A Framework for Comparing Web-Based Learning Environments

by

Christian Pantel

B.C.Sc. (Hons.), University of Manitoba, 1995

A THESIS SUBMITTED IN PARTIAL FULFILLMENT
OF THE REQUIREMENTS FOR THE DEGREE OF
MASTER OF SCIENCE
in the School
of
Computing Science

SIMON FRASER UNIVERSITY
October 1997

0-612-24220-X

APPROVAL

Name: Christian Pantel

Degree: Master of Science

Title of thesis: A Framework for Comparing Web-Based Learning Environments

Examining Committee: M. Stella Atkins, Chairman
Professor of Computing Science, SFU

Thomas W. Calvert, Senior Supervisor
Professor of Computing Science, SFU

F. David Fracchia, Supervisor
Assistant Professor of Computing Science, SFU

Philip H. Winne, Supervisor
Professor of Education, SFU

Lucio Teles, External Examiner
Co-Director of the LohnLab, SFU

Date Approved:

Abstract

As notions such as the information revolution and knowledge-based economies become increasingly realistic, there is considerable pressure on academic, corporate and governmental decision makers to improve both the accessibility and the quality of learning opportunities that they provide to those they serve. Many are turning to World Wide Web-based technologies as part of the solution.

A Web-based learning environment is a networked computer application that enables people to learn from a distance. Learners can be physically separated from teachers and from each other, and they can participate in the learning environment at their convenience. Choosing which Web-based learning environment to adopt is an important decision as it often involves a substantial investment for organizations and significantly impacts how people will learn.

In order to better understand how people learn and the possibilities for Web-based learning environments to support learning, we review the relevant educational theory. We also review the human-computer interaction literature to better understand how people can effectively use software.

Intuitive and ad-hoc comparisons between competing products are not likely to lead to adopting the optimal Web-based learning environment for a particular organization. Therefore, we propose a comparison framework for Web-based learning environments which is based primarily on educational theory and human-computer interaction research. The comparison framework consists of a large number of comparison dimensions covering a broad range of issues relevant to the adoption of a Web-based learning environment. The comparison framework serves two main audiences. The primary audience consists of decision makers of organizations considering

the adoption of a Web-based learning environment. We propose a methodology that these decision makers can follow when using the comparison framework to guide them in selecting the most appropriate Web-based learning environment for their organization. The secondary audience consists of developers of Web-based learning environments. By examining the comparison dimensions and by reading the literature reviews, developers may identify strengths and weaknesses of their product which can be useful in marketing and in planning future product releases.

Acknowledgments

First, I thank the Lord Jesus Christ. He made this work possible and deserves all the glory.

I thank the members of my supervisory committee for their interest in my research and their feedback on my thesis. I would like to particularly thank my senior supervisor, Tom Calvert. I am very pleased that he accepted me as one of his students and that he encouraged me to get involved in TeleLearning and the Virtual-U project. I am also grateful for the opportunities he provided me to meet other researchers in the field of human computer interaction and to make industry contacts. I enjoyed working with him. Dave Fracchia has always been encouraging and friendly, and he thoroughly reviewed my thesis. Phil Winne also thoroughly reviewed my thesis. His input on educational matters was particularly valuable and appreciated.

Thanks to the Virtual-U team for getting me started in TeleLearning. I particularly benefited from working with Lyn Bartram on user interfaces and from discussions with Denise Stockley: thank you both.

I am grateful to the Natural Sciences and Engineering Research Council of Canada and to the British Columbia Advanced Systems Institute for their scholarships. They helped make graduate studies feasible.

IBM's Pacific Development Center has been very supportive of my work. I greatly appreciated the Graduate Student Internship placement I obtained with them and am pleased to be joining them as a regular employee. Dave Hunter and Sal Visca deserve special thanks for giving me the opportunity to gain valuable industry experience while providing me with the flexibility I needed to balance work, university, and family.

My parents, Roland and Édith, have been excellent role models. They have sacrificed a lot for me to be where I am, and who I am, today. Their support has been very valuable and I appreciate it. I love them both. Finally, I am very thankful to my wife Sarah, whom I love and to whom this thesis is dedicated. Without her input, her patience, her encouragement, and her love, this thesis would not have happened and my life would not be complete.

Dedication

To my wife Sarah

Contents

List of Tables

List of Figures

Chapter 1

Introduction

As notions such as the information revolution and knowledge-based economies become increasingly realistic, there is considerable pressure on academic, corporate and governmental decision makers to improve both the accessibility and the quality of learning opportunities that they provide to those they serve. Many are turning to World Wide Web-based technologies as part of the solution.

1.1 Fundamentals

Wilson defines a *learning environment* as "a place where learners may work together and support each other as they use a variety of tools and information resources in their pursuit of learning goals and problem-solving activities" [117]. A *Web-based learning environment* is a hypermedia-based program that uses the attributes and resources of the World Wide Web to facilitate learning [63, 86]. It can be considered the *place* where learners and teachers interact [110]. We will restrict this definition to learning environments that support distributed learning. *Distributed learning* is place- or time-independent; that is, the teacher and the individual learners may be in different physical locations, or may engage in learning at different times [51, 115]. Distributed learning lessens the possible disruption to students' lives [81] by giving them more control over their learning environment. Some of the problems that learners report with traditional or correspondence type distributed learning, such as lack of motivation and a feeling of isolation, may be remedied by well designed technologically enhanced learning environments. Although similar technology can be used to augment face-to-face learning, they will not be the focus of this work.

1.2 The World Wide Web Medium

Technology supported learning environments are not exclusive to the World Wide Web. Web-based learning environments, however, have been singled out for study in this thesis. The simple fact that a new technology and medium like the Web now exists does not necessarily mean that it can be effectively used for the purpose of education. Many authors raise the possibility of bad or even dangerous uses of technology [5, 19, 51, 70, 81, 107]. Some of the issues raised include:

- additional stress caused by having to learn and manage a new technology;
- eye-candy — the gratuitous use of multimedia objects;
- the use of technology for its own sake;
- technology may not be customized for education;
- learners may spend more time and cognitive effort on learning the technology rather than their chosen field of study; and
- people tend to attribute undue authority to materials presented over an electronic medium by the simple virtue that they emanate from a computer which is assumed always to be correct.

From these general concerns over technologically supported learning environments, one gets a sense of the importance of properly designing a learning environment. Misuse of technology must be avoided. Even if the Web is not actually misused, it may not be easy to use it to its full potential. For example, Bostock [19] observes that if all we do is broadcast via the Web then "Human contact is lost, students are isolated, and the educational experience is passive, limited, and alienating."

What follows is a brief history of the World Wide Web and a discussion of some of the advantages and disadvantages of the World Wide Web as compared to other media.

1.2.1 History of the Web

The World Wide Web was invented by Tim Berners-Lee at CERN, the European Particle Physics Laboratory in Geneva, Switzerland. He first proposed the global hypertext project in 1989, based on earlier work on an unpublished program named *Enquire*. In October 1990 he began writing the first World Wide Web server and client. The client was a *WYSIWYG* ("what you see is what you get") hypertext browser and editor. These first Web products were made available on the Internet in the summer of 1991. Since 1994, Berners-Lee has been the director of the W3 Consortium which has teams at the Massachusetts Institute of Technology in the USA and at INRIA in France. The W3 Consortium is responsible for coordinating the development of the Web worldwide.

The Web's backbone is the Internet, the world's largest public network. In 1996, the network was made up of over 30 000 interconnected networks spanning over 70 countries [79]. The Web provides the Hypertext Transfer Protocol (HTTP) for accessing and retrieving multimedia resources identified by Unified Resource Locators (URL). The Web subsumes most of the existing Internet protocols including the File Transfer Protocol (FTP), Gopher, the Network News Transfer Protocol (NNTP) and the Simple Mail Transfer Protocol (SMTP) [79].

In 1994, the Web was a publishing medium. Clients could access static information stored on networked servers. The Web was therefore not very interactive. In 1995, submittable forms and the Common Gateway Interface (CGI) were added to the Web. These additions enabled a more interactive environment where the server could respond to user input. The year 1995 also saw the introduction of the Secure Socket Layer (SSL) and the Secure Hypertext Transfer Protocol (S-HTTP) which provided for secure transactions over the Web. In 1996, Sun Microsystems' Java was popularized. Sun describes Java as "a simple, object-oriented, network-savvy, interpreted, robust, secure, architecture-neutral, portable, high-performance, multithreaded, and dynamic language" [108]. Java allows developers to program components called applets that can be distributed over the Web. Since most current Web browsers can interpret Java code, much richer interactions between users and Web pages are now possible and the inherent statelessness of HTTP can be overcome.

The Web is usually thought of as *pull* technology. That is, end users request (pull) the transmission of documents from servers. An alternative technology is *push* technology. With push technology, servers, rather than end-users, initiate (push) the transmission of documents. PointCast [83] is an example of push technology which allows end-users to subscribe to news channels that interest them and then to receive relevant articles as they become available. To our knowledge, push technology has not yet been used in Web-based learning environments but we expect that this will change. Broadcasting of announcements is a simple example of how push technology could be useful in learning environments.

1.2.2 Why the Web?

The World Wide Web as a learning medium has a number of advantages and disadvantages as compared to other learning media. Although recent interest in Web-based learning environments seems to suggest that advantages sufficiently outweigh disadvantages, it is a useful exercise to familiarize oneself with the differences between various media.

Advantages

One significant advantage of the Web is that it causes us to expand the notion of a document by enabling multimedia documents. Whereas print is better at text, photography better at images, radio better at sound, and television better at video, the Web's advantage is that it combines all of these media reasonably well into a single package. There have been long debates in educational literature over the effect of different media on learning. Although there does not seem to be any significant advantage in using computer-based media over book-based media, Mayer [70] found that it is advantageous for learners to use multimedia learning materials. Since people seem to learn some content better through a combination of modalities and since the Web makes it easy to provide contiguous multimedia objects, the Web affords designing good multimedia learning.

The Web also causes us to expand the notion of documents by enabling hypermedia documents. Hypermedia documents can be linked together to form a non-linear information space with multiple paths between individual documents. The advantage of hypermedia for learning over other media such as print is not yet well established [33]. Some authors favour hypermedia for education. McManus [72], for example, states that because of the Web's non-linear nature and because it can also be considered as a content provider, it "could easily be considered the ultimate constructivist learning environment" (see Chapter 2 for a discussion of constructivism). Others are more cautious about the Web. Scardamalia and Bereiter [92], for example, state that hypermedia has "the potential to make matters worse, creating a sprawling network of items lacking even topical organization."

One clear advantage of the Web over media like books and CD-ROMs is the immediacy of changes. It is far easier to keep learning materials current on the Web then it is in traditional publications. The time lag and expense involved in printing a new version of learning materials is often prohibitive whereas changing a few Web pages is quite easy and relatively inexpensive.

Web-based learning environments remove the need for co-presence. Learners and instructors can be separated both in time and place. As compared to traditional distance education, Web-based learning environments can enable some much needed interaction between students and instructors as well as class collaboration [81].

As compared to other custom commercial software that could be used to support learning, Web clients and servers are relatively inexpensive. In fact, many are free of charge. The clients and servers are also relatively well standardized and therefore operate on many computer platforms.

The Web server's ability to track learning events is an important advantage for Web-based learning environments. Jonassen et al. [58] affirm: "Technologies provide us with the option of capturing and compiling learning events and making them available to other learners. Indeed, this may turn out to be the most effective use of computers in education."

Disadvantages

The World Wide Web is not a perfect medium for education. There are disadvantages to using it. Some of these disadvantages are of a technical nature whereas others are more people-centered.

The Web was originally designed for browsing and retrieving information as opposed to supporting an active learning system [41]. The implication is that developers of interactive Web-based learning environments must force the Web to perform a task for which it was not originally designed. As compared to other technologies, the Web entails a certain amount of loss of technical control for developers [61]. Since Web browsers are under the control of the end-user, developers must struggle to overcome differences in end-users' screen resolution, number of colours, network connection speed, window size, and available fonts. Furthermore, the somewhat different behaviour of competing Web browsers also must be accounted for by developers.

Quality of service can be a problem for the Web. Whereas it is possible to guarantee quality of service for networks such as ATM, it is impossible to do so for the Internet [20]. This means that the response speed for requests over the network can only be estimated, not guaranteed.

Synchronizing media on the Web is difficult. This makes it difficult to deliver rich prescribed presentations as one might give in a traditional lecture hall.

The fact that the Web is a nonlinear information space and that the size of the information space is often unknown to end-users can be problematic. Users may experience a sense of disorientation and become frustrated. This, combined with the fact that the interaction mechanisms of the Web are significantly different from those of print media, implies that learners must acquire a new set of skills to function adequately in this type of environment. Until they do, they may feel disoriented and overwhelmed in the information space [33].

Two significant disadvantages of the Web over non-technological vehicles for learning are technical and logistical problems. Hiltz [52] makes the point that although the specific nature of these problems changes, they are ever present. Resolving these problems consumes valuable resources that might otherwise be better utilized for teaching and learning.

Although important disadvantages do exist in using the Web for education as compared to other media, people seem willing to accept them, or at least to ignore them. This is evidenced by the growing popularity of the Web. According to the Spring 1997 Internet Demographic Study by Nielson Media Research and CommerceNet [26], 23% of people over the age of 16 in Canada and the USA use the Internet and 17% are on the Web. Matrix Maps Quarterly: State of the Internet, January 1997 [67] reports that a total of 194 countries, representing roughly 80% of all countries, have at least one internet domain. North-American governments are investing considerable funds to equip schools and libraries with Internet-connected computers.

Turoff [112] expresses his optimism about the potential benefit of technologies like the Web for education by stating: "Once we free ourselves from the mental limits of viewing this technology as a weak sister to face-to-face synchronous education, the potentials to revolutionize education and learning become readily apparent."

In limiting the comparison of learning environments to Web-based learning environments, we do not mean to imply that these are necessarily better than other learning environments. We simply level the playing field between competing solutions and focus our attention on a manageable problem.

1.3 Goals

The main goal of this thesis is to design a generalizable framework for comparing Web-based learning environments. The framework should guide and structure the comparison of competing products. Furthermore, the framework should be objective and broad enough to account for the needs of most users of a Web-based learning environment, including learners, teachers, system administrators and organization administrators. The comparison framework must account for the fact that different organizations will have different needs and priorities. It is therefore important to emphasize that based on the comparison framework, one can not declare that a particular Web-based learning environment is the best. Blanket statements about the merits of competing solutions are not useful nor encouraged, and they will not be found in this work. The most that can be said is that a given learning environment is likely the most suitable for a particular organization.

Since the framework's domain is educational, it is important to have a grounding in educational research. To better understand how people learn and the possibilities for Web-based learning environments to support learning, we review some of the relevant educational theory. The framework is informed by constructivism and sociocultural theories of learning. Although other theories of learning such as the information processing model are also valuable and could be used to inform comparison frameworks, constructivism and sociocultural theories of learning have been singled out in this study.

Since Web-based learning environments are software products and since human-computer interaction (HCI) research treats issues important to the success of software products, the comparison framework will benefit from a solid grounding in human-computer interaction research. As the field of HCI has been maturing it has been gaining respect. Many software giants such as IBM and Microsoft now employ HCI experts in research and product development. The HCI literature raises critical issues that need to be considered by software developers to maximize the usability, the usefulness and the customer acceptance of their products.

The comparison framework serves two main audiences. The primary audience consists of decision makers in organizations considering the adoption of a Web-based learning environment. The secondary audience consists of developers of Web-based learning environments.

Intuitive and ad-hoc comparisons between competing products are not likely to lead to adopting the optimal Web-based learning environment for a particular organization. The comparison framework proposed herein can serve those who have to make informed decisions about adopting a particular Web-based learning environment for their organizations. We propose a methodology decision makers can follow when using the comparison framework to guide them in selecting the best Web-based learning environment for their organization.

Developers of Web-based learning environments can also benefit from this work. With the excitement and the opportunities that the Web has created in recent years, many vendors have been rushing products to market without having much knowledge of the research that could guide and support them in creating better solutions. The review of educational and human-computer interaction research included herein can be a good starting point. Eales and Welsh emphasize the importance of educational theory by stating: "In order to make anything more learnable we believe that one needs at least a working understanding of what it means to learn" [34]. How can one solve a problem effectively without first making an effort at understanding it? By examining the framework for comparing Web-based learning environments, developers may identify relative strengths and weaknesses of their product which can be useful in marketing, and in planning future product releases.

1.4 Thesis Organization

Chapter 2 reviews the educational theory that will inform the design of the comparison framework. Chapter 3 reviews important human-computer interaction literature and highlights issues that need to be considered in evaluating interactive software products. In Chapter 4, we present the framework for comparing Web-based learning environments and suggest ways it can be used by organizations considering the adoption of a Web-based learning environment as well as by developers of these environments. The final chapter discusses the comparison framework and conclusions.

Chapter 2

Educational Theory

There is a rich body of educational research that deals with what it means to learn and how technology can be used to support effective learning. The challenge to the uninitiated is to first understand the language used by researchers in this field and then to begin to comprehend the many learning theories and models that are proposed in the literature. When trying to do this, one quickly comes to the realization that there is considerable controversy in the field and that researchers do not yet use consistent terminology nor agree on the best theories.

The goal of this chapter is to distill and synthesize research in the field of educational theory to inform the design of a framework for comparing Web-based learning environments. Although this is not by any means a complete review of all educational theory research, it is a thorough review of a selection of salient research. The reader will find that several of the main theories of learning are introduced and that a few of the them are examined in more detail, namely, constructivism and sociocultural theories of learning. It is hoped that this chapter will provide a solid grounding for the framework presented in Chapter 4 so that it will be useful in comparing not only the current state of the art in Web-based learning environments but also distributed learning solutions that will be developed in the future. This chapter should also be a good starting point for those who develop technological tools to support learning and who are not familiar with educational research. By familiarizing themselves with the relevant research, developers can better understand users of their software (learners and teachers) and discover opportunities to apply technology to support cutting edge learning models.

2.1 Introduction to Learning Theories

Learning theories are based on assumptions about what knowledge is and how people learn. The models and theories of learning that will be introduced in this section are the response strengthening model, the information processing model, constructivism, and sociocultural theories of learning. The latter two will be explored in more detail later in the chapter and will inform the design of the framework for comparing Web-based learning environments presented in Chapter 4.

2.1.1 Information Processing Model

Variations of the information processing model are employed, as they have been for decades, in today's typical North American classrooms. The model is therefore often intuitively familiar as many people have been educated under it. It is important, however, for those who develop learning environments to rely not only on their intuition but also to develop an understanding of the information processing model. This will enable them to form their own informed opinions of the model, to reflect on what a learning environment should support, and to articulate their reasoning for designing a learning environment that supports the information processing model. Understanding some of the premises of the information processing model can also serve as a benchmark for thoughtful consideration of alternative learning theories. What follows is a brief introduction to the information processing model. Readers are encouraged to refer to other sources in educational theory literature for a more complete account of the model.

The information processing model is based on the literal interpretation of the metaphor of learners as information processors [69]. People are believed to be processors of information. They take information as input, apply mental operators to it, and produce information as output [69].

Conception of knowledge

Knowledge is assumed to be an entity that exists independently from a person's mind [12]. As such, it is something that can be transferred from one person to another. Di Vesta [31] explains that "knowledge consists of *copying* elements from the world into a sensory store in a mechanistic way."

Conception of learning

Many definitions of learning are given in the literature but some typically associated with the information processing model include: a quantitative increase in knowledge, memorizing, and the acquisition of facts, skills and methods [5]. Learning is believed to be the process by which knowledge is transferred from the teacher's head into the head of the learner [69, 12, 34]. The accurate storage and retrieval of information which has inherent meaning is stressed, and emphasis is placed on the ability to produce a predetermined measurable final outcome [36].

Teacher's role

The above assumptions about knowledge and learning have given rise to didactic instruction and classical instructional design. Under the information processing model, "the focus of instruction is the curriculum ... the goal of instruction is to increase the amount of knowledge in the learner's repertoire [and] ... learning outcomes can be evaluated by measuring the amount of knowledge acquired" [68]. Teachers are considered dispensers of information [68]. The role of the teacher under the information processing model consists of setting learning goals, planning learning outcomes, preparing and sequencing learning materials, delivering instruction to the learners, assigning activities, evaluating the students' products and giving feedback. A popular instructional technique under the information processing model is lecturing.

2.1.2 Response Strengthening Model

The information processing model is predated by other models of learning. To put it in context, let us briefly touch on the response strengthening model that was the most generally accepted learning model in the first half of the twentieth century [68, 69].

Conception of knowledge

Under this older model, knowledge was considered to be the associations people made between stimuli and responses.

Conception of learning

Learning was the strengthening of desirable associations and the weakening of undesirable ones. An association was believed to be strengthened by repeated combined exposure to the stimulus and the response. Also, an association between a stimulus and response was believed to be strengthened when a learner experienced positive feedback for exhibiting the response and weakened when the learner experienced negative feedback.

Teacher's role

Under the response strengthening model, teachers dispensed rewards and punishments as a function of associations learners developed. Drill and practice was the instructional method of choice for providing learners with the repetition needed to strengthen associations.

2.1.3 Constructivism

Constructivism is an alternative to traditional learning theory that was founded in the 1950's by researchers including Piaget [82] and Vygotsky [114], and that has gained popularity in the 1980s and 1990s [69]. Many flavours of constructivism now exist. In recent years, new learning models have been developed that are rooted in constructivism but which have extended some of its fundamental ideas.

Constructivism, like the other learning theories, is predicated on assumptions about what knowledge is and how people learn. The epistemological differences that exist between theories are significant because they are the foundation on which teaching and learning models are built. What follows in an outline of the theoretical underpinnings of constructivism.

Conception of knowledge

Constructivist theorists do not believe knowledge is a constant for each object or event but rather that it is constructed by individuals as they interact with an object or an event, in relation to their past experiences, their beliefs and their current mental structures [15, 54]. Constructivists, like proponents of the information processing model but in contrast to supporters of the response strengthening model, argue that learners have a mental model whose structure and linkages represent their current knowledge. Jonassen et al. [58, 59] stress the importance of the mind's interpretation of the external world and that knowledge is a personal and individualistic thing, grounded in physical and social experiences. Jonassen et al. go on to argue that the mental models that are produced by the mind are used to explain, predict or infer phenomena in the world.

Conception of learning

For constructivists, learning is the process by which accessed information is transformed into personal knowledge [60]. It involves an evaluation of the new information based on existing mental models, and an augmentation and reorganization of these models to reflect the new knowledge [104]. This is a process of internal negotiation of meaning. Since all learners engage in a learning experience with differing beliefs and background knowledge, and since learning changes individuals' mental structures, it must be a personal experience.

