

The Contribution of Thumbnail Image, Mouse-over Text and Spatial Location Memory to Web Page Retrieval in 3D

Mary P. Czerwinski, Maarten van Dantzich, George Robertson & Hunter Hoffman

Microsoft Research, One Microsoft Way, Redmond, WA 98052, USA.

{marycz,ggr,maartenv,a-hunhof}@microsoft.com

Abstract: We present an empirical evaluation of the contribution of pictorial image and spatial location information on the retrieval of previously stored Web pages. Subjects were given 100 snapshots of Web pages that they stored in spatial locations on an inclined plane in a desktop 3D environment (Data Mountain). We had them return and try to retrieve their pages again, using a variety of retrieval cues. Even though users had not seen their Web page layout for several months, their retrieval times were not significantly slower. In addition, on half of the trials, stored pages were not presented as thumbnail images of the Web pages but as blank icons. Taking the pictorial thumbnail images away initially led to a significant drop in subjects' ability to find the pages, although within a short period of time subjects were able to find the pages equally fast without the thumbnail information. These results indicate that the use of 3D visualization techniques such as those described in this paper can lead to improved user memory for where favourite or frequently used information is stored in an electronic environment.

Keywords: 3D information visualization, information retrieval, thumbnail images, spatial location memory.

1 Introduction

In a recent paper we described a desktop 3D environment for document management (Robertson et al., 1998), and reported that users retrieved documents reliably faster in this environment than with a traditional Web browser's bookmarking mechanism. Since our novel environment (which we call the Data Mountain) had a number of new features that might benefit the user, this paper attempts to tease apart the relative contributions of these features. In particular, we wanted to determine whether retrieval time is helped primarily by users' spatial memory for the location of a document or by the ability to do a visual match using thumbnail miniatures. We also wanted to determine if spatial memory is retained after a long period of time.

1.1 Spatial Cognition

There has been much research investigating the role of spatial memory on our ability to navigate and retrieve information in virtual environments (Darken, 1993; Darken & Sibert, 1996; Jones, 1986; Ruddle et al., 1997; Waller et al., 1998). Darken & Sibert (1996) empirically validated several principles that designers can follow if they strive to design easily navigable

worlds that promote the acquisition of a mental map of the space. These include the use of directional landmarks, gridlines, paths, boundaries and maps to help the user navigate more efficiently to targets. In addition, they recommend combining sensory modalities, like adding 3D sound cues to visual worlds, to aid the navigator. Interestingly, Darken & Sibert (1996) observed that if the space is not divided using these simple, organizing principles, then users will impose their own, conceptual organization upon the space. While we will not explore navigation in head-mounted 3D in this paper, we have leveraged several of the design principles espoused by Darken & Sibert (1996). For instance, we have included left-right and front-back spatialized audio as indicative of where information is in the 3D space. We have also included the use of passive landmarks and allowed users the ability to build a personal organization of their information over time.

Research on navigation has also shown that subjects build up spatial representations of environments by interacting with those environments over time (Montello & Pick, 1993; Franklin et al., 1992; Ruddle et al., 1997). In the study presented in this paper, we brought back a group of subjects to re-

experience their spatial layout of Web pages that they themselves manually arranged in a 3D environment approximately 6 months earlier. These subjects had seen their layouts one additional time prior to this return visit. They returned 6 weeks after their first visit to retrieve their Web pages in a pilot study. There was no significant change in their speed at retrieving Web pages at that time, compared to the session in which the subjects created their layouts. In this study (their 3rd visit to our laboratory), an additional 4 months had gone by since the subjects had last seen their stored Web pages in the 3D environment. We were uncertain whether or not subjects would remember much about their layouts after such an extended period of absence. However, it was our hypothesis that the benefits of the 3D environment (3D perceptual cues and spatial arrangements) would enable subjects to either find their targets just as quickly as when they originally stored them, or at the very least quickly re-learn their layouts and find pages increasingly quickly over trials in the return session.

