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

Andrew Johnson, Thomas Moher, Mark Gillingham, Stellan Ohlsson
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 Round Earth Project is a collaboration between Computer Science, Education and Psychology to study the use of collaborative virtual reality technology in 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, which influence how both direct experience and discourse within that domain are conceptualized. These are deep ideas.

The acquisition of deep ideas, deep learning, involves a paradox. While the intent is to replace the learner's existing ideas, those prior ideas are the learner's main tools for understanding the new idea. The typical outcome of trying to replace these ideas with deeper or more fundamental ideas is distortion; the novel idea is misunderstood in the process of being assimilated to prior knowledge.

A simple and compelling example of this children being taught that the Earth is spherical. Empirical studies have demonstrated that when children are told that the Earth is round they often react to the novel information 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.

Deep learning may require alternative cognitive starting points. Virtual Reality technology can give the learner such a starting point. In the Round Earth Project the children explore a small asteroid. This allows them To 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. The virtual experience of actually 'visiting' a visibly spherical world generates a mental model that depicts the learner as located on such a world. Once established, this mental model can grow and subsume other experiences. This alternative model can eventually subsume the everyday experience of walking on a seemingly flat Earth.

In section 2 we will briefly describe current work in collaborative immersive virtual reality 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. 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 Education

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 [11]. There are currently only two multi-user virtual educational worlds: Virtual Physics at the University of Lancaster, and NICE (Narrative, Immersive, Constructionist / Collaborative Environments) at the University of Illinois at Chicago. Other current educational VR worlds such as the ScienceSpace [2] worlds are being extended to support multiple users.

NICE [3, 8], 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, 4]

2. The learning goal must be hard.

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

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

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 [5, 9, 10].

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 achievement 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. Vosniadou, Brewer and Nussbaum have investigated this problem so there is existing literature.

3. Immersive VR is well suited to create a sense of presence of walking on a spherical object with small diameter.

4. collaborative etc. from the aborted KDI proposal

3 Deep Learning

Underneath the extensive systems of domain-specific knowledge that a person brings to bear on problems and situations, there are oranizing 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 [6, 7]

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 transformational and selectionist explanations of cognitive change. A '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 'selectionist' 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 that is established outside the learner's existing system of domain knowledge. Initially such an alternative representation might be rudimentary and hence be 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 selectionist 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 allows us 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 to assimilate new information is that each aspect of 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

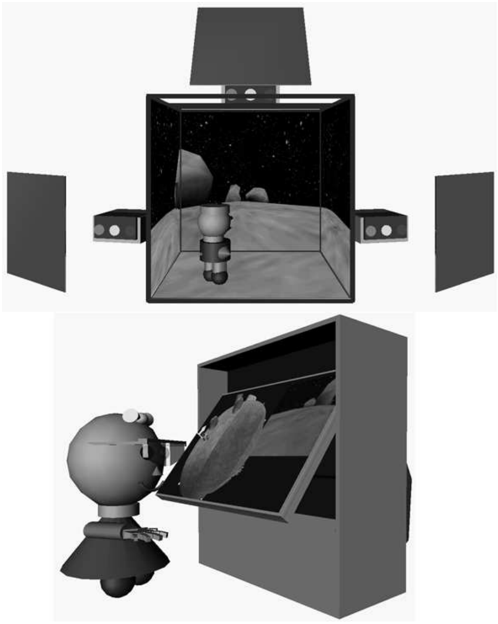

Figure 1: The two views of the asteroid world. The top image shows the astronaut exploring the surface of the asteroid within the CAVE. The bottom image shows the mission control view of the asteroid on an Immersadesk

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 describe the two worlds that we created.

4.1 Goals for the two Experiences

In order to compare the 'selectionist' to the 'transformationalist' approach we constructed two virtual worlds. For the 'selectionist' approach we wanted to start the children off at an alternative cognitive starting point, a small asteroid, where they could learn about walking around a spherical body different than the Earth. For the 'transformationalist approach' we wanted to start the children 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 'rightside up' on the surface. We wanted the collaborative task to foster positive interdependence where neither child could perform the task alone; they had to cooperate and communicate with each other. Through this communication 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 get tired.

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.' Each child gets to perform both roles during the experience. 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. 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.

In both worlds the controls are the same. The generic controller for the CAVE and the Immersadesk is the 'wand' - a six degree of freedom mouse with three buttons and a joystick. In the CAVE with the astronaut, the left button turns the child to the left, the middle button moves them forward, and the right button turns them to the right. In order to pick up an object, 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 (either the asteroid or the Earth.) The world can be turned completely around horizontally but there is limited tilt of the world. This allows mission control to keep the astronaut in view at all times, but allows the astronaut to be 'sideways' near the equator and 'upside down' in the Southern Hemisphere.

