The Round Earth Project: Deep Learning in a Collaborative Virtual World

Andrew Johnson, Thomas Moher, Stellan Ohlsson, Mark Gillingham
University of Illinois at Chicago, Chicago, IL 60607, USA
(312) 996-3002 voice, (312) 413-7585 fax
roundearth@evl.uic.edu, www.evl.uic.edu/roundearth

Abstract

The Round Earth Project is investigating how virtual reality technology can be used to help teach concepts that are counter-intuitive to a learner's currently held mental model. Virtual reality can be used to provide an alternative cognitive starting point that does not carry the baggage of past experiences. In particular this paper describes our work in comparing two strategies for teaching young children that the Earth is spherical when their everyday experiences tell them it is flat.

1 Introduction

The concept of a round Earth is not a simple one for children to acquire. Their everyday experience reinforces their deeply held notion that the Earth is flat. Told by adults that the Earth is round, they often react by constructing a mental model of the Earth as a pancake, or a terrarium-like structure with people living on the flat dirt layer inside, or even a dual model with a spherical Earth and a flat Earth coexisting simultaneously [7, 11, 12]. In effect, children attempt to accommodate the new knowledge within the framework of their existing conceptual models, while holding tight to the features of those prior models, thereby inhibiting fundamental conceptual change.

The Round Earth Project is a collaboration among researchers in computer science, education and psychology to investigate two alternative pedagogical strategies for teaching children that the Earth is spherical, and the implications of that fact. One strategy, which we term the *transformationalist* approach, attempts to effect conceptual change by evidencing a breakdown in the children's prior models. The alternative *displacement* strategy, in contrast, attempts to effect learning in an alternative setting (in our case, a small diameter asteroid), free of preexisting biases, and to relate that learning back to the target domain: the Earth.

Virtual reality (VR) technologies are used to support both pedagogical Strategies. In the transformationalist approach, VR is used to simulate the launching of a spacecraft from the Earth's surface and subsequent exploration within a fixed-height orbit. In the displacement approach, VR is used to simulate a small diameter asteroid upon which the learner may walk on a body with a curved horizon, see objects 'appear' from 'below' the horizon, take a long walk around the entire globe and come back to where they started. In both environments, distributed VR technologies are used to provide a collaborative learning environment promoting positive interdependence among pairs of learners.

In section 2 we will briefly discuss the current role of collaborative VR environments for concept learning. In section 3 we will discuss deep learning in more detail and then in section 4 describe the two virtual worlds created for this study. In section 5 we describe how we set up the pilot studies and in section 6 we will discuss the results of those studies. Finally, section 7 gives our conclusions and directions for future work.

2 VR and Learning

Research in conceptual learning and virtual reality is a relatively young field, but growing rapidly. In a recent report by the Institute for Defense Analysis, Christine Youngblut comprehensively surveys work over the past few years in the area, citing approximately 50 VR-based learning applications, and 35 studies which include desktop but exclude text-based virtual environments [13]. There are currently very few VR-based learning environments designed for young children, and only two multiuser virtual educational worlds: Virtual Physics at the University of Lancaster [2], and NICE (Narrative, Immersive, Constructionist/Collaborative Environments) at the University of Illinois at Chicago [4, 10]. Other current educational VR worlds such as the ScienceSpace [3] worlds are being extended to support multiple users. Collaboration encourages conversation, and conversation serves learning by presenting each learner with a slightly different view of the subject matter. An individual is forced to enrich his or her own representation in order to assimilate their partner's discourse. Conversation also improves evaluation. Rather than 'thinking aloud' the participants are talking to each other.

NICE, an exploratory learning environment for children between the ages of 6 and 10, was an explicit attempt to blend several learning and pedagogical themes within a single application. These themes: constructionism, exploratory learning, collaboration, and the primacy of narrative reflect several of the most important educational reform themes of the past three decades. The NICE garden was originally designed as an environment for young children to learn about the effects of sunlight and rainfall on plants, the 'spontaneous' growth of weeds, the ability to recycle dead vegetation, and similar simple biological concepts that are a part of the life cycle of a garden. NICE supported real-time distributed collaboration with voice communication enabled by a real-time audio connection.

While NICE was successful as an engaging social space, and as a driver for collaborative VR, the cooperative learning was unstructured and undirected. As its successor, the Round Earth Project built on the experience gained from NICE and seeks to remedy these deficiencies.

