Issues in the Design and Evaluation
of a Virtual Reality Learning Environment

by
Maria Roussos
B.S., Athens University of Economics, Greece, 1992
M.F.A., University of Illinois at Chicago, Chicago, IL, 1997

Submitted as partial fulfillment of the requirements
for the Master of Science in Electrical Engineering and Computer Science
in the Graduate College of the
University of Illinois at Chicago, May 1997

Committee: Thomas G. Moher (EECS), Thomas A. DeFanti (EECS), Mark G. Gillingham (Education)

Acknowledgments

This thesis is based on a complete project, from designing to developing to evaluating a virtual reality learning environment. The amount of work involved could not have been carried out without the support, in many different levels, of many individuals.

I would like to express my gratitude to my thesis committee: Tom Moher for being such an inspirational advisor, Tom DeFanti for providing the support and resources which enabled me to pursue this research, and Mark Gillingham for the valuable educational direction.

There are no words to thank the NICE team for a wonderful collaborative experience: Andy Johnson, Jason Leigh, Chris Vasilakis, and Craig Barnes. Their creative ideas, talent, and help has been more than I can possibly thank for.

I would also like to thank the many other people who helped me while I worked on this thesis. In particular:
Jim Costigan, for his tremendous help, willingness, and expertise in tackling all the media; Vicki Pierce, for being a fantastic partner in evaluation; my friends at the Open University in England, Martin Oliver and Cathy Lewin, for their help in forming the evaluation framework; Karl Steiner, whose work has influenced the development of this thesis; the members of the Electronic Visualization Laboratory and the Interactive Computing Environments laboratory for many insights, instances of patience, and help with the children.

Warm and special thanks go to the teachers, Laurie Pierce and Barb Montcalm from St. Gertrude's Elementary, the parents and, above all, the 52 children who participated in the studies, for investing their time and energy.

Last but not least, I would like to thank Alan Millman, for being a wonderful companion, and my most loved ones: my sister Katie, my parents, Chris and Dena, and Fanis.

\initials{MR}

SUMMARY

This thesis describes the design, development, and evaluation of the NICE (Narrative Immersive Constructionist/Collaborative Environments) project, a virtual learning environment implemented for the CAVE. Two central questions are explored in this work. First, what are the issues that must be addressed in designing immersive virtual environments for education, and second, can virtual reality environments be effective tools for learning?

The question of the design of open-ended, exploratory learning environments, such as NICE, is addressed through the analysis of current educational theories. These theories are presented and related to the design of the virtual learning environment.

The issues of the evaluation of NICE are explored through observational studies conducted with a total of 52 children between the ages of 5 and 11. These studies examine the environment from a technical, orientational, affective, cognitive, and pedagogical perspective. Finally, the thesis presents direction for future studies in these domains.

1.1. Statement of the Problem

In designing any kind of learning environment, there are a number of issues one has to consider. In designing a learning environment based on new and relatively unexplored tools, such as virtual reality (VR), these issues become even more complex. Two central questions have motivated this work. First, what are the issues that must be addressed in designing immersive virtual environments for education? And second, can virtual reality be an effective learning tool?

The research described in this thesis has focused on these questions while in the process of creating a virtual reality learning environment developed mainly for the CAVE(tm)\ref{CruzNeira93}. This focus has informed the design, development, and evaluation of the Narrative Immersive Constructionist/Collaborative Environments (NICE) project. The NICE project is a testbed application for exploring ideas in the direction of interactive, open-ended, informal learning environments for children. In NICE children explore, cultivate, and preserve a healthy virtual garden. In order to achieve this, they select seeds, plant vegetables, pull weeds, and use symbolic representations to control environmental factors, such as a raincloud to provide rain, a sun to give sunlight, or a compost heap to recycle weeds. To facilitate their understanding of the underlying principles of this micro-ecosystem, they can change their size or observe the root system. In addition to their activities with the garden, multiple children can collaborate synchronously within the virtual environment, through avatar representations and audio, or asynchronously by observing the visual narrative generated by the system and the WWW page.

The research objectives of this work have been threefold:

First, to contribute to the design of an original architecture for a virtual learning environment, which combines many different and complex educational components. NICE can serve this direction as a prototype for future development of VR applications for education. This model of virtual learning environments is informed by current educational theories, such as constructivism, collaborative learning, storytelling, and aware of human-computer interaction metaphors that work best. Metaphors that are applied successfully with children, may easily be extended to other areas of study and training in VR.

Second, to approach the issue of evaluation with the development of a theoretical framework. The purpose of this theoretical framework is to impose a structure for the future evaluation of virtual learning environments. As virtual worlds usually provide complex, multi-sensory experiences, evaluation should take into account multiple factors, from the technical to the cognitive. The question is how should one proceed with the evaluation of these kinds of environments? This framework attempts to address this question in a way that is generalizable and easily applied to the assessment of any immersive virtual environment for learning.

And third, to bring up the issues of evaluation in VR by performing studies based on the theoretical framework. We need empirical evidence on how the virtual environment affects learning. This thesis work is not so much about evaluating a specific design (the NICE project), as it is about paying attention to the issues involved in the evaluation of not so highly structured and guided learning environments in VR. It is hoped that these observations may provide a focus from which future studies in VR and education can benefit.

1.2. Structure of the Thesis

This thesis is divided into four primary chapters: the Background and Related Work chapter, the chapter on the Design and Implementation of the NICE project, the Evaluation chapter, and Conclusions.

The Background and Related Work chapter (Chapter 2) focuses on three general areas related to the research of this thesis. The first section discusses learning theories that are influencing the design of current educational environments, and their relationship to computer-based learning. The second section outlines the current state of virtual reality in education, by reviewing text-based, desktop, and immersive VR environments for learning, as well as highlighting the value of VR for education. The third section, briefly mentions the different assessment approaches taken in educational evaluation, and discusses whether VR requires new and different assessment techniques.

The next chapter (Chapter 3) describes the design and implementation of the NICE project, the specific virtual learning environment on which this thesis is based on. The three sections discuss the objectives, the ongoing design process of the project, and finally the architecture which ties together all the components of the system.

The chapter on Evaluation (Chapter 4) is divided into four main sections. The theoretical framework section proposes a structure for the evaluation of immersive VR learning environments. The evaluation methodology section outlines the methods and setting for the experiments. This is followed by the next two sections, which describe the initial and main studies conducted with children and conclude with an analysis of the observed results.

The final chapter of the thesis (Chapter 5) covers the conclusions drawn from this work and describes possible future additions.

2.1. Learning Theories

In the following sections learning theories which are influencing the design of current educational environments will be reviewed. The purpose of this chapter is to provide an overview of what is being stressed in contemporary educational literature, so it can then be related to the trends in the development of recent computer projects for learning.

2.1.1. Constructivism and Constructionism

2.1.2. Sociocultural Theories

2.1.3. Cooperative Learning

2.1.4. Motivation

2.1.5. Narrative

2.1.6. Summary Of Learning Theories

2.1. Learning Theories

2.1.1. Constructivism and Constructionism

The traditional approach to education reflects the behaviorist and programmed instruction model, where students learn by accumulating pieces of information transmitted by the teacher. This approach to learning still dominates curriculum and instruction, and is supported by the standardized tests. It is recognized in the field of education as instructional design \cite{Reigeluth83}. Collins (1996) notes that the goal of instructional delivery is to design an educational system that transmits content and skills in a clear, well structured, and efficient manner \cite {Collins96}. Similarly, Dick defines instructional design as an educational intervention ``driven by specific outcome objectives, materials, or procedures that are targeted on these objectives, and assessments that determine if the desired changes in behaviors (learning) have occurred" \cite {Duffy92}. A few cognitive scientists call for a ``mental model," where understanding is seen as an acquisition of a knowledge base as in an expert system model \cite {Duffy92}. The goal of the instruction is to help the learner acquire structure. The goal of evaluation is to examine if mastery has been achieved and that everyone has acquired the same information. Evaluation is thus performed with tests that stand separate from the instruction.

An alternative epistemological basis to the objectivist tradition is constructivism \cite {Duffy92}. The constructivist view argues that the goal of education is to help students construct their own understandings \cite {Collins96}. Constructivism is concerned with learner creation of meaning and linking of new ideas to existing knowledge, and therefore involves a large degree of student autonomy and initiative. The emphasis is on facilitative environments, rather than instructional goals, where the teacher takes on the role of mentor, or facilitator.

Piaget, Dewey, and Vygotsky have influenced the development of constructivist theory. Piaget's constructivism is rooted in discovery, play, and imagination as fundamental activities for the development of the child's learning \cite{Piaget73,Labinowicz80}. Piaget, famous for his research on the psychological development of the child, believed that children have to go through stages in which they accept ideas they may later see as wrong. Understanding is gradually built up step by step through active involvement. Piaget called for teachers to understand these steps in the development of the child's mind. The teacher's role becomes one of guiding mentor stimulating initiative, play, experimentation, reasoning, and social collaboration \cite {DeVries87}. Dewey argued that education depended on action, that children must actively construct knowledge by drawing it out of experiences that have meaning and importance to them \cite {Dewey66}.

Although not considered necessarily a constructivist, Vygotsky's theory combines both instructional design and constructivism. Effective instruction, according to Vygotsky, takes place in the ``zone of proximal development," which he defines as the discrepancy between the child's actual development as determined by independent problem solving and the higher level of ``potential development as determined through problem solving under adult guidance or in collaboration with more capable peers" \cite{Vygotsky78}. His views of social negotiation, the social nature of meaning supported by collaborative construction of knowledge, will be discussed in the next section.

Seymour Papert calls for further distinction of the constructivist views, by focusing on the involvement of the student in the actual designing, constructing and erecting of ``external" products or artifacts \cite {Papert80,Papert93,HarelPapert}. The idea behind using raw data, primary sources, physical, and interactive materials in real-world possibilities is to help learners generate the abstractions that bind phenomena together. Researchers at MIT use the word ``constructionism" to describe the knowledge construction process that arises from the physical creation of objects \cite{Harel91}.

On the negative side, there are extreme implementations of constructivism, which have provoked reactions against it. Constructivism is considered by the behaviorists ``a label for fuzzy, unscientific thinking'', as Cunningham notes in his dialogue assessing the two teaching approaches \cite{Duffy92}. The breadth of applicability of constructivism led some to believe that it supports spontaneous, uncontrolled learning, in contrast to the systematic, organized instruction of knowledge employed by the objectivist tradition. The open-endedness of constructivist problems can be daunting for the entry-level learner. Similarly, it might be difficult for the teacher to incorporate constructivism into her teaching methods. Required course content and externally applied assessments are realities that teachers must accommodate. Changing to constructivist instruction will mean finding the appropriate balance between the existing instructional methods and this new educational practice. Finally, the efficiency and reliability of evaluation methods are questioned, as constructivist learning environments are difficult to evaluate.

Constructivism has emerged in the last decade as an alternate pedagogy closely related to the advances in educational technology. Interest in constructivism has blossomed considerably while conventional instruction and assessment techniques have been criticized for their inflexibility. There is a turn to more flexible, open-ended, adaptive, and multi-dimensional instructional techniques as well as more qualitative, observation-based methods of evaluation. As a result, constructivism is embraced by many educational technologists and this is reflected in the plethora of multimedia and computer-based software that draw from the constructivist premises. As such, it makes an ideal basis for building a theory of learning for open, informal, and virtual learning environments.

2.1. Learning Theories

Sociocultural Theories

Various sociocultural theories have emerged in recent years, expanding upon elements taken from the aforementioned learning paradigm of constructivism. Learning, as derived from the sociocultural perspective, is contextualized, situated, and culturally-based \cite {Soloway96}, recognizing the importance of social interaction. Some of the paradigms developed through this approach include situated cognition, cognitive apprenticeship, and activity theory.

The theory of situated cognition encourages learning and understanding through continued situated use in the context of an authentic situation. Brown, Collins, and Duguid (1989) assert that learning is most effective in authentic, or real world, contexts with problems that allow students to generate their own solution paths. The goal is to prepare students to do the kinds of complex tasks that occur in life \cite{Brown89}. As a result of bringing students into more authentic learning environments, it is hoped that students will evolve from novices to experts. Using technology to build communities of practice and apprenticeship is popular in many recent approaches to project-based science teaching. Other similar approaches include ``learning by doing'', ``guided practice'' \cite{Rogoff90}, and The Cognition \& Technology group's ``anchored instruction''.

The sociocultural perspective emphasizes on helping learners socially construct knowledge through the use of language. Influential in this direction are Vygotsky's views of cognitive development, which focus on discussion and dialogue, as important pedagogical strategies for constructivist design. As mentioned in the previous section, Vygotsky stressed that children use concepts taught to them by adults, and link these concepts to their own ideas and generalization \cite{Vygotsky78}. Important, however, in Vygotsky's thesis is the social context in learning. His emphasis on the significance of the learning culture and the nature of the society within which teaching takes place has been used as the foundation for social constructivism.

The questions that need to be addressed in these perspectives involve the definitions of terms such as ``authentic'' or ``learning by doing''. What constitutes an authentic learning environment and how can one be provided in the context of a classroom? Learning by doing and other theories based on activity and practice are significant, but in many cases can foster automatization and thoughtlessness. These concerns are important to consider when developing synthetic learning environments which seek to provide authentic, project-based contexts that can transfer to real world problems.

2.1. Learning Theories

Cooperative Learning

One of the most important components of educational research over the past two decades has been the emergence of cooperative learning as a dominant learning paradigm. This paradigm emphasizes learning as an active process best facilitated by interaction with peers, teachers, and other learning resources \cite{Slavin91}. The focus of a cooperative setting is on group directed, rather than individually directed effort, which presumes positive interdependencies among the members of the group. Learners work together to discuss learning problems and share responsibility for the learning of the group, as well as for each individual member.

There is general consensus that individuals working in groups produce higher achievement than those working alone, due to the increase in cognitive resources. Other advantages of collaboration in education include co-construction, social facilitation, verbalization, conflict, observation, increased motivation and self-esteem \cite {Slavin91,Issroff96}.

A collaborative setting may, however, also include negative interdependencies amongst its members. Cooperation in this case may mean competition, where one student's gain is the other's loss, or be differently balanced for each member of a group: some learners do not learn well in collaborative settings, while others show increased productivity \cite{Johnson85}. Collaborative learning doesn't ensure that everyone will achieve the same level or type of mastery. Simply placing learners together does not mean that they will collaborate. As shown through classroom practice, collaborative learning often degenerates into working in groups with highly non-uniform participation and effort by group members.

As long as students need help in developing skills for cooperative work, the development of collaborative learning environments requires a framework which emphasizes structure and clearly defined goals, so to encourage individual responsibility in such a way that students are concerned about the performance of the whole group, not just their own \cite{Johnson85}. Computers certainly afford opportunities for improved control, explicit support for differentiated roles, mechanisms for planning group activities and providing feedback on groups' progress towards a goal. A number of projects in the field of Computer Supported Cooperative Learning (CSCL) study the use of computers in augmenting and supporting group work \cite {Johnson86}. A central theme in many of these studies is the notion of verbal interaction between pairs or groups of computer users \cite {Steiner95,Steiner94,Steiner94a}.

The idea of computer-supported cooperative learning is recently being extended through the emergence of telepresence. The integration of virtual reality and high-speed networking \cite{Leigh96a} expands physical boundaries and opens up new possibilities for collaboration beyond the classroom. Through the use of avatars, or virtual personas, geographically separated learners perceive simultaneous presence in a virtual environment \cite{Dede95}. The ability to connect with learners at distant locations, enhanced by visual, gestural, and verbal interaction is key to the development of unique educational experiences. Some of the advantages here include the ability for referential communication (to ``put myself in your shoes"), distributed cognition amongst learners, as well as issues of verbal interaction and cultural awareness.

