

ABSTRACT

Title of Document: WEB NAVIGATION STRATEGY AND PERFORMANCE

Susan Grace Campbell, M.A., 2007

Directed By: Associate Professor Kent L. Norman,
Department of Psychology

The task of web navigation, or finding information on the World Wide Web, appears to depend on spatial cognition and problem solving. Spatial visualization ability is commonly considered to determine efficiency of performance on web search and navigation tasks. In order to investigate the mechanism for this improved efficiency, we developed two conceptual models of the relationship between strategy choice and spatial visualization ability. We found mixed results in three experiments. Of the first two, one suggested that spatial visualization ability predicts performance on web navigation tasks, and one suggested that there was no relationship. In both of these experiments, we also found that web navigation task performance was heavily dependent on strategy. The third experiment showed a relationship between strategy choice and performance as well as between spatial visualization ability and performance, but it did not suggest that spatial visualization ability determines strategy choice.

WEB NAVIGATION STRATEGY AND PERFORMANCE

By

Susan Grace Campbell

Thesis submitted to the Faculty of the Graduate School of the
University of Maryland, College Park, in partial fulfillment
of the requirements for the degree of
Master of Arts
2007

Advisory Committee:
Associate Professor Kent Norman, Chair
Professor Thomas S. Wallsten
Associate Professor Michael R. Dougherty

© Copyright by
Susan Grace Campbell
2007

Dedication

This thesis is dedicated to my family and to John, for not telling me to just shut up about my thesis already.

Acknowledgements

Kent Norman contributed enormously to this work, and the rest of my committee suggested important changes. All of the Cognitive Area faculty and students made helpful comments at various seminar presentations. The Laboratory for Automation Psychology undergraduate research assistants, especially Matt Eskowitz, Lauren Hensley, and Sukanya Swetharanyan, did great work in various phases of data collection. My officemates, Ben Smith and Walky Rivadeneira, also helped me verbally debug the software and experiments.

Table of Contents

Dedication	ii
Acknowledgements	iii
Table of Contents	iv
List of Tables.....	v
List of Figures	vi
Chapter 1: Web Navigation Strategy and Performance.....	1
Navigation	2
Spatial Visualization Ability	4
Chapter 2: Experiment 1	16
Method	16
Materials.....	16
Participants	18
Procedure.....	19
Results.....	20
Discussion	22
Chapter 3: Experiment 2	25
Method	25
Materials.....	25
Participants	27
Procedure.....	27
Results.....	27
Discussion	31
Chapter 4: Experiment 3	33
Method	33
Materials.....	33
Participants	33
Procedure.....	34
Results.....	34
Discussion	37
Chapter 5: General Discussion and Conclusion.....	39
Conclusion.....	39
Appendices.....	42
Appendix 1: 3-Option Site Structures.....	42
Appendix 2: 4-Option Site Structures.....	43
References	44

List of Tables

Table 1: Common and different subtasks of Web Navigation and Spatial

Visualization

6

List of Figures

Figure 1: Spatial Visualization Ability (SVA), Central Executive Attention (CE), strategy (strat), VZ-2 result (VZ-2), and navigation performance (nav).	10
Figure 2: Spatial Visualization Ability (SVA), Central Executive Attention (CE), strategy (strat), VZ-2 result (VZ-2), and navigation performance (nav).	11
Figure 3: Example screens for Experiment 1.	17
Figure 4: Link structure diagram of an example site.	18
Figure 5: Pre-test questionnaire.	19
Figure 6: Experiment 1, SVA and average moves to target	20
Figure 7: Spatial Visualization Ability and Average Moves to Target, 3-Option sites (left) and 4-Option sites (right).	21
Figure 8: Spatial Visualization Ability and Proportion Non-Numeric Strategies.	23
Figure 9: Example of Visited Links Condition in Experiments 2 and 3.	26
Figure 10: Average Moves per site in the Visited Links and No Visited Links Conditions	29
Figure 11: Spatial Visualization Ability and Proportion of Revisitations that were the Root Node	29
Figure 12: Spatial Visualization Ability and Average Moves to Target	30
Figure 13: Spatial Visualization Ability and Average Moves to Target by Condition	31
Figure 14: Proportion of Alphabetic Strategies and Average Moves per Site.	32
Figure 15: Spatial Visualization Ability and Average Moves to Target by Condition	35
Figure 16: Spatial Visualization and Average Moves to Target by Condition and Primary Strategy	37

Chapter 1: Web Navigation Strategy and Performance

The World Wide Web (WWW) has exploded in popularity in recent years. As of 2005, the WWW was estimated to have over 11.5 billion pages that could be indexed by search engines, up from around 800 million in 1998 (Gulli and Signorini, 2005). Large amounts of information that, in years past, would have been placed in books or brochures are now available in hypertext format, which allows for cross-linking between documents (Conklin, 1987). Formally, hypertext systems include a set of nodes, which contain information, and a set of links between these information nodes. The links and information are intermingled in particular nodes.

A website (or Web site) is defined by the World Wide Web Consortium as a collection of interlinked pages which reside at the same network location; it is necessary for this definition that each page be accessible by following zero or more links from the home page of the site (Lavoie and Nielsen, 1999). Most websites are part of a larger structure supported by a Local Area Network (LAN) or Wide Area Network (WAN), such as the Internet. Some sites are cut off from access to the world outside a particular network, however, and are thus referred to as intranets. A third type of site is self-contained on removable media, such as a hypertext learning system. People who use websites and other hypertext systems generally access them using a web browser, such as Mozilla's Firefox™ or Microsoft's Internet Explorer™, which interprets the links and renders the content in the nodes.

The Internet itself is intractably large, and most people only use small parts of it at any given time, relying on revisiting familiar pages or on search engines to provide access to particular content (Cockburn and McKenzie, 2001). It also appears

that many people use remembered paths from known sites to locate information over and over (Cockburn and McKenzie, 2001). Once a user has followed a link to a particular site, however, looking through that site becomes the relevant task, rather than looking through the entire Internet. The World Wide Web refers to the collection of hypertext nodes, linked to each other, that populate the Internet.

Understanding how people use the World Wide Web to find information has broader implications for problem-solving behavior in general, as well as being important in the applied domain of automated information retrieval. Knowing the types of ability required to perform hypertext navigation tasks would allow predictions about which task domains are similar, while knowing the problem-solving strategies employed would allow a better understanding of information retrieval performance.

Navigation

Web users spend a large amount of time and effort to navigate the sites they choose to use by following links from information nodes to other information nodes. A verbal protocol study found that as the amount of processing required for navigation on websites increases, the amount of time that can be spent on learning or evaluating the content decreases (Eveland and Dunwoody, 2000). The same study found that users expended an average of 39% of their effort on navigation-related tasks. This finding means that difficult or obtrusive navigation can lead to user problems with learning or using the information found on websites. A verbal protocol approach to analyzing web behavior, however, may overstate the importance of navigation. Requiring users to speak about generally un verbalized spatial tasks may

slow down processing and make those tasks seem to account for more of the total time than they would in practice.

Any tool or method which makes finding information on a particular website easier makes that site more accessible to potential customers, information-seekers, and decision makers. Understanding the cognitive processes involved in hypertext navigation should lead designers to create systems that are more accessible to the majority of users (Boechler, 2001). Properly supporting web exploration and target-finding requires understanding what strategies people employ and why they choose those strategies.

Navigating hypertext is different from navigating through linear text in several ways because there are multiple options for moving from one information node to another. For instance, a user navigating an online index could move from one topic to another related topic by clicking on a link within the first topic's page or by returning to the main index and finding the new topic from there. Hypertext documents and web sites also have the potential to be much larger than printed documents, though users may only traverse limited sections of a document.

