PART

I

Introduction to Artificial Intelligence and Education

CHAPTER

Introduction

1

People need a lifetime to become skilled members of society; a high school diploma no longer guarantees lifelong job prospects. Now that the economy has shifted from manual workers to knowledge workers, job skills need to be updated every few years, and people must be prepared to change jobs as many as five times in a lifetime. Lifelong learning implies lifelong education, which in turn requires supportive teachers, good resources, and focused time. Traditional education (classroom lectures, texts, and individual assignments) is clearly not up to the task. Current educational practices are strained to their breaking point.

The driving force of the knowledge society is information and increased human productivity. Knowledge workers use more information and perform more operations (e.g., compose a letter, check its content and format, send it, and receive a reply within a few moments) than did office workers who required secretarial assistance to accomplish the same task. Similarly, researchers now locate information more quickly using the Internet than did teams of researchers working for several months using conventional methods. Marketing is facilitated by online client lists and digital advertising created by a single person acting as author, graphic designer, layout artist, and publisher. To prepare for this society, people need education that begins with the broadest possible knowledge base; knowledge workers need to have more general knowledge and to learn with less support.

Information technology has generated profound changes in society, but thus far it has only subtly changed education. Earlier technologies (e.g., movies, radio, television) were touted as saviors for education, yet nearly all had limited impact, in part because they did not improve on prior educational tools but often only automated or replicated existing teaching strategies (e.g., radio and television reproduced lectures) (McArthur et al., 1994).

On the other hand, the confluence of the Internet, artificial intelligence, and cognitive science provides an opportunity that is qualitatively different from that of preceding technologies and moves beyond simply duplicating existing teaching processes. The Internet is a flexible medium that merges numerous communication devices (audio, video, and two-way communication), has changed how educational content is produced, reduced its cost, and improved its efficiency. For example, several new teaching methods (collaboration and inquiry learning) are now possible through technology. Multiuser activities and online chat offer opportunities not possible before in the classroom.

> *What one knows is, in youth, of little moment; they know enough who know how to learn.*
>
> **Henry Adams (1907)**

We do not propose that technology alone can revolutionize education. Rather, changes in society, knowledge access, teacher training, the organization of education, and computer agents help propel this revolution.

This book offers a critical view of the opportunities afforded by a specific genre of information technology that uses artificial intelligence and cognitive science as its base. The audience for this book includes people involved in computer science, psychology and education, from teachers and students to instructional designers, programmers, psychologists, technology developers, policymakers, and corporate leaders, who need a well-educated workforce. This chapter introduces an inflection point in education, discusses issues to be addressed, examines the state of the art and education, and provides an overview of the book.

1.1 AN INFLECTION POINT IN EDUCATION

In human history, one technology has produced a salient and long-lasting educational change: the printing press invented by Johannes Gutenberg around 1450. This printing press propelled a transfer from oral to written knowledge and supported radical changes in how people thought and worked (Ong and Walter, 1958). However, the advances in human literacy resulting from this printing press were slow to take hold, taking hundreds of years as people first learned to read and then changed their practices.

Now computers, a protean and once-in-several-centuries innovation, have produced changes in nearly every industry, culture, and community. It has produced more than incremental changes in most disciplines; it has revolutionized science, communication, economics, and commerce in a matter of decades. Information technology, including software, hardware, and networks, seems poised to generate another *inflection point* in education. An inflection point is a full-scale change in the way an enterprise operates. Strategic inflection points are times of extreme change; they can be caused by technological change but are more than technological change (Grove, 1996). By changing the way business is conducted, an inflection point creates opportunities for players who are adept at operating in the new environment (e.g., software vendors and e-learning companies) to take advantage of an opportunity for new growth.

One example of a business inflection point is the Japanese manufacture of smaller and cheaper memory products, which created an inflection point for other manufacturers of memory products. Intel and others were forced out of the memory chip business and into the relatively new field of microprocessors (Grove, 1996). This microprocessor business then created another inflection point for other companies, bringing difficult times to the classical mainframe computer industry. Another example of an inflection point is the automated teller machine, which changed the banking industry. One more example is the capacity to digitally create, store, transmit, and display entertainment content, which changed the entire media industry. In short, strategic inflection points may be caused by technology, but they fundamentally change enterprise.

Education is a fertile market within the space of global knowledge, in which the key factors are knowledge, educated people, and knowledge workers. The knowledge economy depends on productive and motivated workers who are technologically literate and positioned to contribute ideas and information and to think creatively. Like other industries (e.g., health care or communications), education combines large size (approximately the same size as health care in number of clients served), disgruntled users, lower utilization of technology, and possibly the highest strategic importance of any activity in a global economy (Dunderstadt, 1998).