We are focusing our efforts on learning problems that meet four criteria:

1. The learning goal must be important. That is, it must be identified as a component of adult scientific (or other) 'literacy,' as reflected in national learning goals, standards, or benchmarks, such as those published by the National Council of Teachers of Mathematics or the American Association for the Advancement of Science [1, 6]

2. The learning goal must be hard. That is, it must be reflected in the literature of researchers and/or practitioners as difficult to effect and resistant to conventional pedagogical strategies.

3. The learning goal must be plausibly enhanced by the introduction of immersive VR technologies.

4. Finally, VR-based learning environments must be informed by contemporary research in the learning sciences and educational practice.

Young children believe the Earth is flat. More precisely, their mental model of the world separates 'sky' and 'earth' into two parallel layers, one 'above' the other; the two directions 'up' and 'down' are absolute. Empirical studies have demonstrated that telling young children that the Earth is round does not cause their intuitive model to be replaced by a spherical conception of the Earth. Instead children assimilate the new information into their prior knowledge and conclude that the earth is flat and circular [7, 11, 12].

Children's intuitive model of the Earth is natural because it is consistent with, and strongly supported by, everyday experience. Discourse has little impact, not only because words like 'round' are ambiguous but also because talk about the Earth is abstract and cannot compete with the vividness if looking upwards when looking at the sky or seeing ocean with its straight horizon. Pictures have little impact because they require a complex mapping between experience and the 2D plane. To understand a picture of a person on a spherical body, the viewer must project himself or herself into the picture, a cognitive capability beyond very young children. The same is true of a 3D representation such as a globe. However, in VR the children can be immersed in the experience if walking on the spherical surface of a small planetary body such as an asteroid.

Teaching young children that the Earth is spherical seemed a good match for our four criteria:

1. In AAAS Project 2061: Benchmarks for Science Literacy, 5th grade graduates should know "things on or near the Earth are pulled toward it by the Earth's gravity" and "the Earth is approximately spherical in shape." 8th grade graduates should know "everything on or anywhere near the Earth is pulled toward the Earth's center by gravitational force."

2. The existing literature by Vosniadou, Brewer and Nussbaum discuss the difficulty of this learning problem [7, 11, 12].

3. Immersive VR is well suited to giving a person the sense of walking on a spherical object with small diameter, seeing objects 'appear' from 'below' the horizon and returning to the starting point after circumnavigating the sphere.

4. The environments designed in this project emphasize role differentiation with positive interdependence and collaborative learning [5].

3 Deep Learning

Underneath the extensive systems of domain-specific knowledge that a person brings to bear on problems and situations, there are organizing concepts - fundamental ideas - that influence how both direct experience and discourse within that domain are conceptualized. Such deep ideas form the axiomatic core of entire systems of knowledge [8, 9]. When experience or discourse attempts to communicate a deep idea that is both different from and more fundamental than the learner's existing ideas, a paradox occurs. Although the intent is to replace the learner's existing ideas, those existing ideas are the learner's only tools by which to acquire the new idea.

If this learning paradox is real then how is anything new ever learned? Our approach to this question is founded on a distinction between 'transformationalist' and 'displacement' explanations of cognitive change. The transformationalist account assumes that new knowledge is created via operations on prior knowledge. Prior knowledge serves as raw material and the new knowledge is the result of generalization, specialization, or some other cognitive operator, applied to the raw material.

The displacement account of cognitive change assumes that a new understanding of a domain or phenomenon begins by establishing an alternative cognitive starting point - an idea or concept is established outside the learner's existing system of domain knowledge. Initially such an alternative representation might be rudimentary and hence dominated by the prior well-established representation. However, over time, all available representations compete and a representation that is useful in dealing with certain types of situations or problems gradually gains strength and may even displace the previous representation.

The displacement framework suggests a particular instructional strategy for supporting deep conceptual learning - fundamental ideas which contrast with the learner's current ideas need to be established on their own terms before they are brought into contact with the learner's prior ideas. VR, we believe, provides a powerful tool in helping to create such alternative cognitive starting points. We can also use VR to juxtapose and switch between multiple interlinked representations of the same experience. Our natural tendency in assimilating new information is that each facet of reality tends to be conceptualized in only one way, within a single perspective. Impasses on simple problems occur because the thinker assimilates or subsumes the problem under a prior conceptualization that does not support the solution. Switching representations is difficult, but deep learning may require precisely such shifts between alternative representations.