Because hypertext can be conceptualized as a multidimensional semantic space, the predominant metaphor for moving between pieces of information is navigation (Boechler 2001). Note that this metaphor, while widespread, is not strictly correct; users are finding and retrieving information rather than literally moving in space (Marchionini, 1995). Though the sense in which the spatial metaphor is used today is not the original sense in which it was intended, the nomenclature has stuck because, on some level, hypertext navigation does seem to share characteristics with

spatial navigation in the real world. The metaphor may not be completely accurate, but Hsu and Schwen (2003) found that including multiple metaphors for a task improved performance for users; a metaphor may be useful even if it is not accurate. Spatial processing is used to find and keep one's place in hypertext spaces, just as it is in physical spaces, and losing one's place leads to disorientation and frustration, just as getting lost in a physical space does (Edwards and Hardman, 1988).

Spatial Visualization Ability

The task of hypertext, or web, navigation can be decomposed in several different ways, leading to different lists of cognitive factors that are important for hypertext performance. The earliest explorations of the task demands for hypertext navigation actually referred to the task demands of hierarchical file systems and databases. Vicente, Hayes, and Williges (1987) found that vocabulary and spatial ability were the best predictors for finding particular files in a hierarchical file system, and that users with low spatial ability were getting "lost" in the structure. Hypertext is slightly more complex in that it is not strictly hierarchical and allows for cross-linking between items, increasing the opportunities for disorientation, but also changing the task demands.

Spatial visualization ability is one of the two main components that have been teased out of general spatial ability. Ekstrom and colleagues defined Spatial Visualization as "The ability to manipulate or transform the image of spatial patterns into other arrangements" in the Kit of Factor-Referenced Cognitive Tests (Ekstrom et al., 1976).

The construct of spatial ability in its modern form appears to have first been identified by Thurstone in the 1930s as a separable part of general intelligence (Thurstone, 1934; Thurstone, 1938). Originally, spatial abilities were largely defined by being distinct from verbal or mathematical abilities, though that definition has been revised since (Pellegrino et al., 1984). Subsequent latent factor analytic studies separated the spatial factor into spatial relations ability and spatial visualization ability. Tests of these two factors can generally be distinguished by degree of difficulty and degree of speededness; the less speeded and more difficult tests tend to measure spatial visualization rather than spatial relations (Pellegrino et al., 1984). Another way to distinguish the two factors is to consider that spatial relations tends to refer to rotation and displacement of the same figure or set of figures, while spatial visualization tends to refer to the folding or other transformative manipulation of stimuli (Michael et al., 1957; Ekstrom et al., 1976). The two factors are highly correlated, but the factors themselves are still separable using factor analysis, and the tasks used to measure them are separable using the criteria above.

Pellegrino and colleagues also suggested that performance on tests of spatial visualization, which tend to be fairly heterogeneous and complex, may include a component of strategy selection, as those with higher measured spatial visualization ability tended to show more consistent data on spatial visualization tasks (Pellegrino et al., 1984).

Therefore, spatial visualization ability can be narrowly defined as the ability to mentally manipulate (not just rotate) possible physical objects. A more broad

definition might include the memory for spatial configurations and the ability to select and implement a strategy on mental manipulation tasks.

Though Vicente and colleagues chose spatial visualization as just one of a set of cognitive ability measures, their explanation of finding a spatial visualization effect appeared to be more spatial-memory related – participants with lower spatial visualization ability were getting lost. Low spatial visualization participants appeared to forget where they were both in the file and in the file hierarchy, though their visual memory scores and spatial scanning performance were not significantly different from the high spatial visualization participants (Vicente et al., 1987).

It is possible, based on these and other results, that tests of spatial visualization ability are an indicator for certain types of processing which are required for navigating websites and finding information on them. Table 1 shows a partial list of similarities and differences between Web Navigation and Spatial Visualization tasks. Of particular note are memory for problem state and representing hidden structure, both of which can apply either to Web Navigation or to Spatial Visualization. The requirements of Web Navigation appear to be similar to those of Spatial Visualization, not in the strictly spatial domain, but in the domain of mental operations required to perform certain tasks.

Table 1: Common and different subtasks of Web Navigation and Spatial Visualization

	Web Navigation	Spatial Visualization
Common Elements	Inferring structure Hypothesizing locations Selecting from a list Remembering problem state	Representing hidden structure Hypothesizing locations Selecting from a list Remembering problem state
Different Elements	Evaluating semantic content Encoding semantic web	Manipulating drawn figures Encoding drawn figures

A requirements analysis of navigating networked user interfaces such as web sites finds that situational awareness, task-set switching, and spatial ability are the most important factors in performance (Neerincx et al, 2001). Situational awareness in this context refers to an awareness of one's current state while searching, whereas task-set switching refers to the ability to switch from a navigational task to a comprehension or learning task.

Chen and Rada (1996), in a meta-analysis, found that users with more "active" cognitive styles, which they defined as having an internal locus of control, field independence, and high spatial ability, performed better on hypertext tasks. Therefore, users who felt that they were in control of their browsing, who could find information regardless of its context, and who could think spatially were more efficient in performing hypertext-related tasks.

Zhang and Salvendy (2001) found that users with high spatial-visualization ability took fewer steps to find targets within an online encyclopedia than users with low spatial-visualization ability. However, in this case, the items were located predictably within an encyclopedic structure. The verbal and spatial elements both predicted the same target locations.

The task decomposition that informs this investigation consists of a verbal-semantic task and a spatial task. Since word semantics are individually variable and hard to control, controlling the semantic information in a task without changing its spatial character should lead to a more complete understanding of the spatial tasks involved.

One must be careful declaring that any task for which spatial processing is used is a spatial task, as spatial processing correlates highly with tasks that are supposed to measure Executive Attention (Miyake et al, 2001). However, it appears that spatial ability is a good predictor of performance on web navigation tasks, especially in efficiency. Spatial navigation ability may or may not predict the absorption of content from websites, but people with more spatial visualization ability tend to navigate information spaces more quickly (Chen and Rada, 1996).

Perhaps recall is predicted by cognitive style rather than spatial ability. Graff (2005) suggested that the important factor for recall of hypertext information, as assessed using hand-drawn concept maps, was the analyst-intuition aspect of cognitive style. Participants who scored at the analyst end of the spectrum tended to recall more elements in the hierarchical condition, while those who scored at the intuitive end recalled more elements in the linear condition, suggesting that people with differences in this particular measure of cognitive style pay attention to different elements of a hypertext structure.

Dünser and Jirasko (2005) found that sequential learners, who like step-by-step instruction, tend to recall more and report less disorientation with a suggested path through a document, rather than a strictly hierarchical organization. Global learners, who approach the task in a more holistic fashion, do not use the suggested path if it is present, and thus its presence or absence does not affect their performance; a global learning style predicts more effective learning from a hierarchical structure.

Task goals can also affect the performance on recall as well as spatial tasks, though different types of goal would lead to different performance predictions. Oulasvirta (2004) found that the best predictor for remembering content was whether the participant was focused on the content or on the navigational task, and attributed this difference to semantic matching involving a higher level of processing than the feature matching that goes on in navigational tasks.

Learning tasks should also be affected by metacognition and metacognitive strategies. Chavez (2001) investigated whether metacognition, verbal working memory, and visuo-spatial working memory affected performance on a hypertext learning task. She found that gamma for metacognitive judgments on domain knowledge questions and visuo-spatial working memory, but not verbal working memory, predicted participants' accuracy of learning from hypertext. This finding suggests that spatial working memory does predict recall, but spatial visualization ability was not measured in this study. It is also unclear whether metacognitive judgments of learning on unrelated questions are an appropriate measure of metacognitive ability.

In a meta-analysis of experimental studies, Chen and Rada (1996) found that spatial ability correlated with measures of efficiency, such as the time to navigate to a particular target, with an effect size, measured with Pearson's r , of approximately $r=0.45$. People with higher measured spatial ability, however, did not seem to be correct more often. Evidence exists that spatial ability predicts efficiency but not accuracy on web navigation tasks, suggesting that it affects the spatial components of the task more than the semantic ones.

