Revisiting 2D vs 3D Implications on Spatial Memory

Andy Cockburn

Department of Computer Science
University of Canterbury, Christchurch
New Zealand

andy@cosc.canterbury.ac.nz

Abstract

Prior research has shown that the efficient use of graphical user interfaces strongly depends on human capabilities for spatial cognition. Although it is tempting to believe that moving from two- to three-dimensional user interfaces will enhance user performance through natural support for spatial memory, it remains unclear whether 3D displays provide these benefits. An experiment by Tavanti and Lind, reported at InfoVis 2001, provides the most compelling result in favour of 3D—their participants recalled the location of letters of the alphabet more effectively when using a 3D interface than when using a 2D one. The experiment reported in this paper is based on Tavanti and Lind's, but it controls some previously uncontrolled factors. The results strongly suggest that the effectiveness of spatial memory is unaffected by the presence or absence of three-dimensional perspective effects in monocular static displays.

Keywords: 3D user interfaces, spatial memory, location learning, evaluation.

1 Introduction

The efficient use of graphical user interfaces relies heavily on human capabilities for spatial cognition. Several research projects, summarised in the following section, have shown that measures of spatial cognition are strongly correlated with performance in a variety of user interface tasks. This correlation raises the question "what can be done to better exploit human spatial capabilities in user interfaces?"

The powerful 3D graphics hardware available in desktop computers provides an attractive opportunity for enhancing interaction. It may be possible to leverage human spatial capabilities by providing computer generated 3D scenes that better reflect the way we perceive our natural environment. Systems such the 'Data Mountain' (Robertson, Czerwinski, Larson, Robbins, Thiel and vanDantzich 1998), the 'Task Gallary' (Robertson, vanDantzich, Czerwinski, Hinckley, Thiel, Robbins, Risden and Gorokhovsky 2000), and Win3D (www.clockwise3d.com) all work towards this goal by providing 3D alternatives to the 'flat' desktop metaphor.

There is some evidence supporting improved spatial memory in 3D. Robertson et al (1998) showed that task times and error rates were lower when retrieving web pages using their 3D Data Mountain than when using the standard 2D 'Favorites' mechanism of Internet Explorer. In our prior work, however, we compared two versions of the Data Mountain, one with and one without perspective effects, and showed no reliable difference between 2D and 3D (Cockburn and McKenzie 2001). We assume that the original Data Mountain outperformed the Favorites mechanism because it provided a spatial layout allowing immediate access to every item, while the Favorites interface did not (the users had to scroll to most items). Robertson et al's study, and their subsequent studies (Section 2), provides many important observations on the power of spatial cognition, but they say little about the role of 3D in its effectiveness.

Tavanti and Lind (2001) described an experiment comparing the effectiveness of spatial memory in computer generated 2D and 3D displays. Their tasks involved recalling the location of letters of the alphabet hidden behind 'cards' depicted in hierarchical 2D and 3D displays, as shown in Figure 1. The participants' spatial memory was much better in the 3D condition. They concluded, "a realistic 3D display better supports a specific spatial memory task, namely learning the place of an object". There were, however, several potential confounding factors in their experiment, some of which they acknowledged. These include the vertical versus horizontal orientation of the windows, the use of letters (which are normally arranged horizontally in written language), and the sizes and separation of the individual icons. Another important difference between the two interfaces is that the 2D version (Figure 1a) effectively provides a linear arrangement, with no two icons overlapping on the y-axis. This constraint is not a requirement of a hierarchical 2D display, as Figure 1c shows.

The experiment reported in this paper uses a method similar to that of Tavanti and Lind to compare the effectiveness of spatial memory in 2D and 3D monocular static computer displays. The implications of the research are clear: if the 3D interface allows improved spatial memory, then there is strong reason for suspecting that user interface performance can be improved by incorporating perspective effects.