While we believe it is crucial to construct these alternative mental representations, our overall strategy requires a second step. The alternative representations must be brought into contact with the learner's prior knowledge of the domain and absorb or subsume it. Unless the learner brings his or her new experience on the asteroid into contact with the everyday experience of walking on a seemingly flat Earth, the learning objective is not reached. The point is not just to know what it would be like to walk on a spherical planetary body, but to understand that the Earth is such a body. We call this second step 'bridging activities.'

4 Asteroid World & Earth World

In this section we will discuss what features we wanted in the VR experience and then describe the two sets of worlds that we created in this project to compare the displacement and transformationalist approaches, the asteroid world, and Earth world respectively.

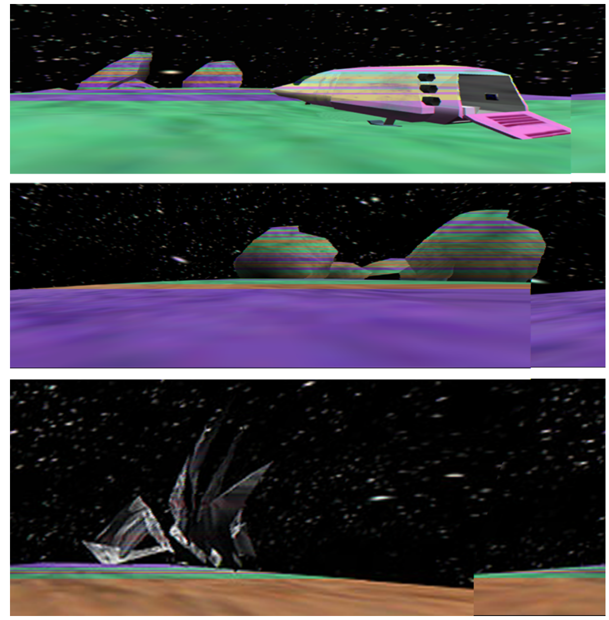

Figure 1: In the Asteroid World, the astronaut explores the surface of the asteroid within the CAVE, moving around the surface collecting fuel cells, guided by Mission Control. This figure shows three locations on the asteroid: the spaceship, some rocks, and a crystal plant

4.1 Goals for the Two Experiences

For the displacement approach the children start off at an alternative cognitive starting point: a small asteroid. Here they can learn about walking around a spherical body different from the Earth. For the transformationalist approach the children start off on the Earth and attempt to transform their current flat Earth model into a spherical Earth model.

We wanted the children to see both the spherical representation of the planetary body as well as the 'flat' view from the surface and integrate these two views. Because of this we made the world collaborative, with one child experiencing the surface of the world and the other seeing the avatar of the first child on the spherical world. We wanted to give the kids a task to perform so that the child on the surface needed to move around the spherical body. This way one of the participants would often be 'upside down' on the sphere but 'right side up' on the surface. We wanted the collaborative task to foster positive interdependence where neither child could perform the task alone; they would need to cooperate and communicate with each other. Through this communication