2.1. Learning Theories

2.1.4. Motivation

The concept of motivation plays a significant role in educational activity, and is gaining acceptance in current educational literature. Anderson et al. (1975) define motivation as the ``arousal, direction, and continuance of behavior." In other words, motivation is concerned with behavior regulation: what drives action, how action is directed, and what action is under voluntary control \cite{Chan96}.

In contrast to cognitive and psychomotor factors, which have always been the focus of learning research, the affective parameter is often overlooked in educational practice and evaluation. Motivation, however, certainly relates to learning and should be a very important characteristic of learning environments designed for students. It concerns whether a student is willing to learn, not whether the student is able to learn \cite{Chan96}. The learning environment should be designed to foster intrinsic, self-initiated interest as well as continuing engagement. Interest and knowledge come hand in hand, but the correlation is not always a straight line. Over a period of time, if learners have increased interest, their knowledge might also increase. It is not always the case, however, that students' interest will remain high if they continue acquiring knowledge. Of particular significance, in this case, is the study of continuing motivation: the student returning to the use of a system without external constraints to do so.

In the last few years there has been a considerable shift towards providing motivating learning settings and the building of learning environments that nurture motivation. The entrance of the computer onto the learning scene and its impact on young users has resulted in the research of motivation in relation to computers, especially computer and video games \cite{Provenzo91,Malone80,Kafai95}. The computer is considered by many as an important motivator for learning, and this affective quality has lead many to believe that computers are more advantageous to educational activity than other, more traditional, pedagogical methods. Malone (1980) describes three primary characteristics of computer games that directly contribute to motivation: providing a challenge (a goal), creating fantasy (an environment or situation not otherwise available), and provoking curiosity (to make the strange familiar). Malone and Lepper (1987) expand this framework to add confidence and control as aspects of motivation. Related to confidence, Papert talks about empowerment. The role of the computer, he notes, has less to do with information and more with giving children a greater sense of empowerment \cite{Papert80}.

As important as motivation is, it is not an end unto itself. For example, many popular computer games containing violence are engaging and motivating for a large population of children \cite{Provenzo91}, but they are not usually educational. Issroff and delSoldato \cite{Issroff96} make a distinction between motivation to use a tool and motivation to learn a subject. The study of affective parameters, such as motivation, should strive to gain a deeper understanding of what is engaging in learning without compromising educational value.

2.1. Learning Theories

2.1.5. Narrative

Storytelling is a familiar and ubiquitous part of everyday life; a form of communication and shared experience since antiquity. Stories may serve social, cognitive, and emotional functions. With respect to its social function, the communicative nature of storytelling is evident in the discourse which takes place in a learning context. In this sense, educational theories which support narrative are highly related to the collaborative and sociocultural learning theories described in previous sections. A model for teaching that draws on the power of storytelling and narrative is not only important for language development but can enhance the recall and natural understanding of a variety of topics \cite {Egan88}. In terms of its cognitive function, the structure and dramatic tension of a narrative creates expectation which is satisfied upon resolution of the story, and aids in planning, reconstructing, illustrating, and summarizing abstract concepts. Stories are an important component of childrens' playcrafting and scaffolding, as children develop understanding when advancing from simple relatively plotless stories to more complex stories \cite {Applebee78,Steiner94a}. In addition to serving as an effective representation for learning, narrative serves an affective function. Through stories, tales, or myths we can assume other perspectives, become more culturally aware, and learn about values, habits, and traditions. Storytelling is also a very motivating and engaging learning tool, whether it involves experiencing a pre-determined narrative script or constructing ones own story.

While many two-dimensional computer-supported storytelling environments have been developed for use by children \cite{Steiner92,Lewin95}, narrative has just begun appearing as a theme in virtual environments. These virtual worlds are almost exclusively developed by artists or linked to the area of Interactive Fiction \cite {Laurel93}. These systems explore narrative in the sense of space and time \cite {Laurel94}, plot and character development \cite {Lebowitz84}, and emotional interactive drama \cite {Bates92}, and do not directly address educational issues.

Characters and story structures are not unexplored concepts in interactive media. Virtual reality worlds, however, focus on the construction of objects and spaces, but not stories that tie them together. Incorporating story and characters requires the development of more ``intelligent'' computational models. Recent advances in the field of Artificial Intelligence include the development of agents, artificial creatures incorporating a set of human-like behaviors, as well as the exploration of plot and story structures which may emerge from the interaction between these agents \cite {Maes89}. Despite the interesting developments, these programmed agents have far to go before they can successfully simulate a perceptual, cognitive, or emotional level that may produce meaningful stories for learners.

2.1. Learning Theories

2.1.6. Summary of Learning Theories

In summary, most of the aforementioned theories reflect student-centered learning practices, recently emerging to counter behavioristic models of education environments. This shift from the highly guided to the more open learning curriculum is also appropriated in the development of the technology for learning, namely in the way new resources are formed to support new teaching and learning methods.

2.2. Virtual Reality in Education

There are different types of virtual worlds ranging from a simple virtual environment with text, no audio, and one observer, to environments which are more complex containing sophisticated computer graphics, audio, and interaction with objects and humans. This variety of virtual worlds represents the difference between ``presence'', ``immersion'', and ``interactivity''. Presence is the (mental) feeling of being in a virtual space, such as MOOs, virtual walk through's, or even books and films. Immersion is the complete visual and auditory submersion into the virtual world through VR systems such as the helmet-mounted display (HMD) or the CAVE(tm). Interactivity refers to how reactive the system is in response to the user's actions or how much power the user has to modify the environment.

Based on the above, the virtual reality environments which are designed specifically for education typically fall into three categories. The first includes networked text-based virtual environments, which are highly interactive but not immersive. The second involves desktop virtual reality simulations, where interactivity is usually limited but varies according to the control given by the program, and immersion also varies but is not easily provided. The third category includes the immersive VR environments, where immersion is high, but interactivity may be limited, depending on the complexity of the virtual world. The following sections will present an overview of the current state of educational VR projects which fall into these categories.

2.2.1. Networked Text-based Virtual Worlds

2.2.2. Desktop Virtual Reality

2.2.3. Immersive VR projects

2.2.4. The Value of Virtual Reality for Education

2.2. Virtual Reality in Education

2.2.1. Networked Text-based Virtual Worlds

Text-based virtual worlds are real-time multi-person virtual environments created through text descriptions rather than graphics. These environments are commonly known as MUDs, or Multi-User Dungeons, and MOOs (MUD Object Oriented), a MUD where users define spaces and objects. Developed from on-line game environments that allow several participants to play a computer game of Dungeons \& Dragons together, MOO's and MUD's support real-time interactive use among a large number of multiple users that may be sitting at remotely located computers but sharing the same virtual world.

These text-based virtual environments were primarily designed for entertainment, but an increasing number of MUDs and MOOs now exist for education. One of the earlier ones is ExploreNet, a networked virtual world where children interact with each other via avatars in settings filled with computer-controlled props. ExploreNet employs a story-like two-dimensional graphical interface to combine networked use of text and graphics \cite{Hughes96}. The Grassroots project is a popular text-based simulation of a neighborhood in a MOO combined with a WWW page. Students from around the world recreate their neighborhoods and share them with others \cite{Parsons96}. A more ``traditional'' MOO is Moose Crossing, an improved version of the earlier MediaMOO project, designed for kids to program objects and their behaviors \cite{Bruckman95}. A few other MOOs have been developed for teachers, such as Diversity University and Deadalus MOO, but most of them are more loosely defined and function as forums, or ``chat rooms'', where issues of self expression, identity, social interaction, and cultural awareness may be explored.

With respect to education, the advocates of textual VR cite two virtues of text-dependence: the creation of a community context, and the encouragement of reading and writing. The sense of community inspired by these environments allows the exploration of social codes, power and group dynamics. The fact that the world is built through text rather than graphics supports the second argument. Text-based VR has developed a strong connection to narrative and the area of Interactive Fiction (IF), as participants can assume various identities, develop situations cast in story-like or adventure-based formats, and interact with objects and characters.

On the other hand, the lack of a specific learning structure and direction, may leave many participants floundering in the virtual environments without specific goals. As with the Internet, making information available isn't enough to qualify the setting as an educational experience. Additionally, the lack of multi-sensory representation in the physical sense (visual, auditory, and tactile), may render MUDs incomplete learning environments in and of themselves. Moreover, students and teachers have to learn the necessary design and construction conventions to create specific text-based worlds and objects. The detailed mastery of these rules is in many cases difficult and irrelevant to learning.

2.2. Virtual Reality in Education

2.2.2. Desktop Virtual Reality

Classroom uses of virtual reality are limited, primarily due to the high cost of VR systems. For this reason, the term ``virtual reality'' has been applied more widely to include desktop virtual reality \cite{Winn93}. VR applications on the personal computer allow users to walk through simulated environments created via readily available commercial software, such as Virtus WalkThrough. Some slightly more expensive systems add peripheral devices, such as datagloves, joysticks, headmounts, or tracking, which plug directly into the desktop to provide a higher degree of interactivity.

Consequently, the VR projects developed using these low-end technologies are limited in size and complexity, and lack immersive and interactive qualities. Typical projects resemble 3D multimedia simulations and may involve the creation of simple models and minimalist worlds, which can be experienced via ``walkthroughs''. Certain efforts are worth noting, such as Grove's attempt to recreate historical contexts in VR \cite{Grove96}, but continue to remain rather passive simulations of limited detail. Perhaps the most significant use of desktop VR systems is in the education of learners with special needs or learning disabilities, who may not otherwise be able to experience certain aspects of the physical world. In this case, other types of VR systems may be unusable due to health, safety, and ethical reasons \cite {Cromby95}.

The integration of Internet and VRML software into the desktop configuration have expanded the usage of virtual reality in the classroom and added access to networked worlds, such as the ones described in the previous section. Research institutes such as The Virtual Reality and Education Laboratory at East Carolina University are also helping expand the use of VR at the level of the individual teachers by providing information about VR software for use in the classroom \cite {Pantelidis93}. Although educators may be willing to incorporate VR in their teaching practices, the introduction of VR in the classroom may bring with it multiple problems, including issues of funding, safety, teacher and student training, reshaping the curriculum, technological anxiety and confusion \cite {Bricken91}. Any efforts rely on the individual action of the teachers, who may need to convince reluctant administrators, colleagues, and parents of the value of VR in their classroom.

2.2. Virtual Reality in Education

2.2.3. Immersive Virtual Environments

In contrast to the virtual learning environments described above, immersive educational environments are being developed using high-end equipment. They are consequently limited to situations with special funding, such as academic and research environments. Most of the projects are developed especially for head-mounted display systems (HMD).

A large part of this educational research has been focused on science education, as in the ScienceSpace projects developed by Dede and his colleagues \cite {Dede96}, which set out to explore motion, electrostatic forces and other physics concepts. The initial formative evaluation reports on learners' engagement, surprise and understanding of the alternative representations of the concepts provided in the ScienceSpace worlds \cite {Dede96}.

The Human Interface Technology Laboratory (HITL) at the University of Washington has been one of the early educational seedbeds for VR, with projects such as the Virtual Reality Roving Vehicle (VRRV) \cite{Rose95,Winn93}, Water on Tap \cite{Byrne96}, and summer camp programs in VR for students \cite{Bricken93}. The VRRV and summer camp projects focus on "world-building" activity, where students conceive and create the objects of their own virtual worlds, by using 3D modeling software on desktop computers. Although this gives the opportunity for students to understand the process involved in creating a virtual setting, the actual immersive experience is limited to a short visit of the pre-designed virtual worlds (4 to 10-minute VR experiences). The concept that virtual reality is a process and not a product is important \cite {Bricken91}, but does not take advantage of the potential educational benefits of the technology and does not justify its use. One of the reasons students are not more actively involved with the actual virtual experience within the virtual reality system is the fact that the systems used by these projects (HMDs) are not flexible enough to allow more than one participant at a time. The VRRV project attempts to overcome such restrictions in an interesting way, by travelling to schools and giving students (of grades 4-12) the possibility of experiencing VR, although still one at a time and for a short time.

Virtual reality technologies such as the CAVE \cite{CruzNeira93} and the ImmersaDesk, are freed from the limitations of HMDs (unwieldy hardware, single-user participation, short and infreqent immersive experiences), but not the limitations of size and cost. Although these systems have been used successfully in the area of scientific visualization, the development of CAVE applications for education has been almost non-existent. CitySpace, a learning project where children build virtual cities, has been demonstrated in the CAVE, but not in an educational context. This project, as with HITL's projects, emphasized the children's modeling activity prior to incorporating the models in a virtual city. The more natural physical setup of the CAVE (a 10x10x10 ft. room rather than a device), the relatively non-intrusive interface hardware (no helmets to wear), and its multi-user support (more than one learner can share the experience at once) show promise as a VR tool for the development of learning environments. The NICE project, described in the next chapter, seeks to explore this potential.

Finally, a study performed in a non-academic setting is worth mentioning. The Computer Museum in Boston set out to compare the immersive VR, two-dimensional desktop, and video instruction of cell biology. The subject population was large and varied, consisting of randomly chosen museum visitors \cite{Gay94}. The study reports that the VR users scored higher on the factual post-tests than the video users, but lower than the monitor users. However, the immersive users enjoyed their experience the most and reported the most desire to continue learning about the subject after the experience. A very similar but limited study was also performed with Water on Tap \cite {Byrne96}, a chemistry application developed at the HITL, but the quantitative approach to this study produced results of questionable value.

2.2. Virtual Reality in Education

2.2.4. The Value of Virtual Reality for Education

Evident in the descriptions of the various projects by their researchers is the belief that VR affords unique learning opportunities. Setting aside the media hyperboly, which promises VR as the solution to all real world problems, there are reasons to believe that virtual reality can be a useful tool for learning. Unfortunately, research in the area has been unable to prove the need for VR. None of the mostly anecdotal findings indicate that something can be learned in VR that cannot be learned elsewhere. Researchers in the field assert that VR offers promise as a learning tool but struggle to justify its value for learning. As VR is already making inroads into the home in the form of inexpensive game systems, and advanced systems are becoming more accessible to the public, virtual learning environments become inevitable. If alternatives are not provided, the only developer of virtual learning environments will be the video game industry.

The following points summarize the ideas of researchers regarding the educational value of virtual reality.

1. Accessing the unreachable or the unrealizable.
   In virtual reality environments, students can have access to apparatus which in real life is too expensive, too delicate, too big, or too dangerous to keep in the classroom. This can be particularly helpful for students with physical or perceptual disabilities.

2. Multiple or alternative representations.
   The ability to create and explore from multiple or different perspectives is a major attraction in complex multi-sensory virtual environments. VR is seen as a new lens through which arbitrary levels of scale can be applied to facilitate observation \cite {Winn93}; alternative, even physically impossible, points of view can be chosen; time can be easily shortened or stretched out \cite{McLellan94}, with flexibility for repetition and self-pacing; objects and processes can be studied in isolation, in detail, in close-up or at a distance. VR provides a high-bandwidth medium for multiple representations of data (e.g., graphical, auditory, annotations), all within the same environment, which can be translated to other formats, reviewed, and analyzed as it is generated. These multiple observational records can be further analyzed together to evaluate models of learning and thinking.

3. Abstractions become more concrete.
   Cognitive scientists point out the ability of VR to make the abstract more concrete and visible by providing symbols not available in the non-symbolic real world. Focus can be given on salient aspects of a situation, so learners don't get lost in complexity.

VR advocates also tout examples of the value of VR as a medium for interactive, collaborative, and engaged learning. While highly responsive interactive virtual environments can provide immediate feedback and be extremely motivating, they may also foster automaticity, thoughtlessness, and lack of problem solving \cite {Collins96}. Telepresence, the integration of VR and high-speed networking, creates new possibilities for distance learning and the exploration of social behavior (e.g., cooperation, conflict) amongst remotely located learners who share the same space. However, once millions of students are connected, issues of control and the structure of learning will emerge. Virtual environments are also valued as being extremely motivating for learners, especially for those with non-traditional learning styles. Much of this engagement, however, is due in part to the novelty effect and not necessarily inherent in the design of the learning environment.

