

INTERACT INTEGRATE IMPACT

Proceedings of the 20th Annual Conference
of the Australasian Society for Computers in
Learning in Tertiary Education (ASCILITE)

Adelaide, Australia
7–10 December 2003

Editors

Geoffrey Crisp, Di Thiele, Ingrid Scholten, Sandra Barker, Judi Baron

Citations of works should have the following format:

Author, A. & Writer B. (2003). Paper title: What it's called. In G.Crisp, D.Thiele, I.Scholten, S.Barker and J.Baron (Eds), *Interact, Integrate, Impact: Proceedings of the 20th Annual Conference of the Australasian Society for Computers in Learning in Tertiary Education*. Adelaide, 7-10 December 2003.

ISBN CDROM 0-9751702-1-X WEB 0-9751702-2-8



Published by ASCILITE www.ascilite.org.au

3D ENVIRONMENTS FOR SPATIAL LEARNING: THE IMPORTANCE OF LEARNING TASK DESIGN

Barney Dalgarno

Charles Sturt University, AUSTRALIA
bdalgarno@csu.edu.au

Barry Harper

Faculty of Education
University of Wollongong, AUSTRALIA
bharper@uow.edu.au

Abstract

This paper describes two empirical studies which investigated the importance for spatial learning of view control and object manipulation within 3D environments. A 3D virtual chemistry laboratory was used as the research instrument. Subjects, who were undergraduate University students (34 in the first study and 80 in the second study), undertook tasks in the virtual laboratory and were tested on their spatial knowledge through written tests. The results of the study indicate that view control and object manipulation enhance spatial learning but only if the tasks that the learner undertakes require this learning. These results have implications for educational designers making a choice between video or animation and interactive 3D technologies. The results are discussed within the framework of Piaget's theories on active learning and Gibson's ecological theory of perception and action.

Keywords

3D environments, virtual environments, spatial perception, spatial learning, active learning

Introduction

In a paper at last year's ASCILITE conference we argued that the potential learning benefits of 3D Learning Environments (3DLEs) depend very much on the degree to which they can facilitate spatial learning. We argued that this is the case whether the environment is designed for learning about a real space or learning of abstract concepts through spatial metaphors (Dalgarno, Hedberg and Harper, 2002). We then went on to identify the unique characteristics of such environments. These are presented in Table 1. We concluded by establishing a research agenda towards establishing the degree to which each of these characteristics is important for spatial learning. This paper focuses on the degree to which control over view position and direction and the ability to manipulate objects in the environment contribute to spatial learning. Results from two studies of student spatial learning in a virtual environment are presented.

The focus on view control and object manipulation is of particular importance for educators considering the development of an interactive 3D environment as a learning tool. If the focus of the learning were on a real environment the alternative would typically be the use of a video. If the focus was on abstract concepts, then the alternative might be the use of an animation. View control and object manipulation are the key features that distinguish a 3D environment from a video or animation. Given that the cost of developing an interactive 3D environment is normally greater than the cost of developing a video or animation, it is important that the learning benefits of view control and object manipulation can be demonstrated.

Category	Characteristic
Fidelity	Realistic display, including 3D perspective, lighting and occlusion
	Smooth update of views showing viewer motion or panning
	Smooth display of object motion
	Consistent modelling of object behaviour
Learner activity	Control over view position and direction
	Object manipulation
	Control over object model and simulation parameters

Table 1. Unique Characteristics of 3D Learning Environments

The two studies described in this paper were broad in scope, investigating the importance of a number of the unique aspects of 3D environments for spatial learning. Due to space restrictions this paper reports only part of the results from these studies, that is, those parts that address the importance of view control and object manipulation. For example in the first study, there were three groups of students, a user control group, a dynamic views group, who's visual experience was similar but without the degree of control, and a third group who undertook tasks in the real laboratory. The results from this third group are not presented, as they do not contribute to the outcomes addressed in this paper. In the second study, there were again three groups, a user control group, a dynamic views group and a static views group, who viewed a series of still images of the laboratory. Again the results from the third group are not presented in this paper. In each of the two studies a range of written tests were used in addition to questionnaires and, in the case of the second study, tests in the real laboratory. This paper reports the results of two parts of the written tests only, parts D and F. These parts focus on particular aspects of spatial learning that contribute specifically to this paper. The other results from these studies will be reported in further publications.