The test that we used to assess spatial ability was the VZ-2, a task in which participants are asked to visualize what a folded piece of paper with holes would look like were it to be unfolded (Ekstrom et al, 1976). This test corresponds to a test originally used by Thurstone to measure spatial abilities and later incorporated into test batteries in the 1950s (Michael et al., 1957), though the response mode is now multiple-choice rather than drawing the punched holes on a square.

Were we to measure that participants with higher measured levels of spatial ability were more efficient in traversing hypertext structures, that would replicate established findings. But why do people with higher levels of spatial ability perform more efficiently on hypertext navigation tasks? Do they pick better strategies to begin with, or do they use the same strategies as other people but execute them more effectively? Strategy choice may be part of what is measured in Spatial Visualization Ability, so the answer is not obvious.

Two competing models for this relationship are illustrated in Figures 1 and 2. In Figure 1, both Central Executive Attention and Spatial Visualization Ability predict strategy choice, which influences performance on Web Navigation tasks. In

Figure 1: Spatial Visualization Ability (SVA), Central Executive Attention (CE), strategy (strat), VZ-2 result (VZ-2), and navigation performance (nav).





Figure 2, only Central Executive Attention predicts strategy choice, but both strategy choice and Spatial Visualization Ability predict performance on the Web Navigation task. Unfortunately, Executive Attention is generally measured with spatial tasks, so the results of, for instance, a Tower of London task, would not necessarily measure only Executive Attention in such a way that it could be disambiguated from Spatial Visualization Ability.

Mislevy and colleagues proposed a Bayesian method for distinguishing between simple strategies in a mental rotation task given speed and accuracy responses to a set of test items, provided that the candidate strategies can be mathematically defined (Mislevy et al., 1991). This method of measurement informs the data analysis strategies used in following sections.

Figure 2: Spatial Visualization Ability (SVA), Central Executive Attention (CE), strategy (strat), VZ-2 result (VZ-2), and navigation performance (nav).



We cannot state that the only factor affecting the ability of participants to traverse a website is spatial visualization ability. The selection of a strategy may not be related to spatial ability, but to problem-solving ability, with which it is correlated (Miyake et al., 2001). Early problem-solving theorists suggested that human rationality is bounded, and that participants would select strategies that would require

fewer resources rather than those that were computationally intensive (Simon, 1990). It is possible, however, that people who have more resources in a particular domain, such as spatial manipulation, may find it less necessary to conserve those resources by prudent strategy choice; if a person can recall something without writing it down, why expend the energy to write it down in the first place?

Problem solving strategies influence performance on any number of tasks. The selection of an inefficient strategy can impair performance, while the selection of a more efficient strategy can lead to faster performance. From earlier findings about the speed of web navigation (Norman and Butler, 1989; Chen and Rada, 1996; Zhang and Salvendy, 2001), we expected that people with higher measured spatial visualization ability would be faster at traversing the menus, and hypothesized that this speed increase came from more efficient uses of strategy.

For instance, people with better spatial visualization ability generally do not revisit the top-level node in a hierarchy, and therefore they perform better on navigation tasks (Campagnoni and Ehrlich, 1989; Norman and Butler, 1989). Campagnoni and Ehrlich initially thought that people with higher levels of spatial visualization ability might use prior knowledge about semantic structures to find information in a hypertext help system. However, it appeared that people with higher levels of spatial visualization ability were similarly likely to other participants in the study to use a strategy of reading the category labels and responding to them; spatial visualization ability does not seem to determine whether one matches results to a query based on keywords or based on prior knowledge (Campagnoni and Ehrlich, 1989).

Performance on any task relies on some combination of task difficulty, strategy use, and participant ability. In the case of web navigation, the task difficulty can be operationalized in the structure and number of options in the site to be navigated and the location of the target in the structure. Strategy use and selection refer to choosing between using all of the information presented and minimizing cognitive effort, as well as the more local choice to revisit top-level nodes or more proximal nodes to one's current position. In this case, participant ability refers to spatial visualization ability.

The task we used in this set of experiments involved asking participants to navigate a set of web pages that were organized strictly hierarchically. The task was based on an experiment developed in 1988 to assess menu navigation performance in HyperCard systems (Norman and Butler, 1989). The task used in that study was implemented as a HyperCard navigation system that used a file cabinet metaphor. Items were said to be located in a particular cabinet, particular drawer, and particular folder within that drawer. Each cabinet, drawer, and folder was marked with the probability that it contained the target item.

The task adapted for use in these experiments involved giving participants a well-defined target, links labeled with a non-informative name, such as "Section A," and a conditional probability of finding the target down a particular branch. Names were non-informative in order to standardize and control the semantic navigation component of the task as much as possible. For this specific implementation, the spatial structure of each site was constant – there were no cross-links between

sections, and all items were strictly hierarchical. Once a participant found the target in a particular site, she could go on to the next site.

Results from Norman and Butler (1989) suggested that participants with higher levels of spatial ability were more likely to use breadth-first strategies, visiting all of the child nodes of a particular node before moving on to the next, rather than using depth-first strategies, restarting from a higher-level node each time. Participants with a higher spatial visualization ability were less likely to revisit the initial node and less likely to engage in random search of the database. Note that revisiting the initial node may lead to finding the target faster or more efficiently than exploring lower-probability results in the same branch as high-probability results, but that it is generally slower. Norman and Butler did not report any use of strategies that did not take the stated probability into account. They also capped the number of steps taken by each participant in a given menu system at 50.

Evidence from our pilot studies and first experiment, however, suggest that some participants use systematic, non-spatial strategies which do not take the given relevance information into account in order to reduce the memory load required for the task of navigating the World Wide Web. Strategies of this type tend to involve starting at the first section (Section A), first subsection (Subsection A), and first page (Page A), then moving to the second page (Page B), then the third (Page C), though some participants in the pilot study also started from the last section (Section C) and so forth. No attempt is made to take into account probability information in choosing what pages to visit. These alphabetic strategies may not be the most efficient in number of pages visited, but may be similarly quick to informed strategies (depth-

first or breadth-first) when the number of nodes to be visited is small, as it is in this probabilistic navigation task. Systematic, alphabetic strategies have the advantage of allowing participants to be reasonably certain that they have not missed a node in their exploration and do not require much cognitive effort.

In our second experiment, we automatically marked which nodes had been visited for half of the participants. We expected that when the influence of memory was removed from the task, not only would performance be faster, but the spatial nature of the task would also lead participants' spatial ability to better predict performance, measured both in nodes visited and in time. We further expected that participants who were given information about which nodes they had already visited would be less likely to use alphabetic strategies because they could be more certain of which nodes they had already visited.

The specific hypotheses for this second experiment were that people with higher spatial visualization ability would be more likely to choose a strategy that took into account the probabilities and better able to apply that strategy, and that marking links that had been visited would allow spatial visualization ability to correlate more strongly with performance, as the need for memory effort would be reduced.

Because of methodological problems with Experiment 2, Experiment 3 was conducted. Experiment 3 was identical in design to Experiment 2, with tighter experimenter control of the testing environment. Experiment 3 also included some different data analysis techniques, though the task was nearly identical.

Chapter 2: Experiment 1

In Experiment 1, we investigated the relationship between Spatial Visualization Ability, measured with the VZ-2 scale, and number of moves required to find the target in a site. We also investigated the strategies used by participants to accomplish this task and how those strategies related to the number of moves required to find the target in a site.