In developing virtual learning environments, one should not forget that access to information in innovative ways means little without guidance as to the use of that information. Additionally, many interactive VR systems, as is the case with interactive multimedia, can easily become glorified multiple choice systems. Interactivity, in this case, has nothing to do with responsiveness and certainly nothing to do with learning, at least in the constructivist sense.

2.3. Evaluation of Learning Environments

2.3.1. Qualitative vs. Quantitative Research

The area of educational research is widely divided when it comes to the issue of evaluation. Broadly defined, this division is mainly a dichotomy between the two methods of qualitative and quantitative evaluation. On one hand, quantitative evaluation, also known as empirical or positivistic, is based on controlled experiments: a research problem is selected, a hypothesis is formulated and then studied by gathering data specific to the problem. The evaluation setting is usually constrained and restricted to established conditions and assumptions that do not change during the study. The purpose of imposing constraints is to preserve the validity of the study.

Qualitative research, on the other hand, is primarily based on observation. Also known as naturalistic, ethnographic, or interpretive, it is a means for describing and attempting to understand the observed regularities in what people do, say and report of their experience as a whole, not as separate variables \cite{Ely91}. In qualitative research, the focus of attention is on the perceptions and experiences of the participants, on what they believe, the feelings they express, and the explanations they give. It is also focused on the process rather than the end result, assessing what individuals understand, through observation and dialog. Due to this approach, qualitative evaluation can be accused of yielding primarily anecdotal data that do not support general conclusions. It is suspected by many to be biased, subjective, and less rigorous than quantitative because it can loose precise answers and produce ambiguous results. It is also time-consuming.

What differentiates qualitative from quantitative inquiry are the assumptions held by the investigator, the researcher's beliefs rather than the methods of data collection. In this sense, qualitative evaluation is more relativistic, in comparison to quantitative assessment, which seeks the objective and ``absolute truth''. However, the difference between quantitative and qualitative evaluation is not just methodological, but involves a fundamental difference in the quality of learning itself. Knowing that the evaluation will be test-based, the material may be taught in a way to guarantee students' successful performance in the tests. By basing decisions about what is learned on final accomplishments, we change the motivation structure for students in school \cite{Collins96}.

Although qualitative research has a history of use in the social sciences, it has not been, until recently, a significant part of mainstream scholarship or research training. The education community is comprised mostly by researchers performing quantitative assessment and the formal educational practices employ exclusively quantitative achievement tests to evaluate student performance. The value of qualitative study continues to be an object of debate within the education community \cite{Ely91}.

For reasons described in the following sections, the quantitative approach was not used in the evaluation of the learning environment described in this thesis.

2.3. Evaluation of Learning Environments

2.3.2. Alternative Assessment

Although traditional assessment has long been a target of criticism within the context of education, it has proven to be remarkably resilient \cite{Wilson96}. Over the past decade, however, conventional assessment techniques have fallen into disfavor among education researchers, and the development of ``authentic'' assessment techniques remains an active but immature research area today. Significant to this switch has been the move to more flexible educational theories, such as the ones mentioned previously, which require new methods of evaluation that are directly related to the learning approaches themselves.

The increasing interest in alternative assessment is reflected in the proliferation of terms, such as authentic assessment, performance assessment, and portfolio assessment \cite{Wilson96}. These methods of assessment focus primarily on the process rather than just the product of learning. Authentic assessment, for instance, involves students in the evaluation of their own learning. Self assessment is an important component here. Similarly, performance assessment requires the learner to demonstrate knowledge by creating something. Portfolio assessment is focused on the process, as well as the final product. All of these evaluation methods emphasize common themes, such as problem solving and complex learning, which entail a wide range of responses and challenging tasks with multiple steps, time, and effort.

Alternative assessment affords us the ability to include motivation as an important factor in the evaluation process. This is especially relevant to constructivist learning environments which rely much more on intrinsic motivation than traditional learning environments \cite{Wilson96}.

The critics of alternative assessment complain that it is time and labor intensive, and heterogeneous. The outcome of the evaluation can vary widely in the specific knowledge domain being assessed. It can also vary widely because individual students' performance varies. It relies on students' verbal and communication abilities and there's no easy comparison among students.
Perhaps the most common critique states that alternative assessment lacks generlizability to other contexts. Some cognitive psychologists, however, don't consider this a disadvantage, as they believe that the nature of knowledge itself is highly contextualized with limited generalizability \cite{Brown89}.

2.3. Evaluation of Learning Environments

2.3.3. Assessment of virtual learning environments

While it is possible to identify areas in which VR could afford unique, effective learning experiences, to date there is little principled empirical work to support this belief. As Dede states, ``one of the biggest stumbling blocks in VR research right now is the lack of concrete data on the usefulness of VR" \cite{Dede96}. As VR becomes more widespread, it is important to investigate the educational efficacy of the medium in specific learning situations or broader learning domains, and develop new rubrics of educational efficacy that compare it to other approaches.

Dede believes that the efficacy of VR can be truly established only by rigorously comparing VR's benefits to traditional educational methods and only ``through careful analysis that can accurately diagnose the weaknesses and limitations of the technology" \cite {Dede96}. Whitelock et al. (1996) argue that effective evaluation methods need to be established to discover if conceptual learning takes place in VR \cite{Whitelock96}. In practice, however, the assessment of VR technology has been focused primarily on its usefulness for training and less on its efficacy for supporting learning in domains with a high conceptual and social content.

The question if VR requires new and different assessment techniques beyond those in development today, remains relatively unexplored. Although there have been some attempts at constructing theoretical frameworks for the evaluation of virtual worlds \cite{Rose95,Whitelock96}, very few working examples or reports on the practical use of these frameworks exist. The evaluation framework developed as part of this thesis is presented in Chapter 4.

The education world would argue that using just paper and pencil, in the form of standardized tests, is not an effective way to evaluate a virtual learning experience. As VR is a dynamic learning tool, evaluation should be inextricably coupled with the actual learning process. In essence, what applies to the evaluation of constructivist learning environments, should also apply to the evaluation of virtual learning environments, as they both place learners in positions where they can explore, experiment, and actively solve problems. Following the authentic assessment model, learning in constructivist learning environments, whether open or virtual, is directly related to its evaluation \cite{Wilson96}. Moreover, considering the incomplete nature of the field at this time, the key to conducting meaningful assessment will be to apply multiple measures of learning and performance \cite{Rose95}.

Moving further, in the direction of the use of virtual reality itself as a tool for assessment, VR has shown to stretch the technology for supporting the analysis of events. It has been used as a tool for recording, annotating, synchronizing, and analysing the results. Virtual learning environments can be used to extend this idea of VR as an evaluation medium in many ways. Networked virtual reality systems can embed methods for facilitating learner's discourse while in the environment. Actual mentors, disguised as virtual characters, serve multiple purposes as guides and evaluators: to answer questions, direct action, ask for clarification, prompt for interpretation. In addition to recording data while in the virtual environment, it is also easy to ``playback" the recorded actions or experiences for further reflection and interpretation. This form of assessment, embedded in the learning process, can provide meaningful reflections on learners' skills and knowledge.

3. THE N.I.C.E. PROJECT: DESIGN AND IMPLEMENTATION

An obvious step in the exploration and testing of the educational ideas, systems, and evaluation approaches mentioned in the previous chapter was to develop a working system that is informed by these issues. This lead to the development of the Narrative Immersive Constructionist / Collaborative Environments project, or NICE.

The NICE project started in 1995 as a collaboration between students from the Interactive Computing Environments laboratory and the Electronic Visualization Laboratory at the University of Illinois at Chicago. Inspired by work conducted at both laboratories, especially by the Graphic StoryWriter \cite{Steiner92,Steiner95,Steiner94,Steiner94a}, NICE set out to explore the possibilities of creating a virtual learning environment based on storytelling and narrative. As the group was comprised of members with diverse interests, this original idea was extended to include and combine many other thematic areas into a family of educational environments.

As currently implemented, NICE is an experiential learning environment in which children explore, cultivate, and tend a healthy virtual garden. In order to achieve this, they select seeds, plant vegetables, pull weeds, and control symbolic representations of environmental factors, such as a raincloud to provide rain, a sun to give sunlight, or a compost heap to recycle weeds. To facilitate their understanding of the underlying principles of this micro-ecosystem, they can change their size or observe the plants' roots. The garden is part of a larger island that includes a volcano, a bridge, the beach, and an underground maze. In addition to their activities in the garden, multiple children can collaborate synchronously within the virtual environment, through avatar representations and audio, or asynchronously by observing the visual narrative generated by the system as a WWW page \cite{Roussos96,Roussos97,Roussos97b}.

Figure: A child interacting with NICE in the CAVE(tm).

3. THE N.I.C.E. PROJECT: DESIGN AND IMPLEMENTATION

3.1. Goals in the Design of NICE

Guided by the premise that the design of environments for learners must be guided by educational theory, the principle goal of NICE is to realize current theories of learning in the design of the virtual environment. This implementation encompasses a combination of the theoretical underpinnings mentioned in the previous chapter.

In order to put a number of important theoretical principles to work, we had to determine who would interact with the system and what ideas or skills we want our users to acquire. The subject matter chosen was a qualitative idea: ecological models and the control of different environmental factors. The interest was in studying the value of VR as a medium for conceptual learning of a subject matter rather than a training platform. The choice of a qualitative idea allows for better examination of conceptual learning, because it is not linked to symbols or other notation. Moreover, qualitative ideas are usually more abstract, requiring simplification in order for them to be understood, and their instruction can be compatible with a constructivist view of learning. The target group chosen was children. We were particularly interested in studying the behavior of children in virtual environments and their reactions towards networked participants represented by avatars. The choice of young users seemed to be the most appropriate for studying the effectiveness of virtual environments for teaching and cooperation.

One of our concerns before developing NICE regarded the uniqueness of the project in comparison to other learning projects in general, and other virtual reality projects in particular. What could NICE provide that other learning projects did not? The answer to this question was sought by identifying the value of VR for learning. What, besides the use of the specific CAVE technology, could make NICE a unique virtual reality application? The general consensus was to create an alternative to the VR applications developed for learning. The issue of a persistent project grounded on theory, that moved beyond the scientific visualization of a dataset (a common application of VR) was the starting point for NICE.

The principal objectives of NICE are outlined in the following points:

- Learning from multiple perspectives which can be realized only in VR: engaging graphical representation, the change of time, the ability to manipulate objects that are bigger or smaller, or objects that do not exist.

- Learning how to collaborate with other people. An important goal of the NICE environment is to provide a collaborative setting in which students can learn about role differentiation, social interaction, social interaction in ``silent" networked worlds, and the notion of self in situations of telepresence. Learning, in this case, is through collaboration and can be particularly useful for communication-apprehensive children that do not thrive socially.

- Learning by actively exploring and controlling environmental variables, as a way to better understand them. What happens to a plant if it is not watered or if it gets too much water? The learners should be able to actively observe the effects of their actions.

- Taking the above a step further is programming by demonstration. An interest of the NICE project is to allow learners to specify the characteristics of the ecological model. The children can learn more about modeling when they have the ability to construct their own viable model. The notion of end-user authoring has been implemented in other learning projects \cite {Cypher93,Cypher95,Kahn96} but has never been explored in VR.

- Exploring story structures. Alan Kay has coined the term ``active essay'' to describe a new form of narrative expression in which manipulable computational objects are embedded within text. A goal for NICE is to provide an interesting testbed for the exploration of active narrative, a persistent narrative with multiple threads that the child can create.

- Creating a final product: the learner should be able to walk away with a program, a printout, a videotape, a map, etc.

Each of these goals has influenced the design of NICE in areas that will be described below. Additionally, a description is included of how some of these goals have been accomplished in the current implementation of NICE.

3. THE N.I.C.E. PROJECT: DESIGN AND IMPLEMENTATION

3.2. The Design Process

Both technical and educational design issues were taken into consideration in the process of developing NICE. The main areas of development presented below include constructionism, collaboration, narrative, the user interface, persistence, and intelligent agents.

3.2.1. Constructionism

The approach to constructionism taken by NICE echoes Papert's ideas in two ways: first, the learners can craft the environment within the virtual world (construction inside the VR environment), and second, the children can construct something meaningful to them.

The form of construction originally conceived for the NICE environment resembled the ``incredible machines'' or the Lego projects at MIT. Children can understand how complicated machines can be by constructing them. The environment provides the primitives, building blocks, or raw materials, as well as compositional rules for these objects. Although the rules include a well structured set of semantics for the behavior of each object, the objects should have separate structure from function; otherwise they would be overloaded with functions. Attached to each object is a description of its physical structure. The functionality of the object is not in the object itself. An object can be made to interact by being selected or brought into physical proximity. It is the person that interacts with it that constructs the affordances of the object.

Some of these thoughts have been implemented in NICE but not to the level of building blocks. In the garden one can shrink, jump, select and plant, but cannot construct from primitives. While other VR projects support ``world-building'' activity, the construction takes place outside of the virtual space. It is a technical challenge to provide a construction environment within virtual reality. The focus has thus moved to the construction of the underlying ecological models from within the virtual world through end-user authoring techniques. Learners may work in a ``greenhouse'' to qualitatively specify plant growth rules or isolate individual controlling variables. In NICE, the garden itself is the artifact, the object that the children design on their own. Additionally, a more physical product of the interaction with the virtual space, is the story published on the web page.

3.2.2. Collaboration

The construction of the environment may be in support of fostering the collaboration. Our goal in NICE is to promote social interaction. As with the Graphic StoryWriter \cite{Steiner92}, learning environments may teach interaction and social skills in an implicit way. In the case of virtual reality, additional opportunities for the support of role differentiation are afforded but additional issues also arise. Each person has a mental model of the other person and the world. How does the mental model change when one can not see the person? What do you do when other people don't behave the way you expect?

The physical structure of the CAVE offers the advantage of supporting a group of users that can experience the virtual world simultaneously. Support for remote collaboration has been achieved in NICE through the use of networking and the representation of participants as avatars. The avatars handle a wide range of gestures that allow close-to-real-life communication (\ref{figavatars}).

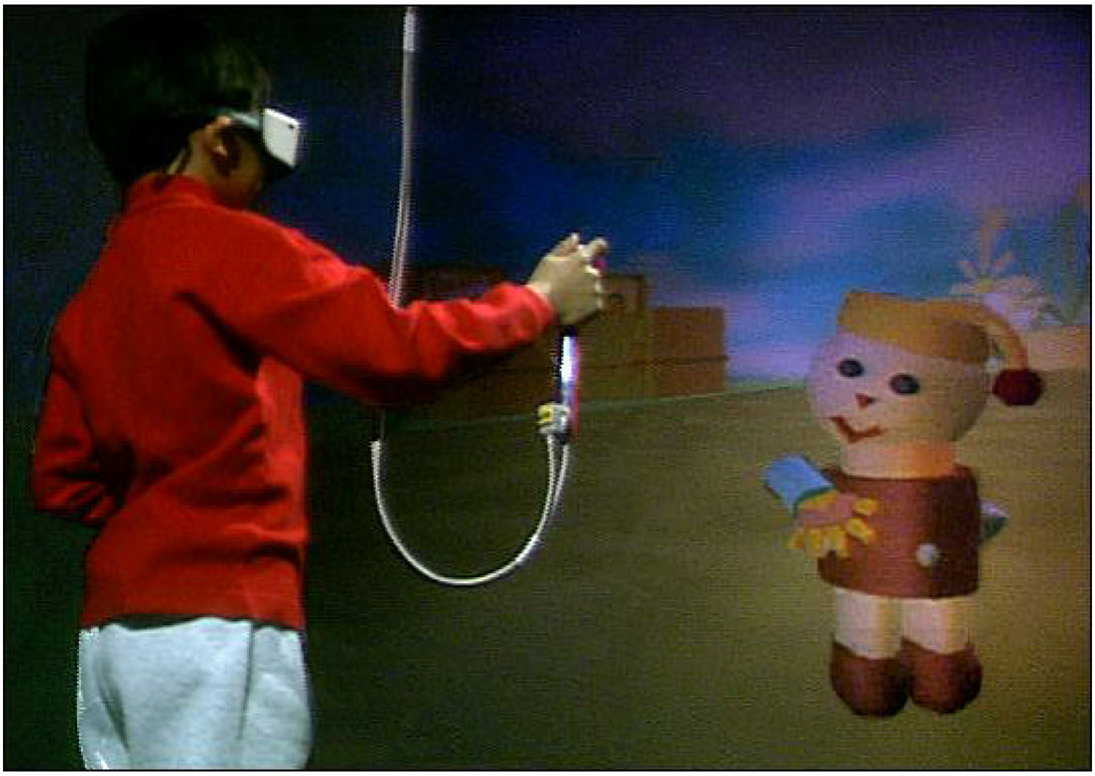

Figure: Remote collaboration in NICE enabled via avatars.

