LEARNING BY EXPLAINING: THE EFFECTS OF SOFTWARE AGENTS AS LEARNING PARTNERS

By

Jeffrey T. G. Holmes

Dissertation
Submitted to the Faculty of the
Graduate School of Vanderbilt University
in partial fulfillment of the requirements
for the degree of
DOCTOR OF PHILOSOPHY
in
Teaching and Learning

December, 2003
Nashville, Tennessee

Approved by

Robert D. Sherwood
Deborah W. Rowe
Xiaodong Lin
Charles K. Kinzer
John D. Bransford

ACKNOWLEDGEMENTS

I have been very fortunate over the five years of my graduate work to have been the recipient of support, guidance, and encouragement from a large number of friends, colleagues, and family. Their generosity and wise words have helped me in innumerable ways and made my experience at Vanderbilt both memorable and rewarding.

Much of my research work was guided by Xiaodong Lin who both mentored me and accepted me as a colleague. Although I was, on occasion, accused of being somewhat of a stubborn learner, I am hopeful that her perseverance and toughness helped instill in me a few of the qualities necessary to do good research.

In typical fashion, Bob Sherwood generously gave of his time and energy throughout my graduate program and provided me with several opportunities that would not have been possible otherwise. I am also appreciative and thankful for the advice and careful thought of Chuck Kinzer, Debbie Rowe, and John Bransford, all of whom challenged me to think more deeply about my work.

I owe my friends and colleagues who worked and played along side of me through graduate school a great deal. They made my life as a graduate student an incredibly rich and fulfilling experience that will forever influence how I think, work, and live. In particular I want to thank Sashank Varma, Joan Walker, and Jay Konopniki for their marvelously diverse discussions and debates over endless cups of Fido brew.

If it is true that family is the greatest influence in the life of a person, then I am indeed a lucky individual. My parents provided my siblings and me opportunities and experiences that most people only dream about and their dedication to lifelong learning has inspired me to continually look for ways to broaden my education. I am also incredibly fortunate to have the encouragement and unflagging support of my brother and sister, to whom I owe so much.

Finally, I have had the great fortune of meeting my future wife during graduate school. Her unwavering support and complete faith in my abilities has been a powerful source of comfort and strength for me even though we have been separated by thousands of miles. As I look ahead to a new set of opportunities and adventures, I am content in the knowledge that I will be able to share them with a most remarkable partner in learning and in life.

TABLE OF CONTENTS

Appendix

LIST OF TABLES

LIST OF FIGURES

CHAPTER I

INTRODUCTION

Research over the last fifteen years has provided solid evidence that generating explanations can lead to deeper understanding when learning new material. There remain, however, many unanswered questions about the factors that influence the effectiveness of an explanation. One of those factors is the presence of a learning partner. This dissertation explores the interactions between learning partners and their effects on the quality of explanations generated in a virtual river ecosystem environment.

Although it is possible to generate explanations in isolation, such as when self-explaining, there are many reasons why it is important to study this activity as part of a social learning process. The inclusion of learning partners provides opportunities to influence not just the quality of the explanations but also the goals and methods of learning. For example, a learning partner in the role of an expert can act as a source of knowledge and provide examples about how to explain. The nature of a learning partner can also influence the goals of an explainer and whether the explanations generated are superficial and fact-oriented or deep and connected to underlying principles. Finally, partners in learning can have an impact on self-regulatory processes such as comprehension monitoring and focusing on the specified task.

The remaining sections of the introduction provide an overview of the purpose of this study, a description of the research questions, and an outline of the theoretical framework that guided the study design.

Purpose

The purpose of this dissertation is to extend explanation theory and inform learning environment design by examining the influence of different types of learning partners in a middle school science activity. In the study, programmable software agents were used as partners in an advisor role, an approach designed to address the specific goal of providing support to the student explainers. Support is an important issue because although the promise of improving learning through explanations is appealing, the reality is that the process of generating an effective explanation is neither intuitive nor quickly learned by students. Prior research has found that students who are not adequately trained or scaffolded in generating explanations do not construct explanations in ways that engage them in complex knowledge building (Coleman, 1998; King, 1994).

One reason for poor explanation outcomes in novices is that the effectiveness of the explanation depends on how it is framed. Novice explainers usually adopt a descriptive perspective that is geared more to answering questions such as "What is…" rather than deeper, integrative questions such as "How are … and … related?" or "What would happen if…" (King, 1994). Students do not naturally ask themselves or others the kind of questions that lead to more thoughtful explanations. More importantly, students are often not encouraged to question their own understanding when they explain. The main goal of this study is to explore how learning partners might be used to address these challenges by offering interactive advice and examples in a structured format. The study takes advantage of the capability of a virtual world environment to provide agents that can be programmed with specific characteristics and personalities.

A secondary goal of this study is to explore ways in which a summary report of the explanations generated during the experiment might be useful to teachers as a record of student thinking.

Research Questions

The research questions were designed to help focus the study on the effects of learning partners during an explanation activity. Of specific interest was how a learning partner might act as an advisor in the process of generating explanations and what kind of influence they might exert on novice explainers. It was hypothesized that software agents with human-like personas will encourage students to better attend to and utilize advice about how to explain. In light of this, the following research question was asked:

- Can an animated software agent advisor increase the use and value of explanation resources over a text-based help system?

In addition to providing expert advice, software agents can also take on the role of someone who can listen to an explanation and provide feedback, albeit in a limited form. Evidence from previous research suggests that there may be advantages in having a partner to explain to as opposed to simply generating self-explanations (Coleman et al., 1997; Web, 1989). What is unknown, however, is whether the nature of the partner has any effect on what or how students explain. Another goal of this study was to investigate the differences between software agents and student peers as partners who listen, give feedback, and generate their own explanations too. One possibility is that students will engage software agents in a more focused manner than a classmate who might be more inclined to encourage off-topic conversation. The following question addressed this issue:

- Can an animated software agent evoke more focused and thoughtful explanations than another student as a learning partner?

Finally, the ability of the computer to record explanations might also be beneficial from a teacher's perspective. A pervasive problem with educational software programs is that they provide little in the way of feedback that might be useful to the teacher. If a teacher had access to a summary report of their students' explanations, they would be in a better position to address specific problems and identify areas that might need further discussion or clarification. To explore these possibilities, a third research question was put forward:

- How might a summary report of student explanations be useful to teachers?

Theoretical Framework

A theoretical framework acts as a guide for designing both experiments and learning environments. In broad terms, the designs of this study and the learning environment on which it is based are influenced by the constructivist view that knowledge is meaning. This relationship between knowledge and meaning is significant because it implies that in order to learn, individuals must encode new information in a way that is personally meaningful. At its core, this study is about how individuals learn new concepts and new ways of thinking through the application of meaning.

Individuals, however, do not exist in isolation. This point was emphasized by Vygotsky (1978), who recognized that meaning is applied within the bounds of social and cultural practice. Since practice is to a large extent mediated by language, the use of language has a direct effect on the formation of meanings. An important aspect of learning by explaining is that it is an activity that encourages students to operationalize language. This concept has been increasingly promoted for science education by theorists such as Lemke (1990) who claim that students need to "talk science" in order to understand the field in a meaningful way. In this study, language plays a primary role as students think about and connect ideas through explanations.

A more focused constructivist perspective is put forward by Chi (2000) to describe the cognitive processing that occurs when one explains. Chi brings together evidence from several years of work on self-explanations to support an argument for mental model revision. Chi (2000) proposes that the act of self-explaining causes students to recognize gaps or conflicts in their mental models and therefore creates opportunities for mental model revision. From this theoretical standpoint, learning is very much the responsibility of the learner. Chi et al. (1994) have demonstrated that significant learning gains can be made when students are simply prompted to think more about whether new material makes sense to them and how it relates to what they already know. The key is that learner's thoughts are directed internally - not at an external and less personally relevant space. To clarify this point, note how the following prompt (adapted from Chi, 2000) directs the learner to think inwardly:

Does this new information give you insight into your understanding of how the circulatory system works?

The phrases "give you insight" and "your understanding" highlight the metacognitive nature of this prompt. This format is very different from the type of question that simply asks students to summarize a paragraph, for example. One can easily summarize material without ever contemplating how it fits with one's own mental model. An externally-oriented question does not encourage students to construct personal meanings. Consequently, the design of the explanation prompts used in this study were guided by a goal of inducing personally meaningful thought.

Another important implication of this theoretical perspective is that support structures do not have to rely on an extensive understanding of the learner's state of knowledge. They simply must assist the learner to think inwardly. This suggests that using software agents that do not possess sophisticated abilities such as diagnosing misconceptions is a reasonable approach. In this study, the agents were programmed with the main goal of helping students explain in a way that encourages them to construct and reflect on meaning.

Virtual worlds that support software agents as avatars (personal representations) open up new possibilities for interactive learning partners. Theories on persuasion suggest that learners will more likely accept and incorporate into their own thinking arguments from partners who are more credible and understandable (McGuire, 1985). Therefore, it is expected that software agents in the form of avatars will be more effective in helping students learn about explanations than plain text. Although research is limited, there is some evidence that animated agents can be more influential than text and lead to greater understanding when acting as discussion partners (Burgoon et al., 2000).

In summary, this study is based on an explanation activity that promotes the construction of personally relevant meanings. Students are supported by software agents that possess human-like features designed to increase authenticity and interest. As both advisors and participants, it was hypothesized that these agents would provide a more focused and motivating experience.

Importance of the Study

A dissertation in educational technology should aim to make a theoretical contribution as well as provide insights to designers using technology to improve education. Studying learning partners in an explanation activity provides such an opportunity because explanation theory lacks a complete picture of how social interaction affects not just what is explained but also how an explanation is generated and even the reason for making an explanation. This study represents one of the first attempts to manipulate the social aspect of learning by explaining through the use of software agents acting as advisors and participants in an explanation activity.

Designers of learning environments are in the enviable position of selecting from an ever increasing array of tools that present new opportunities to enhance learning. With these choice, however, come the challenge of implementation. Software agents are one example of an exciting new technology for education but little is yet known about how to incorporate these versatile instruments into pedagogically sound software. It is hoped that insights from this study will help guide the development of effective new education environments.

Overview of Chapters

The following chapter is a review of research literature on learning by explaining with a focus on support mechanisms and the nature of learning partners. Chapter three provides a detailed description of the methodology and design for the study. Results, general discussion, and conclusions complete the final two chapters.

CHAPTER II

REVIEW OF THE LITERATURE

Introduction

The goal of this review is to inform the design of the experiment described in the following chapter through an analysis of research on learning by explaining. Examining the methodologies and outcomes of prior studies is one way to identify key components of the explanation process and suggest areas where a technological intervention might provide distinct benefits. Although most studies focus on the learning outcomes of a particular method of explaining, a literature review provides an opportunity to gain insight into the learning process itself by comparing outcomes across a range of studies. The challenge is to develop an analytic framework to highlight specific areas of explaining that can be affected by learning partners. For the purposes of this review, this framework is build around the nature of resources and support that have been used in explanation studies and on the specific roles of learning partners.

The scope of this review covers studies published in peer-reviewed journals that have featured learning by explaining as a major component. While some work on peer interaction (e.g. Skon et al., 1981) and cooperative learning (e.g. Slavin, 1983) certainly pre-dates many of the studies discussed here, it was not until the late 1980s and early 1990s that researchers focused more acutely on the role of explanations in learning. This began with the work of Chi et al. (1989) on self-explanations and has continued through to the present with researchers expanding the field to include strategy training (Bielaczyc et al., 1995), collaborative explaining (Coleman, 1998), and peer tutoring (King, 1998) among others.

The next section provides an introduction to learning by explaining and includes some definitions and an overview of theoretical development. This is followed by an analysis of support strategies and how they have been implemented in explanation activities and on the role of learning partners. A discussion of the implications for learning environment design and suggestion for future research concludes the chapter.

Learning by Explaining

Explanations are commonly thought of as the stuff of textbooks and lectures. Typically, it is considered the role of the expert or teacher to provide explanations to their students (Dagher & Cossman, 1992). However, there is a widespread intuitive belief supported by a growing number of empirical studies suggesting that teaching, or explaining to others, can be a powerful medium for learning (e.g. Bielaczyc et al., 1995; Bransford et al., 1999; Coleman, 1998).

Despite general acknowledgement about the potential benefits of student generated explanations, there have been surprisingly few research studies and classroom activities developed to explore this idea. This is especially true with regards to technology-based learning environments. Shared "knowledge building" environments usually focus on the construction of a collaborative artifact. Group work is often considered to support explanatory activities but in many cases it simply results in a division of labor exercise rather than an opportunity for peer teaching (see Cohen, 1994 for issues regarding positive goal and resource interdependence). Because of the many possible ways to interpret explanation-based learning, the following discussion attempts to constrain the concept for the purposes of this paper.

A Definition of Explaining

Teaching, and explaining in particular, are multifaceted concepts that are discussed in a wide variety of contexts. The definition of "explaining" for the purposes of this paper is borrowed partly from Chi's (2000) discussion of the term with respect to self-explanations, although a broader interpretation is adopted here. First and foremost, explaining is considered a constructive activity in the sense that it requires the explainer to bring together knowledge in new ways in the form of an explanation. It is not summarizing because there is a component of addition or "filling in" of material that is often tacit or implicit. Chi (2000) differentiates self-explaining from elaborating with the claim that elaboration implies a process of generating relationships between concepts but without necessarily making them meaningful. Therefore, explanation may be considered a form of elaboration but one with a specific goal of sense-making. Another important characteristic of this activity is that it incorporates prior and current contextual knowledge. Finally, it is recognized that there may be some important differences between explaining to oneself and explaining to someone else, as Chi (2000) notes. These distinctions will be discussed in more depth later. In sum, the act of explaining is an intentional learning process that involves the explicit generation of new and meaningful relationships among knowledge structures.

Theoretical Development

Cognitive psychologists have theorized about what basic cognitive mechanisms might be responsible for improved learning through teaching or explaining. Slavin (1995) uses the term *cognitive elaboration perspective* to describe the approach that researchers adopt to investigate this form of learning. Siegler (1995) groups explaining with other activities that promote "meaningful learning" such as analyzing problem semantics, defining problem solving goals, and determining problem solving strategies. Siegler refers to deep processing, elaboration, structure mapping, mental modeling, and metacognitive analysis as constructs that describe meaningful learning and suggests that the underlying key is inducing the learner to "think about the task more deeply than usual" (p. 265). Other researchers have designed studies that aim to reveal more precisely the cognitive processes involved in learning by explaining. Chi et al.'s (1989) study on self-explanations is widely recognized as the first significant work to reveal the role of self-explanations during the study of worked examples and solving related problems.

The work of Chi and her colleagues (1989, 1994) on self-explanations was based on the conjecture that effective problem solving depends on the completeness of the encoding process. While some theorists (e.g. Anderson, 1987) highlighted the importance of the conversion of already-encoded declarative knowledge into useful procedural knowledge, Chi et al. (1989) believed that a crucial component of learning (at least in the sense of converting text and examples into usable skills) involves learner dependent processing on the front end, during the initial encoding stage. Anderson's ACT* theory (1987), on the other hand, assumed that encoding is rather straightforward and automatic, and that the key learning stage occurs later during the compilation of encoded declarative knowledge into procedural knowledge.

Chi et al. (1989) hypothesized that the differences between good and poor learners who studied worked-out examples was that the good students were learning the examples with understanding (i.e. their encoding process was richer) while the poor students simply working through them. In order to examine this process in detail, the 1989 study analyzed verbal protocols of college students who studied worked examples and then solved problems in the domain of Newtonian physics. Chi et al. (1989) based their analysis on an "idea hierarchy" that broke down the ideas that students generated into three main categories: explanations, monitoring statements, and others.

In their analysis Chi et al. (1989) split subjects into either a Good or Poor group based on posttest performances. Good students were found generate more statements in each of the three idea categories during the worked-example study period. Further, the Good students generated a higher proportion of explanations that contained inferences involving the principles and concepts of the domain. Chi et al. (1989) proposed that because the text materials omitted a substantial amount of information, students needed to generate these inferences in achieve a deeper level of understanding. They concluded that generating inferences was the key factor because it ensured that incomplete text would be encoded in a richer, or more complete, manner.

The results from Chi et al's 1989 study provided a basis for a theory of learning by explaining. The study suggested an interactive process involving the perception of new knowledge and the generation of elaborative explanations that provide additional information through the production of inferences. An important implication of this theory is that typical well-structured learning resources provide only some of the "knowledge pieces" that are required for deep understanding. A learner must internalize the resource and create new connections between ideas through the generation of an inference. The act of explaining both generates this new inference and makes it verbally explicit.

The previous discussion suggests that in order to learn the material more deeply, the student must construct a new link to bring more meaning to their mental model. In their physics study Chi et al. (1989) provide the example, "Ummm, this would make sense, because they're connected by a string that doesn't stretch." (p. 161). In this case, the student has generated an explanation that adds new information to what was provided in the example - namely that the link between the two objects referred to in the problem is non-elastic.

Other researchers extended the self-explanation work initiated by Chi et al. (1989). Pirolli and Recker (1994) replicated the previous findings in the domain of Lisp programming and showed that the acquisition of programming skills was improved when explanations linked instructional examples to abstract concepts. Other studies in the domain of electricity (Ferguson-Hessler & de Jong, 1990) and mathematics (Nathan, Mertz, & Ryan, 1994) also positively correlated problem solving success with the number of self-explanations generated.

Eliciting Self-Explanations

Chi, de Leeuw, Chiu, & LaVancher (1994) performed an important follow-up study on self-explanations with eighth grade biology students to investigate whether the benefits of self-explanations could be realized if students were prompted to self-explain. Their study confirmed that prompting did lead to more generated self-explanations although individual difference were still present. Since closer analysis of the self-explanations revealed that 30% were produced by integrating the new information with prior knowledge, it was hypothesized that the variability within the prompted group might be related to either differences in prior "general world" knowledge or strategies for integrating new knowledge. Early work implied that the materials used in the study were basically operated on and "filled-in" during learning before being stored in long-term memory. The biology study highlighted the contribution of prior knowledge and demonstrated how it is integrated with new information when students learn with deeper understanding.

With mounting evidence supporting the notion that self-explanation and self-regulation play a key role in learning, Bielaczyc, Pirolli, and Brown (1995) undertook a study to investigate the causal role of these strategies in problem solving performance. The work was part of an ongoing effort to understand complex problem solving in the domain of computer programming (Lisp). Strategy training interventions were developed in order to teach students, through modeling and scaffolded instruction, several types of self-explanation strategies designed to identify main ideas and connect underlying concepts. In addition, students were trained to monitor their comprehension, application of self-explanation strategies, and progress towards stated goals. Their results showed that not only was training to self-explain possible, but that the specific strategies used in the training were associated with improved performance.

Types of Explainers

Renkl (1997a) extended the work of Chi et al. (1994) and Pirolli and Recker (1994) by studying a larger sample size (36 subjects) and by looking more closely at the importance of specific self-explanation characteristics and the effects of prior knowledge. Renkl (1997a) described 7 categories that were used to analyze think-aloud protocols as students studied the worked-out examples. These are displayed in Table 1 along with categories used in the studies by Chi et al. (1989) and Bielaczyc et al. (1995) for comparison. Renkl's study is important because it provides evidence to suggest that effective self-explaining consists of multiple non-correlated factors. In other words, there are varying levels of effective self-explanation. One student may, for example, employ situation elaboration and notice example coherence but not make use of anticipative reasoning. Renkl (1997a) identified four types of learners in his study based on the relationship between the quality (characteristics) of their explanations and learning outcomes. Anticipative and Principle-based learners were successful learners that employed anticipative reasoning and principle-based explanations respectively. Passive learners noticed comprehension failures (self-monitoring) more than other groups but like the fourth group termed Superficial, they did not utilize anticipative reasoning and principle-based explanations and were less successful on posttest problem solving. This study also showed that anticipative reasoning and principle-based explanations are predictive of medium transfer task success (defined as items with different surface and deep structures from the worked examples).

Table 1. Categories for Analyzing Self-explanation Protocols

Self-Explanation Protocol Analyses		
Chi et al. (1989)	Bielaczyc et al. (1995)	Renkl (1997a)
References to domain principles Impose goal or purpose for actions Refine or expand conditions of an action Monitoring-negative Monitoring-positive	Determining the form and meaning of examples Identifying and elaborating main ideas Connecting concepts between the texts and examples Monitoring-general Monitoring-comprehension Monitoring-responses to	Principle-based explanations Goal-operator combinations Anticipative reasoning Elaboration of situation Noticing coherence among examples Monitoring-negative Monitoring-positive

Mental Model Revision

Work through the 1990s contributed to an evolving model of the cognitive processes involved in self-explanation. Chi's earlier work (1989, 1994) assumed that these processes were based on inference generation. The original hypothesis was that the inferences would fill in gaps in the students' mental models. Chi (2000) undertook another study of self-explanations that looked more closely at the self-explanation process of one student studying text materials about the circulatory system. Based on this work and a re-analysis of previous findings, Chi modified her self-explanation hypothesis to one involving a process of mental model revision. This change was driven by findings that were inconsistent with the theory that generating inferences was the only mechanism behind self-explanations. One such inconsistency was that self-explanations did not occur uniformly throughout the studying of the text materials. Chi (2000) expected that, since information is omitted in a uniform manner within the text, there should be a uniform occurrence of self-explanations. This was not the case. Further, the self-explanations that did occur varied in location between students. Second, the inference theory predicted that self-explanations should be semantically similar between students (when they occurred at the same location) since they would be "filling in" the same information. This was also disproved. Finally, the form of the explanations themselves was inconsistent with a theory of inference generation. Analysis revealed that, unlike elaborations, the self-explanations were often non-scientific, fragmented, and sometimes simply incorrect - and yet still helpful for learning. Part of what this suggests is that the self-explanations are, in fact, serving a purpose other than just inferences. As Chi (2000) points out, these inconsistencies in the "when" and "what" of self-explanations suggest that they are somehow more closely related to the individual. This move to the idea of revising one's mental model meant that self-explanations needed to be reinterpreted in the context of a learner's prior knowledge rather than the material contained within the text itself.