The future impact of information technology on education and schools is not clear, but it is likely to create an inflection point that affects all quadrants. Educators can augment and redefine the learning process by taking advantage of advances in artificial intelligence and cognitive science and by harnessing the full power of the Internet. Computing power coupled with decreased hardware costs result in increased use of computation in all academic disciplines (Marlino et al., 2004). In addition, technological advances have improved the analysis of both real-time observational and computer-based data. For example, the science community now has tools of greater computational power (e.g., higher resolution, better systems for physical representation and modeling, and data assimilation techniques), facilitating their understanding of complex problems. Science educators are incorporating these tools into classrooms to stimulate motivation and curiosity and to support more sophisticated student understanding of science. Learners at all levels have responded to computational simulations that make concepts more engaging and less abstract (Manduca and Mogk, 2002). Students who use this technology think more deeply about complex skills, use enhanced reasoning, and have better comprehension and design skills (Roschelle et al., 2000). Computers improve students' attitudes and interests through more interactive, enjoyable, and customizable learning (Valdez et al., 2000).

Formal public education is big business in terms of the numbers of students served and the requisite infrastructure (Marlino et al., 2004); during the 1990s, public education in the United States was a $200 billion-a-year business (Dunderstadt, 1998). More than 2.1 million K-12 teachers in 91,380 schools across the United States teach 47 million public school students (Gerald and Hussar, 2002; Hoffman, 2003). More than 3,700 schools of higher education in the United States prepare the next generation of scientific and educational workers (National Science Board [NSB], 2003).

A major component of the educational inflection point is the Internet, which is now the world's largest and most flexible repository of education material. As such, the Internet moves education from a loosely federated system of state institutions and colleges constrained by space and time into a knowledge-and-learning industry.

This technological innovation signals the beginning of the end of traditional education in which lectures are fixed in time and space.

One billion people, or more than 16.7% of all people worldwide, use the Internet (Internetworldstats, 2006). In some countries, this percentage is much higher (70% of the citizens in the United States are web users, 75% in Sweden, and 70% in Denmark) and is growing astronomically (Almanac, 2005). The Internet links more than 10 billion pages, creating an opportunity to adapt millions of instructional resources for individual learners.

Three components drive this educational inflection point. They are artificial intelligence (AI), cognitive science, and the Internet:

- AI, the science of building computers to do things that would be considered intelligent if done by people, leads to *a deeper understanding of* knowledge, especially representing and reasoning about "how to" knowledge, such as procedural knowledge.
- Cognitive science, or research into understanding how people behave intelligently, leads to a deeper understanding of how people think, solve problems, and learn.
- The Internet provides an unlimited source of information, available anytime, anywhere.

These three drivers share a powerful synergy. Two of them, AI and cognitive science, are two sides of the same coin—that is, understanding the nature of intelligent action, in whatever entity it is manifest. Frequently, AI techniques are used to build software models of cognitive processes, whereas results from cognitive science are used to develop more AI techniques to emulate human behavior. AI techniques are used in education to model student knowledge, academic topics, and teaching strategies. Add to this mix the Internet, which makes more content and reasoning available for more hours than ever before, and the potential inflection point leads to unimaginable activities supporting more students to learn in less time.

Education is no longer perceived as "one size fits all." Cognitive research has shown that the learning process is influenced by individual differences and preferred learning styles (Bransford et al., 2000b). Simultaneously, learning populations have undergone major demographic shifts (Marlino et al., 2004). Educators at all levels need to address their pupils' many different learning styles, broad ranges of abilities, and diverse socioeconomic and cultural backgrounds. Teachers are called on to tailor educational activities for an increasingly heterogeneous student population (Jonassen and Grabowski, 1993).

1.2 ISSUES ADDRESSED BY THIS BOOK

The inflection point will likely produce a rocky revolution in education. Profound innovations generally lead to a sequence of disruptive events as society incorporates them (McArthur et al., 1994). An innovation is typically first used to enhance, enable, or more efficiently accomplish traditional practices (e.g., the car duplicated the functionality of the horse-drawn carriage). Later, the innovation transforms society as it engenders new practices and products, not simply better versions of the original practice. Innovations might require additional expertise, expense, and possibly legislative or political changes (cars required paved roads, parking lots, service stations, and new driving laws). Thus, innovations are often resisted at first, even though they solve important problems in the long term (cars improved transportation over carriages). Similarly, educational innovations are not just fixes or add-ons; they require the educational community to think hard about its mission, organization, and willingness to invest in change.

One proposition of this book is that the inflection point in education is supported by *intelligent educational software* that is opportunistic and responsive. Under the rubric of intelligent educational software, we include a variety of software (e.g., simulations; advisory, reminder, or collaborative systems; or games) that use intelligent techniques to model and reason about learners. One example of this approach, which is based on student-centered rather than teacher-centered strategies, is the *intelligent tutor*.[1] Intelligent tutors contain rich, dynamic models of student knowledge that depict the key ideas learners should understand as well as common learner conceptions and misconceptions. They have embedded models of how students and teachers reason and can adapt their model over time as student understanding becomes increasingly sophisticated (American Association for the Advancement of Science [AAAS], 1993; Corbett and Anderson, 2001; Marlino et al., 2004). They have embedded student models that reason about how people learn, specifically how new knowledge is filtered and integrated into a person's existing cognitive structure (Voss and Silfies, 1996; Yekovich et al., 1990) and reshapes existing structures (Ferstl and Kintsch, 1999). Within intelligent tutors, students move at their own pace, obtain their own knowledge, and engage in self- or group-directed learning.