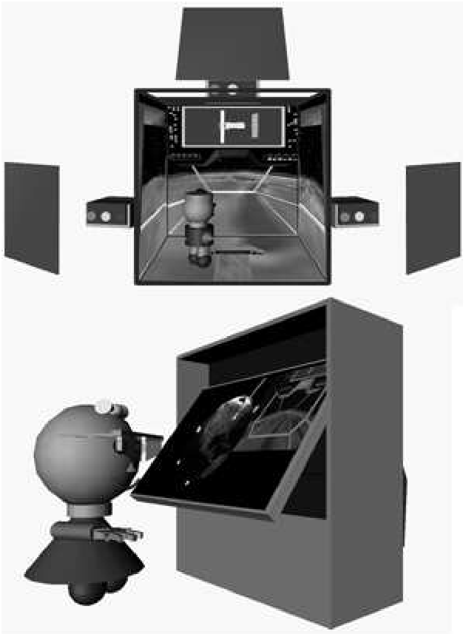

Figure 2: The two views of the Earth world. The top image shows the astronaut orbiting the Earth within the CAVE. The bottom image shows the mission control view of the Earth on an Immersadesk

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

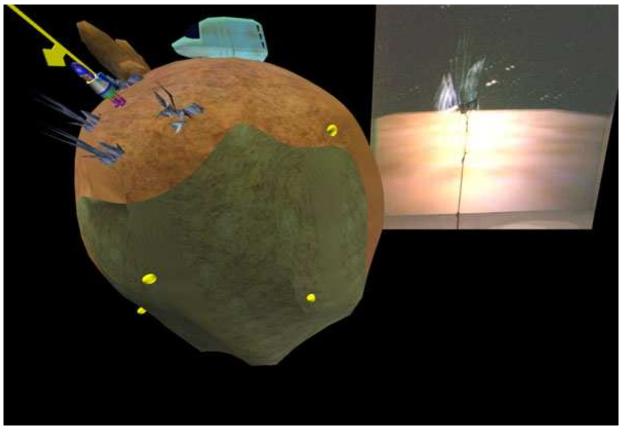

Figure 3: The Mission Control screen for the asteroid world. Mission control simultaneously sees the avatar of the astronaut moving around the spherical asteroid and the view from the astronaut's perspective of the slightly rounded surface of the asteroid.

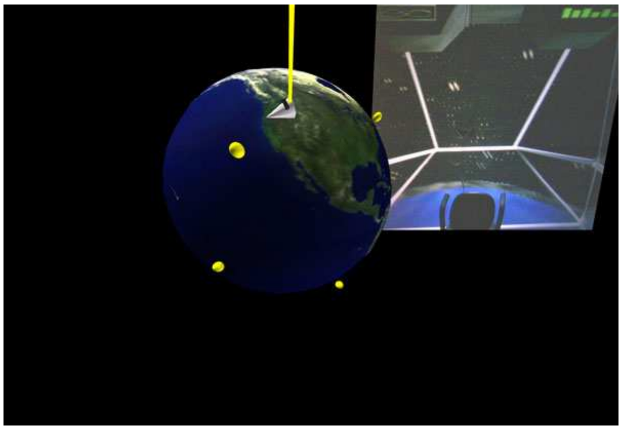

Figure 4: The Mission Control screen for the Earth world. Mission control simultaneously sees the avatar of the astronaut's spaceship moving around the spherical Earth and the view from the astronaut's perspective of the slightly rounded surface of the Earth.

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 what both of the participants see. 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 Chicago. Since our students are from downtown 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, and the city, and the Earth fall away as they rise 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 as a 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 both of the participants see. Figure 4 gives a closer look at what mission control sees. This view shows the spherical Earth with the avatar of the astronaut's spaceship orbiting above the surface as well as the view from the CAVE out the front window of the orbiting ship.

5 Experimental setup

In this section we will describe the experimental setup.

The CAVE is a standard 10' cube, three walls and one floor CAVE. Its graphics are driven by a rack onyx with one Infinite Reality Engine for each pair of walls. The screen resolution is 1024 x 768 with an average frame rate of Z 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 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.

