Running head: INTELLIGENT TUTORING SYSTEMS

Intelligent Tutoring Systems

Arthur C. Graesser, Mark W. Conley, and Andrew Olney

University of Memphis

Send correspondence to:

Art Graesser
Department of Psychology & Institute for Intelligent Systems
202 Psychology Building
University of Memphis
Memphis, TN 38152-3230
901-678-4857
901-678-2579 (fax)
a-graesser@memphis.edu

Intelligent Tutoring Systems (ITS) are computerized learning environments that incorporate computational models from the cognitive sciences, learning sciences, computational linguistics, artificial intelligence, mathematics, and other fields. An ITS tracks the psychological states of learners in fine detail, a process called student modelling. The psychological states may include subject matter knowledge, skills, strategies, motivation, emotions, and other student attributes. An ITS adaptively responds with activities that are both sensitive to these states and that advance its instructional agenda. The interaction between student and computer evolves in a flexible fashion that caters to the constraints of both the student and the instructional agenda. This is a marked departure from a book or a lecture, which unfold in a rigid linear order and are not tailored to individual students.

The history of ITS was christened with the edited book by Sleeman and Brown (1982) entitled *Intelligent Tutoring Systems*. The contributors to the book were an interdisciplinary mixture of researchers in artificial intelligence, cognitive science, and education. There was the vision that learning could substantially improve by developing adaptive learning environments that took advantage of the latest advances in intelligent systems. Some of these ITS attempted to model characteristics of human tutoring whereas others adopted ideal, rational models of intelligent systems. By the early 1990's there were two conferences that directly focused on ITS development and testing: Intelligent Tutoring Systems and Artificial Intelligence in Education. A monograph by Woolf (2009), *Building Intelligent Interactive Tutors*, describes ITS architectures and many of the landmark contributions during the 30 history of the field.

Advances in ITS have progressed to the point of being used in school systems. One noteworthy example is the $Cognitive\ Tutor_{TM}$ developed by Carnegie Learning (Anderson, Corbett,

Koedinger, & Pelletier, 1995; Koedinger, Anderson, Hadley, & Mark, 1997; Ritter, Koedinger, Anderson, & Corbett, 2007). The Cognitive Tutor helps students learn algebra, geometry, and programming languages by applying learning principles inspired by a general model of cognition called ACT-R (Anderson, 1990, 2007). There is a textbook and curriculum to provide the content and the context of learning, but the salient contribution of the Cognitive Tutor is to help students solve problems. The Cognitive Tutor is now used in over 2000 school systems throughout the country and is among the methods accepted in the *What Works Clearinghouse*. According to Ritter, Anderson, Koedinger, and Corbett (2007), standardized tests show improvements over suitable control conditions and are particularly effective for the more challenging subcomponents of problem solving and use of multiple representations.

This chapter reviews the research on different classes of ITS. This includes a description of the computational mechanisms of each type of ITS and available empirical assessments of the impact of these systems on learning gains. However, before diving into the details of these major ITS advances, it is important to give some highlights on what we know about human tutoring, including the practice of human tutoring, relevant pedagogical theories, and empirical evidence for the effectiveness of human tutoring. It is important to take stock of the activities of human tutors because some ITS are built with the guidance of what human tutors know about effective pedagogy and because comparisons are sometimes made between human and computer tutors. The final section identifies some future directions for the ITS field to pursue.

What Do We Know about Human Tutoring?

Tutoring is the typical solution that students, parents, teachers, principles and school systems turn to when the students are not achieving expected grades and educational standards.

Stakeholders in the United States are worried when a school is underperforming according to the

criteria and laws adopted by the individual states or by the Federal government, such as *No Child Left Behind* in the Bush Administration and *Race to the Top* under President Obama. There are salient implications for teacher employment and salaries in addition to the more abstract concerns about the students' educational needs and the global image of the community. Human tutors step in to help under these conditions.

This section analyzes human tutoring from the perspective of three questions. First, what is the nature of human tutoring in actual schools systems and the community? Second, what do the research studies reveal about the effectiveness of human tutoring on learning gains? Third, what are some common tutoring patterns when human tutors tutor?