Summary

Initial work on the self-explanation effect provided some valuable insights about how individuals encode new information. The research suggests that learning by explaining involves a process of mental model revision where the integration of new and prior knowledge is mediated by self-monitoring and cognitive strategies. What is striking about this form of learning is how almost all of the key components (with the exception of the learning materials and task itself) are dependent on the learner. Learners must possess adequate monitoring and cognitive skills in order to successfully construct and revise useful mental models. Even though prior domain knowledge is not necessarily required, it appears that successful learners depend on making connections to "general world" knowledge when attempting to give meaning to new information. An important implication is that even providing specific instruction about what is essential to know will not be effective unless learners encode new material in a way that is meaningful to them personally.

The study described in the next chapter examines how software agents might be used to support an explanation activity. Insights from this review of theoretical work on explaining suggest that the agents must engage the students in such a way as to promote the use of monitoring and cognitive strategies. Further, it is important that the students use this strategies within a context that is personally meaningful and relevant. The next section examines how prior researchers have supported explanation activities and identifies specific monitoring and cognitive strategies.

Support for Explainers

There is general agreement that learning by explaining can be made more effective for novices when certain forms of support are provided (Bielaczyc et al., 1995; Chi, 1994; Coleman, 1998; King, 1998). This section examines the various ways in which researchers have supported the explanation process with the goal of identifying key components. The categories, shown in Table 2, reveal the wide range of options available to designers of support systems.

Table 2. Types and Presentation Format for Support Strategies

	Presentation Format		
Support Strategy	Resource	Expert	Learning Partner
Instructional video	√		
Modeling		√	
Classroom training		√	
Cues and prompts	√	√	√
Feedback			√

Note that while some forms of support involve pre-explanation measures, such as an instructional video, most support strategies are designed for use during the actual explanation period. There is also a large variation in how the support is presented. In some cases support is offered in the form of a resource, such as written instructions, while in other cases help might come from an expert (e.g. teacher) or a learning partner.

The Goals of Support Strategies

Researchers have found that most students do not spontaneously generate effective explanations whether self-explaining (Chi, 1989; Renkl, 1997b) or when explaining to a partner (King, 1998). A novice explainer tends to generate descriptive or fact-based explanations that do not lead to deeper understanding (King, 1998). One reason why these superficial explanations are not effective is that they do not involve higher forms of thinking, such as anticipative reasoning (Renkl, 1997a). Superficial explanations also do not make underlying principles explicit. Therefore, an important goal of many support strategies is to help students frame their explanations in such a way as to promote higher levels of thinking.

Cognitive Framing

I use the term cognitive framing to refer to the use of strategies that promote the kind of thinking that leads to deeper understanding. These cognitive strategies are deliberate, generally domain-free ways of thinking that act as catalysts in the construction of new knowledge. Cognitive strategies are essential to developing deeper understanding through "knowing why". Because many traditional learning contexts focus on "knowing what", students perform well on tests of recall and recognition but are unable to complete tasks that require knowledge transfer. Table 3 below highlights the general principles behind these strategies.

Table 3. Principles of Cognitive Strategies

Principles of Cognitive Strategies
• Identify new information • Identify relationships between new ideas • Relate new ideas to prior knowledge • Impose a goal on the new content • Give meaning to new content • Anticipate outcomes

Researchers have employed various methods to help students understand and use cognitive strategies. Coleman (1998), for example, conducted a study involving 48 middle school students working in small groups of three with a unit on photosynthesis. Twelve students were identified as "high intentional learners" (HIL) based on responses to a modified Implicit Learning Theory Interview (Steinbach, Scardamalia, Burtis, & Bereiter, 1987). These students scored 80% or higher in the interview and were assigned to the HIL group. The other 36 students were divided equally and randomly into an average control group (AC) or an average intervention group (AI). The materials consisted of a four chapter unit on photosynthesis and instruction consisted of lectures, class discussions, and group activities (problem solving and experiments) over a 4 week period. Explanation scaffolds, in the form of questions or prompts were given to the students in the intervention group. The scaffolds were adapted from Scardamalia, Bereiter, and Steinbach (1984) and Scardamalia and Bereiter (1985) and designed to facilitate higher level scientific thinking processes. An example is "Can you explain this using the scientific information that we learned in class?" (Coleman, 1998, p394).

Results showed that the AI groups performed significantly better than the AC groups and as well as the HIL groups in both the posttest and concept-mapping tasks. Perhaps the most interesting results from the study, however, were the findings on student explanations. Students in the AI groups constructed conceptually more advanced explanations than the AC groups. Since there was no significant difference between the AI and HIL groups in this measure, it suggests that better students are composing these explanations as a matter of routine - a finding that agrees with the work of Chi (1989) on self-explanations.

Another study that utilized cognitive strategies involved undergraduate students learning Lisp programming (Bielaczyc et al., 1995). In preparation for a problem solving exercise, one group of students were taught to: identify and elaborate main ideas in texts; determine the form and meaning of Lisp code examples; and connect concepts between the texts and examples. Bielaczyc et al. (1995) showed that students in the intervention group used the strategies and performed better on the problem solving exercise.

Cognitive strategies have also been used successfully with peer tutoring approaches, such as the one developed by Palincsar and Brown (1984). Their model, known as Reciprocal Teaching, aimed to improve how students learn from texts. The format was based on individual members of a group taking turns to read text passages and then to invoke various teaching strategies. Palincsar and Brown (1984) hypothesized that there are general comprehension strategies that can be taught and that these are transferable to other learning situations. In developing reciprocal teaching, they identified four concrete strategies that are key to comprehension: summarizing, questioning, clarifying, and predicting. Results from their work showed an increase in reading comprehension and reliable maintenance of strategy use over time.

Self-monitoring

One of the reasons that Palincsar and Brown (1984) chose the particular strategies mentioned above is that they could also be used for self-monitoring as well as facilitating comprehension. Promoting self-monitoring or metacognitive skills is the other major goal of many explanation support strategies.

Metacognition, or knowing about knowing, plays a crucial role in understanding when and where to apply specific ways of thinking. Table 4 summarizes self-monitoring skills that have been incorporated in various learning by explaining studies.

Table 4. Self-monitoring Skills

Self-monitoring Skills
• Identify personal learning goals • Identify comprehension successes • Identify comprehension failures • Identify knowledge gaps • Identify shifts in perspective • Develop knowledge of self as learner

Citing Campione and Brown (1990), Bielaczyc et al. (1995) point out that the use of self-monitoring skills increases the ability of students to "understand and take control of their own learning process" (p. 224). This understanding appears to take place on three levels. First, students are able to assess their comprehension of the immediate problem at hand. An example would be recognizing a comprehension failure. Second, on a more macro level, the problem can be situated within a larger "world view". This requires that students reflect on how what they learn might lead to shifts in perspective. For example, learning the specifics of a local economy can influence how students think about world trade. The third level of understanding involves knowing about the self as learner. Lin (2001) argues that students' knowledge about themselves as learners is closely linked to social context because of the fluidity of self-concept within a social environment. The implication is that development of positive self-knowledge depends on how learners perceive their role within the learning community. This is discussed in more detail later in the section on learning partners.

As with cognitive strategies, researchers have used a variety of methods to help students adopt a metacognitive approach. In their study involving students learning about human biology, for example, Chi et al. (1994) used a brief verbal instruction as a prompt for self-explanations. It is worth reproducing their instruction here to demonstrate the degree to which it directs students to reflect on their own knowledge:

> *We would like you to read each sentence out loud and then explain what it means to you. That is, what new information does each line provide for you, how does it relate to what you've already read, does it give you a new insight into your understanding of how the circulatory system works, or does it raise a question in your mind? Tell us whatever is going through your mind - even if it seems unimportant.* (Chi, 2000, p. 15).

Other researchers have developed more elaborate means of helping students self-monitor. Palincsar and Brown (1984) incorporated a teacher modeling component to guide the development of sophisticated comprehension fostering and monitoring activities. In their system, an adult teacher initially models appropriate actions such as asking clarification questions or summarizing what was just stated. As students become familiar with the procedure, the teacher gradually becomes less involved and the students take on more teaching responsibilities.

In another approach taken by Bielaczyc et al. (1995), an instructional video was included as one component of a three part training session. In the session the experimenter gave each student a one-on-one introduction to cognitive and self-monitoring strategies. This was followed by the instructional video which modeled the use of the strategies. Finally, the experimenter verified that each student understood and could apply the strategies to a set of instructional materials.

Summary

Despite the wide variety of approaches taken to support students generating explanations, most researchers agree that students need to be taught to understand and use both cognitive and metacognitive strategies. Cognitive strategies are intended to promote deeper ways of thinking by having learners identify underlying principles and connect ideas in meaningful ways. Explanations that incorporate these strategies tend to answer "why" and "how" questions as opposed to those that are more superficial and descriptive. Self-monitoring or metacognitive strategies encourage learners to think inwardly and consider their own state of knowledge. Reflecting on their own knowledge helps students learn not just about what they know but also about themselves as learners.

The next chapter describes how software agents can be used to promote the use of monitoring and cognitive strategies within a virtual learning environment. This section of the literature review provides the framework and specific strategies for designing prompts to be used by the agents. In this way, the agents can be situated as “explanation experts” and are able to both prompt students and model correct strategy use.

Learning Partners

Learning partners have played an important role in many interventions that incorporate learning by explaining. Although there have been several examples of successful self-explanation activities (Chi, 1994; Pirolli & Recker, 1994; Renkl, 1998), the inclusion of a partner as "explainee" provides certain advantages and opportunities that are not available to students who simply explain to themselves.

Roles of Learning Partners

The range of involvement of a partner extends from merely being anticipated through to a fully interactive participant who provides support and feedback to the explainer. Table 5 below outlines the role and activities of learning partners that have been used in explanation studies. The roles are ordered from least interactive (i.e. not present) to most interactive where the partner provides prompts and feedback to the explainer.

Table 5. Learning Partner Roles as Explainees

Learning Partner Roles	
Involvement	**Actions**
Anticipated	None (not present)
Passive	Listen only (no intentional feedback)
Interactive	Listen Encourage Prompt Question

Anticipated Learning Partners

The effects of anticipating an explainee have been investigated in a limited number of studies. In the early 1980s, Bargh and Schul (1980) and Benware and Deci (1984) conducted experiments to determine whether the expectation of having to teach someone else would improve learning.

In the first of two experiments, Bargh and Schul (1980) compared two groups of undergraduate students studying texts in the domain of psychology. One group was told to study the materials in order to answer a set of questions while the other group was told that they would be teaching other students to be able to answer the questions. A monetary incentive based on the final test performance (of the tutees in the case of the second group) was used. Once the study session was complete, the students in the preparation for teaching group were told that their tutee had not shown up and that they would take the final test themselves. Results showed that the students who studied in order to teach another performed significantly better on recall and recognition items than students who had studied for themselves.

Coleman, Brown, and Rivkin (1997) added to these findings in a study involving undergraduate students learning evolutionary biology material. In their study, 83 students were randomly assigned to six groups described as follows. A Teach condition involved two subgroups. One group was told to learn the material with the expectation of having to explain it to a partner and explicitly requested to make inferences for any information that the text left out. The other teaching group was told to concentrate on summarizing the main points of the text and to not worry about details or make inferences. Two Self conditions were also divided into either

Summarizing or Explaining conditions but in this case they were simply told to summarize or explain aloud to the experimenter after they had studied the text. Finally, there were two Hear conditions that did not read the text. These participants were told that they would be answering questions after they listened to either an explanation or summarization from their partner. The results of the experiment revealed that while the Explanation and Summary groups did not differ in measures of text comprehension, the Explainers performed significantly better than the Summarizers and Hearers in a far transfer problem solving task. Participants in the Explaining to Teach condition performed significantly better than the Self-explainers or Hearers across all of the far transfer problems.

The results of the Coleman et al. (1997) study indicate that an important issue raised by at least two authors (Chi, 2000; Plötzner, Dillenbourg, Preier, & Traum, 1999) regarding whether there is a distinction between self-explaining that involves verbalizing to an experimenter and explaining to another student. Plötzner et al. (1999) point out that students instructed to self-explain aloud to a researcher may construct more careful explanations than they would under normal study conditions. Coleman et al. (1997) found that explaining with the intent to teach another was more beneficial than explaining for one's own learning. They reasoned that the expectation of teaching another encourages the explainer to identify misunderstandings and to reorganize the material so as to be more easily comprehended by the learner. This more careful processing leads to a richer and more complete mental model.

Benware and Deci (1984) hypothesized that preparing to teach others would facilitate greater intrinsic motivation that preparing for a traditional test. Benware and Deci (1984) also attempted to address the issues of short study duration and rote learning assessment that arose from the Bargh and Schul (1980) study. Forty-three college students were instructed to study an article on brain functioning over a period of 2 weeks. The 22 students in the control group were told to use whatever methods they preferred and to spend about 3 hours total on the task in order to prepare for a test. The remaining 21 students were placed in the experimental group and told to study the article in order to be able to teach another student enough to take a test on the material. After the 2-week period, students in the experimental group were told that they had been selected at random to take the exam before it was administered to their own students. The test was designed to measure both rote memory and conceptual understanding. All participants also completed a questionnaire on intrinsic motivation before taking the test.

Results of the Benware and Deci (1984) study supported the hypothesis that preparing to teach others would promote a higher degree of active learning and intrinsic motivation (as measured by interest, enjoyment, and participation responses) and greater conceptual understanding of the content.

Results from some pilot work on teachable software agents by Biswas, Schwartz, Bransford, and the Teachable Agents Group at Vanderbilt (2001) provide some insight into why anticipating having to teach someone else has benefits on learning. The researchers found that students studying a psychology text on memory experiments with the understanding that they would need to teach the material to rest of the class spent twice as much time studying the material as a control group who were told simply to prepare for a test. Additionally, the data indicated that students were more likely to try to grasp the "why" of the experiments. This was seen in attempts to consider the larger context of the studies, by questioning their purpose, and by noticing flaws or alternative approaches.

It seems clear that students who prepare to teach or explain to someone else have increased intrinsic motivation and are encouraged to think more deeply about the learning materials. The next section explores the nature of learning when a partner is present but in the limited role of a listener.

Learning Partners as Listeners

Although having a listener is somewhat similar to the self-explanation activities that require students to verbalize their thinking, the presence of an audience may provide some benefit from what has been variously referred to as the "audience effect" and "social facilitation". For example, in a review of social facilitation studies, Zajonc (1965) concluded that the presence of others might have positive motivational effects. Although a number of studies have investigated the effects of explaining to others with no intentional feedback, many involve a preparation stage making it impossible to partial out the effect of a listener being present from that of preparing to teach. For example, it is unclear from the Coleman et al. (1997) study described above how much of an effect explaining to a live audience had on learning since their study did not separate out preparing to explain from the actual explanation experience. One study that does make this distinction was conducted by Bargh and Schul (1980).

Bargh and Schul (1980) contrasted students working alone silently, working alone and thinking aloud, and teaching another while working on a Tower of Hanoi problem and learning from text. Performance was measured on the Tower puzzle by calculating the ratio of minimum possible moves to actual moves. The text passage scores were derived from recall and recognition questions. Although Bargh and Schul (1980) failed to obtain significant performance differences between the conditions, the researchers speculated that significant performance differences on the Tower of Hanoi problem might have been prevented by the short duration of the teaching time (8 minutes per session) and artificial context. While lack of teaching time may indeed have contributed to the negative results in the Bargh and Schul (1980) study, anxiety may have also contributed to the problem.

Renkl (1997b) investigated motivation and an anxiety in two experiments involving worked-out examples and text materials. In the first study, students were placed in dyads where one member was assigned the role of explainer and the other the role of listener. Data were collected from a questionnaire on motivation and anxiety and from performance on a posttest. Surprisingly, listeners achieved better learning outcomes and felt more motivated than explainers. Concerned that the experimental conditions may have somehow favored the listener role, Renkl (1997b) conducted a second identical experiment except that during the explaining phase, students were told to explain to each other in an unstructured cooperative manner. Similar to listeners in the first experiment, learners in the unstructured cooperative condition outperformed students in the explainer role. Renkl (1997b) concluded that these unexpected results might be due to several factors including increased stress on the part of the explainer, the one-way nature of information flow, or the dual task nature of the explainer's role (learning and teaching). It should also be noted that the learning phase of these experiments was limited to 25 minutes, which may have contributed to the explainers' poor results.

In summary, there is little evidence to support the use of learning partners in listener-only roles when students have limited time to prepare and teach. In studies where students do have adequate time to prepare, such as in Coleman et al.'s (1997) experiment, learning gains can be achieved by explainers but the benefits of a live listener remain unclear. Advantages of having a learning partner present become more evident in the next section where the explanation context is interactive.

Interactive Learning Partners

An important outcome of explaining to an interactive learner is that the process becomes one of communication rather than simply transmission. In education research this is usually discussed with the goal of avoiding situations where students are passive listeners of teachers' explanations. However, in this context, the focus is on the learner as teacher or explainer. Therefore, the aim is to encourage active participation from the listener in order to help the explainer.

Several researchers have explored how interactivity affects learning by explaining. In a review of peer interaction studies, Webb (1989) found that elaborative explanations were positively associated with achievement. Significantly, most of the suggested reasons for the beneficial effects are a direct result of an interaction between the explainer and explainee. Webb (1989) points out that the necessity to attend to the needs of a learning partner create opportunities for clarifying and organizing material, responding to feedback, and searching for new information. Indeed, the need to seek common ground, or co-construct knowledge, increases the likelihood of beneficial outcomes such as the production of abstract representations as reported by Schwartz (1995). The distributed nature of learning in interactive situations, however, makes it a difficult task to isolate the contributions of learning partners. As Webb (1989) notes, the interdependency of the learning relationship must be examined intact: "…student's learning cannot be understood in isolation from the group context" (p. 36). One context that provides some insight into this complex relationship is peer tutoring.

Graesser, Person, and Magliano (1995) conducted an informative naturalistic investigation with untrained graduate and high school tutors. Since most of the previous studies looking at the dialog during tutoring used relatively skilled tutors, these researchers sought to obtain a detailed account of discourse patterns and their relationship to learning outcomes with a more typical population of "normal" non-expert tutors. Prior work by Graesser and his colleagues (Graesser, 1992; Graesser & Person, 1994; Person, Graesser, Magliano, and Kreuz, 1994) identified various tutoring protocols (e.g. pedagogical strategies, feedback mechanisms, questioning, etc.). The 1995 study attempted to relate these protocols to eight components of modern learning theories discussed below. While this work was focused on the effectiveness of tutoring from a tutee point of view, it is useful because it reveals something about which components result from unskilled tutoring. Graesser et al. (1995) found that only some of the learning components were prevalent in normal tutoring situations. These included: anchoring learning in specific examples; collaborative problem solving and question answering; and deep explanatory reasoning. Deficient or absent components included: active student learning; sophisticated pedagogical strategies; convergence towards shared meanings; feedback, error diagnosis, and remediation; and affect and emotion. One conclusion that can be drawn from this work is that the learning gains achieved by tutors is not a function of developing shared meanings or diagnosing student misconceptions. The benefit for the tutor seems to be more related to generating explanations based on the tutor's notion of what needs to be taught. Therefore, from a tutoring perspective, the benefits of explaining to a live partner do not seem to depend as much on the specific feedback that an explainee might provide as simply having the goal to teach. So what advantages can an interactive learning partner provide?

Several researchers have tried to leverage advantages of interactive partners with the recognition that novice explainers require scaffolding. Alison King's work (1994, 1998) on peer tutoring, for example, is based on principles of sociocognitive theory. King describes peer-mediated learning as *transactive* in nature, where there is interdependence and reciprocity in the exchanges between the learners. For example, learners' responses may be guided or framed by their partner's comment or question. King (1998) claims that because peer mediated learning "depends upon mutual scaffolding and guiding" (p. 61) learners need to be taught these skills. King's (1994) study on Guided Peer Questioning revealed that students who were not trained to ask experience-based or lesson-based constructive questions tended to formulate questions that were fact-oriented. In turn, this created the situation were explanations followed a similar fact-based format. King found that students generating and answering fact-oriented questions were out-performed by students using more integration and comprehension questions. In other words, the type of questions asked tended to drive the responses and thus impact learning. King (1998) was most interested in promoting higher-level learning (e.g. constructing new knowledge, complex problem solving) as opposed to comprehension-level learning (e.g. reviewing material discussed in class). In order to foster interactions that mediate high-level learning (asking questions, elaborating explanations, allowing sufficient time for thinking, and supportive communication skills), King (1998) developed a model for peer tutoring named "ASK to THINK-TEL WHY".