1.2 Textual Titles

In addition to spatial layouts in 3D, our subjects were provided with the title of a Web page whenever they moved their cursor over one. The work of Jones & Dumais (1986) has shown that the power of semantic labels as cues during retrieval is very strong, but document retrieval can benefit to some degree from the addition of spatial location knowledge. In their study, a condition that included both a semantic label and a specific spatial location for a document was superior to conditions that simply included the label only or the spatial location only. Therefore, we attached the text titles to the pages stored in our environment, and used 3D spatial locations to help users create an organization that made sense to them.

1.3 Clustering in Memory and Retrieval

According to the memory literature, people have a natural inclination to cluster items in memory when allowed to do so, and the amount of organization created (with or without experimenter-supplied category structure) strongly influences how easily the information is later remembered. Bousfield (1953) used a 60-item list comprised of 15 words from four distinct categories: animals, personal names, professions and vegetables. Although presented with these list items in a randomized order, subjects had a strong tendency to recall the items and write them down by category (e.g. owl, hawk, sparrow, carrots, cucumbers broccoli), evidence of what Tulving (1962) calls 'subjective organization' (as opposed to organization imposed by the experimenter or Web

browser). The Data Mountain facilitates a computer user's natural tendency to cluster information into meaningful, informal chunks (Miller, 1956) by making it easy for people to represent their Web pages in clusters on the mountain's surface. According to Mandler:

"memory and organization are not only correlated, but organization is a necessary condition for memory ... all organizations are mnemonic devices" (Mandler, 1967, p.328-9).

We hypothesized that the clustered Web page layouts built by users would provide a strong mnemonic aid that might hold over long periods of absence from the user interface.

Bower et al. (1969) showed how such organization can improve the storage of information into long term memory. Words were presented in a hierarchy that was organized in a meaningful way for one group, or randomly for a second group. Subjects were given four trials to learn all the words in the hierarchy. On the fourth trial, subjects who studied the meaningful hierarchy performed with 100% accuracy, while those in the random condition only reached 62% recall accuracy on the fourth trial.

The memory literature is informative about what sorts of Web page representations are likely to lend themselves to this powerful mnemonic of clustering and which ones will not. Gollin & Sharps (1988) gave participants items derived from four categories (animals, medical things, vehicles, and kitchen things) with 10 items per category. For one group of subjects, the stimuli were photographs of the items (e.g. a photo of a lion). Another group got printed names of the items presented verbally only (e.g. LION). The category superiority effect was only found for the VERBAL items. Thus, the text titles and spatial segregation of clusters we used may prove essential for exploiting the organizational structure of items in memory during retrieval in 3D environments.

On the other hand, retrieval cues must be easily associated with the information being sought in a personal information space. The greater the number of features (cues) provided to a user during search that might benefit activation in memory for where that information is stored, the higher the probability of recall (Underwood & Schulz, 1960). Classical mnemonic research has documented that mental cues in the form of visual images are an excellent way to enhance memory for items. As a result, classical mnemonic procedures emphasize using vivid visual images to enhance memorability of information

(Patten, 1990). To increase ease of retrieval, the Data Mountain includes the thumbnail images of Web pages, in addition to the spatial and textual cues mentioned above. Unfortunately, little is known about the contribution of pictorial images during the storage and retrieval of Web pages. For instance, do pictorial cues provide better cues than either semantic labels or knowledge of the target page's spatial location? How do the various cues combine during the retrieval process?

In this study, we used a desktop 3D application to help tease apart what it is that subjects actually do remember about the spatial location of Web documents, as well as the contributions of Web page titles and their pictorial images (thumbnails) on the retrieval of those pages. It was our hypothesis that the Web page title, its unique spatial location on a user's desktop, and its pictorial thumbnail image would combine to form a powerful retrieval cue that would aid users trying to return to a Web site after a long period of time has passed since their first visit.