Human Tutoring in the Wild

Human tutoring consists of one-to-one educational interventions with a human tutor. There are many different versions or purposes for human tutoring, including assignment assistance, instructional tutoring, and strategic tutoring (Hock, Pulvers, Deshler, & Schumaker, 2001). Instructional tutoring consists of direct instruction in such skills as literacy, mathematics and writing. Assignment assistance is supplemental to basic instruction and involves a tutor working with students who have difficulty completing course assignments. In strategic tutoring, students are taught to "learn how to learn" while they engage in classroom instruction and complete course assignments. Strategic tutoring combines instructional tutoring and assignment tutoring, with a focus on a particular learning strategy such as problem solving, summarization or questioning.

There can be a great deal of disagreement about who does the tutoring, the focus of the tutoring with respect to knowledge and skill, and even the intended clients or recipients of the tutoring. Under ideal conditions, human tutors are knowledgeable or skilled teachers or other

adults who provide one-to-one support for students who need to improve in some academic area (Conley, Kerner, & Reynolds, 2005; Hock, et al., 2001; Pressley & McCormick, 1995; Tollefson, 1997). Sometimes human tutors who are not expert and yet tutor in academic areas can do more harm than good. For example, special education teachers without specific academic expertise in areas of tutoring fail to help students develop the knowledge and skills to become independent thinkers and learners in those areas (Carlson, 1985). Novice human tutors mayfail to effectively engage their students, may ignore important components of tutoring such as modeling or providing effective feedback, and may abandon research-based instruction in favor of personal views grounded in their own experiences with schooling, all despite the most intense training and monitoring (Hock, Schumaker, & Deshler, 1995)

So-called professional tutoring programs respond to the need for expert tutors by using certified teachers or by putting the tutors through intensive, specialized training (Slavin, Karweit, & Madden, 1989; Wasik, 1998; Wasik & Slavin, 1993). Wasik and Slavin's (1993) comparative review of tutoring programs that used varying kinds of skilled personnel and differing kinds of training (including volunteers) revealed that the greatest impact on reading achievement comes from the use of certified professional teachers trained in specific tutoring interventions.

Unfortunately, the most effective human tutoring program also was the most costly tutoring program. It is not uncommon in the research literature to see cost analyses accompanying research on human tutoring, with a similar finding reinforced each time: effective human tutoring requires expertise, but that expertise is almost always a formidable expense.

There is an obvious trade-off between the expertise of the tutor and the cost of professional human tutoring programs. In light of this tradeoff, there is a wide range of tutoring programs that employ community volunteers, college students, and student peers. To

compensate for the lack of innate expertise with many human tutors, researchers embed tutoring scripts or lesson plans in the tutoring, provide extensive or scaffolded tutor training, or use a variety of multimedia resources to cue the tutors to do the right things at the right time. One of the most successful versions of this approach is the Charlottesville Volunteer Tutorial, popularly known as Book Buddies (Invernizzi, Rosemary, Juel, & Richards, 1997). Book Buddies is a joint effort from the University of Virginia and the Charlottesville City Schools. Tutors in Book Buddies are community volunteers who tutor first grade students in emergent and early literacy. Tutors are provided a tutoring manual and are supervised in local schools by expert teachers and site coordinators. The content of the tutoring focuses on early reading skills including alphabetic knowledge, phonemic awareness, word recognition and fluency. The scripted nature of the program, along with the site-based supervision, has resulted in considerable consistency among the tutors in their implementation of the program. Researchers found that a minimum of 30 onehour tutoring sessions was necessary to see gains in alphabetic knowledge, phoneme-grapheme knowledge, word recognition and concept of a word. The found that 40 sessions yielded an even greater impact. The researchers compared the cost of the volunteer tutoring program with other school-based funded tutoring programs such as *Reading Recovery*. The volunteer tutoring program costs (e.g., program coordinator, on-site supervision of tutors, training manuals) was one-sixth the cost of school based funded programs whereas effect sizes for word recognition in Book Buddies far outpaced effect sizes using professionally trained teachers in more expensive programs. Proponents cite these factors (i.e., low cost and high impact) as key reasons to use volunteer tutors working under carefully structured and scripted conditions rather than professional tutors who are expensive to train.