Copyright © 2004, Australian Computer Society, Inc. This paper appeared at the 5th Australasian User Interface Conference (AUIC2004), Dunedin. Conferences in Research and Practice in Information Technology, Vol. 28. A. Cockburn, Ed. Reproduction for academic, not-for profit purposes permitted provided this text is included.

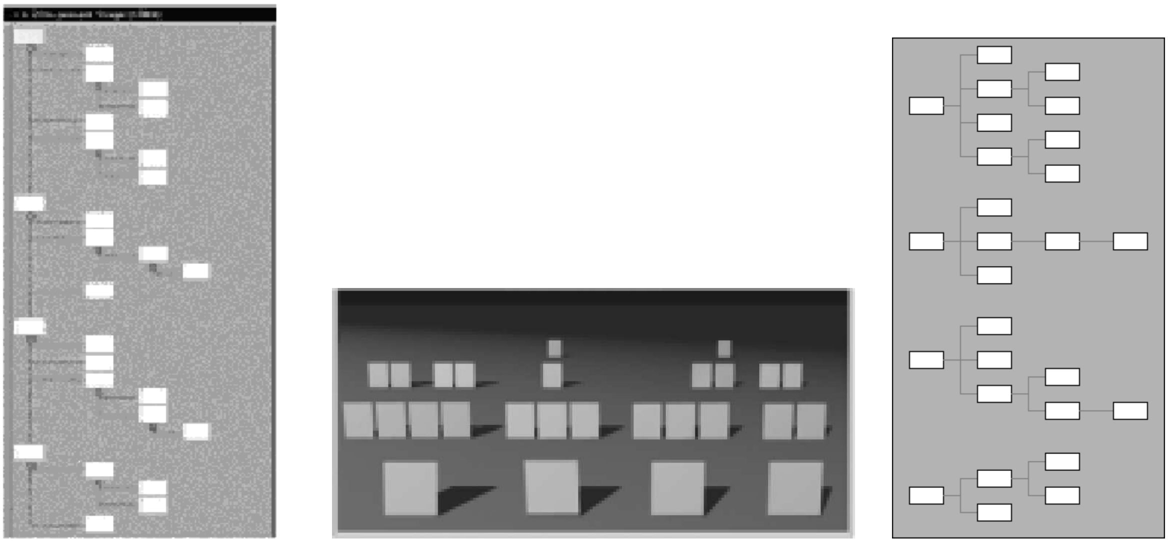

(a) Tavanti and Lind's 2D interface. **(b) Tavanti & Lind's 3D interface.** **(c) A 2D structure, equivalent to (a), but non-linear on the y-axis.**

Figure 1: The interfaces used by Tavanti and Lind (2001), and.a 2D-structure equivalent to (a).

2 Related Work

Two areas of related work are particularly relevant to this investigation. First, several researchers have shown a relationship between users' spatial capabilities and their performance with user interfaces. Second, there has been extensive prior research comparing the effectiveness of 2D and 3D user interfaces, but relatively little investigating whether spatial memory differs in 2D and 3D displays.

2.1 Spatial Memory and User Interface Performance

Performance with user interfaces is strongly predicted by spatial aptitude. This result has been confirmed in many separate experiments and with varied interface types. Egan and Gomez (1985) showed that measures of spatial memory and age provided the best predictors of how well participants learned to use a text editor. Gagnon (1985) reported the surprising result that computer game scores were not correlated with measures of hand-eye coordination, but were correlated with scores on a spatial memory test. Vicente, Hayes and Williges (1987) and Leitheiser and Munro (1995) also concur that measures of spatial ability predict performance in hierarchical file browsing tasks and in a variety of file management tasks.

As mentioned earlier, the Data Mountain's spatial arrangement of webpage thumbnail images allowed more rapid and accurate page retrieval than the 'Favorites' mechanism in Microsoft Internet Explorer (Robertson et al. 1998). More remarkably, a follow-up evaluation showed that participants were able to rapidly retrieve pages from their spatial arrangements four months after creating them (Czerwinski, vanDantzich, Robertson and Hoffman 1999). The strength of the spatial cue is dramatically demonstrated by the fact that retrieval times did not significantly worsen when the thumbnail images were replaced with blank outlines.