2 The Desktop 3D Application

The desktop 3D environment used in this study is referred to as the Data Mountain, and has been described in detail elsewhere (Robertson et al., 1998), so only a short overview of the environment will be provided here. Currently the Data Mountain is being used as a 3D alternative to a Web browser's favourites or bookmark mechanism.

The Data Mountain uses a planar surface (a plane tilted at 65 degrees; see Figure 1), on which documents are dragged. A document being dragged remains visible so that the user is always aware of the surrounding pages. The user can place her Web pages (or documents) anywhere on the mountain. In practice, the user creates meaning by organizing the space in which she lays out these documents. In our study, there were many ways to organize the Web pages in the space. Each user was allowed to freely choose a method and alter it at any point throughout the study. It is our belief that the manual layout of Web pages in 3D space is a strong determinant of spatial location memory during storage.

The Data Mountain allows users to quickly see which Web pages have been stored in a particular cluster via inspection of each Web page's thumbnail, and its title in 'mouse-over' text. A user need simply rest the mouse over a Web page and the mouse-over text is popped up immediately. Note that mouse-over text has some similarity to hover tool-tips (where causing the mouse to hover over an object for some period of time brings up a tool-tip), except that in our

case the appearance delay is zero. In addition, a yellow halo surrounds the Web page and the title so that it is clear to which page the title text belongs.



Figure 1: Data mountain with 100 user-organized Web pages.

There are a number of 3D depth cues designed to facilitate spatial cognition. The most obvious are the perspective view, accompanying size differences, and occlusion, particularly when pages are being moved. Simple, circular landmarks on the surface of the Data Mountain also offer obvious cues, which may or may not be utilized during page placement as well as retrieval. Less obvious, but also quite important, are the shadows cast by the Web pages, and the spatialized audio cues that provide the user with right-left, front-back differential audio feedback as Web pages are moved in the space.

Previous research (Robertson et al., 1998) showed that the Data Mountain was a significant improvement in design over the standard, hierarchical tree mechanisms used to store Web pages today (e.g. Microsoft's Internet Explorer 4.0). However, in that study users stored and retrieved 100 Web pages in a same-day session. It was not clear whether the observed advantage for the Data Mountain environment would hold over a long period of absence. In addition, it was not clear from the earlier research which features were contributing to superior performance in the Data Mountain. To explore the contribution in performance from the thumbnail images, we did not show the thumbnails for half of the retrieval trials. Subjects were only shown a white placeholder outlined in black where each thumbnail would have occurred on these trials, as shown in Figure 2. If the thumbnail image was one of the stronger contributors to performance when subjects retrieved Web pages, we would hypothesize that performance would be especially bad on these trials after 6 months. However, if subjects could leverage

their spatial location memory for Web pages, perhaps performance would be no slower in this condition.




Figure 2: Same data mountain without thumbnails.

3 User Study

3.1 Subjects

Nine subjects, all of whom participated in a previous study investigating the Data Mountain, returned after approximately 6 months. All subjects had created a layout of 100 Web pages during a storage phase in the first session, and then retrieved each of those pages in the same session.

In order to qualify for participation in the initial study, all users had to successfully answer a series of screening questions pertaining to Web browser and Internet knowledge. Subject ages ranged from 18 through 50 years old, and all had normal or corrected-to-normal vision. There were 5 females in the return sample.

3.2 Stimuli and Equipment

The study was run on high-end Pentium machines (P6-300s), with 128MB of memory, and a 17-inch display. The machines had either an Intergraph Intense 3D Pro 1000 or 2200 graphics accelerator card and ran WindowsNT4. One hundred Web pages were used in this study; fifty pages were selected randomly from PC Magazine's list of top Web sites and fifty pages selected randomly from the Yahoo! database. Figure 1 shows an example of a subject's Data Mountain layout that was stored and reused for this study. Each subject was provided with the layout created in a session 6 months prior to the current session. For the trials in which the Web page thumbnails were turned off, a simple outline of a white placeholder image was viewable by the subject.