Teacher's role

Mayer [68] explains that under constructivism, the focus of instruction is the learner's cognition, the goal of instruction is to help learners "develop learning and thinking strategies" and the evaluation of learning outcomes consists of "determining how the student structures and processes knowledge."

As explained in [49], a constructivist teacher's role is quite different from that of a teacher in a traditional learning environment. In a constructivist learning environment the teacher relinquishes the role of *sage on the stage* for one of *guide on the side* [85]. Teachers assume a role of coach or facilitator, guiding learners as opposed to being the keepers and dispensers of knowledge [19, 86, 103]. Much of the cognitive labour that is offloaded onto teachers in more traditional learning models is returned to the learners [50]. This includes tasks such as summarizing and synthesizing and may also include goal setting and evaluation. Constructivist teachers are often models of the expert learner.

2.1.4 Sociocultural Theories of Learning

Educational researchers have continued their study of learning and have developed sociocultural theories of learning. These theories resemble constructivism more than the information processing model and the response strengthening model. As Vargo [113] remarks, they are in many ways the children of constructivism.

Oakeshott's 1962 work is recognized as being the first to introduce the idea that knowledge does not exist apart from its application [111]. His idea has been elaborated in recent years by Lave and Wenger [65], and by Brown et al. [21]. Stockley and Roth [107] provide an interesting literature review on the topic. Two theories that are proposed in the literature and that will be introduced in this section are situated learning and socioculturalism. These theories are similar but the latter is more extreme in its beliefs.

Conception of knowledge

The sociocultural theories of learning stress that knowledge resides in communities and therefore that objects and events do not carry inherent meaning. Knowledge is built by a social negotiation of meaning [18], as opposed to a constructivist internal negotiation of meaning. Social negotiations of meaning take place between members of a learning community and involve the surrounding environment. The environment includes common practices, language, use of tools, values and beliefs of the community [104]. This focus on context and on a community of practice is primarily what distinguishes sociocultural theories of learning from constructivism, which is more focused on an individual's learning.

Brown et al. [21] compare knowledge to a tool. They note that it is quite possible to acquire knowledge, like a tool, and to talk about it, and yet be unable to use it. It is through the active use of tools that people may gain better understanding of the tools themselves and the environment in which they are used. They stress that meaning is not invariant but is defined by the community of practitioners that uses it. Concepts, culture and activity are therefore argued to be interdependent.

Conception of learning

Supporters of the sociocultural theories consider learning to be a process of enculturation [21, 104]. Brown et al. [21] call it a "life-long process resulting from acting in situations." As people observe and participate in a community of practice they want to enter, they learn to act in accordance to its norms. According to this model, people learn best when they are either immersed in the actual activities of a community of practice or an accurate simulation thereof. Meaning-making is the result of active participation in socially, culturally, historically, and politically situated contexts [106]. Learners can internalize a culture without explicit teaching; they do so by participating in it with others.

Teacher's role

The teacher's role in sociocultural learning environments is similar to that of constructivist teachers. They serve as coaches and guides to learners for whom they are responsible. Socioculturalism is more extreme in its beliefs than situated learning in that it focuses on the development of the collective knowledge of a community as opposed to the development of individuals' knowledge within a community.

Learning theories and models are numerous and much has been written about them. The ones that were introduced here represent the most influential models of the twentieth century. As Mayer [69] describes, the stimulus response model was generally accepted in the first half of the century, the information processing model was popular in the 1960s and 1970s, the constructivist model has been popular since the 1980s, and sociocultural models may be emerging as the next generally accepted learning theory. Since constructivism and sociocultural theories have recently generated considerable interest in both educational researchers and learning environment developers, and since the scope of the educational theory considered in this thesis had to be restricted to be manageable, we will delve only into the constructivist and sociocultural theories of learning.

2.2 Values and Practices of Constructivism

In this section, we will expand on the introduction to constructivism in Section 2.1.3 to expose constructivistic values and practices. The main points that we make are constructivists' commitment to the development of higher-order thinking and constructivists' emphasis on empowering learners to take ownership of their learning. We also present a disagreement within the constructivist community with respect to the application of constructivist learning environments.

2.2.1 Higher-Order Thinking

The design of learning environments for constructivists is not so much concerned with sequencing prescribed material. Rather, it is concerned with creating a learning environment that facilitates the development of higher-order thinking processes such as critical thinking, creative thinking, reflective thinking and metacognition. In such an environment, learners are given significant cognitive responsibilities. They engage in analysis, synthesis, problem solving, experimentation and creativity. They also learn to examine topics from multiple perspectives [60].

Analysis and synthesis

Analysis is a skill that allows a person to determine the relevance and relative importance of accessed information with respect to the task at hand. It is an ability which is becoming increasingly important in a world where it is a significant challenge to sift through the masses of easily accessible information to find the pearls. Synthesis is the process by which a person takes existing ideas and associates them in innovative ways to generate knowledge [87]. Jonassen et al. [58] call this generative processing and note that material is generally better remembered if it is generated by the learner rather than by someone else.

Problem solving

The development of problem solving skills is critical to transferring knowledge from the learning environment to the work environment. The ability to utilize existing knowledge to solve new problems in differing contexts is very much sought after in the workplace. Problem solving should therefore be experienced and learned. Considering problems from multiple perspectives can help in problem solving as well as in being more creative. Learners will therefore benefit from learning environments where they are exposed to concepts many times and in many different ways.

Reflective thinking and metacognition

Reflective thinking and metacognition are advocated by constructivists. Reflective thinking is managing and monitoring one's learning and performance [54]. It involves both planning and evaluation. Metacognition, on the other hand, involves becoming aware of cognitive strategies that one possesses and how they can be applied in new situations. Metacognition is well defined by Di Vesta as: "a learner's awareness about what skills or strategies are useful or detrimental for performing, understanding, comprehending, or reasoning about particular tasks, given the learner's knowledge about her advantages or limitations in capacity, knowledge, motivations, or other characteristics" [31].

The ability to evaluate information and one's own progress, as well as the ability to generate new knowledge and solve problems creatively based on multiple perspectives, form a powerful toolkit of thinking skills for the constructivist learner. Note that although these skills are emphasized by constructivists, they are also often advocated by supporters of other models, such as the information processing model, and sometimes exercised in more traditional learning environments.

2.2.2 Ownership of Learning

Constructivists advocate that learners be allowed and encouraged to take more ownership of their learning; that is, they should be given more responsibility over their learning. Schank [93] stresses the point that learning should be based on one's goals and interests. As Soloway et al. put it: "as students have more input into the choice and control of their environments, their motivation for pursuing cognitively challenging problems increases. Allowing students to decide how to plan, design, and work on their models can engage them in the learning process" [104].

To help learners take ownership of their learning, they can be encouraged to generate hypotheses and to pursue their resolution [58]. In doing so, learners develop the ability to set goals and to articulate their reasoning; important skills to develop for today's workplace. Giving learners ownership of their learning requires that learners not only participate in initial goal-setting but also in the evaluation of their progress.

As proposed in [15], learners should be given the opportunity to set the goals of evaluation for their learning.

Giving learners ownership of their learning empowers them [50]. This requires greater flexibility from the learning environment and can contribute to learning activities that are more authentic and meaningful to the learners.

2.2.3 Application

There is disagreement within the constructivist community with respect to the question of what types of learners benefit from constructivist learning environments over more traditional learning environments. Some researchers promote constructivism for advanced learning [58, 105] while still relying on more traditional methods for introductory knowledge acquisition. Others like Scardamalia and Bereiter [92] have observed that elementary school children can successfully build knowledge in a constructivist environment. Honebein [54] suggests that, from the learners' perspective, "the domain always involves advanced knowledge acquisition (after all, it is yet to be learned)... Thus we see little reason to suggest different principles for the design of instruction for less advanced learners." Some argue that training is inherently objectivistic because it supports explicit company performance goals [58] while others like Schank disagree and state: "The folks who think there's a difference between training and education have profoundly misunderstood education" [93].

Within the constructivist community, there seems to be agreement that constructivist learning environments are good for advanced knowledge acquisition. Note, however, that there is no such consensus for introductory knowledge acquisition.

2.3 Values and Practices of Sociocultural Theories

As the previous section expanded on the introduction to constructivism in Section 2.1.3, this section will expand on the introduction to sociocultural theories of learning presented in Section 2.1.4. We will explore some of the values and practices of the sociocultural theories' supporters. In doing so, we will introduce concepts and terminology such as authentic activities, communities of practice and legitimate peripheral participation which are important to proponents of this model.

2.3.1 Authentic Activities

Many researchers have stressed the need for providing authentic learning activities for learners. Simply put, authentic activities are the ordinary practices of a culture [21], those one would expect to encounter in real life practice. Jonassen et al. [60] define authentic activities as those which are realistic, meaningful, relevant, complex and information-rich. Note that the latter two characteristics of authentic activities in Jonassen et al.'s definition are not strictly required for an activity to be authentic. Whether an activity is authentic is a relative concept [54]. An activity can only be authentic with respect to an environment, or a community of practice. If a person is learning how to function in a test-taking environment, then drill and practice activities are authentic. If the environment consists of being able to recall and apply knowledge to new situations as they occur in the workplace, then drill and practice is not an authentic activity. Many researchers (e.g. [58, 59, 60, 105]) advocate the use of case-based learning and problem-solving activities to achieve better learning.

2.3.2 Complexity

The level of complexity provided in a learning environment contributes to the authenticity of activities. Adherents to the sociocultural theories of learning, like constructivists, argue that it is important to reflect the complexity of the application domain in the learning environment so they prefer to maintain a level of complexity similar to that encountered in the domain of practice.

The best level of complexity to maintain in a learning environment is debatable. It is clearly not appropriate for middle school children learning about science to encounter the same complexity that practicing chemists or physicists experience. Honebein attempts to settle the question by stating "the complexity of the stimulus environment at the beginning of learning should reflect the complexity of the environment expected in the terminal performance" [54].

2.3.3 Legitimate Peripheral Participation

Researchers have observed that people can learn a lot from observing others who are experiencing and overcoming difficulties [58]. In recognition of this fact, they encourage the use of legitimate peripheral participation in learning environments. Learner participation is legitimate if it involves authentic activities from the community of practice. It is peripheral in the sense that it is not the same as experts in the community [107]. It allows for novices to observe and learn from implicit or explicit demonstrations from experts and from other learners. Novices who start out on the periphery of their community of practice can progressively move towards its center until they themselves become experts.

2.4 Progression of the Theories

It is important to consider each learning theory separately to understand its assumptions about knowledge and learning, as was done in preceding sections. It is also enlightening, however, to consider the shortcomings of earlier theories observed by researchers and practitioners, and how correction of these was attempted when developing new learning models and theories. This more practical approach highlights some of the goals of constructivist and sociocultural theories of learning.

As one might expect, there are also researchers who disagree with aspects of the more recent theories. For example, there is an ongoing debate about the value of situated learning [6, 44, 7] in which Anderson et al. [6] argue that the major claims of situated learning are overstated. The goal here, however, is not to promote constructivism or sociocultural theories of learning over other learning models but rather to understand enough about constructivism and sociocultural theories so that they can inform the design of the comparison framework presented in Chapter 4. It therefore serves our purposes to examine how proponents of these more recent theories attempt to solve problems they identified with earlier learning models.

Although most researchers and practitioners acknowledge that the information processing model has its place, many researchers have argued that it has serious shortcomings. Brown et al. have observed that the behaviour of learners in traditional schools is quite different from that of people outside of schools: "the general strategies for intuitive reasoning, resolving issues, and negotiating meaning that people develop through everyday activity are superseded by the precise, well-defined problems, formal definitions, and symbol manipulation of much school activity" [21]. Learners sometimes feel they have to masquerade their approach at solving a problem and pretend they arrived at their solution by a prescribed method because, although schooling attempts to teach problem-solving, it often discounts students' inventive heuristics.

2.4.1 Conceptual Misunderstanding

Spiro et al. [105] have identified biases and assumptions in learning models which can cause learners to form only a superficial understanding of concepts. The main problem is oversimplification, also called the reductive bias. Problems are sometimes simplified and treated in isolation from their context to expose elegant general patterns. If real complexities and irregularities exist in the application domain and if their mastery is important, then oversimplifications cause learners to misunderstand key concepts.

Oversimplifications can take many forms. The additivity bias is the assumption that parts of complex situations which were studied in isolation retain their behaviour when reintegrated into the context they were taken from. The discreteness bias is the incorrect assumption that continuous measurements like length and time are discrete. The compartmentalization bias is the study of individual parts of a complex situation without regard for their interdependencies. Wilson cautions against these dangers by remarking that "as we simplify and package instruction for consumption, the richness of the subject can be bleached away" [117].

To combat oversimplification and the conceptual misunderstandings that follow, Spiro et al. [105] along with other constructivist researchers [59, 60, 15, 48, 54], suggest that learners should be exposed to content multiple times, in varying contexts and from many different perspectives. Supporters of constructivism also advocate maintaining the complexity and the irregularity of the real world. The emphasis placed on authentic practices by proponents of the sociocultural theories of learning also serves to address the problem of oversimplification.

2.4.2 Transfer of Knowledge

One critical characteristic of effective learning is the ability to transfer the knowledge and skills acquired in a learning environment to the environment of practice. The failure to apply knowledge has been identified by many researchers [105, 13] as one of the common difficulties for learners taught in traditional learning environments. This has motivated proponents of constructivism and sociocultural theories to attempt to improve the transferability of skills and knowledge.

Constructivists have addressed this point in many ways. In their Interpretation Construction (ICON) design model, Black and McClintock [15] achieve better transferability by exposing learners to multiple manifestations of similar interpretations of an object or an event. Observing interpretations in different contexts helps learners apply these interpretations to new situations. Constructivists also propose maintaining the complexity of the domain of practice in the learning environment as a means of increasing transferability.

Proponents of sociocultural theories of learning go further in their attempt to increase the learners' ability to transfer knowledge and skills. Their emphasis on authentic practices exposes learners to situations similar to those that they will encounter in their community of practice. Learners also develop the ability to use their knowledge and skills within a legitimate context as opposed to only in artificial situations like traditional classrooms. This helps to ensure that learners do not become dependent on cues and supports in a classroom setting that, if absent in the community of practice, could hinder their ability to transfer their knowledge and skills.

2.4.3 Motivation

Lack of motivation can be a problem for learners. Soloway et al. observe: "One can not count on the motivation of learners: both students and professionals have a strong tendency to procrastinate, to fritter away time" [103].

Constructivism and sociocultural learning theories place significant importance on what each learner brings to the learning experience because they believe that prior experiences and motivation are factors in how knowledge is built. Constructivists and proponents of the sociocultural theories of learning attempt to remedy the learners' lack of motivation by giving them more ownership and control over their learning. The learners are involved in goal setting, choosing their learning activities and evaluating their progress. Learners are encouraged to be curious and explorative. Based on their interests, learners are encouraged to formulate and evaluate hypotheses.

Sociocultural learning environments address motivational issues by providing learners with tasks that are authentic with respect to their field of study and that are to be performed in a complex, social environment. This tends to make learning more meaningful for learners and potentially more motivating. As pointed out by Anderson et al. in a critique of situated learning: "It seems important both to motivation and to learning to practice one's skills from time to time in full context" [6]. In the context of team sports, for example, players spend the majority of their time practicing individual and group skills, but they are motivated to do so by the prospect of playing in full-scale games [6]. The same is believed to be true for learners. Although critics of situated learning like Anderson et al. do not agree that full context learning should be the principal mechanism of instruction, they do acknowledge its motivational benefits.

Socioculturalists stress the knowledge of the collective and focus students on communal problem-solving while creating a supportive social climate for learning [50]. The social aspect of learning that is encouraged by socioculturalists may be a good motivational factor for learners.

Next, we will consider learning techniques that proponents of constructivist and sociocultural theories advocate. The techniques are based on the proponents' assumptions about knowledge and learning, their values and practices, and the efforts they made to overcome the shortcomings of earlier theories.

2.5 Techniques that Foster Learning

In this section we will examine some of the instructional and learning techniques that are proposed by educational researchers who support constructivism or sociocultural theories of learning. Techniques such as drill and practice for response strengthening and lecturing for knowledge transfer are out of the scope of this section. The techniques that will be considered are facilitating with scaffolding and fading [110], cognitive apprenticeship and collaborative learning.

2.5.1 Facilitating with Scaffolding and Fading

Teachers in constructivist learning environments are facilitators. Instead of being disseminators of knowledge, they guide learners, coaching and supporting them in their knowledge building efforts. They facilitate a learner's personal construction of knowledge. Two of the techniques that teachers can employ to support learners in their learning endeavors are scaffolding and fading.

Scaffolding

Scaffolding is an instructional technique that appropriately uses a metaphor from the construction industry. In construction, scaffolding is a temporary structure that supports builders in the construction of a more permanent structure. Similarly, constructivist scaffolding is a technique by which teachers temporarily support learners as they build their own knowledge structures. The intellectual support can take the form of comments, suggestions, feedback or observations [110].

Scaffolding enables learners to perform tasks they could not otherwise perform [103]. Activities that would normally be out of their reach become accessible [104, 106]. This technique supports learners as they challenge the upper bound of what Vygotsky [114] called their zone of proximal development. This is the zone in which learners need help to perform a task successfully [54].

As learners construct knowledge, their zone of proximal development changes. They develop the ability to successfully perform tasks without assistance where they previously could not. The scaffolding is therefore no longer needed.

Fading

Fading is the technique by which a teacher slowly removes the supports they provided learners in their early attempts at performing a task. As the learner progresses towards mastery, the teacher fades [110] and the learner assumes full control over the task at hand.

In summary, scaffolding and fading are two constructivist teaching techniques for facilitating knowledge building that go hand in hand. Since the zone of proximal development does not disappear until complete mastery is attained, scaffolding and fading can be seen as an iterative process. As partial mastery is attained, some scaffolds fade and new ones are provided to support continued development.

2.5.2 Cognitive Apprenticeship

Supporters of the sociocultural theories of learning strongly advocate Collins et al.'s [25] cognitive apprenticeship model for learning. Teles defines apprenticeship as: "Individual learning combined with the support of knowledgeable others through the processes of modeling, observation and successive approximations" [110]. Traditional craft and trade apprenticeships have had a long and successful history in teaching people to be carpenters, electricians, cabinet makers and masons, for example. Successful examples of apprenticeships are not only limited to the trades. An invaluable part of the training for medical doctors comes in the form of apprenticeships. The same is true for lawyers and accountants who must work in their profession for a period of time before receiving their professional designations. Aspiring university professors are prepared for their positions by learning to conduct research as graduate and post-doctoral students under the guidance of more senior researchers. This is again an example of successful apprenticeship.

The main characteristics common to these educational paths are that learners are immersed in the culture of their chosen profession and that they learn by performing tasks that are authentic and endorsed by the community of practitioners they are trying to enter. For Schank [93], learning by doing is the real issue, as he believes too much learning in the school system is by memorization.

Berryman [13] has identified the following characteristics of traditional apprenticeships, many of which also apply to cognitive apprenticeship:

- work is the prime motivation;
- performance standards are dictated by the community of practice;
- teaching is usually not explicit and teachers are mentors; and
- apprentices are initiated into a community of expert practice.

Apprenticeship involves modeling real world events and the performance of experts [58]. Cognitive apprenticeship extends the apprenticeship techniques associated with education for the professions listed above to the development of knowledge and skills typically targeted in classroom instruction. The term apprenticeship emphasizes the fact that activity is central to learning and that it is situated. Cognitive apprenticeship is a learning model that seeks to develop concepts out of and through continuing authentic activity [21]. Learners acquire, develop and use cognitive tools while engaging in authentic activities.

There is not much concrete evidence yet of how effective cognitive apprenticeship is in practice. Berryman [13], however, makes the point that cognitive apprenticeship is a learning technique that is unusually well-grounded since it inherits from the successful model of traditional apprenticeship and puts into practice the ideas and findings of dedicated educational researchers.

2.5.3 Collaborative Learning

Collaborative learning is a learning process that emphasizes active participation within groups of learners and the teacher(s) [112]. Learners develop their knowledge while sharing ideas, reflecting and interacting in learning groups. Schrage [94] defines collaborative learning as a process of creating shared meaning. Groups of learners develop a shared understanding that no individual previously possessed and that no individual could have constructed alone [94]. This is also known as collective problem solving.

Hewitt and Scardamalia [50] identified two forms of distributed cognition. The first is division of cognitive labour where groups of people collaborate by splitting the task into sub-tasks which are assigned to individual people. It is conceivable that this might cause learners to suffer from the compartmentalization bias, as introduced in Section 2.4.1. The second form of distributed cognition is shared cognitive labour where groups of people work on the same task together. Hewitt and Scardamalia favour the latter and argue that it is more self-scaffolding.

Adherents to sociocultural theories of learning believe that learning is a social negotiation of meaning. Since interaction with other people from one's community is essential for this social negotiation of meaning to take place, proponents of these theories necessarily advocate some form of group learning.

Constructivists make use of collaborative techniques because these encourage learners to consider information from multiple perspectives. In a collaborative learning environment, learners are exposed to many people's opinions and interpretations of objects and events. The learners' own ideas and conceptions are valued and are debated by the learning community. Since learners' misconceptions are challenged they have occasion to modify their knowledge structures. The learners therefore have potential to develop cognitive flexibility [15], that is, the ability to consider knowledge from multiple perspectives and to apply their knowledge to new situations [105]. This improved transfer of knowledge is very important to constructivists. Shaw [97] comments that much of the richness of collaborative learning environments is due to the different knowledge constructions that each participant brings to the interactions.

Aside from facilitating learning the task at hand, collaborative learning environments can help learners develop their collaborative work skills. Interpersonal and social skills are valuable in today's workplace and in society in general. These are therefore desirable for learners. Since much activity in the workplace is of a collaborative nature, a collaborative learning environment also increases the authenticity of the learning activities.

Brown et al. [21] identified four salient features of group learning: collective problem solving, confronting ineffective strategies and misconceptions, providing collaborative work skills, and displaying multiple roles. The first three have been treated above. Let us now consider the latter. In a collaborative learning environment, group members can share the multiple cognitive roles that are needed to address problems. This sharing can encourage reflection about the roles and foster discussion about the cognitive strategies being employed by the group. The difficult alternative in a non-collaborative learning environment is to equip an individual learner with the cognitive skills required to be able to display the multiple cognitive roles that are needed to solve an authentic problem.