Ehret (2002) provides interesting insights into how users learn the location of items in a user interface. His experiment shows that users learn locations more effectively when targets poorly represent their function. In other words, the higher the 'evaluation cost' (the degree of effort the user must put into finding the function of an item), the better the location is learned. Ehret presents a theory that predicts how well users learn the location of interface items as evaluation cost varies.

Jones and Dumais (1986) provide some cautions on over-reliance on spatial organization. Their evaluation indicates that semantic labels provide stronger retrieval cues than spatial organisation alone, but indicate that combinations of semantic and spatial organization enhance performance.

2.2 2D versus 3D Spatial Memory

There has been a great deal of prior work comparing the general effectiveness of 2D and 3D user interfaces, particularly in the military and aviation domains. Many of the findings are dependent on the precise tasks under analysis. Wickens et al (1997) provide a fitting summary for prior work on 2D versus 3D evaluations:

> "whether the benefits of 3D displays outweigh their costs turns out to be a complex issue, depending upon the particular 3D rendering chosen, the nature of the task, and the structure of the information to be displayed."

The specific question addressed in this paper is "do 3D interfaces result in better spatial memory than 2D ones"? Prior work disagrees on this point.

In describing their follow-up Data Mountain evaluation, Czerwinski et al (1999) stated "3D visualization techniques such as those described in this paper can lead to improved user memory…" yet their evaluation did not isolate dimensionality as a factor. Our own evaluation, which compared a 'flat' (no perspective effects) version of the Data Mountain with one that did support perspective showed no significant difference in performance, although the participants strongly preferred the ‘cool’ 3D version (Cockburn and McKenzie 2001). In a follow up study that aimed to control effects due to the use of computer displays (Cockburn and McKenzie 2002), we compared similar tasks using 2D and 3D physical models as well as equivalent computer-supported systems. The 2D physical model was reliably faster than the 3D one, though again there were no significant differences between the three computer-supported displays. Several of the participants using the 3D physical model commented that it was hard to remember where they had placed their ‘clusters’ of related pages, yet similar comments were seldom made about the 2D model.

Finally, as described in the introduction, Tavanti and Lind (2001) produced results showing a strong positive effect for memory in 3D over 2D. The aim of the experiment described in this paper is to control the uncontrolled factors occurring in their experiment, and to determine whether the positive effects of 3D still occur.

3 Evaluation

Before describing the evaluation method, we must stress the scope of the evaluation. In calling our interface ‘3D’, we are consistent with many similar evaluations, as reported in the previous section. There are, however, many other interpretations of the term ‘3D’ in interactive systems: these include the use of immersive technologies such as virtual reality, the use of advanced display technology such as stereoscopic views, and the use of navigable 3D worlds such as game environments. Our evaluation only addresses the presence versus absence of perspective effects in displays that provide a monocular static viewpoint: the user cannot ‘move’ within the 3D environment on any axis. Within this limited domain, prior research has produced divergent results, leading to confusion over the efficacy of 2D versus 3D spatial memory. A negative result for 3D would mean that next generation interfaces using 3D should have some motivation other than enhancing spatial memory as their basis—there are many important benefits of 3D beyond spatial memory, but these are not addressed in this evaluation. A positive result for 3D would mean that next generation interfaces using 3D can expect efficiency enhancements simply from their use of perspective effects because these effects enhance spatial memory, which strongly predicts the efficiency of use.