Drawing from theories of distributed cognition (Pea, 1993; Perkins, 1993), King's (1998) model attempts to take advantage of the cognitions that are distributed between learning partners. King's model structures the tutoring process by training students to "ask thought-provoking questions and respond with elaborated explanations" (King, 1998, p. 63). This idea evolved out of previous work by King (1994) on Guided Peer Questioning where students were trained to ask and answer questions designed to solicit elaborated explanations. The model extends this work by addressing other aspects of the transactive process. For example, students are trained to provide supportive feedback such as verbal encouragement and positive body language (nodding, maintaining eye contact).

Efforts to link support strategies and learning partners has been promoted by other researchers as well. In Reciprocal Teaching (Palincsar & Brown, 1984) described earlier, students in the "teacher" role provide prompts the that enable the explainer think more deeply and monitor their own comprehension. In other words, learning partners help explainers to develop and use the kind of cognitive strategies and self-monitoring skills that were discussed in the section on resources and support.

Coleman (1998) further explored the idea of scaffolding collaborative learning by investigating the process and content of explanations during collaborative discourse. In her study involving a unit on photosynthesis, students were provided with explanation scaffolds in the form of questions or prompts. These resources, adapted from Scardamalia, Bereiter, and Steinbach (1984) and Scardamalia and Bereiter (1985), were designed to facilitate higher level scientific thinking processes in a similar manner as that undertaken by King (1998). Compared to a control group, Coleman (1998) found that explainers who received this kind of support from peers constructed conceptually more advanced explanations and performed better on a posttest and concept-mapping task..

Research on tutoring has downplayed the importance of the explainee's influence on the explainer, suggesting that striving for shared understanding per se may not contribute significantly to learning. However, it does appear that interactive learning partners can play a valuable role in explanation activities when they assume a supportive role. A summary of this section on interactive partners follows.

Summary

The influence of learning partners on explainers has been investigated from three levels of involvement. At the least involved level, where the presence of a partner is anticipated but never realized, there is good evidence that this can increase motivation and encourage learners to think more deeply about content. At the second level, where a partner is present but in the capacity of listener only, there appear to be few advantages over simply anticipating a partner. The third level of involvement, where a partner becomes an active participant in the explanation process, provides some significant opportunities for learning, especially for novice explainers. An interactive partner can provide prompting and feedback designed to help the explainer think more deeply about content. This type of support was referred to as cognitive framing in the previous section. In addition, interactive partners can support explainers by providing metacognitive or self-monitoring prompts.

An important conclusion that can be drawn from this analysis of research on learning by explaining is that an adequate support structure is essential. In cases where learning partners provide prompts and feedback to the explainer, the partner must have access to the proper support to ensure that the interaction is effective. In the study described in the next chapter, software agents are programmed to be “knowledgeable” about effective ways to generate explanations. Specifically, they have the ability to provide examples and specific prompts to demonstrate proper use of monitoring and cognitive strategies. The next section extends this discussion on possible roles for technology in learning by explaining and provides some further suggestions on how computers might be used to support explanations.

The Role of Technology in Learning by Explaining

The preceding analysis of research on learning by explaining has provided some key insights that begin to form a basis for describing how learning partners influence this type of activity. First, however, it should be emphasized that the main goal is to place the responsibility of learning on the explainer. The notion that generating explanations leads to mental model revision, as advocated by Chi (2000), makes this clear. Chi wisely suggests that interventions should focus on providing opportunities for this recognition and repair process to occur rather than concentrating on the optimal way to present learning material. It is vital for students to be given opportunities to make their own thinking explicit. This means promoting not just active learners, but active *aware* learners. One way to accomplish this is through the use of prompts provided by a learning partner.

There is ample evidence to show that prompting students to explain can enable them to learn with deeper understanding and become more knowledgeable about themselves as learners. However, care must be taken when designing the prompts because they must be successful in promoting the use of both cognitive and self-monitoring strategies. The analysis provided a detailed account of how these prompts might be designed.

Previous research has shown that learning partners can play a key role in supporting the explanation process. However, even though partners have figured prominently in many studies, there is very little data available to help determine what kind of characteristics might be optimal. Clearly, simply putting students together in pairs or groups will not guarantee success. Focused intervention studies are needed to fill in this gap in the literature. One way of addressing this issue is through the use of software agents that can be programmed to offer various prompts as well as feature different personalities and roles.

The following section examines how technology has been used to provide both prompts and computer-based partners to enhance learning in previous work. The goal is to identify some insights that will inform the design of software agents that can support learning by explaining.

Technology and Learning Partners

One of the most interesting possibilities that computer environments offer is the introduction of a computer-generated learning partner. At the Learning Sciences Institute (LSI) at Vanderbilt there are several projects underway to explore this concept. Computer-based "teachable agents" at the LSI are being developed as learning partners that provide students with an opportunity to learn by teaching another student (Biswas et al., 2001). Among other advantages, students are able to see the consequences of their teaching without the possibility of hindering the learning of a real person.

An underlying assumption at the LSI is that effective teachable agents can be developed with a very basic level of artificial intelligence. This philosophy runs counter to many learning environment developers. For example, VanLehn, Ohlsson, & Nason (1994) suggest that the level of sophistication required for simulated students may be quite high. Although they point out that it may be difficult for the simulation to possess an accurate representation of the student's mental model, they imply that the simulation must at least maintain a representation of its own knowledge as it learns (on the assumption that it would know approximately as much as its human partner). This may be overestimating the level of sophistication required, however, as there is little evidence that any significant mental model sharing occurs between two human collaborators.

It studies of human tutoring, for example, it has been demonstrated that the tutors do not have a good representation of their tutee's mental model (Graesser, 1995). It may be enough, in fact, for the simulated student to possess the kind of "knowledge about learning" taught in many of the interventions mentioned in the previously discussed research and simply a general framework of the domain. With this model, the simulated learning partner utilizes general prompts to encourage the student to think about and explain concepts more deeply than they would otherwise.

Simulated learning partners may also provide more opportunities for the average student to generate explanations. Webb (1989), points out that certain students in heterogeneous groups may be at a disadvantage in terms of explanatory interactions as special relationships tend to form. For example, middle ability students may not engage in as many interactions as high-low students, who tend to form a teacher-learner relationship. Technology may have the advantage of stabilizing the perceived ability of the learning partner.

Technology and Prompting

One way that computers can provide guidance to students is through the use of learning partners who are capable of providing effective prompts. Usually computer environments support prompting in one of two ways. One technique is to provide a framework for peers or experts to generate prompts as in the case of the Sensemaker tool, developed as part of the Knowledge Integration Environment (Bell, 1997) and CSILE (Hewitt & Scardamalia, 1998). Here, prompts are generated by users in reaction to some information or claim that the learner has shared within a collaborative environment.

Another way is for the software itself to provide prompts. For example, Lin and Lehman (1999) developed an Isopod Simulation program that prompted students to explain their scientific problem solving processes as they planned, performed and interpreted various isopod experiments. In a discussion about designing technology to support reflection, Lin et al. (1999)

identify process prompting as a method to direct students' attention to specific aspects of learning processes. By process prompting, Lin et al. (1999) refer to systems that pose questions and help students monitor their own progress.

Although there are many examples of software that prompts students in various ways and a growing number of studies investigating the use of software agents, there are no interventions that combine these two elements to directly look at the effects of prompting by a learning partner.

Summary and Implications

The analysis of research on learning by explaining provides the foundation for the experimental design described in the next chapter. This summary brings together some key ideas for the design of software agents as learning partners in an explanation activity.

The use of technology to enhance learning by explaining should focus on ways of making student thinking explicit and encourage students to reflect on their own state of knowledge – both about content and themselves as learners. The preceding analysis of research on learning by explaining offers some valuable insights on how software agents as learning partners might be used to accomplish this goal. Although the full potential of learning partners in an explanation setting remains unknown, it is likely that they impact the learning process in several ways.

First, agents as partners can change how student thinking is made explicit. Although it is possible to make thinking visible with self-explanations, it is hypothesized here that the presence of a partner compels the explainer to consider the perspective of the other person (or agent) and therefore encourages the generation of a deeper explanation. One possible reason for this is that explainers may not make the same kind of assumptions about what is known or understood when they have to explain to another.

Second, interaction with a partner provides opportunities for the partner to actively influence how an explanation is generated. One way to do this is through the use of prompts designed to help students understand and develop their own way of learning. These prompts should be specifically designed to encourage the kind of cognitive and self-monitoring strategies identified earlier. It is argued that an agent does not necessarily have to possess specific knowledge about its partner since the emphasis is more on enabling rather than diagnosing.

Third, it is likely that the nature or personality of a learning partner can play an important part in determining the particular attitude a student takes when trying to explain a concept. It is hypothesized that an agent in the role of an advisor will encourage students to take the task more seriously and remain more focused on the learning goals than if they explained to a peer.

The ideas outlined above are used to form the hypothesis that a software agent, infused with knowledge about how to generate explanations and programmed with the ability to provide appropriate prompts, will make an effective learning partner by enabling students to remain focused on the learning task and generate deep explanations. The experiment described in the next chapter is designed to test the validity of this claim.

CHAPTER III

METHODOLOGY AND RATIONALE

Rationale

This study is best described as mixed methods experimental design since both quantitative and qualitative methodologies were employed. The quantitative component helped determine whether the experimental conditions differed in terms of learning outcomes. A pre- and post-test containing a mixture of multiple choice and short answer questions was used for this purpose. The qualitative component was used to explore multiple dimensions including the quality of explanations generated and student perceptions of the use and value of resources provided in the virtual environment. In addition, an interview was conducted with the classroom teacher to investigate how an explanation summary report might benefit teachers.

The rationale for this mixed methods design was to provide insights beyond basic learning outcomes to inform future study designs. Since virtually nothing is known about how teachers might use an explanation summary report, this exploratory approach was designed to provide useful information to help guide further work in this area.

Participants and Location

Participants consisted of approximately 80 5th grade students and one science teacher from a local private school. The school was considered appropriate for this study because the computer technology required was readily available and has been tested in two prior pilot studies. In addition, the researcher had an established working relationship with the science teacher which helped reduce the risk of implementation errors. While concerns about the generalizability of findings are not unjustified when experiments are conducted in private schools, it was felt that rapid advances in technology warrant the testing of this type of software in schools that are perhaps ahead of the curve in technological implementation. It should also be noted that the software itself does not require especially powerful or advanced computers.

Research Design

The student portion of the study had two components: an introduction tutorial period where students learned how to navigate through the virtual world and explanation exercise consisting of two learning modules.

The Environment

At the core of this study was the interaction between students, their learning partners, and the explanation resources. A main goal was to determine whether the presence of an animated software agent (also referred to as a bot) would improve the way students attend to and use explanation resources. The bots were programmed to provide advice, through conversational dialog, about generating elaborative explanations. This form of receiving learning resources was compared to similar information provided in a text-based format.

In order to support and capture the students' interactions, a browser-based 3D chat environment known as Active Worlds was used (see Figure 1). The 3D virtual world has several advantages. First, its function as a chat environment provides the opportunity to capture student-generated explanations and store them for later conversion into summary reports. Second, Active Worlds supports programmed software agents that, in this case, act as learning advisors or coaches who provide information and suggestions for making effective explanations. Third, the virtual world provides a controlled space where many students can complete the same interactive learning module in parallel. Finally, previous pilot work has demonstrated that students find working in Active Worlds to be an exciting and motivating experience.

Figure 1. Active Worlds View and Chat Windows

<u>The Learning Modules</u>

The learning activity was situated within an Active Worlds environment that simulated important aspects of a river ecosystem. Through the use of software characters known as avatars (graphical representations of people), students were able to walk around the environment, chat with their learning partner, and perform various problem solving tasks. All students worked through a set of two modules in the domain of ecology. Ecology is an appropriate subject area as the topics involve complex systems such as river or wetlands ecosystems. Chi (1994) points out that for students to develop deep understanding of complex systems, they need to not only understand each separate component but also the relationships within and among them. One of the primary advantages of generating explanations is that they help students integrate these relationships as they develop more sophisticated mental models (Chi, 1994).

In each module, the task was to help solve a problem, such as determining whether a river was polluted, and to generate explanations for several scientific concepts along the way. It was intended that the material be novel for most students.

Software Agent Design

The design of the software agents was constrained by two factors. First, the look and feel of the agents was pre-determined by the environment because the agents utilized existing Active Worlds avatar forms. This constraint meant that it was not possible to control certain random “expressions” that the standard avatars perform. These include head-turning, arm gestures, and other physical movements that exist to help make the avatars appear more life-like. These characteristics were shared by the avatars used by the students as well. The second constraint on the agent design was the scripting software that was used to control the agents’ actions and dialog. The software enabled the agents to respond to chat messages and other events within the virtual world such as a student clicking on a sign.

The artificial intelligence programmed into the agents was limited to identifying keywords and responding to specific actions by the students. Depending on the condition. the goal was to place the agent in the role of an explanation advisor or an advisor and learning partner. When acting only as an advisor, the agent would simply provide suggestions for how to prompt for and generate an effective explanation based on the strategies outlined in the next section. The agent would also model an explanation. When positioned as a learning partner as well, the agent became a full participant in the explanation exercise. That is, it provided the prompt for its student partner and, in turn, generated an explanation in response to a student prompt. A detailed description of the experimental conditions is presented later.

Because the artificial intelligence capability was quite limited, an important design decision was to place the agents or bots in a leadership role. This had the advantage of simplifying aspects of command and control since the student did not have to direct the agent. As a result, it was the agent that controlled progress through the exercise. When one task was completed, the agent would inform the student that the task was done and where they should move to next one. At a new location the agent would only respond once the student performed a certain action, such as clicking on an information sign. Finally, it is important to emphasize that the main purpose of the agents was to support students in the construction of explanations and reflection on the meaning of specific ecosystem concepts.

Explanation Resources

Resources to help students learn about and generate effective explanations were provided in either a text- or agent-based format, depending on condition (see below). The resources included advice about the characteristics of elaborative explanations and some examples. Although the presentation of these resources varied between groups (text-based or agent-based), the general format and content was the same. Students were introduced to features of effective explanations that have been identified by a number of researchers and are commonly taught to students who receive training on explanation technique (Bielaczyc, 1995; Chi, 1994; Coleman, 1998). Several are listed below:

- Identify new information
- Relate new ideas to prior knowledge
- Identify relationships between new ideas
- Impose a goal on the new content
- Give meaning to new content
- Anticipate outcomes

In addition, students were introduced to self-monitoring strategies help learners develop personally meaningful knowledge and knowledge of self as learner (Chi, 2000; Palincsar & Brown,1984). Some examples are listed below:

- Identify personal learning goals
- Identify comprehension successes
- Identify comprehension failures
- Identify knowledge gaps
- Identify shifts in perspective

An example of a scientific concept around which the explanations were generated is provided below (see Figure 2) along with the accompanying resource in both the text- and agent-based formats. Note that although both resources were presented as text, the wording and tone of the agent-based resource was conversational and was intended to be consistent with informal chat dialog. Another significant difference in the presentation of the agent-based resources was that the information appeared in the chat window. The text-based resource was viewed as a web page document.

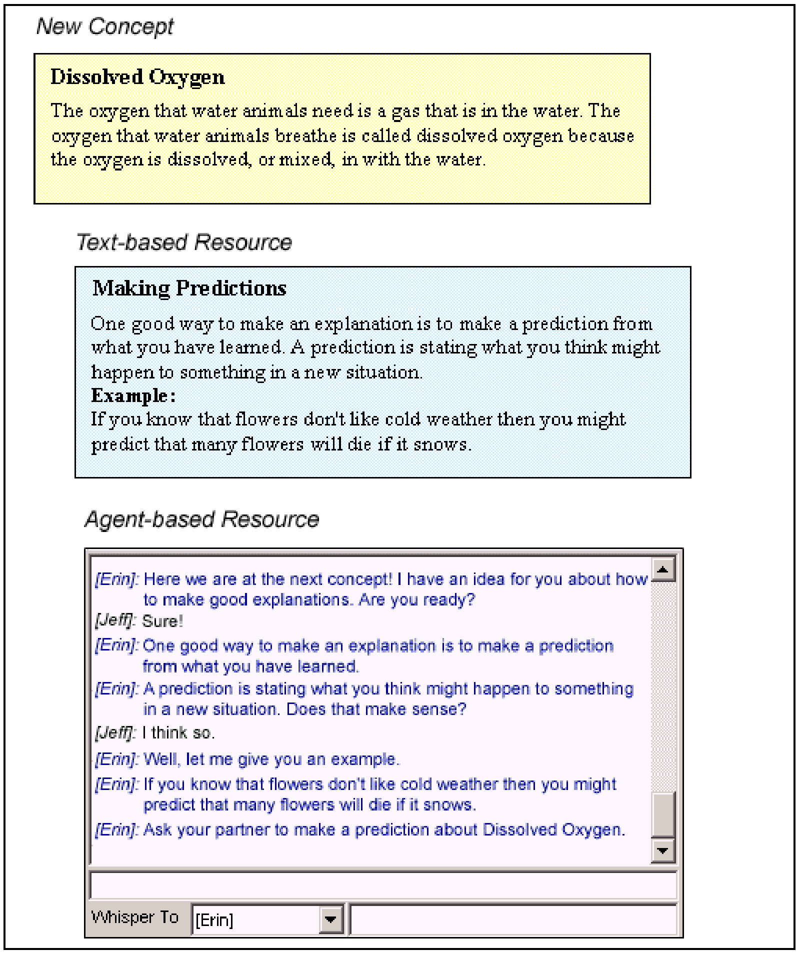

Figure 2. Example Concept and Related Resources.

Experimental Conditions

The three conditions used in this study varied by how the explanation resources were presented and whether the listener was another student or software agent (see Table 6 below). The first two conditions involved pairs of students who worked in collaboration. Students in the first condition (Group SST) were provided with text-based resources while students in the second group (Group SSB) obtained their resources through dialog with a software agent (or *bot*). Therefore, the main difference between these two groups was the interactivity of the resources provided in the environment.

Table 6. Experimental Conditions

Partner	Resources	
	Text	Bot
Student	**SST**	**SSB**
Bot		**SBB**

Students in the third group (Group SBB) did not have a student partner but obtained their resources from an agent similar to the students in Group SSB. Because they did not have a student partner, students in this condition explained to the agent who, in turn, provided an explanation for the student for every second concept. This ensured that students in this condition did not generate twice as many explanations as students in the other two conditions. The main reason for including an individual student condition was to differentiate between the effectiveness of software agents as resource providers and as listeners. If both the groups with software agent outperformed the group with no agent, it would suggest that software agents can be used effectively as resource providers. However, if the individual condition outperformed both the other groups then it would suggest that the advantage of the agent is as an explanation partner rather than resource provider since explaining to the agent is a unique feature of Group SBB. Previous results from a pilot study showed that dyads frequently strayed off-topic during an Active Worlds problem solving exercise so it is anticipated that the software agent as both resource provider and listener will help students stay on task.

A. Pairs with Text-based Resources (Group SST)

Students in Group SST were presented with text information that provided hints and suggestions for making better explanations. At each explanation area in the environment, a new concept was presented and students alternated roles as prompter (and listener) and explainer. It was the responsibility of both students to use the available text-based resources. Communication took place in real-time through a chat window located beneath a world browser window (see Figure 1).

B. Pairs with Agent-based Resources (Group SSB)

Student pairs were presented with explanation resources similar to the Pairs in Group SST except that the information was presented as dialog from their software agent partner (bot). The students typed their explanations in a chat box as part of an ongoing dialog between their student partner and the software agent.

C. Individuals with Agent-based Resources (Group SBB)

Students in this group received resources from a bot as in Group SSB but in this case the students' explanations were directed back to the bot, not to another student. Since each student in the pair conditions took turns generating explanations, the bot generated half the explanations in this condition. This ensured that students in this group generate the same total number of explanations as students in the other two conditions. A potential confound was the possibility that the bot's explanations were more elaborate or different than the average student partner. However, this was of less concern than having students in this group generate twice as many explanations as students in other conditions. It should also be noted that the strength of this learning activity was not in the hearing of explanations but in their generation. It is also important to have a condition where the student explained to the software agent as this helped differentiate the benefits of agents as resource providers and listeners.

<u>Timeline</u>

The study was conducted over a period of three weeks (see Figure 3 below). In the first week students took the pre-test and completed the first virtual world module. In the second week the students completed the other module and took the post-test. Finally, the interviews will be took place in the third week.

Week 1	Week 2	Week 3
Pre-test and Module 1	Module 2 and Post-test	Discussions and Interviews

Figure 3. Study Timeline

Data Collection

Data was collected from a variety of sources in an effort to triangulate results and strengthen conclusions. Sources included pre- and posttests, Active World module transcripts, and data from the follow-up interviews.

Assessments of Learning and Student Perceptions

A pre- and post-test was administered to all student participants. The assessments consisted of both multiple choice/short answer type questions (see Appendix F). Material was based on content used in a pilot study and supplemented with new questions with cooperation from the classroom teacher. In addition, several questions were included in the posttest to assess perceptions of the use and value of explanation resources and learning partners. A Likert scale format was used for this component.

Active World Transcripts

All dialog during the Active World explanation exercises was recorded in log files. These served as transcripts for analyzing the interactions between learning partners and agents.

Interviews

All student interviews and the teacher interview were tape-recorded for later transcription.