One limitation of many human tutoring programs concerns their relatively narrow focus. Book Buddies, for example, involves tutoring with several hundred first graders at a time (Invernizzi, et al., 1997). Strategic tutoring, which employs highly trained tutors working with special education children, often engages 30 or fewer children at a time (Hock, et al., 2001; Vellutino, et al., 1996). In one notable example in which Book Buddies was scaled up and implemented in a more complex school setting, the Bronx, tutors were unable to move students to a reading level necessary for them to succeed in school, compared with non-participating peers (Meier & Invernizzi, 2001). In a recent review of volunteer tutoring programs for elementary and middle school children, 21 studies using 28 different cohorts of students revealed positive effects on achievement, but only with very basic sub-skills in reading, such as phonemic awareness, word recognition and fluency (Ritter, Barnett, Denny, & Albin, 2009). Noteworthy in this meta-analysis was the finding that human tutoring programs, as they were construed, did not have a significant impact on comprehension or other academic areas like mathematics. That is, there was little if any evidence of transfer of skill from human tutoring in basic literacy skills to other academic areas like mathematics or science. In summary, for all of the expense and effort that often goes into human tutoring, both professional and volunteer, the outcomes have been promising but extremely limited in the population served and overall impact. These findings on tutoring in the wild have motivated many to seriously consider ITS and other computer-based technologies as an alternative form of one-on-one tutoring.

Learning Gains with Human Tutoring

The studies cited above present a mixed picture of the impact of human tutoring on learning.

One of the landmark early assessments of human tutoring was the meta-analysis conducted by

Cohen, Kulik, and Kulik (1982) who reported learning gains of approximately 0.4 sigma when compared to classroom controls and other suitable controls. A sigma is a measure in standard deviation units that compares a mean in the experimental treatment to the mean in a comparison condition. According to Cohen (1992), effect sizes of 0.20, 0.50, and 0.80 are considered small, medium, and large, respectively, so the 0.4 sigma is considered a medium effect size.

Many of the tutors in the documented research on tutoring include unskilled tutors. Unskilled tutors are defined in this chapter as tutors who are not experts on subject matter knowledge, are not trained systematically on tutoring skills, are not certified, and are rarely evaluated on their impact on student learning. Unskilled tutors are paraprofessionals, parents, community citizens, cross-age tutors, or same-age peers. In spite of their lack of training and skill, they nevertheless improve learning under some conditions. There are many possible explanations of these learning gains from tutors who are unskilled. Perhaps they have sufficient intuitions to detect deficits in the student and can make sensible recommendations on how the student can improve. Perhaps there is something about one-on-one conversation with another human that helps motivation and learning (Graesser, Person, & Magliano, 1995). Available evidence suggests that the expertise of the tutor does matter, but the evidence is not strong. For example, collaborative peer tutoring shows an effect size advantage of 0.2 to 0.9 sigma (Johnson & Johnson, 1992; Mathes & Fuchs, 1994; Slavin, 1990; Topping, 1996), which is slightly lower than older unskilled human tutors. Peer tutoring is a low-cost effective solution because expert tutors are expensive and hard to find. Unfortunately, systematic studies on learning gains from expert tutors are few in number because the tutors are expensive, they are difficult to recruit in research projects, and they tend to stay in the tutoring profession for a short amount of time (Person, Lehman, & Ozbun, 2007). However, available studies report effect sizes of 0.8 to 2.0 (Bloom, 1984; Chi, Roy, & Hausmann,

2008; VanLehn et al., 2007), which is presumably higher than other forms of tutoring. Nevertheless, there clearly needs to be more research on the impact of tutoring expertise on student learning in mathematics, reading literacy, writing, science, and self-regulated learning strategies.

The impact of tutoring on learning is not determined entirely by the tutor, of course, because there are also constraints of the student and the student-tutor interaction. Many theories of learning have emphasized the importance of getting the student to generate the knowledge rather that expecting an instruction delivery system to provide all of the information (Bransford, Brown & Cocking, 2000; Mayer, 2009; O'Donnell, Volume 1). Students learn by expressing, doing, explaining, and being responsible for their knowledge construction, as opposed to being passive recipients of exposure to information. We know, for example, that the tutors in the same-age and cross-age collaborations tend to learn more than the tutees, even when they start with essentially the same level of mastery (Cohen et al., 1982; Mathes & Fuchs, 1994; Rohrbeck, Ginsburg-Block, Fantuzzo, & Miller, 2003).