The evaluation method is heavily based on that of Tavanti and Lind (2001). The 3D interface used in our experiment (Figure 2b) is a faithful replication of that used by Tavanti and Lind (Figure 1b). Our 2D interface, shown in Figure 2a, diverges from that used by Tavanti and Lind (Figure 1a). In comparing 2D and 3D interfaces of this type, we saw little justification for constraining the 2D interface to a linear arrangement on one axis—no two icons in Tavanti and Lind’s interface overlap on the y-axis, demanding that icons be small, even when closely placed on the y-axis.

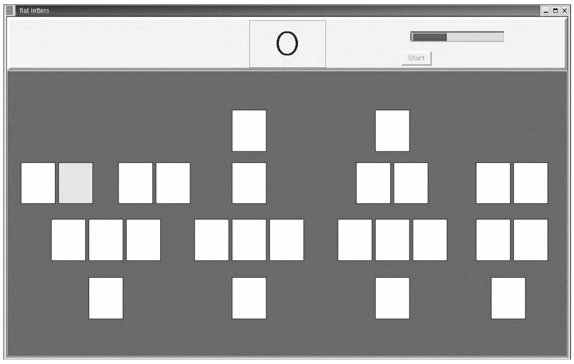

(a) 2D: revealing a letter.

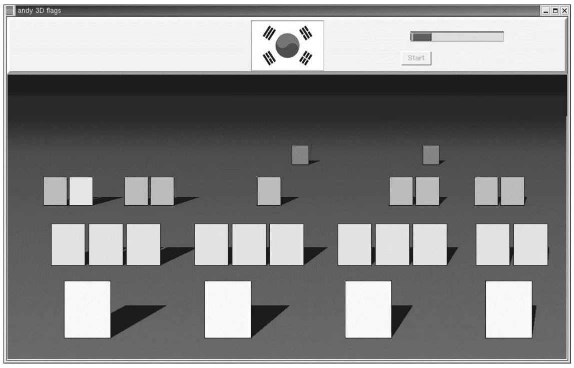

(b) 3D: revealing a flag.

Figure 2: The 2D and 3D interfaces used in the experiment.

Both the 2D and 3D display areas were identically sized at 1000×510 pixels. They differed in that the 3D version provided the following perspective effects: size gradients (‘further’ icons are smaller than ‘near’ ones), shadowing, and proximity luminance covariance, meaning that the icons fade towards the background colour with distance (Dosher, Sperling and Wurst 1986). Other than perspective effects, both interfaces behaved identically.

The experimental tasks involved first memorising and then recalling the location of letters of the alphabet and national flags. Although Tavanti and Lind only evaluated recall of the twenty-seven letters of the Swedish alphabet we also evaluated recall of flags because preliminary trials revealed that participants used mnemonic aids to construct words, sounds, or word sequences from letters. For example, if the bottom row of letters revealed ‘K’, ‘D’, ‘O’, ‘Q’, the participant might form the mnemonic ‘Klingons Don’t Order Quietly’. Mnemonics such as these confound the intended measurement of spatial capabilities. The results support our conjecture that flags are less readily aided by mnemonics than letters.

During the memorisation stage, pressing the mouse over one of the ‘cards’ (blank icons) in the display would highlight the card and reveal the letter or flag ‘hidden behind’ it. When the mouse button was released, the letter/flag and highlighting would disappear. The letter/flag associated with each card was shown in a separate display area at the top of the window (see Figure 2a and 2b). Tavanti and Lind’s systems behaved identically.

The separation between the cards and the display raises important validity concerns. For example, if the separation were removed, then the letters and flags would diminish on the more 'distant' cards in the 3D display. They would also become less legible due to the depth fading. It might be that the sizing variation would reinforce the spatial location in 3D, but alternatively it is possible that the reduced legibility would detrimentally affect it. Despite these concerns, we maintained the separation between cards and display in order to maintain maximum consistency with Tavanti and Lind's original experiment. They found a positive affect for 3D under similar conditions, and we wished to test as similar an interface as possibile, removing the confounding factors identified above.