3.3 Procedure

Subjects were briefly welcomed back to the lab and informed that they would be seeing their old Data Mountains, and would be asked to find the same 100 Web pages they stored previously. Subjects were also told that, on half of the trials, the thumbnail images would be missing from their view of the Data Mountain. Once subjects indicated that they understood the instructions, the study was started. Subjects were allowed 5 minutes to review their Data Mountain before starting the retrieval trials.

During the test session, participants were shown one of four different retrieval cues and asked to find the related page. The four retrieval cueing conditions were: the title of the page, a one or two sentence summary of the page's content, a thumbnail image of the page, and all three cues simultaneously (called the 'All' cue). Participants saw 25 trials of each cueing condition, for a total of 100 retrievals. All pages presented for retrieval were seen in the first visit's storage phase and thus present in the environment. The Web pages to be stored and the subsequent retrieval cues were presented in a random order for each participant. Figure 1 shows an example of each of the four styles of retrieval cues. If a participant could not find the target page within two minutes, a 'time-out' was enacted and the participant was instructed to proceed to the next retrieval task. A Web page retrieval was defined as bringing a page forward to a close-up position by clicking on it. Users were not explicitly discouraged from producing incorrect retrievals.

This version of the study was run in blocks of 10 trials, with the display alternating between having a thumbnails 'on' (Figure 1) and a thumbnails 'off' (Figure 2) view of the subject's organization. The study always started with the thumbnails turned on for each subject, as a way of allowing them to familiarize themselves with their layouts over the first 10 trials.

Four main dependent variables were used in this study:

1. Web page retrieval time;
2. the number of incorrect pages selected prior to finding the correct page;
3. the number of trials for which the participant failed to retrieve the correct page within the two-minute deadline; and
4. the participants' subjective ratings of the software.

These dependent measures are assumed to be powerful indicators of subjects' ability to locate items

Title	Title, Summary & Thumbnail	Summary	Thumbnail
Bezerk — The Free On-line Entertainment Network	Bezerk — The Free On-line Entertainment Network. The new, FREE premier on-line entertainment network features YOU DON'T KNOW JACK the netshow, the on-line version of the irreverent quiz show party game CD-ROM 	The new, FREE premier on-line entertainment network features YOU DON'T KNOW JACK the netshow, the on-line version of the irreverent quiz show party game CD-ROM	

Table 1: Examples of the four cueing conditions used in the study.

in space. If subjects know where their Web pages are, even with the thumbnails turned off, retrieval performance should be efficient. These dependent measures will be compared to the same performance measures collected in the initial session, when subjects were first introduced to the Data Mountain.

4 Results

4.1 Reaction Times — Memory Retention after Absence




Figure 3: Average retrieval times to find 100 pages.

We analysed retrieval time performance comparing Web page retrieval time after a short break vs. after approximately 4 months delay (thumbnail-visible trials only). We found no reliable difference between performance immediately after storing the 100 Web pages compared to having a 4 month interval of not interacting with the layout on the Data Mountain. This is shown in Figure 3. A 2×4 (length of delay $\times$ retrieval cue condition)

ANOVA with repeated measures showed no significant main effect of delay on reaction time, $F(1,8) < 1$, NS. The main effect of retrieval cue condition was highly significant, $F(3,24) = 11.92$, $p < 0.001$, $MSE = 12.11$. No significant interaction between length of delay and retrieval cue type was found, $F(3,24) = 1.72$, $p > 0.10$, NS.

4.2 Number of Incorrect Retrievals




Figure 4: Average number of incorrect pages retrieved before finding the correct page for the first and second visit.

There was no reliable difference in the number of incorrect Web pages visited prior to finding the right page when comparing the two retrieval sessions (thumbnail-visible trials only). This is shown in Figure 4. A 2×4 (length of delay $\times$ Retrieval Cue Condition) ANOVA with repeated measures showed no reliable main effect of delay on number of incorrect pages retrieved, $F(1,8) < 1$, NS. A reliable main effect of retrieval cue condition was found, $F(3,24) = 4.27$, $p < 0.05$, $MSE = 0.04$, and no interaction

between length of delay and retrieval cue type was evident, $F(3,24) < 1$, NS.