The obvious question that learning scientists have been asking over the years is why tutoring is effective in promoting learning? We believe that the best approach to answering this question is to analyze the tutoring process and to explore what it is about the process that lends itself to learning (see Graesser, D'Mello, & Cade, in press). But what are these processes? A number of studies have performed a very detailed analysis of the tutoring session structure, tasks, curriculum content, discourse, actions, and cognitive activities manifested in the sessions and to speculate how these might account for the advantages of tutoring (Chi, Roy, & Hausmann, 2008; Chi et al., 2001; Graesser & Person, 1994; Graesser, Person, & Magliano, 1995; Hacker & Graesser, 2007; Lepper, Drake, & O'Donnell-Johnson, 1997; McArthur, Stasz, & Zmuidzinas, 1990; Merrill, Reiser, Merrill, & Landes, 1995; Person & Graesser, 1999; 2003; Person, Kreuz, Zwaan, & Graesser,

1995; Shah, Evens, Michael, & Rovick, 2002; VanLehn et al., 2003). We believe these detailed analyses have considerable promise. At the same time, however, these observation-based analyses fall short of an experimental approach that assesses causal relations between tutoring practice and learning. One needs to manipulate the tutoring activities through trained human tutors or computer tutors and to observe the impact of the manipulations on learning gains (Chi et al., 2001, 2008; Graesser, Lu et al., 2004; Litman et al., 2006; VanLehn et al., 2003, 2007). Manipulation studies allow us to infer what characteristics of the tutoring directly cause increases in learning gains, barring potential confounding variables.

Tutoring Strategies and Processes

The studies referenced have documented the strategies and processes of tutoring in diverse areas and in tutors with diverse expertise. For example, Graesser and Person analyzed the discourse patterns of 13 unskilled tutors in great detail after analyzing over 100 hours of tutoring sessions (Graesser, & Person, 1994; Graesser et al., 1995; Person & Graesser, 1999). Person et al.'s (2007) meta-analysis of accomplished tutors revealed that the sample sizes of expert tutors have been extremely small (N <3) in empirical investigations of expert tutoring and often the same expert tutors are used in different research studies. To address these deficits, Person et al. (2007) analyzed 12 tutors in 40 hours of tutoring sessions across domains of science and mathematics.

Unfortunately, these studies that analyzed the tutoring process in detail did not have learning outcome measures. There is a large void in the literature on detailed analyses of human tutorial dialogue that are related to outcome measures and that have a large sample of tutors. There are also practical limitations that present serious obstacles to collecting such data. The subject matters of the tutoring sessions are difficult to predict in advance so it is difficult to proactively identify suitable pre-test and post-test measures from normative testbanks. There is also a large attrition rate for both

tutors and students. Nevertheless, these tutoring corpora can be analyzed to identify the tutoring processes.

A number of strategies are not implemented very often by either expert or novice tutors according to the above studies that have investigated the tutoring process. According to systematic analyses of the tutoring process (Graesser et al., 1995; Graesser, D'Mello, & Person, 2009; Person et al., 1995), human tutors are not prone to implement sophisticated tutoring strategies that have been proposed in the fields of education, the learning sciences, and developers of ITS. In particular, tutors rarely implement the following pedagogical techniques:

Socratic tutoring. The tutor asks learners illuminating questions that lead the learners to discover and correct their own misconceptions in an active, self-regulated fashion (Collins et al., 1975).

Modeling-scaffolding-fading. The tutor first models a desired skill. Then the tutor somehow gets the learners to perform the skill while the tutor provides feedback, explanation, and other forms of scaffolding. Finally, the tutor fades from the process until the learners perform the skill all by themselves (Rogoff & Gardner, 1984).

Reciprocal Teaching. The tutor and learner take turns working on problems or performing a skill, with the other person giving feedback, guidance, and advice along the way (Palincsar & Brown, 1984).

Frontier learning. The tutor selects problems and gives guidance in a fashion that slightly extends the boundaries of what the learner already knows or has mastered (Sleeman & Brown, 1982).

Building on prerequisites. The tutor covers the prerequisite concepts or skills in a session before moving to more complex problems and tasks that require mastery of the prerequisites (Gagne, 1985).