3.2.3. Stories

Of great interest in NICE has been the development of a narrative environment that would support multiple, resumable, infinite, interesting, and evolving stories. This could mean a complex story or many story threads which do not limit the participants to having a single overall set of goals or a single storyline. What is important is the potential for stories with a loose goal that give the freedom to the user to create closure or leave it open. Creativity and ingenuity are more engaged by open-ended situations where conclusions cannot be foretold.

The problems with narrative in synthetic environments have been noted by others \cite{Bates92,Maes89}: Where is the story represented? Is there a database or are there rules? Is there a general central definition of the story? Do you need a director? Is there a moral to the story? Is the moral something we can build or, for example, will the selection of the characters define the moral? An idea for defining a personal story is to have the user choose an initial set of characters and objects to interact with. This is a subset of a larger set and does not have to be the same every time. It can be removed or misplaced by the user. Thus the child crafts the story by selecting the initial state of the world. There are also standard objects, embellishments, or requirements loaded by the designers from the start. The rule base operates with these two subsets to provide stories that will require problem solving.

The story in NICE is the long-term multiparticipant environment. In essence, our manifestation of story in NICE follows the stories created by emergent writers. Children typically advance from telling list-like stories to more complex and easily recognizable stories \cite{Applebee78,Steiner92}. The story that the NICE system outputs now is much like the first. The narrative action in the CAVE is captured through a simple transcript generated by the system. The transcript is automatically published on a WWW page after passing through a simple parser which replaces some of the nouns in the text with their iconic representations (\ref{figstories}).

Figure: The WWW story generated by the NICE system.

3.2.4. User Interface

In a learning environment, the interface assumes special importance. Engagement of the user at the level of the interface takes on special importance in the case of virtual reality, as the goal is to induce the participant's complete and willing suspension of disbelief. The main issue is how to build a simple, familiar interface to a complex environment. In the case of NICE, the challenge is in supporting all the actions, especially the construction aspect, with an interface that is natural, flexible, and easy for use by children.

To support multiple actions in VR, various techniques have been used by researchers. The user may have the ``Midas touch'' or wear ``smart costumes'' \cite{Laurel94}. When touching an object it can change colors or functions, or the participant may wear different coats to perform different actions, such as the ``move'' coat, the ``scale'' coat, the ``rotate'' coat and so on. Also important in VR is spatial orientation, to provide the children with a ``sense of place''. In the real world, people are known to leave marks to orient themselves. Brenda Laurel (1994) extends this idea with the virtual equivalent of footprints, graffiti, shadows, or planting flags on the moon.

The NICE world includes several features to enhance the interface. At the level of the hardware, efforts have been made to replace the wand, the CAVE's main interaction device, with a more natural interface method, such as a dataglove or physical gardening tools. Currently we utilize the wand, using all three buttons for the simple actions and the joystick for moving. The virtual world includes the representation of the child's hand holding the wand. Attached to the hand is a small menu as a reminder of the button functions, and a virtual pointer to aid picking and navigation. Visual and audio feedback mechanisms are provided: every object that can be selected is surrounded by a yellow balloon and the plants indicate their state by displaying umbrellas or wearing sunglasses (\ref{figfeedback}). The children can also view their reflection on the surface of the water surrounding the virtual island.

Figure: Visual feedback cues for interaction in NICE.

3.2.5. Persistence

An area we are greatly interested in exploring is the notion of never-ending virtual worlds. That is, virtual worlds that continue to exist and interact either with other participants or intelligent agents, even when no one is currently in the world. This allows users to enter a virtual world, work on extended tasks that may take some time to complete, and have the opportunity to continue at another time. This is significant in terms of learning, as the ideal uses of learning technology should not be episodic but involve a voluntary and longer-term commitment, not an isolated short experience. Challenging issues involve the depiction of persistence: how will the users know that the world has evolved, how will they view what has happened during their absence?

The idea of persistence is constantly being developed in NICE. The garden itself evolves just like a real garden; if it is not tended, the plants grow old and weeds take over. The design of the system is such that it allows the parts which include behavior to operate without running the virtual environment.

3.2.6. Intelligent Agents and Actors

The children entering a virtual world can have an experience that is loosely predetermined: they may choose their level of participation by taking a leading role or not, use their physical abilities, or extend them. The experience involves the world, objects, and other characters, either remote, proximate, real, or artificial. Ideally, smart agents can be active participants, serve as directors, and be capable of picking and manipulating objects, just like real people. Their behavior and interaction can be described through scripts or evolve autonomously, depending on the degree of intelligence.

With respect to learning, artificial characters may be guides or ``mentors'', facilitating the teacher by taking on the load of more automated tasks. Moreover, children with higher levels of communication apprehension may be more comfortable talking to a synthetic, fantasy creature, which may be anthropomorphic but not a real human. They can communicate and act without the fear of a critical adult or teacher.

The developments of Artificial Intelligence in the area of autonomous artificial agents have not reached a level to achieve the flexible, responsive behavior required by a learning environment. For this reason, the ``agents'' and characters in NICE are still real people disguised by their avatars. This technique is called Wizard of Oz \cite{Wilson88} and has been used in other learning projects and VR worlds \cite{Laurel94}, adding the value of performance and aiding the study of social issues.

The NICE project is continuously being expanded to include many of the above ideas. The goal is to create a testbed application that explores and exemplifies these principles in an educational context.

3. THE N.I.C.E. PROJECT: DESIGN AND IMPLEMENTATION

3.3. System Architecture

The architecture of the NICE system is based on techniques derived from CALVIN (Collaborative Architecture Via Immersive Navigation), a networked collaborative environment for designing architectural spaces, developed by Jason Leigh and Andrew Johnson \cite {Leigh96a,Leigh96b,Leigh96c}. The NICE system is organized into three main modules: the behavioral module, the networking module, and the graphics and interaction module. These modules are relatively independent of one another (thus facilitating the distribution of the programming tasks) and communicate using a reliable socket connection (\ref{figsystem}).

The author's programming contribution to NICE is the behavioral module, or central server (LIFE), which handles the actions in the garden, calculates the environmental model, and shapes the output of the story. The vegetables and flowers follow simple rules of evolution and behavior. They share a common set of factors which contribute to their growth, such as their age, the amount of water and sunlight they need, and their proximity to other plants. The combination of these attributes determines the growth and affect of each plant. When the user selects a vegetable seed and plants it in the soil, the graphics and interaction module sends a message to the central server. The server adds this plant to the existing family of plants, calculates its characteristics and continuously transmits the state of each plant to the other modules. The growth of each plant is indicated by a change in its size. Its humidity state is indicated by the display of an umbrella when the plant is too wet, or sunglasses when the plant is too dry (\ref{figfeedback}). The central server operates constantly in the background and attempts to reestablish lost connections in the event of failure. If the server is stopped, there is a choice of resuming its operation by either starting from where it last stopped or beginning a new garden. Additionally, the LIFE server contains another module which automatically outputs the server's activity in the form of a WWW page, thus creating the story described and seen above in \ref{figstories}.

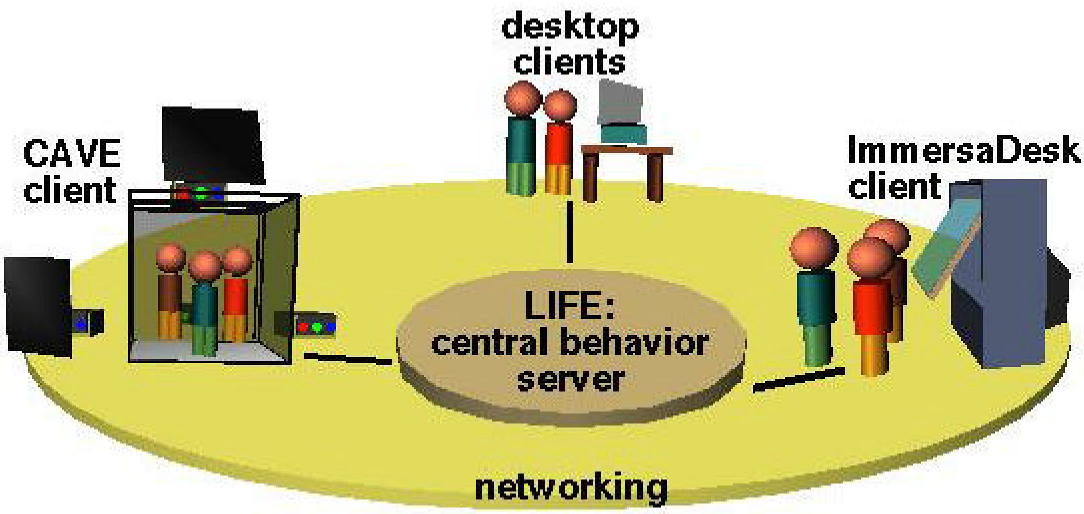

Figure: The NICE system architecture.

The network component in NICE, written by Jason Leigh, allows multiple networked participants to explore the same virtual space. The number of participants is limited only by bandwidth and latency of the network. Multiple distributed applications running on separate VR systems are connected via the central LIFE server to guarantee consistency across all the separate environments. The communications library uses multicasting and/or UDP to broadcast positional and orientation information about each participant's avatar, and uses TCP/IP to broadcast state information between the participants and the central server. New networking middleware optimized for VR is currently being developed, which transparently supports both the networking and database broker services as well as provides richer database support for sustaining persistence in VR. This allows NICE to be viewed easily from Java applets on web browsers.

The avatars also follow the CALVIN model. Each participant's head and hand are tracked, allowing the environment to transmit gestures between the participants, such as nodding, turning, pointing, and waving. As the child waves her hand in the real world, her avatar waves its hand in the virtual world. As these avatars have sufficiently detailed representations, the participants can communicate notions of relative position to one another. The NICE system's graphics and interaction module, programmed by Andrew Johnson, uses the Silicon Graphics' Performer library and the CAVE library. Performer allows maintainance of the necessary frame rate to support VR, and provides functionality for culling, multiprocessing, and intersection testing. All the objects and avatars in NICE are Inventor models. As part of the database design, each of these objects may easily be included or excluded from the world and features may be enabled or disabled. Some of these features include collision detection, calculation of the user's position over an object (so users can climb over it or up stairs), friction (so avatars can pick each other up), animation, and so on.

More recent additions to NICE include a two-dimensional web-based view of the three-dimensional virtual garden and a JAVA-based interface to the virtual world, but these features have not been fully developed and were not used in the evaluation studies described in the following chapter.

4.1. Conceptual Framework

Evaluation involves gathering and processing feedback on the issues that are being studied with the ultimate goal of demonstrating the worth or validity of a design. The purpose of this framework is to impose a structure on the evaluation of open virtual learning environments, such as NICE. The need for a detailed framework to assess the value of a learning environment follows reasons that have been emphasized in previous chapters (chapters 2 and 3).

The exploratory nature of this study requires a sound conceptual framework that encompasses, rather than restricts, the multiple dimensions of the issues that need to be examined in a virtual learning environment. Taking into account the multidimensionality of learning as well as virtual reality as a field, a number of technical, orientational, affective, cognitive, pedagogical, and other issues are examined \cite{Lewin95}. The proposed framework, although used specifically for the evaluation of NICE, is meant to serve as prototype for a general evaluation framework that can be used for assessing the educational efficacy of virtual reality experiences.

- Technical
- Orientation
- Affective
- Cognitive
- Pedagogical
- Other (value-added, etc.)

Table I
SUMMARY OF EVALUATION FRAMEWORK

Framework Category	Issue	Measurement
Technical	Usability	Time to learn an interface, comprehension of instructions, physical and emotional comfort
Orientation	Navigation, spatial orientation, presence and immersion, and feedback	Time to become immersed and comfortable in the environment
Affective	Engagement, preference, and confidence	Length of engagement, time to reach fatigue, reported and perceived enjoyment
Cognitive	Conceptual change, new skill	Performance within and outside the environment, think-aloud and stimulated recall techniques, oral and written surveys, video documentation
Pedagogical	Content general and specific teaching techniques	Collaboration (e.g., turn-taking, conflict, interaction), avatar acceptance, comparison of techniques
Collaborative VR	The added value of collaborative VR to instruction and learning	Comparisons of instruction and learning within and outside of collaborative VR environments

1. **Technical**

The technical aspect examines usability issues, in terms of interface, physical problems, and system hardware and software.

Interface issues:

a) learning the interface. How long does it take for a child to learn the buttons and joystick of the wand, or the menu? How long and clear do instructions need to be before the child is independent?

b) using the interface. Does the child feel comfort when using the interface device? Does the child use all features and consult the menu? Does the child collaborate with the other children to learn/use the interface?

Physical issues:

a) how many users display motion sickness?

b) how long can the child participate without feeling uncomfortable?

Evaluation of the system hardware and software:

a) robustness for museum/school use.

b) cost/effectiveness for museum/school use

c) flexibility of the program. Can it be expandable and by whom? Can it be authored, modified, customized?

2. **Orientation**

The orientation aspect examines the relationship of the user to the virtual environment, including navigation, spatial orientation, presence and immersion, and feedback issues.

Interface issues:
- a) is navigation simple, understandable ?
- b) how long does it take the child to understand his/her virtual relationship to the environment?

Spatial orientation:
- a) how is left, right, up, and down perceived?
- b) how quickly do the children orient themselves?
- c) are there any differences between ages or genders?

Presence/immersion: how ``real'' or how three-dimensional does the environment seem to the child?

Feedback: types of feedback (glowing sphere, sound, etc.) and their effect.

3. **Affective**

The affective parameter looks at the user's engagement, likes and dislikes, and confidence in the virtual environment.

Motivation: a) is the environment engaging? b) what captivates learners' attention the most? c) for how long? To ensure accurate responses, multiple methods may be employed: observe time spent doing or using something, interrupt children in certain instances based on certain focal points rather than time, ask what five things were liked and disliked about NICE, and use a survey with diagrammatic representations of their feelings (``happy and sad faces'').

Likes and dislikes: a taxonomy of features that the children liked or disliked can be created. For example, likes or dislikes may include the control over the environment, the graphics and animation, the sound, the difficulty in movement, and so on.

Confidence: does the amount of control over the environment, for example, empower the child?

4. **Cognitive**

The cognitive aspect identifies any improvement of the subject's internal concepts through this learning experience. Evaluating the cognitive aspect means measuring learning. Cognitive evaluation might start within the environment with the given learning task built into the experience. In NICE, for example, an agent or avatar can give goals to the users or ask them questions (e.g., plant and harvest a row of tomatoes). The responses to these activities could reveal what the user understands about the environment while inside it.

