Education and Information Technologies **3** 247–259 (1998)

Collaboration in a virtual world: support for conceptual learning?

PAUL BRNA
Computer Based Learning Unit, Leeds University Leeds LS2 9JT, England, UK
E-mail: paul@cbl.leeds.ac.uk

ROB ASPIN
Computing Department, Lancaster University, Lancaster LA1 4YR, England, UK
E-mail: aspinr@comp.lancs.ac.uk

Immersive and semi-immersive Virtual Reality (VR) systems have been used for training in the execution of procedures, in exploring (often static) 3D structures such as architectural designs or geographical features, and in designing buildings or constructing molecules. In a separate line of technological development, the availability of distributed computing capabilities has led to VR systems that provide facilities for groups of students that are geographically separated to learn together in a collaborative manner. However, relatively little work has been done to investigate the advantages of such Collaborative Virtual Environments (CVEs) for learning the underlying conceptual content.
A pilot study is described which features several worlds designed as part of the Distributed Extensible Virtual Reality Laboratory (DEVRL). The basic results are presented along with a discussion as to how the research could be moved forward to provide improved support for conceptual learning. The discussion also raises the issues of how the interfaces design affects conceptual learning; of navigation and conceptual learning; of the role of collaboration in learning; and of the difficulties associated with constructing dynamic VR worlds.

KEYWORDS: collaborative learning; Human-Computer Interface (HCI).

INTRODUCTION

Desktop and immersive Virtual Reality Environments (VEs) provide new opportunities for learning which are based on (a) first person experience, and (b) increased reliance on sensory information. These two aspects provide advantages over standard 2D simulation or modelling environments from a Piagetian perspective on conceptual learning. Distributed (immersive or semi-immersive) Virtual Reality environments provide the further benefit of improved computer-based opportunities for collaboration in learning about the world. This provides improved support for a Vygotskian view in which learning is mediated through interaction between a person, the objects in their world and the activities shared with a group of learners.

Distributed VR systems intended for collaborative learning are here termed Collaborative Virtual Environments (CVEs). These are of importance owing to the increased interest in both collaborative learning and Virtual Reality. However, there are a number of questions about the value of such systems for conceptual learning:

- Can learners communicate with each other effectively in relation to the conceptual issues?
- Can systems be designed that provide accessible pathways from the phenomenological world into the underlying theory (or theories)?
- Can learners take advantage of the freedom of movement often provided by CVEs to find viewpoints that provide information of optimal value for learning?
- Do CVEs provide the tools to allow learners to construct their own worlds in a simple manner?

In this paper, research into the principles for the design of improved computer-based support for conceptual learning of physics is described. This research has been informed by work on: learning in small groups; computer-support for cooperative work (CSCW), recent work on computer supported collaborative learning (CSCL); and a few examples of systems designed to support collaborative learning in immersive virtual reality environments (VEs). It is also based on a growing understanding of the potential of VR for education and training, and on developments in our understanding of the various kinds of interaction between perception and cognition (Whitelock *et al.*, 1996).

First, the issue of how collaboration can promote conceptual learning is addressed. Then short reviews are provided of previous work on VEs and CVEs for conceptual learning. The Distributed Extensible Virtual Reality Laboratory (DEVRL) is introduced and a pilot study is outlined. Qualitative results are then provided along with a discussion of their significance.

VES AND CONCEPTUAL LEARNING

There is very little detailed research on the effectiveness of VEs on conceptual learning. The most relevant work to date is that of Dede and his colleagues on the ScienceSpace project (Dede *et al.*, 1996b). Byrne has also studied the connection between interactivity and immersion in learning about chemical structure, with the result that immersion seemed less important than interactivity for her context (Byrne, 1996).

Dede *et al.* have performed a formative evaluation of the effectiveness of VEs for the remediation of misconceptions (Dede *et al.*, 1996b). NewtonWorld, one of the three

VEs produced, makes use of multisensory cues (i.e. visual and tactile cues) indicating the presence of potential energy[1], friction etc. Students can also ‘become’ one of the balls in NewtonWorld, or be located at the centre of mass and so on. The other two VEs were MaxwellWorld – an electrostatic world featuring visible equipotential surfaces – and PaulingWorld, a molecular modelling environment. Dede *et al.* have plans to extend this work to look at collaborative problem solving (Dede *et al.*, 1996b) but there is no current report on this aspect as far as we are aware.