Study 1

Method

Overview

In this study, participants explored either a version of a virtual environment based on a chemistry laboratory or the real laboratory itself and were then tested on their spatial knowledge. The participants were divided into three groups, a real lab group, a user control group who explored a virtual environment with control over their position and view direction and with the ability to examine objects, and a dynamic views group who viewed an animated tour of the laboratory with control only over the pace. In this paper only the results from the user control and dynamic views groups are discussed.

Participants

The participants were University Information Technology students, who volunteered to participate. In all, 34 people participated in the study, 20 males and 14 females. The participants were randomly allocated to three groups. 6 males and 5 females were allocated to the dynamic views group, 7 males and 5 females were allocated to the user control group, and 7 males and 4 females were allocated to the real lab group. One male member of the user control group struggled with the navigation interface and consequently viewed only part of the virtual laboratory. This participant's data was excluded from the study, leaving 6 males and 5 females in the user control group.

Virtual Laboratory

The study used a 3D environment based on the Charles Sturt University Wagga Wagga campus undergraduate chemistry laboratory. We refer to this environment as the virtual laboratory. Both the virtual laboratory and the environment used for training the participants were developed using the Virtual Reality Modelling Language (VRML) and were explored through Internet Explorer 5.5 and the Blaxxun Contact VRML browser version 5.104, using a PC with a 15 inch screen and a standard keyboard and

mouse, running Windows 2000. The PCs had basic hardware acceleration, allowing a frame rate of between 5 and 15 frames per second (depending on the part of the virtual environment visible at the time). The virtual environments used the default VRML field of view of 45 degrees. The participants were allowed to choose their distance from the monitor, so the effective field of view was not controlled. Internet Explorer was configured to run full-screen, so that none of the Internet Explorer options, or the Windows taskbar were visible.

The version of the virtual laboratory used by the user control group contained a text area and a virtual environment area. Participants were able to move through the environment and look to the left or right by using the arrow keys. If they clicked on a cupboard door or a draw, it opened and their viewpoint was adjusted so that they could see into it. It would then shut when they clicked on it again. If they clicked on an item of apparatus the name of the item was displayed in the text area. If they double-clicked on an item of apparatus the item would move up close so that it could be inspected. If the object was clicked while being inspected, it would automatically rotate left-right and if clicked again it would rotate up-down. They could then put the item back down by double-clicking on it again. See figure 1 for a screen shot showing the virtual laboratory.



*Figure 1. The Virtual Chemistry Laboratory used in Study 1 by the User Control Group
This shows a view after clicking on the fume cupboards*

The dynamic views group were provided with a similar virtual environment, containing a text area, and a virtual environment area, along with Next View and Previous View options. These two options were their only way to move through the environment. They began outside the laboratory and each time they clicked on the Next View option they moved to a new location on a tour of the laboratory. The name of the location or the part of the laboratory they were shown was displayed in the text area. Sometimes when they clicked on Next View they were taken to an item of apparatus and the name of that item was displayed in the text area. In some cases a draw or cupboard was first opened. The subsequent time they clicked on Next View, the item was moved up close for inspection and rotated around so that they could view it from all sides, before returning to its position. See figure 2 for a screen shot showing the controls provided to the dynamic views group.



Figure 2. The Virtual Chemistry Laboratory used in Study 1 by the Dynamic Views Group
This shows a part of the screen only. Notice the Next and Previous options.

Training Procedure

The participants began with 10 minutes of guided exploration of a 3D environment, modelled on an art gallery, with screen layout and navigation options the same as in the corresponding virtual environment used by that group.