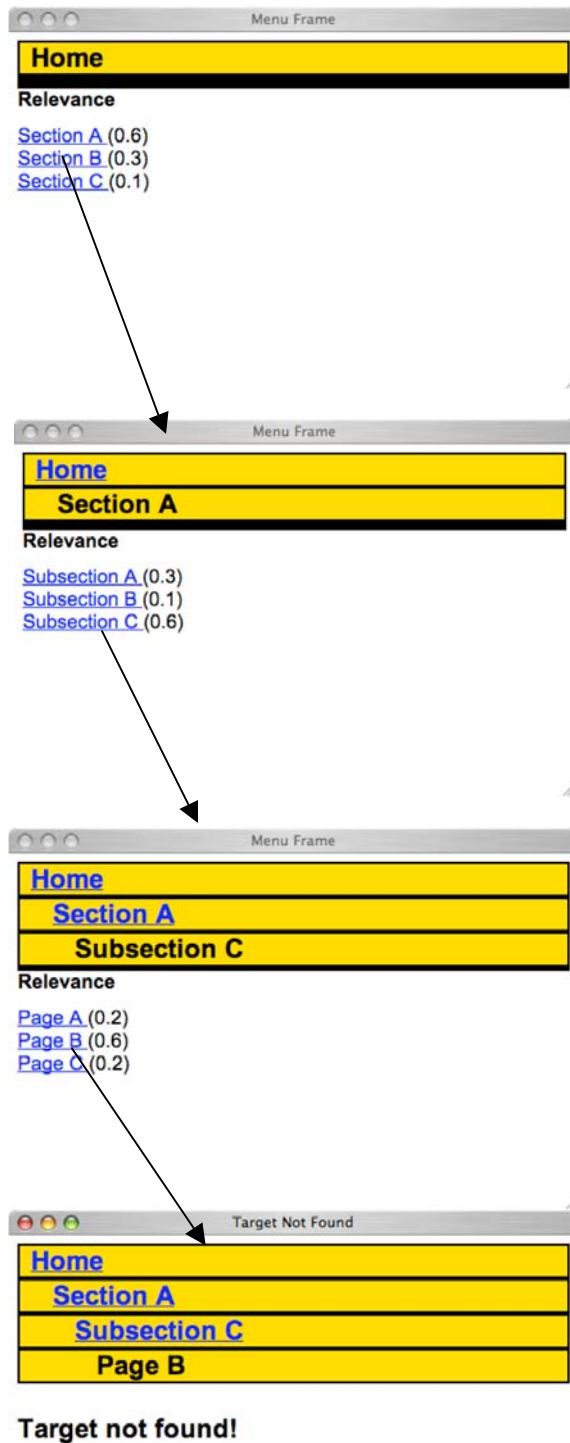
Method

Materials

The materials for this experiment were 20 menu structures adapted from Norman and Butler (1989) (See Appendix 1 and 2). Each structure contained a set of conditional probabilities and a target location. The order of structures was randomized for each participant. The participants were also given an online version of the VZ-2 (Ekstrom et al., 1976).

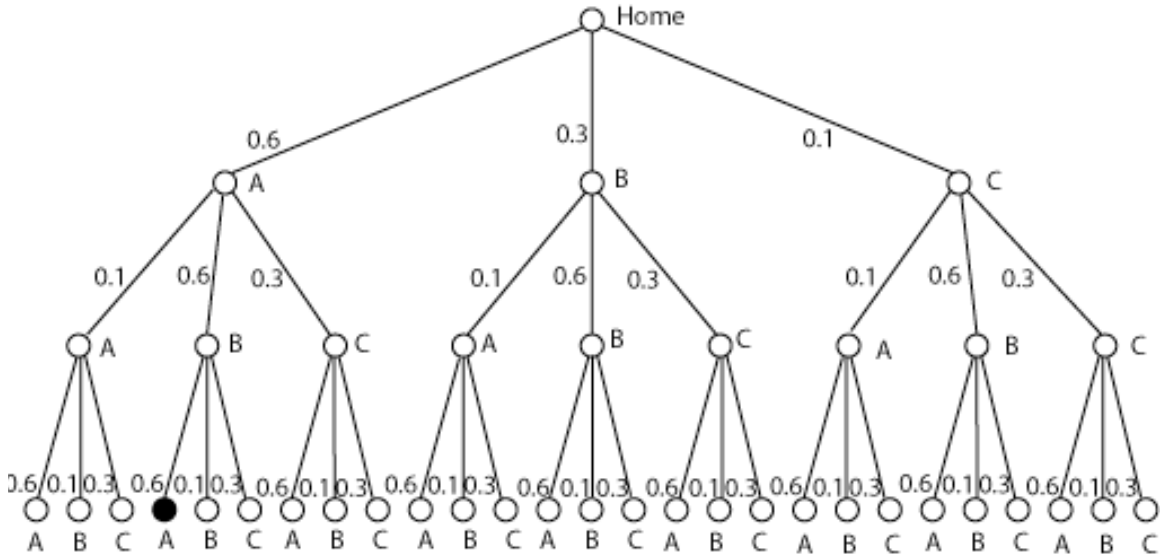
Filemaker™ Pro 6.0 was used as to serve the web pages to the client computers. Participants viewed the web pages on Macintosh™ computers with two 15-inch LCD monitors running the Safari™ (Version 2.0.3) web browser. The browser was placed on one monitor and was configured not to display the “forward” or “back” buttons or the address bar (See Figure 3). The menu structures were presented as web pages, such that each menu structure represented a web site. On each page, participants were allowed either to choose a node that was a child of the page displayed, or to move up in the structure to any parent node, including the Home

Figure 3: Example screens for Experiment 1.



node. A sample structure with the probability information included can be seen in Figure 4, though participants never saw the structures in this format.

Figure 4: Link structure diagram of an example site.



Participants

Twenty undergraduate students at the University of Maryland received optional course credit for participating in this experiment. One participant did not finish the experiment in the time allotted and was removed from the analysis because of incomplete data, including missing a VZ-2 score. The participants that did finish ranged in age from 18 to 20, with a mean age of 18.6. Twelve were female and seven were male. When asked to rate their experience with “overall use of computers” on a 10-point Likert scale with the endpoints “no experience” (1) and “very experienced” (10), their mean response was 6.7. When asked to rate their experience with “use of the World Wide Web” on the same scale, their mean response was 6.9 (for exact format of pre-test questionnaire, see Figure 5).

Figure 5: Pre-test questionnaire.

Pre Questionnaire

Pre-Test Questionnaire

Please provide the following information:

Your age:

Your sex: Female
 Male

Your major/occupation:

Please rate your experience with the following:

	no experience									very experienced
Overall use of computers	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Use of the World Wide Web	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>


Procedure

Participants came into the laboratory for up to an hour. As few as one or as many as three participants could be run at a time. After completing the consent form, participants were first given an explanation of how to navigate the sample web pages and instructed that the probability information was important. They were then given the opportunity to look for the target in each of the 20 menus in random order. The system recorded each page the participants visited and how long they spent on that page. They were not allowed the option to stop browsing before the target was found, and each site search terminated when the target was found. The 20 menus were followed by an online version of the VZ-2 Spatial Visualization Ability test (Ekstrom et al., 1976). After they completed the VZ-2, participants were asked what kind of strategy they had used to traverse the pages.

Results


We found that VZ-2 scores may have correlated with the total number of moves, but the correlation was non-significant ($r=-0.392$, $p=0.10$; see Figure 7). This

Figure 6: SVA and average moves to target



effect appears to have been due to participants with higher VZ-2 scores using fewer moves on the smaller (3-options per level) menus ($r=-0.573$, $p<0.01$; see Figure 7), but not necessarily on the larger (4-options per level) menus ($r=-0.219$, $p=0.37$; see Figure 7). Note that there was one outlying score at (11, 113) not shown on the graph. Removing the outlier would produce very little change in the overall correlation ($r=-0.389$), the correlation on the 3-option menus ($r=-0.611$), or the correlation on the 4-option menus ($r=-0.185$).

Figure 7: Spatial Visualization Ability and Average Moves to Target, 3-Option sites (left) and 4-Option sites (right).



The strategy used was a more useful predictor than VZ-2 score for the number of moves used in a particular menu. Each traversal was coded by hand as either a numeric or alphabetic strategy. A numeric strategy referred to any strategy where the participant appeared to have used the probability information provided, whether that use was effective or not, while an alphabetic strategy referred to the specific strategy of going through the menu systematically in alphabetical order (AAA, AAB, AAC, ABA, ABB, etc.) from either the first link option or the last link option. This coding yields a proportion of menus traversed by each participant using an alphabetic strategy. Because of the limitations of the data analysis, this estimate does not take into account situations where the participant started in a numeric strategy and switched partway through to an alphabetic strategy; the proportion is therefore likely to be a conservative estimate of the proportion of the time that the participants used alphabetic strategies. This analysis suggested 3 participants out of 17 used an

alphabetic strategy all of the time, and a further 7 out of 17 used an alphabetic strategy at least part of the time.