4.3 Failed Trials




Figure 5: Average number of trials on which a Web page could not be found for first and second visit.

We observed no reliable difference between the number of trials in which the target Web page could not be found between the immediate and delayed retrieval sessions (thumbnail-visible trials only). This is shown in Figure 5. A 2×4 (length of delay between study and test $\times$ retrieval cue condition) ANOVA with repeated measures showed no significant main effect of delay on failed retrieval rates, $F(1,8) < 1$, NS. The main effect of retrieval cue type was highly significant, $F(3,24) = 10.34$, $p < 0.001$, $MSE = 0.11$. No reliable interaction between visit and retrieval cue type was found, $F(3,24) < 1$, NS.

4.4 Contribution of Thumbnails, Text Title and Spatial Location Memory




Figure 6: Average retrieval time by thumbnail condition and trial block (1 block = 10 trials).

In the next set of analyses, we examined the contribution of the thumbnail image to the speed and success of finding the target Web page. Only the 6 month delayed retrieval condition included trials when the thumbnails were turned off. We observed an initial significant slowdown in retrieval times when the thumbnail images were not available for inspection. However, this difference went away quickly after the first 2 blocks of the session (see Figure 6). Also, it is clear from looking at Figures 3–5 that the significant drop in performance was due to the thumbnail cueing condition only. Looking only at the results of the 6 month delayed retrieval phase, a 2×4 (thumbnail visible vs. thumbnail not visible $\times$ retrieval cue condition) ANOVA with repeated measures showed that removing the thumbnail images led to significantly longer reaction times, $F(1,8) = 10.97$, $p < 0.01$, $MSE = 42.69$. A significant effect of retrieval cue type was evident, $F(3,24) = 6.42$, $p < 0.002$, $MSE = 30.77$, and a significant interaction between visibility of the thumbnail and retrieval cue type was found, $F(3,24) = 6.03$, $p < 0.003$, $MSE = 33.68$. A separate repeated measures ANOVA examining block (5) $\times$ thumbnail condition (2) showed that there was a significant effect of block on retrieval times, $F(4,32) = 3.8$, $p = 0.013$, and no other effects were significant. In other words, thumbnail images were more important for speedy retrieval in the first few blocks, but over time subjects learned to rely on spatial location memory so effectively that there was no difference between seeing the thumbnail or not.

In addition, looking only at the results of the second visit, a 2×4 (thumbnail visible vs. thumbnail not visible $\times$ retrieval cue condition) ANOVA with repeated measures showed that removing thumbnail images led to significantly more failed retrievals, $F(1,8) = 12.42$, $p = 0.008$, $MSE = 0.02$. A reliable effect of retrieval cue type was found, $F(3,24) = 15.40$, $p < 0.001$, $MSE = 0.02$. And there was a reliable interaction between the visibility of the thumbnail and the retrieval cue type, $F(3,24) = 30.36$, $p < 0.001$, $MSE = 0.33$.

Finally, looking only at the results of the second visit, a 2×4 (thumbnail-visible vs. thumbnail-not-visible $\times$ retrieval cue condition) ANOVA with repeated measures showed that removing thumbnail images led to a marginally significant increase in number of incorrect pages retrieved, $F(1,8) = 4.35$, $p = 0.07$, $MSE = 0.32$. Retrieval cue type had no reliable effect on incorrect retrieval rates, $F(3,24) = 2.04$, $p = 0.14$, $MSE = 0.75$. A marginally significant interaction between visibility of the thumbnail and

retrieval cue type was found, $F(3, 24) = 2.50, p = 0.08, \text{MSE} = 0.43$.