1.2.1 Computational Issues

The software discussed in this book supports teachers in classrooms and impacts both formal and informal learning environments for people at all levels (K to gray). Creation of a rich and effective education fabric is developed through sophisticated software, AI technology, and seamless education (accessible, mobile, and handheld devices). This book discusses global resources that target computational models and experimentation; it explores the development of software, artificial intelligence, databases, and human-computer interfaces.

Software development. The old model of education in which teachers present students with prepackaged and ready-to-use nuggets of information has had limited impact on children in the past and will have limited success for both

[1] The term *intelligent tutor* describes the engineering result of building tutors. This entity has also been described as knowledge-based tutor, intelligent computer-aided instruction (ICAI), and intelligent tutoring system (ITS).

adults and children in the future. The new educational model is based on understanding human cognition, learning, and interactive styles. Observation of students and teachers in interaction, especially through the Internet, has led to new software development and networks based on new pedagogy. Innovative approaches to education depend on breakthroughs in storing methods and processes about teaching (strategies for presenting topics and rules about how teachers behave). Intelligent tutors use virtual organizations for collaboration and shared control, models and simulations of natural and built complex systems, and interdisciplinary approaches to complexity that help students understand the relevance of learning to daily life. Software responds to student motivation and diversity; it teaches in various contexts (workplace, home, school), for all students (professionals, workers, adults, and children), and addresses many goals (individual, external, grade, or use). Intelligent tutors include test beds for mobile and e-learning, technology-enabled teamwork, wearable and contextual computing, location aware personal digital assistants (PDA), and mobile wireless web-casting.

Artificial intelligence. The artificial intelligence (AI) vision for education is central to this book and characterized by customized teaching. AI tutors work with differently enabled students, make collaboration possible and transparent, and integrate agents that are aware of students' cognitive, affective, and social characteristics. Intelligent agents sense, communicate, measure, and respond appropriately to each student. They might detect learning disability and modify the pace and content of existing pedagogical resources. Agents coach students and scaffold collaboration and learning. They reason about student discussions, argumentations, and dialogue and support students in resolving differences and agreeing on a conclusion. They monitor and coach students based on representations of both content and social issues and reason about the probability of student actions. Probability theory (reinforcement learning, Bayesian networks) defines the likelihood of an event occurring during learning. AI techniques contribute to self-improving tutors, in which tutors evaluate their own teaching.

Databases. The database vision for education includes servers with digital libraries of materials for every school that store what children and teachers create, as well as hold collections from every subject area. The libraries are windows into a repository of content larger than an individual school server can hold. Educational data mining (EDM) explores the unique types of data coming from web-based education. It focuses on algorithms that comb through data of how students work with electronic resources to better understand students and the settings in which they learn. EDM is used to inform design decisions and answer research questions. One project modeled how male and female students differentially navigate problem spaces and suggested strategic problem-solving differences. Another determined that student control (when students select their own problems or stories) increased engagement and thus improved learning.

Human-computer interfaces. New paradigms for interface design minimize the barrier between a student's cognitive model of what he or she wants to accomplish and the computer's understanding of the student's task. The interface is optimized for effective and efficient learning, given a domain and a class of student. New interaction techniques, descriptive and predictive models, and theories of interaction take detailed records of student learning and performance, comment about student activities, and advise about the next instructional material. Formative assessment data on an individual or classwide basis are used to adjust instructional strategies and modify topics.

> *The frequency of computer use [in education] is surprisingly low, with only about 1 in 10 lessons incorporating their use. The explanation for this situation is far more likely lack of teacher preparedness than lack of computer equipment, given that 79% of secondary earth science teachers reported a moderate or substantial need for learning how to use technology in science instruction (versus only 3% of teachers needing computers made available to them).*
>
> **Horizon Research, Inc. (2000)**

1.2.2 Professional Issues

Managing an inflection point in education requires full participation of many stakeholders, including teachers, policy makers, and industry leaders. Changes inevitably produce both constructive and destructive forces (Grove, 1996). With technology, whatever can be done will likely be done. Because technological change cannot be stopped, stakeholders must instead focus on preparing for changes. Educational changes cannot be anticipated by any amount of formal planning. Stakeholders need to prepare, similar to fire department leaders who cannot anticipate where the next fire will be, by shaping an energetic and efficient team capable of responding to the expected as well as to the unanticipated. Understanding the nature of teaching and learning will help ensure that the primary beneficiaries of the impending changes are students. Stakeholders should consider the following major issues:

Teachers as technology leaders. Rather than actively participating in research, teachers are too often marginalized and limited to passively receiving research or technology that has been converted for educational consumption (Marlino et al., 2004). Among K-5 science teachers recently surveyed nationwide, only 1 in 10 reported directly interacting with scientists in professional development activities. For those with such contact, the experience overwhelmingly improved their understanding of needs for the next-generation scientific and educational workforce (National Science Board [NSB], 2003). Historically, large-scale systemic support for science teachers and scientific curricula has increased student interest in science (Seymour, 2002).