The evaluation of ScienceSpace that has taken place was organised around four basic aspects: usability, learnability, usability vs learnability and educational utility (Dede *et al.*, 1996a). The main empirical results for Newton World are that participants given a combination of haptic, audio and visual cues perform the best at learning about velocity and acceleration, while participants receiving haptic and audio feedback were relatively worse at predicting the system’s behaviour. Usability studies have indicated that a gesture-based interface is the least satisfactory compared with (simulated) voice command, menus and a multimodal interface (Salzman *et al.*, 1995). Generally, the multimodal interface (voice and menus) was preferred.

Some of the implications discussed within the ScienceSpace project include: the advantages of multiple perspectives, the use of multisensory cues, 3D visualisations, the use of talk aloud protocols, and the utilisation of a cycle of predict-observe-compare to support the learning process.

Whitelock, Brna and Holland have proposed a research framework for studying the interactions between representational fidelity, immediacy of control (i.e. autonomy and interaction) and presence, for studying conceptual learning (Whitelock *et al.*, 1996). They hypothesise that a high presence value and a high degree of immediacy of control leads to a high degree of tacit learning, while a low value for immediacy of control is more likely to be associated with explicit conceptual learning – assuming that the learner is familiar with the interface!

CVES AND CONCEPTUAL LEARNING

It was pointed out above that adding support for collaboration is in line with Vygotskian ideas of learning – i.e. there is an interaction between social aspects of collaboration and cognition. The process of collaboration involves the maintenance of the groups mutual understanding of the set of goals and the ways in which the goals may be solved (Roschelle and Teasley, 1995; Burton and Brna, 1996). In terms of how collaboration interacts with the conceptual content of what is being learned, Roschelle and Teasley point out that collaboration involves ‘a coordinated, synchronous activity that is the result of a continued attempt to construct and maintain a shared conception of a problem’ (Roschelle and Teasley, 1995). In this view, a

[1]Potential energy was communicated through a form of haptic feedback -essentially a shirt that vibrated!

measure of conceptual learning can be expected during effective collaboration when the problem has been chosen to have some desired conceptual content.

This argument does imply that it is quite possible for groups of collaborating learners to miss opportunities for conceptual learning. Collaboration is not always effective: research on the broad educational utility of collaboration (e.g. Webb, 1991; Cohen, 1994) shows mixed results. The consensus is that collaboration can be productive but Harwood has pointed out the need for learners to develop collaborative skills which include 'listening, questioning, challenging, supporting, giving explanations and evidence, summarising and checking for consensus' (Harwood, 1995). This has implications for the kinds of support that CVEs need to provide.

In short, the potential benefits of collaboration include: learning to develop the social skills necessary to manage such situations more effectively; help in learning to interpret the task, and break it down effectively into its constituents; help in learning to decide on, and maintain the procedures necessary for the performance of the tasks; learning how to maintain a shared understanding of the goals; and determining whether the mutual goals have been achieved. Given appropriately chosen tasks, effective collaboration also provides for some degree of conceptual learning.

Turning to how both collaboration and VEs combine to effect conceptual learning, there are very few research results. This is, in part, due to the absence of readily available environments. Probably the most relevant CVE is being developed by the NICE group at the University of Illinois, Chicago (Roussos *et al.*, 1996).

The GULLIVR environment utilises CAVE technology (Cruz-Neira *et al.*, 1993) to create an immersive VR environment. The CAVE (Cave Automatic Virtual Environment) is an enclosed room designed so that (currently) three walls and the floor portray the environment. Several people can inhabit a CAVE at one time but only one privileged person may act upon the environment. Since the view of the world is from the privileged perspective, others will tend to see increasingly distorted views as they move away from the privileged person. The CAVE also provides a form of distributed VR – it is therefore possible for multiple people to act on the virtual environment if each person inhabits a different CAVE. Educationally, GULLIVR has been used to support a constructivist perspective on learning aspects of environmental science through doing, with an emphasis on encouraging students to construct stories of their experiences in GULLIVR. It remains to be seen what the NICE team can learn from their recordings of student activities, as there are not yet any results in terms of conceptual learning.