Student Interviews

To further investigate student perceptions of learning partners and explanation resources, a brief interview was conducted with several participants from each of the three conditions following the group discussions. The goal was to explore the students' perceptions of both learning partners and resources in terms of credibility, understanding, and value. The interviews were conducted in an intentional, semi-structured format (Bogdan & Biklen, 1992) designed to provide a focused conversation with enough flexibility to explore unanticipated topics. The following questions were used:

Resources:

- How did you like the help that you received during the virtual world exercise?
- What advice would you give to another student about using the resources?
- Did the resources (or bot advice) help you to make better explanations? If so, how?

Learning Partners:

- How did you like working with your partner?
- If you were to do this exercise again, would you prefer to work with a partner or work alone? Why?

Teacher Interview

An interview with the participating teacher focused on the summary report of the explanations and related dialog for each student. The reports showed the prompt (whether from an agent or student) and explanations generated by each student. The teacher report also included student names. The goal was to investigate how this teacher perceived the utility and value of such a report. The following questions were used in a semi-structured interview:

- Did you learn anything new about your students from these reports?
- Did anything surprise you?
- Could you list a few ways that might you use these reports?
- Which report do you think best highlights students with comprehension problems?
- How could we design these reports to be more useful in the future?

Data Analysis

Pre- and Posttests

Data analysis for the study involved both quantitative and qualitative methods. Learning, as measured by performance on the pre- and post-tests, was analyzed quantitatively. Descriptive statistics of the scores for the three groups was computed for both tests and a repeated measures analysis of variance (ANOVA) provided inferential statistics to determine whether the group means differed statistically. An F value was calculated to provide a measure of significance. The ANOVA was also used to determine whether an interaction existed between resource format and type of learning partner. Responses to perceptual questions regarding resources and learning partners in the posttest was partialled out and analyzed separately.

Explanation Exercises and Student Interviews

One of the advantages of a mixed methods design is that it provides a way to triangulate or seek convergence of results (Greene et al., 1989). To this end, qualitative analyses were conducted on data collected from the online explanation exercise and student interviews. Student perceptions of their understanding and the value of the resources were analyzed from the Active World chat logs and student interview data. Chat logs were coded by dividing the messages from each student into pre-defined categories. Comments that related to explanations were coded by explanation type, in a similar manner to the scheme developed by Chi et al. (1994). Other categories included statements relating to help-seeking, self-monitoring, and off-topic conversation.

Student interviews were coded under the two main categories of Resources and Learning Partners. Within each categories, sub-categories were developed that related to the data collected from the chat logs. The goal was to build a detailed account of learning partner relationships and how students perceived and used the explanation resources.

Teacher Interview

The goal of the teacher interview was to explore the participating teacher's perceptions of the utility and value of an explanation summary report. Rather than interviewing several teachers, it was felt that an in-depth interview with the participating teacher would be more enlightening as this teacher was more familiar with the explanation process and the participating students. The transcription was coded initially around three main themes: 1) what the teacher learned from the reports; 2) how the reports might be used and; 3) how the reports could be modified to be more useful.

Ethical Considerations

All appropriate Internal Review Board (IRB) documentation was completed prior to the study and the teacher was provided with all appropriate consent forms and letters to obtain parental permission for each student. In order to protect the anonymity of participants, pseudonyms were used in the final report. Students who did not obtain prior consent to participate were given another exercise by their teacher while the study took place. Although the intervention involved an online virtual world, the use of a login procedure with password access only ensured that the students were working within a closed environment at all times.

CHAPTER IV

RESULTS

Learning by Explaining Study Results

The primary purpose of this study was to examine the role of software agents as learning partners and to determine their effects on student learning. Both quantitative and qualitative data were collected as students completed pre- and posttests and worked through a river monitoring exercise in the virtual world.

Pre-Post Test Results

A posttest was administered following the virtual world exercise to help answer the first research question that asked if students could learn ecological concepts by working through a virtual river monitoring exercise and would the type of learning partner and way resources are presented affect the amount learned. The first analysis focuses on a pre-post comparison of scores from questions that were identical on both tests.

Questions on the pre- and posttest consisted primarily of declarative knowledge type questions. Extra questions on the posttest attempted to measure transfer and were analyzed separately. A one-way ANOVA conducted on pretest scores showed that there was no significant difference between the three groups prior to the exercise, $F(2, 67) = .117$, $p<.89$. A repeated measures ANOVA showed that there were significant gains in pre- to posttest scores for all three groups, $F(1,67) = 58.9$, $p<.001$ (see Figure 4). Although the students working individually with a bot had greater gains (SBB=20.0%) than the other two groups (SSB=18.6% and SST=15.9%), this difference was not significant, $F(2,67) = .28$, $p<.76$. The trends were, however, in the expected direction with the SBB group gaining the most and the SST group (students working together without a bot) achieving the least gain.

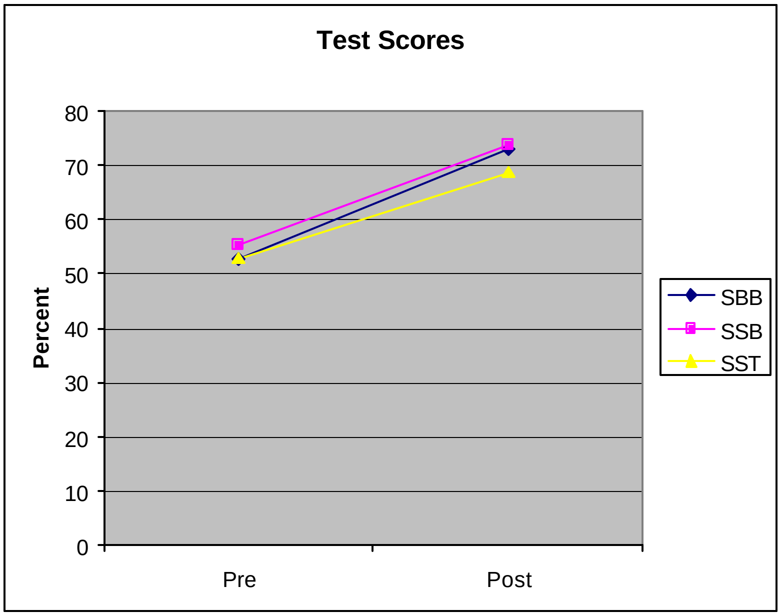

Figure 4. Pre- and Post Test Results

Table 7. Descriptive Statistics and Repeated Measures ANOVA Results

Descriptive Statistics

	GROUP	Mean	Std. Deviation	N
PRE	SBB	52.8000	20.51828	25
	SSB	55.2381	18.33550	21
	SST	52.9167	17.56458	24
	Total	53.5714	18.65287	70
POST	SBB	72.8000	18.37571	25
	SSB	73.8095	15.64487	21
	SST	68.7500	15.12628	24
	Total	71.7143	16.41592	70

Tests of Within-Subjects Effects

Measure: MEASURE_1

Source		Type III Sum of Squares	df	Mean Square	F	Sig.
SCORE	Sphericity Assumed	11447.042	1	11447.042	58.905	.000
SCORE * GROUP	Sphericity Assumed	109.048	2	54.524	.281	.756
Error(SCORE)	Sphericity Assumed	13020.238	67	194.332		

The posttest also included 5 extra questions that were designed to measure near to medium transfer. An example of one transfer type question is:

A friend of yours is interested in river monitoring but also likes to fish. He asks you why you can't use fish instead of macroinvertebrates to monitor water quality.

A one-way ANOVA was used to compare transfer question scores between groups (Table 8). This analysis showed that there were no significant differences on transfer question performance between the three groups, $F(1,69) = 1.278$, $p>.28$.

Table 8. Comparison of Posttest Transfer Questions

Descriptives

TRANSFER

	N	Mean	Std. Deviation	Std. Error	95% Confidence Interval for Mean		Min.	Max.
					Lower Bound	Upper Bound		
SBB	24	70.8333	26.23749	5.35570	59.7542	81.9125	0.00	100.00
SSB	21	78.5714	16.36634	3.57143	71.1216	86.0213	50.00	100.00
SST	24	77.0833	20.74256	4.23406	68.3245	85.8421	25.00	100.00
Total	69	75.3623	21.64756	2.60606	70.1620	80.5626	0.00	100.00

ANOVA

TRANSFER

	Sum of Squares	df	Mean Square	F	Sig.
Between Groups	779.633	2	389.816	0.828	0.442
Within Groups	31,086.310	66	471.005		
Total	31,865.942	68			

Table 9. Transfer Question Performance by Group

Average Group Performance on Transfer Questions

Group	N	Q1	Q2	Q3	Q4	Q5
SBB	24	70.83	75.00	54.17	33.33	83.33
SSB	21	85.71	80.95	61.90	28.57	85.71
SST	24	79.17	70.83	79.17	41.67	79.17

The results from the pre-post test analyses were somewhat surprising because informal classroom observation during the exercise suggested that there was an appreciable difference between the groups in how they worked through the river monitoring exercise. In order to obtain a clearer picture of the post-test results, a follow-up analysis was conducted.

Follow-up Analysis of Pre- Post-test Performance

Because the results in the initial analysis were so unexpected, a more in-depth analysis was conducted to examine the outcomes on each test question. The average scores for each question are shown in Table 10 below.

Table 10. Average Post-test Scores by Question.

Group	Q1	Q2	**Q3**	**Q4**	Q5	Q6	Q7	Q8	Q9	**Q10**
SBB	100.0	87.5	**58.3**	**54.2**	87.5	83.3	79.2	75.0	87.5	**37.5**
SSB	90.5	85.7	**57.1**	**66.7**	90.5	81.0	57.1	81.0	95.2	**38.1**
SST	87.5	70.8	**66.7**	**66.7**	83.3	87.5	62.5	58.3	75.0	**25.0**

Three questions in particular stood out on the post-test where the average score for all groups was low. Since the questions were designed to assess knowledge of the concepts that students explained, a closer look at questions 3, 4, and 10 and the specific prompts for the concepts related to these questions was warranted. This examination revealed the following:

Question 3 was determined to be a rather difficult transfer type question. The question stem read as follows: "If you find many pollution tolerant macroinvertebrates in your river samples then…" The correct answer was "the river might be polluted but you need more information." However, another answer read, "you know the river is polluted." To answer this question correctly, students needed to recognize that the question did not specify if any other types of macroinvertebrates were present or not. Without knowing how many intolerant macroinvertebrates were in the sample, it is impossible to determine the pollution status of the river. Therefore, this question was more difficult than others on the post-test.

Question 4 was based on the concept of macroinvertebrate habitats. Unfortunately, the prompt used for this concept in the exercise was did not match well with one of the planned prompt types. Recall that the prompts were designed to promote deeper thinking by promoting the use of cognitive strategies (such as relating material to prior knowledge) or by monitoring strategies (such as reflecting on one's misconceptions). Although the prompt for this concept was intended to encourage students to reflect on their knowledge about habitats, it is likely that it failed to elicit deep thinking because the wording was too vague. The prompt (the final prompt for day 1) read as follows: "Step back and consider what you have learned today. Does everything make sense? Do you have any questions about this material?" Although the intent of the prompt was to encourage students to reflect on their knowledge about habitats, an examination of the transcripts for this concept revealed that students (understandably) interpreted the prompt to refer to their understanding of the entire day's material. No student reflected specifically on their knowledge about macroinvertebrate habitats.

Question 10 referred to the concept entitled "What are Macroinvertebrates?" A re-evaluation of the material provided for this concept revealed that it contained information on both the physical nature of the organisms as well as their relationship to the health of a river ecosystem. The prompt for this concept asked students to make a prediction about the ecosystem if there were no macroinvertebrates present. Because the prompt was oriented around the ecosystem relationship students focused on that aspect of the concept. The question in the post-test, however, asked students about the physical nature of the organisms. Therefore, the question did not align well to the topic covered by the students' explanations.

Because these three questions had specific problems and performance on each was low across all groups, a second analysis of the test scores was carried out with questions 3, 4, and 10 removed. Figure 5 below reveals the post-test gains for the three groups when these questions are factored out of the results. A follow-up repeated measures ANOVA with the three questions removed revealed that there were significant gains in pre- to post-test scores for all three groups, $F(1,66) = 48.025$, $p<.001$ as was found in the first analysis. In addition, a significant interaction was now revealed between scores and group, $F(2,66) = 3.397$, $p<.039$ (see Table 11).

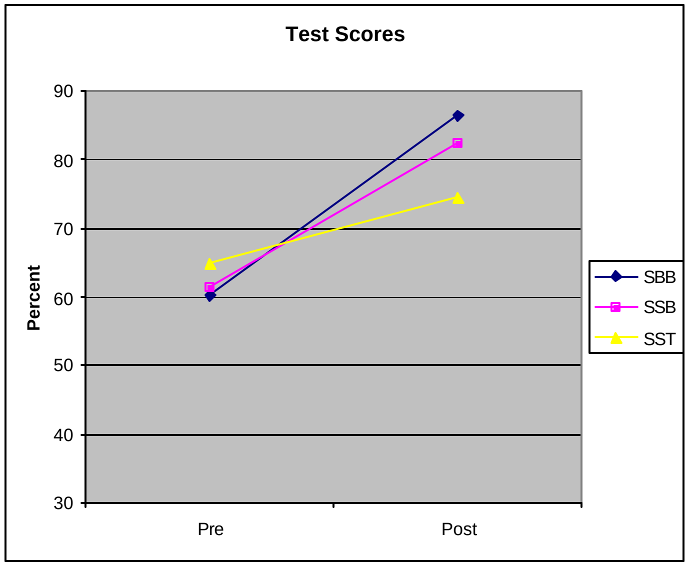

Figure 5. Test Performance with Questions 3,4, and 10 Factored Out.

Table 11. Descriptive Statistics and Repeated Measures ANOVA on Revised Results.

Descriptive Statistics				
	GROUPCOM	Mean	Std. Deviation	N
Pre-test Avg.	SBB	60.1190	24.19206	24
	SSB	61.2245	18.15065	21
	SST	64.8810	21.05436	24
	Total	62.1118	21.19283	69
Posttest Avg.	SBB	86.3077	12.27119	24
	SSB	82.3129	16.22845	21
	SST	74.4048	17.35846	24
	Total	80.9518	16.00414	69

Source		Sum of Sqrs	df	Mean Square	F	Sig.
Test Score	Sphericity Assumed	12318.777	1	12318.777	48.025	.000
Test Score * Group	Sphericity Assumed	1742.618	2	871.309	3.397	.039
Error(FACTOR1)	Sphericity Assumed	16929.568	66	256.509		

The results from this follow-up analysis suggest that although all three groups made significant gains in their test scores, students who worked with the bots had an advantage over students who only worked with one of their classmates. The next analysis focuses in on the chat transcripts in an effort to gain further insight into how software agent learning partners might have affected learning.

River Monitoring Exercise Transcript Analysis

Analysis of the Active Worlds chat messages revealed several striking differences between the treatment groups. An initial analysis was conducted to determine how students in each group differed in terms of how they communicated with their partners. The chat messages were coded into several categories and reported as a percentage of the total dialog averaged across group members. Figure 6 below shows the types of dialog that occurred during the river monitoring exercise.

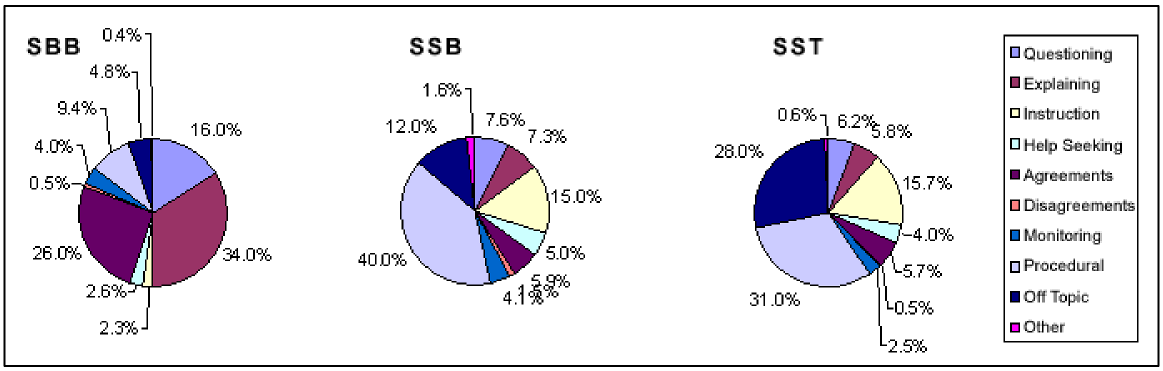

Figure 6. Dialog Patterns During the Virtual World Exercise

Questioning and Explaining

Chat messages that included either asking a question or making an explanation about one of the ten concepts presented in the virtual world were coded into these categories. Because asking questions and making explanations was the main purpose of the exercise, these two categories are representative of time on task. As expected, students who worked alone with a bot partner stayed more on task than students in either of the other two groups. 50% of the SBB group's dialog fell into these two categories while this type of dialog from student pairs who worked with bots and pairs who had text-only resources made up only 15% and 12% of their total dialog respectively.

As was noted earlier, students in the SBB group generated significantly more "deep" explanations that students in the other groups. To a large degree, this effect appears to be a result of the better prompts that were provided by the software agents. Students in the SBB group had the advantage that the bot was their learning partner as well as advisor and this meant that the bot both prompted the student and generated an explanation, therefore modeling good explanatory skills. Below is one example of an SBB student responding to a bot prompt:

> Bot: "Do you think you would find a lot or very few pollution intolerant macroinvertebrates in a polluted river?"
>
> Student 6: "Very few… if they cannot stand it they wouldn't be there."
>
> Bot: "Try to explain why…"
>
> Student 6: "… they would die if they couldn't breathe."

The success of the bot prompting often depended on a follow-up prompt as in the example above. While students would often initially make an incomplete explanation, a simple prompt like "Try to explain why" encouraged them to think harder and make causal connections. This type of follow-up prompt was largely absent in the SSB and SST group interactions were students appeared satisfied by single sentence explanations from their partners. In this example from a pair in the SST group, notice also that the format of the question is designed simply to elicit a right or wrong answer:

> Student 17: "What do macroinvertebrates eat?"
>
> Student 16: "macroinvertebrates eat bacteria, dead decaying plants and animals"
>
> Student 17: "Right"

The pattern was typical in the SST group even though they were told to read about how to ask deep questions on the resource page that accompanied each concept. The SSB student pairs who explained to each other but had a bot advisor did not appear to do much better. In some cases the students left their bot advisor behind or essentially ignored the bots advice. In some cases the students would respond to the bot but then not follow its instructions:

> Bot: "So for this concept, ask your partner how this information about pollution goes with something they already know."
>
> Student 8: "[to bot] Okay"
>
> Student 8: "[to Student 9] Are you done?"

Student 9: “Yes, I’ll wait for you at sign #2”;

In this case, the students appeared to simply be reading the information for themselves, and so did not ask questions or give explanations to each other.

Although the SBB students working alone with the bots generally did better, the bot prompts were not always successful in eliciting a good explanation. The following exchange demonstrates how one student became tired of his bot’s questions:

Bot: “A good technique for explaining is to say what you think is new information and how that changes your thinking.”

Bot: “So, for this concept, what new thing did you learn?”

Student 2: “nothing, don’t ask me again.”

The SBB students were not very successful in asking good questions. Despite the fact that their bot advisor modeled a good prompt first, they typically asked very shallow declarative knowledge type questions. The students’ questions were, in fact, very similar across groups and generally of poor quality.

A disadvantage in using software agents with limited artificial intelligence is that they are unable to interpret what a student writes and so often miss good opportunities for an extended discussion. In one dialog about pollution a student made the following observation:

Bot: “How does this information about pollution go with something you already know?”

Student 1: “I knew that pollution upsetted the environment but not that changes such as warm water are really that bad.”

Bot: “Okay, do you have anything else to add?”

Student 1: “Pollution could simply be a log messing up the river. Does that mean beaver dams could be pollution?”

Unfortunately, the bot was incapable of responding to this interesting thought and simply moved on to the next concept.

Instruction

This category included any dialog that involved giving an instruction to either a student or bot learning partner (e.g. “Come over here.”). This category made up around 15% of the dialog in both the SSB and SST groups. Students working alone with a bot only issued these command type statements less than 3% of the time.

It was expected that students in the SBB group would give fewer instructions since their bot advisor was programmed to take a leadership role. Although the bots in the SSB group were programmed similarly, they were much less successful in controlling progress through the world because the students could also interact with their student partner. Both SSB and SST (pairs without a bot) students spent a considerable amount of time telling each other where to go or what to do. In many cases one student would want to move quickly and efficiently through the activity but his or her partner would not cooperate very well. Here is a typical exchange:

Student 15: “Ask me a question… please?!”
Student 15: “ask me!”
Student 15: “ask me a question!”
Student 14: “I haven’t gotten to read the sign yet.”
Student 15: “What are you doing??”

Sometimes, however, students would have more computer experience than their partner and offer advice on how to move around the world:

Student 19: “What do I do?”
Student 18: “Come to me”
Student 19: “I don’t know how… where?”
Student 18: “Use the arrow keys [to move around]”
Student 19: “Ok”
Student 18: “Use your mouse to click on the blue sign”

Help Seeking

Seeking help included requests for help in both content and environment (procedural) areas. Students working alone with a bot (SBB) asked for help about half as much as the groups with two students (2.6% versus 5% and 4%). Most of the help requests involved questions about how to move around in the environment (when students were first entering the world) or procedural questions pertaining to the river monitoring and explanation activities. Although all of the students received a general introduction on how to complete the exercise, many students were unsure of how to proceed and so much of the question asking centered on what steps to follow. This generally occurred within the SST and SSB groups. An example of one student asking her partner for help during the river sampling stage follows:

Student 10b: "How do we take samples?"
Student 11b: "You click on the blue sign"
Student 10b: "I am just getting the rifver sample"
Student 10b: "so now we mark them on our sheet, right?"
Student 11b: "scroll down…"
Student 10b: "Ohhhh.. I was confused."