4.5 Subjective Ratings

After retrieving all 100 Web pages, subjects filled out a subjective questionnaire in each session. Figure 2 shows subjects' subjective ratings for their first and second retrieval sessions in this study. In general, the ratings were very favourable of the 3D environment. In two cases, subjects were reliably more positive after the 6 month delayed session. Specifically, subjects rated the software as feeling more familiar in the second session, and they also felt that the spatial layout of their Web pages aided their performance 6 months later more than in the initial session. All significant differences were determined by a paired t-test, one-tailed, $\alpha = 0.05$. (We used a one-tailed test because we were looking for improvements of subjects' ratings over time.)

In addition, we had subjects rank the relative effectiveness of four of the features in the Data Mountain that we hypothesized might benefit retrieval of the Web pages. A rank of 1 meant that a feature was considered the best for retrieval, 4 was the worst ranking. On average, subjects ranked the thumbnail images as the most helpful (avg. = 1.8), followed closely by the mouse-over text (avg. = 2.0) and the spatial location of the Web page on the Data Mountain (avg.rank = 2.2). The spatialized audio feedback was consistently ranked the least helpful cue (avg.rank = 4.0). These differences were significant as determined by a Kruskal-Wallis non-parametric test, $\chi^2(3) = 24.02, p < 0.001$.

Questionnaire Item (Scale: 1 = Disagree, 5 = Agree)	1st Session	2nd Session
I like the software I used today.	3.87	3.8
The software was easy to use.	4.1	4.2
The software feels familiar.	3.4	**4.1
It was easy to find the page I was looking for with the software.	3.67	3.56
Laying items out spatially helped me find them later.	3.56	**4.2

Table 2: Average satisfaction ratings for 2 retrieval sessions, on scale 1 (worst) to 5 (best). (** indicate significant differences between the 2 sessions).

We asked subjects to tell us how they found Web pages when the thumbnails were turned off in a multiple-choice format (guessed the location, knew the relative location or knew the exact location). Of the nine subjects, only 2 said that they simply guessed where the Web page was. The remaining 7 subjects said that they remembered the general vicinity of the

Web page. We often observed subjects honing in on the cluster of Web pages they knew to contain the target page, after which they would use the mouse-over text to find the specific target page.

5 Conclusion

We developed a 3D environment for storing favourite Web pages, called the Data Mountain. Previous studies had shown the benefits of using the Data Mountain for retrieving stored Web pages over traditional, hierarchical tree user interfaces. However, it was unclear from those studies what particular features of the 3D environment were benefiting performance. This study explored the contribution of the thumbnail images, spatial location, and mouse-over text on the Data Mountain. With the thumbnail images viewable, subjects were no slower after an absence of approximately 4 months in retrieving Web pages previously stored. When the thumbnail images were turned off, performance was disrupted, but only initially and only for the thumbnail retrieval cueing condition. After 2 blocks of trials, subjects were just as fast at retrieving their Web pages with absolutely no pictorial cues. This is quite interesting, given that subjects often chose the thumbnail image as their preferred cue. Of course having the thumbnail images turned off when the thumbnail cue was presented was an extremely difficult trial type, and accounted for the majority of the disruption in subjects' retrieval performance in the no thumbnail condition. Although subjects preferred to use the thumbnail image during retrieval, they could perform just as well using the mousing-over text title and the spatial location of the page. We think that this is an important finding for the design of 3D environments, especially for instances in which no visual features are available to distinguish one item from another. For example, if a subject is looking through a variety of electronic documents that are primarily text, there will be little distinguishing visual detail in a thumbnail or iconic representation. Our findings suggest that, by allowing the subject to manually lay out their data in 3D space, and by providing a title on mouse-over, subjects will be able to remember where they have stored their non-pictorial information quite well. Especially if one considers the fact that our subjects had not seen their 3D layouts for a 4-month period, the present results indicate that this combination of pictorial image, spatial location and mouse-over text title can be quite effective. We encourage designers to consider these issues when designing information spaces.