Other systems exist that can be viewed loosely as CVEs: many come from the so-called text-based Virtual Environments based on Curtis's MOO (Curtis and Nichols, 1993). MOOS are primarily concerned with interpersonal communication as multiple users are provided with a variety of tools that allow them to discuss the environment and other issues (they are also able to affect the world). Increasingly, these text-based environments have a visual interface. However there has been little research on issues connected with conceptual learning.

The factors present in CVEs and their influence on conceptual learning have not been studied to any great extent. This applies as much to text-based VR as to immersive and semi-immersive VR. However, our intuitions are that the key combination of factors present in CVEs are related to multiple perspectives, multisensory cues, and the process of maintaining a shared conception of the problem through collaboration. The next section outlines the progress we have made in designing environments that might provide the basic level of support required for conceptual learning.

THE DEVRL VIRTUAL CLASSROOM

The Distributed Extensible Virtual Reality Laboratory (DEVRL) project was an EPSRC funded joint project between University College London, Nottingham University and Lancaster University. The research reported here is the development of a Virtual Physics Laboratory containing a variety of virtual environments designed to support collaborative problem solving within an educational context.

The CVEs that have been developed are dynamic, exploiting physical simulations. These VEs pose several problems for students of physics. Firstly, as with 2D simulations, learning to visualise complex systems is difficult. The skill usually needs to be developed through constant practice and application which may lead to the development of an implicit understanding of the system with little or no connection between the body of formal physics underlying the VE and the student's own tacit knowledge.

Secondly, novices have different degrees of prior knowledge, different levels of ability at visualising complex systems, and different capabilities for communicating their knowledge. Students need to develop their understanding of the visualisation at the same time as their growing understanding of the domain, and the need to collaborate is an additional strain.

However, one of the strengths of VEs is that the student is free to find a frame of reference from which the problem can be viewed, and solved more effectively. There is a skill in learning to reperceive the problem which is complementary to that of problem rerepresentation (Amarel, 1968).

The emphasis on collaboration requires the selection of tasks which are hard to perform without collaboration. This rules out, for example, tasks which are very easy to perform for individuals. A brief overview is given of the environments and tasks developed so far.

- Cannon: The world consists of a wall; a target; cannon and cannonball; and two participants. One is initially next to the cannon while the other is initially 'bound' to the cannonball. The task is to hit a target when the participant on the ground cannot see it.
- Table: consists of a pivot; a table; and a number of objects (3) on the table. The

task is to level the table by moving the objects. This is an extension of a simple balance system. While the task is not impossible for a single person, it is quite difficult.

- **Friction:** The environment is a 'snooker table' with 10 differently coloured balls rolling around. Collisions cause changes in velocity but these changes do not necessarily respect the Conservation of Momentum or the Conservation of Energy. As with O'Shea's experiments with the Alternative Reality Kit (O'Shea, 1989), unlabelled sliders control the extent to which the world satisfies the laws of Conservation of Momentum and Energy – Fig. 1 provides more details on this aspect.

Several different functionalities are provided to help users. By clicking the mouse button on a ball the user may jump from ball to ball allowing the user to adopt a number of different frames of reference. During the time they are attached to the ball they may influence the ball's motion by inducing an impulse force along their direction of sight. The user can also bind or unbind two balls. When unbound, the user can select a repulsion force which acts between the two balls. Additionally, the elasticity of each object, and the friction between each ball and the ground can be controlled. In these ways the user is given limited control over the environment.