The leadership provided by the bot advisor in the SBB group reduced the need for students to ask procedural type questions since the bot would provide instructions at each stage. However, there were occasions when students would be frustrated by the bot's inability to provide specific assistance. In this exchange, the student tries to elicit help from his bot advisor:

Bot: "Great, you should be able to record the first sample on your data sheet now."
Student 7: "what datasheet?"
Student 7: "hello"
Student 7: "why won't you answer??"

Agreements and Disagreements

Statements of agreement were especially prevalent in the SBB group were these messages made up 26% of the dialog. This compared to around 6% for the SSB and SST groups. Chat messages coded into this category included answers to questions where the response was similar to "Yes" or "I agree" and statements of acknowledgement such as "Ok" and "fine". Disagreements included dialog where students voiced a dissenting opinion to their partner. These made up very little of the dialog in any group (1.5% or less).

Monitoring

This category included messages where students made statements that appeared to involve some form of self-monitoring such as "Okay, I see", "I meant to say…" or "I'm not sure". Dialog of this type made up around 4.5% of the conversation in groups SBB and SSB and 2.5% in group SST. It is interesting that this activity occurred more often in the presence of a bot. It is possible that since the bots were programmed as "experts", the students expected to receive more help if they acknowledged their misunderstandings or lack of knowledge. Here is one example:

> Bot: "How does this information go with something you already know?"
> Student 2: "I don't know that's a tough question".
> Bot: "Well, let me ask it in a different way."
> Bot: "What other animals do you know that breathe underwater?"
> Student 2: "Fish am I right?"

This dialog also provides an example where the limited intelligence of the bot worked quite well. The bot was programmed to re-phrase the question if the word pattern "I don't know" was found.

Procedural

This category included any chat messages that referred to an activity or procedure in which the students were engaged (e.g. "I am taking a sample now."). These statements made up a large percentage of both the SSB (40%) and the SST (31%) groups but much less so in the SBB group (9.4%). Because the students were free to move around anywhere within their section of the virtual world, they spent a significant amount of time informing each other about where they were or what they were doing. Although the pairs were supposed to move through the activity together, this did not always work out since one student was often a faster reader or more comfortable moving around in the computer environment. Even when the students were in close proximity, however, the avatars lacked visual cues that would be present in a real world environment so that students felt a need to keep their partner informed of their activity. For example, just because a student could see his or her partner standing in front of a sign, there was no way of telling if he or she was actually reading the information:

Student 11c: "Are you done?"
Student 10c: "yup"
Student 10c: "I'm done"
Student 11c: "meet me at sign 2"

<u>Off-topic</u>

Comments that were unrelated to the explanation or river monitoring tasks were included in this category. As expected, the students in the SBB group remained the most on task with only 4.8% of their statements coded as off-topic. Students in the SSB group made off-topic statements 12% of the time while students working in pairs without a bot (SST) were off-topic around 28% of the time. Off-topic statements ranged from chat messages about features in the environment (such as changing the appearance of an avatar) to non-school related topics and "nonsense" dialog. Because the Active World environment was new to almost every student, it was expected that there would be a significant amount of discussion related to the physical appearance of things in the environment. A typical exchange about avatar appearance follows:

Student 15d: "Who are you?"
Student 14d: "Bike" [referring to a specific avatar appearance]
Student 15d: "Look, i'm a bird"
Student 15d: "birdie"
Student 14d: "How did you change that?"
Student 15d: "You can use the menu…"

It appears then, that the presence of a bot helped keep students focused on the activity since SBB students spent far less time typing off-topic messages. When two students were present there was an obvious temptation for them to chat about other subjects but the inclusion of a bot (group SSB) reduced this activity by more than half.

Off-topic statements that were found in the SBB transcripts tended to occur when a student discovered that their bot partner had no real ability to detect what they were typing. Although this did not happen frequently, it suggests that strategies such as telling the students that their chat logs were being recorded might be an important deterrent. Here is one example of an SBB student typing "nonsense" chat:

Bot: "So, what was the key point?"

Student 1c: "bbb"

Bot: "I'm not sure I understood, can you tell me more?"

Student: "cvfgj"

Bot: "Ok, when you're ready, join me at our last sampling area."

Summary

Analysis of the chat transcripts indicated several major differences between students who worked with a software agent and those who worked with another student. The largest discrepancy between the groups appears to be time on task. The analysis showed that students working alone with a bot were asking questions or making explanations 50% of the time. Student pairs, working with or without a bot partner, spent 15% or less time asking questions and making explanations. It was also clear that students working alone with a bot spent a lot less time giving instructions to their partners. This was expected as the bot was programmed to lead the student through the activities. Also expected was the increased time that student pairs spent on procedural types of conversation.

An Evaluation of Explanation Types

An analysis of explanation type was conducted to gain a better understanding of the quantity and quality of explanations generated by each experimental group. The prediction was that the students working with the software agents should produce a larger number of higher quality explanations than students who only worked with another student. In addition, it was predicted that students who worked alone with a bot (group SBB) would generate better explanations than students who received information from a bot but explained to a student partner (group SSB) because in the SBB group the bot also modeled prompts and explanations.

Chat logs from 30 students (10 from each group) were randomly selected for a detailed analysis of explanation type. Explanation categories based on the system developed by Chi et al. (1989) were used to classify the types of explanations made by the students. In their study, explanations included statements that contained inferences, imposed a goal or purpose, or gave meaning to expressions. Although Chi et al. (1989) reserved the word *explanation* for statements that said something substantive about the material under study, a somewhat broader definition is used here. For this analysis, explanations were divided into three main categories: Deep Explanations, Shallow Explanations, and Monitoring Statements (see table 12 for examples). Deep Explanations included statements that stated a causal relationship, made a connection to prior knowledge, or gave meaning to the content. Shallow explanations, on the other hand, included declarative statements or comments that were simply paraphrases of the material. Monitoring statements included any references to the student's own state of knowledge (either understandings or misunderstandings).

Table 12. Examples of Different Explanation Types

	Example Statements
Deep	[In discussing consequences to a river if there were no macroinvertebrates this student relates an increase in bacteria to the concept of river health] "What would happen is that the health would drop in the river… because the macroinvertebrates wouldn't be there to eat (the bacteria)."
Shallow	[Paraphrasing the content provided] "Macroinvertebrates are important to the ecosystem because they eat bacteria."
Monitoring	"I didn't know that some animals eat bacteria."

Table 13 provides a summary of the average number of different explanation types generated by group.

Table 13. Explanation Distribution in SBB Group

	Frequency of Explanation Types (Mean (SD))		
	Shallow	Deep	Monitoring
SBB	8.0 (3.8)	5.2 (4.5)	1.1 (0.6)
SSB	5.7 (2.1)	1.7 (1.5)	1.0 (0.7)
SST	4.6 (3.6)	0.9 (1.1)	0.8 (0.8)

One-way ANOVAs conducted on each of the explanation categories revealed that there was a significant difference between groups on the number of deep explanations generated ($F_{(2,29)} = 5.821$, $p<.008$) but no significant difference in the number of shallow explanations ($F_{(2,29)} = 2.444$ $p<.106$) or monitoring statements made ($F_{(2,29)} = 0.103$, $p<.902$). A Tukey HSD comparison between groups performed on the deep explanation data showed that the SBB group generated significantly more deep explanations than either the SSB or SST groups. No other differences were significant.

An Exploration of Individual Differences

Although the students who worked alone with the software agents outperformed the other groups on the posttest and in the number of deep explanations generated, there was a sufficiently large amount of variation in the results to warrant future investigation. The goal of this analysis was to determine whether some students were able to benefit more than others from their software agent partners. One possibility is that below average students might benefit more from a software agent learning partner than above average students since the bot was programmed to take the lead, provide help, and not to stray off task. Therefore, students who normally need more guidance and support might perform better with the aid of a software agent. A competing hypothesis is that above average students might actually benefit more from a bot partner since they would be able to remain focused on the task and not be disturbed by another student – i.e. it is more like working alone. In addition, students who had higher verbal abilities might succeed more in an environment where the explanations had to be communicated as typed messages.

Since no class ranking information was collected prior to the study, an initial analysis was performed based on students' pre-test scores.

Results from students in groups SBB (students working alone with a bot) and SST (students working with a student partner only) were included in this analysis. Students were divided into High and Low students based on their pretest scores. Students scoring 60% or above on the pretest were included in the High group and the rest were included into a Low group. This breakdown provided a fairly even number of students among the four groups (see Table 14).

Table 14. Pre and Posttest Averages by High/Low Groups

Group	N	Pre-test Average	Post-test Average
SBB High	12	75.00 (SD=16.26)	94.05 (SD=7.36)
SBB Low	12	45.24 (SD=21.82)	78.57 (SD=11.40)
SST High	10	84.29 (SD=10.54)	77.14 (SD=19.28)
SST Low	14	51.02 (SD=14.52)	72.45 (SD=16.30)

Surprisingly, the results show that the High students in the SST group actually decreased in performance from pre- to posttest (See Figure 7). Based on this grouping, High SST students lost an average of about 5% from pretest to posttest while the High SBB students gained around 19%.

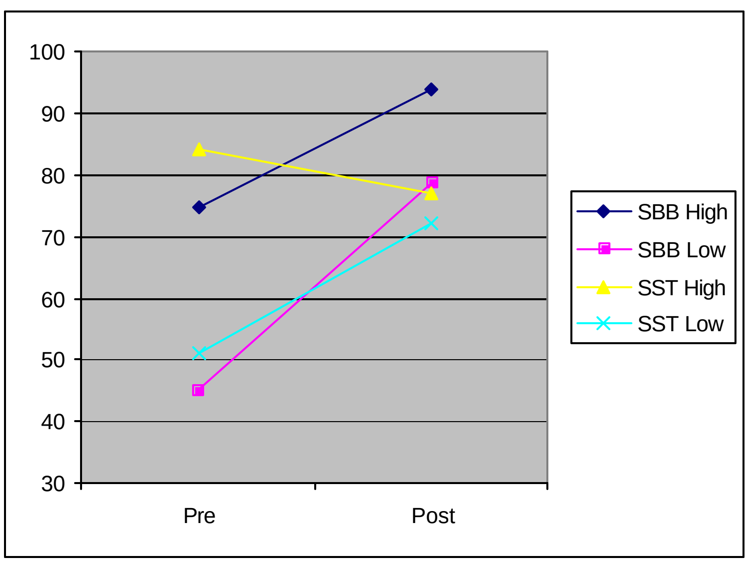

Figure 7: High and Low Student Performance by Pre-test Score

A repeated measures ANOVA was conducted to determine whether the differences between the High/Low groups were significant. Results of this analysis indicated that there was an overall significant effect for group F (1,44) = 26.40, p<.001 (see Table 9).

Table 15. Repeated Measures ANOVA Based on High/Low Student Groups

Source	Sum of Squares	df	Mean Square	F	Sig.
GROUP	6572.066	1	6572.066	26.401	0.000
GROUP * HIGH_LOW	4659.201	3	1553.067	6.239	0.001
Error(FACTOR1)	10953.198	44	248.936		

In addition, a significant Group by High/Low interaction was found, F (3,44) = 6.24, p<.001. A Tukey test, performed to compare difference between groups, revealed that there was a significant difference between the SBB High group and both Low groups. There was also a significant difference between the SST High group and the SBB Low group (see Table 16).

Table 16. Tukey Test Comparisons Between Groups

Multiple Comparisons

HIGH_LOW	HI_LOW	Mean Diff	Std. Error	Sig.	95% Confidence Interval	
					Lower Bound	Upper Bound
1 (SBB High)	2	22.6208(*)	4.29962	.000	11.1408	34.1008
	3	3.8095	4.50948	.833	-8.2308	15.8499
	4	22.7891(*)	4.14321	.000	11.7267	33.8515
2 (SBB Low)	1	-22.6208(*)	4.29962	.000	-34.1008	-11.1408
	3	-18.8113(*)	4.50948	.001	-30.8516	-6.7710
	4	.1683	4.14321	1.000	-10.8941	11.2307
3 (SST High)	1	-3.8095	4.50948	.833	-15.8499	8.2308
	2	18.8113(*)	4.50948	.001	6.7710	30.8516
	4	18.9796(*)	4.36061	.000	7.3367	30.6224
4 (SST Low)	1	-22.7891(*)	4.14321	.000	-33.8515	-11.7267
	2	-.1683	4.14321	1.000	-11.2307	10.8941
	3	-18.9796(*)	4.36061	.000	-30.6224	-7.3367

* The mean difference is significant at the .05 level.

Student Interviews

Several students were randomly selected from each of the experimental groups to participate in an interview following the river monitoring activity. The purpose of the interview was to gain further insight into how the type of partner (bot versus student) affected learning in the virtual world and whether bots were able to provide adequate support and guidance as students generated explanations. Analysis of the interview data revealed that the bots achieved mixed results as learning partners, at times effective and helpful and at other times frustrating and limiting. Three main themes emerged as being important aspects of how students perceived and interacted with their partners.

Expectations

One consequence of using a 3d virtual world with human-like avatars is that it may elevate students' expectations of the intelligence and behavior of the software agents. For students working alone with a software agent (SBB) it became evident that they saw the bots not simply as guides or resources but also as companions with whom they wanted to interact with beyond a simple exchange of information. Several students were disappointed that they were unable to discuss their activities in more depth when they had a bot partner:

> "You can't talk about the [concepts] together. With the robot you have to say your point then and there."

Some students commented on their desire to communicate with the software agents about more than just the task at hand. For example, one student commented:

> "I didn't like it because you couldn't like, talk to it. It wouldn't like answer you. Like, I mean you couldn't talk [about] more than the questions it asked."

Interestingly, there was also evidence that students may have higher expectations for the intelligence of software agents through exposure to other programs and even advertisements. One student compared his bot to an agent from a reading program that he had used. In the program the agent provides advice and assistance in looking up words in a dictionary. Despite the limited functionality of the agent in this case (i.e. it remains in one location and simply draws from a dictionary database) the student perceived this agent to be more intelligent than the Active Worlds agents. It is likely that expectations increase as more realistic visual cues are added. If a bot appears in the same form as the student and has the ability to walk around in a 3-

dimensional environment then it is reasonable to assume that students will expect it to interact with them in a more human-like manner. Another student, in comparing his agent in Active Worlds to another "intelligent" computer program mentioned the ability of the other program to "locate all the red cars." This was reference to a television advertisement where a customer is interested in buying a car and wants to see all the red ones that are available. The inherent simplicity of this particular task is obscured by the speed and accuracy of the software so that the student is mislead into thinking that the program is more intelligent than it actually is. This misconception of what constitutes smart behavior remains a daunting challenge for designers of software agents.

Because students working with a classmate and a bot (group SSB) had the choice of interacting with either a student or a bot it became clear that they viewed each type of learning partner differently. Students tended to view the bots as more of a source of information than learning partner. Interactions with software agents tended to occur only when a student was having difficulty:

"… I didn't know what to do for a while so I asked [my bot] what to do and it answered so I got on track."

Another student commented:

"I do think it is a good option to have a computer partner because if someone says 'What do we do now?' then you can ask…"

Although the bots were programmed to have a greater role than simply answering questions when asked, students expressed irritation when the bots provided unsolicited advice:

"He was kind of annoying; he would say things we already knew how to do."

"I thought it was annoying because it kept saying 'okay, you have to do this… you have to do that'… personally, I wish I could walk out on it."

Therefore, it seemed that the students accepted the bots much more readily in the role of advisor rather than director. This point is supported later in the discussion around the issue of control.

Students working only with another student (group SST) viewed their partners both positively and negatively. On the positive side, these students commented that their partners provided advice when they needed help and were fun to have around for social reasons:

"I liked having the student person because he'd like, he'd wait for you and you know him… it's just better."

But these students also recognized the negative aspects of working with another student. Several commented on how another student could get them off track or slow them down:

"[Working alone] you probably won't get off track as easy and you would probably be able to go through and do it faster."

"… sometimes you partner wouldn't ask you questions… because they wouldn't read the signs or something."

Despite these frustrations, students appeared to enjoy working with another student and many commented on the how much they enjoyed the experience:

"… it was really fun working with a student partner and I learned stuff I didn't know."

"I would tell everybody that it was really, really fun."

<u>Functionality</u>

The limited artificial intelligence of the software agents was often a source of frustration to students. The most common complaint was that the bots would not respond to specific questions that the students asked. Unfortunately, the bots were only capable of replying to a limited selection of keywords in relation to a prior question. This meant that if a student asked a clarification question, for example, the bot might actually record it as the student's explanation and then move on to the next activity.

Some students also commented that the amount of "chat" from their bot partner was too extensive. This was in relation to information provided by the bot when they first arrived at a new concept. The intention was that the bot would introduce the concept and then prompt the student for an explanation. Unfortunately, some students found these introductions too long to read:

"I mean, they would say all this information… they were like, 'macroinvertebrates bla, bla, bla…'" (the student gestures that she was overwhelmed).

As mentioned above, students also became impatient when the bots repeated similar instructions at several points during the exercise. Several comments indicated that they would have preferred the bots to have a more limited role, perhaps only providing advice when called upon for help.

Finally, it did not take students long to figure out the limitations of the bots' abilities to interpret chat messages:

"I typed in gibberish at the very end and it accepted it cause it said, 'do you have any other questions?'"

"…someone could just type in keywords."

An attempt was made to discourage this behavior by telling the students that their explanations were being recorded for use in later activities. It also helped that the exercise was new to the students. However, if the program was to be used over a longer period this might become a serious issue as students would quickly become familiar with the bots' weaknesses.

<u>Control</u>

Perhaps the most interesting theme that emerged from the student interviews centered on issues of control. It became clear that students appreciated the freedom to move about in the world and did not appreciate being lead around by the software agents. In comparing student partners to bots, one student commented that benefit of another student was that they "don't make you go with them." The issue of control over one's actions is undoubtedly more significant in a virtual world setting than other computer environments since students see representations of themselves in the form of avatars that have the ability to walk around a 3-dimensional space. Additionally, the presentation of material can be much more flexible because students can choose their own paths through the world. Although students in this particular exercise were limited to walking or running, the Active Worlds environment does support flying and other novel modes of travel, including the ability to move through solid objects. Several students asked how they might change their avatars to allow them to fly or walk through walls.

Students also expressed an interest in having the ability to choose their own bots:

"I wish we could pick our own robot if we had to work with them."

"We should be able to choose our own bots."

This desire to control their bot partners extended to being able to specify certain characteristics such as gender, appearance, and name:

"Could you make it like, a boy for a boy… and not a girl" (all the bots were female)

"It would be awesome, then I could make it like a space fighter…"

"They had ugly names."

Summary

The student interviews proved valuable in highlighting some of the strengths and weaknesses of using software agents compared to other students as learning partners. One downside of using agents that look and act much like human-controlled avatars is that students have high expectations of their abilities. Their expectations are further raised when the agents take on the role of leader as opposed to a more passive role such as advisor. In general, students were disappointed at the level of interaction that was possible with the bots and often expressed annoyance at the limited way in which a bot could respond to their questions. Students also wanted to be more in control of their virtual world – including having the ability to choose and modify characteristics of their bot partners.

Teacher Interview

An interview focusing on an explanation summary report was conducted with the participating teacher in order to explore how such a report might benefit teachers. The report contained copies of the explanations students had generated for each scientific concept in the environment along with their names. Although this report was produced manually, the goal would be for future reports to be generated automatically at the request of the teacher. The semi-structured interview was organized around three themes: initial impressions, potential uses for this type of report and, other information that might be collected during the virtual world exercises to assist teachers.

The teacher, Mr. H., was pleased to be able to review a printed summary of his students' explanations. Although he acknowledged that the better students wrote more clearly, he noted that the poorer students also had important ideas to contribute:

> "I think it's really important for kids to be able to put comments in their own words. I'm impressed by how they struggle to do that… I think it's more important that they do that than [simply] learn an explanation that I give them. The better students tend to state things more clearly but the kids who are having trouble still have good ideas – they're just not expressed as well. They all seem to get the basics. Everyone seems to have the basic understanding…"

When asked how having to type their explanations might have affected what the students generated, Mr. H. noted that this requirement likely did have an impact. He pointed out several students from whom he had expected longer or more complete explanations:

> “There were a couple of people who gave short responses here who I think are not very good typists. But there are people who gave longer responses because they were faster typists.”

This discussion led to further dialog on how students might be better prepared for working through science modules that required the students to generate explanations. Mr. H. recognized the value of helping the students understand how to formulate an explanation and suggested that the report might be used to assess this aspect of their knowledge by comparing explanations generated during the exercise to ones generated prior to beginning the module. Mr. H. also suggested that a joint training exercise with an English teacher might be a good way to teach the students the basics of generating effective explanations:

> “One of the things that would be helpful would be to help them learn how you write a response, how you do a written response, so I can see us coordinating something with the English teachers. How do you write [an explanation]? Is there some type of organizational system or way to structure your thinking that would help these kids write clearer responses?”

He also noted that reading comprehension for science was not currently addressed in their curriculum:

> “Reading comprehension that I know English teachers do is about, given a story, reading for facts in the story. But I don’t know that any of the texts they used were particularly science related.”

