $n = 17$.

Table 16: ANOVA summary table for WebLSA

Variable	<i>df</i>	<i>F</i>	η^2	<i>p</i>
Goal Orientation	1	0.947	.016	.335
Hypertext	1	0.216	.004	.644
Interaction	1	0.266	.004	.608
Error	56	(0.00)		

Note. Values enclosed in parentheses represent mean square errors

* $p < .05$. ** $p < .01$.

The ANOVA statistical procedure with the CustomLSA scores as its dependent variable was performed in order to determine whether the data corpus which was based on reading materials purportedly more relevant to the hypertext content (Appendix D.2) would be more successful in reflecting conceptual change undergone by the participants as a consequence of studying the hypertext. The evidence of homogeneity of variances ($p = .36$) for the CustomLSA variable and the normality of its distribution, indicates the legitimacy of employing this statistical procedure.

The summary tables of the results of the analyses of variance indicate that no main or interaction effects were statistically significant. Given the demonstration of

Table 17: ANOVA summary table for CustomLSA

Variable	<i>df</i>	<i>F</i>	η^2	<i>p</i>
Goal Orientation	1	1.228	.024	.273
Hypertext	1	0.34	.004	.56
Interaction	1	3.02	.058	.088
Error	56	(0.00)		

Note. Values enclosed in parentheses represent mean square errors

* $p < .05$. ** $p < .01$.

the lack of reliability of the data should not be interpreted as evidence of the absence of an effect of goal orientation on conceptual change in hypertexts, but merely as the absence of evidence of such an effect.

The influence of mental models on factual knowledge gain

The evidence of the nature of the relationship between the mental model indicators and factual knowledge gain was not encouraging. The correlation between the performance on the structural knowledge measure and change in knowledge brought about by the study session was not significant ($r = -.11, p < .42$).⁹² However, despite the finding of no association between the motivational groups and either path stratum or path compactness, variables which it is postulated affect the mental model derived from a hypertext, the evidence of a nominally significant relationship ($r = -.26, p < .1, N = 46$) between Lostness and knowledge gain provides some support for the notion of the importance of mental models for the later recall of factual knowledge.

⁹²The conceptual change values were not included in the correlational analysis, due to their demonstration of poor reliability, and in the interests of minimising the probability of committing a Type I error.

Table 18: Descriptive statistics for analysis of the recall of factual knowledge

Group	Variable	<i>SD</i>	Means	Adj. Means
Hierarchy Hypertext				
NLO _a	Post-know	2.64	8.96	8.91
	Pre-know	2.04	4.48	—
LO _b	Post-know	2.91	9.2	9.19
	Pre-know	2.03	4.13	—
Network Hypertext				
NLO _c	Post-know	2.12	6.59	6.65
	Pre-know	1.51	3.47	—
LO _d	Post-know	3.13	7	7.01
	Pre-know	1.28	3.95	—

Note. Pre-know stands for the base levels of the participants' knowledge prior to the study phase, and its contribution to subsequent knowledge gain is controlled for through its inclusion in this analysis as covariate to the Post-know variable. NLO = Non Learning Oriented students. LO = Learning Oriented students.

^a*n* = 23. ^b*n* = 15. ^c*n* = 17 ^d*n* = 20

The influence of goal orientation on recall

In the light of the finding that goal orientation affected both hypertext navigation (Table 11) and mental models (Table 13), and for the reasons provided on page 77, evidence was sought for a relationship between goal orientation and the recall of factual knowledge. This aim was pursued by means of conducting an ANCOVA. The non-equivalence of motivational groups in terms of their initial levels of knowledge was controlled for by means of using the scores they obtained prior to the study session as a covariate of the dependent variable. A Levene's test for Homogeneity of Variances provided evidence for the appropriateness of using an ANCOVA, as both the Post-Knowledge ($F = 1.36, p > .1$) variable, and the Pre-Knowledge covariate ($F = 1.96, p > .1$) failed to attain significance. In addition,

the distributions of the participant's scores within the four groups formed by the hypertext and goal orientation independent variables did not deviate substantially from normality. There was no evidence of an interaction or linear relationship between the dependent variable and the covariate (Appendix A.3). The failure to find a significant correlation between pre and post - knowledge levels, while interpretable as an indictment against the claim that background knowledge⁹³ affects knowledge gain, will remain in the subsequent analysis for purposes of theoretical consistency.

Table 19: ANCOVA summary table of specific effects for analysis of interaction of goal orientation and hypertext on recall

Source	<i>df</i>	<i>F</i>	η^2	<i>p</i>
Goal Orientation	1	0.25	.00	.62
Hypertext	1	11.61*	.14	.01
Interaction	1	0.00	.00	.95
Error	70	(7.53)		

Note. Values enclosed in parentheses represent mean square errors

* $p < .05$. ** $p < .01$.


The results of the ANCOVA analysis, as presented in Table 19, provide evidence of the influence of the hypertext structure on recall. While the mean recall levels in the Hierarchical condition are significantly larger than those observed in the Network condition (refer to Table 18), neither goal orientation, nor the interaction of the hypertext or goal orientation, significantly affected recall scores.

Knowledge gain vs. recall

The finding of no relationship between SK and the knowledge difference scores ($r = -.11, p > .1$) and knowledge gain and Lostness ($r = -.26, p = .09$) when the

⁹³In this case of the subject matter found in the hypertext.

Figure 7. The effect of goal orientation on recall in the Network and Hierarchical conditions




influence of hypertext condition is ignored ($N = 46$), does not appear to support the original contention that the integration of the structural and conceptual elements of a mental model of the hypertext facilitates learning. However, the fact that the correlation between Lostness and knowledge gain approaches significance is attributable to the strength of the relationship between these variables in the Network condition ($N = 12, r = -.53, p = .08$), as opposed to its weakness in the Hierarchical condition ($N = 34, r = -.17; p = .33$). This finding can be related to the observation that the absolute amount of knowledge possessed by participants following the study session in both hypertext conditions was significantly related to the Lostness metric ($N = 46, r = -.37, p = .011$). In other words, while the recall of factual knowledge presented in hypertext seems in general to be related to ones degree of disorientation, irrespective of how much of the material presented is new to one, the possession of background knowledge

affects how disorientated a person becomes when provided with the opportunity to use hyperlinks. This is suggested by Figure 7, where it is only in the representation of the network conditions that the knowledge levels prior to the study phase and those following it run parallel to one another.

University of Cape Town

Discussion

Figure 8. Revised model of relationships between goal orientation, navigation, mental models and knowledge acquisition.



Many of the aspects of the theoretical model which has been developed in this paper were supported by the results. The diagrammatic representation of this model has been re-presented in Figure 8, with minor alterations resulting from the findings made. Each of the relationships which it was not possible to demonstrate as having a significant impact in the theoretical framework constructed, for whatever reason, are represented by stippled lines.

There is evidence that goal orientation does affect hypertext navigation, with the caveat that as learning orientation was determined to be the most distinguishing characteristic between the participants with respect to traits which are associated with goal orientation, it is only possible to make this statement with regard to the presence or absence of learning orientation. Learning orientated individuals made more extensive use of the hyperlinks in the Network condition, presumably as a result of the greater degree to which they were inclined to expend effort. The effect of this form of goal orientation on hypertext navigation would have been amplified by their response to the disorientation they suffered, which, in line with the theoretical characterisation of learning orientation, would have taken the form of additional use of the hyperlinks. The finding that neither background knowledge nor levels of hypertext self-efficacy impacted on their navigational behaviour, was surprising. While the lack of the finding of an effect with regards to background

knowledge can be attributed to the poor reliability of the scale used, this is not the case for self-efficacy.

Learning orientated individuals were revealed, moreover, as navigating through the hypertext in a way which indicates that they possessed more knowledge of its structure, while at the same time demonstrating less behavioural symptoms of disorientation than those students who were not prepared to study in order to master the material. Significantly, these differences were only apparent in the Network condition as far as structural knowledge was concerned, while the degree of disorientation suffered within this condition was exaggerated relative to that experienced in the Hierarchical condition. This is consistent with the notion of the Network hypertext representing a more challenging learning environment, in which differences resulting from a person's motivational orientation are likely to be most apparent.

The argument that the relationship between hypertext navigation and both disorientation and recall is mediated by a person's mental model, was difficult to verify. This was primarily due to the poor reliability of the conceptual knowledge measure in determining the role of mental models of hypertext characterised as being composed of both knowledge of the structure of a hypertext and of the conceptual relations embedded within it. However, by employing disorientation as a surrogate measure, it has been possible to interpret the significant negative correlation between structural knowledge and disorientation as supporting the notion that mental models are especially important when hypertexts contain hyperlinks. This was as predicted, given the greater non-coherence which these hyperlinks introduce into a hypertext, and which would result in the greatest displacement of hypertext structure and the conceptual relations of its content.

The diagrammatic indication of a relationship between hypertext disorientation and factual knowledge gain was partially supported by the finding of a nominally significant relationship between disorientation and change in knowledge levels, subsequent to the completion of the reading task. This relationship was once again only apparent in the Network condition. Participants who showed evidence of being lost in hyperspace were less able to recall information that they had read in the hypertext. The observation that students were better able to recall information in the Hierarchical than the Network condition, can be explained in terms of the greater degree of cognitive overload which they were likely to experience in the

latter condition. The question as to whether disorientation causes poorer recall, or vica versa, as well as the extent of the influence of cognitive overload on recall levels, will be discussed in greater depth in a later section.

The difficulty of separating discussion of the measures used from phenomena which they were designed to measure, and for constructs which have been operationally defined, such as conceptual knowledge, the impossibility of doing so, has led to the strategy of using the examination of these metrics as the context for the discussion of related issues. This mode of presentation will provide the vehicle by means of which the nature of the constructs introduced in this study, as well as their relationships to one another, can be explicated. Moreover, it will serve as a springboard for suggestions as to alternative methods which might be employed in research into the same types of questions as those addressed here. Finally, this section will close with a discussion of the limitations of this study, as well as potentially productive avenues of future research which have emerged while conducting it.

Before commencing, it should be noted that the discussion shall contain a speculative element, particularly with regards to the section on mental models. This was unavoidable, both because of factors which have made the determination of the role of mental models problematic, and as a result of the exploratory nature of a large component of the study.

Goal orientation

The validity of goal orientation as a construct was verified through the results obtained in this study. The dimension of learning orientation was revealed as important in discriminating between individuals with respect to their performance on measures associated with goal orientation. The decision to focus on two as opposed to the results of the three cluster solution, although convenient in terms of sample size, received additional justification from the finding by Meece (1993), that 40% of the work avoidant group would have been classified as performance orientated according to the traditional median split method of identifying goal orientation groups.⁹⁴ However, although the speculation that it is fear of failure which translates into work avoidant behaviour (Meece et al., 1988) could be

⁹⁴This is substantially larger than the 25% probability that these participants would have fallen into any one of the four groups into which students would have been divided on the basis of the median split procedure.

interpreted as supporting the notion of convergence between the performance orientation and work avoidant clusters, as the former are theorised as being particularly prone to anxiety, it was the learning orientated individuals who demonstrated greater anxiety in this study. This anomalous finding was revealed as caused by the influence of a work-avoidant motivational group, whose low levels of anxiety and whose memberships within the performance orientated cluster in the two group solution resulted in the evidence of apparent greater relative anxiety amongst the learning orientated individuals.

The greater relative constancy of the learning orientation cluster when converting from the three to the two cluster solution can be interpreted as evidence of the importance of the presence of learning orientation in distinguishing between groups of people. Ames (1988) discovered that perceived learning orientation, as mediated by classroom instruction, was a significant predictor of positive attitudes towards the class, the number of learning strategies employed, and the school students' proclivity for challenging tasks, irrespective of whether they were classified as performance orientated or not. Learning orientation has, moreover, been identified by Elliot (1988) as overriding the effects of perceived competence, while this is not the case for performance orientated individuals.⁹⁵ This finding may be attributable to the belief in the malleability of intelligence amongst those who were learning orientated, as it is this fundamental cornerstone of the belief system which differentiates learning from performance orientated individuals which has been demonstrated as overriding perceptions of self-efficacy (Braten & Olaussen, 1998). Finally Dweck (1986), in her original conceptualisation of goal orientation, hinted at the primacy of learning orientation by indicating that self-efficacy would influence performance orientated individuals, while those who were learning orientation would remain impervious to its effects. This may partially account for the fact that while the observations in other studies that female participants are more performance orientated than men (Dweck, 1986) was not replicated here, their lower levels of self-efficacy may fail to protect them from the worst effects of performance orientation.

Although there is evidence of the dominance of learning orientation with

⁹⁵Although the performance orientated subjects with a high degree of self-efficacy in the same study still tended not to take advantage of learning opportunities which they presumably felt could expose their failure to others.

regards to goal orientation, its importance relative to performance orientation is not quite settled. Methodological shortcomings undermine the confidence with which one can claim prominence for learning orientation in the Ames study. The conflation within the goal orientation scale used, of the emphasis on learning versus performance in the classroom, with statements reflecting on how performance or learning orientated the students themselves were, jeopardised the validity of the goal orientation construct used. The generalisability of the findings are also called into question, given the homogenous nature of their advanced grade 8 to 11 student sample base.⁹⁶ Moreover, the study which called the appropriateness of splitting goal orientation scores on the median into question (Meece & Holt, 1993), found that the achievement scores of those subjects who demonstrated "mastery" goals were dependent on whether they possessed "ego-social" goals as well.