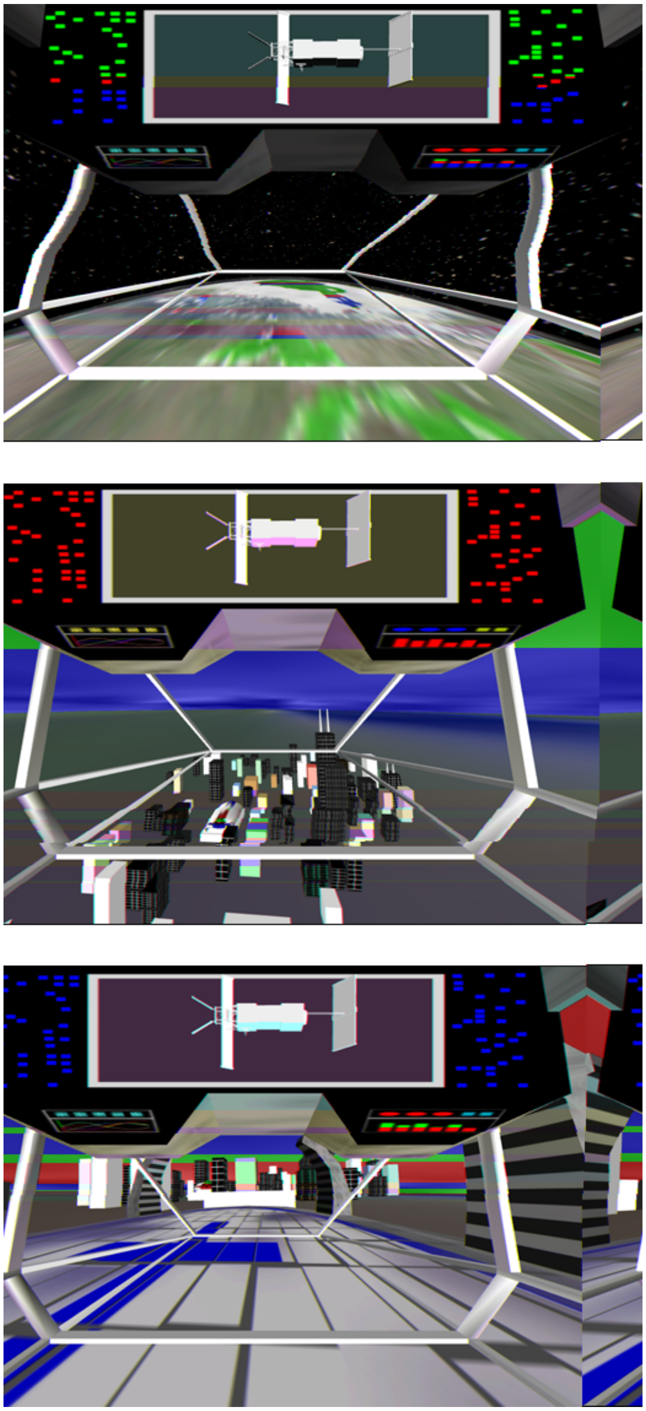

Figure 2: In the Earth world, the astronaut launches from Chicago up into orbit in the cockpit of a spaceship within the CAVE. The astronaut flies around the Earth collecting parts of a broken satellite, guided by Mission Control. This figure shows the Earth falling away as the astronaut launches into orbit.

the children would need to reconcile their different views. We wanted the controls to be simple so there was little training time involved and the children could concentrate on the experience. We wanted to keep them engaged, giving them a long enough experience to grasp the concept but not so long that they became fatigued.

In both worlds the children are given the task of finding ten objects scattered around the planetary body. The roles of the two children are 'the Astronaut' and 'Mission Control.' The astronaut moves around the planetary body collecting each of the parts, guided by the child in mission control. Mission Control sees a spherical view of the planetary body, as seen from an orbiting satellite, and can see the location of each of the ten objects. Each child performs both roles during the experience to see both views.

The astronaut experiences this shared virtual world from a CAVE, increasing their sense of immersion; mission control uses an ImmersaDesk to monitor the 3D planetary body in 3D. The generic controller for the CAVE and the ImmersaDesk is the 'wand' - a six degree of freedom mouse with three buttons and an isometric thumb-controlled joystick. In the CAVE with the astronaut, the left button turns the child to the left, the middle button moves forward, and the right button turns to the right. In order to pick up each of the ten objects, the child simply needs to get within 5 feet of the object and it is automatically grabbed. At the ImmersaDesk with mission control, the joystick is used to spin the world. The world can be turned completely around horizontally with limited tilt of the world. This allows mission control to keep the astronaut in view at all times, but allows the astronaut to be 'rightside up' in the Northern Hemisphere, 'sideways' near the equator and 'upside down' in the Southern Hemisphere.

4.2 Asteroid World Description

In the Asteroid world, the two children find themselves marooned on the surface of a small asteroid and they need to retrieve ten fuel cells from the surface and bring them back to the ship. The astronaut starts out in the airlock of the marooned spaceship and has ten minutes to explore the surface in search of the fuel cells. The child can carry up to four cells and then must return to the ship to drop them off. Mission control sees the astronaut as a person in a spacesuit walking on the surface of the asteroid. After ten minutes, the astronaut is automatically teleported back to the ship. The children then switch roles. When both have had their time on the surface they are told that they successfully completed their mission and they both stand in front of the ImmersaDesk to see their spaceship liftoff from the surface of the asteroid and begin their journey home.

Figure 1 shows a typical view from the astronaut's perspective on the surface of the asteroid. Figure 3 gives a closer look at what mission control sees. This view shows the spherical asteroid with the avatar of the astronaut moving about the surface as well as the view from the CAVE of the astronaut moving on the surface.