Approaches to the assessment of cognitive abilities include :
qualitative methods of observation

a) users demonstrate their understanding through self explanation with ``think-out-loud'' while doing (can you explain what you are doing?) and ``meta-explanation'' (can you explain what you did/learn?).

b) open-ended written surveys or open-ended oral questions, depending on user's age.

c) continuous video documentation to be viewed and transcribed later.

d) observer as a teacher-avatar in the environment.

quantitative

pre/post tests can be valuable in some situations. In the constructivist approach, the process of knowledge acquisition is equally important to the final product, so evaluation should be embedded in the actual learning process. However, testing skills before and after might also be useful.

5. Pedagogical

Teaching approach: how to effectively gain knowledge about the environment and the concepts that are being taught -in this case, ecology or earth science. With respect to virtual reality and NICE, this seems like a second level task once the technical and motivational issues are better known. A collaboration with a teacher who wants to teach something specific would add validity.

Collaboration: do the children learn how to collaborate? How does cooperation aid or hinder learning? Specifically,

a) observe local interaction between 2 or more users (turn-taking, conflict),

b) observe interaction with remote participants,

c) observe differences in communication apprehension and social behavior after the experience.

Narrative : use it as a tool to enhance concept mapping. Does it offer any insights, improve spatial orientation, affect motivation?

6. Other

``Value added'': what does VR technology add to the instruction that can not be achieved otherwise? This concept could fit into the technical or the pedagogical categories. When children plant a natural garden and then a virtual one, what is it in the virtual that complements the learning experience of the real? Where and how do we measure what the advantages of VR are versus the real thing?

In the case of NICE, for example, what does the ability to shrink and examine the plants, manipulate the environment, or collaborate with children across distance, add to the learning process?

4.2. Evaluation Methodology

In carrying out the research for this thesis, several studies were conducted, based on a qualitative approach to evaluation. The studies were exploratory in nature (as is the NICE project) and attempted to test out qualitative ideas and issues rather than performance. According to Lincoln and Guba (1985), successful qualitative research is performed through prolonged and persistent observation and through a triangulation of methods, each of which is flawed in some way (cited in Ely, 1991).

Following this model, a wide range of methods for data collection were employed to ensure the assessment of different aspects of this study. These include observation, survey, and interview methods, all of which are interpreted and cross-verified. Several other elements, such as the reports and portfolios produced by the children after the study, were taken into account.

Observation is perhaps one of the most important methods used for the evaluation of VR, as it can provide a first-hand view of the children's behavior and interactions in the virtual world. Both real time observation as well as video and audio recordings were used. All the sessions were videotaped for later observation and transcription. A significant part of the observation is an examination of the kinds of discourse, or ``exploratory talk'', used by the kids during their experience. As a first step, experience from observing NICE users of different ages at various events and conferences helped to refine the methods of observation that would be used for the evaluation of NICE. These users included the attendees of large conferences such as SIGGRAPH '96, computer-literate teenagers at the ThinkQuest event, students ages 12-15 as part of a pre-freshman enrichment program at UIC, and teachers interested in incorporating technology in their instruction at the NCSA Education and Outreach workshop.

The children participating in the main studies were asked to complete written questionnaires before and after their experience in virtual reality. The initial questions attempted to gather as much information as possible about each child's relationship to technology, knowledge of gardening, likes and dislikes. We wanted to establish what knowledge and understanding of the concepts displayed in the environment the children brought with them before the study. After their experience in NICE, the children completed an additional set of questions that related to their impressions and understanding of the environmental relationships in the NICE garden. All of these questions included space for open-ended responses, as they attempted to gather the children's actual thoughts and feelings written in their own words, rather than quantified responses. Additionally, a set of surveys involving the effect of VR on children's level of communication apprehension were conducted by Victoria Pierce and followed the standard pre-test / post-test structure \cite{Pierce97}.

Finally, an open-ended set of interviews was conducted with the children after their experience in virtual reality. The children talked about what they did while in the environment, what they liked or disliked, and what they thought they learned.

The parents were also asked to complete a short questionnaire, pertaining mostly to demographic information -name, age, gender, school information- as well as computer usage information. Both parents and children were asked how much time the kids spend watching television, playing video games, or using a computer per day. This was to avoid biased answers (parents like to think that their children don't spend too much time in front of a screen, so when they say 1 to 2 hours, they really mean 2).

The evaluation included a total of fifty two children: case studies of eight children and two classrooms of twenty two students each. All students volunteered to be part of the study. The ones participating in the case studies were accompanied by their parents, while the classrooms took a school field trip and were accompanied by their teachers. All children returned a parent permission form in compliance with the human subject review panel (IRB \#H-97-078) which is included in Appendix B.

The activities at each evaluation session of NICE took approximately one to three hours to complete, depending upon whether the tests were conducted with groups or pairs of children. This included time to introduce the activity and organize the students, give them time to plan the activity beforehand, perform the activity inside the VR environment, and have some time for post-activity questions and discussion. With the children in the two classrooms there was a familiarization day at their school a few weeks before the activity. This included some pre-test questionnaires to help identify the strength and weaknesses of the children, and their degree of communication apprehension \cite {Pierce97}. The VR setting in all studies included the CAVE and one or two Immersadesks, all linked by an audio connection.

The use of all these techniques should be interpreted as exploratory in nature, involving many qualitative judgments suggested by what came naturally when the research was conducted. The next sections will detail the setting, questionnaires, interviews, and results of each set of studies.

4.3. EMPIRICAL STUDIES

Initial and main studies were performed as part of the evaluation of NICE. Both these studies are described in the following setions.

4.3.1. Initial Study

Pilot studies are used as precursors to larger, more formal evaluation studies. The initial study conducted for NICE was limited in comparison to usual pilot studies. Two case studies took place about a month apart from each other; the first was the study of a 9-year old boy, while the second included two girls, one 9 and the other 7 years old.

The research objectives of these studies were twofold: First, to examine issues regarding the usability of the system, to become familiar with the procedure for conducting studies in the CAVE, to solve any technical problems early on, and to gain experience, so as to refine the testing process for the main studies. Second, to collect some preliminary observations on the differences in the childrens' interaction with the environment. We wanted to see if the differences in the habits, behavior, and attitude of the children affect their ability to interact in the virtual environment. The observation focused on what questions were asked by the kids, what they liked or disliked the most, where they kept their focus, how well they understood the goal. An important point in the design of the study was the focus on the differences between children that play video games and the ones that do not. In this sense, the initial studies were intended as comparative studies.

4.3.1.1. Method

The subjects were selected according to gender and level of computer literacy. Both selection criteria are related, as differences with respect to the use of technology have been reported between boys and girls. This is mostly due to stereotypes and the focus of the video game market on boys. Consequently, boys start using technology earlier and are more comfortable and confident with it than girls. We wanted to see if this behavior extends to VR. The boy selected is an experienced video game player and computer user, whereas the two girls spend most of their time reading books and hardly ever use a computer. None of the children had previous experience with the NICE environment. The boy, however, had experienced a commercial entertainment-based immersive virtual reality environment.

The CAVE and one Immersadesk were used for the initial studies. The boy was studied interacting alone in the CAVE for 5 minutes and then with a remote character, controlled by the researcher, for the remaining 65 minutes. The two girls were sisters and preferred to start out in the CAVE together. An avatar appeared after 30 minutes, and then another 30 minutes later the girls split and interacted with each other remotely. The younger one remained in the CAVE while the older continued from the Immersadesk. Their virtual reality experience lasted a total of 1 hour and 30 minutes. No audio connection had been set up for these initial studies, so the avatars communicated only with gestures.

After the experience the children were interviewed. Additionally, a set of pre and post study questions were asked, to identify communication apprehension problems.

4.3.1.2. Observations

Noticeable differences between the two case studies were observed during this first use of the CAVE environment with children. The boy, an avid video game player, learned how to use the interface very quickly and with hardly any instruction. His movement in the environment was confident, displaying complete immersion. He was very dynamic and animated, using his full body and turning his head a lot. Many times he kneeled or sat on the CAVE floor when the task required it. He accomplished his tasks in a stable and precise manner steadily picking and planting vegetables with a clear pre-planned sense of organization. He asked and answered his questions on his own but had to be reminded to speak ``out loud''. The boy's favorite part was the interaction with the remote character. He waved and followed the girl-avatar that was controlled by the researcher, or tried to keep the bird-avatar, which was trying to steal plants, out of the garden by moving his hands around and commenting out loud that he wished he had a scarecrow.

The girls, on the other hand, although engaged, stood fairly still when using the wand and faced straight ahead when moving around. They were not as talkative or comfortable at the beginning. Towards the end they were considerably more talkative, but did not seem to get better at navigating or feeling more confident when moving around in the environment. As the older girl noted, ``you want to go one way and it moves the other way''. Their collaboration, however, was natural: they switched glasses often to give each other a chance, or to help solve a problem. While the boy was much more adept at picking up objects, even those that were moving, the girls were not. Indicative is the older girl's comment about planting a vegetable: ``sometimes when you want to plant it doesn't get yellow''. She had not perceived that the yellow bounding sphere, indicating that the object may be picked or placed, would appear when pointing at the soil. The boy had understood this feedback cue right away. The girls' sense of organization was also very different from the boy's, indicated by the way they planted their garden. Their decisions lacked a clear structured manner and were fairly random. Their favorite part was the underground. They were both more animated and motivated when searching for the passage to the underground. The most notable difference was their behavior towards the other avatar. Unlike the boy, they almost completely ignored the researcher-controlled boy-avatar who entered the environment and was trying to get their attention. This doesn't mean that they didn't understand what he was doing, because while he was trying to get their attention, we asked one of the girls what the character was doing. She replied: ``He's telling me to come here,'' and continued doing what she was previously doing without paying attention to the character.

4.3.1.3. Identified Problems

These early tests pointed out or reconfirmed some problems regarding the user interface of the VR system. First of all, the size of the stereo glasses turned out to be a problem. The use of glass ``ties'' to tighten the large glasses on the head did not help in the case of the girls because they were switching the pair of tracker glasses constantly; each one wanted to make sure that the other had the same shared experience. Unfortunately, no better solution can be found for this problem until the industry produces new glasses.

Another problem, related mainly to testing, was the lack of an audio connection between the two VR systems. NICE's environmental sounds overpower the children's voices in the CAVE, making the extraction of information from the recorded video source very difficult. The addition of a wireless microphone is a solution for louder and clearer sound, but adds another piece of equipment for the kids to wear. This problem was solved in the final studies by using an ambient sound connection.

An attempt to have the children demonstrate their understanding during the experience, by thinking out loud, did not produce the desired results, as the children were a little shy and spoke softly. In the formal studies we used a ``teacher avatar'' to overcome this problem.

Some technical problems which occurred during both tests, such as the failure of the logging program to record the time, were corrected in time for the main studies.

In terms of the experiment design, the main problem was that of too many differences between the children: this study confirmed that a comparison between the children was difficult and its validity questionable. While there was an attempt to provide almost identical testing conditions, so the studies would be smoother and the comparison more valid, there were three major difference factors between the two case studies: gender, computer literacy, and the fact that the boy was alone, whereas the girls were together in the CAVE. Moreover, differences between the individuals in a comparison study of this size may overwhelm other factors.

4.3.2.2. Observed Results

The observed results from the case and classroom studies have been grouped based on the theoretical framework defined previously. These observations have been collected by converging the multiple pieces of data gathered through observation, interviews and questionnaires, and are presented below.

Technical issues
As with the initial study, the children in the main studies exhibited diversity in their use of the wand, the interaction device. The instructions given were not exactly the same for everyone but depended on the situation and the environmental or personal distractions. Generally, these instructions started with showing the representation of the virtual hand to the leader, then the use of the joystick for navigation, and finally, once the child was able to move into the garden, the function of the buttons. Learning the functions of the wand lasted from 2 minutes, for the children that learned quickly, to 7 minutes.

After learning how to use the wand, the children's effort was focused on orientation, as noted in the following section. Limitations of the physical design of the wand caused discomfort to young users, as both hands were needed to reach the buttons and press the joystick at the same time. The joystick of the current wand is difficult even for adult users and requires applying considerable force when navigating. The children obviously had a problem doing this. It was expected that the boys would generally be better at using the wand, partly because of their strength and partly because of their familiarity with similar input devices from playing video and computer games. According to both parents' and kids' reports, 92\% of the boys play electronic games weekly, as opposed to 42\% of the girls (see \ref{tablevideogame}). The majority of these games have joystick-based interface devices. We did not notice, however, any gender differences in learning to use the wand.

A larger problem was the size of the stereo glasses. Despite the glass-ties used to tighten the glasses on the children's heads, the glasses would still fall off. Most children had to hold the glasses with their free hand and, when tired holding them, would just take them off. Not only did this contribute to the subjects' fatigue, but also to their level of motivation and excitement. Since the stereo glasses and the wand are an integral part of the virtual experience, these limitations are a current hindrance not only to usability but also to learning.

Some of the best ``drivers'' consulted the menu that was attached to the virtual hand. These instances, however, were very few and sometimes inconsistent: some of the drivers who consulted the menu at the beginning forgot about it later, while others remembered to use it only part of the time.

The children's susceptibility to simulator sickness was not as large as expected. Less than 5\% of the subjects complained about getting a headache or being dizzy during or after the experience, and for most it was so slight that they had not noticed until asked. Only a few girls, mainly ImmersaDesk users, felt dizzy during the experience and for longer afterwards. One girl felt slight nausea commencing about 15 minutes into the experience, and lasting for about a half hour.

Evaluation of the system with respect to its robustness and cost effectiveness for broader use must be put off until the system is in a public locale. Nevertheless, it is necessary to mention the large number of human staffing required for these studies. Approximately eleven people helped each day of the classroom evaluations. Part of this number relates to the evaluation procedure itself (videocamera person, guides, interviewer, etc.), but a core number of at least three people, including a technical person, an instructor to handle the glasses and teach the use of the interface devices, and a teacher-avatar are needed for even the simplest case study. The NICE software is flexible enough to eventually expand into a user-authoring system. To be effective, however, it needs to be used by a small number of learners for an extended period of time.

VIDEO & COMPUTER GAME EXPOSURE OF SUBJECTS

Activity	Boys	Girls
Video-game playing	3.44	1.34
Computer use for school	4.14	2.4
Parent computer usage	5.85	6.16

in average hours per week

Orientation

After learning how to use the wand, the children focused on trying to navigate and orient themselves in the virtual environment. With respect to the classroom groups, this proved to be the effort of the leader and not of the other children in the group, although their mission was to help the leader. The drivers were the only ones focused on the orientation task at hand, as they were the ones navigating, while the other children were distracted by the movement and the three-dimensional graphics. The girls seemed slightly better at orienting themselves in the environment, possibly because they were generally more focused and reserved compared to the boys. Even with the case studies, although not nearly to the same extent, there were times when the other child (the one not using the wand) would wander around, instead of observing or directing the driver's actions. While it was not expected that all children's full attention would be given at orientation, the result in these studies was that each child came up with their own version of the right direction, voiced them at the same time as the other children and confused the leader, who then individually decided which was the right path to take. As a result, apart from the difficulty in using the joystick for navigation, the leaders exhibited noticeable individual differences in their abilities to interact with the 3-D environment. These differences seemed to relate to their level of ``independence'': the ones pursuing their own goals did well, while the ones that attempted to listen to the others in their group ended up confused and disoriented.

A test for spatial orientation was the ability to find areas in the space, such as the hole that leads to the area under the garden. This was a relatively hard task, although there were spatial clues: the passage was located near the only set of trees behind one of the garden fences. These were the only instances where verbal interaction seemed to work well, largely because the goal was very specific and required the kids' complete attention.