Other information that Mr. H. thought might be useful to show on the teacher summary report included time on task (perhaps collected on a per-concept basis) and some kind of ability rating (pre-assigned by the teacher). The ability rating would be used to evaluate the program in terms of its effectiveness on low, medium, and high performing students.

Finally, Mr. H. made some general comments about the program, noting that an intentional pairing system would be better than random assignment since some students did not work effectively together. He also concurred with the findings from the student interviews, observing that students like to have control over their environment (such as being able to change the appearance of their avatars and having more freedom to move about and help other students).

Summary

The teacher appreciated having a summary of the students’ activities following the virtual world explanation exercise. From his perspective, the report was most valuable as an assessment tool, providing not only a measure of student learning but also a gauge of the program's effectiveness on students with differing abilities. Mr. H. suggested that tracking how long students spent on particular tasks and more information about each student (including some kind of ability ranking) would enhance the report's usefulness. He also recognized that learning how to make effective explanations requires more extensive training and suggested that it be should be included as part of an interdisciplinary curriculum, perhaps in cooperation with English teachers.

Chapter Summary

The results from this experiment provide evidence that software agents can have an impact as learning partners in virtual world activities. The pre-post test findings revealed a significant difference in the amount learned between treatment groups which agreed with the prediction that students working with bots would outperform students paired with other students. There were large qualitative differences in the type of dialog that occurred within each of the three treatment groups. When a bot was present, students remained more focused on the learning task. Students working only with other students tended to spend much more time giving instructions to each other and discussing procedural aspects of the task.

A closer look at learning outcomes for students split into High and Low ability levels provided evidence that High students gained some advantage when working with the software agent. High students working with another student, however, did not show gains on their pre- to posttest scores. When the explanations of High and Low students working with a bot (group SBB) were analyzed, it was found that the High students generated significantly more deep explanations than the Low students. This suggests that the bots did not provide the necessary support to students who needed more guidance in how to generate explanations.

The student interviews indicated that while students enjoyed the virtual world experience regardless of which partner they were assigned to, they would have appreciated more control over the world and the characteristics of the software agents. It also seemed evident that students' expectations of the bot capabilities were not met. Students were frustrated that the agents had only limited ability to answer questions and to chat about other topics.

The interview with the participating teacher provided some ideas for future teacher reports and the exercise in general. For the reports, Mr. H. suggested that including more information on the ability of each student would help evaluate the software's effectiveness on different types of students. He also offered that having access to data such as time spent on specific tasks would help assess how the students worked through the module. For the overall exercise, Mr. H. recommended more extensive training for the students in cooperation with an English teacher so that they might develop a greater understanding of how to generate effective explanations.

CHAPTER V

DISCUSSION AND CONCLUSIONS

Summary of the Study

The aim of this study was to investigate the use of software agents as learning partners in an activity where students generated explanations for scientific concepts involved in river ecosystems. An explanation exercise was chosen because previous research has demonstrated that students who make explanations tend to encode new information in a richer manner, leading to deeper understanding and more accessible knowledge (Chi et al., 1994; Coleman, 1998). The theory is that the process of generating explanations encourages students to identify underlying principles, link new information to prior knowledge, and to monitor their own comprehension more effectively (Chi et al., 1994). Unfortunately, eliciting explanations that stimulate these deeper ways of thinking from students is not straightforward. Most students, who do not receive training or specific prompts, tend to generate shallow explanations that do not lead to deeper understanding (King, 1994). Because of this problem, education researchers have explored different ways of supporting the explanation process. Successful interventions that have been developed for classroom use usually involve extensive training periods for the students and require a heavy commitment of teacher time. Through the use of software agents (or bots), this study explored one way in which technology might help overcome these challenges.

This chapter begins with a review of the original research questions and examines how they might be answered based on the findings of the study.

Effectiveness of Agents on Resources

The first research question asked if an animated software agent could increase the use and value of explanation resources over a text-based help system. The point of this question was to help focus the study on the investigation of learning partners as a means to support an explanation activity. In other words, if the agents provided resources, would students use them more than if the resources were presented as text-only web pages?

While there was no direct way to determine if a student in the text-only condition actually read or skipped over the advice on generating explanations, there was evidence from the chat transcripts that revealed a greater use of the advice by students who interacted with the agents. Thus, utilization, rather than direct measure provides support for the greater effectiveness of agents over text. This use was measured both in the number of explanations generated as well as the quality of the explanations.

The analysis of explanations revealed that SBB students generated significantly more "deep" explanations than the other groups. It should be noted that this difference was not significant for the SSB group where the students worked in pairs with an agent. One possibility is that the presence of another student made if less likely that either of them would attend to the agent. This contention is supported by the higher percentage of chat message coded as non-explanation-related by the SSB versus the SBB group. One might conclude then, that agents can encourage the use of resources but that other factors, such as distractions from other students, may have a negative impact.

The question also asked if the value of the resources would be greater if they were agent-based. Although students who interacted with agents did incorporate more of the resource advice into their explanations, this does not provide insight into the value of the resources. One measure of value is how much students learned. Since students who obtained resources from the agents performed better on the post-test, it is reasonable to conclude that indeed the agents did have an added value in their supporting role.

Effectiveness of Agents on Explanations

The second research question asked if an animated software agent could evoke more focused and thoughtful explanations than another student as a learning partner. This question was designed to help determine what benefits an agent might provide when positioned as an active partner in the learning process. There is certainly evidence that the agents helped the students stay on task and so encouraged a more focused dialog. The chat transcripts revealed that students working alone with an agent spent about 50% of their time prompting or making explanations. This compares with about 15% and 12% for the SSB and SST groups respectively. It might also be argued that the SBB students generated more thoughtful explanations as well since they made significantly more "deep" explanations than the other groups. However, it is important to point out that this analysis only provides a partial picture of a student's thinking. The limitations of the agents' intelligence definitely had an impact on the breadth of conversation possible between the student and the agent so that there were many occasions when a student's thinking was curtailed. The tradeoffs between shorter, more focused conversations versus those that have less focus but fewer constraints remain unclear. This topic, as well as a discussion about the perceived role of an agent, are further elaborated in the sections below on implications for research and design.

Summary Reports for Teachers

The third research question asked how a summary report of student explanations might be useful to teachers. The purpose of this question was to direct part of the research toward ways of supporting teachers. Although this preliminary effort was meant to be exploratory in nature, the interview with the participating teacher provided some useful insights. The teacher's enthusiasm for the report and his request for more information indicates the potential value of providing more feedback for teachers when students engage in computer-based learning environments. Too often software is designed entirely around the student and the teacher is left on the sidelines. This approach encourages teachers to adopt the view that computer programs are effective "substitutes" instead of an integral part of a larger learning process. Designers need to examine ways of providing more feedback to teachers, not just in the form of student tracking and performance measures, but creative ways in which student generated output can be used in further classroom activities.

Implications

The results of the study suggest that the use of software agents as learning partners and resource providers can be successful in helping students learn by explaining. The goal of this section is to discuss these results, along with insights gained from the virtual world chat transcripts and interviews, within a framework of theory, research, and design. The first section on implications for theory discusses how the findings of this study fit within and add to existing theories on learning by explaining. In the second section, implications for research, the discussion focuses on three questions for future work that emerged from this study. Finally, implications for software design are discussed in the third section where both technological and pedagogical issues surrounding the use of software agents as learning partners are considered.

Implications for Theory

Self-explanation theory, as promoted by Chi et al. (1994) and others, contends that learning with understanding requires new information to be integrated with existing knowledge. The activity of self-explaining, it is thought, facilitates this process because it is constructive in nature, promotes links between prior and new knowledge, and is performed on a piecemeal basis (Chi et al., 1994). Several researchers, however, have noted that most students do not spontaneously engage in the kind of explanatory activity required for deep understanding (King, 1990; Coleman, 1998; Chi et al, 1994). When asked to explain by an untrained questioner, most students will summarize information that has been provided or simply make a declarative statement. One problem is that the explainers do not have the necessary skills to formulate their explanations in an effective way. The other problem is that untrained prompters tend to formulate questions that encourage a shallow response (King, 1990). Results from this study support these contentions.

Although all participants were prompted for explanations, students in two of the conditions (SSB and SST) were prompted by their student partners, who had not received any specific training in formulating prompts. Students in these groups tended to simply restate information provided in the webpage resources. For example, for the concept about the importance of macroinvertebrates in river systems, the following exchange took place between one SST pair:

Student 1 (prompt): “What do macroinvertebrates eat?”

Student 2 (explanation): “Macroinvertebrates eat bacteria, dead decaying plants and animals.”

In contrast, students working with the bots were prompted with a question designed to promote deeper thinking. Here is part of one SBB pair exchange:

Bot prompt: “What would happen to the bacteria in a stream if there were no macroinvertebrates?”

Student Explanation: “What would happen is the health would drop in the river.”

In this example, the bot has asked the student to make a prediction based on information contained in the concept webpage. The student, in creating an explanation based on a prediction, has made the causal connection explicit between the presence of macroinvertebrates and the idea of river health. In general, it was found in this study that students who received carefully constructed prompts from software agents were more likely to generate deeper explanations (SBB=5.2 versus SST=0.8) and learn more. Although these findings match what current explanation theory would predict, there was considerable variation in the effectiveness of the intervention. Some of the factors contributing to differences found within the treatment groups are discussed next.

Individual Differences

In any random sample of students there will be a certain percentage of high performers. Researchers who have explored the role of explanations in learning have pointed out that some students already possess the necessary skills to engage in the types of cognitive processes (e.g. linking new information to prior knowledge) and self-monitoring that promote deeper understanding (Chi et al., 1989). Coleman (1998) refers to these students as high intentional learners. In her study, participants were assigned to groups based on their responses to an interview designed to identify students with strong problem solving skills (Individual Implicit Learning Theory Interview). Students scoring 80% or higher in the interview were placed in a separate high intentional learner group (HIL). Coleman (1998) obtained significant results based on treatment and control groups made up of *non*-HIL students, so that she had effectively factored out the high intentional learners beforehand. In this study, no such pre-selection was performed so that it is likely that the presence of some high intentional learners diluted the effects of the intervention and produced a larger variation in the quality of explanations generated.

The previous argument points out that some students will not benefit especially from specific instruction on making explanations. For most students, however, generating explanations is not a mastered skill and they do generally benefit from carefully designed prompts. Unfortunately, success is not guaranteed. Some of the reasons why an intervention may lose effectiveness are discussed below.

Support Issues

In the SBB group where students worked alone with a bot partner, some students did not generate good explanations even though they received carefully designed prompts and remained more or less “on-task”. One cause of failure was a result of an inadequate help system. The bots had no ability to provide specific help and instead either attempted to encourage the student to say more or, simply moved on to the next topic. The following example illustrates this problem:

> Bot: “For this explanation, try to explain what is most important and why.”
> Student: “I don’t understand.”
> Bot: “So, what was the key point?”
> Student: “about what?”
> Bot: “Well, I guess we are off to sign 4 [the next concept]. I’ll meet you there!”

In some instances, the failure occurred because the explanations were very short or incomplete. For example, when one student was asked to make a prediction about the bacteria population in a river with no macroinvertebrates, the student stated correctly that “the number of bacteria will increase.” However, this explanation is inadequate without an explanation for why the population would increase. Although the bot asked the student if she had anything else to add, the student simply replied, “No.” In this case, while the student may well have been able to expand on what she had said, the opportunity was lost because the prompt was non-specific. Again, the situation could have been avoided if the bot had been able to interpret what the student had typed. These failures highlight the potential benefit of targeted feedback for students who may require extra help in formulating complete explanations.

Learning Partner Effects

What impact does the type of prompter/listener have on the effectiveness of interventions that use explanation prompts? In their 1994 study, Chi et al. achieved successful learning outcomes when students were prompted to explain to themselves. Although the participants in the study were self-prompting, some have pointed out that the presence of the researcher constitutes at least some form of audience (Renkl et al., 1998). It is possible that students were more thoughtful about their prompts and explanations because the researcher was present.

One study that did attempt to examine the influence of the listener was performed by Coleman et al. (1997). In the study, they compared the effects of explaining and summarizing to oneself versus a fellow student. Their design did not, however, take into account feedback from the listener nor did it explore the effects of different types of listeners on the explainer.

The software agents in this study were positioned as experts and there was some evidence that this relationship might actually impede learning because students had the impression that their bot partners "knew" the correct answers. There was a tendency among some students to simply ask for the "right" answer instead of trying to struggle with understanding the concept themselves. One student, who was prompted by her bot partner to make a prediction appealed to the bot for help, "Well, I'm not sure… do you have any info?" Another student, when asked to explain the key point of a concept said, "Oh, um the key point is that, well, I'm not sure. Can I ask you what the key point is?" Clearly, these students were that the expert would simply tell them the right answer. While a human tutor or mentor would be able to redirect the student and likely help them find the answer themselves, a limitation of the software agents, as pointed out earlier, is that they do not have the ability to provide this kind of sophisticated feedback. Regardless of the type of feedback that an explainer receives, however, it seems evident that the nature or status of the learning partner may have an effect on how students approach an explanation task. This topic is examined further in the section on implications for research.

Training

Although most researchers agree that training is an important part of helping students understand how to learn through explanations, there is no general consensus on how much training is required or what form it should take (c.f. Chi et al., 1994; King, 1994). Solutions range from providing only prompts to elaborate training sessions that attempt to teach students not just how to explain but also how to formulate deep questions that can elicit explanations. This study compared the use of web pages verses software agents as modes of delivering information on how to explain. Both the SBB and SSB groups interacted with bots that were programmed to provide advice on how to ask questions and make explanations. For the SBB students, the bot also acted as their explanation partner, whereas the SSB students asked questions and explained to their student partner.

Students in the SBB group obtained advice from the bot but also had the advantage of being exposed to the bot modeling an example question and explanation. These students were much more successful in generating more effective explanations than either of the other groups. Although the SSB students also received advice from their bot, they were much less inclined to use effective explanation strategies than the SBB group (1.1 versus 5.2) and only marginally better than students in the SST group (1.1 versus 0.8). An important difference between the SBB and SSB groups was that the SSB students had a student partner to talk to as well as the bot. This meant that they could (and often did) ignore the bot and so did not make use of the advice that it provided. SBB students, on the other hand, had no one else to talk to and so were obligated to pay closer attention to their bot advisor.

The student pairs in the SST group did not interact with a bot but instead were told to read about making explanations on a webpage. Providing explanation resources via a webpage did not prove useful as most students did not read or chose to ignore the text provided on the webpage. These students tended to ask questions and generate explanations on a superficial level (see table 9).

The failure of many students to attend to instructional resources on how to make explanations highlights the importance of creating an atmosphere where there is some incentive to apply new skills. Researchers who have had students successfully use explanation strategies invariably set up an environment where students feel compelled to pay attention to and follow instructional advice. This may be through pressures introduced unintentionally, such as self-explaining in the presence of a researcher (e.g. Chi et al., 1994), or by establishing a standard of practice that is reinforced through extensive training and frequent use (e.g. King, 1994). In these situations, it is the social nature of the explanation task that is critical. Unfortunately, students were not given the opportunity in this experiment to establish a meaningful relationship with their bot advisors and so did not experience a sense of obligation or commitment that might have otherwise developed. The sociocognitive aspects of student-agent interactions are explored further in the section on implications for design.

<u>Summary</u>

In general, this study produced results that fit well with established theory on learning by explaining. However, the significant amount of variation in the results suggests that several factors may have considerable influence on the success of this type of learning activity. Aside from a normal variation in students' abilities to explain, the degree of success in learning by explaining is likely affected by the availability of individualized support, the nature of the explanation partner (e.g. expert versus novice), and the social conditions under which the activity takes place. These issues are explored further in the following sections on implications for research and design.

Implications for Research

Advances in both hardware and software are enabling programmers to develop much more sophisticated and realistic software agents than was possible even a few years ago. However, there remain many questions about how these agents might be used effectively in learning environments. The following discussion outlines some areas of research that may provide useful insights based on the results of this study and the work of others.

The Social Role of Agents

One of the most compelling aspects of software agents is that they can embody much more than simple instructions or advice. By their very nature, software agents have the ability to possess personalities and to exhibit specific behaviors. Indeed, this capability can be exploited for specific purposes. In this study, for example, the agents were presented as experts so that they could act as advisors in order to provide students with information about making explanations. The goal was to determine whether providing resources in this manner would be more effective than offering the same information on static web pages. The agents did, in fact, help students utilize the information since the students who worked with bot partners made better explanations. However, even though there are potential benefits to the use of "expert" agents, there is some evidence that this role may not always be ideal.

One clue comes from prior research that has examined how students study in preparation to teach others. The Teachable Agents Group at Vanderbilt (Biswas et al., 2001) found that students who learned by preparing to teach their peers spent more time studying and felt obligated to gain a deeper understanding of the material. These findings suggest that students experience an increase in motivation when faced with the prospect of teaching a peer or novice. A future study is needed to explore the differences in motivation of students who have to explain to an expert versus a novice agent to see if these effects carry over to virtual learning partners.

There may also be benefits in how well explanations are constructed when students explain to a novice instead of an expert. In classroom settings, students are not normally given opportunities to explain (Graesser et al., 1995). Instead, teachers typically ask students questions designed to elicit only a right or wrong answer. It is likely then, that this classroom practice leads students to think that an explanation is simply an attempt to convey the "right answer". Would explanations crafted for novices draw more upon underlying principles? Cawsey (1993), who works on human-computer interactions, makes the following observation:

> "When explaining how something works to a novice, you don't expect to get across all the subtleties that you might explain to an expert, yet the explanation might take at least as long as it needs to be explained more carefully." (p. 19)

There is some evidence that when students explain to novices, and have the opportunity to obtain feedback, they do generate explanations that are more carefully constructed, involve more causal relationships, and contain less irrelevant information (Leelawong et al., 2002). In their work with a teachable agent program known as Betty's Brain, Leelawong et al. (2002) have shown how students learning about ecosystem concepts need to develop rather detailed explanations for the agent Betty, in order for her to answer questions correctly. Since Betty is a novice, the students are unable to avoid explicitly stating important underlying causal relationships. When explaining to an expert, on the other hand, it is easy to gloss over important details since students presume that the expert really understands the material anyway. More research needs to be done to explore the relationship between explainers and listeners.

If classroom practice has indeed affected how students think about explaining, it is important to examine ways in which their perceptions can be made explicit and, hopefully, broadened. Analysis of the interactions that occurred when students explained in this study supports the notion that students view explaining as a test of factual knowledge. The following exchanges illustrate this point:

(Student pair 18-19)

Student 19: "What are macroinvertebrates?"

Student 18: "They are things without a backbone and you can see without an aided eye."

Student 18: "Did I get it right??"

Student 19: "Ya mon, you did!"

(Student pair 16-17)

Student 16: "What is pollution?"

Student 17: “Pollution is garbage that gets thrown into the water.”

Student 16: “No”

Student 17: “It’s something… it’s anything that can hurt the water let’s say.”

Student 16: “Yes.”

Making Thinking Explicit

One challenge in getting students to change how they think about a concept is simply getting them to recognize their own thinking in the first place. Lin et al. (2001) developed a system to explore students’ perceptions of what constitutes an ideal student. In that exercise, students were asked to write down five characteristics of an ideal student for two different types of classrooms. Results from that study revealed a remarkable trend for public schools students to associate behavioral characteristics with being a good student. A similar approach might be used to probe students’ perceptions of the use and value of explanations. More importantly, if the exercise involved different types of scenarios, such as explaining something to a friend versus a teacher, the results might be used to help students develop a deeper understanding of how explanations can be used for learning, not just assessment.

Positioning Agents for Effective Support

More research is needed to develop effective ways of supporting students who learn by explaining. In this study there was a large amount of variation in how well students explained, even among those in the SBB group who received the most direct support. Results from both the student interviews and the Active Worlds chat logs suggest that students become frustrated with a one-size-fits-all approach to support. Many students, whether they needed help or not, became annoyed when the bots offered unsolicited advice. Recall from the interview data how one student commented, “I thought it was annoying because it kept saying ‘okay, you have to do this… you have to do that’… personally, I wish I could walk out on it.” Another student expressed his feelings more abruptly: “I hate you, Beth.” Clearly, these students were not happy with their advisors. The challenge is to provide appropriate levels of support to students whose needs vary.

Baylor (2002) suggests a couple of ways to avoid the problem of advisors that are too intrusive. One possibility is to take advantage of the ability of agents to provide more subtle forms of feedback. For instance, simple gestures like winking or nodding can convey a significant amount of meaning without the need for excess verbiage. Having the ability to select a level of support might also alleviate frustration caused by excessive interference.

Providing adequate support for learning goes beyond controlling how much help is provided. The type of support deserves equal attention. Some designers have concentrated on developing systems that recognize precisely where students are experiencing difficulties. Programs that attempt to diagnose errors, such as that found in expert tutoring systems like the Algebra Tutor (Anderson et al., 1995), are highly complex and require a complete mapping of the domain space and sophisticated diagnostic algorithms. A different approach to support is through the use of multiple perspectives. This method may also be more appropriate for domains that involve ill-defined problems and where subjective opinions are valuable. Presenting different viewpoints during problem solving helps students appreciate different facets of a problem and lets them compare their thinking to others (Schwartz et al., 1999).