4.3 Earth World Description

In the Earth world, the two children must retrieve ten parts from a broken satellite scattered in orbit around the Earth and bring them back for reassembly. The astronaut sits in the command chair of a spaceship on a launching pad surrounded by skyscrapers in downtown Chicago. Since our students are from Chicago, this gives them familiar starting point on a very flat looking Earth. As the engines roar, the astronaut is then launched into space. The astronaut sees the buildings, then the city, then the Earth fall away as he/she rises into orbit to see an Earth with a curved horizon. Once in orbit the astronaut maneuvers the ship close to the satellite parts so that they can be retrieved. Mission control sees the astronaut's pointy space capsule flying over the surface of the Earth. After ten minutes the autopilot engages, maneuvers the ship back over the city of Chicago, and the ship lands back on the launch pad. The children then switch roles. When both have had their time in orbit they are told that they successfully completed their mission and they both stand in front of the ImmersaDesk to see the satellite that has been reconstructed.

Figure 2 shows what the astronaut sees on liftoff. Figure 4 gives a closer look at what mission control sees. This view shows the spherical Earth with the astronaut's spaceship orbiting above the surface as well as the view from the CAVE out the front window of the orbiting ship.

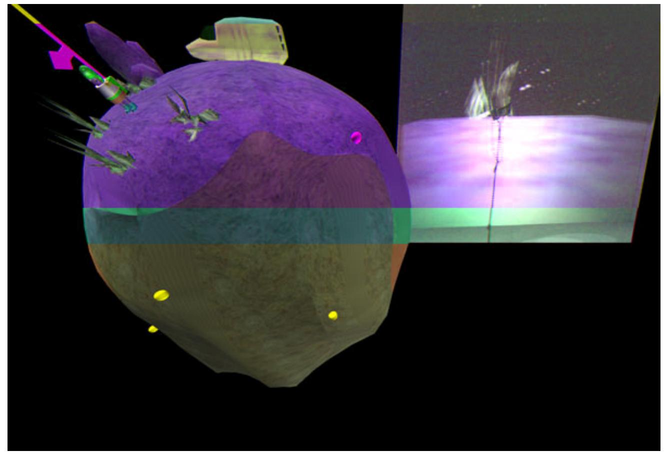

Figure 3: The Mission Control screen for the asteroid world. Mission control simultaneously sees the avatar of the astronaut moving around the spherical asteroid and the view from the astronaut's perspective of the slightly rounded surface of the asteroid.

5 Experimental Setup

The CAVE used is a standard 10' cube, with three walls and one floor. Its graphics are driven by a rack SGI Onyx with one Infinite Reality Engine for each pair of walls. The screen resolution is 1280 x 800 per screen with an average stereo frame rate of 10 frames per second. The astronaut wears a pair of Stereographics glasses to mediate the stereo imagery containing a position sensor for the Flock of Birds tracker, and carries the standard CAVE wand. The astronaut's speech is picked up via an ambient microphone mounted on the top of the front wall of the CAVE. Audio from the application and from the mission control are mixed and sent through the CAVE's four speakers. A low light color CCD camera mounted outside the entrance to the CAVE sends the image of the astronaut and the front CAVE screen into the computer for the ImmersaDesk and into the VCR for recording.

The ImmersaDesk is a standard ImmersaDesk with one 6' by 4' screen. Its graphics are driven by a rack SGI Onyx with one Reality Engine 2 for the screen. The screen resolution is 1280 x 1024 with an average frame rate of 8 stereo frames per second. Mission control wears a pair of Stereographics glasses to mediate the 3D images, but mission control is not head tracked. This was done so that the 3D image of the spherical planet would always be completely on the screen, no matter how active the child was. The camera image from the CAVE is sent in through a Sirius video board and is placed onto the screen. Mission control's speech is picked up via a headworn microphone. Audio from the application and from the astronaut are mixed and sent through the ImmersaDesk's two speakers. A low light color CCD camera mounted behind the child sends the image of mission to the VCR for recording. Audio from both the CAVE and ImmersaDesk microphones are also sent to the VCR.

The pre- and post-tests involve aural, 2D drawing, and 3D sculpting questions to try and elicit the model, or models, the child is using. These questions were adapted from published questions used in earlier studies of children's models of Earth[7, 11, 12].

A human guide was on hand to help the children at the CAVE and ImmersaDesk. Initially the guides were there to simply help with the equipment and the initial setup of the task, though their role became larger as the pilot studies continued. The bridging activities were also modified as the pilot studies progressed. These changes and other issues raised by the pilot studies are described in the following section.