The fact that there is contention surrounding the nature of the relative importance of different aspects of goal orientation for learning indicates that the finding in this study, that learning orientation is the more influential dimension, is non-trivial. Recognition, however, needs to be given to the possible effects of the encouragement of participants to exploit the hyperlinks in distorting the results which have been attributed to goal orientation. Although goal disorientation has been characterised as a relatively stable personality trait, studies conducted to determine the degree to which it is affected by the instructional context, such as a classroom (Ames & Archer, 1988; Elliot & Dweck, 1988), betrays the belief that it can be modified by the environment in which a person finds him/herself, at least on a temporary basis. There is no reason to believe, however, that the encouragement provided the participants reduces the validity of the interpretations given of the effects of learning orientation, as these instructions were given to both those who fell in the learning oriented group, as well as those who fell without.

The effect of goal orientation on hypertext navigation

The two-pronged approach taken in determining the nature of the relationship between goal orientation and hypertext navigation appeared to pay dividends. While an examination of the pathways followed through the hypertext did not clearly discriminate between orientational groups in a manner consistent with their

⁹⁶Though refer to the study for reasons given as to why the nature of the sample was interpreted as strengthening the conclusions reached.

performance on variables associated with goal orientation, the top-down examination of differences in the navigational patterns of the LO and NLO groups was more informative. The finding that those individuals who are learning orientated follow a more densely inter-connected path through a hyperlinked hypertext is attributable to their greater exploitation of those hyperlinks. This is most probably due to their greater willingness to exert effort, as it was this variable which was most significantly different between the orientational groups in the two cluster solution. Although the goal orientation theoretical framework would also predict levels of anxiety to be pivotal in the extent to which individuals employ hyperlinks, a finding supported where three orientational groups were identified, the actual results obtained were based on the two group scenario, and prohibits one from drawing any conclusions with respect to the possible influence of anxiety.

The measurement of navigational pathways

The finding that link density was demonstrably different for members of the different orientational groups can be interpreted as supporting the proposition that learning orientated individuals make substantially greater use of hyperlinks in their navigation. In contrast to expectations, however, the metrics of path stratum and compactness were not significantly different for learning orientated individuals across both conditions. The interpretation of this finding should be approached with caution. These metrics have only been applied twice in the past in the context of user navigation in a hypertext, and in both cases using the same hypertext system (McEneaney, 1999, 2001). Moreover, there are a number of issues surrounding their use which need to be considered.

One of the major areas of concern revolves around the drastic change in the nature of the relationship between path compactness and stratum which occurs when the number of nodes in the user's path is substituted for normalisation purposes for the number of nodes in the hypertext itself. As noted by Botafogo (1992) in his original discussion of the metrics of stratum and compactness within the context of the analysis of hypertext systems, this is an issue which requires further research. Although this study meets both conditions which McEneaney (2001) lists as providing reason to use the number of nodes accessed as the normalisation measure, namely that the Hierarchical and Network hypertexts are the same size, and that the study is interested in differences on the level of

participants, there are good grounds to argue that the total number of hypertext nodes would be a more appropriate normalisation measure. While not an issue in this study, as there was no evidence that goal orientation affected the number of distinct nodes visited in the hypertext, the comparison of measures which are insensitive to differences in the number of nodes accessed, such as is the case when the user path node count is used for normalisation, would have little power in detecting differences when examining a construct which in other circumstances might be predicted to influence the number of nodes the participant exposes him/herself to.

Density and link density provided more intuitive measures than path compactness and stratum, in that they consist of a ratio of the total number of links and hyperlinks which are available to those which have been used, respectively. Link density, in particular, might be considered a more appropriate measure than path compactness within the context of this study, given the focus on hyperlink usage. Although the extent to which hyperlinks are made use of would be predicted to correlate strongly with the inter-connectedness of their path through the system, this was not the case. This could be attributable to the modular structure of the hypertext used in this study, which ensures that the majority of the hyperlinks in the system lead to pages within the same module. The use of the fewer number of hyperlinks between modules would accordingly have a disproportionate influence on the measure of inter-connectivity, and it is this weighting of these particular hyperlinks which could account for the deviation between the density and path compactness scores. Such discrepancy might be a useful indicator of the willingness to explore simple hierarchical hypertexts, with content of a self-contained nature, as the utilisation of hyperlinks which lead to other sections of a hypertext could safely be regarded as intentional in systems in which it is simple to predict where a hyperlink might lead on the basis of its immediate salience to the topic discussed in a particular node. This was, however, not the case with this hypertext employed in this study, in which deliberate efforts were made to establish less coherent and predictable hyperlinks between modules. An additional advantage to using density scores over path compactness and stratum is that the upper limit for these measures is dictated by the structural relations within the hypertext, and as it is trivial to calculate their value as a percentage, they do not require the introduction of an arbitrary normalisation measure in obtaining a value which falls between 0 and 1.

The issue of mental models

The concept of a mental model, as used in this study, has been characterised as consisting of both a structural and a conceptual component. It is difficult, however, on the basis of the results obtained in this study, to determine whether the proposed integrative conception of mental models has greater explanatory potential than the naive conception of mental models of hypertexts as being composed in their entirety of knowledge of where the nodes lie in relation to one another. While this research program observed significant differences in the interaction of structural knowledge and disorientation across hypertext and goal orientation conditions, the substitution of the Lostness measure for the measure of conceptual change, due to the latter's lack of reliability, left the conceptual dimension of mental models largely untested.⁹⁷ The results obtained are as consistent with the argument that the more people exploit hyperlinks, the better their structural knowledge of the hypertext, and accordingly, the less likely they are to suffer disorientation, as with the integrative characterisation of mental models.

The impossibility of obtaining direct access to the nature of the mental model construct has, moreover, led to the characterisation of mental models in general terms, such as "well-developed", "flexible", "accurate", "implicit", "under-specified" and "abstract", which are by themselves not easy to validate or disconfirm. This indirect characterisation of mental models implies that it is only by means of the concept's utility as an explanatory mechanism in binding together other theoretical constructs that the integrative account of mental models can be evaluated. In particular, as diagrammatically represented in Figure 8, it needs to be able to establish a causal connection between hypertext navigation and disorientation, and hypertext navigation and recall. There is reason, on the basis of the results obtained, to believe that the integrative conception of mental models is better able to accomplish this than the perception of mental models as mirroring the structural layout of a hypertext.

The findings indicate that the hypertexts employed in the Network and Hierarchical conditions differ from one another in a manner which has implications for the type of mental model best suited for both their navigation, and their use as

⁹⁷The inclusion of the Lostness metric was justified by its conceptualisation as combining the influences of both knowledge of the structure of the hypertext, as well as how the concepts within it relate to one another.

an educational medium. Evidence of a relationship between the recall of factual information and lostness is only found for the network hypertexts, a result which is particularly significant, given the definition of the latter as encompassing both the structural and conceptual aspects of mental models.⁹⁸ In lieu of a reliable measure of knowledge of the conceptual relatedness of terms in the hypertext, the Lostness measure has been employed as a stand-in for the conceptual knowledge measure in determining the extent of development of a participant's mental model.

The finding of a correlation between a measure which has been characterised as consisting of a combination of structural and conceptual knowledge, and recall, provides some support for the contention that it is the integration of these two which facilitates recall. Although no detailed mechanism via which this effect may occur has been provided, two possibilities have been mentioned. The first is the notion that a well-developed conceptual model of hypertext material provides a multitude of routes between knowledge nodes, thereby reducing recall to the navigation down one of these pathways (Britton & Gulgoz, 1991). This is a more naive version of Kintsch's associationist hypothesis (1998) in which the integration of new and background knowledge allows for the stronger connection established between knowledge nodes to increase the probability that the activation of one node will lead to the recall of the other. Another possibility is that the structure of the hypertext serves a mnemonic function in a similar fashion to the evidence cited by Dillon (1995), of the memory of information in a page being facilitated by memory of its location. These opposing scenarios will be discussed in greater depth elsewhere.

The observation that learning orientation made more than a 100% difference in structural knowledge scores obtained by participants in the Network condition, as well as exaggerating the gap between the disorientation scores found between

⁹⁸ An alternative explanation for the fact that a relationship between lostness and knowledge was only found in the network condition needs to be considered. The lack of such a relationship in the hierarchical hypertext could be a consequence of differences in the causes underlying the degree to which nodes are revisited in the different hypertexts. In the Hierarchical hypertext, students are constrained by the sequential arrangements of nodes to either move back along the path they have come if they wish to return to a previous node, or alternatively, to return to the modules base-page, both strategies which would result in a high frequency of the re-occurrence of nodes in a student's pathway. The revisitation of nodes in the Network hypertext, on the other hand, might more appropriately be characterised as resulting from an inability to remember where one has been, or where particular nodes lead. If this is indeed the case, then the Lostness measure could be considered a more accurate indicator of hypertext disorientation in the Network than the Hierarchical hypertext conditions.

orientational groups in the Hierarchical condition, provides evidence that the extent to which a mental model is of assistance in a hypertext is largely dependent on one's motivation for studying in the first place. This is due to the fact that it is in the Network conditions in particular that conceptual-structural knowledge integration is required, and in which the factor which most clearly distinguished orientational groups in this study, namely willingness to exert effort, would be expected to make a difference.

That the conceptual component of one's mental model of a hypertext should possess greater importance when hyperlinks are present, is consistent with the definition of conceptual information as the semantic value acquired by a term through its co-occurrence with other terms.⁹⁹ It is exactly the co-occurrence of terms usually not associated,¹⁰⁰ through their connection to one another by means of hyperlinks, which is the avenue through which non-coherence is introduced into a hypertext. The less coherent a hypertext, the greater the conceptual change required on the part of the user to re-associate terms that are linked in such a way as to bind the hyperlinked texts to a topic which best preserves the nature of the conceptual relationship. The more successful the hypertext user is in accomplishing this, the more effective the topic would stand as a cue for recall, thereby accounting for the relationship observed between disorientation and recall.

Disorientation

Although Otter and Johnson (2000) claim that Smith's measure of lostness is the only published measure of hypertext disorientation, this is not entirely accurate. There have been numerous less rigorous formulaic expressions of hypertext disorientation. These generally take the form of a ratio of the number of nodes through which the user must pass in order to complete a task to the number which they actually visited (Chou & Lin, 1998; Calvi, 1997; Dias & Sousa, 1997; Stanton et al., 2000). This formula provides a more valid measure of search efficiency than disorientation, and accounts for the fact that it is similar to the nature of the formula used in this study to determine structural knowledge levels, as this type of knowledge would be expected to greatly assist in finding information in the

⁹⁹A definition which matches the operational definition of semantics as applied by LSA

¹⁰⁰A "term" is used here as a grammatical entity which is identified with its semantic value, thereby allowing for the possibility that identical words are considered distinct terms, when used in different semantic contexts

hypertext. That it might not be adequate by itself is hinted at by the finding that the effects of the provision of graphical browsers are felt most keenly with regards to search efficiency (MacDonald & Stevenson, 1998), but not search effectiveness (Boechler, 2001). The finding of information of relevance to a particular individual is only likely to be aided through the individualised mapping of knowledge to the hypertext structure.

Given the definition of the statistic as non-directional, it is not possible to determine on the basis of the correlation between lostness and recall alone, whether hypertext disorientation affects recall, or *vica versa*. It is conceivable that the failure to integrate conceptual and structural information accounts for the inability to recall factual information, in a similar fashion to that proposed by Kintsch with regards to the integration of background knowledge and new information, or alternatively, that the memory for the hypertext's factual content may be tied to memory of its structural layout. The isolation of this effect to the Network condition does not help distinguish between these two possibilities, as both may in fact be facilitated by the extra degree of processing which this condition required.

Learning and Recall

The observation that participants are significantly better at recalling factual information within the Hierarchical than the Network conditions is consistent with the finding in the CFT research (Jacobson & Spiro, 1995). The possibility that participants in the Network condition experienced a greater degree of cognitive overload than their counterparts in the Hierarchical condition would account for this difference. Cognitive overload has been characterised as a factor which can seriously circumscribe the educational potential of hypertext which allows one to "criss-cross the conceptual landscape" (Niederhauser, Reynolds, Salmen, & Skolmoski, 2000; Kaplan, 2001). Given the fact that the hypertext content used was extracted from a post-graduate course, it is probable that the students possessed very few cognitive resources in reserve to those used in reading it. The additional processing that it has been argued is necessitated through the introduction of hyperlinks could have over-taxed the abilities of the students, especially given the limited time period they had in which to process the material, and the fact that the study session occurred at a time of the day, and year,¹⁰¹ at which the students were unlikely to be at their

¹⁰¹Refer to the comments on page 122 for further explanation.

mental best.

Limitations

There are numerous limitations to this study, which despite many of the theoretically consistent results which have been obtained, suggest that it should be regarded as possessing utility primarily in terms of the avenues of future research it suggests. A blatant shortcoming was that of the small sample size employed. This had implications both for the resolution at which the theoretical constructs of interest in this study could be examined, as well as the statistical procedures which were employed in analysing the data. The decision to focus exclusively on the distinction between the LO and NLO groups of students, as opposed to the three clusters whose constitution was more consistent with the literature, was partially an attempt to compensate for the sample size.