This strategy factor correlates with moves taken to traverse the entire set of structures ($r=0.476$, $p=0.04$), as well as with moves to traverse the smaller 3-Option ($r=0.459$, $p=0.048$) and larger 4-Option ($r=0.520$, $p=0.02$) hypertext structures; the more often a person used the alphabetic strategy, the more moves the person took on average to traverse a menu structure. Because these data were proportions and thus do not conform to the normality assumptions of correlational analysis, they were also analyzed after being subjected to the arcsine transformation for proportions (Cohen et al., 2003). The transformed data showed the same relationship, with the transformed strategy variable correlating with the number of moves taken on average overall ($r=0.459$, $p=0.048$), non-significantly with the number of moves taken on average in the 3-Option menus ($r=0.443$, $p=0.057$), and the number of moves taken in the 4-Option menus ($r=0.502$, $p=0.029$).


The VZ-2 score and strategy use were also correlated; the higher the VZ-2 score, the less likely the person was to use non-informed strategies to choose alternatives ($r=-0.487$, $p=0.034$; See Figure 8). This correlation also held with the transformed proportions ($r=-0.476$, $p=0.039$).

Discussion

It appears that spatial visualization ability predicts performance on simple navigational tasks where a certain number of options can be eliminated immediately. In situations like the 3-level menus, where there are only twenty-seven possible target locations, spatial visualization ability appears to predict performance. In larger

structures or structures where fewer options can be eliminated, spatial visualization ability does not appear to predict performance as well. It is unclear why this is the

Figure 8: Spatial Visualization Ability and Proportion Non-Numeric Strategies



case, but it is possible that memory or strategy choice is more important than spatial visualization ability as structures get larger.

Strategy use, in the narrow sense of using versus not using the relevance information provided, did predict efficiency of performance on the task. The less a participant used the relevance information, the more moves she used to traverse the menu. This finding seems self-evident, since the average number of moves required to traverse the menus is higher in the alphabetic case, even if that strategy is applied perfectly, than it is in the informed case. Participants switched between the two types of strategy freely, but it is unclear what criteria they used to choose those strategies. Some suggested in their verbal responses that they used the relevance information only when the most likely option was considerably more likely than the others.

VZ-2 score and strategy use also appeared to be related in this experiment. The higher the participants' VZ-2 scores were, the less likely they were to use alphabetic strategies. This finding supports the idea that people with more spatial visualization ability are more likely to choose good, but potentially more taxing, strategies in this task.

Chapter 3: Experiment 2

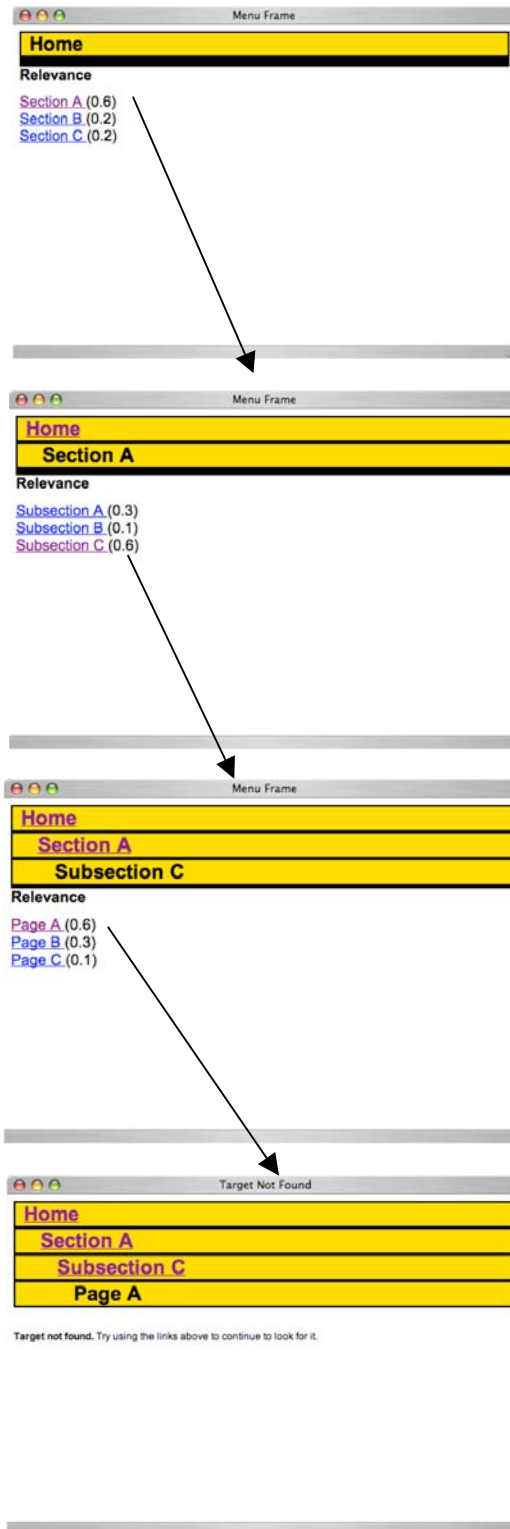
In this experiment, we intended to follow-up Experiment 1 by examining whether making the task less demanding on memory would increase the effect of spatial ability on performance. We therefore marked the branches of the hypertext structure that had been visited to reduce the need to remember which branches had been visited within the current section. Note that memory and spatial ability would still be required to remember whether branches had been visited in other sections. The link-marking scheme was a standard scheme used by many websites, where unvisited links were blue and visited links were purple. This scheme has been the default since the earliest web browsers, and is generally understood by most web users, though it is not the only possible scheme for providing link information (Weinreich and Lamersdorf, 2000). We also hoped to replicate the findings of Experiment 1.

Method

Materials

The same materials were used for Experiment 2 as Experiment 1, with the exception that participants in Condition 1 were shown versions of the pages that marked visited links in purple and there were some cosmetic changes in the page design (See Figure 9). The question at the end of the experiment was also reworded to ask more specifically about the web navigation task and not the paper-folding VZ-2 task.

Figure 9: Example of Visited Links Condition in Experiments 2 and 3.



Participants

Twenty-nine undergraduate students at the University of Maryland received optional course credit for participating in this experiment. They ranged in age from 18 to 22, with a mean age of 19.7 years. Nineteen were female, and ten were male. When asked to rate their experience with “overall use of computers” on a 10-point Likert scale with the endpoints “no experience” (1) and “very experienced” (10), their mean response was 7.9. When asked to rate their experience with “use of the World Wide Web” on the same scale, their mean response was 7.9. Their median VZ-2 score was 13. Of the participants, 16 participated in the Visited Links Condition, and 13 participated in the No Visited Links Condition. There were no significant differences in gender, age, or reported computer or web experience between the two groups.

Procedure

The method was the same as for the first experiment, except that the participants in the Visited Links Condition were given instructions that the links to particular pages would change color when those pages had been visited. The question about strategy use was also clarified to refer to the web navigation task, rather than the aptitude test, as some participants in the first study had stated that their strategy involved visualizing the paper unfolding.

Results


Due to an experimenter error in which the browser history was not reset between participants, an indeterminate number of participants in the Visited Links Condition were given information (in the form of colored links) about what paths previous participants had pursued. The marked links may have included paths which did not lead to the target, but they always included the target location. After a couple

of menus, many participants probably realized that any link which was not colored purple did not contain the target. It was unclear what effect this problem with the manipulation had on the data for the Visited Links Condition, but it was clear that the experiment needed to be rerun. The problem should not have affected the first few participants in the Visited Links Condition, but had the potential to affect at least 10 of the participants in that condition. It could have produced positive performance effects (higher efficiency) if the participants realized that the link information highlighted a subset of the hypertext structure that contained the target. If they believed that the information was not meaningful, on the other hand, that would have led to performance more similar to the No Visited Links Condition.