After the memorisation period, the interfaces prompted the user to find all of the letters/flags, one at a time. A randomly selected letter/flag would be shown in the display area, and the user would have to press the card associated with it. Pressing the mouse button on a card highlighted it, but no other feedback was provided, hence the participants did not know whether their selection was correct or not. A time-bar (top-right of Figures 2a and 2b) showed the remaining time in both the memorisation and recall activities.

3.1 Experimental Method

Participants were randomly assigned to either the 2D or 3D condition for the experiment. They carried out three separate memorise/recall tasks.

The first memorise/recall task was used to familiarise the participants with the experimental procedure, and involved learning the location of four punctuation characters ('!', '@', '#', and '$') hidden behind the four cards in the display. They were allowed 20 seconds to memorise the location of the four characters, and 30 seconds to find them. Data from this training exercise was discarded.

The second and third tasks both involved memorising 26 items. The cards always appeared in the locations shown in Figure 2 (a similar layout to that used by Tavanti and Lind, shown in Figure 1). Participants had three minutes for memorisation (the complete time period was always used), and a maximum of five minutes for recollection (usually only one minute was required). The items were the twenty-six characters of the English alphabet, or the national flags of the 26 most heavily populated countries in the world. The order in which the participants used letters/flags was randomly assigned, as was the relationship between individual letters/flags and the cards that hid them.

During memorisation, software automatically logged the time, location, and item displayed for all mouse-button presses. During recall, software logged the time and location of all mouse-button presses, whether the card was the correct one, and if not, the distance between the correct one and the one selected (on both the x and y-axes). Software administered questionnaires after each memorisation stage, and after the entire experiment, recorded a variety of subjective measures and comments.

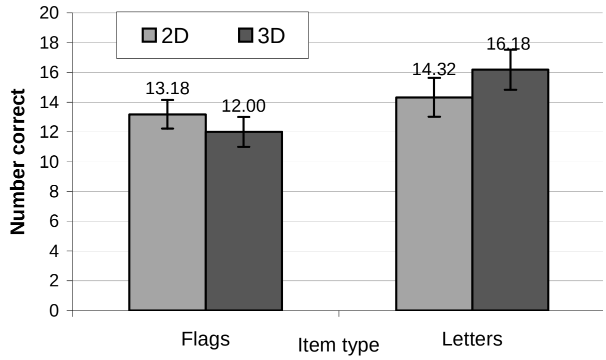

Figure 3: Mean number of correctly recalled items per condition. Error bars ±1 standard error.

3.2 Participant and Equipment Details

Forty-four (forty male, four female) Computer Science undergraduate students participated in the experiment. They were randomly assigned (gender-balanced 20-2) to either the 2D or 3D condition. Participation in the experiment lasted approximately thirty-five minutes and was rewarded with a $5 shopping voucher. Forty percent of the 2D participants and 36% of the 3D participants stated that they regularly played 3D computer games. Only one of the 2D and two of the 3D participants stated that they had never played 3D computer games.

The experiment was run on Pentium III computers with 17inch displays running at 1600×1200 resolution.

3.3 Data Analysis

The dependent measures in the experiment were the number of items correctly recalled and the mean miss distance (distance between the correct item and the one selected). It is unsurprising that these two dependent measures are strongly correlated, so the results focus on the number of correctly recalled items. Tavanti and Lind's experimental results also focus on the number of correctly recalled items.

The dependent measure is analysed in a 2×2 mixed-factors analysis of variance (ANOVA) for factors interface-type (between-subjects, 2D and 3D) and item-type (within-subjects, Letters and Flags).

4 Results

The participants generally enjoyed the experiment. The memorisation stage was clearly demanding, with all using the full three minutes allowed. One participant's comment that "as the counter approaches the three minute mark, it feels like last minute cramming for exams" indicates the pressure that most participants felt. The recall stage generally required only one or two minutes, with no one needing more than the five minutes allocated.