Baylor (2002) is conducting some initial research along these lines with an environment known as MIMIC (Multiple Intelligent Mentors Instructing Collaboratively). The preliminary version of MIMIC includes two types of agent-mentors designed to help pre-service teachers learn about instructional planning. One agent is characterized as an *instructivist* and, as such, presents a teacher-driven approach. The other agent is presented as a *constructivist* and focuses on the learning context, the importance of active participation, and emphasizing learning processes rather than products. Results from her initial study suggest that pre-service teachers were able to appreciate differences in the pedagogical approaches and that the agents were capable of taking on the role of knowledgeable and believable mentors with differing views. This multi-agent approach offers some intriguing possibilities for future study. For instance, it would be interesting to explore how students react to agents that have a different perspective than their own. Would these types of interactions help students develop a better understanding of their own ways of thinking?

Summary

The ability of software agents to take on different personas provides new opportunities and challenges for designers of learning environments. Future studies will need to provide insights on when and how different types of agents should be implemented. Results from this study indicate that it may not always be a good idea to combine different roles in a single agent. There may actually be advantages to developing environments with a variety of agents with distinct roles and personas, such as being able to offer several levels of support or a variety of perspectives. Observing the effects of different viewpoints exhibited in the behavior and actions of multiple agents might help students develop a deeper understanding of their own ways of thinking.

Implications for Design

In this section, three aspects of learning environment and software agent design are considered. The first involves the concept of user-control and how it relates to motivation, self-regulation, and instruction. The second aspect relates to the goals of the explanation task and examines the importance of intentionality. Finally, the relationship between software agent design and student expectations is examined in the context of student-agent interactions.

User Control

One of the challenges in designing software learning environments is achieving the right balance of instructional guidance and learner control. The literature on self-regulation suggests that providing sufficient opportunities for students to take control of their own learning is an essential component of successful environments (Baylor, 2002). For example, Zimmerman (2002) points out that students need to be given the chance to make choices about learning tasks, methods of completing assignments, and who they work with, in order to encourage the development and use of self-regulation strategies. How does the design of software agents affect the ability of students to self-regulate?

If software agents are placed in the role of mentor or tutor, there is a temptation for designers to create agents with a large degree of control over the learning process. In this study, for example, the bots were designed to have a significant amount of control over the dialog and general procedure in the explanation and river monitoring activities. In part, this design decision was driven by the limitations imposed by the simple artificial intelligence available to the bots.

Since the bots had very limited ways of responding to students, it was important that their messages controlled the conversation. Another reason for putting more control in the hands of the bots was to encourage students to remain on task. This was deemed important since the learning environment was a new and exciting place and so likely to incite a significant amount of off-task behavior (such as exploring, changing avatar appearances, etc.). However, in retrospect, the high degree of control exerted by the bots likely had an adverse effect on learning because it discouraged self-regulatory processes. In particular, the bot design suppressed several dimensions of learning that Zimmerman and Risemberg (1997) propose as key elements of self-regulation. These include motivation, methods of learning, use of time, physical environment, social environment, and performance.

Agents that control the learning process can have the same adverse affect on motivation as many teachers who do not prepare or assist students to learn on their own (Zimmerman, Bonner, & Kovach, 1996). For instance, it has been shown that students who are given opportunity to set their own goals exhibit higher levels of intrinsic motivation (Dembo and Eaton, 2000). If the teacher or agent is always setting learning goals, then a valuable opportunity to build intrinsic motivation is lost. A similar argument can be made for the other learning dimensions mentioned above. The key then is to develop agents that do not dominate the learning process but, rather, facilitate and encourage students to engage in self-regulatory practices.

In her work with intelligent agents at MIT, Maes (1997) advocates an approach where the agent begins a relationship with a new student assuming very little control and displaying only limited abilities. As the student works with the agent, it gradually changes the way it behaves based on the student’s actions and feedback. In this way, the student develops a sense of trust and understanding for the agent and the agent develops appropriately to meet the needs of the student. This co-developmental approach represents an exciting direction in agent-based research and opens up some interesting possibilities for future work.

The second aspect of learning environment design under consideration relates to the particular learning task in which are students engaged. The activity in this study really was a combination of two related, but separate tasks; river monitoring and explaining ecosystem concepts. Although both tasks were oriented around river ecosystems, it is significant that they were not connected in a more meaningful way from the perspective of the students. That is, there was no obvious way in which explaining the scientific concepts would directly help them with the river monitoring task. The monitoring task was presented as an activity that simply required the students to follow a series of logical steps in order to determine the pollution status of the river. Being able to explain the concepts to either another student or a software agent provided no direct connection to successfully assessing the river's status. Thus, the explanation task lacked what Bereiter and Scardamalia (1989) refer to as intentionality.

Intentionality

Intentionality, as described by Bereiter and Scardamalia (1989), refers to a particular goal-oriented approach that a learner can take when acquiring new knowledge. Coleman (1998) points out that "intentional learners tend to identify problems that need to be solved or explained, ask insightful and critical questions, and attempt to explain unexpected events… instead of directly assimilating new knowledge into what was previously known" (p. 389). Because having intent usually means learning for the purpose of solving a problem, intentional learners are more likely to organize new knowledge around a problem space rather than a topic and in doing so, engage in a learning process that leads to deeper understanding and more accessible knowledge (Bereiter, 1992).

Although some researchers, such as Chi et al. (1994), have successfully used prompts in a non-problem oriented way, it seems likely that relating explanations to a particular goal would improve learning outcomes. In this study, the lack of intentionality in the explanation process meant that students focused more on answering the prompt question rather than attempting to convey meaning. A better approach might be to organize the activity around a problem that the agent must solve after being taught by the student. In this way, the student would attach value and meaning to the level of understanding achieved by their partner. The challenge, however, is creating software agents that are able to interpret the meaning of explanations provided by the students. One possible method would be to provide a set of pre-defined "explanation components" so that students could construct an explanation in a way that could be easily interpreted by the software. These components might be in the form of selectable phrases that the students would use to form complete sentences. More work is needed to explore the potential of this component-ware design approach.

Expectations

Finally, how does the design of software agents influence what students expect in terms of interactivity? Shneiderman (1997) warns that it is unwise to encourage software agents to be compared to humans because users will automatically make unrealistic assumptions about their intellectual abilities. There is evidence from this study that the use of sophisticated graphics and software agents with human-like appearances and behaviors might lead students to believe that the software is highly intelligent. In the Active Worlds environment, software agents have the same physical characteristics and mode of communication (through chat messages) as any avatar that represents a human. In light of this, many students seemed disappointed when their bot learning partner could not respond to questions or comments that deviated from the its scripted dialog. Some students tried to get their bot to make value judgments on their explanations. For example, on student asked, “Do you like my explanations?” Another student wondered if her bot was ready to move on to the next concept: “I’m finished reading, are you?” Several students asked their bot personal questions like “Hi, are you a girl?” and “Are you married?” Questions such as these suggest that students expected their bots to be able to interpret natural language and carry on a normal conversation. Comments made in the interviews following the exercise support this point: “I didn't like it because you couldn't like, talk to it.” and “You can’t talk about the [concepts] together.”

How then, can agents be designed so that students accept a more limited and realistic level of agent intelligence? Baylor (2000) suggests that successful student-agent interactions may depend on ensuring that students understand how the agents “think”. She promotes a “glass box” rather than “black box” design where the agent’s reasoning more visible to users. With this approach, students can gain an understanding of the agent’s limitations and can modify their interactions accordingly to achieve more useful results. For the type of avatar agents used in this study, one possible approach would be to provide more feedback so that the bots would explicitly tell students that they didn’t understand a question or comment. It would also be useful if the bot explained their limitations upfront so that students would have a better initial understanding of the bot’s capabilities. This method might encourage students to interact with their bot partners in a more realistic and understanding way.

Summary

Because the personas exhibited by software agents can have an important influence on student motivation, designers must be careful to ensure that agents facilitate, rather than dominate, the learning process. Agents that have the ability to develop or grow intellectually along with their student partner may provide one way fostering intrinsic motivation and help students take control of their own learning. Intentionality is also an important factor to consider when designing an explanation activity. Orienting the activity around a problem-based teaching exercise may promote a more intentional experience. Finally, avoiding unrealistic expectations for the abilities of an agent may be possible by ensuring that an agent's thought processes are made explicit and not hidden in a "black box".

Conclusions

The main goal of this study was to explore the animated software agents as learning partners in a virtual world environment. An explanation task was chosen as the activity around which to investigate the agents because it is both a social task and one that requires a significant amount of support. To focus the study, data analysis concentrated on answering several research questions regarding the use of agents and their effect on learning. Overall results showed that the agents did add value and encouraged the use of explanation resources albeit with certain caveats. Students generated better explanations more frequently when they received prompts from the software agents. The agents added value to the resources because these students also performed better on the posttest. Based on the analysis of explanation content, it was concluded that the agents did promote more carefully thought out explanations. Unfortunately, averages alone do not tell the complete story when evaluating educational interventions.

Detailed analysis revealed that not all students benefited equally when partnered with an agent and there remain significant support challenges that must be overcome before a system similar to the one used in this study could be incorporated into regular classroom use. Although some might argue that a system lacking sophisticated artificial intelligence will never be able to meet the support needs of a typically diverse classroom, there are several design considerations that may refute this view. Generally, the goal of supporting students needs to shift from a notion of leading to one of facilitating and enabling. This means designing agents that are not necessarily libraries of information but, rather, promoters self-regulation. Instead of attempting to create an agent that will always know the correct answer, designers need to invent agents that encourage students to *want* to know the correct answer. Agents used in this study were presented as experts whose ways of thinking were constrained and hidden. This approach dampened intrinsic motivation and often led to shallow thinking and frustration. The suggestions discussed in the preceding sections, such as providing agents that co-develop along with their student partners and making the agents' thinking visible, hold possible answers to some of design challenges encountered here and are worth exploring in future studies.

Finally, a significant problem with much of today's educational software is that little is provided for the teacher in the way of feedback or other information about student progress. An additional question investigated in this study was if an explanation report generated following the activity might be beneficial for the participating teacher. Results from an interview indicated that the teacher very much appreciated having more insight into his students' work and he was quick to suggest additional features. It is critical that future efforts to develop learning environments recognize the important role of teachers and ensure that they have access to tools and information that will help promote learning among their students.

APPENDIX A

IRB LETTER OF CONSENT FOR STUDENTS

Dear Student,

The following information is provided to inform you about the research project and your participation in the study. Please read this form carefully. Please feel free to ask any questions you may have about this study and the information given below. You will be given an opportunity to ask questions, and your questions will be answered. You will be given a copy of this consent form.

We are researchers at Vanderbilt University interested in how students learn. With this letter, we are asking you to take part in a research project about how students solve science problems. You do not have to take part if you do not want to. If you agree to be in this study, but decide later change your mind, you can drop out any time you want.

So we can learn more about how you solve science problems, we will ask you to use a computer-based virtual learning environment. Specifically, we will ask you to enter the world and determine how a river has become polluted. Some students will solve this problem in pairs, others will solve the problem individually. If you decide to participate, the study will take place during the school day in your science classroom. . You will also participate in groups discussions following the exercise and several of you will be asked to take part in a brief interview so that we can find out more about what you felt about your experience.

The study will take place during two of your regular science class periods. Students who decide not to participate in the study, will be assigned regular class work by your teacher.

As part of the study, we are asking your permission to look at your science grades for the current school year. Whether you participate in the study or not, no difference will be made in your grades. If you decide to participate, you will be asked to take home a letter similar to this one for your parents to sign and return. Your teacher and your parents will not see your individual answers. All information you provide will be labeled with an identification number rather than your name. We will keep a list matching your number and your name in a locked cabinet at Vanderbilt University. Only researchers at Vanderbilt University see your answers. Any audio recordings will be confidential and only used by the researcher (Jeff Holmes) and his advisor. The tapes will be stored in a locked cabinet in the researcher's office.

Participation in the study will give you a chance to experience a "virtual field trip" an experience not usually available in the classroom. For this reason, we believe the study will be helpful in understanding how students solve science problems and how technology can help.

If you should have any questions about this research study, please feel free to contact Jeff Holmes (343-2614) or my Faculty Advisor, Bob Sherwood at (343-2596). For additional information about giving consent or your rights as a participant in this study, please feel free to contact the Vanderbilt University Institutional Review Board Office at (615) 322-2918 or toll free at (866-224-8273). If you are participating in a research project at the VA Medical Center, please contact the Research and Development office at (615) 327-5346.

APPENDIX B

IRB LETTER OF CONSENT FOR PARTICIPATING TEACHER

The following information is provided to inform you about the research project and your participation in the study. Please read this form carefully. Please feel free to ask any questions you may have about this study and the information given below. You will be given an opportunity to ask questions, and your questions will be answered. You will be given a copy of this consent form.

We are researchers at Vanderbilt University interested in how students learn. With this letter, we are asking you to allow the students in your classroom to participate in a research project investigating how students solve science problems. You do not have to participate if you do not want to. If you agree to be in this study, but decide later change your mind, you can drop out any time you want.

So we can learn more about how your students solve science problems, we will use a computer-based virtual learning environment. Specifically, we will ask approximately 80 students to enter the world and determine how a river has become polluted. Because one goal of the study is to understand how explanation facilitates problem-solving, some students will work on the problem in pairs, others will work on the problem individually. Students will also participate in groups discussions following the exercise and several students will take part in a brief interview so that we can find out more about what they felt about their experience in the learning environment.

If you agree to participate, the study will take place during the school day in your classroom. Initially, we will need instructional time to describe the study to students and answer their questions about participation. We will also need your help in sending home and collecting letters of informed consent for a parent to sign. Additionally, the study will require access to computers. Ideally, these should be located in the science classroom. Once the study is underway, participation should not take more than 2 class periods away from students' attendance in science class. Finally, we are asking for permission to look at participating students' science grades for the current school year.

All information you and the students provide will be labeled with an identification number rather than your name. We will keep a list of matching student ID numbers and student names in a locked cabinet at Vanderbilt University. Identification of students by name will only be possible by matching subject identification numbers to letters of informed consent. Any audio recordings will be confidential and only used by the researcher and his advisor. The tapes will be stored in a locked cabinet in the researcher's office.

We believe the study will be helpful in understanding how students solve science problems and how technology can help. Furthermore, participation gives students a chance to experience a "virtual field trip," an educational experience not usually available in the classroom.

If you should have any questions about this research study, please feel free to contact Jeff Holmes (343-2614) or my Faculty Advisor, Bob Sherwood at (343-2596). For additional information about giving consent or your rights as a participant in this study, please feel free to contact the Vanderbilt University Institutional Review Board Office at (615) 322-2918 or toll free at (866-224-8273). If you are participating in a research project at the VA Medical Center, please contact the Research and Development office at (615) 327-5346.

Sincerely, Jeff Holmes

APPENDIX C

IRB LETTER OF CONSENT FOR PARENTS

Name of student ____ Age__________

Please read this form carefully. Please feel free to ask any questions you may have about this study and the information given below. You will be given an opportunity to ask questions, and your questions will be answered. You will be given a copy of this consent form.

We are researchers at Vanderbilt University interested in how students learn. With this letter, we are asking you to allow your child to participate in a research project investigating how students solve science problems. Your child does not have to participate if s/he does not want to. If your child agrees to be in this study, but later change her/his mind, s/he can drop out any time s/he wants. So we can learn more about how students solve science problems, we will use a computer-based virtual learning environment. Specifically, we will ask your child to "take a virtual field trip" in order to determine how a river has become polluted. Because one goal of the study is to understand how explanation facilitates problem-solving, some students will work on the problem in pairs, others will work on the problem individually. Students will also participate in groups discussions following the exercise and several students will take part in a brief interview so that we can find out more about what they felt about their experience in the learning environment.

If you agree to your child's participation, the study will take place during the school day in your child's science classroom. Students who do not participate in the study will simply attend their regular science class. Whether or not your child participates in the study, no difference will be made in his/her grades. However, as part of the study, we are asking for permission to look at students' science grades for the current school year. Only researchers at Vanderbilt University see your child's answers.

We will keep a list of matching student ID numbers and student names in a locked cabinet at Vanderbilt University. Identification of students by name will only be possible by matching subject identification numbers to letters of informed consent. Any audio recordings will be confidential and only used by the researcher and his advisor. The tapes and consent forms will be stored in a locked cabinet in the researcher's office.

We believe the study will be helpful in understanding how students solve science problems and how technology can help. Furthermore, participation gives students a chance to experience a "virtual field trip," an educational experience not usually available in the classroom.

If you should have any questions about this research study, please feel free to contact Jeff Holmes (343-2614) or my Faculty Advisor, Bob Sherwood at (343-2596). For additional information about giving consent or your rights as a participant in this study, please feel free to contact the Vanderbilt University Institutional Review Board Office at (615) 322-2918 or toll free at (866-224-8273).

Note: Immediate necessary care for adverse events will be provided at Vanderbilt University without charge if you are injured because of participation in this research project. Vanderbilt will neither provide for the costs of further treatment beyond immediate necessary care nor provide monetary compensation for such injury.

Sincerely,

Jeff Holmes

APPENDIX D

STUDENT INTERVIEW QUESTIONS

Questions Regarding Resources:

- How did you like the help that you received during the virtual world exercise?
- What advice would you give to another student about using the resources?
- Did the resources (or bot advice) help you to make better explanations? If so, how?

Questions Regarding Learning Partners:

- How did you like working with your partner?
- If you were to do this exercise again, would you prefer to work with a partner or work alone? Why?

APPENDIX E

TEACHER INTERVIEW QUESTIONS

The following questions will be used in a semi-structured interview:

- Did you learn anything new about your students from these reports?
- Did anything surprise you?
- Could you list a few ways that might you use these reports?
- Which report do you think best highlights students with comprehension problems?
- How could we design these reports to be more useful in the future?

APPENDIX F

SAMPLE EXPLANATION DIALOG

GROUP SBB (Erin=Bot; Student=expert5)

[Erin]: When you are ready, click on Sign 1 and read about our first concept in the web page!
[Erin]: Explaining something to another person is a really good way to learn.
[Erin]: So we are going to make up questions and explanations!
[Erin]: The kind of question you ask someone is important too. We want to ask a question that makes the other person think hard about the concept. Got it so far?
expert5: yes
[Erin]: Great!
[Erin]: One good kind of explanation is one that makes a prediction.
[Erin]: For example, after reading about Benthic Macroinvertebrates I want you to explain
[Erin]: what would happen to the bacteria in a stream if there were no macroinvertebrates?
[Erin]: So, now you try to answer that and I'll start recording now!
expert5: No one would know there was bacteria in the stream, and the river would get even more polluted
[Erin]: Oh, so I think you are saying there would be more bacteria?
expert5: no, the bacteria would, well I'm guessing multiply in time
[Erin]: Okay, I think I got it. Anything else to add?
expert5: nope
[Erin]: Ok, now you make up a question about a prediction. Try to ask something about what might happen in a certain situation. Go ahead.
expert5: ummm... what happens when there are no macroinvertabrates in a polluted river
[Erin]: I'll just record your question for the discussions later.
[Erin]: But I'll make an explanation for this concept too. I know it might not answer your question, but I hope it helps us understand this concept!
[Erin]: Ok?
expert5: thanks
[Erin]: I predict that the number of dead things in a river would increase if there were no macroinvertebrates because the macroinvertebrates help clean up dead things by eating them.
expert5: ok
[Erin]: Do you like my explanation?
expert5: yah, that's what I thought would happen! Thnx
[Erin]: Let's move on to the first sampling site. Follow me!

GROUP SSB (Ingrid=Bot; Students=expert10 and expert11)

[Ingrid]: Hmmm... this concept is about how scientists use macroinvertebrates.
[Ingrid]: A good technique for explaining is to say what is new information and how that changes your thinking.

[Ingrid]: So, for this concept, ask each other about what you think is new and how that changes your thinking.
expert11: click on sign 2
[Ingrid]: I'll wait over by the second sample area.
expert11: How does it show that the river is polluted
expert10: if the intolerent types of mac. in the river die then the water is polluted.
expert10: why do scientists collect and count mac.
expert11: Some types of macs are sensitive to pollution.
expert11: lets go to sign 2
[Ingrid]: How macroinvertebrates breathe. Interesting!
[Ingrid]: Well, here is some more advice about explaining.
[Ingrid]: A good explanation can link together new information with something we know already.
[Ingrid]: For this concept, go ahead and ask your partner a question about linking something about breathing to something they already know.
[Ingrid]: Well, I'll meet you at sign 4.
expert10: What is called dissolved oxegen and why?
expert11: What kind of oxygen do macs breathe.
expert11: ?
expert10: they breathe the same kind of oxegen as any other animal.
expert11: ask me a question
expert10: oh and thier oxegen contains water
expert11: It is oxygen that you breathe underwater, and it is called dissolved oxygen because it is dissolved in the water.

GROUP SST (No Bot; Students=expert16 and expert17)

expert17: how do experts know if a river is polluted?
expert17: nm
expert17: what?
expert16: if it has lots of intolerent ones it is not polluted.
expert17: yeah
expert16: let me think
expert17: ok
expert16: you stole my question
expert17: sorry
expert16: its oaky
expert17: n
expert17: oops
expert16: i was thinking what
expert16: what does water quality reflect?
expert17: which types of organisms can survivein a body of water
expert16: yep
expert17: ok
expert17: lets go sample
expert16: just a min

expert17: ok
expert16: come here
expert17: ok
expert16: no to my desk
expert16: k
expert16: we need to go to area 3
expert17: i goota sample
expert16: not yet
expert17: come here
expert17: I will give u a question
expert16: whatdo M.'s breath?
expert17: oxygen
expert16: yep
expert17: what is the oxygen they breath called
expert16: disolved oxygen
expert17: ya
expert17: lets go

APPENDIX G

PRE- AND POST-TEST QUESTIONS

First Name: ________________ Last Name: __________________ Expert #: ____ Class Letter: ___

Circle the best answer.