Diagnosis/remediation of deep misconceptions. The tutor identifies fundamental and persistent misconceptions of the learner. Once identified, the tutor diagnoses the causes of each misconception and implements methods to repair it (Lesgold et al., 1992).

According to Graesser, D'Mello, and Person (2009), tutors do not have a deep understanding of what the student knows, i.e., the student model, so their responses run the risk of not being on the mark in advancing learning. Tutors have only an approximate understanding of the students' profile so there is an inherent limitation in how well they can adapt (Graesser, D'Mello, & Person, 2009). This opens to door for an ITS to prove its metal. An ITS has the potential to show improvements over human tutors to the extent that it can perform more accurate student modeling and can intelligently implement a broader array of tutoring strategies. Graesser and Person's analyses of tutorial dialogue uncovered a number of frequent dialogue structures that presumably have some potential in advancing learning (Graesser & Person, 1994; Graesser et al., 1995; Graesser, Hu, & McNamara, 2005). These structures are also prominent in the work of other researchers who have conducted fine-grained analyses of tutoring (Chi et al., 2004; Graesser, Hu, & McNamara, 2005; Litman et al., 2006; Shah et al., 2002). The following dialogue structures are prominent: (a) a curriculum script with didactic content and problems (i.e., difficult tasks or questions), (b) a 5-step Tutoring Frame, (c) Expectation and Misconception Tailored (EMT) dialogue, and (d) Conversational Turn Management.

(a) Curriculum script. The tutor tries to cover a curriculum with substantive content and a set of problems that address the content. The content can be presented in a mini-lecture,

hopefully at the right time for the individual learner. The problems are tasks, activities, and difficult questions that reveal whether the student is understanding. Tutors tend to present the mini-lectures and problems in a rigid way rather than tailoring to the individual learner.

- (b) 5-Step Tutoring Frame. After a problem or difficult main question is selected to work on, the problem is solved through an interaction that is structured by a 5-Step Tutoring Frame. The 5 steps are: (1) Tutor presents problem, (2) Student gives an initial answer, (3) Tutor gives short feedback on the quality of the Student's initial answer, (4) the Tutor and Student collaboratively improve on the answer in a turn-by-turn dialogue that may be lengthy, and (5) the Tutor evaluates whether the Student understands (e.g., asking "Do you understand?" or testing with a follow-up task). This 5-step tutoring frame involves collaborative discussion, joint action, and encouragement for the student to construct knowledge rather than merely receiving knowledge.
- (c) Expectation and Misconception Tailored (EMT) Dialogue. Human tutors typically have a list of expectations (anticipated good answers, steps in a procedure) and a list of anticipated *misconceptions* associated with each main question. They want this content covered in order to handle the problem that is posted. The tutor guides the student in articulating the expectations through a number of dialogue moves, namely *pumps* ("What else"), *hints*, *questions* to extract specific information from students, and *answers* to students' questions. The correct answers are eventually covered and the misconceptions are hopefully corrected.
- (d) **Conversational Turn Management.** Human tutors structure their conversational turns systematically. Nearly every turn of the tutor has three information slots (i.e., units or constituents). The first slot of most turns is feedback on the quality of the learner's last turn. This feedback is either positive (*very good, yeah*), neutral (*uh huh, I see*), or negative (*not quite, not really*). The second slot advances the interaction with either prompts for specific information,

hints, assertions with correct information, corrections of misconceptions, or answers to student questions. The third slot is a cue for the floor to shift from the tutor as the speaker to the learner. For example, the human ends each turn with a question or a gesture to cue the learner to do the talking.

Student Modeling. One of the central questions is how well the tutor can track the psychological states of the student as they implement these strategies. Available evidence suggests that human tutors are not able to conduct student modeling at a fine-grained level (Chi et al., 2004; Graesser et al., 2009). They can perform approximate assessments, but not fine-grain assessments. This of course limits how well humans can adaptively respond. Moreover, there is the possibility that computers might show advantages over humans to the extent that computers can accurately conduct student modeling and generate intelligent responses.