The mean number of items correctly recalled across both interface types and both item types was 13.9 (standard deviation 5.6). The range of correctly recalled items extends from a minimum of three by a participant in the 3D Flags condition to a maximum of 26 by three

Table 1: Percentage of correct recalls for each letter and flag. Note the wide distribution for flags compared to letters.

Percent	Flags	Letters
11-20%	Egypt (19) 	
21-30%	Ethiopia (27)  Indonesia(30) 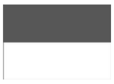 Iran (30)  Mexico (30)  Philipines(30) 	
31-40%	Bolivia (32)  Russia (32)  Turkey (35)  India (38) 	
41-50%	France (41)  Italy (41)  Pakistan(41) 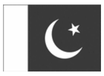 China (46) 	**M** **N** **T** 49 49 49
51-60%	Congo (51) 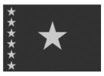 Thailand(51)  B'ladesh (54)  Ukraine (54)  Vietnam (54)  Nigeria (57) 	**H** **R** **V** **O** **F** **J** 51 51 51 54 57 57 **S** **D** **K** **Y** **Z** 57 59 59 59 59
61-70%	Germany(68)  Brazil (70) 	**C** **E** **G** **P** **Q** **U** 62 62 62 62 62 62 **W** B **I** **L** **X** 62 65 65 65 68
71-80%	Japan (73) 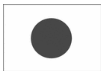 USA (76) 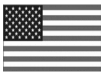	**A** 78
81-90%	UK (81) 	
91-100%	S.Korea (92) 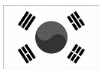	

participants (two 3D, one 2D) all using Letters. The mean miss distance across all conditions was 128 pixels (s.d. 71).

The means for the 2D and 3D conditions were very similar at 13.8 (s.d. 5.3) and 14.1 (s.d. 5.9), yielding no significant difference: $F_{1,42}$=.06, p=0.8. This result disagrees with Tavanti and Lind. The only data value generated under similar conditions in our experiment and theirs is the 3D Letters condition, which yielded a mean of 16.2 (s.d. 6.4) in our experiment, and 19.3 in theirs (when scaled from 27 to 26 items). Figure 1 shows the mean recall counts in our experiment for each of the four conditions. The mean miss distance also showed no significant differences between dimensions with 2D and 3D means of 128 (s.d. 63) and 128 (s.d. 78) pixels: $F_{1,42}$=0.001, p=0.9.

There was a significant difference between the number of items correctly recalled when using Letters and Flags. The mean number of items recalled with Letters and Flags were 15.3 (s.d. 6.2) and 12.6 (s.d. 4.5): $F_{1,42}$=8.2, p<0.01. Comments from the participants supported our conjecture that the recall of letters is readily aided by mnemonics. Several participants made comments in agreement with the following statement:

> "Letters were easy. For example, the second row contained NUL HBI TOY PM. They're all words, more or less, except for the HBI bit. Flags were a nightmare. I had to remember where each flag was. It was easier when I recognised the flag, like the UK was right there (points) and America there, but I didn't know most of the flags."

Although the numbers correctly identified letters and flags differed significantly, the mean miss distances showed only a marginal statistical difference, with mean distances of 115 (s.d. 84) and 141 (s.d. 50) pixels for Letters and Flags: $F_{1,42}$=3.9, p=0.06. The difference between the significant result for counts, but marginal result for distances, is probably best explained by participants guessing the location of items when uncertain about their location.

There was no significant interaction between factors interface-type and item-type: $F_{1,42}$=2.7, p=0.11, indicating no reliable differences in the way that letters and flags were memorised using the 2D and 3D interfaces.

Differences between the recall of letters and flags are further illustrated by analysing how often each particular letter or flag was correctly recalled. Table 1 shows that there is a narrow 29% range between the least and most frequently recalled letter: from letters M, N, and T which were correctly found in 49% of tasks to letter A which was correctly found in 78% of tasks. All other letters were in a narrow range between 51% and 68% of tasks. Successful recall of flag locations was much more varied, with a 73% range from the Egyptian flag (recalled in only 19% of tasks) to the South Korean flag (recalled in 92% of tasks). Twenty two percent of the participants were of Asian descent.