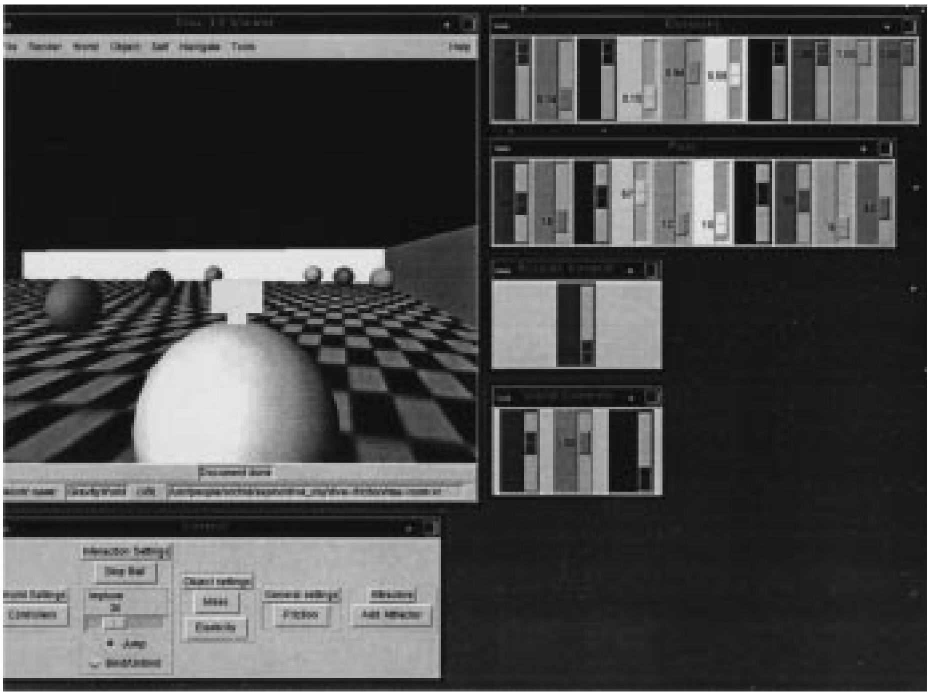

Figure 1. A view of the Friction World
The DIVE environment provides facilities for viewing the VR: from different perspectives (1st person, 3rd person); using different representations (rendering using flat textured, wireframe etc.); with different methods of navigation (keyboard, mouse etc). The avatar is currently attached to the ball in the immediate foreground.

- Bowls: The world contains a green; a ramp, for which the slope and direction can be changed; a jack ball; three bowls for each player; and two sliders, one governing the acceleration due to gravity, and one controlling the frictional coefficient of the green. Balls are rolled down the ramp and, once on the green, sliding friction acts. Users are asked to try to get the bowls as near to the jack as possible. Two concepts of value here are that of an optimal ramp slope, and the relationship between friction and weight.

These initial environments together with the selected tasks were aimed primarily at providing the motivation for collaboration. This approach was complemented with audio support for symbolic (verbal) communication. However, other possible mechanisms for providing symbolic communication have not been developed up to the current time.

The Virtual Physics Laboratory has been built using the SICS Distributed Interactive Virtual Environment (DIVE) on a network of Silicon Graphics workstations. DIVE provides for the development of 3D immersive environments specially designed for multiple users working over the internet. Since it is also specialised for meetings, each participant has a visible representation in the VE – an avatar.

PILOT STUDY

The immersive interfaces consisted of buttons and sliders that the participants could operate within the worlds. Buttons controlled scalar values, such as the elevation of the cannon – see Fig. 2. Sliders in the Bowls World were implemented as a ball the participants could drag up and down, with no numerical representation. The Cannon World only provided for the control of visible properties but the Bowls World featured friction and gravity controllers. However, these provided primarily symbolic rather than iconic information.

The GUIs consisted of sets of forms displayed outside the world – see Fig. 3. These enabled the display of much more complex data, and allowed far greater numbers of widgets to be used. Initially, the participants found these daunting and would start rapidly manipulating many of the widgets together. However, after a familiarisation period, a more disciplined manipulation of the controls started to take place.

Of the two interface types, the immersive one seemed the most easily understood: users quickly grasped the actions necessary to utilise the interface tools. As the controls were always present, the users were naturally reminded as to the scope of their influence over the world. The GUI controls allowed much more detail to be displayed, but required a greater familiarisation with the interface, and provided none of the visual cues given through the immersive interface.