Professional development of teachers. A teacher's professional development in technology has been significantly associated with increased student achievement. How teachers use technology is impacted by factors such as their age, computer expertise, length of and access to pertinent training, perceived value of using computers, and views of constructivist beliefs and practices (Maloy et al., in press; Valdez et al., 2000). To strongly influence workforce preparedness, technology must address issues of teacher training, awareness, and general educational infrastructure. Technology is more likely to be used as an effective learning tool when embedded in a broader educational reform, including teacher training, curriculum, student assessment, and school capacity for change (Roschelle et al., 2000).

Hardware issues. A decent benchmark of classroom computers and connectivity suggests one computer for every three students (diSessa, 2000). This metric is achievable as 95% of U.S. schools,[2] and 98% of British schools are connected to the web (National Center for Education Statistics [NCES], 2003; Jervis and Steeg, 2000).

Software issues. Schools need software programs that actively engage students, collaborate with them, provide feedback, and connect them to real-world contexts. The software goal is to develop instructionally sound and flexible environments. Unprincipled software will not work (e.g., boring slides and repetitive pages).

> *Rather than using technology to imitate or supplement conventional classroom-based approaches, exploiting the full potential of next-generation technologies is likely to require fundamental, rather than incremental reform.... Content, teaching, assessment, student-teacher relationships and even the concept of an education and training institution may all need to be rethought ... we cannot afford to leave education and training behind in the technology revolution. But unless something changes, the gap between technology's potential and its use in education and training will only grow as technological change accelerates in the years ahead.*
>
> **Phillip Bond (2004)**

1.3 STATE OF THE ART IN ARTIFICIAL INTELLIGENCE AND EDUCATION

This book describes research, development, and deployment efforts in AI and education designed to address the needs of students with a wide range of abilities, disabilities, intents, backgrounds, and other characteristics. Deployment means using educational software with learners in the targeted venue (e.g., classroom or training department). This section briefly describes the field in terms of its research questions and vision.

1.3.1 Foundations of the Field

The field of artificial intelligence and education is well established, with its own theory, technology, and pedagogy. One of its goals is to develop software that captures

[2] However, only 74% and 39% of classrooms in low-poverty and high-poverty schools, respectively, have Internet access.

the reasoning of teachers and the learning of students. This process begins by representing expert knowledge (e.g., as a collection of heuristic rules) capable of answering questions and solving problems presented to the student. For example, an expert system inside a good *algebra tutor*[3] represents each algebra problem and approximates how the "ideal" student solves those problems (McArthur and Lewis, 1998). Student models, the student systems inside the tutor, examine a student's reasoning, find the exact step at which he or she went astray, diagnose the reasons for the error, and suggest ways to overcome the impasse.

The potential value of intelligent tutors is obvious. Indeed, supplying students with their own automated tutor, capable of finely tailoring learning experiences to students' needs, has long been the holy grail of teaching technology (McArthur and Lewis, 1998). One-on-one tutoring is well documented as the best way to learn (Bloom, 1984), a human-tutor standard nearly matched by intelligent tutors, which have helped to raise students' scores one letter grade or more (Koedinger et al., 1997; VanLehn et al., 2005). Over time, intelligent tutors will become smarter and smarter. Advances in cognitive science will ensure that they capture an increasing share of human-teaching expertise and cover a wider range of subjects (McArthur et al., 1994). However, evidence suggests progress will be slow. Although the speed of computer hardware roughly doubles every two years, the intelligence of computer software, however measured, creeps ahead at a snail's pace.

The field of artificial intelligence and education has many goals. One goal is to match the needs of individual students by providing alternative representations of content, alternative paths through material, and alternative means of interaction. The field moves toward generating highly individualized, pedagogically sound, and accessible lifelong educational material. Another goal is to understand how human emotion influences individual learning differences and the extent to which emotion, cognitive ability, and gender impact learning.

The field is both derivative and innovative. On the one hand, it brings theories and methodologies from related fields such as AI, cognitive science, and education. On the other hand, it generates its own larger research issues and questions (Self, 1988):

- What is the nature of knowledge, and how is it represented?
- How can an individual student be helped to learn?
- Which styles of teaching interaction are effective, and when should they be used?
- What misconceptions do learners have?

In developing answers to some of these questions, the field has adopted a range of theories, such as task analysis, modeling instructional engineering, and cognitive modeling. Although the field has produced numerous tutors, it is not limited to producing functional systems. Research also examines how individual differences and preferred learning styles influence learning outcomes. Teachers who use these tutors

[3]An algebra tutor refers to an intelligent tutor specializing in algebra.

gain insight into students' learning processes, spend more time with individual students, and save time by letting the tutor correct homework.