The participants mentioned two factors that affected whether they recalled flags: first, whether they knew the country associated with the flag; and second how visually distinct it was from the others in the set. These observations are consistent with extensive psychology research showing that recall improves when items are meaningful, concrete, and form an image in the mind (Paivio, Rogers and Smythe 1968; Paivio, Yuille and Madigan 1968), and when images are simple rather than complex (Attneave 1955).

Finally, the participants' responses to five-point Likert-scale questions, ranked from one (disagree) to five (agree), were consistent with the results reported above. Responses to the question "The display provides a sense of depth (some icons seem further away than others)" showed a strongly significant difference between 2D (mean 1.3, s.d. 0.7) and 3D (mean 4.2, s.d. 1.1): Mann-Whitney U=19.5, $p<.001$. Responses to questions regarding how well the participants expected to recall pages after memorisation ("I will accurately recall the location of pages") and how well they thought they performed after recall ("I did accurately recall the location of pages") were similar across 2D and 3D, with no significant differences.

5 Discussion

To summarise the results, the presence of perspective effects made no difference to how well participants recalled the location of letters or flags. Letters were recalled better than flags—an effect probably best explained by the easy use of mnemonics to enhance spatial memory with letters. Finally, the percentage of successful recalls of specific flags varied widely, with the participants' knowledge of the flag's country and its visual distinctiveness providing the most likely explanation.

Why, then, do our results differ from those of the experiment by Tavanti and Lind? We suspect that the vertical orientation of Tavanti and Lind's 2D display (Figure 1a) made the formation of effective letter mnemonics more difficult than the horizontal 3D layout—words and word combinations normally run horizontally left to right (in English and Swedish).

The results support our prior work showing that perspective effects in monocular static computer displays do not significantly influence the effectiveness of spatial memory (Cockburn and McKenzie 2001; Cockburn and McKenzie 2002). This does not imply, however, that spatial memory is the same in 'real' 2D and 3D. There has been little work comparing 2D and 3D spatial memory in the real world, so it remains unclear whether a 'perfect' computer-based implementation of 3D would produce spatial memory advantages or disadvantages for 3D. We intend to continue addressing this question in our further work.

The observation that flags were best recalled when personally meaningful and when visually distinctive supports prior memory research on human memory (Attneave 1955; Paivio et al. 1968; Paivio et al. 1968). Recent results show that users learn locations better when the cost of determining the association between items and their meaning is high (2002). Our result adds to this by confirming that when the costs are equal (all our icons were blank, as shown in Figure 2), locations are better learned when the underlying information is meaningful to the user and when it is visually distinctive.

6 Conclusions

The graphics hardware available in standard desktop computers makes a new range of rapidly interactive three-dimensional interfaces for office work technically feasible. Several research and commercial systems are already demonstrating 3D 'office' environments that might replace the current desktop metaphor.

Related work has claimed that 3D interfaces improve users' spatial memory for the location of objects in the interface. If correct, it is likely that 3D interfaces would improve user performance due to a well-established correlation between spatial memory and efficient use of graphical user interfaces.

The experiment described in this paper was based on an experiment providing the main evidence for improved spatial memory in monocular static 3D displays. The experiment attempted to constrain some previously uncontrolled factors.

The results disagree with the prior work, and strongly suggest that these 3D effects make no difference to the effectiveness of spatial memory in monocular static displays.

References

Attneave, F. (1955). "Symmetry, Information, and Memory for Patterns." American Journal of Psychology **68**: 209-222.

Cockburn, A. and B. McKenzie (2001). 3D or Not 3D? Evaluating the Effect of the Third Dimension in a Document Management System. Proceedings of CHI'2001 Conference on Human Factors in Computing Systems, Seattle, Washington, March 31-April 6.