A low sample size was the most important factor mitigating against the use of a technique known as Structural Equation Modelling (SEM) to analyse the data. SEM is particularly well - suited for the analysis of the validity of theories such as this one, which postulate the existence of a complex arrangement of relationships between a multitude of latent variables. Through a procedure of estimation and fitting, it determines whether the relationships which have been theorised to exist match the actual pattern of covariation as observed in the data collected. If the theoretical model and the data are consistent with one another,¹⁰² one can consider the model a candidate "explanation" of the data. Although SEM theorists are quick to point out that even models which fit the data can only be considered one of many possible reflections of the underlying reality (Hoyle, 1995), it is regarded as perfectly acceptable to put forward a number of alternate models, prior to the running of the analysis, so as to determine which conceptualisation of the relationships between the variables one is examining fits the data most accurately. It is for this reason that SEM can be considered a useful exploratory tool.

On the face of it, SEM thus appears to be ideally suited as an instrument for the verification of the validity of the model that has been proposed. It would have had the potential to discriminate between alternative versions of the model, and

¹⁰²Whether or not one accepts a model as providing an adequate fit of the data depends on conventions associated with the particular fitting algorithm one employs. SEM can therefore not be regarded as a statistical procedure, as it does not discriminate between models on the basis of probability levels.

might have assisted in determining whether it was necessary to postulate that mental models are composed of both conceptual and structural elements, through which they exert their influence over disorientation and recall, or, alternatively, whether it is sufficient to argue that the density of an individual's path affects their knowledge of the structural layout of the hypertext, which consequently determines how disorientated they are. Other possibilities which could be tested is whether disorientation is influenced through both a mental model's conceptual and structural components, or whether they exert control indirectly via their effects on levels of recall, or whether both routes of influence operate simultaneously. Finally, it would have provided insight into an aspect of the model proposed which it has not been possible to delve into in any great depth, namely the possible existence of feedback loops between the elements of disorientation, mental models, and hypertext navigation.¹⁰³ These would be of the most interesting facets of the model which has been proposed, as they would lend an iterative and dynamic aspect to the process of hypertext learning and navigation. As it is, on the basis of the analyses conducted, it is only possible to state that goal orientation most probably operates on navigation directly, as well as via its effect on the mental model formed, with its attendant consequences for the degree of disorientation experienced, and the reaction to that disorientation in further exaggerating the navigational behaviours observed.

Unfortunately, unless certain special conditions are met, SEM is generally only considered appropriate for sample sizes which fall in the hundreds. The sample size requirements become even more stringent for the algorithms which have been employed to overcome non-normality of the data (West, Finch, & Curran, 1995),¹⁰⁴ a feature of certain of the variables used in this study. The problem of the limited total number of participants in this study is compounded by the fluctuation of the sample size as a result of the particular comparisons which are made. The parsimony of SEM has therefore had to be substituted by a number of separate analytic procedures, with the attendant danger of committing a Type I error.

Another factor which may be construed as reflecting negatively on this study, is

¹⁰³ Although see MacCallum (1993) for a discussion of the exponential increase in the numbers of equivalent models which emerge with the theorisation of a bi-directional influence between two components of a model.

¹⁰⁴ Although West, Finch and Curran (1995) do suggest that transformations of the data might allow one to employ the more conventional fitting and estimation algorithms.

that of the relatively untested nature of many of the procedures employed. This is particularly problematic from the point of view of validating the results obtained through their comparison with the findings of other studies. While every attempt was made to ensure that the measures used in this study were not of an entirely ad-hoc nature, the relative novelty of the area of hypertext research, in combination with the myriad of systems which have been developed, has resulted in a paucity of standardised instruments from which to choose. The magnitude of this problem was revealed by Liao (1999), who, in a meta-review of the educational benefits of hypermedia systems, discovered that 75% of the instruments used were developed specifically for the studies in which they were applied.

The theoretical framework within which the hypotheses tested were developed also necessitated the development of novel forms of data collection. A case in point is that of the word-sorting task, in which the constructivist aspect of the model that has been proposed dictated that the conventional task of sorting words into pre-conceived categories be rejected in favour of a more free-form sorting procedure. The exploratory nature of substantial portions of this study, was, moreover, a significant factor in determining, as well as justifying, the degree to which the measures used were specific to it. Exploratory analyses, by their very nature, are not designed for purposes of replication or cross-validation.

The effect of the period of time over which the experiment was conducted needs to be considered, as well. The student spent on average only 63 minutes reading the material. This may not have been long enough to allow the degree of processing of the material needed for learning on a conceptual level to take place. This is especially the case for the acquisition of incidental knowledge, or knowledge which students were not asked to memorise in order to answer particular questions. It is this form of knowledge which the students in this study would have acquired, and which research indicates would be poor for hypermedia education in a comparable time frame (Astleitner & Leutner, 1995). It was, moreover, only possible to schedule the tutorial sessions for a stage during the year in which the students had a number of other academic commitments.¹⁰⁵ It seems likely that they would not have been in a mental or psychological state conducive to the degree of concentration required for the processing of additional learning materials. Furthermore, the sample size

¹⁰⁵Many of the students were required to write two tests in the same week that they attended the second experimental session.

employed was not sufficiently large to compensate for this by providing greater power to the measure of conceptual change to detect a change. Finally, it is probable that the contribution to learning of any dynamic processes of feedback which occurs between mental models, disorientation, and hypertext navigation would only become evident with research of a more longitudinal nature.

Finally, there is some question as to the degree to which one can generalise the findings obtained in this study to other research programs. Participants who participated in this study, were all registered for a second year Psychology course in a tertiary educational institution, and as such constitute a fairly limited demographic cross-section of society. While this is true of most, if not all, of the studies reviewed in this paper, and while the ethnic composition of the sample employed in this study was in all likelihood more varied than that found in the majority of studies referred to,¹⁰⁶ it is true that in research into the influence of goal orientation, the sample most commonly used consists of primary or secondary school children. Whether or not the fact that participants in this study have undergone a more rigorous process of selection than most, and represent individuals who should all be relatively competent learners, alters the role that goal orientation plays with respect to how they learn, is a topic requiring further investigation.

Measure of conceptual knowledge

There are reasons, in addition to those mentioned, which could account for the failure of this study to detect conceptual change. The decision to provide participants with greater flexibility in the conceptual categories that they used to sort the words provided the means of determining the understanding that people gain from a hypertext with regards to the conceptual relations between more abstract concepts. While this helped separate conceptual knowledge gain from that measured by tests of factual knowledge, a necessary step given the dissociation observed in the literature between the influence of hypertext characteristics on different knowledge types (Jacobson & Spiro, 1995; Kaplan, 2001), it has introduced certain methodological complications. The lack of a priori categories mitigated against the possibility of alternative means of determining conceptual change, in the face of the failure of LSA for this purpose. This is particularly the case given that no systematic approach was taken in determining which concept terms should be

¹⁰⁶Most of which have been conducted in the United States.

included in the list.

A different and possibly more effective approach would have been the separation by “experts” of concept terms into a number of categories, prior to the study, and the identification by those experts of the categories they had used. A small number of the most salient categories, as determined by correspondence of categories across experts, could subsequently have been extracted, with each expert being instructed to select a certain number of concept terms which most clearly belonged in each category. The degree of inter-rater agreement for terms selected for particular categories would serve as the basis for keeping a smaller sub-set of the categories, and the word selection - category selection process could reiterate until a suitable degree of inter-rater agreement, or number of categories, had been arrived at.¹⁰⁷

While the procedure outlined above might be suited for material which falls within a single knowledge domain, the deliberate selection of multi-disciplinary content matter for the hypertext employed in this study severely reduced the chances of finding people who were knowledgeable in all of the disciplines included. It was therefore, for all practical purposes, not possible to determine which words were the most coherently related across domains prior to the actual application of the LSA procedure. Moreover, the effectiveness of the LSA procedure was itself sensitive to a number of factors, one of the most important of which was that it was trained on appropriate material. It was doubts as to the ability of the generic WebLSA database to capture the associations between words in the somewhat more specialised hypertext material, which motivated the development of a customised LSA training set. This was a time-consuming procedure, involving as it did the collection of appropriate text, as well as the removal of variations in spelling caused by differences of dialect. Articles composed in American English had to be converted to the South African English spelling, prior to their input into the training database.

Faith in the accuracy with which the CustomLSA dataset provided a true reflection of the contextual meaning of terms used in the hypertext was compromised by a number of factors. For one thing, as is evident from an examination of the readings used (Appendix D.2), much of the training material took the form of published articles. Their inclusion was judged necessary given the academic nature of the material within the hypertext itself. The standard of

¹⁰⁷This procedure has the advantage that one could trade off the inter-rater agreement against the number of categories most suited to one's research

academic english in which such articles are conventionally written would be predicted to pose problems in the LSA training process, given the relatively large vocabulary used. Synonymy, or the representation of the same concept using multiple terms could severely handicap a process which relies on the contextual co-occurrence of particular terms, despite the fact that the higher-order relations detected by LSA should offer some degree of protection against this (Foltz, 1996). Inversely, the nuanced uses of a term imbues it with multiple meanings, a phenomenon referred to as polysemy, such that a term may occupy an ambiguous position within semantic space (Landauer & Dumais, 1997). These observations are as pertinent to the actual reading material presented in the hypertext as to the articles which constituted the text corpus, and might be regarded as representing an argument for having used content which was both more accessible, and which could more easily have been pigeon-holed into a specific academic discipline. This would also have had the advantage of separating the effects of the introduction of hyperlinks on the coherence of the hypertext from a lack of coherence brought about through the presence of multi-disciplinary content.

While the failure to detect a conceptual change could be interpreted as evidence that no such change took place, when all of the factors which have been mentioned are considered, it is perhaps not surprising that the word-sorting task did not provide any indication of conceptual change following the reading phase of the study. Researchers who wish to replicate this study in the future might consider constructing a hypertext around material extracted from a particular discipline. They would then be able to focus exclusively on the effect of non-coherence resulting from the introduction of hyperlinks on conceptual knowledge. The discipline of Psychology would be a good choice, given the availability of LSA datasets based on psychological texts.¹⁰⁸

¹⁰⁸Obtainable from the LSA @ Boulder Website.

General discussion

It has been proposed that mental models, by their very nature, are a form of knowledge representation which is well suited as a means of gaining insight into hypertext education and navigation. It would perhaps be advisable at this point to consider some of the reasons why they represent a superior explanatory mechanism for this purpose than their closest theoretical rival, the cognitive map. Cognitive maps are the mental representations which it has been assumed underly effective navigation of a hypertext system (Boechler, 2001), with their reconstruction interpreted as an adequate identification of the knowledge a student possesses of the hypertext structure and the material contained within it. Defined as the “internalised analogy in the human mind to the physical layout of the environment” (Chen 1998a, p 80), they are only appropriate representations of hypertext in cases in which students are presented with simple, hierarchically structured hypertexts, in which there exists a one-to-one mapping between the nodes of the hypertext and the concepts presented in it. It is exactly these types of hypertexts which have been the focus of concept mapping tasks (Chou & Lin, 1998; MacDonald & Stevenson, 1998).

Concepts maps fall short as an explanatory tool in hypertext education, however. The assumption made that semantic relationships are adequately captured by the placement of hypertext nodes is dubious in the extreme, except possibly for the simplest of hypertexts (Dillon et al., 1993). Not only does it side-step the prickly issue of the applicability of spatial metrics of similarity to the 2 dimensional arrangement of nodes, but the bodies of knowledge which are imparted in online courses are complex enough that links between concepts are not likely to be entirely unambiguous. This is especially the case when multi-disciplinary material is employed, as a particular concept’s semantic connotations will tend to fluctuate across the conceptual boundaries established by the specialised areas of knowledge included in the hypertext.¹⁰⁹ This complexity is increased manyfold by the introduction of hyperlinks, as the links themselves can be considered to convey information (Landow, 1990).

Once the conceptual richness of hypertext material reaches a point where the non-alignment of the hypertext’s structure with its content becomes a possibility, it

¹⁰⁹Refer to section for a discussion of the implications of this on the LSA - derived word - similarity ratings.

is necessary to substitute cognitive maps with a construct such as the mental model, which has the capability of accommodating both structural and content-based forms of knowledge. Given that the ambiguity of links is a function of the background knowledge of the hypertext user, in addition to the characteristics of the hypertext, researchers have only been able to avoid problematising the relationship between the structural layout of a hypertext and its content by not taking background knowledge into account. Ironically, it is exactly those participants who have little background knowledge who appear to benefit the most from the provision of graphical maps and systems which are hierarchically structured (Balcytiene, 1999). As Spiro and his colleagues have demonstrated, however, there is danger in over-simplifying the presentation of ill-structured material in order to make it digestible to introductory students. To do so could lead to a species of cognitive distortion referred to as reductive bias (Spiro et al., 1996), which only serves to impede subsequent comprehension of more advanced material.