The number of moves required, on average, to traverse a menu was different between the Visited Links and No Visited Links Conditions ($t(26.6)=-3.91, p=0.001$). A 95% confidence interval of the difference suggests that people in the Visited Links Condition found the target with between 5.4 and 17.9 fewer moves than those in the No Visited Links Condition. This difference appears to be driven by a more pronounced difference in the larger (4-option) structures than in the smaller (3-option) structures ($t(27)=-5.02, p<0.001$; $t(27)=-1.34, p=0.19$; See Figure 10). The amount of time required to traverse each menu was also lower in the Visited Links Condition ($t(27)=-3.99, p<0.001$).


We also found that participants in the Visited Links Condition were less likely to revisit possible target locations than participants in the no Visited Links Condition, as shown by the difference in their mean number of possible target locations revisited ($t(24.6)=-3.97, p=0.001$).

Figure 10: Average Moves per site in the Visited Links and No Visited Links Conditions




Norman and Butler (1989) found that participants with higher VZ-2 scores were less likely to revisit the root node. In investigating whether this was the case in the current experiment, we found that it was possible that participants with higher VZ-2 scores were less likely to revisit the Home Page, but that the correlation was not significant, so no conclusions could be drawn (see Figure 11).

Figure 11: Spatial Visualization Ability and Proportion of Revisitations that were the Root Node



Unlike in the first experiment, we found no correlation whatsoever between the VZ-2 and performance on any measure in the second experiment. The correlation between VZ-2 score and number of moves was positive and near zero ($r=0.084$, $p=0.67$; see Figure 12), and the correlation between VZ-2 score and time to traverse each menu was also positive and near zero ($r=0.118$, $p=0.54$). A positive correlation would suggest that participants whose measured spatial ability appeared higher were slower and took more steps to find the target on this task.


Figure 12: Spatial Visualization Ability and Average Moves to Target



The correlation between VZ-2 and number of moves required did appear to be stronger in the Visited Links Condition ($r=-0.158$) than in the No Visited Links Condition ($r=0.059$), but we cannot draw any conclusions based on this weak finding, since neither was significantly different from zero (See Figure 13).

The correlation between observed strategy and number of moves required remained higher than the correlation between the VZ-2 and number of moves required. In this experiment, the correlation between observed alphabetic strategy

Figure 13: Spatial Visualization Ability and Average Moves to Target by Condition



uses and number of moves required to traverse a structure was quite high ($r=0.629$, $p<0.001$; See Figure 14). The transformed correlation was not quite as high ($r=0.619$, $p<0.001$). Note that Figure 14 appears to show two groups, the smaller of which used more alphabetic strategies and was also less efficient, so this correlation cannot be assumed to show a linear effect.

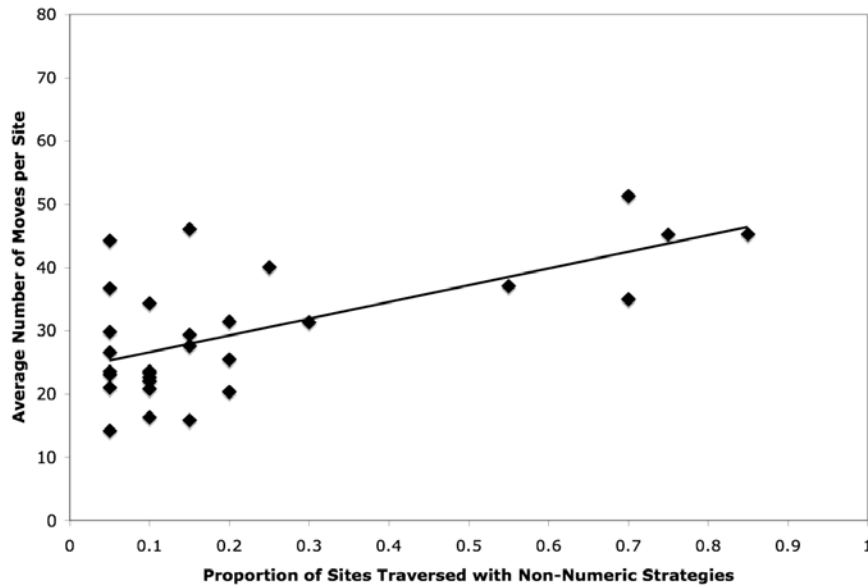
Unlike in Experiment 1, there was no correlation in this data set between strategy use and VZ-2 score ($r=0.016$, $p=0.936$). The transformed value was even lower ($r=0.003$, $p=0.989$).

Discussion

The hypothesis suggested by the first experiment, that spatial visualization ability affects performance in the 3-option structures more than in the 4-option structures, was not supported in this second experiment.

We did not replicate the findings in the literature or from the first experiment. It is conceivable that the methodological issues contributed to this situation, but it is

Figure 14: Proportion of Alphabetic Strategies and Average Moves per Site



also possible that there is no actual relationship between spatial ability and performance on this task.

It does appear that the manipulation worked, inasmuch as marking which links had been visited reduced the number of revisits of previous pages. Participants were faster and more efficient in the Visited Links Condition, though it is possible that some of that was due to information which carried over from prior participants who had already successfully performed each task.

The use of strategy did predict performance in this task as much as or more than it did in Experiment 1, but it did not appear that people with more spatial visualization ability picked alphabetic strategies less frequently. This effect could partially be due to the data contamination; alphabetic strategies which are restricted to sets of nodes visited by someone who was using an informed strategy are difficult to distinguish from informed strategies, and without knowing what the pages those participants saw looked like, we can draw no conclusions about their behavior.

Chapter 4: Experiment 3

Experiment 3 was a replication of Experiment 2 without the data contamination problem from the sharing of other participants' visited link history. Strategy use was also classified more rigorously than in previous experiments, and slightly different data analysis was performed.

Method

Materials

The same materials were used for Experiment 3 as Experiment 2. The web browser cache was removed between participants, and the stated strategies were checked to verify that the participants were not using previous link information.

Participants

Fifty-seven undergraduate students at the University of Maryland received optional course credit for participating in this experiment. They ranged in age from 18 to 24 and had a mean age of 20.2 years; twenty-six were female and thirty were male. When asked to rate their experience with "overall use of computers" on a 10-point Likert scale with the endpoints "no experience" (1) and "very experienced" (10), their mean response was 7.9. When asked to rate their experience with "use of the World Wide Web" on the same scale, their mean response was 8.1. Their mean VZ-2 score was 13.2. Of the participants, 33 participated in the Visited Links Condition, and 24 participated in the No Visited Links Condition. There were no significant differences in gender, age, VZ-2 score, or reported computer or web experience between the two groups.

Procedure

The procedure for Experiment 3 was the same as that for Experiment 2.


Results

Participants in the Visited Links Condition found the target in significantly fewer moves than those in the No Visited Links Condition ($t(55)=2.63$, $p<0.05$), though a 95% confidence interval finds that they were between 1.6 and 11.9 moves faster. The mean number of moves to find the targets in the Visited Links Condition was 26.5, while in the No Visited Links Condition the mean was 33.3.

Much of the difference between the Visited Links Condition and the No-Visited Links Condition can be attributed to the difference in the number of possible target locations that participants revisited. Participants in the Visited Links Condition revisited a mean of 0.2 target locations per site, whereas participants in the other condition revisited a mean of 2.1 target locations per site, giving a 95% confidence interval of the difference as between 1.5 and 2.3 visits per site ($t(26.3)=9.53$, $p<0.01$). In terms of proportions of possible target locations visited, 2% of visits to possible target locations in the Visited Links Condition were revisits, as opposed to 17% in the No-Visited Links Condition.

The overall correlation between participants' scores on the VZ-2 and the mean number of moves that they took to get to targets was a modest -0.239 ($p>0.05$), though the correlation in both conditions was higher: -0.318 ($p>0.05$) in the condition without link markings and -0.353 ($p<0.05$) in the Visited Links Condition (See Figure 15). This finding is not out of line with the -0.45 correlation reported in the literature in more complex tasks.