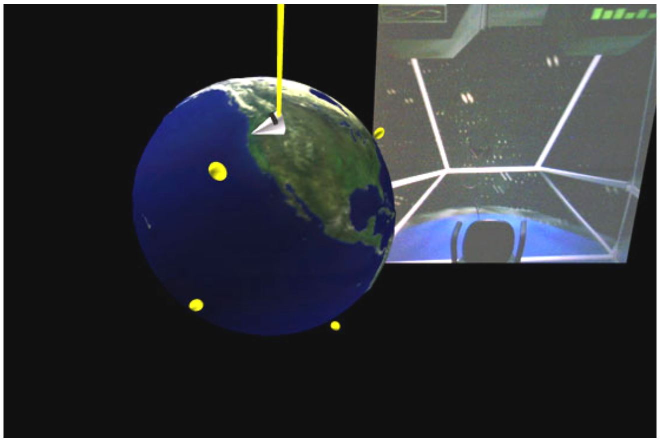

Figure 4: The Mission Control screen for the Earth world. Mission control simultaneously sees the astronaut's spaceship moving around the spherical Earth and the view from the astronaut's perspective of the slightly rounded surface of the Earth.

6 User Studies

To date, we have conducted three pilot studies. The first study consisted of four pairs of children, looking primarily at interface and usability issues in the two worlds. The second study consisted of eight pairs of children concentrating on learning in the two worlds, the effectiveness of the bridging activities, and the ability of the pre- and post-testing to reflect change in the children's models. These led to changes in our procedures; a third pilot study of five pairs of children in the asteroid world was undertaken to evaluate our modified design.

The children came from a small urban Chicago public school. The total K-8 enrollment at the school is about 550. The school is 99 percent African-American and 93 percent of the students come from families below the poverty line. The third grade students score significantly below the state and district averages in reading and math and below the state average in writing. In the first pilot study the children chosen were children of the teachers and administrators at the school, allowing us to familiarize their parents with our procedures. In the second and third pilot studies the children were summer school students who did not pass the Iowa Test of Basic Skills Grade 3 exam.

The first pilot study showed us several things about the usability of the VR worlds. As with the NICE studies, the Stereographics glasses are too big for small children. The only satisfactory solution we have found is to tie the glasses on. Our initial 'realistic' asteroid world interface of having the astronaut reach down to pick up a fuel cell was replaced with automatic grabbing when the astronaut got close enough to the cell. This allowed the kids to concentrate on the important task of moving about the asteroid rather than on the unimportant skill of picking fuel cells. The navigation was also simplified from using the analog joystick to move about the asteroid, to using the three buttons to perform turn left / move forward / turn right which was easier for smaller hands to control. The representations of the astronaut in the Mission Control view were also enlarged to eight times their actual size to make the direction they were facing more obvious. Once these changes were made, the children were very effective in their use of the VR technology, and the application remained virtually unchanged for the rest of the studies.

The second pilot study involved four sets of two children in each of the worlds. After an introduction by the guides, the children split up, spending ten minutes on each task twice, e.g. one child would be mission control for ten minutes, then astronaut for ten, then mission control for ten, then astronaut for ten. The children were very engaged in the activity, and their sense of presence seemed high. Several said that they initially felt they would fall off the asteroid if they walked over the nearby horizon but once they walked over that horizon they became comfortable with moving over the surface. Most of the children were actively talking to each other. Unfortunately, the children seemed too engaged in the task, focusing on the goal of collecting the fuel cells or satellite parts and only conversing on that specific topic. The children seemed to treat the experience as a big enjoyable videogame, only interested in achieving the goal of the game.

We had thought that the mission control child would comment about the astronaut child being 'upside down' and that the two children would need to integrate mission control's directions of "go up" and "go down" with the astronaut's directions of "go left" and "go right." The kids did not talk about the other child being 'upside down', didn't use any of the available landmarks to aid in navigation, and most mission controllers eventually adopted a strategy of telling the astronaut to remain in place and turn in either direction until told to stop, then move forward. Several times mission control incorrectly told the astronaut to turn to the left, meaning right, but the astronaut incorrectly turned to the right anyway, briefly proving that two wrongs can make a right.

The mission control children rarely looked at the live video image of the astronaut in the CAVE. Instead they concentrated almost exclusively on the computer-generated image of the avatar moving around the sphere. As mentioned previously, the children were focused on their goal of collecting the ten objects and the computer-generated spherical view helped them achieve that goal, while the live view did not.