Learning Tasks and Procedure

After the training phase, participants explored their allocated version of the virtual laboratory for a period of 40 minutes. They were asked to learn the layout of the lab, locate as many items of apparatus as they could, and to learn the structure of each item. They were asked to indicate which items of apparatus they located by ticking a list provided and to make any additional notes that they thought would help them to remember the layout of the lab and its apparatus.

Test Tasks and Procedure

Participants undertook a test on completion of the learning phase of the study. The test was delivered on paper but with supporting materials on computer (through a simple web site) to allow for the use of full colour images. Participants recorded their responses on paper and each part was completed and submitted before commencing the next part. Only the results of parts D and F of the test are discussed in this paper.

Part D required participants to draw a plan of the laboratory, given an outline of the lab, and given a list of 20 items of furniture to include and a colour photo of each. If there were multiples of a particular item of furniture this number was listed next to its name. This test was designed to determine the participants' spatial cognitive model of the relative positions of the laboratory. They were given 7 minutes to complete this task. The ability to draw a map is a common way to test the degree to which a person has formed a configurational model of the space, as distinct from a view-dependant model. Hunt & Waller (1999) note that "a good map is always evidence of a good [cognitive] representation, but a bad map may simply be a sign of a poor artist" (p.8). They suggest that requiring the learner to fill in blank areas of a map can be better than requiring the learner to draw a map on a blank sheet of paper. Providing the learner with an outline of the lab including the external walkway, as a starting point, is consistent with this. This is very similar to the test used by Wilson, Foreman, & Tlauka (1997) where they required participants to draw a map of a building given a piece of paper showing a map with one room already drawn. Witmer, Bailey & Knerr (1996) also advocate the use of a sketch map for measuring configurational knowledge, but note that it can be difficult to score the learner's maps. The scoring mechanism is discussed in the results section.

Part F of the test required participants to indicate the location of a list of 10 items of apparatus, given a plan of the laboratory, including labelled furniture, and given a colour photo of each item. This tested recall of apparatus locations in relation to a topological lab representation. They were given 5 minutes to complete this task. The positioning of objects on a topographical map is a common way to test a person's recall of landmarks within an environment. For example Ruddle, Payne & Jones (1997), in replicating a famous study by Thorndyke & Hayes-Roth (1982), required participants to indicate the position of named landmarks on a piece of paper which showed the position of other named locations.

Results

Recall of Laboratory Layout

Part D of the test required participants to draw a labelled plan of the laboratory showing furniture, doors and windows, given an outline of the lab with the walkway indicated to show the orientation, and given a list of items to include. As discussed above, the scoring of participant's maps can be problematic. Some researchers have come up with complex measures of the overall error in distance or orientation. For example Patrick, Cosgrove, Slavkovic, Rode, Verratti, & Chiselko (2000) scored participants maps as follows. "For each landmark pair (${}_{10}C_2 = 45$), we oriented and scaled the entire reported map until the pair matched its analog in the virtual environment. Distance error (in meters) was then calculated for the remaining eight transformed landmarks. The cumulative placement score (360 distance measurements per participant) evenly weights every landmark relationship." (p.482-483). The problem with approaches such as this one is that the resulting numbers have no intrinsic meaning. For this study it was desirable to get a sense for how many objects were recalled correctly and thus a scoring mechanism, which defined whether an object was correctly or incorrectly placed, was used. Specifically, a mark was given for each item placed within 2m of the correct location, as long as it was between 50% and 200% of the correct size and within 2m of its correct size in each direction, in the correct room and orientated correctly. There were 43 items to include on the plan so the maximum possible score was 43.

The user control group had a mean score of 28.0 items, as compared to the dynamic views group with a mean of 23.0 items. An ANOVA comparing the three group means suggested that there may have been a group effect on performance on this part of the test ($p=0.07$). Post Hoc analysis using Tukey's HSD test showed that there was no significant difference between the user control and dynamic views groups ($p = 0.29$).

This is contrary to the expected results. It was expected that the control of position and viewpoint of the user control participants would have allowed them to develop a more complete spatial cognitive model. The possible reasons for this result are discussed below in the discussion section.