1.3.2 Visions of the Field

One vision of artificial intelligence and education is to produce a "teacher for every student" or a "community of teachers for every student." This vision includes making learning a social activity, accepting multimodal input from students (handwriting, speech, facial expression, body language) and supporting multiple teaching strategies (collaboration, inquiry, and dialogue).

We present several vignettes of successful intelligent tutors in use. The first is a child reading text from a screen who comes across an unfamiliar word. She speaks it into a microphone and doesn't have to worry about a teacher's disapproval if she says it wrong. The tutor might not interrupt the student, yet at the end of the sentence it provides her the correct pronunciation (Mostow and Beck, 2003).

Now we shift to a military classroom at a United States General Staff Headquarters. This time an officer, being deployed to Iraq, speaks into a microphone, practicing the Iraqi language. He is represented as an avatar, a character in a computer game, and is role-playing, requesting information from local Iraqi inhabitants in a cafe. The officer respectfully greets the Iraqis by placing his right hand over his heart while saying "as-salaamu alaykum." Sometime later he is inadvertently rude and the three avatars representing Iraqi locals jump up and challenge the officer with questions (Johnson et al., 2004).

Now we shift to a classroom at a medical school. First-year students are learning how the barometric (blood pressure) response works. Their conversation with a computer tutor does not involve a microphone or avatar, yet they discuss the qualitative analysis of a cardiophysiological feedback system and the tutor understands their short answers (Freedman and Evens, 1997).

Consider the likely scenarios when such intelligent tutors are available any time, from any place, and on any topic. Student privacy will be critical and a heavily protected portfolio for each student, including grades, learning level, past activities, and special needs will be maintained:

Intelligent tutors know individual student differences. Tutors have knowledge of each student's background, learning style, and current needs and choose multimedia material at the proper teaching level and style. For example, some students solve fraction problems while learning about endangered species; premed students practice fundamental procedures for cardiac arrest; and legal students argue points against a tutor that role-plays as a prosecutor.

Such systems infer student emotion and leverage knowledge to increase performance. They might determine each student's affective state and then respond appropriately to student emotion. Systems recognize a frustrated student (based on facial images, posture detectors, and conductance sensors) and respond in a supportive way with an animated agent that uses appropriate head and body gestures to express caring behavior. Such systems can also recognize bored students (based on slow response and lack of engagement) and suggest more challenging problems.

Intelligent tutors work with students who have various abilities. If a student has dyslexia, the tutor might note that he is disorganized, unable to plan, poorly motivated, and not confident. For students who react well to spoken text messages, natural language techniques simplify the tutor's responses until the student exhibits confidence and sufficient background knowledge. During each interaction, the tutor updates its model of presumed student knowledge and current misconceptions.

Students work independently or in teams. Groups of learners, separated in space and time, collaborate on open-ended problems, generate writing or musical compositions, and are generally in control of their own learning. In team activities, they work with remote partners, explaining their reasoning and offering suggestions. They continue learning as long as they are engaged in productive activities. Teachers easily modify topics, reproduce tutors, at an infinitesimal cost to students and schools and have detailed records of student performance.

Necessary hardware and software. Students work on personal computers or with sophisticated servers managed within a school district. Using high-speed Internet connections, they explore topics in any order and are supported in their different learning styles (e.g., as holists and serialists) (Pask, 1976; Self, 1985). They ask questions (perhaps in spoken language), practice fundamental skills, and move to new topics based on their interests and abilities. Tutors generate natural language responses. Metacognitive strategies identify each student's learning strengths (e.g., the student requests hints and knows how to self-explain new topics).

Intelligent tutors know how to teach. Academic material stored in intelligent systems is not just data about a topic (i.e., questions and answers about facts and procedures). Rather, such software contains qualitative models of each domain to be taught, including objects and processes that characterize trends and causal relations among topics. Each model also reasons about knowledge in the domain, follows a student's reasoning about that knowledge, engages in discussions, and answers questions on various topics. New tutors are easily built and added onto existing tutors, thus augmenting a system's teaching ability. Tutors store teaching methods and processes (e.g., strategies for presenting topics, feedback, and assessment). This knowledge contains rules about how outstanding teachers behave and teaching strategies suggested by learning theories.

These scenarios describe visions of fully developed, intelligent instructional software. Every feature described above exists in existing intelligent tutors. Some tutors are used in classrooms in several instructional forms (simulations, games, open-learning environments), teaching concepts and procedures from several disciplines (physics, cardiac disease, art history).

These educational scenarios are not just fixes or add-ons to education. They may challenge and possibly threaten existing teaching and learning practices by suggesting new ways to learn and offering new support for students to acquire knowledge (McArthur et al., 1994). Technology provides individualized attention and augments a teacher's ability to respond. It helps lifelong learners who are daily called on to integrate and absorb vast amounts of knowledge and to communicate with multitudes of people. The educational community needs to think hard about its mission and its organization:

- *School structure.* What happens to school structures (temporal and physical) once students choose what and when to study and work on projects by themselves or with remote teammates independent of time and physical structure?
- *Teachers and administrators.* How do teachers and administrators react when their role changes from that of lecturer/source to coach/guide?
- *Classrooms.* What happens to lectures and structured classrooms when teachers and students freely select online modules? What is the impact once teachers reproduce tutors at will and at infinitesimal cost?
- *Student privacy.* How can students' privacy be protected once records (academic and emotional) are maintained and available over the Internet?