It is proposed that the same phenomenon occurs with respect to navigation and comprehension of hypertext material. The simplistic representation of a complex domain yields cognitive maps which not only hinder more in-depth understandings of the material, but due to their non-congruence with the multiple meanings imbedded in the nodes and the relationships between them, actually accounts for the frequency with which novice students are reported to suffer from disorientation. It is proposed that the failure of students to resolve the ambiguities that they encounter prevents their ability to consolidate the information within a mental model, and as Johnson-Laird (1990) has argued, results in the knowledge representation reflecting the propositional form in which it is presented. It is this tying of the student's understanding of the material to the form in which it was encoded (Jacobson & Spiro, 1995) which explains why they are less able than students who have been exposed to hypertext systems which take full advantage of the medium's hyperlinking capabilities, to transfer knowledge they gain to novel situations.

This argument, moreover, has implications for the theorisation of recall as well. The suggestion that recall in a hypertext may be facilitated as a result of the proliferation of pathways between nodes is circular, as these pathways themselves need to be remembered. It is argued instead, that the development of the associational complex which emerges from the integration of newly acquired information (which may include general survey-type knowledge of the entire

hypertext structure) with background knowledge would have to take place in some degree first, before active recall of the pathways is possible. With the establishment of a mental model, retrieval of information contained in the hypertext could be accomplished by means of two routes. The connectionist aspects of the Construction-Integration theory would predict the automatic recall of information which is strongly associated with a recall cue, with a more tightly integrated situation/mental model enabling more concepts or terms to serve as cues. Knowledge of the particularities of a specific hypertext's structural layout of the hypertext may then be built on top of this, with one's own reading of the semantics of over-determined hyperlinks enabling one to tag the relative positions of the nodes in a manner which does not require the additional cognitive burden of remembering the hypertext structure as presented.

The requirement that the formulation of mental models takes place before recall of ill-structured or hyperlinked hypertext is facilitated, is consistent with evidence provided by Graesser (1980) that while recall for information which is inconsistent with a stereotypical script is facilitated immediately after that information is presented, when the recall task is delayed, information which is more consistent with the script is more easily recalled. This allows one to argue that textual systems which are less coherent therefore offer the inconsistencies which could bootstrap the representation of knowledge to a state in which it benefits learning, but only to the extent students are motivated to process that inconsistent information in the first place. This accords well with the argument that inferences are remembered when they provide for the resolution of inconsistencies, but not otherwise (Obrein Myers, as cited in Albrecht & O'Brein, 1993) , as it is this process which stands as a pre-condition for mental model formation.

The pay-off of making the studying experience more difficult in the short-term is that it facilitates learning in the long-term. While this finds support in Reynold's observation (as cited in Beasley, 1995) that participants who used a hypertext system which only possessed hyperlinks, while finding the experience more frustrating than those students who had a map of the system at their disposal, were better able to recall the information presented to them, it does imply that students are initially likely to experience cognitive overload. Cognitive overload has been a motivating factor behind the construction of many of the navigational aids which have so proliferated. It provides, in addition, a likely explanation for the finding in

this study that students who were exposed to the simpler hypertext system possessed greater recall of the facts contained within it. Jonassen (1990) found that students who were required to generate the type of relationship signified by links in a hypertext had lower levels of recall than those who had the relationships provided to them, a finding which parallels what was observed in this study, especially given the fact that participants in the Network condition were asked to try to determine why nodes were linked to one another. However, while cognitive overload will probably always be a factor in hypermedia learning, it is less of an issue when one substitutes the traditionally endorsed view of hypermedia learning as instantaneous, with the conception of it as a more gradual, dynamic process.

The integrative theory of mental models also offers an account of mental models which is more consistent with the archetypal characterisation of hypertext disorientation as provided by Edwards and Hardman (1989), than is their equation with cognitive maps. More specifically, the description of disorientation as a state of not knowing where to go next in the system, accords with the effect of the lack of structural knowledge in reducing one's ability to locate facts in the hypertext, at least with respect to an information retrieval scenario. The conceptualisation of being lost in hyperspace as further involving knowing where to go but not how to get there, and not knowing where one is at any particular time in relation to the entire hypermedia system, is related to route and survey knowledge, respectively, and would be expected to be reduced through the alignment of hypertext content with structure. The greater ability of hypertext users with well developed mental models to engage in the rote memorisation of the routes between nodes, in combination with the facilitation of survey knowledge through the re-establishment of coherence between general topics and the concepts linked in a hypertext, suggests that the integration of conceptual and structural knowledge is required to account for these aspects of disorientation.

Future directions for research

The theoretical model as outlined, in keeping with the nature of models in general, represents a gross over-simplification. Proper cognisance of the dynamic aspects involved in learning and navigation within a hypertext would merit a research project of its own. A central issue which has emerged within this study is the relationship between mental model formation and recall. The contention that

the integration of conceptual and structural information is a gradual process, carries with it the implication that the benefits it confers to recall might only make themselves apparent over a relatively long period of time. A possible avenue of future research into these dynamics would involve the identification of cliques, or ensembles of nodes within which each node is fully inter-connected to every other.¹¹⁰ By comparing the cliques within the hypertext to the constellations visible in the grouping of concepts by the students, it should be possible to trace the development of student's mental models as a function of the hypertext. This methodology would allow one to test predictions made concerning the convergence between a person's mental model of the conceptual relations embedded in a hypertext with the spatial layout of the hypertext. It would also provide one with a means of quantifying the relationship between the qualities of a mental model and recall.

Additional research would also be productively applied to the issue of the extent to which cognitive overload is a product of the fact that navigation and comprehension/learning within a hypertext are based on the same visual modalities (Boechler, 2001). Work on systems which attempt to assist navigation by means of auditory representations of depth, such as the current development of a system by Microsoft (Czerwinski & Larson, in press), in which the levels of a hierarchical website in which a person finds him/herself are denoted by means of variations in the pitch and number of times a tone repeats itself, can be seen as being based on the notion that limiting interference between the sensory modalities used for navigation and reading might reduce the combined cognitive demands made by these tasks. However, consideration has to be given of the evidence that the use of different modalities is only advantageous to the extent that their features co-operate in the fulfilment of a particular task (Schonotz & Bannert, in press), as mental model theory claims that mental models are independent of the perceptual qualities of the information represented.

A final issue which merits further attention, as it is pivotal to the design of many computer-based educational systems, is that of the role of metaphor in hypertext (Boechler, 2001). Metaphors can be described as "the poor man's mental model" in their assistance of the comprehension and utilisation of novel phenomena

¹¹⁰This requirement can be relaxed somewhat through the use of K-cliques, or groups of nodes which are required to have a certain minimum (K) number of connections to one another. The interested reader can refer to (Scott, 2000) for further discussion of these metrics.

through reference to experiences which are familiar.¹¹¹ However, in common with all static representational devices which are used as a crutch in aiding novices gain expertise in a new area, such as graphical browsers, the drawbacks in their use often outweigh the benefits. It is not always apparent to the novice how the representations being compared relate to one another, while the attributes which are not shared between the phenomena compared often only serve to confuse the user. For instance, the recognition that novice hypertext users frequently employ the schema of a book in their initial dependence on indexes and table of contents in a hypertext (Rouet & Levonen, 1996), can be seen as maladaptive in the extent to which it prohibits the user from taking full advantage of the navigational possibilities offered by the newer technology.

In the same way, the unexamined acceptance of the metaphor of hypertext as a hyper-"space", has also proved maladaptive, in that it has led to the development of systems of representation which have been assumed for decades as providing invaluable assistance to hypertext users, despite the recognition that psychological and hypertext notions of space do not adhere to the criteria which have been set for similarity metrics (Tversky, 1977). This has led to the complacent assumption that spatial distance, and the unchanging representation of that space in the form of a map (Calvi, 1997), is an appropriate visualisation of the semantic and conceptual relations in a hypertext, a notion which, it has been argued, is highly problematic. It was the contention in this paper that it is exactly the failure of this class of representation to capture the richness of semantic relations in hypertexts which employ hyperlinks, which contributes to hypertext disorientation. Productive terrain for future research might therefore be the development of more psychologically plausible metaphors for hypertext, and the basis of graphical displays on these alternative conceptions. An example of one such metaphor might be that of hypertext as more closely resembling the dynamic nature of Einstein's space-time continuum, in which the mass of an object [node] determines its gravity [meaning], which in turn warps the space-time in its vicinity, so as to draw neighbouring nodes within closer proximity [semantic relatedness]. Although in their service to accuracy, metaphors such as these might sacrifice something in terms of simplicity, it has become patently clear in the course of this study that this

¹¹¹Boroditsky (2000) uses this argument with respect to the comprehension of abstract concepts to test the claim that notions of time are informed by spatial metaphors.

would by no means be a bad thing.

Conclusion

This research study represents an attempt to construct a theoretical framework which is effective in accommodating the role of the student as well as the hypertext in education. Studies assessing the design aspects of the hypertext medium have been accused of fixating on ad-hoc features which have been appended to systems with minimal theoretical justification, and demonstrate the same lack of concern for the theoretical underpinnings of those alleged dangers of hypertext, such as disorientation, which have been documented in this paper. Approaching from the opposite direction, research into the interaction between personality and/or cognitive characteristics, and hyper/multimedia, has had its impact somewhat ameliorated by failure to specify those qualities of the hypertext which could effect its educational potential. Where it has done so, it has rarely been in a form which allows comparison with the results obtained by other studies, a problem which has been compounded by the creation of measurement instruments on a per study basis.

Pains were accordingly taken in the description of the research question presented herein, and of the method employed in answering it, to be as explicit as possible, so as to counter-act this trend. The motivational qualities of the hypertext user were identified as of paramount importance in negotiating the most characteristic of hypertext features, its hyperlinks, and attempts were made to capture the multiple ways in which their presence and use expresses itself. The attempt at a comprehensive specification of the user characteristics and the hypertext qualities of interest provided the freedom with which to examine a construct with a somewhat more nebulous nature. The examination of hypertext education within the context of a theoretically guided discussion of mental models represented an attempt to put the mechanism through which the hypertext medium might achieve its educational potential on a firmer theoretical footing.

The gung-ho attitude which was initially adopted by researchers in response to the discussion of the virtues of hypertext as an educational instrument is finally showing signs of giving way to more considered perspective. With the ever-increasing rapidity with which new technologies are being developed and deployed, it is becoming increasingly important that educators demonstrate moderation and wisdom in the reasons they select educational media. Without

proper consideration of how user characteristics interact with those of the system, predictions made regarding the educational potential of a medium will at best, ring hollow, and at worst, lead to a backlash against the use of the technology which will only serve to retard further empirical validation of its attributes.

University of Cape Town

References

- Ahuja, J. S., & Webster, J. (2001). Percieved disorientation: an examination of a new measure to assess web design effectiveness. *Interacting with Computers*, 14, 15-29.
- Albers, M. (1997). Cognitive strain as a factor in effective document design. In *Conference proceedings, SIGDOC 97 : the 15th annual international conference on computer documentation : Crossroads in communication* (p. 1-6). New York: ACM.
- Albrecht, J. E., & O'Brein, E. J. (1993). Updating a mental model: Maintaining both local and global coherence. *Journal of Experimental Psychology: Learning, Memory and Cognition*, 19, 1061-1070.
- Aldenderfer, M. S., & Blashfield, R. K. (1984). *Cluster analysis*. London: Sage Publications.
- Altun, A. (2000). Patterns in cognitive processes and strategies in hypertext reading: A case study of two experienced computer users. *Journal of Educational Multimedia and Hypermedia*, 9(1), 35-55.
- Ames, C., & Archer, J. (1988). Achievement goals in the classroom: Students' learning strategies and motivational processes. *Journal of Educational Psychology*, 80(3), 260-267.
- Astleitner, H., & Leutner, D. (1995). Learning strategies for unstructured hypermedia - a framework for theory, research, and practise. *Journal of Educational Computing Research*, 13(4), 387-400.
- Astleitner, H., & Leutner, D. (1996). Applying standard network analysis to hypermedia systems: Implications for learning. *Journal of Educational Computing Research*, 14(3), 285-303.
- AT&T Lab-Research. (2002). Graphviz - open source graph drawing software [Computer software]. <http://www.research.att.com/sw/tools/graphviz/>.
- Balcytiene, A. (1999). Exploring individual processes of knowledge construction with hypertext. *Instructional Science*, 27, 303-328.
- Bartasis, J., & Palumbo, D. (1995). *Theory and technology: Design consideration for hypermedia/discovery learning environments*. Retrieved April 03, 2001, from http://129.7.160.115/inst.5931/Discovery_Learning.