Blumenfeld et al. [16] advocate group learning. They do, however, caution that success with group learning is dependent on informed implementation. Peer interaction is not the cure-all for educational problems. Improper implementation of group learning "can stigmatize low-achievers, exacerbate status differences, and create dysfunctional interactions among students" [16]. Collaborative learning environments therefore require careful planning. In order for group work to succeed, Blumenfeld et al. argue that educators must attend to the following:

- establishing group norms;
- choosing learning tasks that are conducive to meaningful student interchanges;
- judiciously forming groups of learners; and
- providing mechanisms to ensure accountability and fair evaluation.

Collaborative learning techniques seem to be an effective means for learning, when properly implemented [16]. Harasim et al. affirm that "collaborative groups facilitate greater cognitive development than the same individuals achieve when working alone ... [learners] appreciated the exposure to a diverse range of perspectives in the group learning" [48]. In a controlled experiment, Schutte [95] found that learners who had no face-to-face contact with the teacher but who collaborated together scored 20%

higher than the learners in a traditional classroom. In an empirical study, Scott et al. [96] found that learners collaborating over large distances can produce good quality artifacts and that virtual teams can be as effective as face-to-face teams. Turoff [112] has found that collaborative learning can be more interactive and effective than traditional learning, particularly for mature and motivated learners. These results suggest that it is desirable for learning environments to support collaborative learning.

This chapter has introduced the most popular learning theories of the twentieth century. More details have been presented about constructivism and sociocultural learning theories including their values and practices, how they addressed shortcomings observed in earlier theories, and instructional techniques that they advocate. Having acquired a reasonable understanding of these theories, it will now be possible to consider how well Web-based learning environments support them. Furthermore, the knowledge gained in this chapter provides the necessary grounding to understand why certain characteristics of learning environments might be desirable. For example, collaboration features are often considered an asset to Web-based learning environments but, without this chapter, we could not have understood that they are desirable because they enable distributed collaborative learning. Moreover, we would not have known why collaborative learning is believed to be effective. The valuable knowledge gained from this chapter will be used to inform the design of the framework for comparing Web-based learning environments presented in Chapter 4.

Chapter 3

Human-Computer Interaction Issues

The discipline of human-computer interaction (HCI) is concerned with the effective use of computing machines by people. Human-computer interaction practitioners are a diverse and multidisciplinary group whose primary fields of interest range from computing science to psychology, graphic design, sociology, engineering, et cetera. HCI is a broad discipline whose mandates include designing software that is: useful for the task for which it was designed, easily usable, and well accepted by its target audience. As such, HCI is critical to the success of any software application.

Prominent in much HCI work is a determination to know and understand a product's users to better support their wants and needs, and to accommodate their differences. This focus has led to user-centered design [78] and then to participatory design methodologies (e.g. [17]). User-centered design has more recently been extended to learner-centered design by Soloway et al. [103, 104, 106, 84]. HCI practitioners consistently advocate an iterative approach to software development with early and significant input from users and with frequent testing.

The design of user interfaces is an important aspect of human-computer interaction and is becoming increasingly important in the software industry as more and more resources are devoted to designing and evaluating innovative user interfaces. Designing software for the World Wide Web poses particular challenges to the design of user interfaces due in part to the limitation of bandwidth, and to the very nature of hypertext transfer protocol (HTTP) browsers as document retrieval systems. This latter restriction is a significant one for those trying to develop interactive software for the Web. The advent of programming languages like Sun Microsystems' Java will help remedy this situation [40]. Another difficulty posed by the Web is one of navigation. Providing users with an efficient hypermedia navigation system that is intuitive and that helps users visualize the information space is very important. This has proved to be a difficult task as many users of Web systems still report feeling disoriented and being lost in cyberspace.

One of the goals that some software developers set for themselves is creating a sense of presence for their users. This is what is meant, for example, when people talk about a Web-based learning environment being the *place* [110] where learners and teachers interact. This can be accomplished by providing a high degree of interactivity [27] and by using a metaphoric user interface.

The design of groupware presents some particular challenges to application developers. Users must be supported in their efforts to maintain awareness of others and of their shared workspace so that they can work together effectively. In order for most groupware applications to be successful, a high percentage of the intended users must adopt and use the product, and the benefits of using the product must outweigh the additional work required for the vast majority of users.

The remainder of this chapter will elaborate on some of these human-computer interaction issues and will serve along with Chapter 2 as the base upon which the comparison framework detailed in Chapter 4 will be built.

3.1 User-Centered Design

It is generally acknowledged that the term *user-centered design* was coined by Norman whose 1986 book *User-Centered System Design* [78] contains a collection of seminal papers on the topic. User-centered design is an iterative approach to developing human-computer interfaces and entire software systems that is driven by users' needs. The three main tenets of user-centered design are the following:

- users should be consulted and involved in all phases of a product's life cycle;
- a multidisciplinary team should be responsible for a product's development; and
- there should be a focus on competitiveness.

The term user-centered design itself implies the first of the above principles. First of all, software developers should identify precisely who will use their products. Next, product developers should determine the users' goals and needs, their problems, the tasks they wish to perform, the environment in which they work, the tools they require, and the methods they prefer to use. It is critical to know and understand the users. Users' input and feedback should be sought at all stages of product development and should drive each iteration of a product's development.

A multidisciplinary team should be responsible for designing the total user experience with a product: that is, everything the user sees, hears and touches. The team should consist of representatives from all the disciplines relevant to the product. This includes domain experts, human-computer interaction specialists, application developers, training and support people, as well as marketing specialists.

User-centered design recommends a relentless focus on the competition. In order to determine whether it is worthwhile to invest in the development of a product, it is necessary to determine whether the product will offer users significant advantages over the competition. It is recommended that developers precisely identify who their main competitors are. Some will be tempted to say that there is no competition or that the competition is every other product on the market. This is not acceptable. Those who are tempted to say that there is no competitor should ask themselves how users would get on in life without their product. This should identify who the competition is, even if it is a manual process. Those who say that every product on the market is their competitor should probably focus on the few products that have the most market share in the market segment they are targeting.

Developers should set measurable and trackable usability goals. Good usability goals are often defined in relation to the competition. For example, users will perform tasks x, y and z faster and more accurately with our product than while using products A, B and C, and will report better satisfaction.

Table 3.1: Economic benefits of user-centered design[1]

• Typical return on investment for user-centered design is 50 to 100 times the cost.
• User-centered design has caused reduction in the development cycle by 33% to 50%.
• IBM found that 87% of all calls to their help desks were related to usability.
• 80% of maintenance is due to missed user requirements.
• Design changes due to usability work were estimated to save IBM $6 800 000 in 1991 alone.

[1]Adapted from [102]

For user-centered design to be successful, management must be committed to managing for users. This means that managers must invest in assembling a skillful multidisciplinary team and that they must empower the team to design the product. Managers should ensure that measurable usability goals are set and that these are met before product release. Time and resources must be allocated in the project plan to allow for modifications and corrections subsequent to user feedback.

Soderston and Rauch [102] make a compelling case for user-centered design by surveying a number of studies that have shown the economic benefits of such design. User-centered design has been found to reduce total development time and costs, decrease training and support costs, reduce maintenance costs, increase sales and revenues, and increase productivity. A sampling of results reported by Soderston and Rauch are presented in Table 3.1.

The human-computer interaction community has generally accepted the principles of user-centered design. Scandinavian researchers, however, have been credited with developing an alternative software design methodology that has proven successful and that has recently been receiving increased attention from the HCI community in the rest of the world. The methodology is often called *participatory design* [11, 17]. With user-centered design, users may be consulted extensively but their role is typically limited to reacting to designs and providing feedback. As the name implies, participatory design gives users more responsibility and makes them equal partners on the design team. This enables users to work together with designers to build applications that best meet their needs. Some researchers such as Scaife et al. [91] agree that user-centered design does not empower users enough but find that participatory design is impractical to implement in some circumstances. Scaife et al. have proposed a middle-of-the-road methodology called *informant design* that they have employed in the development of learning environments. Calling users who participate in this process informants emphasizes the fact that they possess information that is valuable but as yet unknown to designers. Although the informants are not treated as equals, they do have substantial input into the design of the product.

Interested readers should monitor the developments of the proposed *ISO 13407 Human-Centered Design Process for Interactive Systems* standard [57] which is currently at the status of committee draft.

3.2 Learner-Centered Design

Soloway et al. of the University of Michigan's Highly Interactive Computing Group have embraced user-centered design and have extended its principles by creating learner-centered design. They make the point that while ease of use is an important goal, it should not be the sole objective of human-computer interaction practitioners [103].

Learning is no longer just for students in classrooms. With the acceptance of notions such as life-long learning and learning organizations, professionals are being called upon to constantly learn, change and grow. Since learning in the context of doing is often effective [103, 93] and since workers in knowledge-intensive industries almost all have personal computers, these increasingly powerful computers could be used to support learning and learners.

If users are at the centre of the design process, then it is natural to focus on a product's ease of use. If learners are at the centre of the design process, then the focus shifts to the development of understanding, performance and expertise. Learner-centered design therefore extends user-centered design by proposing the following four additional principles to address learners' special needs:

- software should help learners understand what they are doing;
- software should help motivate learners;
- software should accommodate the diversity of learners; and
- software should account for the fact that people change and grow as they learn.

The key to learner-centered design is scaffolding, as presented in Section 2.5.1. Soloway et al. [103] have proposed the TILT Model (Tools, Interfaces, Learner's needs, Tasks) to guide the learner-centered design of software. The TILT Model begins to identify particular scaffolding strategies that are effective in supporting the special needs of learners. Coaching and modeling are scaffolding strategies commonly used in intelligent tutoring systems that can help learners develop their understanding of concepts as they perform tasks. Tools need to be adaptable to accommodate learners as they change and grow. Interface scaffolding can flatten the learning curves that are common in technology-oriented contexts thereby increasing learners' motivation.

In a CHI '96 paper, Soloway et al. [104] further refined their learner-centered design framework. Four important elements of a learning environment were defined: context, tasks, tools, and interface. For each of these four elements, scaffolding strategies are proposed that support learners' special needs: growth, diversity and motivation. It was recommended that software be designed to support an authentic, project-based learning environment. Learners' tasks could be redefined to reduce and manage complexity thereby supporting learners' growth of understanding. In order to accommodate the diversity of learners, the tools that perform the task could be designed to support different learning styles and levels of expertise. The interface to tools could help learners with their wavering motivation by initially being approachable and personalizable, so as to facilitate initiating engagement. The interface could also help sustain engagement by being highly interactive and malleable, as well as by encouraging articulation.

Soloway et al. [84, 103, 104, 106] have put learner-centered design into practice to develop, test and even commercialize several software applications that make significant use of scaffolding strategies. Two important lessons they have learned about providing scaffolding in software, aside from being very difficult to accomplish, are the following. First, scaffolding must fade as learners develop more confidence and expertise in using a product. Second, scaffolding is not all or nothing. The scaffolding needs to be flexible enough to be present when the learner wants and needs it but to disappear when the learner wants to act independently.

3.3 Navigation

Navigating a large information space like the World Wide Web is a difficult task for users. Users commonly report feeling disoriented and lost in hyperspace while using the Web. To find their way on the Web, users ask themselves a number of navigational questions: Where am I? Where have I been? Where can I go? How big is the information space? Where do I want to go? How do I get there? Am I on the right track? Am I there yet? Users should have tools at their disposal that can quickly supply the information needed to answer these types of questions.

Other media must also support navigation. Books, for example, provide explicit and implicit navigational aids that readers understand. Readers expect books to be linear in nature. Readers therefore know which page they read before the current one and they have reasonable expectations of what will be on the following page. The thickness of a book and its page numbers give clues as to the size of the information space as well as to a reader's current position in the information space. The hierarchical structure of books that are broken down into chapters or sections can help readers navigate them. Tables of contents give readers a quick summary of the contents and the structure of a book as well as providing a mechanism for directly accessing portions of the book that may be particularly interesting. Indices allow readers to quickly find references to key words or phrases. Readers may also fold corners of pages or make use of bookmarks to tag pages for future reference.

The potential benefits of efficient navigation are many and include the following [109]. Good navigational tools can facilitate the traversal of information spaces by providing structure for otherwise poorly structured spaces. Efficient and natural navigation can reduce the cognitive and physical burdens of users thereby freeing up personal resources that can be applied to the task at hand. Improved navigation is also beneficial in that it can reduce the load on system resources since fewer requests need to be served when users find what they are looking for with less intermediate steps.

Darken and Sibert [28] have studied *wayfinding* in large virtual worlds. They define skilled wayfinding behaviour to be "purposeful, oriented movement during navigation." Bachiochi et al. [10] distinguish between the wayfinding needed for navigating between websites and for navigating within websites, calling the latter *structural wayfinding*. They state that inter-website navigation should be supported by the browser whereas structural wayfinding should be supported by the design of each web site. Web-based applications such as learning environments are a special case of websites whose navigation is even more so the responsibility of the application designers.

Few studies have been conducted to inform the design of navigational aids for Web-browsing; fewer still for navigating Web-based applications. Bachiochi et al. [10] made some recommendations for Web page layout based on a usability study they carried out and they established usability criteria for efficiently finding information on the Web. They determined that users should not have to traverse more than four pages nor should they require more than sixty seconds to find the information for which they are looking.

Tauscher and Greenberg [109] have done empirical research on how people revisit pages when browsing the Web that can inform the design of some types of navigational support tools such as history mechanisms. They found that there is a 58% probability that the next page a user will visit has previously been accessed by that user. This means that there is a high likelihood that users will revisit pages and that there is therefore considerable opportunity to ameliorate navigation by improving the mechanism by which users return to previously visited pages. In their research, Tauscher and Greenberg have considered specific navigational patterns and have used these to evaluate different history mechanisms that browsers could support. They found that the stack based history mechanism that is prevalent in today's browsers is not optimal in its ability to predict which page a user may want to revisit. A superior mechanism is one which remembers the pages visited sorted by recency with duplicates positioned only in the most recent position. It is conjectured that this mechanism is based on a simpler conceptual model and that it would prove to be more usable than the stack based mechanism.

The navigation of large and complex information spaces is a difficult problem that has not yet been solved [22]. CZWeb is an example of a prototypical navigational tool that employs the innovative continuous zoom algorithm [32] to help people navigate complex information spaces [24]. Further research and development are needed to refine tools like CZWeb and to inform the design of other tools for navigational support in general and for Web-based applications in particular.

3.4 Metaphors

Metaphors are ubiquitous in our everyday language. They help people understand a target domain in terms of a more familiar source domain. Metaphors usually do not provide a direct mapping between the source and target domain but rather emphasize and suppress certain details of the target domain by drawing comparisons with the source domain. Since the popularization of the Apple Macintosh's desktop metaphor, metaphors have been commonly used in user interfaces.

Metaphors in user interfaces are often conveyed with graphics and icons. It is important, however, to realize that user interface metaphors can exist without any associated graphics or icons. An example of this is the UNIX *pipe* command.

Erickson [38], from Apple Computer, proposes five questions to pose in evaluating user interface metaphors. The first question to consider is the amount of structure provided by a metaphor. Does the metaphor evoke enough related ideas and concepts in users' minds so that it can be used to represent a complex target domain? The second aspect to assess is the applicability of the metaphor. Are there misleading associations that users might make between the metaphor and the target domain? The third question deals with the representability of the metaphor. Are there distinctive images, sounds and vocabulary that are representative of the metaphor? The fourth question asks whether or not the metaphor is suitable to the intended audience. If the metaphor is not understood by the members of the target audience then it is worthless. The final aspect to assess is a metaphor's extensibility. Can the metaphor grow with the system it represents and does the metaphor suggest unintended uses of the system to users? Erickson emphatically makes the point that end-users must be consulted and observed throughout the design process of user interface metaphors. They are therefore necessarily involved in the evaluation of the above five questions.

Education researchers have found that teaching with metaphors, comparisons and analogies can be effective in helping people to learn [100]. Although the use of metaphors in user interfaces is very common and is advocated in many design guideline documents including those from Apple Computer and IBM, there is very little empirical research that supports the folk view that user interface metaphors can facilitate the use of software.

Smilowitz's [100] work is an exception. Experiments were conducted to determine whether or not metaphoric user interfaces aided user performance and led to better user satisfaction. Experiments were also conducted to identify features of good user interface metaphors. In the cases studied, the experiments clearly showed that user interfaces with good metaphors were superior to non-metaphoric interfaces. This supports the common beliefs about the use of metaphors in user interfaces. The experiments also showed, however, that user interfaces designed around poor or inappropriate metaphors provide no advantage over non-metaphoric interfaces. Usability testing is needed to determine whether or not a particular metaphor is effective.

Smilowitz's experiments also inform us about how a metaphor is successfully carried by a user interface. The results show that the effect of the metaphor is not at all provided through the graphics or icons, but rather through the user interface's terminology. This is a surprising finding since software developers typically tend to value the look of a user interface much more than its terminology. The consequences are that user interface designers should concentrate on using appropriate terminology in their user interfaces and that they should invest resources in conducting usability tests to evaluate how well the chosen words help users understand the metaphor. This should be done even at the expense of graphical design.

Researchers have debated whether integral metaphors are superior to composite (or mixed) metaphors. Some researchers have shown a preference for using a single metaphor that provides a unified structure to the user interface (an integral metaphor), even at the expense of additional mismatches between the metaphor and the target domain. Others advocate the use of composite metaphors because they provide better coverage of the target domain and because they cause the user to have multiple perspectives on the target domain. Little empirical research exists to support either of these views. Smilowitz's experiments suggest that user interfaces with good integral metaphors yield better user performance and satisfaction than do user interfaces with composite metaphors. If the integral metaphor is poor, there is no advantage over an interface with composite metaphors.

The empirical research on the effectiveness of metaphors in user interfaces is only preliminary. More research is needed to validate Smilowitz's results, to identify more desirable characteristics of user interface metaphors, and to inform the effective design of user interfaces around good metaphors.

3.5 Groupware

Distributed collaborative learning supported by technology makes use of groupware. Groupware is simply defined as software that is used not by individuals but by groups of people. The people who use groupware are usually working together but are geographically separated. Developing successful software that will be used by groups of people tends to be much more difficult then developing software for individuals. While it is not recommended that one rely only on intuition when designing user interfaces for single user applications, doing so for groupware applications is even more dangerous. The software community has less experience with developing groupware then it does with single user applications and the supporting research is less well established.

Much of the research that has been done on groupware has focussed on computer-supported collaborative work (CSCW) and it deals almost exclusively with synchronous, or real-time, collaboration. The groupware research that is particularly relevant to this thesis is that which is focussed on computer-supported collaborative learning (CSCL).

The goal for educational groupware should be to support the existing practices and processes of group learning which are effective. However, groupware applications do not yet provide the richness of face-to-face interactions [47]. Human-computer interfaces that have this deficiency may limit the collaboration opportunities for learners and increase the cognitive overhead on learners who have to compensate for the deficient interaction.

3.5.1 Awareness

Greenberg et al., from the University of Calgary's GroupLab, have done an impressive amount of research on groupware. Their research has led them to develop *GroupKit* [89] which is a toolkit of groupware widgets that help support people as they learn and work together.

Greenberg et al. have identified users' lack of awareness as one of the key problems with groupware. People must have general awareness of the people around them and their availability to be able to initiate electronic meetings. Greenberg [43] highlights the need for better mechanisms to support this informal awareness. Learners who participate in a collaborative learning environment must also have an awareness of what is going on around them in order to participate fruitfully in collaborative tasks. Greenberg et al. have proposed prototypical solutions to groupware awareness problems.

Gutwin et al. [47] identify four types of learner awareness in their framework of awareness in collaborative learning: social, task, concept and workspace awareness. These are defined in Table 3.2.

Learners can gain social awareness through informal communication and interaction with other learners. Task awareness and concept awareness are often supported

Table 3.2: Types of learner awareness

Social awareness	Awareness of the social connections within the group including members' roles.
Task awareness	Awareness of how the task at hand will be completed.
Concept awareness	Awareness of one's knowledge about a concept and how it relates to the task at hand.
Workspace awareness	Awareness of other learners' interactions with the shared workspace.

with scaffolding as in CSILE [50] and in CoVis' Collaboratory Notebook [35, 37, 36]. Workspace awareness is the responsibility of the user interface.

Workspace awareness is up-to-the-minute knowledge of other learners' interactions with the shared workspace. In face-to-face interactions, learners have visual and auditory cues of other people's activities which help them know who else is participating, where they are, what they are doing, what they have done, what they plan to do next, and how they can be helped. The many ways that people learn collaboratively all require that learners have a clear understanding of other learners' interactions with the workspace. Workspace awareness helps people coordinate their tasks and the use of resources [46]. Workspace awareness is also important because it facilitates the use of workspace artifacts as conversational props. This allows learners to converse more naturally about artifacts in the shared workspace by gesturing to the particular items in question.

Collaborating learners necessarily share an overall objective. The tasks undertaken to achieve the overall goal may be mostly individual or they be group tasks. Gutwin et al. [47] call this *task separation*. If learners are closely working together on specific tasks then they need fine-grained workspace awareness. They must be aware of the other learners' precise location and activity. If collaborating learners do not share individual tasks then they need a coarser-grained workspace awareness which gives them just enough knowledge of other people's location and activity to know whether it affects their own activities.

The type of awareness required also depends on what Gutwin et al. [47] call *view separation*. If collaborating learners have the same view of the shared workspace then they will be interested in the precise objects that others are manipulating. If collaborating learners have different views on the shared workspace than they will be more interested in knowing where others are in the workspace rather than the precise artifacts with which they are working.

A difficult challenge for developers of tools that support workspace awareness is understanding how people maintain workspace awareness in face-to-face environments. They then need to find ways to support this in distributed environments. A further challenge is striking a balance between providing learners with enough information about other people's interactions within the shared workspace, and not distracting them from their own tasks by providing too much information. Only careful evaluation with users can determine how well widgets and tools support workspace awareness.

3.5.2 Adoption

For a software application to be successful, it must be used. Software purchasing decisions are often in the hands of administrators who are not the intended users, and these decisions are often made without sufficiently consulting the eventual end-users. End-users do, however, usually have the power to accept or reject software applications by choosing whether to use them.

The adoption of software is important for single user applications but it is critical for groupware. If 20% of home computer users make use of a certain vendor's application for tracking their personal finances, then that vendor has successfully captured 20% of that market. On the other hand, if even a few people in a company do not consistently update the information in their electronic calendars, then the company's electronic meeting scheduling application becomes unusable and useless. *Critical mass* is the term usually used to designate the minimum percentage of adopters required for a groupware application to be successful. Adoption is critical to groupware because relatively few dissenters can nullify most of the benefits of a groupware application.