> *We are not going to succeed [in education] unless we really turn the problem … around and first specify the kinds of things students ought to be doing: what are the cost-effective and time-effective ways by which students can proceed to learn. We need to carry out the analysis that is required to understand what they have to do—what activities will produce the learning—and then ask ourselves how the technology can help us do that.*
>
> **Herbert A. Simon (1997)**

1.3.3 Effective Teaching Methods

For hundreds of years, the predominant forms of teaching have included books, classrooms, and lectures. Scholars and teachers present information carefully organized into digestible packages; passive students receive this information and work in isolation to learn from fixed assignments stored in old curricula. These passive methods suggest that a student's task is to absorb explicit concepts and exhibit this understanding in largely factual and definition-based multiple-choice examinations. In this approach, teachers in the classroom typically ask 95% of the questions, requiring short answers or problem-solving activities (Graesser and Person, 1994; Hmlo-Silver, 2002).

These teaching approaches are not very effective (Waterman et al., 1993); they have succeeded only with the top fourth of each class, often motivated and gifted students. Student achievement in classroom instruction (1:30 teacher/student ratio) was found to differ from achievement based on individual tutoring (1:1 teacher/student ratio) by about two standard deviations (Bloom, 1984). That is, the typical bell curve of achievement was centered on the 50th percentile for traditional lecture-based teaching and was raised to the 98th percentile by one-to-one human teachers (see Figure 1.1).

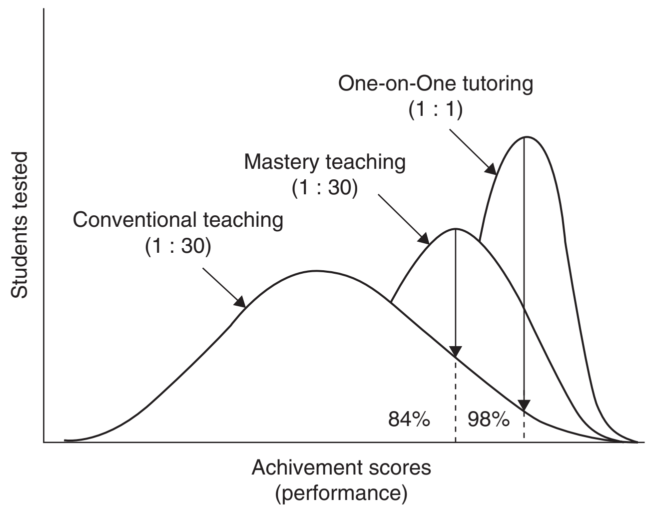

FIGURE 1.1

Advantages of one-to-one tutoring. (Adapted from Bloom, 1984.)

Other effective teaching methods (e.g., *collaboration, inquiry*, and *teaching meta-cognition*) actively engage students (including the disadvantaged, financially insecure, and unmotivated) to create their own learning. However, these methods are nearly impossible to implement in classrooms without technology, as they are so time and resource intensive. For example, one-to-one tutoring (adapting teaching to each learner's needs) requires one teacher for each student (Bloom, 1984). Collaboration (facilitating students to work in groups and explain their work to each other) often results in students learning more than the best student in the group but requires individual attention for each group of one to three students. Inquiry learning (supporting students to ask their own questions, generate research hypotheses, and collect data) is powerful because students are engaged in authentic and active work and use information in a variety of ways. However, inquiry learning requires teachers to guide students in asking their own questions and gathering and analyzing evidence. Teaching about *metacognitive skills* (students focus on their own learning approaches, e.g., asking for hints, self-explanation) sometimes results in more effective learning, yet requires teachers to individually guide each student. Most schools in most disciplines cannot provide this individual attention, although many nations support one teacher for every student in high-risk professions (airplane pilots or controlling nuclear reactors) or in the arts (music or painting).

One example of ineffective teaching methods is the tradition of only transmitting facts to students. Understanding the components and data of a discipline is not as effective as understanding its structure. This distinction is particularly true in fields such as science, mathematics, and engineering, where students need to know the processes by which the discipline's claims are generated, evaluated, and revised.

Information technology is effective in teaching and improves productivity in industry and the military. Intelligent tutors produce the same improvements as one-to-one tutoring and effectively reduce learning time by one-third to one-half (Regian et al., 1996). Recall that one-to-one human tutoring increases classroom performance to around the 98th percentile (Bloom, 1984). Intelligent tutors are 30% more effective than traditional instruction (Fletcher, 1995; Regian et al., 1996), and networked versions reduce the need for training support personnel by about 70% and operating costs by about 92%.