1. Macroinvertebrates can help us decide a river is polluted because

A. It is easy to see when they are sick
B. Certain types of macroinvertebrates will die in polluted water
C. A polluted river will have no macroinvertebrates in it
D. There will be lots of them because they can eat the pollution

2. The small size of macroinvertebrates makes them useful to measure pollution because

A. Small animals get sick more easily than larger animals when the water is polluted
B. You can keep them in a fish tank to watch them
C. You can use a microscope to tell if they are sick
D. It is easy to get many different types when you collect a river sample

3. If you find many pollution tolerant macroinvertebrates in your river samples then

A. You know the river is polluted
B. You know the river is clean
C. The river might be polluted but you need more information
D. The river was polluted before but is getting cleaner now

4. The habitat of the macroinvertebrates is important because

A. They stay in one place for a long time so they can tell us more about the pollution
B. They can hide in the river bottom to get away from the pollution
C. They will move to a new home if the river becomes polluted
D. They will move to the polluted areas of the river

5. Pollution intolerant macroinvertebrates will die in polluted water because

A. The water becomes too dark for the macroinvertebrates to find their food
B. The water will become too thick and the macroinvertebrates won't be able to swim
C. The sun cannot reach the macroinvertebrates because the water is too dirty
D. The level of oxygen in polluted water is too low for the macroinvertebrates to live

6. What is the most complete definition of pollution?

A. Pollution is waste from humans
B. Pollution is anything that upsets the balance of an ecosystem
C. Pollution is garbage that either looks or smells bad
D. Pollution is something that kills everything in an ecosystem

7. Baseline data is

A. Data that you have collected at several times over a long period
B. The first data that you collect in the year
C. The last data that you collect in the year
D. Data that you have collected from sampling the bottom of a river

8. When you sample a river for water quality you should

A. Collect one large sample that contains many different macroinvertebrates
B. Collect at least three samples from the area that you are interested in
C. Collect the sample in the morning because the macroinvertebrates will be sleeping
D. Collect one sample in the daytime and one at night

9. Different types of macroinvertebrates help us figure out if a river is polluted. This is because

A. They eat different things
B. They are eaten by many types of fish
C. They live in different areas along the river
D. They require different amounts of oxygen in the water

10. The name macroinvertebrate is used to describe animals that

A. Are very small and have a backbone
B. Are too small to see without a microscope and don't have a backbone
C. Are big enough to see with your eyes and don't have a backbone
D. Are any plants or animals that live in the bottom of rivers

11. You go to help a scientist collect samples from a river that she thinks is polluted. If the scientist is right, what would you expect to find in your sample?

A. There would be more macroinvertebrates than usual
B. Most of the macroinvertebrates would be pollution intolerant
C. There would be fewer pollution intolerant macroinvertebrates than usual
D. There would be no macroinvertebrates in the sample

12. The manager of a steam power plant tells everyone in town not to worry about the waste water that flows into the river because they make sure it is clean before it leaves the plant. We've tested it and it's just pure warm water, he says. Should the town be worried?

A. Yes, because the water probably has some chemicals in it
B. No, because the water is pure so it will not poison any animals
C. Yes, because the waste water will not have any macroinvertebrates in it
D. Yes, because even pure warm water can upset the balance of an ecosystem

13. A friend of yours is interested in river monitoring but also likes to fish. He asks you why you can't use fish instead of macroinvertebrates to monitor water quality. What would be the best answer?

A. Fish are stronger and don't get sick as easily as smaller animals like macroinvertebrates
B. You could use them except that they are "No Fishing" areas on some rivers
C. Fish are less numerous and move around a lot so it is hard to collect a good sample
D. It’s too hard to tell when fish are sick

14. A friend calls you up because he is worried that his river is becoming polluted. He tells you that the number of pollution intolerant macroinvertebrates is now less than when he sampled a month ago. Should he be worried?

A. Yes, because fewer pollution intolerant macroinvertebrates means the water is polluted
B. Yes, because any change in the number of macroinvertebrates is bad news
C. Maybe, but he should compare his results to a sample taken at the same time last year
D. Maybe, but he should check to see if any fish are sick

15. If you went to another country to help monitor river pollution, what kinds of animals would you try to find for your work?

A. Animals that would be strong and swim fast
B. Animals that would all be able to survive in water with very little dissolved oxygen
C. A group of animals where some types were more sensitive to oxygen levels than others
D. A group of animals that would eat the pollution

REFERENCES

Allen, V. L. (1976). Children as teachers: Theory and research on tutoring. New York: Academic Press.

Anderson, J. R. (1987). Skill acquisition: Compilation of weak-method problem solutions. Psychological Review, 94(2), 192-210.

Anderson, J. R., Corbett, A. T., Koedinger, K., & Pelletier, R. (1995). Cognitive tutors: Lessons learned. The Journal of Learning Sciences, 4 (2), 167-207.

Aronson, E. & Patnoe, S. (1997). The jigsaw classroom: Building cooperation in the classroom (2nd ed.). Addison Wesley Longman, New York.

Arreaga-Mayer, C. A., Terry, B. J., & Greenwood, C. R. (1998). Classwide peer tutoring. In K. Topping & S. Ehly (Eds.), Peer-Assisted Learning, (pp. 105-119) Mahwah, NY: Erlbaum.

Atkinson, R. K., Derry, S. J., Renkl, A., & Wortham, D. W. (2000). Learning from examples: Instructional principles from the worked examples research. Review of Educational Research, 70.

Azmitia, M. (1988). Peer interaction and problem solving: When are two heads better than one? Child Development, 59, 87-96.

Bargh, J. A. & Schul, Y. (1980). On the cognitive benefits of teaching. Journal of Educational Psychology, 72, 593-604.

Baylor, A. L. (2002). Agent-based learning environments for investigating teaching and learning. Journal of Educational Computing Research, 26(3), 249-270.

Baylor, A. L. (2000). Beyond butlers: Intelligent agents as mentors. Journal of Educational Computing Research, 22(4), 373-382.

Bell, P. (1997). Using argument representations to make thinking visible for individuals and groups. In R. Hall, N. Miyake, & N. Enyedy (Eds.), Proceedings of CSCL '97: The second international conference on computer support for collaborative learning, (p. 10-19). Toronto: University of Toronto Press.

Bell, P., Davis, E. A., & Linn, M. C. (1995). The Knowledge Integration Environment: Theory and Design. In Proceedings of the Computer Supported Collaborative Learning Conference (CSCL '95: Bloomington, IN) (pp. 14-21). Mahwah, NJ: Lawrence Erlbaum Associates.

Benware, C. A. & Deci, E. L. (1984). Quality of learning with an active versus passive motivational set. American Educational Research Journal, 21, 755-765.

Berry, D. C. (1983). Metacognitive experience and transfer of logical reasoning. Quarterly Journal of Experimental Psychology, 35A, 39-49.

Berry, D.C. & Broadbent, D.E. (1987). Explanation and verbalization in a computer- assisted search task. Quarterly Journal of Experimental Psychology, 39A, 585-609.

Bielaczyc, K., Pirolli, P. L., & Brown, A. L. (1995). Training in self-explanation and self-regulation strategies: Investigating the effects of knowledge acquisition activities on problem solving. Cognition and Instruction, 13(2), 221-252.

Biswas, G., Schwartz, D. L., Bransford, J. D., & Teachable Agents Group at Vanderbilt (2001). Technology support for complex problem solving: From SAD environments to AI, smart machines in education: The coming revolution in education technology. In K.D. Forbus and P.J. Feltovich (eds.), AAAI/MIT Press, Menlo Park, CA.

Bransford, J.D., Brown, A.L. & Cockings, R.R. (Eds.). (1999). How people learn: Brain, mind, experience, and school. Washington, DC: National Academy Press.

Bransford, J.D., Franks, J., Vye, N. & Sherwood, R. (1989). New approaches to instruction: Because wisdom can't be told. In Vosniadou, St. & Ortony, A. (Eds.), Similarity and Analogical Reasoning. Cambridge Unversity Press , Cambridge, Mass. 470-497.

Bogdan, R. & Biklen, S. (1992). Qualitative research for education: An introduction to theory and methods. Allyn and Bacon, Needham Heights, MA.

Campione, J.C., & Brown, A.L. (1990). Guided learning and transfer: Implications for approaches to assessment. In N. Frederiksen, R. Glaser, A. Lesgold, & M. Shafto (Eds.), Diagnostic Monitoring of Skill and Knowledge Acquisition (pp. 141-172). Hillsdale, NJ: Erlbaum.

Carey, S. & Smith, C. (1993). On understanding the nature of scientific knowledge. Educational Psychologist, 28(3), 235-252.

Cawsey, A. (1993). Explanation and interaction: The computer generation of explanatory dialogues. MIT Press, Cambridge, Massachusetts.

Chi, M.T.H. (2000). Self-explaining: The dual processes of generating inferences and repairing mental models. In R. Glaser (Ed.), Advances in Instructional Psychology Volume 5. Lawrence Earlbaum Associates, Mahway, NJ. 161-238.

Chi, M.T.H., Siler, S., Jeong, H., Yamauchi, T., & Hausmann, R. G. (2001). Learning from human tutoring. Cognitive Science, 25(4), 471-533.

Chi, M. T. H., de Leeuw, N., Chiu, M. H., & LaVancher, C. (1994). Eliciting self-explanations improves understanding. Cognitive Science, 18, 439-477.

Chi, M.T.H., Bassok, M., Lewis, M.W., Reimann, P., & Glaser, R. (1989). Self-explanations: How students study and use examples in learning to solve problems. Cognitive Science, 13, 145-182.

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

Coleman, E.B. (1998). Using explanatory knowledge during collaborative problem solving in science. The Journal of Learning Sciences, 7(3&4), 387-427.

Coleman, E.B., Brown, A.L., Rivkin, I.D. (1997). The effect of instructional explanations on learning from scientific texts. The Journal of the Learning Sciences, 6(4), 347-365.

Dagher, Z. & Cossman, G. (1992). Verbal explanations given by science teachers: Their nature and implications. Journal of Research in Science Teaching, 29, 361-374.

Dembo, M.H., & Eaton, M.J. (2000). Self-regulation of academic learning in middle-level schools. Elementary School Journal, 100, 473-490.

Dillenbourg, P., Baker, M., Blaye, A. & O'Malley, C. (1995). The evolution of research on collaborative learning. In P. Reimann & H. Spada (Eds.), Learning in Humans and Machines: Towards an Interdisciplinary Learning Science. El-sevier , Oxford. 189 - 211.

Fantuzzo, J.W., King, J.A., & Heller, L.R. (1992). Effects of reciprocal peer tutoring on mathematics and school adjustment: A component analysis. Journal of Educational Psychology, 84, 331-339.

Fantuzzo, J.W., Riggio, R.E., Connely, S., & Dimeff, L.A. (1989). Effects of reciprocal peer tutoring on academic achievement and psychological adjustment: A component analysis. Journal of educational psychology, 81, 173-177.

Ferguson-Hessler, M. and de Jong, T. (1990). Studying physics texts: Differences in study processes between good and poor solvers. Cognition and Instruction, 7, 41-54.

Forman, E.A. & Larreamendy-Joerns, J. (1998). Making explicit the implicit: Classroom explanations and conversational implicatures. Mind, Culture, and Activity, 5 (2), 105-113.

Fuchs, L.S., Fuchs, D., Bentz, J., Bilshop, N., & Hamlett, C.L. (1994). The nature of student interactions during peer tutoring with and without prior training and experience. American Educational Research Journal, 31(1), 75-103.

Gabbert, B., Johnson, D.W. & Johnson, R. (1986). Cooperative learning, group-to-individual transfer, process gain and the acqusition of cognitive reasoning strategies. Journal of Psychology, 120 (3) 265-278.

Gagne, R.M. & Smith, E.C. (1962). A study of the effects of verbalization on problem solving. Journal of Experimental Psychology, 63, 12-18.

Graesser, A.C., Person, N.K., & Magliano, J.P. (1995). Collaborative dialogue patterns in naturalistic one-on-one tutoring. Applied Cognitive Psychology, vol. 9, pp495-522.

Hall, R. J. (1988). Learning by failing to explain: Using partial explanations to learn in incomplete or intractable domains. Machine Learning, 3, 45-77.

Hewitt, J. & Scardamalia, M. (1998). Design principles for distributed knowledge building processes. Educational Psychology Review, 10(1), 75-96.

Hill, G. (1982). Group versus individual performance: Are N + 1 heads better than one? Psychological Bulletin, 91, 517-539.

Katz, S. & Lesgold, A. (1993). The role of the tutor in computer-based collaborative learning situations. In Lajoie, S., and Derry, S., eds., Computers as Cognitive Tools. Lawrence Erlbaum Associates.

King, A. (1998). Transactive Peer Tutoring: Distributing Cognition and Metacognition. Educational Psychology Review, 10, 57-74.

King, A. (1994). Guiding knowledge construction in the classroom: Effects of teaching children how to question and how to explain. American Educational Research Journal, 30, 338-368.

Kintsch, W. (1994). Text comprehension, memory, and learning. American Psychologist, 49, 294-303.

Lajoie, S. P., Lavigne, N. C., Guerrera, C. & Munsie,S. (2001). Constructing knowledge in the context of BioWorld. Instructional Science, 29 (2), pp.155-186.

Lave, J. (1988). Cognition in Practice: Mind, mathematics, and culture in everyday life. Cambridge, UK: Cambridge University Press.

Leelawong, K., Davis, J., Vye, N., Biswas, G., Schwartz, D., Belynne, T., Katzlberger, T., & Bransford, J. (2002). The effects of feedback in supporting learning by teaching in a teachable agent environment. In P. Bell, R. Stevens, & T. Satwicz (Eds.), Keeping Learning Complex: The Proceedings of the Fifth International Conference of the Learning Sciences (ICLS) (pp. 245-252). Mahwah, NJ: Erlbaum.

Lin, X., Holmes, J. T. G., & Schwartz, D. L. (2001). A case study of a virtual learning environment for supporting the joint teaching of science lessons across cultures. Paper presented at American Educational Research Association, Seattle, WA.

Lin, X., Hmelo, C., Kinzer, C. K., & Secules, T. (1999). Designing technology to support reflection. Educational Technology Research and Development, 47 (3), 43-62.

Lin, X., & Lehman, J. (1999). Supporting learning of variable control in a computer-based biology environment: Effects of prompting college students to reflect on their own thinking. Journal of Research In Science Teaching, 36 (7), 837-858.

Mannes, S.M. (1994). Strategic processing of text. Journal of Educational Psychology, 86, 377-388.

Mwangi, W. & Sweller, J. (1998). Learning to solve compare word problems: The effect of example format and generating self-explanations. Cognition & Instruction 16(2): 173-199.

Neuman, Y., & Schwarz, B. (1998). Is self-explanation while solving problems helpful? The case of analogical problem solving. British Journal of Educational Psychology, 68, 15-24.

Okada, T. & Simon, H.A. (1997). Collaborative discovery in a scientific domain. Cognitive Science, 21 (2), 109-146.

Okada, T., Schunn, C. D., Crowley, K., Oshima, J., Miwa, K., Aoki, T., & Ishida, Y. (1995). Collaborative scientific research: Analyses of historical and interview data. Paper presented at the 1995 Meeting of the Japanese Cognitive Science Society.

Pea, R. (1993). Practices of distributed intelligence and designs for education. In G. Salomon (Ed.), Distributed cognitions: Psychological and educational considerations. New York: Cambridge University Press.

Perkins, D.N. (1993). Person-plus: A distributed view of thinking and learning. In G. Salomon (Ed.), Distributed cognitions: Psychological and educational considerations. New York: Cambridge University Press.

Pine, K. J. & Messer, D. J. (2000). The effect of explaining another's actions on children's implicit theories of balance. Cognition and Instruction , 18(1), 35

Pirolli, P. & Recker, M. (1994). Learning strategies and transfer in the domain of programming. Cognition and Instruction, 12, 235-275.

Plötzner, R., Dillenbourg, P., Preier, M., & Traum, D. (1999). Learning by explaining to oneself and to others. In P. Dillenbourg (Ed.), Collaborative Learning: Cognitive and Computational Approaches (pp. 103-121). Oxford: Elsevier Science Publishers.

Renkl, A (2000). Worked-out examples: Instructional explanations support learning by self-explanations. Research Report No. 139. University of Freiburg. Germany. Available: http://www.psychologie.uni-freiburg.de/pi-zentral/fobe-files/139.pdf (October 14, 2001).

Renkl, A. (1997a). Learning from worked-out examples: A study on individual differences. Cognitive Science, 21, 1-29.

Renkl, A. (1997b). Learning by explaining - or better by listening? Paper presented at the Annual Meeting of the American Educational Research Association, Chicago, IL. March, 24-28, 1997.

Renkl, A. (1995). Learning for later teaching: An exploration of mediational links between teaching expectancy and learning results. Learning and Instruction, 5, 21-36.

Renkl, A., Stark, R., Gruber, H., & Mandl, H. (1998). Learning from worked-out examples: The effects of example variability and elicited self-explanations. Contemporary Educational Psychology, 23, 90-108.

Rubin, H. J. & Rubin, I. S. (1995). Qualitative interviewing: The art of hearing data. Thousand Oaks, CA: Sage.

Scardamalia, M. & Bereiter, C. (1996). Computer support for knowledge-building communities. In T. Koschmann (Ed.), CSCL: Theory and practice of an emerging paradigm. Mahwah, NJ: Lawrence Erlbaum Associates.

Scardamalia, M. & Bereiter, C. (1985). Development of dialectical processes in composition. In D.R. Olson, N. Torrance, & A. Hildyard (Eds.) Literacy, language, and learning: The nature and consequences of reading and writing. Cambridge, Cambridge University Press.

Scardamalia, M., Bereiter, C. & Steinbach, R. (1984). Teachability of the reflective processes in written composition. Cognitive Science, (8), 173-190.

Schwartz, D.L. (1995). The emergence of abstract representations in dyad problem solving. The Journal of the Learning Sciences, 4(3), 321-354.

Schwartz, D. L., Lin, X., Brophy, S., & Bransford, J. D. (1999). Toward the development of flexibly adaptive instructional designs. In C. M. Reigeluth (Ed.), Instructional-design theories and models: A new paradigm of instructional theory (Vol. 2). Mahwah, NJ: Lawrence Erlbaum Associates.

Schwarz, B. B., Neuman, Y., & Biezuner, S. (2000). Two wrongs may make it right…If they argue together! Cognition and Instruction, 18(4), 461-494.

Shneiderman, B. (1997). Direct manipulation versus agents: paths to predictable, controllable and comprehensible interfaces. In J. M. Bradshaw (Ed.), Software agents (pp. 97-106). Menlo Park, CA: AAAI Press.

Siegler, R. S. (1995). How does change occur: A microgenetic study of number conservation. Cognitive Psychology, 28, 225-273.

Skon, L., Johnson, D.W., & Johnson, R.T. (1981). Cooperative peer interaction versus individual competition and individualistic efforts: Effects on the acquisition of cognitive reasoning strategies. Journal of Educational Psychology, 73(1), 83-92.

Slavin, R. (1995). Co-operative learning: Theory, research and practice (2nd Ed.). Boston: Allyn & Bacon.

Slavin, R. (1983). Cooperative learning.. New York: Longman.

Solomon, J. (1986). Children's explanations. Oxford Review of Education, 12(1), 41-51.

Steinbach. R., Scardamalia. M., Burti, P. J., & Bereiter, C. (1987, April). Children's implicit theories of knowledge and learning. Paper presented at the annual conference of the American Educational Research Association, Washington, DC.

Sweller, J. (1988). Cognitive load during problem solving: Effects on learning. Cognitive Science, 12, 257-285.

Teasley, S. D. (1995). The role of talk in children's peer collaborations. Developmental Psychology, 51, 207-220.

Ur, S. & VanLehn, K. (1995). STEPS: A simulated, tutorable physics student. Journal of Artificial Intelligence and Education, 6(4), 405-437.

VanLehn, K., Ohlsson, S., & Nason, R. (1994). Applications of simulated students: An exploration. Journal of Artificial Intelligence and Education, 5(2), 135-175.

VanLehn, K., Jones, R. M., & Chi, M. T. H. (1992). A model of the self-explanation effect. Journal of the Learning Sciences, 2(1), 1-60.

Vygotsky, L. S. (1978). Mind in society: The development of higher psychological processes. M. Cole, V. John-Steiner, S. Scribner & E. Souberman (Eds.). Cambridge, MA: Harvard University Press.

Webb, N.M. (1991). Task related verbal interaction and mathematics learning in small groups. Journal for Research in Mathematics Education, 22 (5), 366-389.

Webb, N. M. (1989). Peer interaction and learning in small groups. International Journal of Educational Research, 13, 21-39.

Wegerif, R., Mercer, N., & Dawes, L. (1999). From social interaction to individual reasoning: An empricial investigation of a possible socio-cultural model of cognitive development. Learning and Instruction, 9, 493-516.

Wittrock, M.C. (1990). Generative processes of comprehension. Educational Psychologist, 24(2), 345-376.

Zajonc, R.B. (1965). Social facilitation. Science, 149, 271-274.