The guides only interfered when absolutely necessary, trying to keep the immersion as strong as possible. The bridging activities were done with both children together at the ImmersaDesk where we found that words such as 'horizon' were not in the children's vocabulary, making the bridging activities more difficult than expected.

The 2D paper-based pre- and post-tests originally consisted of 20 open-ended questions on spanning four knowledge components: the shape of the earth, the relativity of up and down, the concept of a horizon, and the Earth's surface's finite but unbounded nature. These were soon supplemented with a 3D sculpting component using Play-Doh to get a better idea of the child's model(s). It became apparent from these questions that there was limited learning reflected under either treatment. We decided to call a halt to this pilot study and to make some significant adjustments to our procedures, focusing for the time being exclusively on the Asteroid world.

Since the children seemed quite able to use the VR hardware and complete the task, we needed some way for them to focus less on the task, and more on the concepts we wanted them to learn. For the third pilot study we modified our approach in several ways.

Instead of a short training time with the guides focusing exclusively on the VR hardware, this initial time was now also used to point out features of the landscape. The guides now gave each child an individual five minute introduction to the astronaut view and a similar introduction to the mission control view, and then introduced the collaborative task. The guide showed that if you kept going in the same direction that you would return where you started, and pointed out objects appearing over the horizon. Since the introductory time was increased, the children spent only ten minutes in each role, rather than two sets of ten as in the previous study. While walking the children between the CAVE and the ImmersaDesk, the guides would try to reinforce the concepts brought up during the training session. The guides also tried to direct the attention of mission control to the video window when appropriate 'rightside up in the video' and 'upside down on the sphere' situations appeared.

Most importantly, the bridging activities were changed from a group debriefing in front of the ImmersaDesk to individual guided inquiry using a physical globe of the Earth and a Styrofoam model of the asteroid. This focused on reminding the subjects of what they had experienced, how their experiences demonstrated the target knowledge components, and how that same knowledge applied to the Earth. While the ImmersaDesk allowed mission control to see the astronaut moving over the surface of the sphere, the physical models allowed more direct manipulation and interaction between the instructor and the student with the model. The instructor could now position a small astronaut figure at any point the sphere and manipulate the orientation of the sphere.

The data from the third pilot study is focusing on a detailed analysis of individual subjects' protocols. The analysis is complicated by the fact that there are few instances of complete, fundamental changes in conceptual models among subjects. Instead, we see some subjects holding strongly to their initial models, some who appear to demonstrate temporary effective learning during the experiment which is not reflected the next day, and some subjects who appear to reflect persistent learning of some, but rarely all, of the target knowledge components.

One obvious outcome of the study has been the sensitivity of the subjects' responses depending on the dimensionality of the media. The children showed little consistency between their 2D and 3D models, often appearing to maintain simultaneous 'separate but equal' representations. Children who demonstrated effective learning when asked to interact with 3D physical models would often revert to 'flat Earth' models when asked to reason on the basis of 2D drawings.

Simulator sickness did not appear to be a significant problem during the pilot studies. In all of the studies, there was only one child reported dizziness during the study, and that child refused to leave the CAVE. Another child expressed concern over sickness prior to the experience, but reported no difficulties during the experiment.

7 Conclusions and Future Work

These three pilot studies with 34 children have shown us that the children were able to use the VR equipment effectively. They were strongly engaged by the nominal tasks, sometimes to the detriment of the target learning, treating the experience like a video game to be 'won' rather than a possible source of learning. The children actively communicated with each other, though again on very task-specific topics.

The effectiveness of the two pedagogical strategies, and the value of distributed virtual reality technologies in support of these strategies, remain open questions. We were encouraged that there were clear instances of learning related to specific knowledge components of the target concept, but it is clear that the methods employed to date have not demonstrated clear, dramatic proof of effectiveness. Part of the problem may have been the highly challenging (albeit important) subject pool which we selected; many of our subjects clearly reflected learning disabilities and attention deficit disorders. Nonetheless, we believe that our experience has given us direction for improving our environments, protocols, and analysis strategies.

We also must mention the difficulty of conducting these studies with young children outside of the school setting. In spite of substantial cooperation on the part of the school and the children, the difficulties in obtaining parental permissions, unanticipated absences, scheduling pullouts, and especially arranging transportation (liability concerns requiring us to employ unreliable and expensive taxi services) combined to make logistic support extremely time-consuming and expensive. The process of running two pairs of children through the experiment typically required six adults for most of a working day.