- Materials: The four virtual environments (Table, Cannon, Bowls and Friction) were set up on two SGI machines to operate in distributed desktop mode.
- Participants: 12 volunteer students were paired together arbitrarily. Their primary study areas included mathematics (2), physics (1), computer science (8)

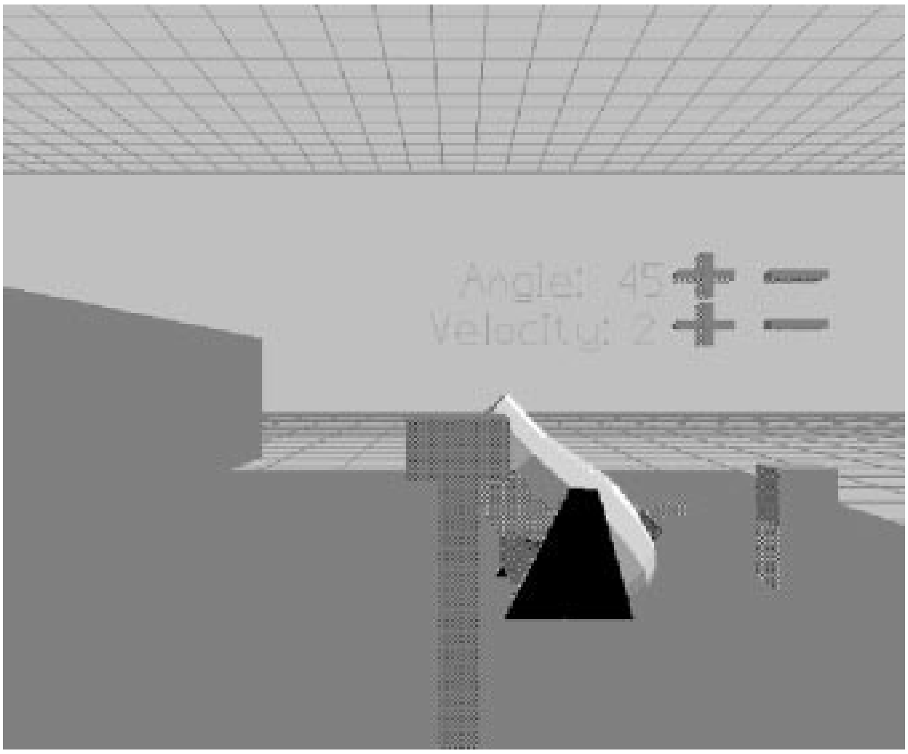

Figure 2. Buttons immersed in the world
Clicking on the + (-) icon increases (decreases) the value of the associated property

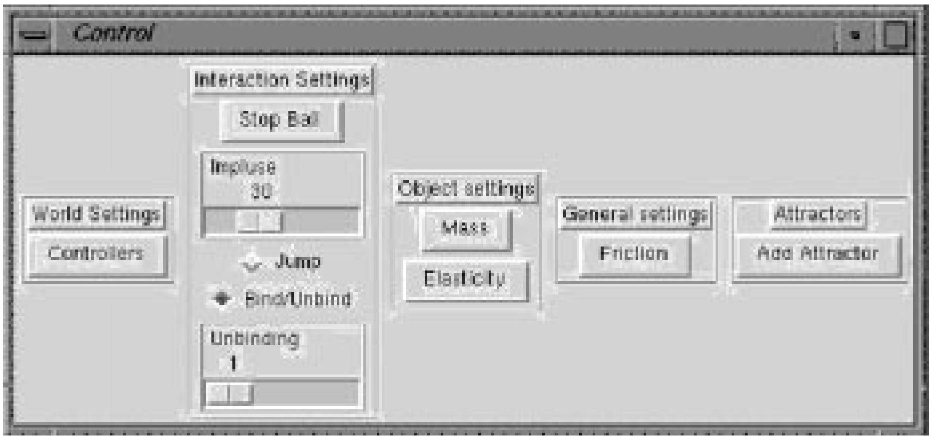

Figure 3. Forms used to define settings

and one participant who was not a computing or science specialist. Two of the participants were female and ten were male.