Location of Apparatus

Part F of the test required participants to indicate the location of a list of 10 items of apparatus, given a plan of the laboratory, including labelled furniture, and given a colour photo of each item. The number of correctly placed items was recorded. Where an item appeared in more than one location, the number of correct locations identified was recorded. The maximum possible score was 13 (three of the ten items appeared in two locations). If the participant's plan showed an item in the adjacent storage location to the correct location they were given a half a mark. For items not found in a cupboard or draw (such as the fire extinguishers) or found in a large storage location (such as the lab benches) they were marked as correct if within 2m of the correct location. If they were placed between 2m and 3m from the correct location, a half a mark was awarded.

Participants in the dynamic views participants had a mean of 6.95 items, which was slightly higher than participants in the user control group who had a mean score of 6.73 items. However an ANOVA comparing the three groups indicated that group was not a significant factor in performance on this test item ($p=0.11$). The results of parts D and F of the test are summarised in table 2.

Again it was expected that the user control group would perform better than the dynamic views group because it was assumed that they would have been more cognitively active in searching for items of apparatus than the dynamic views group who moved through a sequence of views, being shown the location of items without any control. The possible reasons for this finding are discussed below in the discussion section.

	Test D. Recall of Laboratory Layout Structural items and furniture correctly placed on map (out of 43)				Test F. Location of Apparatus Items of apparatus correctly located on map (out of 13)			
	\bar{x}	σ	ANOVA 3 groups	Tukey HSD	\bar{x}	σ	ANOVA 3 groups	Tukey HSD
User Control Group (n=11)	28.0	6.0	p=0.07 marginal	p=0.29 not significant	6.7	3.4	p=0.11 not significant	not required
Dynamic Views Group (n=11)	23.0	10.6			7.0	3.1		

Table 2. Results of Study 1

Discussion

As expected, participants in both the user control and dynamic views groups were able to demonstrate substantial learning as a result of their exploration of the virtual lab. However the finding that there was no significant difference between the performance of the dynamic views and user control groups was surprising. It was expected that control over position and view and the ability to locate and freely explore items of apparatus would have led to a more accurate and more complete spatial cognitive model for user control participants. Two reasons for this finding can be postulated.

Christou & Bulthoff (1999), in attempting to explain a similar finding suggested that the user interface provided for moving around and manipulating items of apparatus may have imposed an additional cognitive load on the users. In their study they used a space-ball, which is a six degree of freedom mouse, held above the desk, and unfamiliar to all participants. Peruch, Vercher & Gauthier (1995) carried out a study comparing the performance of an active and a passive group on a spatial learning task and found that the active participation group performed better on spatial knowledge tests than the passive group. Their study used a joystick, which was likely to be easier to use than a space ball, supporting Christou & Bulthoff's explanation.

However, the motion control interface used in this study was very simple, with movement constrained to ground level and with the use of the arrow keys and the shift key to specify movement or changes in view direction. Its development was informed by comparative studies of desktop 3D motion control (Dalgarno & Scott, 2001) and had been successively simplified as a result of two pilot studies. Observations during this study suggested that, apart from a single participants who encountered significant difficulty and was excluded from the study, all participants found the interface easy to learn and comfortable to use. Additionally, the comparative studies found that an arrow key motion control interface is as easy or easier to use than a joystick interface (Dalgarno and Scott, 2001).

Consequently, it would seem that additional cognitive load was not the reason that the user control group's performance was no better than the dynamic views group. An alternative explanation may be that the tasks that the user control participants carried out did not require that they developed a spatial cognitive model and thus the fact that they were more active in their learning process did not lead to learning benefits. If this is indeed the case then it is an important finding because it means that developers of 3D learning environments may need to be much more careful in designing learning tasks. The second study investigated this, by providing user control participants with a task requiring them to locate items of apparatus, carry the items to a bench, assemble and use them and then return the items to their correct cupboard or draw. It was hypothesised that learners would as a consequence of carrying out this task develop a better spatial cognitive model of the laboratory.