Another factor that affects the successful adoption of groupware is the ratio of additional work to perceived benefit [39, 45]. If some users perceive that the benefits they receive from using a groupware application are not commensurate with the extra time and effort required in using the application, then they may resist using it. This could be the case with electronic calendaring systems, for example, where users who seldom schedule meetings might feel that the time and effort required to maintain their personal calendar outweigh the benefits that they personally receive. Since critical mass can be difficult to attain, groupware developers should attempt to minimize the number of users who might perceive an unfavorable additional work to benefit ratio.

Some human-computer interaction researchers and practitioners [45, 80] concern themselves with the social and cultural setting in which software is used. They point out the necessity of not only trying to understand users in isolation but to understand groups of users in their work environment.

Well designed software that is easy to use may still fail if it conflicts with the users' environment. This is again particularly important for groupware applications where the success of the application hinges on multiple people being able to use it together. Consider, for example, the deployment of a collaborative learning system in an environment where all the performance rewards are individualistic. The conflict between the cooperation expected from the learning system and the individualistic reward system could very well cause stress for users and hinder the acceptance of the software.

The adoption of groupware applications can also be hindered by disruptions it causes to established practices in an organization. Wilson [118] states: "Many innovations fail because they neglect the changed roles demanded of the people expected to use the technology." Web-based learning environment developers, for example, should be cautious of requiring any changes in the roles of instructors and learners. Any required changes in established practice should be made explicit and supported.

Table 3.3: Usability goals

• Easy to buy	• Intuitive
• Easy to set up	• Engaging
• Easy to learn	• Useful
• Easy to use	

3.6 User Interface Design Principles

As summarized in Table 3.3, common usability goals for software products are that the software be easy to buy, easy to set up, easy to learn, easy to use, intuitive, engaging, and useful. Over the years, many researchers and companies have published user interface guidelines (e.g. [56]) to help software developers attain these goals. Furthermore, a large number of printed and online resources about Web design have become available since the popularization of the Web (e.g. [55]). Rather than create another list of user interface guidelines that developers of Web-based learning environments should adhere to, we will review some established user interface principles [11, 55]. The principles to be reviewed are itemized in Table 3.4.

3.6.1 Simplicity

The simplicity principle states that usability should not be sacrificed for function. Users will benefit from a user interface that makes doing the usual tasks easy and the less common tasks possible. Products should only provide functionality that is necessary as determined by user-centered task analysis or use case modeling. Adding unnecessary functionality to a product typically detracts from the user interface because it complicates the interface, adding more objects and possible actions that the user has to understand.

Table 3.4: User-interface design principles

• Simplicity	• Feedback
• Familiarity	• Consistency
• Affordances	• Flexibility
• Context	• Responsiveness
• Visibility	

3.6.2 Familiarity

Users have a conceptual model of software applications they use which captures their goals and beliefs. Software implementors have a different conceptual model of the software application. Their model is in terms of data structures, algorithms and code libraries. User interfaces bridge these two models. The difference between the users' model and the model conveyed by the user interface represents what users need to learn to successfully use the software. The more the user interface resembles the users' model of the software, the easier the software will be to learn.

In general, interfaces should leverage users' existing knowledge. Although innovation in interfaces is sometimes necessary, it often requires that users learn new concepts or methods of interaction. It is usually best to employ slightly inferior user interface artifacts and controls that are familiar to users rather than creating new ones. When innovation is necessary and justifiable, user interface designers should leverage users knowledge of the physical world to inspire their innovations.

3.6.3 Affordances

Good user interfaces provide affordances. Affordances are cues that enable users to reliably infer the meaning of objects and the outcome of actions without needing the support of labels, pictures, or instructions [77]. They are the perceived properties of objects that give users clues about how the objects can be used. For example, knobs afford turning, handles afford pulling and buttons afford pushing. Effective affordances for user-interfaces are often based on real world objects or concepts with which users are familiar that is, metaphors. User interfaces can take advantage of affordances to become more intuitive for users.

3.6.4 Context

The individual tasks that users perform are usually situated in an overall context. The user interface should provide sufficient cues so that users know where they are with respect to a more global context. There are two forms of context that particularly concern users. One is navigational context and the other is task context.

Navigation context gives users a sense of where they are with respect to the entire system, including where they have come from and where they can go next. The navigational context is particularly important for large applications, like the Web, that are relatively unstructured and in which users tend to get disoriented.

Task context is important to reflect to users what task they are working on and how far they have progressed. The current state of the system is thus conveyed to the user. Information on the progress of tasks is particularly important for relatively long and multi-step tasks, such as software installation.

If users are able to walk away from their terminals, return the next day, and quickly understand from the user interface cues where they last left off, then the interface probably provides good context.

3.6.5 Visibility

In user interface parlance, visibility refers to the degree to which the needed cues are available and easily recognizable. When a user wants to perform a task, the controls with which they can do so should be immediately visible and easily identifiable. Users are reluctant to search for occluded information so user interface designers must be meticulous in understanding the tasks users want to perform and ensure that the appropriate information and controls are readily visible. User interfaces with good visibility also help users explore the software application by clearly displaying what actions are available at any particular time. Furthermore, they do not lead users to attempt actions that are not permitted.

3.6.6 Feedback

It is very important that user interfaces consistently provide feedback to users. Feedback is necessary so that users know whether a command they issued has been successfully interpreted by the computer. When abnormal conditions arise, users should be given informative feedback on what happened, how to recover from the situation and how to proceed with their task. Feedback on progress is also very important for users when the application is engaged in time consuming tasks. Feedback can take many forms but it must be designed into the user interface. The best type of feedback is unobtrusive but clear in its meaning.

3.6.7 Consistency

A user interface's terminology, visuals and behaviour should be consistent. The same phrases should always be used to denote interface objects and actions, and these phrases should not be overloaded with multiple meanings. A consistent writing style should be used for all the written materials users encounter. Colour and graphics should carry the same meaning throughout the interface. Furthermore, colours should be associated with concepts in a manner that does not violate accepted conventions and that is consistent with the established cultural practices of the users.

A sure way to frustrate users is to present them a user interface whose behaviour seems to be arbitrary. The navigation style should be consistent throughout the interface and similar commands executed on similar objects should always produce similar results. User interfaces that are predictable in their behaviour give users confidence.

3.6.8 Flexibility

User interfaces should be flexible and versatile. Flexible user interfaces account for system differences like screen resolution and network connection bandwidth as well as for differences in users' goals, abilities and preferences. System differences can be handled by providing user interfaces specifically designed for common system configurations.

Accommodating differences in users is more complex. The simplest way to support different users' needs is to add redundancy to the interface so that user interface artifacts are represented in more than one way. Each of the alternative representations might be more meaningful to different people. Providing alternative means for accomplishing tasks is another type of user interface redundancy that can help accommodate different users' needs and ways of thinking.

Adaptable user interfaces are slightly more complicated. They provide users a set of interface parameters that can be manipulated to tailor the interface's appearance and its behaviour. Interfaces that provide users with a great deal of flexibility by giving them many relatively unconstrained user interface parameters may actually be providing them a disservice. If unchecked, users may choose a poor combination of parameters that moves the interface away from the usability goals of Table 3.3.

Adaptive user interfaces are more sophisticated user interfaces based on artificial intelligence techniques. Adaptive user interfaces monitor individual users' performance with the system and automatically adapt themselves to each user in order to make their experience more efficient, engaging and enjoyable. The changes must be made in a way such as to not surprise users or make them feel like they are not in control of the system.

3.6.9 Responsiveness

User interfaces should support the best interaction styles and the best visual design possible while maintaining acceptable responsiveness. Poor system responsiveness frustrates and confuses users. Web-based applications are particularly susceptible to poor responsiveness due to limited bandwidth. Developers must ensure that their applications are responsive enough for their intended users, even at the expense of other user interface design goals.

Chapter 4

Comparison Framework

There have been a few attempts recently at comparing existing Web-based learning environments [64, 14, 99]. In doing so, the evaluators have determined sets of criteria for comparing the learning environments. The main difficulty that we have observed is that the comparison criteria are usually just a list of features that can be used by learners, instructors and system administrators. These criteria are silent on many other attributes of learning environments that should be considered when evaluating them. Furthermore, the criteria mislead people into believing that more features are better even if the features have not been shown to actually be desirable in general or needed by the adopting organization.

In this chapter we present a practical comparison framework for evaluating Web-based learning environments. We will draw primarily from the educational research presented in Chapter 2 and the human-computer interaction research presented in Chapter 3 to identify desirable attributes and features of learning environments. We will also consider important requirements for learning environments that stem from concerns over other matters including business and economics, hardware and software, maintenance and support, training, and ethics. As such, the comparison framework will structure the evaluation and comparison of competing Web-based learning environments over a much broader and somewhat abstracted set of criteria than just the tools used by learners, instructors and system administrators.

The comparison framework will consist of seventeen dimensions along which learning environments can be compared. All of the dimensions will be explained and some of the reasons why each of them is important will be discussed. It is expected that the relative importance of the dimensions will depend on an individual's opinion and the context into which a candidate learning environment will be deployed.

The development of software always involves tradeoffs so no Web-based learning environment is expected to be the best in all of the identified dimensions. Users of the framework who are considering adopting a learning environment will therefore have to establish their organization's priorities. As the importance of education theory and human-computer interaction requirements have been demonstrated, it is hoped that framework users will highly value the comparison dimensions that capture these requirements.

The comparison framework is designed to serve two main audiences: developers and adopting organizations. Developers of Web-based learning environments can use the framework's dimensions as benchmarks against which to measure their product and can benefit by discovering ways to improve their product. Organizations considering the adoption of a Web-based learning environment can use the framework to guide and structure their evaluation of competing candidate learning environments. We propose a user-centered method in Section 4.2.2 that makes use of the comparison framework to help organizations select the Web-based learning environment that is best for them.

4.1 Comparison Dimensions

In this section we will identify the comparison dimensions that make up the comparison framework. Each dimension will be named and explained. We will also identify a number of system features and attributes that contribute to each dimension and that the framework user can look for when comparing learning environments. Note that the comparison dimensions are not orthogonal because learning environment attributes and features may support or detract from more than one dimension.

Gibson et al. [41] identified six dimensions along which Web-based testing and evaluation systems could be evaluated. These were testing, tracking, grading, tutorial building, implementation issues and security issues. Their dimensions have been subsumed by the dimensions of the comparison framework proposed herein.

The first four comparison dimensions identified below fall out of the discussion on techniques that foster learning in Section 2.5 and on human-computer interaction in Chapter 3. The other dimensions deal primarily with other aspects of Web-based learning environment applications yet they are sensitive to education and HCI goals for Web-based learning environments. The dimensions were not ordered by which type of user they are likely to impact because many dimensions affect more that one role in a learning environment. Table 5.1 on page 111 summarizes our expectations of the relative importance that different groups of users will attribute to each of the comparison dimensions.

4.1.1 Facilitating

Web-based learning environments can be compared as to how well they facilitate learning. There are two aspects to this comparison dimension. One is how well learning environment applications directly facilitate learning and the other is how well learning environments support instructors who choose to take on the role of a facilitator rather than that of the more traditional knowledge transmitter.

Recall from Section 2.5.1 that two instructional techniques for facilitating learning are scaffolding and fading. Instantiations of these techniques can be programmed into learning environments. An example of this is Carroll's [23] user interface training wheels. These metaphorical training wheels support users as they learn an application's user interface and then fade away as the users become more proficient. Learning environment applications with artificial intelligence may be able to offer learners timely support in their learning endeavors.

Learning environments which do not directly facilitate learning may nevertheless support it. While some learning environments might cast instructors in the traditional role of knowledge transmitter, others might make it easier for instructors to assume the role of a facilitator if they choose to do so. Minimally, learning environments'

Table 4.1: Types of facilitating

• Directly facilitating learning
• Support for facilitating learning
• Specialized scaffolding

user interfaces could use language that is sensitive to instructors playing the role of facilitator. Learning environments could provide better support for facilitation by providing tools that help instructors act as facilitators. Tracking tools could help instructors observe learners so that they can provide more appropriate scaffolding to individual learners and time their fading optimally. Communication tools could enable intellectual support in the form of comments, suggestions, feedback and observations, as advocated by Teles [110].

Web-based learning environments might also have specialized features that provide scaffolding for particular learning activities. For example, Edelson et al.'s Collaborative Notebook is said to "provide a scaffold for students as they learn to conduct collaborative, open-ended investigations" [36]. This tool provides scaffolding for learners by giving them a structure within which to perform their scientific inquiries. The Collaborative Notebook resembles a scientist's notebook. Learners assign page types to each page of the notebook they author to explain whether the contents are a question, a conjecture, evidence for, evidence against, plans, steps in plans, or commentaries. They also create associations (hyperlinks) between pages. This structure scaffolds learners in their early attempts at scientific inquiry.

4.1.2 Cognitive Apprenticeship

Web-based learning environments can also be compared based on how well they support cognitive apprenticeship, as described in Section 2.5.2. Since cognitive apprenticeship is collaborative by definition, learning environments that support cognitive apprenticeship should have strong support for collaboration. Collaboration is treated in its own comparison dimension in Section 4.1.3 so it will not be given any further

Table 4.2: Cognitive apprenticeship support

• Interaction with community of practice • Engagement in authentic activities • Observation of others

consideration in the present dimension.

Cognitive apprenticeship requires learners to participate in a community of practice. Web-based learning environments can provide a means for novices from a community of practice (learners) to interact with experts in the community. Alternatively, learning environments can provide realistic simulations of communities of practice in which learners can interact and develop the expertise they need to become masters in their community.

For cognitive apprenticeship, it is critical that learners engage in authentic activities that are endorsed by the community of practice. Since the emphasis is not on typical school activities but on activities relevant to the community of practice, learning environments that support cognitive apprenticeship must usually be highly customized for each different learning domain. Learning environments cannot easily provide generic tools that implement cognitive apprenticeship for all domains.

Legitimate peripheral participation is advocated by supporters of cognitive apprenticeship. For learners to participate on the periphery of a community of practice, they must be able to observe the behavior of others in the community. Web-based learning environments can therefore support legitimate peripheral participation by providing controlled mechanisms that allow learners to observe what others are doing. As learners observe others, they may implicitly learn about the norms of the community of practice and begin the process of enculturation.

4.1.3 Collaboration

Collaboration facilities are elements of the learning environment that enable learners and teachers who are not co-present to communicate and work together. There are many ways in which technology can support collaboration. These can be grouped into two categories depending on whether they are asynchronous or synchronous.

Asynchronous collaboration tools enable people to work together in a time-independent manner: that is, all parties need not be participating at the same time. In this case, the learning environment does not guarantee that all participants' views will be the same at all times (there may be time delays).

Synchronous collaboration tools enable people to work together at the same time, mediated by the software. The learning environment guarantees that changes made in one participant's view will be propagated to the other participants in a timely manner. Although synchronous collaboration may seem more technologically advanced than asynchronous collaboration and may therefore appear to be superior, this is not necessarily the case. We have said that one significant advantage of distributed learning is that it removes the need for co-presence. Notice that synchronous collaboration nullifies half of this advantage because although learning can still be place-independent, it is not time-independent.

Asynchronous

Perhaps the simplest form of collaboration that a learning environment can provide is electronic mail. With email, a person can send a message to a specific person or to a group of people. A standard enhancement to many email systems is the facility to attach multimedia documents to messages. Email can also be encrypted and digitally signed to guarantee privacy and authenticity of messages. Learning environments may provide a directory lookup service so that a person can find someone else's email address.

Mailing lists can also be used for collaborative purposes in learning systems. Email sent to a particular mailing list will be redistributed to all those who have subscribed to that list. Mailing lists can be enhanced with security features that limit who can subscribe to a particular list, who can send messages to the list, and whether list subscribers can know who the identity of other subscribers.

Discussion group systems are frequently used in collaborative learning environments. These are quite varied in sophistication and are known by many names including newsgroups, bulletin boards, and conferencing systems. Discussion groups are repositories of messages that people can read and contribute to. Each discussion group usually has a topic on which discussion is supposed to be centered. Many security models can be added to discussion group systems to control who can read, author, change and delete messages, and who can create, change, and delete discussion groups.

Some discussion group systems can simulate the types of discussions that occur in classrooms and in small groups. Specialized discussion group systems could support structured argumentation, for example, through the use of templates. More advanced discussion group systems may also facilitate certain types of discussion which are not even possible in a face-to-face setting. For example, when a question is posed to a group in a face-to-face interaction only one person has the chance to speak before hearing other people speak. In a specialized online discussion, every participant could be obligated to comment on the question before having access to anyone else's response.

Shared workspaces are important for collaboration in a learning environment. A shared workspace is "a software environment containing learning and work artifacts that can be viewed and manipulated by anyone in the group" [47]. This shared workspace can be used to store any file that a group of collaborators needs to share. File locking is an enhancement to shared workspaces that protects file integrity by allowing a maximum of one group member to make changes to a particular file at any one time. Another enhancement is version control which allows group members to keep a history of the changes made to a file and to roll back changes if needed.

A group review system is a workflow application that manages the review of a document by multiple people. Many different workflows can be supported. An example of such a workflow is the serial review. In a serial review, the author chooses a sequence of reviewers who are sent the document one after the other. Any changes made by earlier reviewers are visible to later reviewers. The author may choose to be notified when a document has completed the review cycle or after each individual review.

An *asynchronous decision support system* is a system that facilitates the process of group decision making. Decision support systems support many types of activities to assist groups of people in making decisions. Three of these are brainstorming, prioritizing and voting. Brainstorming is used for generating ideas. Prioritizing could be used for ranking alternative solutions to a problem. Voting is a facility for voicing one's opinion with respect to a particular question. A decision support system could perhaps be used creatively in a learning environment to enhance collaboration.

Annotations are explicative or critical notes that accompany a document. These notes may either be private, shared with a limited group of people, or accessible by everyone. With shared annotations, people may be allowed to follow-up and critique other people's annotations. Rutherford [90] has labeled shared annotations the key to collaborative learning and teaching on the Internet.

Shared annotations can serve as a powerful collaborative tool because they help centre discussions about particular objects. The document being annotated provides a context for discussion [29]. In their treatment of annotations, Röscheisen et al. [88] stressed that annotations add value to the content of the documents to which they refer. Shared annotations might also be useful for encouraging learners to challenge each other's misconceptions. Rutherford [90] suggested the following six benefits of annotations for education.

- learners benefit from seeing questions they might not have thought of themselves;
- instructors are less likely to have to answer the same questions over and over;
- comments on assignments that instructors may put significant effort into can benefit more than one learner;
- learners have easier access to the wealth of information generated by their learning community;
- learning materials improve over time by adapting to students' needs; and

Table 4.3: Collaboration features

Asynchronous	Email Mailing lists Discussion groups Shared workspace Group review Decision support system Shared annotations Shared semantic associations General awareness
Synchronous	Chat Internet phone Video conferencing Navigation control Group authoring Decision support system Shared knowledge representation General awareness Workspace awareness

- documents are structured meeting places rather than being static one way communication devices.

Learning environments may provide a facility that allows learners to create semantic associations between documents in the learning environment which can be shared with others. This facility can be enhanced by including annotations with the links which allow learners to articulate their reasoning for having associated documents. By reviewing each other's web of associations, learners would be exposed to multiple perspectives on a concept and could challenge each other's misconceptions.

Synchronous

Learning environments may also have a number of synchronous collaborative facilities. Recall that these are collaborative tools which require simultaneous participation.

One such facility is known as chat. Chat enables two or more people to talk to one another by typing text which is displayed simultaneously on all the chat participants' screens. Different chat sessions might be provided to discuss different topics. Various security models can be added to chat facilities which limit who can participate in a particular chat session, who can create a new session and who can delete an existing session.

Internet phone is a telephone-like communication tool that uses the Internet as its carrier rather than long distance telephone lines. Participants with microphones and speakers can talk with other similarly equipped participants who are plugged into the Internet. Internet phone systems may also have a directory service which allows people to discover other people's Internet phone number; that is, their IP address.

Video conferencing systems combine voice communication provided by Internet phone-like tools with video. With video conferencing systems, groups of people equipped with desktop video cameras can see their co-participants as they speak to them.

Synchronous navigation control is a system by which a group master can drive the other group participants' displays to give a demonstration. The navigation control includes "retrieval and display of common documents, synchronous document scrolling, mouse positioning, and common highlights and/or handwriting on documents" [119]. The synchronous navigation control could be enhanced by allowing group members to be peers who take turns behaving as masters.

Learning systems might include a synchronous group authoring system which allows group participants to work together on a single document at the same time. Such a system would have a mechanism to maintain the integrity of the document and for informing every team member of the others' activities. Greenberg's [42] *WYSIWIS* ("what you see is what I see") text editor is a good example of a synchronous group authoring system. A simpler example is shared whiteboards. A shared whiteboard is a canvas displayed on all participants' screens on which standard paint program drawings can be made. It allows groups of people to work on diagrams and text together.

A synchronous decision support system could also be provided to help groups of people in different locations arrive at decisions in real time. The activities that such a system could support are similar to those of the asynchronous decision support system described above.

Synchronous shared knowledge representation tools could be useful groupware applications. These could give learners a mechanism for expressing their understanding of concepts to others and for others to confront their misconceptions. Shared concept mapping tools, as proposed by Niguma [76], are an example.

As discussed in Section 3.5.1, it is important for synchronous groupware applications to provide cues to support awareness. When considering learning environments that claim to provide rich synchronous collaborative tools, evaluators should closely examine to what degree these tools provide awareness of the shared workspace. Features that support awareness of the shared workspace are covered in the user interface comparison dimension in Section 4.1.4 and will therefore not be repeated here.

All collaborative learning environments, whether synchronous or asynchronous in nature, should support general awareness. General awareness is necessary for learners to know who else is in the learning environment and whether or not they are available for collaboration. Collaborative learning environments should provide mechanisms to facilitate making contact with others. Basic support for general awareness can be provided by personal electronic address books which inform learners of who else is in the learning environment and how they can be contacted. General awareness can be further supported with tools like Peepholes, Portholes and Media Spaces [43].

With Peepholes, each participant is assigned a labeled iconic presence indicator. The Peepholes system automatically updates the icons based on people's activity in the workspace. By checking other people's icons, learners get clues as to their availability for collaboration. People who want to collaborate with someone who is temporarily unavailable can ask to receive notification from the system as soon as that person re-enters the learning environment.