- Procedure: Each participant in the dyad was placed in front of an SGI machine. A video camera was trained on each machine. The dyads were placed in the same room but in such a way that they were unable to see the screen of the other participant. The researcher present took notes throughout, and intervened if it seemed that the participants were 'lost' or needed some guidance relating to the operation of the interface.

Initially, participants were informed of the reason for the experiments, to create more effective educational CVEs, and were presented with the Table World which allowed them to familiarise themselves with the controls of the environments, with each other, and the interaction modalities available. Before each task, the participants were given a description of the task, explaining the control interactions available. For the Friction World, they were told that their task was to cooperate in changing the world so that it behaved in as 'real' a fashion as they could manage, and that they should decide what the controls actually did in terms of the 'physics' of the situation.

The participants were then asked to perform the task within the given world, deciding between themselves when the task had been completed. Only in the Cannon World were the participants allocated roles. In this case, the experiment was constructed in such a manner as to require one to become the 'aimer' and the other the cannon-ball rider.

Once the familiarisation phase was over, the worlds were taken in the same order. First, the Cannon World where the task was to hit the target, then the Bowls World, followed by the Friction World. The total time for all the worlds was of the order of one hour. During their usage of a world, if they asked for advice about the nature of the controls these questions were answered; questions about the nature of the physics in the experiment were not answered.

RESULTS

The results are necessarily informal. While all sessions were videotaped, detailed transcription analysis has not been carried out. The observations selected for comment are essentially of two kinds: influences on learning in the context of a VE; and issues connected with using the interface. Three basic issues have been selected for consideration in terms of whether (a) the CVEs support the process of the participants' developing and maintaining a shared conception of the problem; (b) the CVEs provide suitable sensory cues to link theoretical concepts with associated phenomena; and (c) the CVEs support effective 'reperception' of the problem.

In passing, though the study was informal, participants showed little sign of any scientific approach to their exploration of the environments. Amongst the participant pairs, only three spent time investigating the individual interface devices. These all had a science background (2 physicists, and a mathematician), as opposed to the computer science background of most of the others! Participants would invariably manipulate multiple interface devices rather than investigating the effect of each device separately. This was particularly evident in the Bowls world where the gravitational acceleration and frictional coefficient sliders were frequently adjusted together.

- Collaboration and learning: Role playing in collaboration can promote effective collaboration – if the roles can be associated with appropriate cognitive processes (Burton *et al.*, 1997). The appropriate selection of role depends on a mixture of the nature of the task, the current problem solving context and the skills and abilities of the collaborators.

Collaboration in the VEs featured a tendency for participants to persist in a given role. The experiments were always carried out in a set order (Table, Cannon, Bowls, Friction). In the Cannon experiment the participants are given two distinct roles, aimer and observer. In the Bowls experiment, which directly followed the Cannon experiment, it was common, at least initially, for the participants to assume the roles they performed in the previous experiment. One participant, typically the one who performed the observation role previously, would position themselves either above the bowling green, or close to the jack and relate observations to the other participant. The second would commonly assume a position at the 'firing' end of the bowling green and perform the aiming role as directed by the observer. As the experiment continued and the participants discovered that there was no regulation of the roles in this world a more flexible approach was adopted.

- The Interface and learning: To a limited extent, the CVEs provide sensory cues to link theoretical concepts with associated phenomena. So far, more effort has been expended on the design of control mechanisms than on additional feedback mechanisms beyond those provided by the simulated world itself. The issue was explored through both an immersive interface (Cannon, Bowls) and a GUI (Friction).
- Navigation and learning: As part of the support for effective 'reperception' of the problem, the user needs to be able to view the world from different locations. The CVEs provide the general facility to move to a new viewing point and some specific facilities to adopt a moving perspective that supports different frames of reference. However, the desktop version of the CVEs turned out to have a problem relating to navigation when users tried to track moving objects.

It was only possible to move the user's avatar along the line of sight. To track any moving objects not travelling directly towards the participant required a process of rotation, motion along the line of view, stopping, and rotating to check the object's position.