Another test for orientation was the concept map -the plan of the garden on paper. In the planning stage, students developed different strategies for planting (\ref{plannedgardens}, Appendix A). We wanted to see how they were able to implement this plan in VR. The case studies were more focused and, therefore, the children attempted to stick to their plan. With the exception of the boy in the initial study, the children were not successful at completing the task. Most children began planting as planned, but then changed their plans when running into difficulty. A younger girl who tried following the plan, commented that it was very hard to be precise in separating the vegetables. The teacher-avatar helped her with directions, but that ``wasn't enough''. The classrooms, on the other hand, hardly even tried to implement the plan, although constantly reminded by the teacher-avatar. Their entire experience was consumed by dealing with the group's behavior. None of the children admitted that they did not try; rather they stated that implementing the plan was a difficult task. One boy, after seeing the look of the group's final version of the garden asked his group: ``how come we didn't get it right?'' to receive the overwhelming response ``because it was very hard!''.

As perceived through observation, most kids felt immersed. This was indicated by their motion and excitement. Almost all children attempted to ``touch'' the virtual objects by moving and clasping their hands in the air. This was particularly noticable in the case of the virtual beam that extended from the user's hand to help point to and select objects. As the beam was always attached to the hand and close to the user, it felt very ``three-dimensional'' to the children. Many children, however, would take their stereo glasses off and put them back on constantly during the experience. One of the boys in the case studies would take the stereoglasses off every time he needed to accomplish a more demanding task, such as finding his way underground. When asked why, he answered that it was easier to see without the glasses because ``they were heavy and I couldn't do things right''. Technical obstacles, such as the size of the stereo glasses, possibly hindered the perception of presence.

The present feedback seems effective, as most children understood the function of the yellow balloon for picking, and the virtual hand. Many leaders waved at the other avatars with the hand that was holding the wand, indicating that they understood the relationship between the wand, their real hand, and the virtual hand.

Affective

Measuring motivation is difficult, as it is indirect. Anderson et al. (1975) note that we do not see motivation, but behavior. Moreover, in the case of virtual reality, motivation is highly driven by other factors, such as the novelty effect, media hype, and social issues. It is significant to look through these factors and try to identify whether the content taught within this medium is motivating for children, what it is that motivates them, and most importantly, for how long. This was difficult, as all of the children were excited before starting, just by the fact that they would experience virtual reality. Therefore, we had to look at their level of engagement during the actual experience.

The amount of time the children spent in VR ranged from 30 minutes to 1 hour and 30 minutes. Each classroom group, due to time constraints, stayed in the experience for about 30 minutes. The case-study subjects, on the other hand, were allowed to stay until they displayed noticeable fatigue, at which point they were asked if they wished to continue. Most cases wished to remain in NICE for at least 45 minutes and started getting tired after one to one and a half hours.

Interactive activities ranked high amongst the preferences of the children, as shown by their responses in the post study questions. Planting was a favorite. An equal number of responses were in favor of the area under the garden. The fantasy was another fundamental driving force for many of the children. Many liked the water (or ``swimming''), the rain, sun, umbrellas and sunglasses, and the vegetables. The three things that were most disliked by the children included ``the stuff that we had to move with'', the ``glasses falling off'', and the fact that some did not get to drive. Most (73\%) of the children answered ``nothing'' to the question ``what did you dislike the most?''.

The most important issue related to motivation is control. As mentioned in the discussion of orientation, the children that were leading were more on-task and engaged, while all others were distracted and unfocused. This was also perceived, to a lesser degree, with the pairs of children in the case studies: the driver was focused on the task even if that meant only navigation, and was consequently more engaged, while the second child seemed less engaged. The post-experience questions verify these observations: Children that were leaders listed that what they enjoyed the most was being the leader, while most others that did not get that chance were very disappointed.

Cognitive
Examining the cognitive value of a virtual learning environment is very difficult, as there are many other factors which correlate to learning, such as the ones described above. Particularly, distraction, fatigue, and cognitive overhead in mastering the interface influence the outcome. The classroom studies provide good examples of a situation in which all the above took place, and where one cannot derive any conclusions about conceptual learning. The results from the case studies are more promising, as the studies were more focused, prolonged, and with less noise and disorder.

However, even in the case studies, little can be concluded as far as learning is concerned. Confidence in using the interface does not necessarily signify understanding of the subject matter. One of the boys, for example, who reported playing many hours of video games per week, learned the interface very quickly and easily and had very good navigation and picking skills. After interacting with VR for about 40 minutes he was interviewed. During the interview and his post-study questions it was revealed that he had not perceived the effects of the sun and the rain on the plants, nor the function of the umbrellas and sunglasses. This was consistent with his pre-study test, which showed little knowledge of gardening concepts.

To simplify the understanding of the children's knowledge before and after the virtual experience, their responses were grouped into categories.

For the pre-study test, three categories were devised according to the children's understanding of simple ecological relationships. The first category included the responses that displayed a very good understanding of gardening concepts: the plants need water and sunlight (i.e. good temperature), and good soil to grow, they wilt or look brown when they are sick, they wilt if they get too much water and dry out when they get too much sun, and the weeds need to be pulled out. About 12\% of the subjects answered in this way. They were also the ones ranked high in reading/writing skills by the teachers. The second category included most of the above answers except for a few misconceptions (e.g. water is good but sunlight is bad for plants). 42\% of the children's answers fit into this category. The third category included 44\% of the responses, where more than one question included a ``don't know'' response or a wrong answer (such as ``the plants grow down'' when they get sunlight, or that weeds need to be planted and watered). Finally, one child could not answer most of the questions.

The answers to the post-study questions were grouped into categories based on the children's understanding of the NICE model: the plants display umbrellas when they receive too much water and sunglasses when there's too much sun, while the weeds are recycled in the compost heap. The responses here were more difficult to categorize, as many children had trouble synthesizing their learning during post-testing, due to fatigue or excitement, while others misunderstood the questions and answered in the same way as in the pre-test, not understanding that the post-questions pertained to the NICE garden in particular.

Approximately 17 children (35\%) understood, for the most part, the NICE model. Of these 17, 13 were drivers, and all had done well in their pre-study questions. This shows that most of the leaders, children that were actively engaged in the task, understood the model of the NICE garden, whereas only a few of the other children perceived it. Approximately 45\% of the children simply answered ``they grew'' to the questions ``what happenned when you put the rain over the plants'' and ``what happened when you put the sun over the plants''. Five kids answered that they did not know or see what happened while six kids were tired and did not answer at all.

Pedagogical

The children acted naturally while in NICE, just as they would have at a playground. They played, argued, listened, spoke loudly, and even rested. Very few were curious about the technology, excepting a girl who exclaimed that the screens were made of paper. The presence of ``the computer'' was not generally perceived by the children throughout the sessions. As one child put it, ``I thought we were going to play with a computer, but this was different". This indicates that perhaps virtual reality can provide a natural medium for teaching, once technical and technology-specific problems are resolved.

Although children in these studies participated in the VR session longer than in any other educational VR study, it appears that this was not an important factor in the facilitation of learning. We must agree, however, with Dede (1996) who reports that spreading lessons over multiple VR sessions appears to be more effective than covering many topics in a single session, as we attempted to do in our studies. Reviews and post-tests from their studies demonstrated that students were better able to retain and integrate information over multiple lessons. This is usually the case in school-based learning as well as being the main concept of life-long learning.

With respect to their pedagogical function in the NICE studies, collaboration and the narrative are explored further in the following sections.

Collaboration

The classroom studies were set up to encourage intra-group collaboration and inter-group competition, to ensure that each group had an incentive to focus on the task of creating a tended garden. However, none of these forms of cooperation occured. After each group was split, one sybgroup to go to the CAVE and the other to the ImmersaDesk, the children had to be continuously reminded by the teacher-avatar that they were still one group working on a common goal in the same garden. Most children, however, continued not to perceive this and regarded the other (remote) half of their group as their competitors. There were multiple instances of the two drivers fighting over who would grab the raincloud, and children from one location yelling at the ones in the other location to step out of ``their" garden. As far as the classrooms were concerned, competition contributed to the excitement of the children in the group, but kept them off-task and distracted them for nearly the entirety of the experience. Some of the groups even displayed a form of intra-group competition between the leader and other members. This related mainly to the control of the wand. Notable is the case of one girl who caused constant conflict because she was not the one chosen to be in control. The intent during these studies was to have only one child in each group control the wand. Our rationale for this was efficiency: it is easier and quicker to teach one subject than all; it is more efficient for one to control while others direct the activity, and it avoids fighting over who will do it.

On the other hand, this efficiency gain might not be helpful in terms of advancing all the students' learning. In the case of the other students, it was evident that the control over their learning and their experience was in the hands of the leader of the group. It was hoped that, in this way, the students would be able to pay more attention to the subject matter by leaving the control of the learning situation to the leader. For the child controlling, we supposed that this would not be an advantage, as it could lead to less attention to the subject matter and more to the task of controlling. As noted previously, the opposite was observed in these studies: the leader paid more attention to the subject than the other, less active members of the group.

Contrary to the classroom's behavior, the pairs of children in the case studies displayed excellent collaboration and no competition. In most cases, on-task communication was observed and there was general agreement on actions. Based on these observations, issues regarding the selection and number of members in a group must be taken into account for a successful collaborative combination.

For both the classroom as well as the case studies, the teacher-avatar seemed to serve a helpful purpose, especially for giving the kids tips and keeping them on task. In terms of the classroom children, of course, the teacher-avatar consumed most of her time attempting to keep order - not unlike a real classroom.

The Story

The system's visual output (the narrative WWW page) was shown to each group during the interview. Each group was represented in the story by the avatar of the leader. Some children did not understand this until it was explained to them while showing them the narrative. Most were fascinated by the pictorial representations of the characters and vegetables and remembered what they were doing by looking at the story. It is believed that the iconic representation was helpful in giving the groups a general overview of their actions and is worthy of further exploration. An unanticipated function of the story was its use as a spelling aid by two children from different groups. When completing their questions, they asked to look at the WWW page to find the spelling of certain object names.

CHAPTER 5. CONCLUSIONS

The main conclusions drawn from the observations presented in the previous chapters may be summarized in the following points.

Most children needed less than 10 minutes to learn how to use the NICE interface. Nevertheless, technical limitations, such as the wand and the size of the stereoglasses, are a current hindrance not only to usability but also to learning. Hardware issues are one of the most fundamental blocks to learning and these obstacles must be overcome in order for any virtual experience to be effective. Physical fatigue and cognitive overhead in mastering the interface are issues that should be addressed in the design of the virtual learning experience.

Several problems experienced in these studies point to issues that must be addressed to make adequate conclusions with respect to learning. Specifically, group dynamics, in the case of the larger classroom groups, overcame the effects of the virtual reality experience in several instances. Instead of collaborating, the groups in different physical locations seemed to conflict. Also, subdynamics within each group took place; the other members often hindered the leader if he or she paid attention to them at all. The issue of control over the environment, in some cases, became more important to the group members than their experience. We can conclude from this that either larger groups don't work at all in virtual environments, or more likely, that larger groups need to be more carefully orchestrated in order to achieve effective results. These group factors can be more powerful than the experience itself, and therefore in the design of any virtual learning environment careful decisions about the size and selection of the group, as well as its social aspects, should be made.

Perhaps the most important conclusion these studies point toward is the role VR can play in the actual quality of learning. The affective factor is one of the most important characteristics about VR in general, encompassing interaction and engagement, but in these studies was not limited only to increased motivation in the children. The amount of interactivity and engagement directly influenced the outcomes related to the effectiveness of the learning process. Learning was directly tied to the children's level of engagement: the ones actively involved were more engaged and motivated to complete a task and consequently understood the model employed by the NICE garden. The leaders (the ones actively using the control device to drive the interaction) got the most out of the experience while passive viewers did not understand the concepts presented as well as the leaders. This is a tentative conclusion, since the children were chosen as leaders based in part upon their cognitive skills (leadership ability, cognitive skills, and height, in that order). Nonetheless, further study in this area is promising.

5.2. Future Work

Additional research in virtual reality is required to develop and refine the instructional practices and technological tools that serve both a learning and evaluative function. For this work in particular, several directions have been identified for future exploration.

With respect to the design of the educational environment, a number of modifications may be implemented. An immediate modification to NICE would be the enhancement of the central behavior model to support a more complete life cycle. This means adding appropriate feedback cues to indicate the end of the plant's life cycle or other characteristics that follow a more realistic environmental model. Additional enhancements include some features that were originally planned but not implemented, such as the relationship between the growth of a plant and its proximity to the plants surrounding it. The NICE project is a continuing effort with many possibilities for expansion. If our goal with NICE is to provide emerging learning experiences, then the learner should take a more active role. It is necessary to create a model of an environment that will place students more in charge of their own learning and thinking, in a way that can engage more than one learners (in our case more than just the driver). An important step in this direction is the development of an end-user authoring system that will allow children to construct their own ecological models. It seems that this should initially take place outside the VR environment and eventually move in, when technology permits. Other immediate needs include the development of a more familiar and natural interface for use by children, that strives to resolve the technical difficulties observed during this evaluation.

With respect to the evaluation framework, a more detailed and focused framework is needed on the conceptual side. The cognitive and pedagogical aspects of VR learning need to be examined further, as technical factors are currently being resolved by the industry. Specifically, the framework can be expanded to account for the evaluation of deep ideas and concepts that children have a hard time with, and support ways in which their acquisition can be evaluated.

With respect to the evaluation itself, substantial amounts of data have still to be gathered on both the initial questions and the questions that emerge during observation. Future research demands intensive observational studies of children, as well as longitudinal studies on the role of virtual reality in shaping attitudes and learning. To gain a deeper understanding of the issues involved, a more directed evaluation should be conducted with multiple case studies, studied in depth, for longer and multiple times. Case studies of a small number of subjects make prolonged engagement and persistent observation much easier to accomplish, as the subjects can stay longer, come more than once, and be observed under calmer, more relaxed conditions. They also allow to deal with variables that are hard to quantify, such as affective behavior, but which are important in the study of virtual reality for learning.