Study 2

Method

Overview

In this second study participants were again divided into three groups, a user control group, a dynamic views group and a static views group. Each participant explored a virtual environment and was tested on their spatial knowledge through a written test. The user control group explored a virtual environment with control over their position and view direction and the ability to pick up, carry and place objects. The dynamic views group viewed an animated tour of the laboratory with control only over the pace. The static views group viewed a similar tour but made up of still images only. The results from the static views group are not discussed in this paper.

Participants

Participants were 1st year University chemistry students. The study was carried out during class time in an introductory subject and all students were expected to participate. Students were asked for their consent in order for their results to be used in the study and all except one student gave this consent. The results from 11 students were excluded because these students had been in the laboratory prior to the study. There were 31 females and 49 males whose results were used. The participants were randomly allocated to three groups. Because some students did not turn up on the day, the number of students in each group differed slightly. There were 10 females and 14 males in the user control group, 9 females and 17 males in the dynamic views group and 12 females and 18 males in the static views group.

Virtual Laboratory

The virtual environment used in this study was very similar to the environment used in the first study. The main enhancements were some minor improvements to the interface used for moving and looking around the environment and the addition of tools allowing the user to pick up, carry and place objects. See figure 3 for a screen shot showing the new interface.



*Figure 3. The Virtual Chemistry Laboratory used in Study 2 by the User Control Group
This shows a part of the screen only, after picking up a conical flask*

Procedure

Before using the virtual laboratory, participants undertook training similar to that undertaken in the first study. After the training, participants of the user control group were given a printed worksheet listing a series of tasks to complete in the virtual laboratory. The tasks consisted of locating a series of items of apparatus, carrying these items to a bench in the lab, connecting the items together as they would to undertake an experiment and putting them away again. Participants of the dynamic views group instead viewed a series of animated images equivalent to what they would have seen had they undertaken this task. They were given a similar worksheet to the user control participants so that they had a similar sense for the overall task. Participants were allowed a maximum of 45 minutes. Dynamic views participants were encouraged to view the animated images a second time to fill in the 45 minutes.

Test Tasks and Procedure

As with the first study, participants undertook a written test after exploring the virtual environment. This test was divided into 6 parts (parts A to F). Only parts D and F are discussed in this paper. Part D was very

similar to part D in the first study, where participants were required to draw a plan of the laboratory given a list of items of furniture and laboratory structures to include. The specific items listed differed slightly from those used in the first study. Part F was similar to part F in the first study, where participants were required to indicate the position on a plan of the laboratory where they could expect to find each of 11 items of apparatus. The list of items of apparatus differed substantially to those used in the first study. In order to provide an authentic task it was necessary to choose items that a student could realistically expect to carry to a bench as part of an experiment.

Results

Recall of Laboratory Layout

As with the first study part D of the test required participants to draw a labelled plan of the laboratory showing furniture, doors and windows, given an outline of the lab with the walkway indicated to show the orientation, and given a list of items to include. Scoring was done in an identical way to the first study, with a score of 1 for each correctly placed item. There were 41 items to include on the plan so the maximum possible score was 41.

The mean score for user control participants was 22.75 items, which was very similar to the mean for dynamic views participants, which was 21.35 items. An ANOVA comparing the three group means suggested that there was no group effect on performance on this part of the test ($p=0.98$).

The results for the two studies can be compared in percentage terms. Although the mean for the dynamic views group on this study (52%) was similar to the mean for the corresponding group in the first study (53%), the user control group in this study (55%) performed substantially worse than their counterparts in the first study (65%). This is discussed further below.

Location of Apparatus

Part F of the test required participants to indicate the location where each of a list of 11 items of apparatus would normally be found, given a plan of the laboratory, including labelled furniture, and given a colour photo of each item. The scoring used was similar to the first study, with the number of correctly placed items recorded and a half a mark recorded for items in the adjacent storage location to the correct location and within 2.5m of the correct location. In this study only one location of each item was required, so the maximum score was 11.