It is common for motion to be de-coupled in fully immersive VE systems, where the direction of sight is connected to the Head Mounted Display (HMD) and the direction of motion is indicated by the hand controller. However this de-coupling is not usual for desktop VEs. The highly dynamic nature of the experiments ideally requires a more intuitive control interface, which would implement the de-coupling of motion and sight. Without considerable familiarity with the navigation system used in VEs, users can find the overhead of moving around a VE to be so demanding that few resources are available for conceptual learning. This potential problem was mitigated in the pilot study as most participants admitted to familiarity with the 'Doom-like' navigation system used in the DEVRL environments.

DISCUSSION

We address the four fundamental questions raised in the introduction.

- Effective Communication: If CVEs are to be designed to have a genuinely effective communication function for Virtual Classrooms then it is also necessary to consider the nature of the communications that might be useful to the users. These include discussions about how to solve the problem, who is to do what and when, what the conceptual difficulties are and what has been learned, as well as grounding the communication to remove ambiguities.
  Although the primary emphasis has so far been on communications employing the visual modality, sometimes the content is best communicated through a primarily symbolic modality (linguistic and/or diagrammatic) which will support reflection and discussion and the retelling of learner's activities in different modalities. There is therefore a need to develop better communication interfaces that integrate effectively with the simulations of the physical world usually found within VEs.
- Paths to the underlying theory: There is also a need to exploit multimodal communication more effectively. The sensory modality provided by most VR toolkits has been predominantly visual – as opposed to auditory or haptic. Such toolkits typically allow only solid objects such as balls, walls and ramps to be rendered. This makes conceptual information difficult to display. For example, the total heat radiated from an object per second could be represented by the object's colour but it is not very easy to display the heat field created by such an object. Perhaps for this reason many Virtual Environments don't try to develop additional visualisations over and above the 'natural' visible world, though the three worlds in the ScienceSpace project are honourable exceptions.
- Navigation and Learning: generally learners do not take advantage of the possibilities given by the freedom to move in 3D. Certainly this is an aspect that receives little attention in formal education. There are also serious problems in trying to match learners needs to navigational facilities in desktop VR.
- Tools for Construction: The DEVRL project itself has produced a prototype generic Virtual Laboratory that can be used by students to build their own dynamic worlds and manipulate the world's physics. The problem is that, almost without exception, all the current virtual toolkits require the writing of code to make objects move, bounce, and fall. How can users, whose primary concern is to learn about a particular dynamic system, be expected to write code to simulate a system they do not yet understand?

Recent work on Cocoa[2], a 2D environment designed to encourage children to build

[2]Formerly called 'KidSim'.

simulations without any programming is an exciting way to go beyond a primarily mathematical notation which has a specific relevance to VEs (Smith *et al.*, 1994; Smith and Cypher, 1995). In Cocoa, children 'program by example'. A behaviour is programmed by defining a simple condition-action rule – but in a highly visual way. While the approach cannot entirely avoid the use of symbolic expressions, the principles at work are very attractive.

Although the DEVRL project is no longer funded, work is going on at Lancaster to explore the potential of the generic Virtual Laboratory for students to construct their own dynamic virtual environments, and at Leeds to explore issues relating to collaborative learning and reperception. In terms of learning the underlying physics, further work is also needed to develop improved computer-based support for collaborative learning. In terms of providing sensory cues, the current set of interfaces does not provide sufficient support. Much more work would be needed to provide the right kind of sensory cues relating to controlling the world and much more effort is required to develop better methods of providing visual feedback relating to the dynamics of the system being explored.

ACKNOWLEDGEMENTS

The work described in this paper was funded by EPSRC ROPA Grant 'Distributed Extensible Virtual Reality Laboratory'. The authors thank Gareth Smith for his work on setting up the original virtual classroom and building the Cannon and Table Virtual Environments.

REFERENCES

Amarel, S. (1968) On representations of problems of reasoning about actions. In *Machine Intelligence 3*, D. Michie (ed), Edinburgh University Press, Edinburgh.