- Beasley, R. E., & Waugh, M. L. (1995). Cognitive mapping architectures and hypermedia disorientation: An empirical study. *Journal of Educational Multimedia and Hypermedia*, 4(2-3), 239-255.
- Bernstein, M. (1998). *Patterns of hypertext*. Retrieved May 23, 2001, from <http://www.eastgate.com/patterns/Print.html>.
- Bernstein, M. (1999). *Structural patterns and hypertext rhetoric*. Retrieved June 05, 2001, from <http://www1.acm.org/pubs/citations/journals/surveys/1999-31-4es/a19-bernstein/a19-bernstein.pdf>.
- Berry, M. W., & Dumais, S. (2001). *Latent semantic indexing web site - software*. Retrieved August 20, 2001, from <http://www.cs.utk.edu/lsi>.
- Boechler, P. M. (2001). How spatial is hyperspace? Interacting with hypertext documents: Cognitive process and concepts. *CyberPsychology and Behavior*, 4(1), 23-46.
- Borgatti, S., Everett, M. G., & Freeman, L. (1992). UCINET v4.0 for Dos [Computer software]. <http://www.analytictech.com/>.
- Boroditsky, L. (2000). Metaphoric structuring: understanding time through spatial metaphors. *Cognition*, 75, 1-28.
- Botafogo, R. A., Rivlin, E., & Schneiderman, B. (1992). Structural analysis of hypertexts: Identifying heirarchies and useful metrics. *ACM Transactions on Information Systems*, 10(2), 142-180.
- Bouras, C., Konidaris, A., Konidari, E., & Sevasti, A. (2001). *Introducing navigation graphs as a technique for improving www user browsing*. Retrieved February 20, 2002, from <http://citeseer.nj.nec.com/441736.html>.
- Braten, I., & Olaussen, B. S. (1998). The relationship between motivational beliefs and learning strategy use among Norwegian college students. *Contemporary Educational Psychology*, 23, 182-194.
- Bray, T., Paoli, J., Sperberg-McQueen, C. M., & Maler, E. (2000). *Extensible Markup Language (XML) 1.0 (second edition) - W3C Recommendation*. Retrieved December 20, 2001, from <http://www.w3.org/TR/2000/REC-xml-20001006>.
- Brewer, W. F. (1987). Schemas versus mental models in human memory. In P. Morris (Ed.), *Modelling cognition* (p. 187-197). Chichester: John Wiley and Sons.

- Britton, B. K., & Gulgoz, S. (1991). Using Kintsch's computational model to improve instructional text: Effect of repairing inference calls on recall and cognitive structures. *Journal of Educational Psychology*, 83(3), 329-345.
- Bush, V. (1945). *As we may think*. Retrieved September 10, 2001, from <http://www.theatlantic.com/unbound/flashbks/computer/bushf.htm>.
- Calvi, L. (1997). Navigation and disorientation: A case study. *Journal of Educational Multimedia and Hypermedia*, 6(3-4), 305-320.
- Calvi, L., & De Bra, P. (1998). A flexible hypertext courseware on the web based on a dynamic link structure. *Interacting with computers*, 143-154.
- Charney, D. (1994). The effect of hypertext on processes of reading and writing. In C. Selfe & S. Hilligoss (Eds.), *Literacy and computers* (p. 238-263). New York: The Modern Language Association of America.
- Chen, C. (1998). Generalised similarity analysis and pathfinder network scaling. *Interacting with computers*, 10, 107-128.
- Chen, C., & Czerwinski, M. (1998). From latent semantics to spatial hypertext - an integrated approach. In *Proceedings of the ninth acm conference on hypertext and hypermedia : links, objects, time and space structure in hypermedia systems* (p. 77-86). ACM Press.
- Chou, C., & Lin, H. (1998). The effect of navigation map types and cognitive styles on learner's performance in a computer-networked hypertext learning system. *Journal of Educational Multimedia and Hypermedia*, 7(2-3), 151-176.
- Czerwinski, M. P., & Larson, K. (in press). *Human factors and web design book*. Retrieved September 01, 2002, from <http://research.microsoft.com/users/marycz/home.htm>: Erlbaum.
- Deerwester, S., Dumais, S. T., Furnas, G. W., Landauer, T. K., & Harshman, R. (1990). Indexing by latent semantic analysis. *Journal of the American Society For Information Science*, 41, 391-407.
- De Rose, S., Maler, E., & Orchard, D. (2001). *XML Linking Language (XLink) Version 1.0 - W3C Recommendation*. Retrieved January 27, 2002, from <http://www.w3.org/TR/2001/REC-xlink-20010627>.

- Dias, P., & Sousa, P. (1997). Understanding navigation and disorientation in hypermedia learning environments. *Journal of Educational Multimedia and Hypermedia*, 6(2), 173-185.
- Dieberger, A. (1994). *Navigation in textual virtual environments using a city metaphor*. Retrieved November 29, 2001, from <http://citeseer.nj.nec.com/dieberger00navigation.html>.
- Dieberger, A. (1995). *Providing spatial navigation for the World Wide Web*. Retrieved November 29, 2001, from <http://citeseer.nj.nec.com/dieberger95providing.html>.
- Dieberger, A. (1996). *Browsing the WWW by interacting with a textual virtual environment - framework for experimenting with navigational metaphors*. Retrieved November 29, 2001, from citeseer.nj.nec.com/dieberger96browsing.html.
- Dieberger, A., & Frank, A. U. (1998). A city metaphor to support navigation in complex information spaces [Electronic version]. *Journal of Visual Languages and Computing*, 9(6), 597-622.
- Dillon, A., McKnight, C., & Richardson, J. (1993). Space — the final chapter or Why physical representations are not semantic intentions. In C. McKnight, A. Dillon, & J. Richardson (Eds.), *Hypertext - a psychological perspective* (p. 167-191). New York: Ellis Horwood.
- Duda, J. L., & Nicholls, J. G. (1992). Dimensions of achievement motivation in schoolwork and sport. *Journal of Educational Psychology*, 84, 290-299.
- Dweck, C. S. (1986). Motivational processes affecting learning. *American Psychologist*, 41, 1040-1048.
- Dweck, C. S., Chiu, C., & Hong, Y. (1995). Implicit theories and their role in judgements and reactions: A world from two perspectives. *Psychological Inquiry*, 6(4), 267-285.
- Eastin, M. S., & LaRose, R. (2000). Internet self-efficacy and the psychology of the digital divide [Electronic version]. *Journal of Computer-Mediated Communication*, 6(1).
- Edwards, D., & Hardman, L. (1989). Lost in hyperspace: Cognitive mapping and navigation in a hypertext environment. In R. McAleese (Ed.), *Hypertext: Theory into practise* (p. 121-149). Norwood, N.J.: Ablex Publishing.
- Eklund, J. (1995). *Cognitive models for structuring hypermedia and implications for learning from the world-wide web*. Retrieved July 13, 2001, from <http://www.scu.edu.au/sponsored/ausweb/usweb95/papers/hypertext/eklund>.

- Elliot, E. S., & Dweck, C. S. (1988). Goals: An approach to motivation and achievement. *Journal of Personality and Social Psychology*, 54(1), 5-12.
- Fitzgerald, G. E. (1998). Expert and novice performance in a hypermedia learning environment for classroom observation skills: The impact of prior teaching experience. *Journal of Educational Computing Research*, 19(2), 115-131.
- Foltz, P. W. (1996). Hypertext and cognition. In J. Rouet, J. Levonen, A. Dillon, & R. Spiro (Eds.), (p. 109-137). N.J.: Lawrence Erlbaum Associates.
- Foltz, P. W., Kintsch, W., & Landauer, T. K. (1998). The measurement of textual coherence with latent semantic analysis. *Discourse Processes*, 25(2-3), 285-307.
- Foltz, P. W., & Wells, A. D. (1999). Theories and applications of high-dimensional semantic models: Automatically derived reader's knowledge structures from text. *Behaviour, Research Methods, Instruments and Cognition*, 31(2), 208-214.
- Frecon, E., & Smith, G. (1998). *Webpath - a three dimensional web history*. Retrieved June 13, 2001, from <http://www.comp.lancs.ac.uk/computing/users/gbs/webpath/webpath.html>.
- Gamberini, L., & Bussolon, S. (2001). Human navigation in electronic environments [Electronic version]. *CyberPsychology and Behavior*, 4(1), 57-65.
- Garner, R., & Alexander, P. A. (1989). Metacognition: Answered and unanswered questions. *Educational Psychologist*, 24(2), 143-158.
- Garnham, A. (1992). Minimalism versus constructionism: A false dichotomy in theories of inference during reading. *PSYCOLOQUY*, 3(63). (Retrieved from <http://journals.ecs.soton.ac.uk/resource/psychology3.63>)
- Garnham, A. (1997). Representing information in mental models. In M. Conway (Ed.), *Cognitive models of memory* (p. 149-172). Cambridge, Massachusetts: The MIT Press.
- Glasgow, J., & Malton, A. (1999). *A semantics for model-based reasoning*. Retrieved February 15, 2002, from <http://citeseer.nj.nec.com/9877.html>.
- Glenberg, A. M., & Epstein, W. (1987). Inexpert calibration of comprehension. *Memory and Cognition*, 15(1), 84-93.
- Graesser, A. C., & Clark, L. F. (1985). The generation of knowledge-based inferences during narrative comprehension. In G. Rickheit & H. Strohner (Eds.), *Inferences in text processing* (p. 53-94). Amsterdam: Elsevier Science Publishers.

- Graesser, A. C., Gernsbacher, M. A., & Goldman, S. R. (1997). Cognition. In T. van Dijk (Ed.), *Discourse as structure and process - discourse studies: A multidisciplinary introduction, volume 1* (p. 292-319). London: Sage Publications.
- Graesser, A. C., Woll, S. B., Kowalski, D. J., & Smith, D. A. (1980). Memory for typical and atypical actions in scripted activities. *Journal of Experimental Psychology: Human Learning and Memory*, 6(5), 503-515.
- Graesser, A. C., & Zwaan, R. A. (1995). In S. M. C.A. Weaver III & C. Fletcher (Eds.), *Discourse comprehension: Essays in honor of Walter Kintsch* (p. 117-139). Hillsdale, N.J.: Lawrence Erlbaum Associates.
- Gvu's 6th www user surveys. (1996). Retrieved April 12, 2002, from http://www.cc.gatech.edu/gvu/user_surveys/survey-10-1996/.
- Hammond, N. (1993). Learning with hypertext: Problems, principles and prospects. In C. McKnight, A. Dillon, & J. Richardson (Eds.), *Hypertext: A psychological perspective* (p. 51-69). Ellis Horwood Ltd.
- Heylighen, F. (1995). *From world-wide web to super-brain*. Retrieved February 19, 2002, from <http://pespmc1.vub.ac.be/SUPBRAIN.html>.
- Hofer, B. (2000). Dimensionality and disciplinary differences in personal epistemology. *Contemporary Educational Psychology*, 25, 378-405.
- Hofer, B. K., & Pintrich, P. R. (1997). The development of epistemological theories: Beliefs about knowledge and knowing and their relation to learning. *Review of Educational Research*, 67(1), 88-140.
- Howell, D. (Ed.). (1997). *Statistical methods for psychology*. Washington: Duxbury Press.
- Hoyle, R. H. (1995). The structural equation modeling approach: Basic concepts and fundamental issues. In R. H. Hoyle (Ed.), *Structural equation modeling: Concepts, issues, and applications* (p. 16-36). Thousand Oaks, London: Sage Publications, Inc.
- Jacobson, M. J., & Spiro, R. J. (1995). Hypertext learning environments, cognitive flexibility, and the transfer of complex knowledge: An empirical investigation. *Journal of Educational Computing Research*, 12(4), 301-333.
- Johnson-Laird, P. N. (1983). *Mental models: Towards a cognitive science of language, inference, and consciousness*. Cambridge, England: Cambridge University Press.

- Johnson-Laird, P. N. (1990). Mental models. In M. Posner (Ed.), *Foundations of cognitive science* (p. 469-499). London: The MIT Press.
- Johnson-Laird, P. N. (1996). *Language and space*. Cambridge, Massachusetts: The MIT Press.
- Johnson-Laird, P. N., Girotto, V., & Legrenzi, P. (1998). *Mental models: a gentle guide for outsiders*. Retrieved April 20, 2002, from <http://www.si.umich.edu/ICOS/gentleintro.html>.
- Johnson-Laird, P. N., Schaeken, W., & Byrne, R. M. J. (1992). Propositional reasoning by model. *Psychological Review*, 99(3), 418-439.
- Jonassen, D. H. (1993). Effects of semantically structured hypertext knowledge base on user's knowledge structures. In A. D. C. Mcknight & J. Richardson (Eds.), *Hypertext - a psychological perspective* (p. 153-168). New York: Ellis Horwood.
- Jonassen, D. H., & Grabinger, S. (1990). Problems and issues in designing hypertext/hypermedia for learning. In D. H. Jonassen & H. Mandl (Eds.), *Designing Hypermedia for Learning* (p. 3-12). Berlin, Springer-Verlag.
- Kaplan, D. (2001). *Structural equation modeling: foundations and extensions*. Thousand Oaks: Sage Publications.
- Kintsch, E., Steinhart, D., & Stahl, G. (2000). *Developing summarization skills through the use of lsa-based feedback*. Retrieved May 4, 2001, from <http://lsa.colorado.edu/papers/ILE.ekintsch.pdf>.
- Kintsch, W. (1988). The role of knowledge in discourse comprehension: A construction-integration model. *Psychological Review*, 95(2), 163-182.
- Kintsch, W. (1993). Information accretion and reduction in text processing: Inferences. *Discourse Processes*, 16, 193-202.
- Kintsch, W. (1998). *Comprehension: A paradigm for cognition*. Cambridge: Cambridge University Press.
- Kintsch, W., & Franzke, M. (1995). The role of background knowledge in the recall of a news story. In R. Lorsch & E. O'Brein (Eds.), *Sources of coherence in reading* (p. 321-333). Hillsdale, N.J.: Erlbaum.