The mean for user control participants was 5.62 items as compared to dynamic views participants who had a mean of 4.10 items. An ANOVA comparing the three groups indicated that group was a factor in performance on this test item ($p=0.00$). Post Hoc analysis using Tukey's HSD test showed that the difference between the user control and dynamic views group was significant at the 95% level ($p = 0.05$). The results of parts D and F of the test in study 2 are summarised in table 3.

	Test D. Recall of Laboratory Layout Structural items and furniture correctly placed on map (out of 41)				Test F. Location of Apparatus Items of apparatus correctly located on map (out of 11)			
	\bar{x}	σ	ANOVA 3 groups	Tukey HSD	\bar{x}	σ	ANOVA 3 groups	Tukey HSD
User Control Group (n=24)	22.8	11.5	p=0.90 not significant	not required	5.6	2.2	p=0.00 significant	p=0.05 significant
Dynamic Views Group (n=26)	21.4	12.2			4.1	2.8		

Table 3. Results of Study 2

Gender

Because the ratio of males to females in each group differed slightly it is important that we can be confident that gender was not a factor. The overall mean for males on test part D was 20.6 and for females was 24.7. This difference was not significant ($p=0.13$). There was also no significant group-gender

interaction on test part D ($p=0.58$). On test part F the overall mean for males was 4.1 and for females was 3.9. Again the difference was not significant ($p=0.55$) And again there was no significant group-gender interactions on this item ($p=0.46$). Thus we can be confident that gender was not a contributing factor in the differences between groups.

Discussion

The crucial difference between this study and the first study is that user control participants on this study carried out a task requiring them to learn the location of items of apparatus. Consequently it was expected that user control participants would score better than dynamic views participants on recall of apparatus location. The fact that this was in fact the case indicates that the ability to move around and manipulate objects within a 3D environment can lead to greater spatial learning than animated views of the same environment, but only if the learning tasks require this learning. Simply providing a 3D environment and allowing the learner to explore it freely without providing any tasks, goals or problems to solve appears not to result in any learning benefits over animations.

The fact that user control participants in the second study seemed to perform worse relative to the dynamic views group on a test requiring them to recall the layout of the laboratory than their counterparts on the first study may be because they were so focussed on the task of locating items of apparatus that they did not pay as much attention to the layout of the lab.

General Discussion and Conclusion

Much of the literature on spatial perception and cognition comes from an information processing model, which has implicit within it an assumption that the information perceived and the processing that occurs can be treated separately, or factored (Greeno, 1994). Put another way, this model does not consider the task that the learner is attempting to carry out or their motivations to be important factors in determining the information that they will perceive.

Constructivist theories of learning, especially those building on Jean Piaget's stage-independent theories consider the learner's activity to be much more important. These theories suggest that learning is a process of a learner actively testing their existing knowledge against their experience (Piaget, 1958 in Gruber & Voneche, 1977). It could be argued, then, that a consequence of these theories is that what a learner perceives will depend to a large extent on their activity. However, Piaget differentiates between figurative knowledge, which is knowledge about the static world, and operative knowledge, which incorporates the dynamic properties of objects and the ways in which they can be manipulated (Piaget, 1968). Piaget's focus in explaining the learning process was primarily on operative knowledge, with figurative knowledge considered less important (Campbell, 1997). Consequently, it could be argued that it would be inappropriate to apply Piaget's views on active learning to the perception and recall of laboratory layout and location of objects.

James, J. Gibson's theories, on which the discipline of ecological psychology is based, give perception a much higher status. Gibson suggests that perception and action are very heavily intertwined. Greeno (1994) citing Bickhard & Richie (1983) notes that Gibson viewed perception as an aspect of a person's interaction with the environment, rather than simply involving encoding of the information about the environment. Gibson introduced the notion of affordances and suggested that our perception is primarily focussed on identifying the affordances of the objects around us. In other words "in ecological theory, it is assumed that perception exists to facilitate adaptive action" (Stoffregen, 2000, p.18). Consequently, what we perceive will depend on the activities that we are engaged in at a more general level (Greeno, 1994). The results found in these studies are consistent with this theory in that for the learners in the second study perception of the location of items of apparatus was necessary to afford the task of retrieving the items and putting them away again. For the learners in the first study perception of apparatus location was not essential to the task that they undertook and so their perception and retention was less.