Burton, M. and Brna, P. (1996) Clarissa: An exploration of collaboration through agent-based dialogue games. In *Proceedings of the European Conference on Artificial Intelligence in Education*. P. Brna, A. Paiva and J.A. Self (eds), Edicoes Colibri, Lisbon. pp. 393–400.

Burton, M., Brna, P. and Treasure-Jones, T. (1997). Splitting the collaborative atom: how to support learning about collaboration. In *Proceedings of the 8th World Conference on Artificial Intelligence in Education*. B. du Boulay and R. Mizoguchi (eds), IOS, Amsterdam. pp 135–42.

Byrne, C.M. (1996) *Water on Tap: The Use of Virtual Reality as an Educational Tool*. Ph.D. thesis, Human Interface Technology Laboratory, University of Washington.

Cohen, E.G. (1994) Restructuring the classroom: conditions for productive small groups. *Review of Educational Research*, **64**(1), 1–35.

Cruz-Neira, C., Sandin, D.J. and DeFanti, T.A. (1993) Surround-screen projection-based virtual reality: the design and implementation of the CAVE. In *Computer Graphics (SIGGRAPH '93 Proceedings*), J. T. Kajiya (ed), volume 27, pp. 135–42.

Curtis, P. and Nichols, D. (1993) MUDs grow up: social virtual reality in the real world. In *Third International Conference on Cyberspace*, Austin, Texas.

Dede, C., Salzman, M.C. and Loftin, R.B. (1996a) Early research results from using virtual reality to enhance science education. Presented at the 1996 National Conference of the American Educational Research Association.

Dede, C., Salzman, M.C. and Loftin, R.B. (1996b) ScienceSpace: Research on using virtual reality to enhance science education. In *Proceedings of the 1996 ED-MEDIA Conference*. P. Carlson and F. Makedon (eds), AACE, Charlottesville, VA. pp. 172–77.

Harwood, D. (1995) The pedagogy of the world studies 8–13 project: The influence of the presence/absence of the teacher upon primary children's collaborative group work. *British Educational Research Journal*, **21**(5), 587–611.

O'Shea, T. (1989) Magnets, martians and microworlds: Learning with and learning by OOPS. In *Artificial Intelligence and Education: Reflection and Synthesis*. D. Bierman, J. Breuker and J. Sandberg (eds), IOS, Amsterdam. pp. 193–4.

Roschelle, J. and Teasley, S. (1995) The construction of shared knowledge in collaborative problem solving. In *Computer Supported Collaborative Learning*. C.E. O'Malley (ed), Springer-Verlag, Heidelberg.

Roussos, M., Johnson, A.E., Leigh, J., Vasilakis, C.A. and Moher, T.G. (1996) Constructing collaborative stories within virtual learning landscapes. In *Proceedings of the European Conference on Artificial Intelligence in Education*. P. Brna, A. Paiva and J.A. Self (eds), Edicoes Colibri, Lisbon. pp. 129–35.

Salzman, M.C., Dede, C. and Loftin, R.B. (1995) Learner-centered design of sensorily immersive microworlds using a virtual reality interface. In *Proceedings of the Seventh International Conference on Artificial Intelligence and Education*. J. Greer (ed), AACE, Charlottesville, VA. pp. 554–64.

Smith, D.C. and Cypher, A. (1995) KidSim: Child constructible simulations. In *Proceedings of the Imagina'95 Conference*, Monte-Carlo. pp. 87–99,

Smith, D.C., Cypher, A. and Spohrer, J. (1994) KidSim: Programming agents without a programming language. *Communications of the ACM*, **37**(7), 54–67.

Webb, N.M. (1991) Task related verbal interaction and mathematics in small groups. *Journal for Research in Mathematics Education*, **22**(5), 366–89.

Whitelock, D., Brna, P. and Holland, S. (1996) What is the value of virtual reality for conceptual learning? towards a theoretical framework. In *Proceedings of the European Conference on Artificial Intelligence in Education*. P. Brna, A. Paiva and J.A. Self (eds), Edicoes Colibri, Lisbon. pp. 136–41.