1.3.4 Computers in Education

Computers have been used in education since 1959 when PLATO was created at the University of Illinois (Molnar, 1990; Office of Technology Assessment [OTA], 1982). This several thousand–terminal system served elementary school, undergraduate, and community college students. In 1963, another system used a drill-and-practice, self-paced program in mathematics and reading, thus allowing students to take a more active role in the learning process (Suppes, 1981).

The programming language LOGO was developed in the early 1970s to encourage students to think rigorously about mathematics, not by teaching facts and rules but by supporting the use of mathematics to build meaningful products, such as drawings and processes (Papert, 1980). Because LOGO was user-friendly, students could easily express procedures for simple tasks. It was used in various "microworld" environments, including robotic building sets (Lego Mindstorms) that could be used to invent robotics solutions, trucks, spaceships, and mobile artifacts. In building computer-driven LOGO inventions, students defined a problem and developed the skills needed to solve it.

Other engaging uses of computers in education involved project-oriented, case-based, and inquiry-oriented education. For example, the National Geographic Kids Network invited students to measure the quality of their regional water and its relationship to acid rain (Tinker, 1997). Students in more than 10,000 elementary schools at 80 sites in 30 countries gathered data, analyzed trends, and communicated by e-mail with each other and with practicing scientists. Student results were combined with national and international results, leading to the discovery that school drinking water and air pollution standards were not being met. The network provided low-cost devices for measuring ozone, soil moisture, and ultraviolet radiation to calibrate the effects of global warming. In 1991, students measured air and soil temperatures, precipitation, bird and insect presence, and stages of plant growth, thus linking meteorological, physical, and biological observations to a major seasonal event and creating a "snapshot" of the planet. Most teachers using KidsNet (>90%) reported that it significantly increased students' interest in science and that their classes spent almost twice as much time on science than before.

In another example, the Learning Network was used by 35,000 kids from 23 nations to explore social studies and literature issues. It was run by AT&T and managed by state and grassroots organizations.

These projects have powerful significance. Networks permit all students to participate in experiments on socially significant scientific problems and to work with real scientists (Molnar, 1990). Students create maps of a holistic phenomenon drawn from a mosaic of local measurements. Teachers' roles change, and they act as consultants rather than as leaders.

Computer-based education has been well documented to improve learning at the elementary, secondary, higher-, and adult-education levels. A meta-analysis of several hundred well-controlled studies showed that student scores increased by 10% to 20%, the time to achieve goals decreased by one-third, and class performance improved by about one-half standard deviation (Kulik and Kulik, 1991).

However, these early computer-based instructional systems had several drawbacks. Many systems used *frame-based* methods, in which every page, computer response, and sequence of topics was predefined by the author and presented to students in lockstep fashion. Directed learning environments, including tutorials, hypermedia, and tests, typically presented material in careful sequences to elicit correct learner action (Alessi and Trollip, 2000).

In some systems, computer responses were similar for every student, no matter the student's performance, and help was provided as a preworded, noncustomized response. For each student and every situation, "optimal" learning sequences were built in. This approach is similar to playing cops and robbers with predefined paths for chasing robbers. No matter what the robber does, the law enforcer runs down a preset list of streets and crosses specified corners. This model has limited impact; it clearly fails to capture the one-on-one approach of master human teachers who remain opportunistic, dynamically changing topics and teaching methods based on student progress and performance.

Nonetheless, many educational simulations were clearly effective. They allowed students to enter new parameters, watch changing features, start or stop simulations, or change the levels of difficulty, as exemplified by SimCity and SimArt (released by Electronic Arts in 1998) and BioLab Frog (released by Pierian Spring Software in 2000). However, if a student's concept of the modeled interaction differed from that of the author, the student could not ask questions, unless those questions were already programmed into the environment. Students received preformatted responses independent of their current situation or knowledge. They watched the simulation, but typically could not change its nature or learn why the simulation worked as it did.

Other early systems used open-ended learning environments (OLE) that allowed students to experiment, interpret, and learn from errors and to revise their knowledge; these included Jasper Woodbury (Vanderbuilt, Cognition and Technology Group, 1997), EarthTrails (Iowa Public Television, 1997), and Geometer's Sketchpad (Key Curriculum Press, 1995). These systems often provided excellent tools (drawing, measuring, calculating), but many OLEs were very directed and not adaptable. The system pursued a list of problems or cases in a fixed order, and students could not ask questions or alter the program's order. Using these systems required classroom teachers to be trained and remain as central guides during the classroom experience. Teachers divided tasks among collaborative learners and decided which steps each group would tackle and which parameters to enter. OLEs such as Rainforest Researchers or Geography Search (Tom Snyder Productions, 1998, 1995) supported team activities, but did not interact individually with students to help them manage the environment. Neither did they support group creation, group dynamics, role-playing, or planning the next strategy.

1.3.5 Intelligent Tutors: The Formative Years

The field of artificial intelligence and education was established in the 1970s by a dozen leaders, including John Self (1974, 1977, 1985), Jaime Carbonell (1970a, 1970b), and William Clancey (1979). The earliest intelligent tutor was implemented in the 1970 Ph.D. thesis of Jaime Carbonell, who developed Scholar, a system that invited students to explore geographical features of South America. This system differed from traditional computer-based instruction in that it generated individual responses to students' statements by traversing a semantic network of geography knowledge.