Because of these constraints, the first set of actual studies comparing the Asteroid world and Earth world treatments is being undertaken inside a local elementary school. Instead of bringing the kids to the VR equipment in our lab, an ImmersaDesk and stereo monitor along with the associated computers and recording equipment were brought into the elementary school for a 2-week period. As it was unrealistic to set up a CAVE at the school, the ImmersaDesk replaced the CAVE as the astronaut's view while the stereo monitor replaced the ImmersaDesk as the mission control view. This setup inside the school allowed us to test more children within a much shorter period of time with a smaller crew. We are currently analyzing the data from this study.

Acknowledgements

There are many people involved in this research project. In addition to the authors, the Round Earth team includes Joe Alexander, Tom DeFanti, Josh Hemmerich, Jyoti Jain, Mark Orr, Carlos Orrego, Maria Roussos, and Mike Trolio. We also acknowledge the assistance of Josephine Anstey, Jim Costigan, Tom Frisch, Jason Leigh, Dave Pape, Sam Throngrong, and Fang Wang.

We especially wish to thank Shirley Woodard, Program Director, and Anthony Biegler, Principal of South Loop Elementary School, for their Herculean efforts in encouraging and coordinating student participation in the study. To the wonderful children of South Loop Elementary, we owe our deepest gratitude.

This research was supported by funding from the National Science Foundation, award EIA 9720351 - Deep Learning and Visualization Technologies. The virtual reality research, collaborations, and outreach programs at the Electronic Visualization Laboratory are made possible through major funding from the National Science Foundation and the US Department of Energy; specifically NSF awards CDA-9512272, ANI-9712283, CDA-9720351, EIA-9802090, EIA-9871058, and the NSF Partnerships for Advanced Computational Infrastructure program.

References

[1] American Association for the Advancement of Science (AAAS). Science for all Americans: A Project 2061 Report on Literary Goals in Science, Mathematics, and Technology. Technical report, AAAS Publication, Washington D.C., 1992.

[2] P. Brna and R. Aspin. Collaboration in a Virtual World: Support for Conceptual Learning . In *Proceedings of 1st International Conference on Virtual Reality in Education and Training*, Loughborough, UK, June 24-26 1997.

[3] C. Dede, M. Salzman, and R. B. Loftin. ScienceSpace: Virtual Realities for Learning Complex and Abstract Scientific Concepts. In *Proceedings of IEEE Virtual Reality Annual International Symposium '96*, pages 246–253, 1996.

[4] A. Johnson, M. Roussos, J. Leigh, C. Barnes, C. Vasilakis, and T. Moher. The NICE Project: Learning Together in a Virtual World. In *Proceedings of IEEE Virtual Reality Annual International Symposium '98*, pages 176–183, Atlanta, Georgia, Mar 1998.

[5] D. Johnson and R. Johnson. *Learning Together and Alone: Cooperative, Competitive, and Individualistic Learning.* Prentice Hall, 1987.

[6] National Council of Teachers of Mathematics (NCTM). *Curriculum and Evaluation Standards for School Mathematics.* 1989.

[7] J. Nussbaum. *The Earth as a Cosmic Body*, pages 170–192. Children's Ideas in Science. Milton Keynes, UK: Open University Press, 1985.

[8] S. Ohlsson. Abstract Schemas. *Educational Psychologist*, 28(1):51–66, 1993.

[9] S. Ohlsson. *Learning to do and learning to understand: a lesson and a challenge for cognitive modeling.* Learning in Humans and Machines: Towards an interdisciplinary learning science. Oxford, UK: Elsevier, 1995.

[10] M. Roussos, A. Johnson, T. Moher, J. Leigh, C. Barnes, and C. Vasilakis. Learning and Building Together in an Immersive Virtual World. *Presence*, to appear, 1999.

[11] S. Vosniadou. Capturing and Modeling the Process of Conceptual Change. *Learning and Instruction*, 4:45–69, 1994.

[12] S. Vosniadou and W. F. Brewer. Mental Models of the Day / Night Cycle. *Cognitive Science*, 18:123–183, 1994.

[13] C. Youngblut. Educational Uses of Virtual Reality Technology. Technical Report IDA Document D-2128, Institute for Defense Analyses, Alexandria, VA, 1998.