Appendices

Pre-test questions

```
\item   Do you play video games?
\item   What kind?
\item   If you could be anybody or anything, what would you choose to be?
\item   Do you have a garden at home?
\item   Which are your favorite plants?
\item   Which are your favorite flowers?
\item   When is the best time of year to plant a garden?
\item   What do you do to help a sick plant?
\item   What does a plant need to grow?
\item   How do you know that a plant is sick?
\item   What happens to plants when it rains?
\item   Can plants get too much water?
\item   Do plants need sunlight?
\item   What happens to plants when there is too much sunlight?
\item   Can plants get too much sun?
\item   What do you do with the weeds in a garden?
\end{itemize}
```

Post-test questions during the interview

```
\item   Did you drive in the NICE garden?   yes   no
\item   Would you prefer to work alone in the garden?
\item   What did you talk about with the other kids?
\item   What did you think about Vicki?
\item   What 3 things did you like the best?
\item   What 3 things did you dislike the most?
\item   If you could be any character in the VR garden, what would you choose
     to be?
\item   What happened when you put the rain over the plants?
\item   How did you know that a plant needed water?
\item   What happened when you put the sun over the plants?
\item   How did you know that a plant needed sunlight?
\item   What did you do with the weeds in the garden?
\item   Draw what your garden looked like:


\item   Follow-up essay in the classroom:

     "If you could design your own VR world, what would you do?"
```

Garden plans and final gardens

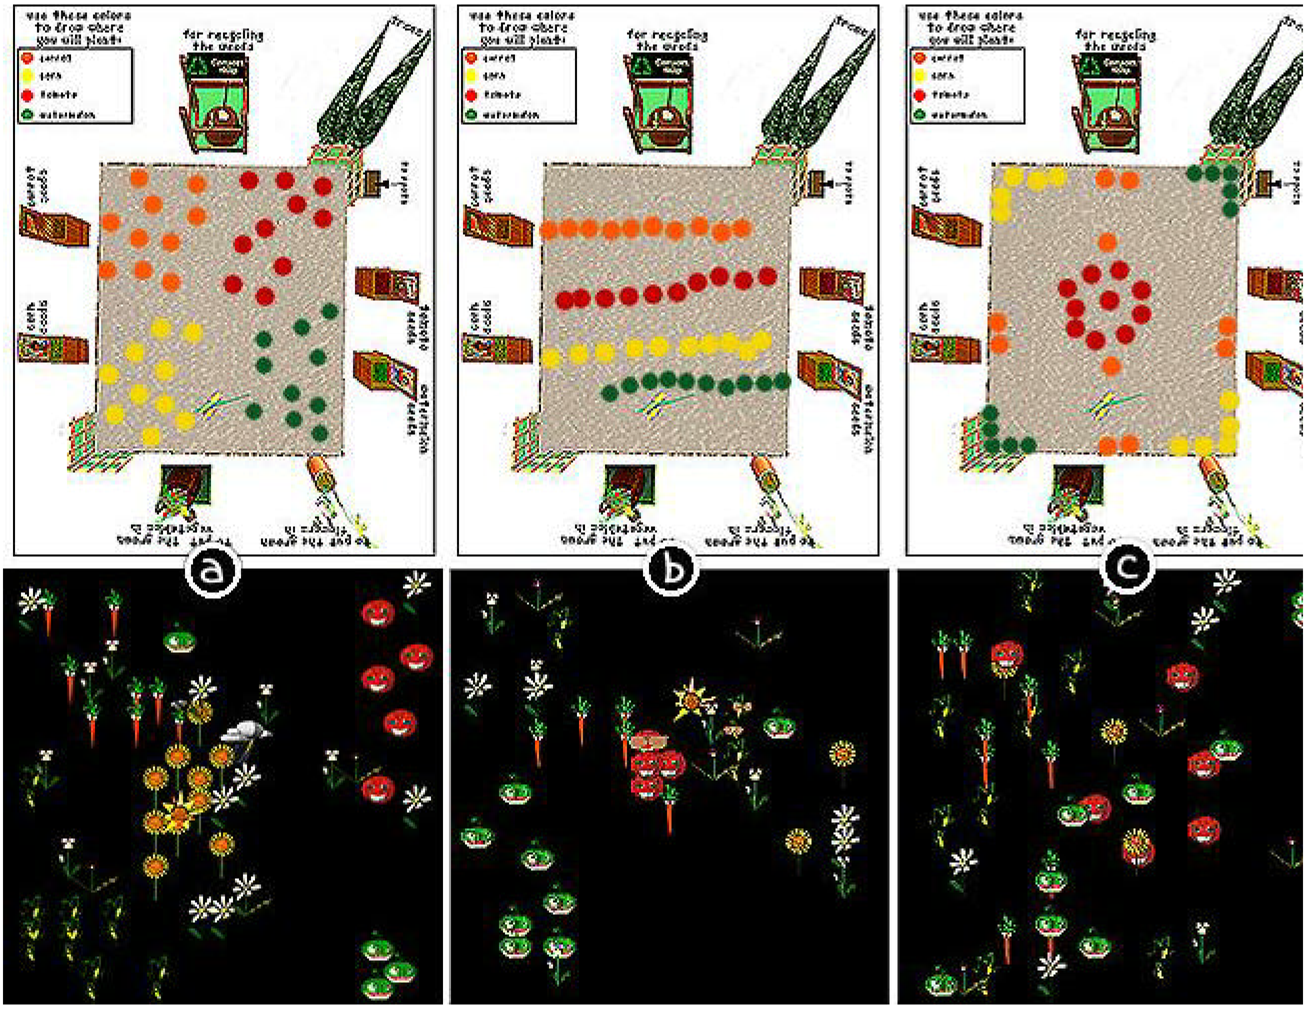

Top row: some of the garden plans created by (a) two boys from the case studies, and (b) and (c) two classroom groups. Bottom row: ... the view of their final gardens generated by the system.