The studies described in this paper provide an initial indication that 3D environments don't necessarily provide for greater spatial learning than alternatives such as video or animation, but they may if the learners undertake tasks for which this spatial learning is central. The interactive capabilities of 3D environments make it possible for authentic tasks to be carried out, but it is important that the expectation that learners undertake such tasks is made explicit either through support materials or within the environment. If instead learners are simply presented with an environment to explore it is likely that there will be no learning advantages over alternatives such as video or animation.

References

- Campbell, R. L. (1997). *Jean Piaget's Genetic Epistemology: Appreciation and Critique* (online). <http://hubcap.clemson.edu/~campber/piaget.html> [Accessed 18 July 2003].
- Christou, C. G. & Bulthoff, H. H. (1999). View dependence in scene recognition after active learning. *Memory and Cognition*, 27(6), pp. 996-1007.
- Dalgarno, B., Hedberg, J. & Harper, B. (2002). The Contribution of 3D Environments to Conceptual Understanding. In A. Williamson, C. Gunn, A. Young and T. Clear (Eds) *Winds of change in the sea of learning: Charting the course of digital education, proceedings of the 19th annual conference of the Australasian Society for Computers in Learning in Tertiary Education* (pp. 149-158). Auckland, NZ: UNITEC Institute of Technology.
- Scott, J. & Dalgarno, B. (2001). Interface Issues for 3D Motion Control. *Proceedings of OzCHI 2001*, Perth.
- Greeno, J. G. (1994). Gibson's affordances. *Psychological Review*, 101(2), pp.336-342.
- Gruber, H.E. & Voneche, J.J. (1977). *The Essential Piaget*. London: Routledge and Kegan Paul.
- Hunt, E. & Waller, D. (1999). Orientation and wayfinding: a review (ONR technical report N00014-96-0380). Office of Naval Research, Arlington, VA.
- Patrick, E., Cosgrove, D., Slavkovic, A., Rode, J. A., Verratti, T. & Chiselko, G. (2000). Using a large projection screen as an alternative to head-mounted displays for virtual environments. *CHI 2000 Conference Proceedings*. ACM Press, New York.
- Peruch, P., Vercher, J. L. & Gauthier, G. M. (1995). Acquisition of spatial knowledge through visual exploration of simulated environments. *Ecological Psychology*, 7, pp. 1-20.
- Piaget, J. (1968). *Genetic Epistemology*. Columbia University Press, New York.
- Ruddle, R. A., Payne, S. J. & Jones, D. M. (1997). Navigating buildings in "desk-top" virtual environments: Experimental investigations using extended navigational experience. *Journal of Experimental Psychology: Applied*, 3(2), pp. 143-159.
- Stoffregen, T. A. (2000). Affordances and Events. *Ecological Psychology*, 12(1), pp.1-28.
- Witmer, B. G., Bailey, J. H. & Knerr, B. W. (1996). Virtual spaces and real world places: transfer of route knowledge. *International Journal of Human-Computer Studies*, 45, pp. 413-428.
- Wilson, P. N., Foreman, N. & Tlauka, M. (1997). Transfer of spatial information from a virtual to a real environment. *Human Factors*, 39(4), pp.526-531.

Acknowledgements

The authors would like to acknowledge the contribution of John Hedberg to the design of the studies reported in this paper.

Copyright © 2003 Barney Dalgarno and Barry Harper

The author(s) assign to ASCILITE and educational non-profit institutions a non-exclusive licence to use this document for personal use and in courses of instruction provided that the article is used in full and this copyright statement is reproduced. The author(s) also grant a non-exclusive licence to ASCILITE to publish this document in full on the World Wide Web (prime sites and mirrors) and in printed form within the ASCILITE 2003 conference proceedings. Any other usage is prohibited without the express permission of the author(s).