The first intelligent tutor based on an expert system was GUIDON developed by William Clancey (Clancey, 1979, 1987). This system was named GUIDON, was also the first to teach medical knowledge (see Section 3.5.2.2). Another knowledge representation, NEOMYCIN, was later designed for use in GUIDON 2 (Clancey and Letsinger, 1981). The GUIDON project became relevant in developing future medical tutors (Crowley et al., 2003) because of key insights: the need to represent implicit knowledge, and the challenges of creating a knowledge representation sufficiently large, complex, and valid to help students learn real medical tasks.

In 1988, Claude Frasson at the University of Montreal, Canada, organized the first conference of the field. The International Conference of Intelligent Tutoring Systems (ITS) provided a forum for researchers and practitioners to exchange ideas, experiments, and techniques in all areas of computer science and human learning. These ITS conferences continued every few years for 20 years under the leadership of Claude Frasson. The first conference of the fledgling field of artificial intelligence and education (AIED), AIED93, was held in Edinburgh, United Kingdom, with Helen Pain as the organizing committee chair. AIED95 was held in Washington, with Jim Greer as the program committee chair, and AIED97, in Osaka directed by Riichiro Mizoguchi. The conference goals are to advance research and development; to support a community from computer science, education, and psychology; and to promote the rigorous research and development of interactive and adaptive learning environments. The *International Journal of Artificial Intelligence and Education (IJAIED)* is the official journal of the AIED Society and contains peered-reviewed journal papers.

1.4 OVERVIEW OF THE BOOK

This book discusses the theory, pedagogy, and technology of intelligent tutors in three parts: "Introduction to Artificial Intelligence and Education," "Representation, Reasoning and Assessment," and "Technologies and Environments."

The first part identifies features of intelligent tutors and includes a framework for exploring the field. Tools and methods for encoding a vast amount of knowledge are described. The term *intelligent tutor* is not just a marketing slogan for conventional computer-assisted instruction but designates technology-based instruction with qualitatively different and improved features of computer-aided instruction.

The second part describes representation issues and various control mechanisms that enable tutors to reason effectively. Tutors encode knowledge about *student* and *domain knowledge, tutoring strategies,* and *communication*. They reason about which teaching styles are most effective in which context.

The third part, extends the narrow range of intelligent tutors and demonstrates their effectiveness in a broad range of applications. For example*, machine learning* enables tutors to reason about uncertainty and to improve their performance based on observed student behavior. Machine learning is used, in part, to reduce the cost per student taught, to decrease development time, and to broaden the range of users for a given tutor. *Collaborative environments* are multiuser environments that mediate learning by using shared workspaces, chat boxes, servers, and modifiable artifacts (e.g., charts, graphs). *Web-based tutors* explore pedagogical and technical issues associated with producing tutors for the web. Such issues include intelligence, adaptability, and development and deployment issues.

In discussing the field, we use a *layered approach* to enable readers to choose a light coverage or deeper consideration. Layers include sections on *what, how*, and *why*:

- The *what* layer defines the current *concept or topic* and serves as a friendly introduction. This level is for readers who seek a cursory description (students, teachers, and administrators).
- The *how* layer explains at a deeper level how this concept or topic works and how it can be implemented.
- The *why* layer describes why this concept or topic is necessary. This layer, which involves theory, is mainly intended for researchers but may interest developers or those who want to know contributing theories and controversies.

SUMMARY

This chapter argued that the rapid rate of change in education, artificial intelligence, cognitive science, and the web has produced an inflection point in educational activities. Information technology clearly narrows the distance among people worldwide; every person is on the verge of becoming both a teacher and a learner to every other person. This technology has the potential to change the fundamental process of education. Managing this inflection point requires that all stakeholders fully participate to ensure that the coming changes benefit not organizations, but students.

This chapter identified specific features that enable intelligent tutors to reason about *what, when*, and *how* to teach. Technology might enhance, though not replace, one-to-one human tutoring, thus extending teaching and learning methods not typically available in traditional classrooms (e.g., collaborative and inquiry learning).

Also discussed were issues to capitalize on the flexibility, impartiality, and patience of intelligent tutors. This technology has the potential to produce highly individualized, pedagogically sound, and accessible educational material as well as match the needs of individual students (e.g., underrepresented minorities and disabled students) and to involve more students in effective learning. Because such systems are sensitive to individual differences, they might unveil the extent to which students of different gender, cognitive abilities, and learning styles learn with different forms of teaching.

The implication is that personalized tutoring can be made widely and inexpensively available just as printed materials and books were made widely available and inexpensive by the invention of the printing press. Like all educational changes, this revolution will proceed slowly but likely faster than the previous ones (e.g., the printing press). Given this age of rapidly changing technology and Internet support of meaningful interactions, intelligent tutors have the potential to provide a skilled teacher, or community of teachers, for every student, anywhere, at any moment.