Portholes and Media Spaces require every participant in the learning environment to have a video camera on their desktop and to always have it turned on. Learners can check each other's availability for collaboration by checking the output from the other's video camera. Media Spaces provide a continuous video stream whereas Portholes economize bandwidth by periodically sending video snapshots.

4.1.4 User Interface

Participants in any learning environment will appreciate a good user interface (UI). The quality of the user interface determines not just the visuals of the learning environment but the total user experience. Norman [78] makes the point: "From the point of view of the user, the interface *is* the system." Organizations considering adopting a learning environment should therefore carefully evaluate each candidate system's user interface to determine whether it is suitable for their members.

The design of user interfaces involves many tradeoffs so different user interfaces are likely to be most effective for different learning circumstances. For example, an interface that is simple and easy to learn is likely to be more effective for short courses where people spend little time using the system, whereas a more elaborate application with advanced features might be more suitable for learners who use it for entire semesters, even if it is harder to learn. Young children and adults have different interaction styles so different user interfaces are required by these two groups. When selecting a learning environment, organizations should ensure that its user interface meets their users' needs.

As was discussed in Chapter 3, there are many factors which contribute to a good user interface. In the remainder of this section, we will identify some user interface attributes that evaluators can consider when comparing Web-based learning environments along the user interface dimension.

User interface principles

Organizations considering the adoption of a Web-based learning environment should do more than just a cursory user interface evaluation based on a few people's intuitions. When comparing the quality of competing Web-based learning environments' user interfaces, evaluators should make a principled evaluation based on established user interface principles such as those discussed in Section 3.6.

The best user interfaces will mask any unneeded complexity and present users

Table 4.4: User interface characteristics

• User interface principles • Metaphors • Workspace awareness

with a simple representation of the system that is consistent with their conceptual model. Interfaces that are intuitively familiar and that provide affordances tend to be easier to learn and more engaging for users. The interface should maintain task context and navigational context to help users remain oriented and aware of their surroundings. The information and controls users need to perform each of their tasks should be clearly visible and easily recognizable at the time they are required. When users initiate an action or when the system is busy performing some processing, the user interface should provide intelligible and timely feedback. The look and feel of the interface should be consistent throughout the application, and the interface should behave predictably. Good user interfaces will take into account the differences between people and will be flexible enough to be customizable for individual users. For an application to be successful, it is a definite requirement that it be responsive. A user interface that suffers from poor responsiveness will frustrate users even if it is simple, familiar, takes advantage of affordances, maintains context, provides good visibility and feedback, is consistent, and is flexible.

Metaphors

User interfaces built around good metaphors tend to be more effective than non-metaphoric user interfaces. Organizations should therefore generally favour Web-based learning environments whose user interface is metaphoric. Organizations should make use of Erickson's [38] five rules of thumb to determine whether or not a candidate learning environment's user interface metaphor is good. Recall from Section 3.4 that those involve determining whether a metaphor:

- provides sufficient structure;

- is applicable to a learning environment interface;
- is representable;
- is understandable by participants of the learning environment; and
- is extensible.

There is some evidence to suggest that composite metaphors can detract from a user interface so organizations should be cautious when considering such interfaces. When evaluating metaphors, remember to not only consider the graphics and the icons that carry the metaphor but to also consider the terminology since some studies have shown that most of the effect of a metaphor is provided through the terminology rather than the graphics.

Workspace awareness

Learning environments that provide synchronous collaboration may provide widgets to support workspace awareness. Different types of widgets should be used depending on task separation and view separation, as discussed in Section 3.5.1. Gutwin et al. [47] give examples of widgets in the University of Calgary GroupLab's GroupKit that support workspace awareness.

To support learners collaborating on the same task and with the same view of the shared workspace, strict-*WYSIWIS* screen linking is provided. Users can configure how closely the screen actions are linked, that is, how often they receive updates on other people's interactions with the workspace. GroupKit also uses multiple cursors to provide fine-grained awareness. Each participant sees everyone's cursor on their screen and the cursors are labelled to indicate to whom they belong. This allows people to monitor each other's movement and activity in the workspace as well as providing a means for people to gesture and to point to workspace artifacts.

To support learners collaborating on the same task but with different views of the shared workspace, relaxed-*WYSIWIS* techniques are used. A multiple scroll-bar widget is included in GroupKit to support coarse-grained workspace awareness. Learners have a standard scroll-bar which indicates their relative position in the workspace plus they have an additional labelled scroll-bar for each of the other group participants. Each additional scroll-bar indicates where one of the other participants is working in the shared workspace. GroupKit also has a global display widget for supporting coarse-grained workspace awareness. The global display widget replaces a standard scroll-bar with an outline of the shared workspace. Each participant's view on the shared workspace is represented with a labelled rectangle on the outline. Participants can change the size or position of their view of the shared workspace by manipulating their rectangle on the outline. Participants can see at a glance where other participants are working in the shared workspace.

In reality, many collaborative activities require people to alternate between working alone and working with others. GroupKit supports this with a history mechanism that allows users to review what other people have done while they were away from the shared workspace.

Greenberg [42] has prototyped the use of fisheye views to support workspace awareness. Fisheye views are usually used in information visualization as a means for examining the details of some portions of the information space while maintaining the context of the rest of the information space (e.g. [22]). This is accomplished by emphasizing the focus areas through magnification while compressing the rest of the information space. Greenberg allocates one focus to each participant in the shared workspace. Each participant can then see some details of what others are doing while maintaining the global context. Participants can control the magnification function used for each of the foci so as to obtain more or less details of other people's activities as needed.

Web-based learning environments that support synchronous collaboration should support workspace awareness. Evaluators can look for widgets as described above or for others which specifically support fine-grained and coarse-grained shared workspace awareness.

Table 4.5: Personal workspace tools and features

• Concept maps	• Tagging
• Personal annotations	• Off-line participation
• Personal semantic associations	• Document upload
• Notification of changes	• Utility tools
• Exploration management	

4.1.5 Personal Workspace

This dimension of the comparison framework is used to compare the tools and features that make up the learners' personal workspace. These are specific tools and features to be used by individual learners to help them learn, and to manage their time and efforts in the learning environment. They are summarized in Table 4.5 and are discussed in the paragraphs that follow. The collaborative tools that were identified in Section 4.1.3 will not be repeated here.

Concept maps are graph-like structures that a learner can build to represent the relationships between their ideas. Learners can update their maps as their understanding of concepts changes. Niguma [76] discusses a Web-based implementation of concept mapping that supports learning under a constructivist learning model. Learning environments may provide a concept mapping tool or they may provide tools to support other forms of knowledge representation. Niguma suggests that concept mapping tools could be extended to become synchronous collaboration tools in which multiple people participate in the construction of shared concept maps.

Various forms of personal annotations could also be useful for learners. A mechanism that allowed highlighting of text could be a useful tool for information analysis. Learners might also benefit from the ability to add their own in-place comments on online learning materials.

The Web is a hypermedia environment where the links between documents are usually defined by the document authors. Edelson et al. [36] have suggested that learners would benefit from being able to create their own semantic associations between documents. Hewitt and Scardamalia [50], and Honebein [54] also identified the ability for learners to create annotated links between documents as a desirable feature of learning environments.

One difficulty with keeping up with online materials is knowing when and to what extent they have changed. A facility that notified learners when resources that concerned them changed would be a useful tool for learners. If the tool could also inform learners of the nature and the extent of the changes it would be even more valuable.

Another potential difficulty for learners with keeping up with online materials is remembering which materials they have already seen. Since hypermedia systems are often non-linear, as opposed to the linear nature of a book for example, tools are required to manage the exploration of resources. These tools should show users where they have already been and also give them a map of the available resources to help guide their exploration.

Since the resource space of a Web-based learning environment can get quite complex, learners may experience difficulty finding resources that they had previously used and found interesting. Learners would therefore benefit from a mechanism that allowed them to tag certain resources that they have seen for future reference. This feature is analogous to folding the corners of pages of books or using bookmarks. This facility could be enhanced by allowing users to annotate the tags they place on resources.

Learning environments may provide a facility for learners to continue to work in the learning environment while disconnected from the network. This may help learners to save on network connection charges and will allow them to work where no network connections are available. Whether or not Web-based learning environments provide a mechanism for working off-line, it is desirable for learning environments to provide a controlled and flexible mechanism for learners to upload documents they have produced or obtained outside of the learning environment.

There are a number of utility tools from which learners could also benefit. Standard utilities like an *electronic note pad*, a calculator and a dictionary could be packaged with the learning environment. Calendaring and scheduling functions may also be provided by some learning environments.

Table 4.6: Aspects of course authoring

• Course structuring
• Creating tests and quizzes
• Creating activities
• Interfacing with simulation software

4.1.6 Course Authoring

Instructors or instructional designers will need to design and develop courses in a format that can be used in the Web-based learning environment their organizations adopt. Learning environments may provide course authoring tools to support some course design activities or they may simply deliver courses in a generic format like the Hypertext Markup Language (HTML).

Although instructional designers and instructors are experts at authoring courses, they are not necessarily trained to design complex human-computer multimedia user interfaces for Web-based learning environments. Learning environments that provide good tools for authoring courses will generally have a more consistent look and feel within and across courses. They also provide more consistent interactions and navigation for the users. The potential disadvantage is that authoring tools may limit a course author's creativity by forcing them to use tools that are not as flexible and powerful as lower level coding.

Course authoring tools can either be complete course authoring packages or they can support a subset of the course authoring activities. Learning environments may provide tools that help instructors structure their courses in units or modules and then sequence them for delivery to the learners. Some learning environments may provide tools for creating tests and quizzes for assessment. Course authoring tools might also include tools to create learning activities like role plays. Interfacing with simulation software is another feature that flexible course authoring tools may provide.

Organizations that have not yet invested much in the development of online courses may favour learning environments that provide good tools to assist in course authoring. Organizations that have many courses already developed for a generic format like HTML may favour learning environments that can deliver these courses as they are or that can automatically translate them into the learning environment's format.

4.1.7 Assessment

A dimension that seems to be very important to organizations considering the adoption of Web-based learning environments is the support for assessment. The assessment dimension contains features and attributes that are used to evaluate learners' performance in learning activities. Assessment includes not only the traditional measurement of learning outcomes such as knowledge acquisition and retention rates but also the evaluation of students' learning processes.

Traditional assessment

Tests and quizzes are a common way of evaluating learners' understanding of concepts. Learning environments may therefore provide a tool for learners to take tests and quizzes electronically. Types of testing questions that might be supported include multiple choice, fill-in the blank, matching, short answer, and essay. More innovative testing questions might be based on simulations. Learning environments may provide automated grading of certain types of questions.

Testing systems need to be very flexible to be used by instructors. There needs to be a great deal of flexibility in the types of questions that can be asked and how the questions are displayed to learners. Systems that provide automated grading must also be very flexible in the grading algorithms that they use. Many instructors are likely to require fine grained control over the entire testing system. There is a lot of variety in the presentation and in the grading schemes even for seemingly simple multiple choice question tests. Good testing systems will account for these and will make it easy for instructors to deliver the tests they want.

Another important aspect of assessment is feedback. Instructors should have a mechanism for providing feedback to learners on their performance during evaluations.

Learning environments may provide instructors with a tool for recording, managing and distributing learners' grades. Minimally, these tools allow instructors to manually

Table 4.7: Assessment features

Traditional	Multiple choice Fill-in the blank Matching Short answer Essay Simulations Automated grading Flexibility Feedback Grade book
Nontraditional	Peer evaluation Portfolio assessment Designing instruction

enter the learners' grades for each activity. More sophisticated tools might allow instructors to chart grade distributions and selectively release grade information to the learners.

One problem that computer based learning environments have not yet been able to solve reliably is that of authentication. There is no feasible electronic mechanism to guarantee that the person taking a test remotely is indeed the learner that they claim to be or that they are working alone. The best solutions to this problem are either to trust the learners or to have testing sessions witnessed by proctors.

Nontraditional assessment

Web-based learning environments that embrace constructivist learning models and those proposed by the sociocultural theories of learning should provide alternative assessment mechanisms which are more consistent with these nontraditional learning models. Attributes of nontraditional assessment methodologies include:

- involving the learners in their own evaluation;
- involving learners in the evaluation of their peers;

- accounting for individual learning goals; and
- evaluating the learning process as well as the learning products.

Peer evaluation is an example of such an alternative assessment methodology. Learning environments that support peer evaluation will provide collaborative tools that make it possible for learners to submit their learning artifacts or a representation of their understanding of some concepts to a panel of reviewers consisting of peers and possibly instructors. The learning environment would have tools such as those discussed in Section 4.1.3 that allow the panel to discuss the merits of the submission amongst themselves and with the learner.

A second example of an alternative assessment methodology is portfolio assessment [71]. Learners are responsible for managing their own portfolio which contains artifacts showcasing their products and process of learning over time. A learning environment could therefore provide tools for managing and reviewing portfolios. A final example of alternative assessment is where learners are given the task of designing instruction for other learners [71]. To support this type of activity, a learning environment could allow protected access to instructor tools in the learning environment and enable the creation of non-production instructional modules.

4.1.8 Tracking

Computer applications can be programmed to leave an electronic trail of all their interactions. These logs of interactions can be a source of all kinds of interesting information. The challenge is to provide tracking of the right data and in the right quantity so that the data can be efficiently analyzed and so that the desired information can be extracted from it.

Data for human consumption

Web-based learning environments could provide sophisticated tracking of individual learners' interactions with the system. This data could help instructors detect trends in learners' behaviours. The data could help to measure learners' participation in

Table 4.8: Tracking features

Human use	Individual interactions Group interactions Data mining Data analysis Reporting
System use	Automatic tuning Context sensitive UI Adaptive UI Dynamic guidance capabilities

activities. It could also be of assistance in evaluating learners' performance of activities and could give insight into their learning processes. Having access to good data might also help instructors detect learners' problems earlier and therefore increase the chances of correcting them. Similar tracking could also be done for group interactions. Data about individual and group interactions with the system could help instructors improve their courses and their teaching. The data could also be used by researchers interested in the dynamics of online learning.

System administrators will also be interested in accessing data about the learning environment. Data about system performance and security violations are examples of data in which system administrators would be interested.

Raw data is not typically very useful. It order to extract information, the data need to be analyzed. Learning environments could provide data mining, data analysis and reporting tools for use by instructors, system administrators and researchers. These tools would help their users validly interpret the data that the learning environment's tracking system captures.

Data for system consumption

Web-based learning environments may also track data for their own use. That is, the learning environment may be programmed to behave in different ways based on the data it captures. For example, a learning environment might be able to do dynamic load balancing based on the usage patterns it detects. User interfaces which "remember" what users have done in the past to improve navigation also depend on tracking. Learning environments which have some artificial intelligence built into them can use data they capture to perform more complex operations.

Web-based learning environments with adaptive user interfaces require complex system logic, and good data capture and analysis capabilities. Based on its artificial intelligence and on the data it captures, the learning environment can modify the user interface it presents to each user to better meet their needs.

Sophisticated tracking could also enable dynamic guidance capabilities [41] in intelligent Web-based learning environments. Dynamic guidance capabilities are an automated system that monitors learner or instructor activities and detects behaviour trends. Based on these trends, the dynamic guidance system can determine better practices. The system can then dynamically intervene and alert learners or instructors of the problems it has noticed and how they can improve their learning or teaching.

For example, a dynamic guidance system could help learners improve their learning by taking advantage of the spacing effect. The spacing effect "refers to the finding that for a given amount of study time, spaced presentations yield substantially better learning than do massed presentations" [30]. If the system could detect that learners were repeatedly reviewing the same materials, it could inform learners of the spacing effect and suggest that their study time could be better spent by spacing the review of particular materials over longer periods of time.

4.1.9 Ethical and Social Issues

Very few of the readings that were consulted in preparing this thesis even acknowledged the fact that there were ethical and social considerations in developing and deploying Web-based learning environments. Hiltz [52] alluded to a few ethical issues and Hill [51] identified ethical issues as one of five areas of consideration in creating distance learning environment via the Web. These were the exceptions rather than the rule. There are important ethical and social questions that need to be asked when considering the adoption of a Web-based learning environment. There are also

Table 4.9: Ethical and social issues

• Depersonalization of learning	• Fair accessibility
• Privacy	• Intellectual property
• Monitoring	

a number of policy decisions with ethical and social consequences that need to be made by organizations that are deploying a Web-based learning environment.

Generally, many people have become excited about the possibilities of technology but they have not given enough consideration to its ethical and social ramifications. The same is true for technology applied to learning environments and in particular for Web-based learning environments. Ethical and social issues are very important in Web-based learning environments because they address considerations such as the depersonalization of learning, privacy, monitoring, fair accessibility and intellectual property. Those deciding on the adoption of a Web-based learning environment for their organizations can use the ethics dimension of the comparison framework to evaluate how much consideration the developers of the applications gave to ethical issues. They can also refer to this dimension to identify some of the issues that they themselves need to consider before the deployment of a Web-based learning environment.

One general way in which learning environment vendors can show that they are sensitive to ethical issues is to defer the final decision on ethical matters to their customers and provide enough flexibility within their products to accommodate their customers' decisions. In this way, organizations can decide on what are the best ethical practices in their circumstances and have the learning environment support these practices.

Depersonalization of learning

Some people fear the depersonalization of learning in technologically based learning environments. The fact that human contact may be lost or only possible through a mediator (the computer) is a legitimate source of concern. The social development of learners, for example, is an important and beneficial consequence of face-to-face learning that might be compromised in learners who primarily interact with machines rather than with peers and instructors. Learners who are isolated by distance may also have difficulty developing relationships with their instructors and peers. Organizations considering the adoption of a Web-based learning environment should favour vendors who have made special efforts to compensate for the lack of direct human contact in computer-mediated learning. Web-based learning environments that provide good collaboration tools and good user interfaces can at least partially address the concern over the depersonalization of learning by enabling instructors and learners to develop a sense of community even though they are separated by distance.

Privacy

Privacy issues are also very important in a digital domain like a Web-based learning environment. Organizations need to make very careful decisions about who is allowed to have access to what information within the learning environment. All those who participate in the environment should know what type of privacy they can expect. Learning environments can support organizations' policy decisions on privacy by providing a flexible and robust security model. Good learning environments will also help protect adopting organizations from criminal prosecution by making it difficult for them to unknowingly violate privacy laws [101].

Monitoring

Related to privacy issues are monitoring issues. In the preceding section about tracking we showed that computer applications can be very good at tracking the activities of their users and that it is desirable to do this for arguably good reasons. These reasons included helping instructors to better understand their students so that they could better teach and guide them. However, some people might consider it an invasion of privacy for others to be able to track their every move through a learning environment. The parallel in a face-to-face environment is having someone constantly looking over your shoulder. This would likely bother many people even if it was supposedly for their own good. If it was possible for monitoring information to fall into the hands of untrusted individuals, then even more people would object to being monitored. Organizations should make a conscious decision about how much monitoring they will use and they should ensure that their users understand the degree to which they are being monitored. Learning environments can support the implementation of policy decisions on monitoring by being flexible in the data they capture and by providing secured control to any monitoring data.

Fair accessibility

Fair accessibility to learning is another important ethical question that many organizations considering the adoption of a Web-based learning environment will have to consider carefully. Web-based learning environments that make use of the most modern technologies sometimes provide attractive functionality but also tend to require more computing resources. If these computing resources are not equitably available to all learners, then some learners may be unfairly disadvantaged. Organizations will have to strike a balance between providing the best learning environment possible for their learners without creating unacceptable disadvantages.

Intellectual property

Questions about intellectual property should be settled before deploying a Web-based learning environment. Clear policy decisions need to be made about who has rights to materials that are manipulated or created within the learning environment. These decisions should cover materials created by learners and by instructors. Learners' and instructors' works may contain those from other sources which must be acknowledged and for which royalties may need to be paid. Care must be taken so that the learning environment only distributes materials for which it has proper distributor rights. The policy decisions that are made should be supported by the organization's management, the instructional designers, the instructors and the learners, and they must respect federal and international laws on intellectual property, copyright and fair use. The difficulty with respecting laws is that they are often unclear with respect to the Web as a medium. Zobel [120] makes the point by stating: "Neither the legal system nor World Wide Web developers nor nominal owners of original works are on solid footing when determining copyright infringement. Therefore, all World Wide Web developers should familiarize themselves with federal law relating to copyright, intellectual property, and fair use."

The CulTech Research Centre [53] has compiled a set of requirements for an ideal intellectual property management system. It has also designed and is testing IVY, an intellectual property management system that satisfies most of the requirements of an ideal system. Organizations should favour learning environments that provide facilities that help track materials and that make it difficult for these to be accessed without proper permission.

Importance of ethical and social questions

It may be tempting for organizations to put off decisions on these matters. They are strongly encouraged to resist this temptation as careful planning and prevention lessen the risks of conflicts further down the road. Organizations that already have policies in place for face-to-face learning should not assume that these policies are adequate for distributed learning via the Web. Some questions will require special consideration for a Web-based learning environment, and some questions will be unique to Web-based learning environments.

To illustrate how difficult some of these questions can be, consider the following seemingly simple one. What information should a learning environment make available to university students about their classmates? There is a wide range of possible answers. The answer that supports privacy the most is: no information at all. This answer, however, is not conducive to building a sense of community or enabling collaboration. If the learning environment made other people's names available then it would be much easier to meet other people. Some people would even argue that they have a right to know who their classmates are because these classmates have access to their contributions to shared workspaces. Others would argue that they should have the right to attend classes anonymously. Although people are seen by others in traditional lecture halls, they do not usually have to make their names known to their classmates. Some of the concerns might be addressed by learning environments that allow learners to use pen names or aliases. This, however, raises other concerns about people impersonating others or people using inappropriate pen names, not to mention the potential problems with people behaving less responsibly in the learning environment because of perceived anonymity.

Organizations will quickly realize that there are often conflicting views on ethical questions and that there are often good reasons to support each of the positions. They should carefully decide on as many of these matters as possible before adopting and deploying a Web-based learning environment.

4.1.10 System Requirements

Web-based learning environments may have quite a wide range of system requirements and these definitely need to be factored into the decision on which learning environment to adopt and deploy. Organizations will generally favour learning environments whose system requirements are similar to their currently installed base of hardware and software. This is important for minimizing costs and maximizing accessibility.

The system requirements can be logically divided into three sections: the client, the server, and the network.