Drawings

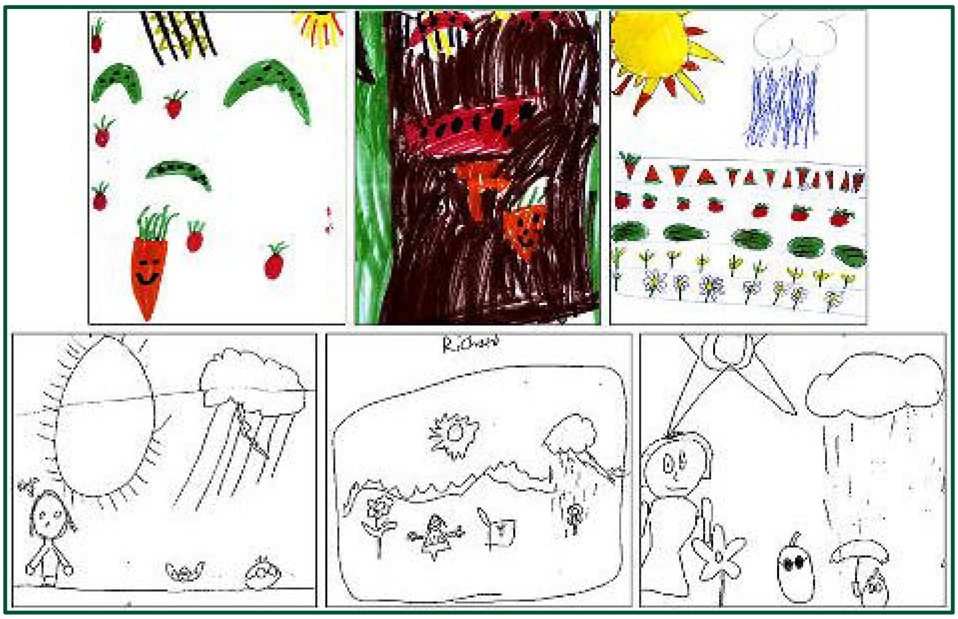

Figure: Some of the children's drawings of the garden.

```
% Maria Roussos - 1996, 1997.
% gathered bibliography for thesis in alphabetical order in bibtex format



@Book{Anderson75,
  author =       "Scarvia B. Anderson and Samuel Ball and Richard T. Murphy",
  title =        "{ Encyclopedia of Educational Evaluation }",
  publisher =    "Jossey Bass Inc. San Francisco",
  year =         1975,
  keywords =     "evaluation"
}

@Book{Applebee78,
      author =       "Arthur N. Applebee",
      title =        "{The Child's Concept of Story : Ages Two to Seventeen }",
      publisher =    "University of Chicago Press",
      year =         1978,
      annote =       "{UIC Main Library       BF723.S74 A6}"
    }

@Book{Bakeman88,
       author = "Roger Bakeman and John M. Gottman",
       title = " { Observing Interaction: An Introduction to Sequential Analysis
       publisher = "Cambridge University Press",
       year = 1988,
       ISBN = "0-521-27593-8",
       note = "{chapter 2: Developing a coding scheme}"
    }


@Techreport{Bates92,
      author =       "Joseph Bates",
      title=         "{The Nature of Characters in Interactive Worlds and The Oz
                      Project}",
      institution =  "Carnegie Mellon University",
      year =         1992,
      number =       "{CMU-CS-92-200}",
      annote = "{http://www.cs.cmu.edu/afs/cs.cmu.edu/project/oz/web/papers.html
    }


@Article{Bricken91,
      author =       "Meredith Bricken",
      title =        "{Virtual Reality Learning Environments: Potentials and
                     Challenges}",
      pages =        "178--184",
      journal =      "Computer Graphics",
      volume =       25,
      number =       3,
      year =         1991,
      keywords =     "education, virtual reality"
    }

@Inbook{Bricken93,
      author =       "Meredith Bricken and Christine M. Byrne",
      title =        "{ Summer students in VR: a pilot study }",
      series =       "{Virtual Reality : Applications and Explorations}",
      pages =        "178--184",
      publisher =    "Academic Publishers Professional",
      year =         1993,
      annote =        "{book Virtual reality : applications and explorations.
                       Main Book Stacks QA76.9.H85 V57 1993 }"
    }


@Inbook{Brown89,
      author =       "J.S. Brown and A. Collins and P. Duguid",
```

```
      title =        "{ Situated Cognition }",
      series =        "{Artificial Intelligence and Education Vol. 2, Lawler and
      pages =        "254--268",
      publisher =    "Ablex Publishing",
      address =      "Norwood, NJ",
      year =         1989

    }


@Article{Bruckman95,
  author =       "Amy Bruckman and Mitchel Resnick",
  title =        "{The MediaMOO Project: Constructionism and
                  Professional Community}",
  journal =      "Convergence",
  year =         1995,
  volume =       1,
  number =       1,
  pages =        "94--109",
  annote =       "{http://asb.www.media.mit.edu/people/asb/convergence.html}"
}


@Phdthesis{Byrne96,
      author =       "Christine M. Byrne",
      title =        "{Water on Tap: The Use of Virtual Reality as
                      an Educational Tool}",
      school =       "University of Washington",
      year =         1996,
      annote =     "{http://www.hitl.washington.edu/publications/dissertations/B
    }

@Article{Chan96,
  author =       "Tak-Wai Chan",
  title =        "{Learning Companion Systems, Social Learning Systems,
                  and the Global Social Learning Club}",
  journal =      "Journal of Artificial Intelligence in Education",
  year =         1996,
  volume =       7,
  number =       2,
  pages =        "125--159"
}



@InCollection{Collins96,
  author =       "Allan Collins and Bolt Beranek and {Newman Inc.}",
  title =        "{Design Issues for Learning Environments}",
  booktitle =    "{International Perspectives on the Design of
                  Technology-Supported Learning Environments}",
  publisher =    "Lawrence Erlbaum Associates",
  year =         1996,
  editor =       "Stella Vosniadou and Erik De Corte and Robert Glaser
                  and Heinz Mandl",
  address =      "Hillsdale, NJ",
  annote =        "{online at http://192.127.252.21/journal/apr-may/ip2learn.htm
}


@Article{Cromby95,
      author =       "John Cromby and Penny Standen and David Brown",
      title =        "{Using Virtual Environments in Special Education}",
      pages =        "1--4",
      journal =      "{VR in the Schools}",
      volume =       1,
      number =       3,
      year =         1995
    }
```

```
@InProceedings{CruzNeira93,
      author =       "Carolina Cruz-Neira and Daniel J. Sandin and Thomas A. DeF
      title =        "{Surround-Screen Projection-Based Virtual Reality: The
                      Design and Implementation of the CAVE}",
      booktitle =    "Proceedings of {ACM SIGGRAPH} '93",
      pages =        "135--142",
      year =         1993,
      publisher =    "ACM Press",
      address =      "New York",
      copyright =    "(c) Copyright 1993 Association for Computing
                     Machinery"
    }


@Book{Cypher93,
      editor =       "Allen Cypher and Daniel C. Halbert and David Kurlander
                     and Henry Lieberman and David Maulsby and Brad A. Myers
                     and Alan Turransky",
      title =        "{Watch What I Do: Programming by Demonstration}",
      pages =        672,
      publisher =    "MIT Press",
      address =      "Cambridge, MA",
      year =         1993,
      price =        "\$45.00",
      ISBN =         "0-262-03213-9",
      contents =     "{not listed}"
    }



@InProceedings{Cypher95,
      author =       "Alan Cypher and D.C. Smith",
      title =        "{KidSim: End User Programming of Simulations}",
      booktitle =    "Proceedings of ACM CHI'95 Conference on Human Factors
                     in Computing Systems",
      pages =        "27--34",
      year =         "1995",
      copyright =    "(c) Copyright 1995 Association for Computing
                     Machinery",
      note =         "{Denver, CO}"
    }


@Book{DeVries87,
  author =       "Rheta DeVries and Lawrence Kohlberg",
  title =        "{Programs of Early Education: The Constructivist View}",
  publisher =    "Longman, New York",
  year =         "1987",

}

@Article{Dede95,
      author =       "Chris Dede",
      title =        "{The Evolution of Constructivist Learning Environments:
                      Immersion in Distributed, Virtual Worlds}",
      pages =        "46--52",
      journal =      "Educational Technology",
      volume =       35,
      number =       5,
      year =         1995,
      keywords =     "education, story grammars",
      annote =       "http://www.virtual.gmu.edu/evolcons.htm"
    }

@InProceedings{Dede96,
      author =       "Chris Dede and Marilyn Salzman and Bowen R. Loftin",
      title =        "{MaxwellWorld: Learning Complex Scientific Concepts Via Im
                       in Virtual Reality} ",
```

```
      booktitle =    "Proceedings of Second International Conference of the Lear
      pages =        "22--29",
      year =         1996
    }



@Book{Dewey66,
      author =       "J. Dewey",
      title =        "{Democracy and Education}",
      year =         1966,
      publisher =    "Free Press",
      address =      "New York",
      annote =       "{ http://www.ilt.columbia.edu/ academic/texts/dewey/
                      d_e/contents.html }"
    }



@Book{Duffy92,
      editor =       "Thomas M. Duffy and David H. Jonassen",
      title =        "{Constructivism and the Technology of Instruction : A
                     Conversation}",
      pages =        221,
      publisher =    "Lawrence Erlbaum Associates",
      address =      "Hillsdale, NJ",
      year =         1992,
      ISBN =         "0-8085-1272-5",
      annote =       "{LB1028.38 .C66 1992}"

    }

@Book{Egan88,
      author=        "Kieran Egan",
      title =        "{Teaching as Storytelling: An Alternative
                      Approach to Teaching and Curriculum }",
      publisher =    "Routledge",
      address =      "London",
      year =         1988,
      annote =         "{UIC Main Library LB1027.E414 1988}"
    }

@Book{Ely91,
  author =       "Margot Ely and M. Anzul and T. Friedman and
                  D. Garner and A.M. Steinmetz",
  title =        "{Doing Qualitative Research: circles within circles }",
  publisher =    "The Falmer Press",
  address =      "London",
  year =         1991
}



@Techreport{Gay94,
      author =       "Eben Gay and David Greschler",
      title =        "{Is Virtual Reality a Good Teaching Tool? }",
      institution =  "Boston Computer Museum",
      year =         1994,
      note =       "{http://www.tcm.org:80/cgi-bin/dynamo/resources/research/vr-
    }


@Article{Grove96,
      author =       "Jonathan Grove",
```

```
      title =        "{VR and History - Some Findings and Thoughts}",
      pages =        "3-9",
      journal =      "{VR} in the Schools",
      volume =       2,
      number =       1,
      year =         1996,
      publisher =    "East Carolina University",
      keywords =     "{VR, education, history, evaluation}"
    }


@Book{Harel91,
      author =       "Idit Harel",
      title =        "{Children Designers. Interdisciplinary Constructions
                      for Learning and Knowing Mathematics in a Computer-Rich
                      School}",
      publisher =     "Ablex Publishing",
      address =      "Norwood, NJ",
      year =          1991,
      annote =         "{a few chapters from this book.
                      It was given to me by Umesch.}"
    }

@Book{HarelPapert,
      editor =       "Idit Harel and Seymour Papert",
      title =        "{Constructionism }",
      publisher =    "Ablex Publishing",
      address =      "Norwood, NJ",
      year =         1991,
      annote =       "UIC main library: BF311 .C657 1991"
    }


@InBook{Hughes96,
  author =       "Hughes, Charles E. and Moshell, J. Michael",
  title =        "{ ExploreNet }",
  series =       "The Virtual Reality Casebook",
  editor =       "C. E. Loeffler and Tim Anderson",
  publisher =    "Van Nostrand Rheinhold",
  address =      "New York",
  year =         1995,
  pages =        "118--122"
}

@InProceedings{Issroff96,
      author =       "Kim Issroff and Teresa DelSoldato",
      title =        "{Incorporating Motivation into Computer-Supported
                     Collaborative Learning}",
      booktitle =    "Proceedings of European Conference on Artificial
                     Intelligence in Education",
      pages =        "284--290",
      year =         1996,
      editor =       "Colibri Editions",
      address =      "Lisbon, Portugal",
      annote =      "{http://www.cbl.leeds.ac.uk/~euroaied/papers/}"
    }

@Article{Johnson85,
  author =       " David W. Johnson and Roger T. Johnson and Mary Beth
                  Stanne ",
  title =        "{Effects of Cooperative, Competitive, and
                  Individualistic Goal Structures on Computer Assisted
                  Instruction}",
  journal =      "Journal of Educational Psychology,",
  year =         1985,
  volume =       77,
  number =       6,
  pages =        "668 -- 677"
}
```

```
@Article{Johnson86,
  author =       " David W. Johnson and Roger T. Johnson ",
  title =        "{Computer Assisted Cooperative Learning}",
  journal =      "Educational Technology ",
  year =         1986,
  pages =        "12 -- 18"
}

@Book{Kafai95,
  author =       "Yasmin Kafai",
  title =        "{Minds in Play}",
  publisher =    "Lawrence Erlbaum Associates",
  address =      "Hillsdale, NJ",
  year =         1995,

}


@Article{Kahn96,
      author =       "Kenneth Kahn",
      title =        "{ToonTalk- An Animated Programming Environment for Childre
      pages =        "175--196",
      journal =      "Journal of Visual Languages and Computing",
      volume =       7,
      number =       2,
      year =         1996,
      keywords =     "visual languages, education",
      annote =       "{http://www.virtual.gmu.edu/evolcons.htm}"
    }



@Book{Labinowicz80,
  author =       "Ed Labinowicz",
  title =        "{The Piaget Primer}",
  publisher =    "Addison-Wesley Publishing Co.",
  year =         1980,
  annote =       "0-201-04090-5"
}


@Book{Laurel93,
      author =       "Brenda Laurel",
      title =        "{Computers as Theatre}",
      publisher =    "Addison-Wesley Publishing Co.",
      year =         1993

}


@InProceedings{Laurel94,
      author =       "Brenda Laurel and Rachel Strickland",
      title =        "{Placeholder: Landscape and Narrative in Virtual Environme
      booktitle =    "Proceedings of ACM Multimedia '94,",
      pages =        "121--132",
      year =         1994,
      copyright =    "(c) Copyright 1992 Association for Computing
                     Machinery",
      annote  =      "{UIC Main library periodicals section 2nd floor south}"
    }

@Book{LawlerYazdani,
      editor =       "Robert W. Lawler and Masoud Yazdani",
      title =        "{Artificial Intelligence and Education}",
      publisher =    "Ablex Publishing",
      address =      "Norwood, NJ",
```

```
      year =         1991,
      annote =         "{UIC main LB1028.3 .A754 1987}"


    }

@Article{Lebowitz84,
      author =       "M. Lebowitz",
      title =        "{Creating Characters in a Story-Telling Universe}",
      pages =        "171--194",
      journal =      "Poetics",
      volume =       13,
      year =         1984,
      annote  =      "{UIC Main library periodicals section 2nd floor south}"
    }

@Article{Leigh96a,
      author =       "Jason Leigh and Andrew E. Johnson",
      title =        "{Supporting Transcontinental Collaborative Work in
                     Persistent Virtual Environments}",
      pages =        "47--51",
      journal =      "IEEE Computer Graphics and Application",
      year =         1996
    }


@Inproceedings{Leigh96b,
      author =       "Jason Leigh and Andrew E. Johnson and Christina A. Vasilak
      title =        "{Multi-perspective Collaborative Design in
                      Persistent Networked Virtual Environments}",
      booktitle =    "IEEE Virtual Reality Annual International Symposium",
      year =         1996,
      pages =        "253--260"

}


@Inproceedings{Leigh96c,
      author =       "Jason Leigh and Andrew E. Johnson and Thomas A. DeFanti",
      title =        "{CALVIN: an Immersimedia Design Environment Utilizing
                      Heterogeneous Perspectives}",
      booktitle =    "Proceedings of the IEEE International Conference on
                      Multimedia Computing and Systems",
      year =        1996,
      pages =        " 20--23 ",
      address =      "Hiroshima, Japan"

   }

@InProceedings{Lewin95,
      author =       "Cathy Lewin",
      title =        "{Test Driving CARS: Addressing the Issues in the
                      Evaluation of Computer Assisted Reading Software}",
      booktitle =    "Proceedings of International Conference on Computers
                      in Education",
      pages =        "452--459",
      year =         1995,
      editor =       "AACE",
      address =      "Charlottesville, USA"
    }


@Article{Maes89,
      author =       "Pattie Maes",
      title =        "{How To Do the Right Thing}",
      pages =        "291--323",
      journal =      "Connection Science",
      volume =       1,
      number =       "3",
```

```
      year =         1989,
      annote  =      "{http://agents.www.media.mit.edu/groups/agents/Publication
    }


@PhdThesis{Malone80,
  author =       "Thomas W. Malone",
  title =        "{What Makes Things Fun to Learn: A Study of
                  Intrinsically Motivating Computer Games}",
  school =       "Stanford University",
  year =         1980
}

@InCollection{MaloneLepper87,
  author =       "Thomas W. Malone and M. R. Lepper",
  title =        "{Making Learning Fun: a Taxonomy of Intrinsic
                  Motivations for Learning}",
  booktitle =    "Aptitude, Learning, and Instruction: Cognitive and
                  Affective Process Analyses ",
  publisher =    "Lawrence Erlbaum Associates,Hillsdale,NJ.",
  year =         1987,
  editor =       "R. Snow and M. Farr"
}




@Article{McLellan94,
      author =       "Hillary McLellan",
      title =        "{Virtual Reality and Multiple Intelligences: Potentials fo
      pages =        "33--66",
      volume =       "5",
      number =       "2",
      month =        "Spring",
      journal =      "Journal of Higher Education",
      year =         1994
    }



@Techreport{Osberg92,
      author =       "Kimberley Osberg",
      title  =       "{Virtual Reality and Education: A Look at Both Sides of th
      institution =  "Human Interface Technology Laboratory, University of Washi
      year =         1992,
      number =       "TR-93-7",
      annote =       "{http://www.itl.washington.edu/publications/r-93-7.html}"

    }

@Article{Pantelidis93,
  author =       "Veronica Pantelidis",
  title =        "{Virtual Reality in the Classroom}",
  journal =      "Educational Tecnology",
  year =         1993,
  volume =       33,
  pages =        "23--27",
  month =        "April"
}

@Book{Papert80,
      author =       "Seymour Papert",
      title =        "{Mindstorms: Children, Computers, and Powerful Ideas}",
      publisher =    "Basic Books, Inc.",
      address =      "New York",
      year =         1980,
      key =          "ai,education, psychology",
```

```
      ISBN =         "0-465-04627-4",
      annote =       "{Mathematics/430 SEO:  QA20.C65 P36 1993 }"
    }


@Inbook{Papert93,
      author =       "Seymour Papert",
      title =        "{The Children's Machine: Rethinking
                     School in the Age of the Computer}",
      pages =        "82--105",
      publisher =    "Basic Books",
      address =      "New York",
      year =         1993,
      annote =       "Chapter called A Word for Learning.
                     I found it in Resnick's dissertation page.
                     the book's call #:LB1028.5 .P325 1993"
    }

@Article{Parsons96,
  author =       "Ann K. Parsons and Bob Zenhausern",
  title =        "{Grassroots, Catalyst for Change}",
  journal =      "{VR} in the Schools",
  year =         1996,
  volume =       2,
  number =       1,
  pages =        "15--17",
  month =        "June",
  annote =       "{ http://eastnet.educ.ecu.edu/vr/Vrits2-1 }"
}

@Book{Piaget73,
       author =     "Jean Piaget",
       title =      "{To Understand is to Invent: The Future of Education}",
       publisher =  "Grossman",
       year =       1973,
       address =    "New York",
    }



@MastersThesis{Pierce97,
  author =       "Victoria Pierce",
  title =        "{ Communication Apprehension in an Immersive Networked
                  Virtual Learning Environment }",
  school =       "{ University of Illinois at Chicago }",
  year =         1997
}

@Book{Provenzo91,
      author =       "Eugene F. Provenzo",
      title =        "{Video Kids: Making Sense of Nintendo }",
      pages =        184,
      publisher =    "Harvard University Press",
      address =      "Cambridge, MA",
      year =         1991,
      ISBN =         "0-674-93709-0",
      key =          "video games, social,psychological aspects",

    }


@Book{Reigeluth83,
  author =       "C. M. Reigeluth",
  title =        "Instructional Design: What is it and why is it?",
  publisher =    "Lawrence Erlbaum Associates",
  address =      "Hillsdale, NJ",
  year =         1983,
  editor =       "C. M. Reigeluth",
  series =       "Instructional Design Theories and Models"
```

```
}



@Book{Rogoff90,
  author =       "B. Rogoff",
  title =        "Apprenticeship in Thinking: Cognitive Development in
                  Social Context",
  publisher =    "Oxford University Press",
  year =         1990,
  address =      "New York"
}

@Techreport{Rose95,
      author =       "Howard Rose",
      title =        "{Assessing Learning in VR: Towards Developing a Paradigm
                      Virtual Reality Roving Vehicles (VRRV) Project }",
      institution =    "Human Interface Technology Laboratory - University of Wa
      year =         1995,
      number =       "TR-95-1",

      annote =       "{http://www.hitl.washington.edu/publications/r-95-1/}"
    }


@InProceedings{Roussos96,
      author =       "Maria Roussos and Andrew E. Johnson and Jason Leigh and
                      Christina A. Vasilakis and Thomas G. Moher",
      title =        "{Constructing Collaborative Stories Within Virtual
                      Learning Landscapes}",
      booktitle =    "Proceedings of European Conference on Artificial
                     Intelligence in Education",
      pages =        "129--135",
      year =         1996,
      publisher =       "Colibri Editions",
      address =      "Lisbon, Portugal",
      annote =         "{http://www.ice.eecs.uic.edu/~nice/}"
    }



@InProceedings{Roussos97,
      author =       "Maria Roussos and Andrew E. Johnson and Jason Leigh and
                      Craig R. Barnes and Christina A. Vasilakis and Thomas G. M
      title =        "{The NICE Project: Narrative, Immersive, Constructionist/C
                      Environments for Learning in Virtual Reality}",
      booktitle =    "Proceedings of ED-MEDIA/ED-TELECOM '97",
      pages =        "0--8",
      year =         1997,
      publisher =       "AACE",
      address =      "Virginia",
      annote =         "{http://www.ice.eecs.uic.edu/~nice/}"
    }


@InProceedings{Roussos97b,
      author =       "Maria Roussos and Andrew E. Johnson and Jason Leigh and
                      Christina A. Vasilakis and Craig R. Barnes and Thomas G. M
      title =        "{NICE: Combining Constructionism, Narrative, and Collabora
                      in a Virtual Learning Environment}",
      booktitle =    "Computer Graphics: Proceedings of ACM SIGGRAPH '97",
      pages =        "1--4",
      year =         1997,
      address =      "New York",
      copyright =    "(c) Copyright 1997 Association for Computing
                     Machinery",
      publisher =       "ACM SIGGRAPH",
      annote =       "{http://www.ice.eecs.uic.edu/~nice/}"
```

```
    }


@Book{Slavin91,
  author =       "R.E. Slavin",
  title =        "{Cooperative Learning: Theory, Research, and Practice}",
  publisher =    "Prentice Hall, Engelwood Cliffs, N.J.",
  year =         1991
}

@InProceedings{Soloway96,
      author =       "Elliot Soloway and S. Jackson and J. Klein and others",
      title =        "{Learning Theory in Practice: Case Studies of
                      Learner-Centered Design }",
      booktitle=     "Proceedings of Computer-Human Interaction '96",
      publisher =    "ACM Press",
      address =      "New York",
      pages =        "189--196",
      year =         1996,
      copyright =    "(c) Copyright 1995 Association for Computing
                      Machinery"
    }



@InProceedings{Steiner92,
      author =       "Karl E. Steiner and Thomas G. Moher",
      title =        "{Graphic StoryWriter: An Interactive Environment for
                     Emergent Storytelling}",
      booktitle =    "Proceedings of ACM CHI'92 Conference on Human Factors
                     in Computing Systems",
      series =       "Use and Evaluation of Learning Environments",
      pages =        "357--364",
      year =         1992,
      copyright =    "(c) Copyright 1992 Association for Computing
                     Machinery",
      keywords =     "User interaction, Story grammars, Educational
                     software"
    }

@InProceedings{Steiner94,
      author =       "Karl E. Steiner and Thomas G. Moher",
      title =        "{A Comparison of Verbal Interaction in Literal and
                     Virtual Shared Learning Environments }",
      booktitle =    "Proceedings of ACM CHI'94 Conference on Human Factors
                     in Computing Systems",
      series =       "INTERACTIVE POSTERS",
      volume =       2,
      pages =        "97--98",
      year =         1994,
      copyright =    "(c) Copyright 1994 Association for Computing
                     Machinery",
      keywords =     "Computer supported cooperative work, Shared
                     workspaces, Computer supported cooperative learning,
                     Interactive learning environments, Constructive
                     learning"
    }


@Article{Steiner94a,
      author =       "Karl E. Steiner and Thomas G. Moher",
      title =        "{Scaffolding Story Construction With Interactive Multimedi
      pages =        "173--196",
      journal =      "Journal of Educational Multimedia and Hypermedia",
      volume =       3,
      number =       2,
      year =         1994,
```

```
      keywords =     "education, story grammars"
    }




@Phdthesis{Steiner95,
      author =       "Karl E. Steiner",
      title =        "{Representation and Interaction with Symbolic Structures in
                      Individual and Shared Learning Environments }",
      school =       "University of Illinois",
      year =         1995,
      annote =        "{UIC Main library TK99999 .S729}"
}

@Book{Vygotsky78,
      author =       "Lev S. Vygotskii",
      title =        "{Mind in Society: The Development of Higher Psychological
      publisher =    "Harvard University Press",
      address =      "Cambridge, MA",
      year =         1978
    }


@Book{Vygotsky78,
      author =       "Lev S. Vygotsky",
      title =        "{Thought and Language}",
      editor =       "A. Kozulin"
      publisher =    "MIT Press",
      address =      "Cambridge, MA",
      year =         1986,
      annote =       "{includes his theory of the Zone f Proximal development}",
   }





@InProceedings{Whitelock96,
      author =       "Denise Whitelock and Paul Brna and S. Holland",
      title =        "{ What is the Value of Virtual Reality for Conceptual
                      Learning? Towards a Theoretical Framework }",
      booktitle =    "Proceedings of European Conference on Artificial
                     Intelligence in Education",
      pages =        "136--141",
      year =         1996,
      editor =       "Colibri Editions",
      address =      "Lisbon, Portugal",
      annote =      "{http://www.cbl.leeds.ac.uk/~euroaied/papers/Whitelock1/}"
    }



@Inbook{Wilson88,
      author =       "J. Wilson and D. Rosenberg",
      title =        "{Rapid Prototyping for User Interface Design }",
      series =       "Handbook of Human-Computer Interaction",
      publisher =    "M. Helander Ed.",
      address =      "North Holland",
      year =         1988,
      pages =        "859--875",
      annote =         "{ }"
    }


@Book{Wilson96,
  author =       "Thomas C. Reeves and James R. Okey",
  title =        "{ Alternative Assessment for Constructivist Learning
                  Environments }",
```

```
  editor =       "Brent Wilson",
  series =       "{ Constructivist Learning Environments: Case Studies
                  in Instructional Design }",
  publisher =    "Educational Technology Publications ",
  year =         1996,
  address =      "Englewood Cliffs, NJ"
}

@Techreport{Winn93,
      author =       "William Winn",
      title =        "{ A Conceptual Basis for Educational
                  Applications of Virtual Reality }",
      institution =  "Human Interface Technology Laboratory - University of Wash
      year =         1993,
      number =       "TR-93-9",
      annote =       "{http://www.hitl.washington.edu/publications/r-93-9.html}"
    }
```