Figure 15: Spatial Visualization Ability and Average Moves to Target by Condition



The Butler and Norman (1989) findings suggested that revisiting the “Home” or top node was more common in participants with lower VZ-2 scores. In this experiment, we found a small negative correlation (-0.223 , $p > 0.05$) between participants’ VZ-2 scores and the number of times that they revisited the Home node in each site. This result appears to be due to participants with lower VZ-2 scores being more likely to use a relevance-based strategy rather than a distance-based strategy, as described below.

To determine what strategy participants were using in each case, we compared the order in which they visited particular nodes. Each node the participant visited in each site was marked with the rank of its first visit. For instance, the Home node always had a rank of 1 because it was always the first node visited. These obtained ranks were compared with generated prototype ranks, and a similarity metric using the correlation between ranks was computed. Using these rank correlations, each participant’s mean similarity to each prototype could be computed. The alphabetic


strategy was not very popular in this replication – only 2 participants out of 57 used an alphabetic strategy. The rest appeared to use either a relevance-based strategy where they attempted to find the next most probable location after a failure to find the target at a leaf node or a distance-based strategy in which they attempted to find the next most probable location one level up from where they failed to find the target. These strategies could also be construed as being a global strategy or a local strategy, taking into account either the relevance scores across the entire site or just those in the local section of the site that the participant was currently browsing. There were 21 participants who appeared to use the relevance-based strategy most often, and 34 who appeared to use the distance-based strategy most often.

Unsurprisingly, adhering well to these strategies appeared to produce better performance, as measured in number of moves taken to find the target. The correlation of participant similarity to the relevance-based strategy with the number of moves they took to reach the target was -0.765, while similarity to the distance-based strategy produced a correlation of -0.569. Similarity to the alphabetic strategy, by contrast, produced only a -0.13 correlation with the number of moves required. Therefore, it appears that the less randomness there was in a participants' strategy, the fewer moves they used to find the target.

Surprisingly, higher VZ-2 scores did not seem to predict a better strategy choice – the relevance-based strategy turned out to be fastest overall, taking a mean number of moves of 26.2, as opposed to the distance-based strategy, which had a mean number of moves of 29.8. The mean VZ-2 score of those using the relevance-based strategy, however, was only 12.6, as opposed to the mean VZ-2 score of those

using the distance-based strategy, which was 13.7. Neither of these differences was significant, however they are suggestive.

Figure 16: Spatial Visualization and Average Moves to Target by Condition and Primary Strategy



Splitting the participants by strategy as well as by condition also improves the observed correlation between VZ-2 score and number of moves taken to reach the target (See Figure 16, which omits the alphabetic strategy because it only has two participants).

Discussion

These results suggest that link-marking is a good idea to prevent site users from getting lost and revisiting pages they have already visited. It appears well suited to preventing repetitions within a particular search.

We found a slightly lower correlation (-0.26 rather than -0.45) than we expected between VZ-2 scores and efficiency, as measured in clicks per site, which improved when the variance accounted for by condition was removed.

The findings from Experiment 3 appear to suggest that high VZ-2 participants tend to succeed in applying strategies despite their strategy-selection ability. This suggests that the efficiency of participants with high Spatial Visualization Ability has more to do with their ability to apply strategies than with their ability to pick strategies.

Chapter 5: General Discussion and Conclusion

In all three experiments, it was clear that strategy choice and implementation acted as stronger determinants of performance in this task than spatial visualization ability. Spatial Visualization Ability and strategy choice correlated with each other strongly in Experiment 1 and not at all in Experiment 2, which led to a difficulty in interpreting whether strategy choice is in fact related to spatial visualization ability or orthogonal to it. In Experiment 3, participants used alphabetic strategies less frequently, removing what may have been an artificial inflation of the correlation between strategy choice and efficiency in Experiments 1 and 2. Experiment 3 was also subjected to a finer-grained analysis of strategy choices, which appeared to show that strategy choice was not significantly related to Spatial Visualization Ability. Strategy implementation, however, did seem to be related to spatial visualization ability.

Conclusion

It appears that, unsurprisingly, using a strategy that takes into account more information leads to better performance on a spatial web navigation task. People with higher levels of spatial visualization ability, when presented with situations like the task described here, tend not to pick strategies that discard information. However, they are not necessarily better at choosing between informed strategies than people with lower levels of spatial visualization ability.

The result that relevance information is discarded in certain situations leads to predictions about whether providing numeric indicators or other indicators would lead to better performance in search engines; for participants who are confident of their

spatial visualization ability or confident of their problem-solving ability, these numbers would help. For participants with lower spatial visualization ability or less experience with computers, however, the numbers do not add any information. Note that current search engines do provide this information in an ordinal fashion by listing the results in order from most likely to be relevant to least likely to be relevant.

We found here that spatial ability does improve efficiency on web navigation tasks to a small degree, and therefore spatial interventions can be appropriate for improving performance on website navigation. Despite the inappropriateness of the original spatial metaphor, spatial processing and spatial ability do appear to be at work in the non-linguistic spaces of websites as well as the linguistic-semantic space.

A useful exploration would involve determining what pushes participants to adopt a less-informed strategy over a more informed one, as well as which sites participants themselves found most difficult. A finer-grained analysis of strategy-switches should provide more information about where and why participants switch their strategy during web navigation. Does it come from a frustration at being unable to find the target, or a memory failure, or is there some other explanation?

It is possible that the mode of presentation of the probability information intimidated or confused some of the participants, leading to the adoption of a non-probabilistic strategy. Norman and Butler (1989) found that a bar graph presentation produced fewer deviations from the optimal path to the target, but it is unclear by what criteria they designated the optimal path, as using the bar graph method did not actually result in fewer moves to find the target.

The next step in investigating the factors leading to strategy choice would be determining what factors lead to choosing a particular strategy, besides spatial visualization ability, which does not seem to predict strategy choice. An example manipulation could be accomplished by systematically varying whether the initial menus showed a lower probability for the actual target location. Varying the number of options at each level, to replicate findings on depth and breadth issues in menu selection and web design, should also yield differences in strategy.

Adding a reward for finding the target quickly would also lead to a better idea of what the participants consider optimal, rather than what they do by default. We could then compare the performance of participants who are rewarded for finding the target with those who are rewarded for finishing quickly.

Another way of manipulating the memory and spatial demands of this task would be to vary the amount of support provided to the participants. One method of accomplishing this manipulation would be to vary the presence or absence of a site map or other methods of viewing the site structure. An alternative to marking particular sections and pages as visited would also be to include the number of unvisited pages remaining in a particular branch of the site.

Once we can determine what factors influence the selection of strategies and the resultant performance with those strategies on web navigation tasks, we can better decide what information to present to users and how much of that information should be presented. We will also be better able to understand performance on web navigation and related tasks.