Client

The client is the hardware and software that the end-user will use to participate in the distributed learning environment. The primary end users are usually the learners. Other end-users may include the instructors and the administrators. Some systems may require different clients for different classes of end-users. The minimum client requirement for all Web-based learning environments is a networked machine with a Web browser. Special care needs to be paid to the system requirements for the client because in many cases organizations will not have complete control over these machines. Although an organization might provide access to some client stations, it

Figure 4.1: Client server

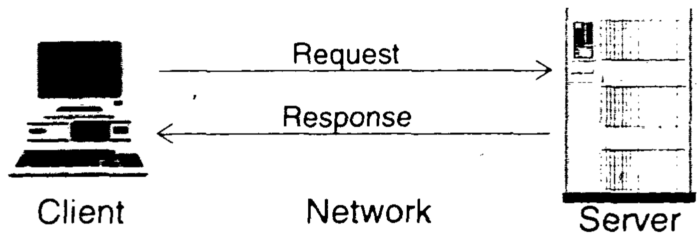

is likely that many end-users will employ other machines such as home computers to participate in the learning environment.

There are currently two leading Web browsers on the market: Microsoft's Internet Explorer [73] and Netscape's Communicator [75]. These browsers have a lot of common functionality but their behaviour does sometimes differ, particularly in the implementation of more advanced features. If an organization wants its members to be able to use either browser, then it is important to select a learning environment that was designed for and tested on both browsers.

The other issues to consider with respect to browsers is the version number. The latest versions of browsers have more functionality and therefore offer more possibilities to the learning environment developers. Learning environments are typically designed for a minimum browser version. The tradeoffs are that end-users may not yet have installed the newest browsers and, more seriously, their machines may not be powerful enough to run them. Organizations will likely choose the most advanced learning environment that runs on as many of their members' machines as possible.

Some learning environments may require installing additional software on client machines, including browser plugins. Organizations must assess whether this further limits how many of their members will be able to run the application on their machines and whether the limitations are reasonable. It is also important to note that installing additional software on client machines will pose an extra burden on end-users and that they will need to be supported in this activity.

A platform independent learning environment will be important to an organization whose members need to use many different types of hardware and operating systems to participate in the learning environment. Note that although the installation of specialized software or hardware may enhance the functionality of a learning environment, it may also significantly limit the number of platforms that can successfully run the learning environment. Consider for example Microsoft's NetMeeting [74]. A Web-based learning environment that makes use of NetMeeting can provide impressive collaborative features. The tradeoff is that a learning environment that makes use of it is limited to running only on Windows platforms.

Web-based learning environments may also have different hardware requirements for the clients. The minimum hardware requirements are a machine capable of running the minimum software, and the network hardware, such as a modem, to connect to the network at the required minimum speed. The hardware requirements typically include the minimum processor, minimum disk speed and minimum memory. Unless an organization controls all the client machines, it is generally preferable to have a platform independent learning environment. This means that the learning environment is not limited to a particular operating system. It could run, for example, on a Unix machine, as well as on a Windows machine, or a Macintosh. With the advent of the Network Computer and set-top boxes, matters will probably be further complicated. Unless there is mass acceptance of a standardized platform, it will be important that learning environments be able to run on these new machines as well as on desktop and portable computers.

Some learning environments may require additional client hardware, much of which is associated with multimedia interactions. A printer might be required for producing hardcopies of online materials. A microphone and speakers might be required to enable audio interactions. A video camera might be required for video conferencing. An enhanced video card might be needed for the fluid navigation of three dimensional virtual worlds. More specialized hardware might also be needed to enable other types of interactions like simulations.

Table 4.10: System requirement attributes

Client	Supported browser type Recommended browser version Extra software required Supported platforms Machine specifications Recommended network connection Multimedia hardware required Extra hardware required
Server	Machine specification Recommended network connection Extra software required
Network	Bandwidth

Server

The learning environment server is the central machine, or cluster of machines, with which the clients interact. Learning environment vendors will usually publish minimum hardware and software requirements for the server machines. The hardware requirements specify the type and speed of processors that are required as well as the needed disk space and random access memory.

Learning environments typically operate on a limited number of server platforms. Since maintaining server machines can be challenging, organizations will typically favour learning environments that operate on server machines that they know how to manage.

A key issue related to server requirements is scalability. The selected learning environment must be able to support the anticipated number of initial users and should scale well if that number increases. It is important for a server to service client requests in a timely manner.

The software required for learning environments is sometimes self-contained in the learning environment application. However, other software is often needed. The other software might include databases, compilers and Web servers. These add complexity to the system, may create bottlenecks and should be factored into the cost of installing and maintaining a Web-based learning environment.

Network

The network is what links the clients and the server together. For Web-based learning environments we usually assume that the network is the Internet but it could also be an organization's intranet. Learning environments which use World Wide Web communication protocols are not typically tied to a particular network type so we will not address the differences in the management of different networks.

The main issue related to networks that is important to this discussion is bandwidth. A network's bandwidth is a measure of the amount of data it can transmit per unit of time. The higher the bandwidth the more data can be pumped over the network. Both the clients and the server need to be connected to the network at a minimum bandwidth. Higher bandwidths will allow richer interactions but will be more expensive. Organizations will choose learning environments that provide the best service for the amount of bandwidth in which the organization is willing to invest. Note that bandwidth requirements necessitate a hardware investment for all end-user client machines.

4.1.11 Robustness and Security

It is important for all software to be robust. It is particularly important for software that multiple people depend on, like a Web-based learning environment, as opposed to a single user word processor, for example. Organizations should only consider adopting well tested Web-based learning environments from vendors who are committed to fixing the more severe problems as soon they occur. Robust systems are not only resistant to problems but they also have mechanisms in place to make graceful recoveries when problems do occur. These systems are said to be fault tolerant. Minimally, Web-based learning environments should enable frequent backing up of all data stored in them and have a mechanism to restore the data from backup.

Every sensible Web-based learning environment will include some form of security.

Table 4.11: Robustness and security features

• Fault tolerant	• Encryption
• Backup functionality	• Digital certificates
• Password protection	• Digital signatures
• Multiple permission levels	• Smart cards

Generally, security mechanisms protect data from unauthorized access and guarantee the identity of those interacting with the system. There are many different ways to implement these security goals with a tradeoff between speed and robustness.

A simple way to handle security in a Web-based learning environment is to supply a unique user name and a secret password to every legitimate participant in the learning environment. The learning environment can then challenge every access to a portion of the environment that is protected. Participants are required to identify themselves to the system with their user names and passwords. The system can then verify whether individual participants have the required permission to use the resources they are attempting to access.

Slightly more sophisticated learning environments will allow control over different levels of access to resources within the learning environment. An access control list can include user names (or groups of user names) and the level of access they are permitted for a particular resource. For example, Lotus Notes [66] provides seven levels of access to its databases which specify whether people can read, create, modify and delete various types of documents. They are shown in Table 4.12.

Encryption can also be used to make a learning environment more secure. Data can be encrypted for storage in the learning environment and it can be encrypted in transit over the network. More robust encryption algorithms with larger decoding keys can be used to obtain the desired security level, at the expense of computation and transmission speed.

Participants in the learning environment can also be issued digital certificates. These certificates provide better security then that provided by user name and password pairs. They can be used as an encryption key which only allows the certificate

Table 4.12: Lotus Notes levels of access

No access	No permissions at all.
Depositor	Permission to submit documents.
Reader	Permission to read documents.
Author	Reader's permissions plus the permission to create documents and edit or delete own documents.
Editor	Author's permissions plus the permission to edit or delete other people's documents.
Designer	Editor's permissions plus the permission to make changes to the database itself.
Manager	Designer's permissions plus the permission to delete the database.

owner to view transmitted information.

Digital signature facilities may also be provided by certain learning environments. Participants can sign documents with a digital signature that proves to other participants with whom documents originated. Digital signatures can also be used for nonrepudiation.

Smart cards are credit card size devices that have memory and some processing power. Each participant in a learning environment could be issued a smart card which would store their identification information as well as all their preferences in the learning environment. Upon inserting the smart card in common access terminals or in kiosks, the Web-based learning environment would recognize the participant and give them the appropriate access and view of the learning environment.

Organization considering the adoption of a Web-based learning environment should favour robust applications. They should also favour learning environments that provide at least the minimum amount of security needed by the organization. In Section 4.1.9 it was stated that organizations must make many decisions on ethical and

Table 4.13: Integration issues

• Internal integration
• Environmental integration

social questions, some of which deal with questions of privacy and access to information. Organizations should therefore favour Web-based learning environments whose security model is flexible enough to accommodate the privacy decisions made by the organization and which can protect intellectual property.

4.1.12 Integration

There are two types of integration that can be considered when comparing Web-based learning environment applications. Firstly, one can consider how well the application's sub-systems are integrated. We will call this internal integration. Secondly, one can consider how well the learning environment would integrate with the software and the human processes already in place in the target organization. We will refer to this type of integration as environmental integration.

Internal integration

Exceptionally well designed applications are built on top of a common architecture which helps ensure seemless internal integration of the application's components. In practice, software applications often do not have such a consistent architecture. Applications often consist of individual tools and components which are pulled together later in the development cycle to form the complete application. Applications developed in this fashion sometimes suffer from poor integration between their components. To evaluate how well an application's components are integrated, one can inspect the user interface for inconsistencies and examine how easily data flow between the different logical parts of the application. Applications which exhibit poor internal integration can be clumsy and frustrating to use. Siegel and Kirkley [98] identified "Access to an *integrated* package of navigational, productivity, communication, collaboration, and knowledge/wisdom creation tools" as one of the desirable features of a digital learning environment.

Environmental integration

The second type of integration that should be considered, environmental integration, is how well the candidate Web-based learning environments would fit with already installed software and hardware, and how well it would accommodate and support established human processes. An organization might find it desirable, for instance, that the learning environment software be able to easily import course registration information from a legacy system. It might also be desirable that an authorized user of the learning environment application, such as an instructor, be able to submit students' final grades to the organization's existing performance tracking system. This type of integration is usually only possible if an application has been purposefully designed to support communication with specific existing applications, or if the application has been designed with an open architecture that provides adequate application programmer interfaces (APIs) that can be exploited by programmers to "glue" together multiple applications.

Some vendors may sell their Web-based learning environment as an integrated part of a larger software package that provides additional valuable functionality. The Web-based functionality that could be provided is almost limitless but examples include registration for courses, payment of fees, ordering of merchandise such as books, and counselling services. Since this additional functionality is somewhat peripheral to learning environments, it is not treated in any more detail here or elsewhere in the framework. However, organizations that are considering the adoption of a Web-based learning environment along with software providing additional functionality should ensure that the products are interoperable and well integrated.

Those responsible for choosing an application for their organization should evaluate how compatible their organization's philosophy and practices are with those of the candidate applications. Organizations will favour applications whose philosophy

Table 4.14: Types of customization

• End-user customization
• System administrator customization
• Vendor customization

about learning is similar to their own. Furthermore, unless an organization has chosen to engage in a business process reengineering exercise, they will favour Web-based learning environments that most closely match their practices.

4.1.13 Customizability

Every computer application offers a certain degree of customizability and Web-based learning environments are no different. There are three distinct levels of customization that an application might support. These are, in order of increasing complexity, end-user customizations, system administrator customizations and vendor customizations. Systems that offer the most customizability tend to be the most flexible but they may also be more complex.

End-user customization

End-user customizations are often presented to the user as finite sets of preferences from which they can select. These preferences are a way for users to exercise some control over the environment in which they learn and work. There is a wide array of variables that end-users might be able to customize. These include user interface preferences and security options.

System administrator customization

System administrators who install and manage a Web-based learning environments may also be able to customize the environment's behaviour and appearance. A well-designed learning environment may have a collection of hooks and configurable variables that system administrators can tweak to customize the learning environment for their particular organizations.

Vendor customization

Some Web-based learning environment vendors may not sell their learning environment as a single, complete product. They may instead have a set of assets which can be combined in various ways to produce customized learning environments for their customers. Vendors may also offer to program custom components for customers and to package these with their other assets to form a complete and unique learning environment.

Those who are considering adopting a learning environment will want to ensure that a learning environment meets their needs as closely as possible. An important issue to keep in mind is whether the business rules programmed into the learning environment agree with their business rules or whether the learning environment can be customized to do so. Examples of business rules that might change from one organization to another and that should be customizable in a learning environment include questions of who has access to students' personal information and how long grades are retained in the system.

Organizations investing substantial sums of money in a large deployment of a learning environment will likely prefer dealing with a vendor who is willing and able to create and support a customized Web-based learning environment. More modest learning environment deployments may be satisfied with off-the-shelf products that offer limited customizability by the system administrator.

4.1.14 Support

The service that may be bundled with any product can greatly increase its value and the satisfaction that customers experience with the product. In Section 1.2 we identified the stress associated with managing new technologies as one of the disadvantages with using technology for learning. Good support makes managing software easier and less stressful. Organizations considering the adoption of learning environments should therefore pay particular attention to the support that vendors offer customers

Table 4.15: Types of support

• Online and printed documentation
• Toll-free call centres
• Consulting services
• Commitment to maintenance

of their products. Support is needed both for administrative tasks and for end-user tasks.

A simple but invaluable form of support is proper documentation. A good learning environment should include either online or hardcopy documentation for both the end-user tasks and the administrative tasks. The end-user documentation should include "getting started" information as well as "how do I?" information on performing common end-user tasks. The administrator documentation should include detailed information on installing, configuring, tuning and troubleshooting the Web-based learning environment.

Toll-free call centres staffed with knowledgeable customer service representatives and technicians are also desirable. These give administrators and end-users easy access to help for resolving their specific difficulties in using the learning environment. This type of personal service also lessens the helplessness that some users feel when they encounter problems in a computer mediated environment. Help by electronic mail is an alternative to this but it lacks the personal touch offered by human contact.

Some learning environment vendors may offer consulting services with their product. These knowledgeable consultants can come to the customer's site to help with the installation and the management of the learning environment. Some consultants might build relationships with customers and continue to be an ongoing resource for the management of the application.

Organizations should favour learning environment vendors who are at very least committed to maintaining their application. This most fundamental level of support includes fixing bugs and receiving customer suggestions for improvements to the application.

Table 4.16: Types of training

• How to use the learning environment
• How to teach and learn effectively online

4.1.15 Training

The training offered with a learning environment is another dimension along which competing solutions can be compared. While it is often desirable that users be able to use software products successfully without training, this is not always achievable. When users will be using a sophisticated application for an extended period of time it is reasonable to assume that some training will be needed. On the other hand, if people use an application infrequently or for only a short period of time it may be more desirable to have a simpler application which most people can use without training. Organizations may favour vendors who offer training with their learning environments if the organizations offer longer courses. Even for organizations offering only short courses, training may be necessary for administrators and instructors.

There are at least two types of training that could be useful for organizations adopting a Web-based learning environment. One is how to use the learning environment application. This is valuable training as it will help demystify the application for new users and teach them the application's capabilities. The more challenging type of training that may be needed is on how to teach and learn effectively in a technology supported distributed learning environment. Instructors who are not experienced in teaching in these learning environments will have to adjust their teaching methods and would benefit from training. Similarly, learners who are used to other forms of instruction might also have to modify their learning methods and could probably benefit from some training on how to best do this.

4.1.16 Business

Web-based learning environment vendors make business decisions that affect how they position their product in the market for distributed learning environments. Vendors

Table 4.17: Business issues

• Target market
• Total costs
• Hosting the learning environment

who try to capture the entire market with a single general offering risk not doing well in any market segment since other vendors who specialize for different market segments are more likely to meet their potential customers' specific needs. One positioning decision that vendors usually make is which academic level to target. The academic levels can be grouped into primary education, secondary education, tertiary education, and professional training. Another decision that vendors should make is how many simultaneous courses their learning environment will support. Similarly, they should decide how many simultaneous participants their learning environment will support.

Organizations considering the adoption of a distributed learning environment via the Web should favour learning environments whose target market closely matches the market segment that they fit into. They should also ensure that the candidate learning environment can support the required number of simultaneous courses and learners, and that it can scale to accommodate anticipated growth.

There may very well be a large variety in the costs associated with competing learning environments. Vendors will often offer a few different options for licensing agreements and may be willing to customize the licensing agreements for their customers. Licensing agreements usually stipulate the cost of the product as a function of the number of servers required, the number of clients and the number of courses. Additional charges may be associated with installation services, ongoing technical support, vendor customizations, consulting, and training. Hardware purchases may also be required in order to accommodate the system requirements.

An additional service offered by some vendors that might be of interest to certain organizations is hosting the Web-based learning environment. This means that the vendor absorbs the costs for the server side hardware and for the technical administration of the learning environment. Some vendors may also be able to offer server or

Table 4.18: Future direction issues

• Planned upgrade
• Vendor's reputation
• Commitment to user-centered design
• Vendor's size

client hardware for lease at attractive prices.

Organizations must determine their budget for a Web-based learning environment and then find the application that best meets their needs while respecting their budget. Organizations are encouraged to look beyond the licensing costs of a Web-based learning environment to try to ascertain the real cost of adopting a Web-based learning environment, realizing that the costs associated with using a Web-based learning environment might outweigh the costs of acquiring it. Organizations should factor into their budgets that a quality distributed learning environment might increase revenue through increased enrollment. Furthermore, some learning environments may also generate additional revenue by carrying advertising. This, however, raises some ethical questions.

4.1.17 Future Direction

It is often difficult to know what future direction a vendor plans to take with its application. Organizations seriously considering the adoption of a Web-based learning environment should, however, make an effort to understand a vendor's plans for its application. A good vendor will be forthright with its intentions.

Minimally, organizations considering the adoption of a learning environment should determine whether a vendor has any future plans for the application. Many planned software projects never are released and many of the released software applications never get upgraded. Organizations should favour learning environment vendors who are committed to supporting their product and releasing at least one upgraded version of the application.

A vendor's reputation is worth considering when assessing its future plans. Vendors with a good reputation for supporting their products and providing reliable upgrade paths are probably more likely to do the same for the current applications they sell. Similarly, one is more likely to receive good customer service from a vendor who has a good track record for its customer service.

Vendors who are committed to user-centered design and who are committed to incorporating an organization's suggestions for improvements in future releases should be favoured as their products are more likely to evolve in the direction that the customer desires. Although smaller or newer vendors may not have an established reputation, they are sometimes more flexible on this point than larger vendors. If one customer represents a larger proportion of market penetration for a vendor then that customer may have a louder voice in the evolution of the product.

Table 4.19 summarizes the dimensions of the comparison framework. The remainder of this chapter proposes ways of using the comparison framework.

4.2 Using the Comparison Framework

This comparison framework serves two main audiences. These are Web-based learning environment developers, and organizations considering the adoption of a Web-based learning environment. While independent evaluators such as trade magazines might find value in referring to the framework's comparison dimensions to understand the range of issues that need to be considered when comparing learning environments, the framework is not specifically designed to support the type of comparisons they usually make. One group of evaluators cannot declare that a learning environment is superior to all others because a fair comparison must be situated in the context of the potential adopter's environment. Different organizations will have different goals and priorities, so blanket statements about the merits of competing solutions are not useful and they will not be found in this work.

Table 4.19: The seventeen dimensions of the comparison framework

1.	**Facilitating**	• directly facilitating learning • support for facilitating learning	• specialized scaffolding
2.	**Cognitive apprenticeship**	• interaction with community of practice • mechanism for observing others	• authentic activities
3.	**Collaboration**	• asynchronous email, mailing lists, discussion groups, shared workspace, group review, decision support system, shared annotations, shared semantic associations, general awareness • synchronous chat, internet phone, video conferencing, navigation control, group authoring, decision support system, shared knowledge representation, general awareness, workspace awareness	
4.	**User interface**	• user interface principles • navigation	• workspace awareness • metaphors
5.	**Personal workspace**	• concept maps • personal annotations • personal semantic associations • notification of changes • exploration management	• tagging • off-line participation • document upload • utility tools

Table 4.19 (cont.): The seventeen dimensions of the comparison framework

6.	**Course authoring** • course structuring • creating tests • interface with simulation software • creating activities
7.	**Assessment** • traditional multiple choice, fill-in the blank, matching, short answer, essay, simulations, automated grading, flexibility, feedback, grade book • nontraditional peer evaluation, portfolio assessment, designing instruction
8.	**Tracking** • human use individual and group interactions, data mining, data analysis, reporting • system use automatic tuning, context sensitive UI, adaptive UI, dynamic guidance
9.	**Ethical and social issues** • depersonalization of learning • fair accessibility • privacy • intellectual property • monitoring
10.	**System requirement** • client supported browser type, recommended browser version, extra software required, supported platforms, machine specifications,recommended network connection, multimedia hardware required, extra hardware required • server machine specification, recommended network connection, extra software required • network bandwidth

Table 4.19 (cont.): The seventeen dimensions of the comparison framework

11.	**Robustness and security**	
	• fault tolerant • backup functionality • password protection • multiple permission levels	• encryption • digital certificates • digital signatures • smart cards
12.	**Integration**	
	• internal integration	• environmental integration
13.	**Customizability**	
	• end-user customization • system administrator customization	• vendor customization
14.	**Support**	
	• documentation • call centres	• consulting services • commitment to maintenance
15.	**Training**	
	• how to use the learning environment	• how to teach and learn online
16.	**Business**	
	• target market • hosting the learning environment	• total costs
17.	**Future direction**	
	• planned upgrade • vendor's reputation	• user-centered design • size of vendor

4.2.1 Developers

Web-based applications are still quite a new concept. Most of the people who have used them so far are early-adopters who typically have considerable technical interest and knowledge, and who have patience for immature products. As Web-based applications become more common, people will expect more of them. Learners and instructors will not be willing to devote too much time or effort to the tool that they use to learn and to teach. They will expect the tools to become transparent.

Since many software developers come from a technical background, they tend to focus on the technology that their products utilize rather than on all of their users' needs. This sometimes leads to feature-rich yet narrow-minded product offerings that neglect important aspects of the total solution that customers desire. This is not good for business.

Ignorance may be partly to blame. Developers of Web-based learning environments are therefore encouraged to familiarize themselves with the framework's comparison dimensions. The framework's many dimensions may help developers realize the possible scope of a complete learning environment solution and lead them to consider broadening their solution. As was said earlier, no Web-based learning environment is expected to be the best in all the comparison dimensions. Market research will help developers determine which dimensions are particularly important for different classes of users and customers. Based on their expertise and the market they wish to target, developers should make deliberate and informed decisions as to which dimensions they will concentrate on and which dimensions they will sacrifice.

Developers who have an existing Web-based learning environment solution can use the comparison framework to assess how well their solution addresses each of the comparison dimensions. This will help them understand their strengths and weaknesses as compared to a complete learning environment solution. Critical weaknesses can then be improved upon.