The Immersadesk is a standard Immersadesk with one 6' by 4' screen. Its graphics are driven by a rack onyx with one Reality Engine 2 for the screen. The screen resolution is 1024 x 768 with an average frame rate of Z 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 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 headset microphone. Audio from the application and from the astronaut are mixed and sent through the Immersadesk's two speakers. A low light CCD camera mounted behind the child sends the image of mission to the VCR for recording.

The pre and post tests involve both aural, 2d drawing, and 3D sculpting questions to try and elicit the model, or in most cases models, the child is using. These questions are similar to those used by Vosniadou. ¡time for the questions¿

To help the children, there was a guide at the CAVE and at the 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 also changed as the pilot studies continued. These changes other issues raised by the pilot studies will be described in the following section.

6 User Studies

At this point in the Round Earth project, we have done three pilot studies. The first study consisted of four pairs of children to look primarily at interface and usability issues as well as refining our pre and post tests. Once these were improved, the second study consisted of eight pairs of children concentrating on the learning in the space, the effectiveness of the bridging activities, and the ability of the pre and post testing to see any change in the children. This led to changes in our procedures so a third pilot study of ? pairs of children was undertaken to evaluate our modified design.

In pilot study one the children chosen were children of the teachers and administrators at the school, allowing us to familiarize their parents with our procedures.

The students we studied came from a small urban Chicago public school. The total K-8 enrollment is about 550. The school is 98.9

¡can we classify the mental models from the pretest¿

The first pilot study showed us several things. As with the NICE user studies, the Stereographics glasses are too big for small children. We still have not found a satisfactory solution to this problem aside from tying the glasses on. In the initial design the astronaut on the asteroid had to reach down and 'touch' a fuel cell to pick it up, while the astronaut orbiting the Earth only had to get close enough for automatic grabbing to occur. The Earth astronauts had a much easier time, so that technique was adopted by the asteroid world as well. The navigation was also simplified from using the analog joystick to move about the asteroid, to using the turn left / move forward / turn right buttons. The representation of the astronaut avatar on mission control's Immersadesk was also enlarged and the direction of travel was made more obvious to make it easier for mission control to direct the astronaut.

The second pilot study involved four sets of two children in the asteroid world and the same number in the Earth world. After an initial introduction by the guides, the children split up, spending ten minutes on each task twice, e.g. mission control for 10, then astronaut for 10, then mission control for 10, then astronaut for 10. The children were engaged in the activity and they were 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. We had hoped the children would comment about the other person being 'upside down' or having 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 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 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, and we found that worlds such as 'horizon' were not in their vocabulary, making the bridging more difficult than expected.

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 second pilot study we changed several things. Instead of a very short training time with the guides focusing on the VR hardware, this initial time was now being used to point out features while teaching about the VR hardware. The guide now showed that if you kept going in the same direction that you would return where you started. The guided also pointed out objects appearing over the horizon. Each child now received a five minute introduction to the astronaut view and a similar introduction to the mission control view with the guide and after this time the task was introduced. The children then spent 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 now try to reinforce the concepts they brought up during the training session. The bridging activities were now done with a globe of the Earth and a styrofoam model of the asteroid.

¡something about how we activated the 'video game' parts of their brain?¿

7 Conclusions and Future Work

Acknowledgements

There are many people involved in this research project. The complete Round Earth team includes: Joe Alexander, Tom DeFanti, Mark Gillingham, Josh Hemmerich, Jyoti Jain, Andy Johnson, Tom Moher, Stellan Ohlsson, Mark Orr, Carlos Orrego and Maria Roussos. We also gratefully acknowledge the assistance of Julieta Aguilera, Josephine Anstey, Jim Costigan, Tom Frisch, Jason Leigh, Dave Pape, Sam Throngrong, and Fang Wang.

The virtual reality research, collaborations, and outreach programs at EVL are made possible through major funding from the National Science Foundation, the Defense Advanced Research Projects Agency, and the US Department of Energy; specifically NSF awards CDA-9303433, CDA-9512272, NCR-9712283, CDA-9720351, and the NSF ASC Partnerships for Advanced Computational Infrastructure program.

This research was supported by funding from the National Science Foundation, award EIA 9720351 - Deep Learning and Visualization Technologies.

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] C. Dede, M. Salzman, and B. R. Loftin. MaxwellWorld: Learning Complex Scientific Concepts Via Immersion in Virtual Reality. In *Proceedings of Second International Conference of the Learning Sciences*, pages 22–29, 1996.

[3] 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 VRAIS '98*, pages 176–183, Atlanta, Georgia, Mar 1998.

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

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

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

[7] 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.

[8] 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, 1998.

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

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

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