- Kleer, J. d., & Brown, J. (1983). Assumptions and ambiguities in mechanistic mental models. In D. Genter & A. L. Stevens (Eds.), *Mental models* (p. 155-190). Hillsdale, New Jersey: Lawrence Erlbaum Associates.
- Knowledge Analysis Technologies. (2002). Intelligent essay assessor [Computer software]. <http://www.knowledge-technologies.com>.
- Laham, D. (2002). Latent Semantic Analysis @ CU Boulder. <http://lsa.colorado.edu>.
- Landauer, T. K., & Dumais, S. T. (1997). A solution to Plato's Problem: The latent semantic analysis theory of acquisition, induction and representation of knowledge. *Psychological Review*, 104(2), 211-240.
- Landauer, T. K., Foltz, P. W., & Laham, D. (1998). Introduction to latent semantic analysis. *Discourse Processes*, 25, 259-284.
- Landauer, T. K., Laham, D., Rehder, B., & Schreiner, M. E. (1997). *How well can passage meaning be derived without using word order? A comparison of latent semantic analysis and humans*. Retrieved 15 April, 2001, from <http://lsa.colorado.edu/papers.html>.
- Landow, G. P. (1990). Popular fallacies about hypertext. In D. H. Jonassen & H. Mandl (Eds.), *Designing hypermedia for learning: Proceedings of the NATO advanced research workshop on designing hypertext/hypermedia for learning* (p. 39-60). Berlin: Springer-Verlag.
- Lawless, K., & Kulikowich, J. M. (1998). Domain knowledge, interest, and hypertext navigation: A study of individual differences. *Journal of Educational Multimedia and Hypermedia*, 7(1), 51-69.
- Lawless, K. A., & Kulikowich, J. M. (1996). Understanding hypertext navigation through cluster analysis. *Journal of Educational Computing Research*, 14(4), 385-399.
- Liao, Y. C. (1999). Effects of hypermedia on student's achievement: A meta-analysis. *Journal of Educational Multimedia and Hypermedia*, 8(3), 255-277.
- Licht, B. G., & Dweck, C. S. (1984). Determinants of academic achievement: The interaction of children's achievement orientations with skill area. *Developmental Psychology*, 20(4), 628-636.
- Lin, L., & Zabucky, K. M. (1998). Calibration of comprehension: Research and implications for education and instruction. *Contemporary Educational Psychology*, 23, 345-391.

- Lindsay, R. K. (1988). Images and inference. *Cognition*, 29, 229-250.
- MacCallum, R., Wegener, D., Uchino, B., & Fabrigar, L. (1993). The problem of equivalent models in applications of covariance structural analysis. *Psychological Bulletin*, 114, 185-199.
- MacDonald, S., & Stevenson, R. (1998). Navigation in hyperspace: An evaluation of the effects of navigational tools and subject matter expertise on browsing and information retrieval in hypertext. *Interacting with Computers*, 10, 129-142.
- Maki, R. H., Foley, J. M., Kajer, W. K., Thompson, R. C., & G., W. M. (1990). Increased processing enhances calibration of comprehension. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, Vol 16(4), 609-616.
- Marshall, C. C., & Rogers, R. A. (1992). *Two years before Mist: Experiences with Aquanet*. Retrieved 21 January, 2002, from <http://www.csdl.tamu.edu/~marshall/echt94.pdf>.
- Marshall, C. C., Shipman, F. M., & Coombs, J. H. (1994). *Viki: Spatial hypertext supporting emergent structures*. Retrieved 21 January, 2002, from <http://www.csdl.tamu.edu/~marshall/echt94.pdf>.
- Marshall, C. C., & Shipman, F. M. C. (1997). *Spatial hypertext and the practise of information triage*. Retrieved 21 January, 2002, from <http://www.csdl.tamu.edu/~shipman/ht97-final.pdf>.
- Mayes, T., Kibby, M. R., & Anderson, T. (1990). Hypertext: State of the art. In C. Green & R. McAleese (Eds.), (p. 121-129). Oxford, England: Intellect Limited.
- McEneaney, J. E. (1999). In K. Tochtermann (Ed.), *Hypertext 99: 10th ACM conference on hypertext and hypermedia* (Vol. 99, p. 61-71). ACM.
- McEneaney, J. E. (2001). Graphic and numerical methods to assess navigation in hypertext. *International Journal of Human-Computer Studies*, 55, 761-786.
- McKoon, G., & Ratcliff, R. (1992). Inference during reading. *Psychological Review*, 99, 440-466.
- McKoon, R., & Ratcliff, R. (1995). The minimalist hypothesis: Directions for research. In C. Weaver, S. Mannes, & C. Fletcher (Eds.), *Discourse comprehension: Essays in honor of walter kintsch* (p. 97-116). Hillsdale, New Jersey: Lawrence Erlbaum Associates.

- McNamara, D. S., Kintsch, E., Songer, N. B., & Kintsch, W. (1996). Are good texts always better? Interactions of text coherence, background knowledge, and levels of understanding in learning from text. *Cognition and Instruction*, 14(1), 1-43.
- Meece, J. L. (1994). The role of motivation in self-regulated learning. In D. Schunk & B. Zimmerman (Eds.), *Self-regulation of learning and performance: Issues in educational applications* (p. 25-44). Hillsdale, NJ: Lawrence Erlbaum Associates.
- Meece, J. L., Blumenfeld, P. C., & Hoyle, R. H. (1988). Student's goal orientations and cognitive engagement in classroom activities. *Journal of Educational Psychology*, 80, 514-523.
- Meece, J. L., & Holt, K. (1993). A pattern analysis of student's achievement goals. *Journal of Educational Psychology*, 85, 582-590.
- Miltiadou, M. (1999). *Motivational constructs as predictors of success in the online classroom*. Retrieved 19 October, 2001, from <http://seamonkey.ed.asu.edu/mcisaac/emc703/mariosf.html>.
- Miltiadou, M. (2000). *Motivational constructs as predictors of success in the online classroom*. Retrieved 10 November, 2001, from <http://jan.ucc.nau.edu/slm/Aero/index.htm>.
- Modjeska, D. (1998). *Parallel worlds: Spatial and textual representations of information structure in VR*. Retrieved June 29, 2001, from <http://citeseer.nj.nec.com/132270.html>.
- Modjeska, D., & Marsh, A. (1997). *Structure and memorability of web sites*. Retrieved June 29, 2001, from <http://citeseer.nj.nec.com/modjeska97structure.html>.
- Moravcsik, J. E., & Kintsch, W. (1993). Writing quality, reading skills, and domain knowledge as factors in text comprehension. *Canadian Journal of Experimental Psychology*, 47(2), 360-374.
- Murphy, C., Coover, D., & Owen, S. (1989). Development and validation of the computer self-efficacy scale. *Educational and Psychological Measurement*, 49, 893-899.
- Niederhauser, D. S., Reynolds, R. E., Salmen, D. J., & Skolmoski, P. (2000). The influence of cognitive load on learning from hypertext. *Journal of Educational Computing Research*, 23(3), 237-255.

- Nielsen, J. (1989). The matters that really matter for hypertext usability. In ACM (Ed.), *Proceedings of Hypertext '89, Second ACM conference on hypertext* (p. 239-248). Pennsylvania.
- Nielsen, J. (2000). *Designing web usability*. Indianapolis, USA: New Riders Publishing.
- Niemiec, R. P., Sikorski, C., & Walberg, H. J. (1996). Learner-control effects: A review of reviews and a meta-analysis. *Journal of Educational Computing Research*, 15(2), 157-174.
- Noel, R. W., & Hunter, C. M. (2000). Mapping the physical world to psychological reality: Creating synthetic environments. In D. Boyarski & A. Kellog (Eds.), *DIS2000 - designing interactive systems: Processes, practices, methods and techniques* (p. 203-207). New York City: ACM.
- Nunes, J. M. B., & Fowell, S. P. (1996). *Hypermedia as an experiential learning tool: a theoretical model*. Retrieved April 10, 2001, from <http://www.shef.ac.uk/is/publications/infres/paper12.html>.
- Otter, M., & Johnson, H. (2000). Lost in hyperspace: Metrics and mental models. *Interacting with computers*, 13, 1-40.
- Oxford. (1989). *Oxford english dictionary*. Oxford: Clarendon Press.
- Panero, J. C. (1995). *Navigation and mental models of the world wide web: An exploratory study*. Retrieved April 21, 2001, from <http://www.oxy.edu/~panero/WWWproj/paper.html>.
- Payne, S. J. (1993). Memory for mental models of spatial descriptions: An episodic-construction-trace hypothesis. *Memory and Cognition*, 21(5), 591-603.
- Perseus Development Corporation. (2000). Perseus Survey Solutions 3.0 [Computer software]. <http://www.perseusdevelopment.com/>.
- Persson, P. (1998). *Supporting navigation in digital environments: A narrative approach*. Retrieved July 15, 2001, from <http://www.sic.se/humle/projects/persona/web/littsurvey/ch12.pdf>.
- Pintrich, P. R., Smith, D. A. F., Garcia, T., & McKeachie, W. J. (1991). *A manual for the use of the motivated strategies for learning questionnaire MSLQ(No. 91-B-004)*. Michigan, USA: National Center for Research to Improve Postsecondary Teaching and Learning.

- Pintrich, P. R., Smith, D. A. F., Garcia, T., & McKeachie, W. J. (1993). Reliability and predictive validity of the motivated strategies for learning questionnaire (MSLQ). *Educational and Psychological Measurement*, 53, 801-813.
- Pitkow, J. E., & Kehoe, C. M. (1996). Emerging trends in the www user population [Electronic version]. *Communications of the ACM*, 39(6).
- Qian, G., & Alvermann, D. (1995). Role of epistemological beliefs and learned helplessness in secondary school students' learning science concepts from text. *Journal of Educational Psychology*, 87(2), 282-292.
- Reed, W. M., Ayersman, D. J., & Liu, M. (1996). The effects of students' computer-based prior experiences and instructional exposures on the application of hypermedia-related mental models. *Journal of Educational Computing Research*, 14(2), 185-207.
- Roedel, T. D., Schraw, G., & Plake, B. S. (1994). Validation of a measure of learning and performance goal orientations. *Educational and Psychological Measurement*, 54(2), 1013-1021.
- Rouet, J., & Levonen, J. J. (1996). Studying and learning with hypertext: Empirical studies and their implications. In J. Rouet, J. Levonen, A. Dillon, & R. Spiro (Eds.), *Hypertext and cognition* (p. 9-24). N.J.: Lawrence Erlbaum Associates.
- Rumelhart, D. E., & Norman, D. A. (1978). Accretion, tuning and restructuring: Three modes of learning. In J. Cotton & R. Klatsky (Eds.), *Semantic factors in cognition* (p. 37-53). Hillsdale, N.J.: Lawrence Erlbaum Associates.
- SAIDE. (1999). Images of africa evaluation [Electronic version].
- Sanders, T. J. M., & Noordman, L. G. M. (2000). The role of coherence relations and their linguistic markers in text processing. *Discourse Processes*, 29(1), 37-60.
- Schommer, M. (1993). Epistemological development and academic performance among secondary students. *Journal of Educational Psychology*, 85(3), 406-411.
- Schommer, M. (1994). An emerging conceptualisation of epistemological beliefs, and their role in learning. In R. Garner & P. A. Alexander (Eds.), *Beliefs about text and instruction with text* (p. 25-40). Hillsdale, NJ: Erlbaum.
- Schonotz, W., & Bannert, M. (in press). Construction and interference in learning form multiple representation. *Learning and Instruction*.

- Schroeder, E. E., & Grabowski, B. L. (1995). Patterns of exploration and learning with hypermedia. *Journal of Educational Computing Research*, 13(4), 313-335.
- Schunk, D. H. (1991). Self-efficacy and academic motivation. *Educational Psychologist*, 26(3-4), 207-231.
- Schwamb, K. B. (1990). *Mental models: A survey*. Retrieved April 29, 2002, from <http://citeseer.nj.nec.com/schwamb90mental.html>.
- Scott, J. (2000). *Social network analysis - A handbook*. London: Sage Publications.
- Shabo, A. (1997). Integrating constructionism and instructionism in educational hypermedia programs. *Journal of Educational Computing Research*, 17(3), 231-247.
- Shipman, F. M., Hsieh, H., Maloor, P. J., & Moore, M. (2001). The visual knowledge builder: A second generation spatial hypertext. In H. Davis, Y. Douglas, & D. G. Durand (Eds.), *Hypertext '01, Proceedings of the Twelfth ACM Conference on Hypertext and Hypermedia* (p. 113-122). University of Aarhus, Aarhus, Denmark.
- Shipman, F. M., & Marshall, C. C. (1999). *Formality considered harmful: Experiences, emerging themes, and directions on the use of formal representations in interactive systems*. Retrieved January 21, 2001, from <http://www.csdl.tamu.edu/shipman/cscw.pdf>.
- Shipman, F. M., & Marshall, C. C. (2000). *Spatial hypertext: An alternative to navigational and semantic links*. Retrieved January 21, 2001, from <http://www.csdl.tamu.edu/shipman/spatialhypertext.html>.
- Shipman, F. M., Marshall, C. C., & LeMere, M. (1999). *Beyond location: Hypertext workspaces and non-linear views*. Retrieved January 21, 2001, from <http://www.csdl.tamu.edu/shipman/ht99.pdf>.
- Shum, S. (1990). Real and virtual spaces: Mapping from spatial cognition to hypertext. *Hypermedia*, 2(2), 133-158.
- Simpson, A., & McKnight, C. (1990). Navigation in hypertext: Structural cues and mental maps. In R. McAleese & C. Green (Eds.), *Hypertext: state of the art* (p. 73-83). Oxford: Intellect Limited.
- Smith, P. (1996). Towards a practical measure of hypertext usability [Electronic version]. *Interacting with Computers*, 8(4), 365-381.