Future work will explore these design issues in larger, 3D information visualizations. We were

surprised to see how robust subjects' memory for their layouts was, and we intend to continue to examine how this spatial memory holds in new designs where the information spaces scale up to 100s and 1000s of Web pages or documents. In addition, we would like to explore subjects' spatial memories with and without the benefit of mousing-over text titles, in order to better isolate just how much each of those particular cues are contributing to retrieval performance in new designs. Finally, the thumbnails used in this study were snapshots of Web pages. Future research will explore whether it is possible to design more effective 'smart icons' as one means of improving ease of memory retrieval in the Data Mountain.

Acknowledgements

We would like to thank Sue Dumais for her feedback on the study, Scott Tiernan for help with the statistical analysis of this study, and the user interface research group at Microsoft for their work on the Data Mountain visualization.

References

- Bousfield, W. A. (1953), "The Occurrence of Clustering in the Recall of Randomly Arranged Associates", *Journal of General Psychology* **49**, 229-40.
- Bower, G. H., Clark, M. C., Lesgold, A. M. & Winzenz, D. (1969), "Hierarchical Retrieval Schemes in Recall of Categorical Word lists", *Journal of Verbal Learning and Verbal Behaviour* **8**, 323-43.
- Darken, R. P. & Sibert, J. L. (1996), Wayfinding Strategies and Behaviours in Large Virtual Worlds, in G. van der Veer & B. Nardi (eds.), *Proceedings of CHI'96: Human Factors in Computing Systems*, ACM Press, pp.142-50.
- Darken, R. P. and. Sibert, J. L. (1993), A Toolset for Navigation in Virtual Environments, in S. Ashlund, K. Mullet, A. Henderson, E. Hollnagel & T. White (eds.), *Proceedings of INTERCHI'93*, ACM Press/IOS Press, pp.157-65.
- Franklin, N., Tversky, B. & Coon, V. (1992), "Switching Points of View in Spatial Mental Models", *Memory and Cognition* **20**, 507-18.
- Gollin, E. S. & Sharps, M. J. (1988), "Facilitation of Free Recall by Categorical Blocking Depends Stimulus Type", *Memory and Cognition* **16**, 539-44.
- Jones, C. B. (1986), *Systematic Software Development Using VDM*, Prentice-Hall International.
- Jones, W. & Dumais, S. (1986), "The Spatial Metaphor for User Interfaces: Experimental Tests of Reference by Location versus Name", *ACM Transactions on Office Information Systems* **4**(1), 42-63.
- Mandler, G. (1967), Organization and Memory, in K. W. Spence & J. T. Spence (eds.), *The Psychology of Learning and Motivation*, Vol. 1, Academic Press, pp.327-72.
- Miller, G. A. (1956), "The Magical Number Seven, Plus or Minus Two: Some Limits on our Capacity for Processing Information", *Psychological Review* **63**, 81-97.
- Montello, D. R. & Pick, H. L. (1993), "Integrating Knowledge of Vertically-Alligned Large-Scale Spaces", *Environment and Behavior* **25**, 457-84.
- Patten, B. M. (1990), "The History of Memory Arts", *Neurology* **40**, 346-52.
- Robertson, G., Czerwinski, M., Larson, K., Robbins, D., Thiel, D. & van Dantzich, M. (1998), Data Mountain: Using Spatial Memory for Document Management., in *Proceedings of the ACM Symposium on User Interface Software and Technology, UIST'98*, ACM Press.
- Ruddle, R. A., Payne, S. J. & Jones, D. M. (1997), "Navigating Buildings in Desk-top Virtual Environments: Experimental Investigations using Extended Navigational Experience", *Experimental Psychology* **3**, 143-157.
- Tulving, E. (1962), "Subjective Organization in Free Recall of Unrelated Words", *Psychological Review* **69**, 344-54.
- Underwood, B. J. & Schulz, R. W. (1960), *Meaningfulness and Verbal Learning*, Lippincott.
- Waller, D., Hunt, E. & Knapp, D. (1998), "The Transfer of Spatial Knowledge in Virtual Environment Training", *Presence* **7**, 129-44.