Cockburn, A. and B. McKenzie (2002). Evaluating the Effectiveness of Spatial Memory in 2D and 3D Physical and Virtual Environments. Proceedings of CHI'2002 Conference on Human Factors in Computing Systems, Minneapolis, Minnesota, 20--25 April.

Czerwinski, M., M. vanDantzich, G. Robertson and H. Hoffman (1999). The contribution of thumbnail image, mouse-over text and spatial location memory to web page retrieval in 3D. Proceedings of Interact '99, Edinburgh, Scotland, IOS press.

Dosher, B., G. Sperling and S. Wurst (1986). "Trade-offs Between Stereopsis and Proximity Luminance Covariance as Determinants of Perceived 3D Structure." Vision Research **26**(6): 973--990.

Egan, D. and M. Gomez (1985). Assaying, Isolating, and Accommodating Individual Differences in Learning a Complex Skill. Individual Differences in Cognition. R. Dillon**:** 173--217.

Ehret, B. (2002). Learning Where to Look: Location Learning in Graphical User Interfaces. Proceedings of CHI'2002 Conference on Human Factors in Computing Systems. Minneapolis, Minnesota, 20--25 April**:** 211--218.

Gagnon, D. (1985). "Videogames and Spatial Skills: An Exploratory Study." Educational Communication and Technology **33**(4): 263--275.

Jones, W. and S. Dumais (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.

Leitheiser, B. and D. Munro (1995). An Experimental Study of the Relationship Between Spatial Ability and the Learning of a Graphical User Interface. Proceedings of The Inaugural Americas Conference on Information Systems, Pittsburgh, PA, August 25--27.

Paivio, A., T. Rogers and P. Smythe (1968). "Why are Pictures Easier to Recall than Words?" Psychonomic Science **11**: 137-138.

Paivio, A., J. Yuille and S. Madigan (1968). "Concreteness, Imagery, and Meaningfulness Values for 925 Nouns." Journal of Experimental Psychology **76**(1, part 2.).

Robertson, G., M. Czerwinski, K. Larson, D. Robbins, D. Thiel and M. vanDantzich (1998). Data Mountain: Using Spatial Memory for Document Management. Proceedings of UIST'98 ACM Conference on User Interface Software and Technology, November 1--4., San Francisco, California.

Robertson, G., M. vanDantzich, M. Czerwinski, K. Hinckley, D. Thiel, D. Robbins, K. Risden and V. Gorokhovsky (2000). The Task Gallery: A 3D Window Manager. Proceedings of CHI'2000 Conference on Human Factors in Computing Systems. The Hague, The Netherlands, April 1--6**:** 494--501.

StJohn, M., H. Oonk and M. Cowen (2000). Using Two-Dimensional and Perspective Views of Terrain. San Diego, SPAWAR Systems Center.

Tavanti, M. and M. Lind (2001). 2D vs 3D, Implications on Spatial Memory. Proceedings of IEEE InfoVis 2001 Symposium on Information Visualization, San Diego, 22--23 October.

Vicente, K., B. Hayes and R. Williges (1987). "Assaying and Isolating Individual Differences in Searching a Hierarchical File System." Human Factors **29**(3): 349--359.

Ware, C. and G. Franck (1996). "Evaluating Stereo and Motion Cues for Visualizing Information Nets in Three Dimensions." ACM Transactions on Graphics **15**(2): 121--139.

Wickens, C., C. Liang, T. Prevett and O. Olmos (1995). Egocentric and Exocentric Dipslays for Terminal Area Navigation. Proceedings of the Human Factors and Ergonomics Society 38th Annual Meeting**:** 16--20.

Wickens, C., O. Olmos, A. Chudy and C. Davenport (1997). Aviation Display Support for Situation Awareness, Aviation Research Lab, Institute of Aviation, University of Illinois at Urbana-Champaign.