Appendices

Appendix 1: 3-Option Site Structures

Menu	Level	Option A (01)	Option B (02)	Option C (03)	Target Location
M01	Section	0.6	0.3	0.1	M01-01-02-01
	Subsection	0.1	0.6	0.3	
	Page	0.6	0.1	0.3	
M02	Section	0.6	0.3	0.1	M02-01-02-03
	Subsection	0.1	0.6	0.3	
	Page	0.6	0.1	0.3	
M03	Section	0.3	0.6	0.1	M03-02-03-03
	Subsection	0.1	0.3	0.6	
	Page	0.6	0.1	0.3	
M04	Section	0.6	0.3	0.1	M04-02-02-01
	Subsection	0.3	0.6	0.1	
	Page	0.6	0.3	0.1	
M05	Section	0.6	0.2	0.2	M05-01-03-01
	Subsection	0.1	0.6	0.3	
	Page	0.6	0.1	0.3	
M06	Section	0.6	0.3	0.1	M06-02-02-01
	Subsection	0.2	0.6	0.2	
	Page	0.6	0.1	0.3	
M07	Section	0.6	0.2	0.2	M07-03-03-01
	Subsection	0.3	0.1	0.6	
	Page	0.6	0.3	0.1	
M08	Section	0.3	0.6	0.1	M08-02-02-01
	Subsection	0.2	0.2	0.6	
	Page	0.6	0.1	0.3	
M09	Section	0.6	0.3	0.1	M09-01-01-03
	Subsection	0.3	0.1	0.6	
	Page	0.2	0.6	0.2	
M10	Section	0.6	0.3	0.1	M10-01-02-03
	Subsection	0.3	0.6	0.1	
	Page	0.6	0.2	0.2	

Appendix 2: 4-Option Site Structures

Menu	Level	Option A (01)	Option B (02)	Option C (03)	Option D (04)	Target Location
M11	Section	0.7	0.1	0.1	0.1	M11-01-02-02
	Subsection	0.4	0.4	0.1	0.1	
	Page	0.7	0.1	0.1	0.1	
M12	Section	0.4	0.1	0.1	0.4	M12-04-01-02
	Subsection	0.7	0.1	0.1	0.1	
	Page	0.1	0.7	0.1	0.1	
M13	Section	0.7	0.1	0.1	0.1	M13-01-04-03
	Subsection	0.1	0.1	0.1	0.7	
	Page	0.4	0.1	0.4	0.1	
M14	Section	0.4	0.3	0.2	0.1	M14-01-02-03
	Subsection	0.1	0.7	0.1	0.1	
	Page	0.1	0.1	0.7	0.1	
M15	Section	0.7	0.1	0.1	0.1	M15-01-02-02
	Subsection	0.4	0.3	0.1	0.2	
	Page	0.1	0.7	0.1	0.1	
M16	Section	0.1	0.7	0.1	0.1	M16-02-04-03
	Subsection	0.1	0.1	0.1	0.7	
	Page	0.4	0.3	0.2	0.1	
M17	Section	0.1	0.3	0.4	0.2	M17-02-01-02
	Subsection	0.7	0.1	0.1	0.1	
	Page	0.1	0.7	0.1	0.1	
M18	Section	0.4	0.3	0.2	0.1	M18-01-02-02
	Subsection	0.3	0.4	0.1	0.2	
	Page	0.1	0.3	0.2	0.4	
M19	Section	0.4	0.2	0.3	0.1	M19-02-03-03
	Subsection	0.1	0.2	0.4	0.3	
	Page	0.1	0.4	0.3	0.2	
M20	Section	0.1	0.2	0.4	0.3	M20-02-02-01
	Subsection	0.4	0.2	0.1	0.3	
	Page	0.3	0.1	0.4	0.2	

References

- Boechler, P.M. (2001). "How spatial is hyperspace? Interacting with hypertext documents: Cognitive processes and concepts." *CyberPsychology & Behavior* 4(1), 23-46.
- Campagnoni, F.R., and Ehrlich, K. (1989). "Information retrieval using a hypertext-based help system." *ACM Transactions on Information Systems* 7(3), 271-291.
- Chavez, N.M. (2002). "Individual differences in verbal working memory, visuo-spatial working memory, and metacognition: Learning from text in a hypertext environment." Unpublished Dissertation AAI3056101, Texas Tech University.
- Chen, C., & Rada, R. (1996). "Interacting with hypertext: A meta-analysis of experimental studies." *Human-Computer Interaction* 11(2), 125-156.
- Cockburn, A., and McKenzie, B. (2001). "What do web users do? An empirical analysis of web use." *International Journal of Human-Computer Studies* 54, 903-922.
- Cohen, J., Cohen, P., West, S.G., and Aiken, L.S. (2003). *Applied Multiple Regression/Correlation Analysis for the Behavioral Sciences, Third Edition*. Mahwah, NJ: Lawrence Erlbaum Associates.
- Conklin, J. (1987), "A survey of hypertext." *Hypertext on Hypertext*. Association for Computing Machinery.
- Dünser, A., and Jirasko, M. (2005). "Interaction of hypertext forms and global versus sequential learning styles." *Journal of Educational Computing Research* 32(1), 79-91.
- Edwards, D. M. and Hardman, L. (1988). "Lost in hyperspace: cognitive mapping and navigation in a hypertext environment." In *Hypertext: theory into practice*, R. McAleese, Ed. Intellect Books, Exeter, UK, 90-105.
- Ekstrom, R.B., French, J.W., & Harman, H.H. (1976). *Manual for kit of factor-referenced cognitive tests*. Princeton, NJ: Educational Testing Service.
- Eveland, W.P. Jr., and Dunwoody, S. (2000). "Examining information processing on the world wide web using think-aloud protocols." *Media Psychology* 2, 219-244.
- Graff, M. (2005). "Differences in concept mapping, hypertext architecture, and the analyst-intuition dimension of cognitive style." *Educational Psychology* 25(4), 409-422.

- Gulli, A., and Signorini, A., "The indexable web is more than 11.5 billion pages." *WWW 2005*.
- Hsu, Y., and Schwen, T.M. (2003). "The effects of structural cues from multiple metaphors on computer users' information search performance." *International Journal of Human-Computer Studies* 58, 39-55.
- Lavoie, B., and Nielsen, H.F., Eds. (1999). "Web characterization terminology & definitions sheet: W3C Working Draft 24-May-1999." W3C (MIT, INRIA, Keio). URL: <http://www.w3.org/1999/05/WCA-terms/01>
- Marchionini, G. (1995). *Information Seeking in Electronic Environments*. New York: Cambridge University Press.
- Michael, W.B., Guilford, J.P., Furchter, B., and Zimmerman, W.S. (1957). "The description of spatial-visualization abilities." *Educational and Psychological Measurement* 17, 185-199.
- Mislevy, R.J., Wingersky, M.S., Irvine, S.H., Dann, P.L. (1991). "Resolving mixtures of strategies in spatial visualization tasks." *British Journal of Mathematical and Statistical Psychology* 44, 265-288.
- Miyake, A., Friedman, N.P., Rettinger, D., Shah, P., and Hegarty, M. (2001). "How are visuospatial working memory, executive functioning, and spatial abilities related? A latent-variable analysis." *Journal of Experimental Psychology: General* 130(4), 621-640.
- Neerincx, M.A., Lindenberg, J., and Pemberton, S. (2001). "Support concepts for web navigation: A cognitive engineering approach." *WWW10*, 119-128.
- Norman, K.L. (1991). *The Psychology of Menu Selection: Designing Cognitive Control at the Human/Computer Interface*, Ablex Publishing.
- Norman, K.L., and Butler, S.A. (1989). "Search by uncertainty: Menu selection by target probability." University of Maryland Center for Automation Research Technical Report CAR-TR-432.
- Oulasvirta, A. (2004). "Task demands and memory in web interaction: a levels of processing approach." *Interacting with Computers* 16, 217-241.
- Pellegrino, J.W., Alderton, D.L., and Shute, V.J. (1984). "Understanding spatial ability." *Educational Psychologist* 19(4), 239-253.
- Simon, H.A. (1990). "Alternative visions of rationality." In P.K. Moser (Ed.) *Rationality in action: Contemporary approaches* (pp. 189-204). New York: Cambridge University Press (original work published 1983).
- Thurstone, L.L. (1934). "The vectors of mind." *Psychological Review* 41, 1-32.

- Thurstone, L.L. (1938). "The perceptual factor." *Psychometrika* 3(1), 1-17.
- Vicente, K.J., Hayes, B.C., & Williges, R.C. (1987). "Assaying and isolating individual differences in searching a hierarchical file system." *Human Factors* 29(3), 349-359.
- Weinreich, H., and Lamersdorf, W. (2000). "Concepts for improved visualization of Web link attributes." *Computer Networks* 33, 403-416.
- Zhang, H., and Salvendy, G. (2001). "The implications of visualization ability and structure preview design for web information search tasks." *International Journal of Human-Computer Interaction* 13(1), 75-95.