- Spiro, R. J., Coulson, R. L., Feltovich, P. J., & Anderson, D. K. (1989). Cognitive Flexibility Theory: Advanced knowledge acquisition in ill-structured domains (tech. rep. no. 441). In V. Patel (Ed.), *Proceedings of the 10th annual conference of the cognitive science society*. (p. 123-137). Hillsdale, NJ: Erlbaum.
- Spiro, R. J., Feltovich, P. J., Jacobson, M. J., & Coulson, R. L. (1996). *Cognitive flexibility, constructivism, and hypertext: Random access instruction for advanced knowledge acquisition in ill-structured domains*. Retrieved June 10, 2001, from http://phoenix.sce.fct.unl.pt/simposio/Rand_Spiro.htm.
- Stanton, N., Correia, A. P., & Dias, P. (2000). Efficacy of a map on search, orientation and access behaviour in a hypermedia system [Electronic version]. *Computers & Education*, 35, 263-279.
- StatSoft, I. (2000). *STATISTICA for Windows [Computer program manual]*. Tulsa, OK: StatSoft, Inc.
- Strommen, E. F. (1995). *Constructivism, technology, and the future of classroom learning*. Retrieved April 16, 2001, from <http://www.ilt.columbia.edu/ilt/papers/construct.html>.
- Tergan, S. (1997a). Conceptual and methodological shortcomings in hypertext/hypermedia design and research. *Journal of Educational Computing Research*, 16(3), 209-235.
- Tergan, S. (1997b). Misleading theoretical assumptions in hypertext/hypermedia research. *Journal of Educational Multimedia and Hypermedia*, 6(3-4), 257-283.
- Thury, E. M. (1998). Analysis of student web browsing behavior: Implications for designing and evaluating web sites. In *Scaling the heights: The future of information technology. The sixteenth annual international conference of computer documentation: Conference Proceedings - September 23-26* (p. 265-270). ACM.
- Tversky, A. (1977). Features of similarity. *Psychological Review*, 84, 327-352.
- West, S. G., Finch, J. F., & Curran, P. J. (1995). Structural equation models with nonnormal variables: Problems and remedies. In R. Hoyle (Ed.), *Structural equation modelling: Concepts, issues and applications* (p. 56-75). New York: Sage Publications, Inc.
- Whalley, P. (1990). Models of hypertext structure and learning. In D. Jonassen & H. Mandl (Eds.), *Designing hypermedia for learning: Proceedings of the NATO*

- Advanced Research Workshop on designing Hypertext/Hypermedia for learning* (p. 61-67). Berlin: Springer-Verlag.
- Winne, P. H. (1995). Inherent details in self-regulated learning. *Educational Psychologist*, 30(4), 173-187.
- Wolfe, M. B., Kintsch, W., & Landauer, T. K. (1998). Learning from text: Matching readers and text by latent semantic analysis. *Discourse Processes*, 25, 309-336.
- Woude, A. Van der. (2002). Probe [Computer software]. <http://probe.sourceforge.net/>.
- Young, R. M. (1983). Surrogates and mappings: Two kinds of conceptual models for interactive devices. In D. Gentner & A. L. Stevens (Eds.), *Mental models* (p. 35-52). Hillsdale, New Jersey: Lawrence Erlbaum Associates.

A Appendix

University of Cape Town


A.1 Graphical representation and descriptive statistics for the clusters derived from the navigational behaviour of participants in the Hierarchy condition



	Means and Standard Deviations			
	Group 1	Group 2	Group 3	Group 4
Performance Orientation	10.8(1.9)	10.7(4)	9.4(4.4)	11.9(2.2)
Learning Orientation	33.4(2.9)	29.2(4.9)	28(6.7)	32.9(5.8)
Efficacy	98.6(18.2)	90.8(22)	106.4(21.4)	101.1(14.7)
Extrinsic	21.4(2.2)	21.9(4.3)	17.4(7.9)	19.7(4.3)
Intrinsic	18.6(2.1)	17.8(2.5)	18.6(5.7)	19.6(3.5)
Anxiety	17.2(3.9)	17.4(5.9)	16(8.5)	20.9(3.8)
Effort	19(3.9)	16.4(3.2)	14.6(5.8)	18.4(5.9)
Background Knowledge	4(1.9)	5.1(1.1)	4.8(1.9)	4.6(1.4)

Notes. SD's in parentheses

A.2 Graphical representation and descriptive statistics for the clusters derived from the navigational behaviour of participants in the network hypertexts



	Means and Standard Deviations			
	Group 1	Group 2	Group 3	Group 4
Performance Orientation	11.3(3.4)	9.8(4)	11.1(4.7)	10.8(4.4)
Learning Orientation	31.6(5.7)	31.7(6.7)	36.3(3.8)	28.8(5.5)
Efficacy	98.9(15.4)	106.8(12.4)	117.6(12.9)	95.4(16.4)
Extrinsic	20.1(6.5)	18.9(7.1)	21.9(6.4)	21(6.4)
Intrinsic	17.1(1.5)	19.3(5)	18.1(2.5)	17.8(3.5)
Anxiety	21(2.7)	17.4(7)	20.1(3.6)	18.9(4.9)
Effort	19.3(3.4)	18.5(5.3)	19.1(2.5)	15.1(4.1)
Background Knowledge	4.3(1.1)	5.1(1.1)	5(1.2)	5(1.4)

Notes. *SD*'s in parentheses

A.3 Tests of assumptions for analysis of covariance in investigating the effects of hypertext structure and goal orientation on recall

Parallelism of pre - knowledge levels with post - knowledge levels

Table 1: Univariate tests of Parallelism

	<i>df</i>	<i>F</i>	η^2	<i>p</i>
Effect	3	0.071	.003	.975
Error	67	(7.84)		

Note. Values enclosed in parentheses represent mean square errors

Table 2: Linearity of relationship between dependant variable and covariate

	Post-knowledge
Pre-knowledge	.133

p* < .05 *p* < .01

Note. *N* = 75

B Appendix

University of Cape Town

B.1 Hypertext System diagrams and screenshots

Description of coding sheet for hypertext schematics

The list below represents the different modules, as well as the associated sub-nodes, from which the hypertext was constituted. The letters in the nodes in the diagrams in Appendix B2 signify those sub-sections which are identified by the letters within parentheses.¹

Paradise	(P)re-modern African "Monsters" Representations of the (K)hoekhoen (M)odernising the Image The story of Human (E)volution The rise of the (E)thnographic Monograph
Race	The (O)rigins and Deconstruction of Racial Theories The Modern Biological Concept of Race/(C)lines The Construction and Deconstruction of (W)hiteness
God	(R)eligious Consciousness (C)hristianity and African religion
African Cities	The Legend of the (L)ost City Welcome to (F)es (N)omads The (H)eritage Industry

Screenshot descriptions

The first screenshot in Appendix B3 is of the home page of the hypertext. This page is identical for both the hierarchical and network conditions.

The second screenshot presents an example of a destination node in the network hypertext conditions, upon clicking of a hyperlink. The highlighted text is clearly visible, as are the activated hyperlinks themselves.



Path diagrams

The path diagrams on consecutive pages for the hypertext systems in the Network condition represent the summed paths taken by LO and NLO participants, respectively. Only those links which had been followed on two or more occasions are represented.

¹In the case of nodes labelled with consecutives letters within the Paradise component, the diagram should be read in the same vertical order as they are listed.



B.2 Comparison of navigational structure of hypertexts

University of Cape Town



B.3 Screenshot of homepage and sample hypertext node


University of Cape Town



B.4 Comparison of path diagrams for the LO and NLO participants in the Network condition


University of Cape Town

Path diagram for the NLO students in the network conditions
Each link indicated traversed at least twice





Path diagram for the LO students in the network conditions
Each link indicated traversed at least twice



C Appendix

Questionnaires

University of Cape Town

C.1 Background - Knowledge Questionnaire

1. When did the mass European colonisation of Africa, referred to as "The Scramble for Africa", begin?
 - Towards the end of the nineteenth century
 - At the beginning of the twentieth century
 - At the beginning of the nineteenth century
 - Towards the end of the twentieth century
2. Which of the following nations was not a major player in the colonisation of Africa?
 - The British
 - The Dutch
 - The Portuguese
 - The Swiss
3. Which of the following images did not reflect how the Africans were perceived by the early European explorers?
 - As children
 - As savages
 - As pure
 - As civilised
4. What was the name of one of the most well known 19th Century European explorers of Africa?
 - David Spencer
 - David Livingstone
 - David Churchill
 - Louis Agassiz
5. When did Vasco da Gama first visit the Cape?
 - 16th Century
 - 15th Century
 - 19th Century
 - 18th Century

6. What do you think was regarded by so-called "Salvage missionaries" as the greatest threat to the "African" way of life?
- Conversion by the Africans to Christianity
 - Failure to convert to Christianity
 - Introduction of European immorality
 - Modernisation of Africa
7. To whom did the first European discoverer of the ruins of Great Zimbabwe attribute the building of the original site?
- Prester John
 - Sir John Mandeville
 - The Syrians
 - The Queen of Sheba
8. How many "species" did Carl Linneaus divide the human race up into during the seventeenth Century?
- 3
 - 4
 - 5
 - 2
9. The so-called "Piltdown Hoax" refers to a hoax perpetuated in order to change ideas about:
- Human evolution
 - Slavery
 - The origins of the Kingdom of Great Zimbabwe
 - The location of the legendary "Lost City"
10. The nation known as the Berbers are to be found in:
- Ethiopia
 - Turkey
 - Tibet
 - Morocco

C.2 Goals Inventory Questionnaire

Statement	Strongly Disagree	Disagree	Neutral	Agree	Strongly Agree
1. I enjoy challenging school assignments	☐	☐	☐	☐	☐
2. It is important for me to get better grades than my classmates	☐	☐	☐	☐	☐
3. I persevere even when I am frustrated by a task	☐	☐	☐	☐	☐
4. Academic success is largely due to effort	☐	☐	☐	☐	☐
5. Sticking with a challenging task is rewarding	☐	☐	☐	☐	☐
6. I try even harder after I fail at something	☐	☐	☐	☐	☐
7. I adapt well to challenging circumstances	☐	☐	☐	☐	☐
8. I am willing to cheat to get a good grade	☐	☐	☐	☐	☐
9. I work hard even when I don't like a class	☐	☐	☐	☐	☐
10. I am very determined to reach my goals	☐	☐	☐	☐	☐
11. Personal mastery of a subject is important to me	☐	☐	☐	☐	☐
12. I work very hard to improve myself	☐	☐	☐	☐	☐
13. I like others to think I know a lot	☐	☐	☐	☐	☐
14. It bothers me the whole day when I make a big mistake	☐	☐	☐	☐	☐
15. I feel angry when I do not do as well as others	☐	☐	☐	☐	☐
16. I am naturally motivated to learn	☐	☐	☐	☐	☐
17. I prefer challenging tasks even if I don't do as well at them	☐	☐	☐	☐	☐
18. Every student can learn to be a successful learner	☐	☐	☐	☐	☐
19. Learning can be judged best by the grade one gets	☐	☐	☐	☐	☐
20. My grades do not necessarily reflect how much I learn	☐	☐	☐	☐	☐
21. Mistakes are a healthy part of learning	☐	☐	☐	☐	☐
22. I feel most satisfied when I work hard to achieve something	☐	☐	☐	☐	☐
23. I would rather have people think I am lazy than stupid	☐	☐	☐	☐	☐
24. It is important to me to always do better than others	☐	☐	☐	☐	☐
25. I give up too easily when faced with a difficult task	☐	☐	☐	☐	☐

C.3. Factual - Knowledge Questionnaire

1. What did the Portuguese explorers call Africa during the 15th Century?
 - "The Lost Continent"
 - "Darkest Africa"
 - "The New World"
 - "The Black-Lands"
2. What are the Peixe Muhler?
 - Mermaids
 - Sea Serpents
 - Dragons
 - Khoekhoen
3. What does the phrase "year zero" refer to?
 - The birth of Christ
 - The creation of the world
 - Noah's Ark and the Flood
 - The Apocalypse
4. What did the missionary Henri Junod believe distinguished between different African tribes?
 - Religion
 - Customs
 - Nothing; He believed that all the tribes originated from one group
 - Language
5. What was left out of Henri Junod's photographs of African people?
 - Slavery
 - Warfare
 - Marriage Ceremonies
 - Guns