The comparison framework can also be helpful for developers in assessing their competition. Assuming that developers understand their customer's needs, they can evaluate how well their solution stacks up against the competition. The dimensions where a developer's solution are relatively weak may represent opportunities for improvement and growth. The dimensions where a developer's solution are relatively strong may be good selling points that could be used effectively in marketing the solution.

Organizations that need a Web-based learning environment and are considering building their own would do well to consider carefully the many facets of a complete learning environment solution as captured in the comparison framework. There is much more to building a good learning environment than there is to building a good web site, which in itself is not such an easy task. They should develop their own solution only if they are confident they can adequately address every comparison dimension to the degree that they find it important.

4.2.2 Adopting Organizations

Organizations considering the adoption of a Web-based learning environment need to arm themselves with unbiased information to choose the best Web-based learning environment. Although people would surely like a fail-safe formula that they can apply to determine the best Web-based learning environment for their organization, that is not provided with this framework. Using the framework requires diligence. The benefit of using the framework is that it guides the comparison and ensures good coverage of issues that contribute to the success or failure of a Web-based learning environment. In this section we propose a user-centered approach that utilizes the comparison framework to help select the best possible learning environment for an organization.

There are five major steps to choosing a learning environment. Each of these, except for step 2, is supported by the comparison framework. The five steps are summarized in Table 4.20, depicted in Figure 4.2, and are as follows:

1. determine your organization's priorities;
2. identify candidate learning environments;
3. gather information about the candidate learning environments;

Table 4.20: Five steps of the adoption process

1. *Prioritize:*	Determine your organization's priorities;
2. *Identify:*	Identify candidate learning environments;
3. *Gather:*	Gather information about the candidate learning environments;
4. *Evaluate:*	Evaluate how well each candidate supports your priorities
5. *Compare:*	Compare candidates and eliminate inferior ones.

4. evaluate how well each candidate supports your priorities; and

5. compare candidates and eliminate inferior ones.

Prioritize

Organizations will probably find that the most difficult and critical step in the adoption process is establishing priorities. If a Web-based learning environment represents more than a trivial investment of resources for an organization, it will want to make the best choices possible. To determine its priorities properly, an organization's decision makers must consult a representative sample of every type of eventual user of the system. This point must be stressed because software purchasing decisions are often made by people other than the software's users. Since these people are usually not intimately familiar with users' needs, they may not prioritize their organization's needs accurately without direct input from end-users.

The comparison framework was designed to provide maximum coverage of the issues relevant to adopting a Web-based learning environment. The framework is therefore well suited to guide organizations as they establish their priorities. Organizations could use the framework in many ways to guide their decision making but they are strongly encouraged to involve their users in whatever method they choose. The

Figure 4.2: Adoption process data flow diagram

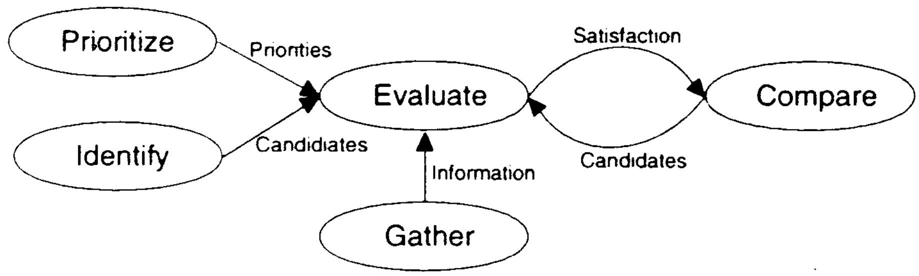

comparison framework also can guide organizations in making the necessary policy decisions before deploying a Web-based learning environment.

One manageable user-centered technique that organizations could employ is focus groups. Small focus groups (e.g. 3–10 people) can be assembled for each role in the organization that is affected by the learning environment. Two important roles that are obviously affected are learners and instructors. Other roles that are affected by a learning environment may include system administrators, computer services people, organization administrators, financial people, recruiters, and legal council. Each of the focus groups should be led by a skilled facilitator. The focus groups could discuss each of the dimensions in the comparison framework and make recommendations as to their prioritization, attributing to each dimension an importance value in the range of 0–10. Any dimensions or specific attributes of a Web-based learning environment that are considered must-haves, or show-stoppers, should be identified as such by the focus groups. The focus group leaders could then integrate each of the groups' recommendations into one cohesive set of priorities for the entire organization. This set of priorities should be validated in a session involving representatives from every role affected by the adoption of the learning environment, and the priorities should be reworked and validated again as needed.

The benefits of involving all parties affected by the introduction of a Web-based learning environment are many. First of all, the final prioritization is much more likely to represent the true values and needs of everyone concerned then if a small group of disconnected administrators made these decisions. Since the decision on which learning environment to adopt is based on the organization's priorities, it is important to get these right. Having representation from every role in the organization helps ensure that priorities are not skewed in favour of one group of people at the expense of others. Another very important benefit that may easily be overlooked is that involving representatives of all affected parties in the decision making process can increase people's buy-in and their acceptance of the final decision. This is particularly important for decisions that affect groups of people, as discussed in Section 3.5.2.

There are two main benefits to using the comparison framework in the prioritization exercise. The first is that it gives people a starting point and a structure for their deliberations. The second is that it may help people broaden their perspective. It is expected that not every dimension will be important to everyone but that aspects of most dimensions are likely to be important to someone in every organization. By using the comparison framework, people with particular roles in the organization are exposed to potentially important issues in the adoption of a learning environment that they might not have considered because they are not directly affected by them.

Identify

To select the best Web-based learning environment for themselves, organizations need to identify candidate systems. The comparison framework does not support this activity. There are, however, many resources available to organizations to help them initially identify candidate Web-based learning environments. The appendix contains a current list of Web-based learning environments and their contact information. These have not been included in the main body of this thesis because we recognize that the market changes quickly and that the list would therefore quickly be out of date. The World Wide Web contains a vast amount of information which organizations can search to extract current lists of candidate Web-based learning environments. Trade press and organizations that have already adopted learning environments are other sources from which candidate learning environments can be identified. Once candidate learning environments have been identified, fact-finding can begin.

Gather

Organizations must gather information about each of the Web-based learning environments it identified as candidates for adoption to be able to evaluate them. If it is made available, some of this information can be reused by other organizations who use the comparison framework to assist in selecting a Web-based learning environment.

The goal of the fact-finding exercise is to gather enough information about candidate learning environments to be able to compare them accurately. Initially, there is no need to gather all the information necessary for comparing every learning environment along each dimension of the comparison framework. It is most efficient to only gather the information necessary for each iteration of evaluation, as described in the following section. Information gathering therefore proceeds in tandem with evaluation iterations, beginning with gathering enough information to determine whether candidates satisfy the must-have requirements.

Information about candidate learning environments can be obtained from many sources. A good starting point is usually the source from which each candidate was identified. Vendors' sales people can usually provide a great deal of the needed information. It is also very helpful to obtain a trial version of each candidate Web-based learning environment to better evaluate them.

Evaluate

The evaluation step of the adoption process measures how well candidate Web-based learning environments support an organization's priorities. It makes use of the comparison framework as well as the output from the other steps of the adoption process, namely, the organization's priorities, the candidate learning environments, and the information gathered about each candidate learning environment.

There are many methods that organizations could use to evaluate Web-based learning environments. These should, however, be iterative and user-centered.

Iterative methods minimize the amount of work required to determine the suitability of a Web-based learning environment. Rather than fully evaluating every initial candidate learning environment, an iterative approach enables evaluators to eliminate unacceptable candidates early in the process.

Those who will be responsible for the evaluation need to use their best judgment to determine which decisions they can make on their own and which decisions would benefit from the input of users of the learning environment. For example, based on a fixed licensing schedule and a fixed budget, evaluators can reasonably determine on their own whether a candidate learning environment can be eliminated from consideration for fiscal reasons. Conversely, evaluators would likely benefit from the involvement of instructors and instructional designers to determine whether a candidate learning environment provides good support for course authoring.

There are many user-centered feedback techniques that evaluators can employ to obtain additional input to help them make decisions. Facilitated focus groups are an example of an effective technique for evaluating non-functional requirements that were described in the prioritization step. Functional requirements can also be evaluated with user-centered techniques. Two examples of such techniques are surveys and talk-aloud protocols. For talk-aloud protocols, users are given important tasks to perform with candidate learning environments and they voice their reasoning, their concerns and their impressions as they do so. Evaluators can obtain important information by observing people use candidate products.

The first iteration of the evaluation should determine whether each candidate learning environment satisfies the must-have priorities. This allows evaluators to quickly eliminate from consideration unsuitable candidate learning environments before investing too many resources in thoroughly evaluating them. This also allows us to assume for the rest of this process that every candidate is at very least suitable, that is, it satisfies the obligatory requirements.

The remainder of the adoption process consists of iteration cycles of evaluation and comparison. For each candidate still under consideration, evaluate how well the candidate satisfies the next most important comparison dimensions as determined in the prioritization step. For example, a score between 0 10 can be assigned to every dimension for every candidate learning environment representing the integrated view of how well the learning environment satisfies a comparison dimension for that organization.

Compare

Based on the priorities of comparison dimensions and on the degree to which each remaining candidate learning environment satisfies it, competing learning environments can be compared. Evaluators will have to use judgment to decide when a candidate is lagging sufficiently behind others to be eliminated from further consideration. It is not required by this process that candidates be eliminated from further consideration at every iteration. It is, however, most efficient to eliminate candidates as early as possible.

To help evaluators compare learning environments, graphs can be used to visualize the priorities and satisfaction scores of each of the comparison dimensions. For example, create one graph for each candidate learning environment. Each comparison dimension is plotted as a point on a two axis graph. The horizontal axis has a scale of 0 10 and represents the importance of the dimension as determined by the organization in the prioritization step. The vertical axis also has a scale of 0 10 but represents how well the candidate supports the comparison dimension as determined by the organization in the evaluation step. A sample graph is shown in Figure 4.3.

The upper right hand quadrant represents the strengths of a learning environment for an organization because the comparison dimensions that are plotted there are those which are important to the organization and well supported by the candidate learning environment. The lower right hand quadrant represents a learning environment's weaknesses for an organization because the comparison dimensions plotted there are those which are important to the organization but which are not well supported by the candidate learning environment. The comparison dimensions plotted in the left hand quadrants are not as critical to the decision making process. These can be used to distinguish between very closely matched candidate Web-based learning environments.

The sample graph in Figure 4.3 illustrates the strengths and weaknesses of a fictional candidate learning environment X for a particular adopting organization. The learning environment supports collaboration well, provides good support and has a

Figure 4.3: Sample graph for learning environment X

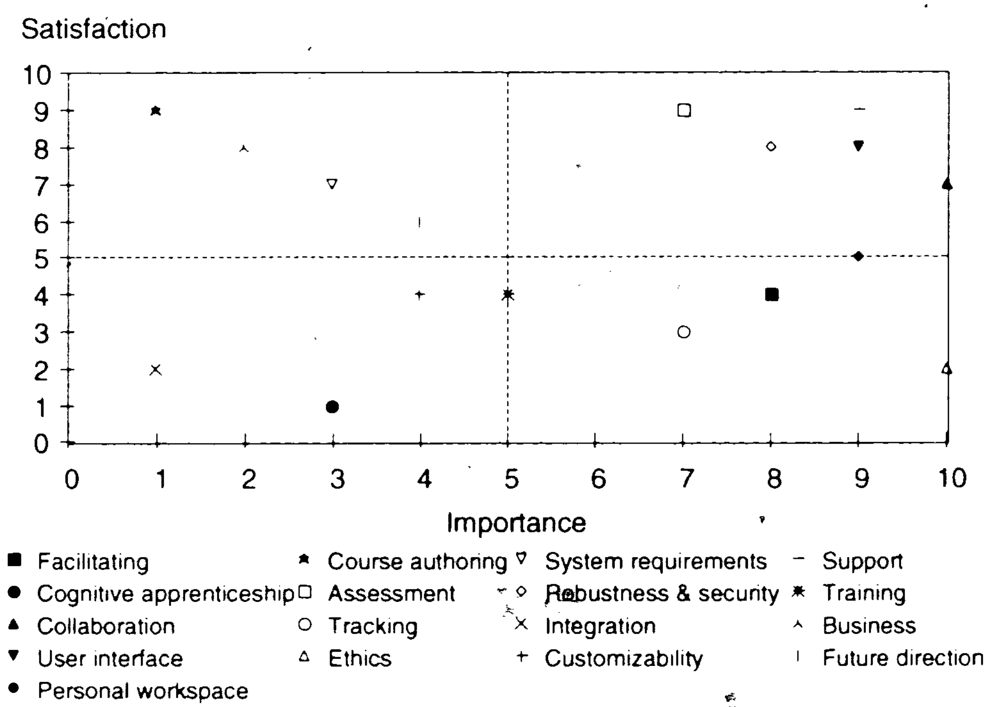

good user interface. These are all important attributes to the organization in question. The product is weak, however, in the support of ethical issues which is another very important issue for the organization. We can observe from the graph that the product has a good course authoring system but that this is not particularly valuable to the organization in question. Evaluators can compare the relative strengths and weaknesses of competing products by comparing their graphs.

With each new iteration, the graphs are completed progressively from right to left. By iterating on the next most important comparison dimension, the graphs are completed with the most critical comparison dimensions plotted first. This is good because it uncovers enough information to eliminate less appropriate candidate learning environment early in the process.

The evaluate and compare iteration cycles end when only one candidate learning environment remains. This is the Web-based learning environment that is likely to be the best overall for the adopting organization because it was selected by user-centered methods and by using the comparison framework.

Chapter 5

Discussion and Conclusions

This thesis presented a framework for comparing competing Web-based learning environments. The framework consists of seventeen dimensions along which Web-based learning environments can be compared covering a broad range of issues relevant to adopting and deploying a Web-based learning environment. Although some dimensions are characterized by software features or system specifications, the framework is not just a list of all the features a learning environment could possibly possess. Each dimension has a reason for existence that is justified either by related literature or by users' needs.

We recognize that many organizations planning on adopting a Web-based learning environment need a solution rather than a product. The framework therefore contains dimensions that treat aspects that are not directly related to software but that are still important for the overall success of a learning environment. Examples of these are the training and the support dimensions. Dimensions such as these could easily be overlooked without a comparison framework.

There are potentially many people playing different roles who must use a distributed learning environment for it to be successful. This is accounted for in the comparison framework. Although some comparison dimensions may not initially seem very important to instructors, for example, they may nevertheless be critical to users of the system who play another role, such as system administrators. Since the framework captures the needs of a wide variety of types of users, evaluators are reminded to consider the virtues of a candidate learning environment from multiple perspectives.

Intuition can help us to predict which dimensions might be particularly important to which people. Table 5.1 summarizes our expectations of the relative importance that would be attributed to each comparison dimension by learners, instructors, system administrators and organization administrators. It is important to note that these are only our general assumptions. The relative importance attributed to each comparison dimension is expected to vary depending on every organization's particular circumstances. Each organization considering the adoption of a Web-based learning environment should therefore involve the product's potential users in order to establish the organization's priorities.

Table 5.1 illustrates the disparity that is expected between the priorities of different types of users. Notice particularly that system administrator priorities are expected to often be opposite to those of instructors. Organizations will have to reconcile the priorities of all the user groups in order to establish their overall priorities.

The comparison framework is not meant to be able to identify the best Web-based learning environment devoid of context. Since different organizations have different needs and priorities, the learning environment that is best for one may not be the most suitable for another. The framework was designed in recognition of this fact. We proposed a user-centered approach to help organizations use the comparison framework to select the learning environment that is most appropriate for them.

The comparison framework also accounts for variations in how people are believed to learn. Based on a review of constructivist and sociocultural learning theory, three comparison dimensions were identified that correspond to innovative learning techniques. Educational researchers would be welcome to expand the comparison framework by considering other learning theories and by identifying other effective learning techniques that Web-based learning environments could attempt to support.

It was decided to not include a comparison of current Web-based learning environments in this thesis. Since we advocated the importance of organizations using a user-centered approach to select the best learning environment for themselves, it would have been inconsistent for us to then include an out-of-context comparison of

Table 5.1: Expected importance attributed to dimensions by user roles

Dimensions	**User Roles**			
	Learner	*Instructor*	*System Admin.*	*Organization Admin.*
Facilitating	low	high	none	low
Cognitive Apprenticeship	low	high	none	low
Collaboration	medium	high	none	low
User Interface	high	high	medium	low
Personal Workspace	high	medium	none	low
Course Authoring	none	high	none	low
Assessment	medium	high	none	low
Tracking	low	medium	high	medium
Ethics and Social Issues	medium	medium	low	high
System Requirements	medium	medium	high	high
Robustness and Security	medium	medium	high	high
Integration	none	none	high	medium
Customizability	low	medium	high	low
Support	medium	medium	high	medium
Training	low	medium	low	medium
Business	none	none	none	high
Future Direction	none	low	medium	high

a few learning environments based on the author's observations. Including a comparison based on a fictional scenario was one possibility that was considered but it was rejected because it did not seem that this would add much value to the thesis since organizations would still have to do their own comparisons. Furthermore, since new releases and new offerings of Web-based learning environments appear so frequently, any comparison included herein would have been very quickly outdated. In contrast, we expect the framework to have lasting value.

The usefulness of the framework would not have been proven by a quick and dirty evaluation of existing products. Its usefulness is conjectured based on the fact that organizations need help in selecting Web-based learning environments and on the fact that the framework is based on sound educational and human-computer interaction research. Its usefulness will be proven when organizations use it. We would be most interested in obtaining feedback from any organization that actually uses the framework to help them choose a Web-based learning environment.

There are a number of ways in which the work presented in this thesis could be extended. As mentioned above, a more thorough investigation of learning theories by experts in the field could suggest enhancements to the comparison framework. Existing comparison dimensions could be expanded or new dimensions could be added to the comparison framework in order to account for important educational characteristics of Web-based learning environments that have not yet been addressed.

The comparison dimensions themselves could be studied more closely to reveal dependencies between them. With a more precise understanding of how the current comparison dimensions affect each other, one might be able to reduce the set of dimensions to an independent set of dimensions that span the entire space of issues relevant to the adoption of a Web-based learning environment. Although it seems unlikely that such a set of independent dimensions could be found, identifying dependencies between dimensions would still be useful to evaluators.

After gaining experience with actually using the comparison framework, both the methodology for using the comparison framework and the framework itself could be improved. The methodology could be made more efficient and flexible. Researchers may identify other metrics that could be useful in measuring how well learning environments support organization's priorities. Other visualization tools might also be developed to visualize the relative strengths and weaknesses of competing candidate Web-based learning environments. The framework's comparison dimensions could be reviewed in order to include important aspects of Web-based learning environments that were initially overlooked.

Although it would be difficult, the scope of the comparison framework could be expanded to include other forms of distributed learning environments as opposed to only Web-based learning environments. Organizations who have a need for a distributed learning solution would probably appreciate assistance in selecting the best learning environment possible regardless of the underlying technology and medium. Many of the comparison dimensions that have been presented herein would still be valid. The economic considerations would, however, be far more complicated and the debate over the advantages and disadvantages of different media would have to be addressed.

A time consuming step in using the comparison framework is gathering information. The comparison framework could therefore be made easier to use if someone were to maintain an up-to-date central database of facts about available Web-based learning environments. The Appendix to this thesis contains a partial list of current Web-based learning environments and the relevant contact information.

In conclusion, we built a framework for the comparison of Web-based learning environment and proposed a methodology for using it. This framework can now be used by organizations considering the adoption of a Web-based learning environment to help them select the most suitable learning environment. It can also be used by developers of Web-based learning environments to discover the relative strengths and weaknesses of their product as compared to the competition.

Appendix: Web-Based Learning Environments

Product	Vendor Contact Information[a]
Athena University	VOU Services International www.athena.edu
Auditorium	PlaceWare www.placeware.com
Authorware	Macromedia www.macromedia.com/software/authorware
CBT Systems	CBT Systems www.cbtsys.com
CSILE [b]	Ontario Institute for Studies in Education csile.oise.on.ca
DigitalThink	DigitalThink www.digitalthink.com
EDROADS [c]	City University www.cityu.edu
eZone	Macmillan Computer Publishing www.mcp.com/waite/ezone/ezone

[a] World-Wide Web addresses were accurate as of August 25, 1997
[b] CSILE is Computer Supported Intentional Learning Environment
[c] EDROADS is Education Resource and Online Academic Degree System

Product	Vendor / Contact Information[a]
IBM Global Campus	IBM ike.engr.washington.edu/igc
Internet Learning Center	Gartner Group www.gglearning.com
LOIS [b]	KnowledgeSoft www.knowledgesoft.com
LearningSpace	Lotus Development www.lotus.com/learningspace
Learning University	Street Technologies www.learninguniversity.com
MOLI [c]	Microsoft moli.microsoft.com
Oracle Learning Architecture	Oracle ola.oracle.com
Personal Learning System	IBM www.uk.ibm.com/learningtechnology
ToolBook II	Asymetrix www.asymetrix.com/products/toolbook2
TopClass	WBT Systems www.wbtsystems.com
Virtual-U	Virtual Learning Environments virtual-u.cs.sfu.ca
WebCT	WebCT Educational Technologies homebrew1.cs.ubc.ca/webct
ZD Net University	Ziff-Davis Publishing Company www.zdu.com

[a] World-Wide Web addresses were accurate as of August 25, 1997

[b] LOIS is Learning Organization Information System

[c] MOLI is Microsoft Online Institute

Bibliography

[1] ACM. *Proceedings of the Conference on Computer Support for Collaborative Learning*. Bloomington, IN, USA, October 1995. Lawrence Erlbaum Associates.

[2] ACM SIGCHI. *CHI '96 Companion: Human Factors in Computing Systems: Common Ground*. Vancouver, BC, Canada, April 1996. Addison-Wesley.

[3] ACM SIGCHI. *CHI '96 Proceedings: Human Factors in Computing Systems: Common Ground*. Vancouver, BC, Canada, April 1996. Addison-Wesley.

[4] ACM SIGCHI. *CHI '97 Proceedings: Human Factors in Computing Systems*, Atlanta, GA, USA, March 1997.

[5] Shirley Alexander. Teaching and learning on the World Wide Web. In *Proceedings of AusWeb 95*, pages 93–102, 1995. Available: elmo.scu.edu.au/sponsored/ausweb/ausweb95/papers/education2/alexander/ [March 14, 1997].

[6] John R. Anderson, Lynne M. Reder, and Herbert A. Simon. Situated learning and education. *Educational Researcher*, 25(4):5–11, May 1996.

[7] John R. Anderson, Lynne M. Reder, and Herbert A. Simon. Situative versus cognitive perspectives: Form versus substance. *Educational Researcher*, 26(1):18–21, 1997.