6. What was Phrenology?
- The measurement of the size and shape of the skull
 - The measurement of facial features
 - The study of magic
 - The classification of racial groups according to body measurements
7. What is Polygenism?
- The classification of diverse superstitions as part of a single religion
 - The theory of the separate origins of racial groups
 - The belief in multiple gods
 - The theory that all of the racial groups are descended from a single ancestral population
8. Which of the following points represents the most novel and important similarity between "Whiteness Studies" and "Masculinity Studies"?
- Both are relatively new fields of study
 - Both have been used to challenge power relationships
 - Both try to explain why one group has power over another
 - Both focus attention on privileges previously taken for granted
9. Which of the following was not regarded as an obstacle by the Europeans in their mission to Christianise and "civilise" the African?
- Myth
 - Magic
 - Tribalism
 - Marriage customs
10. How did Caton-Thompson account for the grandeur of Great Zimbabwe?
- The Africans were assisted by coastal Arabians or Persians
 - It was the work of a "lost civilisation"
 - It was designed by a lone savage genius
 - It had been of great religious significance

11. Who translated the Koran into Berber?

- The Bourghwatas
- The Arabs
- Muarabeen
- French missionaries

12. According to the legend of Sun City's "Lost City", how long after the city was destroyed was it "rediscovered"?

- 300 years
- 500 years
- 250 years
- 400 years

13. Why was there an emphasis put on Africans clothing themselves?

- Nakedness was seen by the colonialists as obscene
- The lack of clothes was seen as a serious contributor to the spread of disease
- Nakedness was perceived as immoral and a source of temptation
- The wearing of European clothes was part of the mission to westernise the African

14. What attracts tourists to the Villa Nouveau in the city of Fes?

- The luxury hotels
- The restaurants
- The night life
- The casino

15. When did the failed Berber coup of Morocco take place?

- 1956
- 1971
- 1974
- 1983

16. What is the pivotal difference between hunting and gathering and farming communities, as far as population genetics is concerned?
- The differences in lifestyle
 - Technological differences
 - Size
 - Social interaction
17. The discovery of an alleged human ancestor in what became known as the "Piltdown Hoax" was so readily accepted because:
- It tied in with theories of human evolution at the time
 - It had an upright stature
 - It possessed a large brain
 - Charles Darwin's "Origin of Species" had recently been published
18. According to David Chidester, when was religion "discovered" amongst the Africans?
- When control was gained over their cultural and social institutions
 - With the establishment in Africa of Christian Missions
 - With the "discovery" of African indigenous religions
 - With the erection of the first African Christian church

C.4 Motivated Strategies for Learning Questionnaire

Original Questions ²

Intrinsic Goal Orientation

1. In a class like this, I prefer course material that really challenges me so I can learn new things.
2. In a class like this, I prefer course material that arouses my curiosity, even if it is difficult to learn.
3. The most satisfying thing for me in this course is trying to understand the content as thoroughly as possible
4. When I have the opportunity in this class, I choose course assignments that I can learn even if they don't guarantee a good grade

Extrinsic Goal Orientation

1. Getting a good grade in this class is the most satisfying thing for me right now.
2. The most important thing for me right now is improving my overall grade point average, so my main concern in this class is getting a good grade
3. If I can, I want to get better grades in this class than most of the other students
4. I want to do well in this class because it is important to show my ability to my family, friends, employer, or others.

Effort Regulation

1. I often feel so lazy or bored when I study for this class that I quit before I finish what I planned to do (REVERSED)
2. I work hard to do well in this class even if I don't like what we are doing
3. When course work is difficult, I give up or only study the easy parts. (REVERSED)

4. Even when course materials are dull and uninteresting, I manage to keep working until I finish.

Test Anxiety

1. When I take a test I think about how poorly I am doing compared with the other students
2. When I take a test I think about other items on other parts of the test I can't answer
3. When I take tests I think of the consequences of failing
4. I have an uneasy, upset feeling when I take an exam
5. I feel my heart beating fast when I take an exam

Revised Questions

Intrinsic Goal Orientation

1. I prefer course material that really challenges me so I can learn new things.
2. I prefer course material that arouses my curiosity, even if it is difficult to learn.
3. The most satisfying thing for me in a course is trying to understand the content as thoroughly as possible
4. When I have the opportunity in a class, I choose course assignments that I can learn even if they don't guarantee a good grade

²Response rating on a 7 point Likert-type scale, ranging from "Not true at all " to "Always true".

Extrinsic Goal Orientation

1. Getting a good grade in a class is the most satisfying thing for me right now.
2. The most important thing for me right now is improving my overall grade point average, so my main concern in a class is getting a good grade
3. If I can, I want to get better grades in a class than most of the other students
4. I want to do well in a class because it is important to show my ability to my family, friends, employer, or others.

Effort Regulation

1. I often feel so lazy or bored when I study for a class that I quit before I finish what I planned to do (REVERSED)
2. I work hard to do well in a class even if I don't like what we are doing

3. When course work is difficult, I give up or only study the easy parts. (REVERSED)
4. Even when course materials are dull and uninteresting, I manage to keep working until I finish.

Test Anxiety

1. When I take a test I think about how poorly I am doing compared with the other students
2. When I take a test I think about other items on other parts of the test I can't answer
3. When I take tests I think of the consequences of failing
4. I have an uneasy, upset feeling when I take an exam
5. I feel my heart beating fast when I take an exam

C.5 Computer and Internet Skills scale

Computer Skills³

I feel confident..

1. ..Working on a personal computer
2. ..Using the users guide when help is needed
3. ..Entering and saving data (numbers or words) into a file
4. ..Escaping (exiting) from the programme (software)
5. ..Understanding terms/words relating to computer hardware
6. ..Understanding terms/words relating to computer software
7. ..Handling a floppy disc correctly
8. ..Making selections from an onscreen menu
9. ..Using a printer to make a 'hard copy' of my work
10. ..Copying a disc
11. ..Copying an individual file
12. ..Adding and deleting information from a data file
13. ..Moving the cursor around the monitor screen
14. ..Writing simple programmes for the computer
15. ..Using the computer to write a letter or essay

16. ..Describing the function of computer hardware (e.g. keyboard, monitor, disc drives, computer processing unit)
17. ..Understanding the 3 stages of data processing: input, processing, output
18. ..Storing software correctly
19. ..Explaining why a programme (software) will or will not run on a given computer
20. ..Getting rid of files when they are no longer needed
21. ..Troubleshooting computer problems

Internet Usage

I feel confident..

1. ..understanding terms/words relating to Internet hardware
2. ..understanding terms/words relating to Internet software
3. ..describing functions of Internet hardware
4. ..trouble shooting Internet problems
5. ..explaining why a task will not run on the Internet
6. ..using the Internet to gather data
7. ..confident learning advanced skills within a specific Internet program
8. ..turning to an on-line discussion group when help is needed

³Response rating on a 5 point Likert-type scale, ranging from "Strongly Disagree" to "Strongly Agree."

C.6 Miscellaneous questions posed at start of second list of questionnaires

Just a few questions - Internet Explorer

File Edit View Favorites Tools Help

Back Forward Stop Refresh Home Search Favorites History Web Feeds Print

Address: [http://www.questcenter.com/quest.html](#) Go Links

Just a few questions

1. How many pages do you think the entire hypertext you've just read consisted of?

- ☐ 40 - 50 pages
- ☐ 50 - 60 pages
- ☐ 60 - 70 pages
- ☐ 70 - 80 pages
- ☐ 80 - 100 pages

2. What was the name (you may use the acronym) of the Hypermedia Instructional System developed by Spiro et. al. in the handout you were given, entitled "Cognitive Flexibility, Constructivism, and Hypertext" - Hint: it was inspired by the name of a classic film

Done My Computer

C.7 Fact Finding Instructions


The final task is known as the Fact - finding task, and as the name implies, you are expected to find the answers to the questions in the text itself. The questions you are asked to answer have accordingly been chosen so that you are unlikely to remember the answers from the text. Even if you are sure that you do know the answer to a question, however, it is highly recommended that you confirm your feelings by tracking the answer down in the text anyway.

The following page will consist of two frames or windows. The window on top contains the actual question, while in the window below you will find the hypertext you have just read through. Once you have managed to track down the answer in the text, fill in the box in the top frame, and press the button labelled "submit". You may then proceed to the next question by pressing on the button labelled "Next Question". Please make sure of the answer, as you will not be able to return to the question once you have moved on.

It is important that you click on the "submit" button before clicking on the "Next Question" button.

Fact Finding Questions

1. Which map of the World was "rediscovered" during the fifteenth century?
2. When was Sir Mandeville's Voyages published?
3. What was the name of the Chief from whose photo Henri Junod removed a chair?
4. Who was both an advocate of evolution, as well as one of Louis Aggasiz's assistants?
5. Name the groups, in addition to the Jews, who have occupied ambiguous position in relation to the dominant white group in the States?
6. Who is the scholar who made a study of the religious hierarchy in Langa, Cape Town?
7. Who discovered Great Zimbabwe?



List of questions which had to be answered and the interface through which they were presented to the students.


C.8 Word sorting task instructions

The next task ahead of you is the same word - sorting task that you encountered in the first session. The procedure is almost exactly the same, the only difference being that once you have finished printing the page out, you must click on the "Next Questionnaire" button at the top of the page. If you do not have credits on your printing account, tell the instructor, and he will save your copy of the task onto stiffy disk. The instructions originally given to you are repeated below.

On pressing the button at the bottom of this page, you will be directed to a page asking for your student number. Once you have entered your number and clicked the submit button you will be directed to a page consisting of a collection of words at the top, with the rest of the page divided up into eight columns. The task asked of you is to move each of the words into groups or categories of related words. The only criteria for success at this task is that the categories you form make sense to you, and that you form at least 3, and at most 8 categories.

You can move the words in a drag-and-drop fashion by positioning the cursor over the word, clicking on the left mouse button, and with the button depressed, moving the word to a particular column. Once you are satisfied with the word clusters that you have formed, please print out the page. You should do this by pressing the Alt and Print Screen button, simultaneously, and pasting what you have copied into Microsoft Word (Go to Edit ->Paste Special, and then just press the OK button.) Please hand the print-out to the instructor.

Should you be unsure about what is required of you whilst proceeding, you can consult the printed copy of these instructions which were handed out.



D Additional Materials

University of Cape Town

D.1 Custom LSA data corpus

1. Boonzaire, E. (1988) Race and the Race Paradigm. in Boonzaire and Sharp (eds.), South African Keywords. Cape Town
2. Brantlinger, P. (1988) The Geneology of the Myth of the 'Dark Continent'. Rule of Darkness: British Literature and Imperialism, 1830-1914. London: Ithaka Press
3. Chidester, David (1996) Savage Systems: Colonialism and Comparative Religion in Southern Africa. Cape Town: University of Cape Town Press, chs 1-2
4. Crafford, D. (1996) "African traditional religions", in Piet Meiring (ed.) A World of Religions: A South African Perspective (Pretoria: Kagiso Publishers), pp 1-26
5. Dyer, R. (1997) White. London: Routledge. pp.145- 183.
6. Foster, D. (1991) On Racism: Virulent mythologies and fragile threads.
7. Fredrickson, G. (1981) White Supremecy. New York: Oxford University Press. ch1-3
8. Greenblatt, S. (1991) Marvellous Possessions. The wonder of the new world. Chicago: University of Chicago Press. pp 26-51.
9. Hall, M. (1995) "The Legend of the Lost City, or, The Man with Golden Balls." Journal of South African Studies 21(2):179-200
10. Landau, M. (1984) Human Evolution as Narrative. American Scientist 72 ,pp 262-268.
11. Pratt, M.L. (1992) Imperial Eyes: Travel Writing and Transculturation. London New York: Routledge. chapters 3 4
12. Robert J. Thornton (1983) "Narrative ethnography in Africa", 1850-1920" Man 18
13. Rosman Rubel, The Tapestry of Culture, Ch.9
14. Smith, A.B. (1993) Different Facets of the Crystal: Early European Images of the Khoikhoi at the Cape, South Africa, in M. Hall and A nn Markell, (eds.) Historical Archaeology in the Western Cape: The S. A. Archaeological Society

D.2 Composition of tasaALL space

The breakdown for samples by academic area (in tasaALL):

	document	paragraphs
LanguageArts	16,044	57,106
Health	1,359	3,396
HomeEconomics	283	403
IndustrialArts	142	462
Science	5,356	15,569
SocialStudies	10,501	29,280
Business	1079	4834
Miscellaneous	675	2272
Unmarked	2,212	6,305
Total	37,651	119,627

specifics for tasaALL				
name	grade	#docs	#terms	#dims
tasaALL	college	37,651	92,409	419

Author Note

Jonathan Ipser, Department of Psychology, University of Cape Town.
Correspondence concerning this article and associated materials should be addressed to Jonathan Ipser, 50 Mountain View Drive, Bellville, 7530. Email: ipsjon001@mail.uct.ac.za

Colophon

This dissertation was formatted in accordance with the APA Publishing Manual (5th Ed.) guidelines. Deviations from APA format are evident in the Front matter and Appendixes, as well as in the use of one and a half spacing throughout the paper and the presence of the Colophon itself. The entire dissertation was typeset using the $\text{\LaTeX}$ document preparation system, and subsequently converted to PDF format for printing. Graphs were produced by means of the Grace graphing software, while the figures were created using both the xfig and Visio drawing packages — the former for higher resolution images at 1200 dpi. The entire dissertation, as well as 41 articles referenced in the paper, and both versions of the hypertext system which was employed in the study, are available on CDROM. This has been bundled with the hard copy of the dissertation, or alternatively, is available from the author on demand.