[8] AusWeb. *Proceedings of the Second Australian World Wide Web Conference*, Gold Coast, Australia, July 1996.

[9] AusWeb. *Proceedings of the Third Australian World Wide Web Conference*. Southern Cross University, New South Wales, Australia, July 1997.

[10] D. Bachiochi, M. Berstene, E. Chouinard, N. Conlan, M. Danchak, T. Furey, C. Neligon, and D. Way. Usability studies and designing navigational aids for the World Wide Web. In *Hyperproceedings of the Sixth International World Wide Web Conference*, Santa Clara, CA, USA, April 1997. Available: ww6.nttlabs.com [May 1, 1997].

[11] Ronald M. Baecker, Jonathan Grudin, William A. S. Buxton, and Saul Greenberg, editors. *Readings in Human Computer Interaction: Towards the Year 2000*. Morgan Kaufmann Publishers, San Francisco, CA, USA, second edition, 1995.

[12] Anne K. Bednar, Donald Cunningham, Thomas M. Duffy, and J. David Perry. Theory into practice: How do we link? In G. Anglin, editor, *Instructional Technology: Past, Present, and Future*. Libraries Unlimited, Denver, CO, USA, 1991.

[13] Sue E. Berryman. Cognitive science: Challenging schools to design effective learning environments. Technical Report T-23, Institute on Education and the Economy, Teachers College, Columbia University, Box 174, 439 Thorndike Hall, 525 West 120th Street, New York, NY, USA 10027, September 1991.

[14] Herb Bethoney. Computer-based training on the Web Learning long distance: The Labs evaluates seven vendors' solutions for delivering training via the Net. *PCWeek Online*, August 18, 1997. Available: www8.zdnet.com/pcweek/reviews/0818/18ibt.htm [August 19, 1997].

[15] John B. Black and Robert O. McClintock. An interpretation construction approach to constructivist design. In Wilson [116].

[16] Phyllis C. Blumenfeld, Ronald W. Marx, Elliot Soloway, and Joseph Krajcik. Learning with peers: From small group cooperation to collaborative communitites. *Educational Researcher*, 25(8):37–40, November 1996.

[17] Susanne Bodker, Kaj Gronbæk, and Morten Kyng. Cooperative design: Techniques and experiences from the Scanadinavian scene. In D. Schuler and A. Namioka, editors, *Participatory Design: Principles and Practices*, pages 157–175. Lawrence Erlbaum Associates, Hillside, NJ, USA, 1993.

[18] Curtis Jay Bonk and Thomas H. Reynolds. Learner-centered Web instruction for higher-order thinking, teamwork, and apprenticeship. In Khan [62], chapter 20, pages 167–178.

[19] Stephen J. Bostock. Designing Web-based instruction for active learning. In Khan [62], chapter 26, pages 225–230.

[20] Patrick W. Brewer, Elizabeth J. Gibson, Ajay Dholakia, Mladen A. Vouk, and Donald L. Bitzer. Advanced computer based education on the World Wide Web. In *Proceedings of the Fourth World Wide Web Conference*, Boston, MA, USA, December 1995. Available: renoir.csc.ncsu.edu/MRA/Reports/AdvEduWeb.html [March 18, 1997].

[21] John Seely Brown, Allan Collins, and Paul Duguid. Situated cognition and the culture of learning. *Educational Researcher*, 18(1):32–42, 1989.

[22] M. Sheelagh T. Carpendale, David Cowperthwaite, and F. David Fracchia. 3-dimensional pliable surfaces: For the effective presentation of visual information. In *Proceedings of the Annual ACM Symposium on User Interface Software and Technology*, pages 217–226. ACM Press, 1995.

[23] Jack M. Carroll and C. Carrithers. Training wheels in a user interface. *Communications of the ACM*, 27(8):800–806, 1984.

[24] Gérald Collaud, John Dill, Christopher V. Jones, and Paul Tan. The continuously zoomed Web — a graphical navigation aid for WWW. In *IEEE Visualization '96 Late Breaking Hot Topics Papers*, 1996.

[25] Allan Collins, John Seely Brown, and S. E. Newman. Cognitive apprenticeship: Teaching the crafts of reading, writing and mathematics. In L. B. Resnick,

editor, *Knowing, Learning, and Instruction: Essays in Honor of Robert Glaser*, pages 453–494. Lawrence Erlbaum Associates, Hillsdale, NJ, USA, 1989.

[26] CommerceNet and Nielson Media Research. "Startling increase" in Internet shopping reported in new CommerceNet/Nielsen Media Research survey. Press Release, March 1997. Available: www.commerce.net/nielsen/press_97.html [April 23, 1997].

[27] Richard H. Cutler. Distributed presence and community in cyberspace. *Interpersonal Computing and Technology: An Electronic Journal for the 21st Century*, 3(2):12–32, April 1995. Available: www.helsinki.fi/science/optek/1995/n2/cutler.txt [March 4, 1997].

[28] Rudolph P. Darken and John L. Sibert. Wayfinding strategies and behaviors in large virtual worlds. In *CHI '96 Proceedings: Human Factors in Computing Systems: Common Ground* [3], pages 142–149.

[29] James R. Davis and Daniel P. Huttenlocher. Shared annotation for cooperative learning. In *Proceedings of the Conference on Computer Support for Collaborative Learning* [1].

[30] Frank N. Dempster. The spacing effect: A case study in the failure to apply the results of psychological research. *American Psychologist*, 43(8):627–634, August 1988.

[31] Francis J. Di Vesta. The cognitive movement and education. In John A. Glover and Royce R. Ronning, editors, *Historical Foundations of Educational Psychology*, chapter 11, pages 203–233. Plenum Press, New York, NY, USA, 1987.

[32] John Dill, Lyn Bartram, Albert Ho, and Frank Henigma. A continuously variable zoom for navigating large hierarchical networks. In *Proceeding of IEEE International Conference on Systems, Man and Cybernetics*, pages 386–390, October 1994.

[33] Andrew Dillon and Erping Zhu. Designing Web-based instruction: A human-computer interaction perspective. In Khan [62], chapter 25, pages 221–224.

[34] R. T. Jim Eales and Jim Welsh. Design for collaborative learnability. In *Proceedings of the Conference on Computer Support for Collaborative Learning* [1].

[35] Daniel C. Edelson and D. Kevin O'Neill. The CoVis Collaboratory Notebook: Computer support for scientific inquiry. Paper presented at the Annual Meeting of the American Educational Research Association, New Orleans, LA, USA, April 1996.

[36] Daniel C. Edelson, D. Kevin O'Neill, Louis M. Gomez, and Laura D'Amico. A design for effective support of inquiry and collaboration. In *Proceedings of the Conference on Computer Support for Collaborative Learning* [1].

[37] Daniel C. Edelson, Roy D. Pea, and Louis Gomez. Constructivism in the collaboratory. In Wilson [116].

[38] Thomas D. Erickson. Working with interface metaphors. In Baecker et al. [11], pages 147–151.

[39] Marjo Favorin. Towards computer support for collaborative learning at work: Six requirements. In *Proceedings of the Conference on Computer Support for Collaborative Learning* [1].

[40] Ulrich Gall and Franz J. Hauck. Promondia: A Java-based framework for real-time group communication in the Web. In *Hyperproceedings of the Sixth International World Wide Web Conference*, Santa Clara, CA, USA, April 1997. Available: ww6.nttlabs.com [July 14, 1997].

[41] Elizabeth J. Gibson, Patrick W. Brewer, Ajay Dholakia, Mladen A. Vouk, and Donald L. Bitzer. A comparative analysis of web-based testing and evaluation systems. In *Proceedings of the Fourth World Wide Web Conference*, Boston, MA, USA, December 1995. Available: renoir.csc.ncsu.edu/MRA/Reports/WebBasedTesting.html [February 26, 1997].

[42] Saul Greenberg. A fisheye text editor for relaxed-WYSIWIS groupware. In *CHI '96 Companion: Human Factors in Computing Systems: Common Ground* [2], pages 212–213.

[43] Saul Greenberg. Peepholes: Low cost awareness of one's community. In *CHI '96 Companion: Human Factors in Computing Systems: Common Ground* [2], pages 206–207.

[44] James G. Greeno. On claims that answer the wrong questions. *Educational Researcher*, 26(1):5–17, 1997.

[45] Jonathan Grudin. Groupware and social dynamics: Eight challenges for developers. In Baecker et al. [11], pages 762–774.

[46] Carl Gutwin and Saul Greenberg. Workspace awareness for groupware. In *CHI '96 Companion: Human Factors in Computing Systems: Common Ground* [2], pages 208–209.

[47] Carl Gutwin, Gwen Stark, and Saul Greenberg. Support for workspace awareness in educational groupware. In *Proceedings of the Conference on Computer Support for Collaborative Learning* [1], pages 147–156.

[48] Linda Harasim, Tom Calvert, and Chris Groeneboer. Virtual-U: A Web-based system to support collaborative learning. In Khan [62], chapter 18, pages 149–158.

[49] Linda Harasim, Starr Roxanne Hiltz, Lucio Teles, and Murray Turoff, editors. *Learning Networks: A Field Guide to Teaching and Learning Online*. MIT Press, Cambridge, MA, USA, 1995.

[50] Jim Hewitt and Marlene Scardamalia. Design principles for the support of distributed processes. Paper presented at the Annual Meeting of the American Educational Research Association, New Orleans, LA, USA, April 1996.

[51] Janette R. Hill. Distance learning environments via the World Wide Web. In Khan [62], chapter 8, pages 75–80.

[52] Starr Roxanne Hiltz. Impacts of college-level courses via asynchronous learning networks: Some preliminary results. *Journal of Asynchronous Learning Networks*, 1997.

[53] Paul Hoffert. Tom Jurenka nad Barry Silverman, Paul Spurgeon, and Lucy White. Managing intellectual property in digital formats: IVY/Intercom Ontario - a case study. US Copyright Office and Interactive Multimedia Association, Electronic Commerce for Content, Forum for Inter-Industry Requirements for Technology-Based Intellectual Property Management, October 1996.

[54] Peter C. Honebein, Thomas M. Duffy, and Barry J. Fishman. Constructivism and the design of learning environments: Context and authentic activities for learning. Available: ftp://ithaca.icbl.hw.ac.uk/pub/nato_asi/duf2.txt.gz [March 17, 1997].

[55] IBM Corporation. HCI guidelines. Available: www.ibm.com/ibm/hci/guidelines/guidelines.html [July 22, 1997].

[56] IBM Corporation. *Object-Oriented Interface Design: IBM Common User Access Guidelines*. QUE, 1992.

[57] International Organization for Stardardization. ISO Online. Available: www.iso.ch [October 2, 1997].

[58] David Jonassen, Terry Mayes, and Ray McAleese. A manifesto for a constructivist approach to technology in higher education. In Thomas M. Duffy, Joost Lowyck, and David H. Jonassen, editors, *Designing Environments for Constructivist Learning*, number 105 in NATO ASI Series F, Computer and Systems Sciences. Springer-Verlag, Berlin, Germany, 1993.

[59] David H. Jonassen. Thinking technology: Toward a constructivist design method. *Educational Technology*, 34(4):34–37, April 1994.

[60] David H. Jonassen, Dean Dyer, Karen Peters, Timothy Robinson, Douglas Harvey, Marsha King, and Pamela Loughner. Cognitive flexibility hypertexts on the

Web: Engaging learners in meaning making. In Khan [62], chapter 15, pages 119–133.

[61] Marshall G. Jones and John D. Farquhar. User interface design for Web-based instruction. In Khan [62], chapter 28, pages 239–244.

[62] Badrul H. Khan, editor. *Web-Based Instruction*. Educational Technology Publications, Englewood Cliffs, NJ,USA, 1997.

[63] Badrul H. Khan. Web-based instruction (WBI): What is it and why is it? In Khan [62], chapter 1, pages 5–18.

[64] Bruce Landon. Comparative analysis of on-line educational delivery applications. Standing Committee on Educational Technology, 1997. Available: www.ctt.bc.ca/landonline [July 31, 1997].

[65] J. Lave and E. Wenger. *Situated Learning: Legitimate Peripheral Participation*. Cambridge University Press, Cambridge, Great Britain, 1991.

[66] Lotus Development Corporation. Lotus Notes. Available: www.lotus.com/notes [September 14, 1997].

[67] Matrix Information & Directory Services. Matrix maps quarterly 401: State of the internet, January 1997. Available: www3.mids.org/mmq/401/ [May 15, 1997].

[68] Richard E. Mayer. Cognition and instruction: Their historic meeting within educational psychology. *Journal of Educational Psychology*, 84(4):405–412, 1992.

[69] Richard E. Mayer. Learners as information processors: Legacies and limitations of educational psychology's second metaphor. *Educational Psychologist*, 31(3/4):151–161, 1996.

[70] Richard E. Mayer. Multimedia learning: Are we asking the right questions? *Educational Psychologist*, 32(1):1–19, 1997.

[71] H. McLellan. Evaluation in a situated learning environment. *Educational Technology*, 33(3):39–45, 1993.

[72] Thomas Fox McManus. Special considerations for designing internet based education. In D. Willis, B. Robin, and J. Willis, editors, *Technology and Teacher Education Annual, 1995*, Charlottesville, VA, USA, 1995. Association for Advancement of Computing in Education. Available: www.edb.utexas.edu/coe/depts/ci/it/projects/papers/special.html [March 18, 1997].

[73] Microsoft Corporation. Internet Explorer. Available: www.microsoft.com/ie [September 14, 1997].

[74] Microsoft Corporation. NetMeeting. Available: www.microsoft.com/netmeeting [September 14, 1997].

[75] Netscape Corporation. Netscape Communicator. Available: www.netscape.com/communicator [September 14, 1997].

[76] Gordon K. Niguma. Concept mapping in a multimedia, World Wide Web environment. Master's thesis, Simon Fraser University, Burnaby, BC, Canada, April 1997.

[77] Donald A. Norman. *The Psychology of Everyday Things*. Basic Books, 1988.

[78] Donald A. Norman and Stephen. W. Draper, editors. *User-Centered System Design*. Laurence Erlbaum Associates, Hillside, NJ, USA, 1986.

[79] Robert Orfali, Dan Harkey, and Jeri Edwards. *The Essential Client/Server Survival Guide*. Wiley Computer Publishing, New York, NY, USA, second edition, 1996.

[80] Wanda J. Orlikowski. Learning from notes: Organizational issues in groupware. In *Proceeding of CSCW '92*. ACM, 1992. 362–369.

[81] Russell Pennell. Managing online learning. In *Proceedings of the Second Australian World Wide Web Conference* [8]. Available: elmo.scu.edu.au/sponsored/ausweb/ausweb96/educn/pennell/papers.html [March 3, 1997].

[82] Jean Piaget. *The Construction of Reality in the Child.* Basic Books, New York, NY, USA, 1954.

[83] PointCast Canada. PointCast. Available: www.pointcast.ca [September 15, 1997].

[84] Chris Quintana, Kasem Abotel, and Elliot Soloway. NoRIS: Supporting computational science activities through learner-centered design. In D. C. Edelson and E. A. Domeshek, editors, *Proceedings of ICLS '96: Second Annual Conference on the Learning Sciences*, pages 272–279. Association for the Advancement of Computing in Education, 1996.

[85] Thomas C. Reeves and Patricia M. Reeves. Effective dimensions of interactive learning on the World Wide Web. In Khan [62], chapter 6, pages 59–66.

[86] Anju Relan and Bijan B. Gillani. Web-based instruction and the traditional classroom: Similarities and differences. In Khan [62], chapter 4, pages 41–46.

[87] Alexander J. Romiszowski. Web-based distance learning and teaching: Revolutionary invention or reaction to necessity? In Khan [62], chapter 3, pages 25–37.

[88] Martin Röscheisen, Christian Mogensen, and Terry Winograd. Beyond browsing: Shared comments, SOAPs, trails, and on-line communities. In *Proceedings of the Third International World-Wide Web Conference*, Darmstadt, Germany, April 1995.

[89] Mark Roseman and Saul Greenberg. Building real time groupware with GroupKit, a groupware toolkit. *ACM Transactions on Computer Human Interaction*, 3(1):66–106, March 1996.

[90] Philip Rutherford. Annotations: The key to collaborative teaching and learning on the Internet. In *Proceedings of the Second Australian World Wide Web Conference* [8]. Available: elmo.scu.edu.au/sponsored/ausweb/ausweb96/educn/rutherford/paper.html [March 18, 1997].

[91] Michaeil Scaife, Yvonne Rogers, Frances Aldrich, and Matt Davies. Designing for or designing with? Informant design for interactive learning environments. In *CHI '97 Proceedings: Human Factors in Computing Systems* [4].

[92] Marlene Scardamalia and Carl Bereiter. An architecture for collaborative knowledge building. In E. DeCorte, M. Linn, H. Mandl, and L. Verschaffel, editors, *Computer-Based Learning Environments and Problem Solving*, volume 84 of *NATO ASI Series F, Computer and Systems Sciences*, pages 41–46. Springer-Verlag, Berlin, Germany, 1992.

[93] Roger Schank. End run to the goal line. *Educom Review*, 30(1), January 1995. Available: educom.edu/web/pubs/review/reviewArticles/30114.html [March 5, 1997].

[94] M. Schrage. *Shared Minds: The New Technologies of Collaboration*. Random House, New York, NY, USA, 1991.

[95] Jerald G. Schutte. Virtual teaching in higher education, 1997. Available: www.csun.edu/sociology/virexp.htm [March 18, 1997].

[96] Don Scott, Catherine Durnell Cramston, Stéphane Gauvin, Beata Lobert, Gerhard Steinke, and Karey Patterson. Internet based collaborative learning: An empirical evaluation. In *Proceedings of the Third Australian World Wide Web Conference* [9].

[97] Mildred L. G. Shaw and Biran R. Gaines. Comparing constructions through the Web. In *Proceedings of the Conference on Computer Support for Collaborative Learning* [1].

[98] Martin A. Siegel and Sonny Kirkley. Moving toward the digital learning environment: The future of Web-based instruction. In Khan [62], chapter 31, pages 263–270.

[99] Jean-Paul Simbandumwe. Tools for developing interactive academic web courses. University of Manitoba, 1997. Available: www.umanitoba.ca/ip/tools/courseware/evalmain.html [July 31, 1997].

[100] Elissa D. Smilowitz. Do metaphors make Web browsers easier to use? In *Conference Proceedings of Designing for the Web: Empirical Studies*, Sandia National Lab, Albuquerque, NM, USA, October 1996. Available: www.microsoft.com/usability/webconf.htm [July 14, 1997].

[101] H. Jefferson Smith. Complying with federal sentencing guidelines. *Beyond Computing*, 6(6), August 1997. Available: www.beyondcomputingmag.com [August 26, 1997].

[102] Candace Soderston and Thyra Rauch. The case for user-centered design. In *Proceedings of the STC 44th Annual Conference*. Society for Technical Communication, 1996.

[103] Elliot Soloway, Mark Guzdial, and Kenneth E. Hay. Learner-centered design: The challenge for HCI in the 21st century. *Interactions*, pages 36–48, April 1994.

[104] Elliot Soloway, Shari L. Jackson, Jonathan Klein, Chris Quintana, James Reed, Jeff Spitulnik, Steven J. Stratford, Scott Struder, Jim Eng, and Nancy Scala. Learning theory in practice: Case studies of learner-centered design. In *CHI '96 Proceedings: Human Factors in Computing Systems: Common Ground* [3], pages 189–196.

[105] Rand J. Spiro, Paul J. Feltovich, Michael J. Jacobson, and Richard L. Coulson. Cognitive flexibility, constructivism, and hypertext: Random access instruction for advanced knowledge acquisition in ill-structured domains. In T. M. Duffy

and D. H. Jonassen, editors, *Constructivism and the Technology of Instruction: A Conversation*. Lawrence Erlbaum Associates, Hillsdale, NJ, USA, 1992.

[106] Jeff Spitulnik, Scott Studer, Elizabeth Finkel, Edwin Gustafson, Jason Laczko, and Elliot Soloway. Toward supporting learners participating in scientifically-informed community discourse. In *Proceedings of the Conference on Computer Support for Collaborative Learning* [1], pages 317–320.

[107] Denise Stockley and Michael Roth. Situated learning: Application to educational technology. In *Ed-Media/Ed-Telecom*, Calgary, AB, Canada, June 1997.

[108] Sun Microsystems. The java language: An overview. White Paper. Available: www.javasoft.com:81/docs/Overviews/java/java-overview-1.html [April 23, 1997].

[109] Linda Tauscher and Saul Greenberg. Revisitation patterns in World Wide Web navigation. In *CHI '97 Proceedings: Human Factors in Computing Systems* [4].

[110] Lucio Teles. Cognitive apprenticeship on global networks. In Linda M. Harasim, editor, *Global Networks: Computers and International Communication*, chapter 16, pages 271–281. MIT Press, Cambridge, MA, USA, 1993.

[111] S. D. Tripp. Theories, traditions and situated learning. *Educational Technology*, 33(3):71–77, 1993.

[112] Murray Turoff. Designing a virtual classroom. In *Invited Papers for the 1995 International Conference on Computer Assisted Instruction*, March 1995.

[113] John Vargo. Evaluating the effectiveness of Internet delivered coursework. In *Proceedings of the Third Australian World Wide Web Conference* [9].

[114] Lem S. Vygotsky. *Thought and Language*. MIT Press, Cambridge, MA, USA, 1962.

[115] Barry Willis and John Dickinson. Distance education and the World Wide Web. In Khan [62], chapter 9, pages 81–84.

[116] Brent G. Wilson, editor. *Constructivist Learning Environments: Case Studies in Instructional Design*. Educational Technology Publications, Englewood Cliffs, NJ, USA, 1995.

[117] Brent G. Wilson. Metaphors for instruction: Why we talk about learning environments. *Educational Technology*, 35(5):25–30, 1995.

[118] Brent G. Wilson. Understanding the design and use of learning technologies. Paper presented at the meeting of the Association for Educational Communications and Technology, Albuquerque, NM, USA, February 1997.

[119] Ping-Jer Yeh, Bih-Horng Chen, Ming-Chih Lai, and Shyan-Ming Yuan. Synchronous navigation control for distance learning on the web. In *Proceedings of the Fifth International World Wide Web Conference*, Paris, France, May 1996. Available: www5conf.inria.fr/fich_html/papers/P28/Overview.html [March 4, 1997].

[120] Stanton M. Zobel. Legal implications of intellectual property and the World Wide Web. In Khan [62], chapter 42, pages 337–340.