Consider, for example, how well human tutors can track knowledge states of learners. Graesser, D'Mello, and Person (2009) documented some of the illusions that typical human tutors have about cognition and communication. These illusions may limit optimal learning. According to the *illusion of grounding*, the tutor mistakenly believes that the speaker and listener have shared knowledge about a word, referent, or idea being discussed in the tutoring session. A tutor should be skeptical of the student's level of understanding and trouble-shoot potential communication breakdowns in common ground between the tutor and student. According to the *illusion of feedback accuracy*, the tutoring mistakenly believes that the feedback is accurate when the student indicates their understanding of the subject matter. For example, tutors incorrectly believe the students' "Yes" answers to their comprehension gauging questions (e.g., Do you understand?). According to the *illusion of student mastery*, the tutor believes that the student has mastered much more than the student has really mastered. For example, the fact that a student

expresses a word or phrase does not mean that the student understands an underlying complex idea. Yet the tutor believes that the students' expression of a word or two is a sufficient proxy for mastering a complex conceptualization. According to the *illusion of knowledge transfer*, the tutor believes that the information they express is accurately transmitted to the mind of the student. In fact, the student absorbs very little. These illusions threaten the fidelity of student modelling in human tutoring, so there is the fundamental question of how useful their dialogue moves will be in advancing learning. A more realistic picture is that the tutor has only an approximate appraisal of the cognitive states of students and that they formulate responses that do not require fine-tuning of the student model (Chi et al., 2004; Graesser et al., 1995).

Another dimension of student modeling consists of student emotions and motivation. Indeed, connections between complex learning and emotions have received increasing attention in the fields of psychology and education (Deci & Ryan, 2002; Dweck, 2002; Gee, 2003; Lepper & Henderlong, 2000; Linnenbrink & Pintrich, 2002; Meyer & Turner, 2006; Pekrun, Volume 2). Studies that have tracked the emotions during tutoring have identified the predominate emotions, namely confusion, frustration, boredom, anxiety, and flow/engagement, with delight and surprise occurring less frequently (Baker, D'Mello, Rodrigo, & Graesser, 2010; Craig, Graesser, Sullins, & Gholson, 2004; D'Mello, Craig, Witherspoon, McDaniel, & Graesser, 2008; D'Mello, Graesser, & Picard, 2007). Aside from detecting these student emotions, it is important for tutors to adopt pedagogical and motivational strategies that are effectively coordinated with the students' emotions. For example, Lepper, Drake, and O'Donnell (1998) proposed an INSPIRE model to promote this integration. This model encourages the tutor to nurture the student by being empathetic and attentive to the student's needs, to assign tasks that are not too easy or too difficult, to give indirect feedback on erroneous student contributions rather than harsh feedback,

to encourage the student to work hard and face challenges, to empower the student with useful skills, and to pursue topics they are curious about.

Cognitive and emotional traits of the students would presumably be important in the tutoring process. Meyer and Turner (2006) identified three theories that are particularly relevant to understanding the links between emotions and learning: academic risk taking, flow, and goals (Meyer & Turner, 2006). The academic risk theory contrasts (a) the adventuresome learners who want to be challenged with difficult tasks, take risks of failure, and manage negative emotions when they occur and (b) the cautious learners who tackle easier tasks, take fewer risks, and minimize failure and the resulting negative emotions (Clifford, 1991). The tutor's feedback and strategies should be very different for these two types of students, but available tutoring research has not systematically examined such contingencies and traced them to differential learning gains or student motivation. According to flow theory (Csikszentmihali, 1990), the learner is in a state of flow when the learner is so deeply engaged in learning the material that time and fatigue disappear. When students are in the flow state, they are at an optimal zone of facing challenges and conquering the challenges by applying their knowledge and skills. Once again, there is a lack of research examining whether tutors are sensitive to students in flow, as opposed to students who are bored, disengaged, or frustrated. Goal theory emphasizes the role of goals in predicting and regulating emotions (Dweck, 2002). Outcomes that achieve challenging goals result in positive emotions whereas outcomes that jeopardize goal accomplishment result in negative emotions. Available research on human tutoring has not explored how tutors monitor the students' goals and outcomes and whether affect-sensitive strategies improve learning gains and student motivation.

In summary, available evidence indicates that human tutors are limited in their ability to conduct student modeling on the psychological states of learners, whether they be cognitive states or emotions (D'Mello & Graesser, in press; Graesser, D'Mello, & Person, 2009). Most human tutors also are not sufficiently trained to implement a wide range of promising pedagogical strategies that are known to help learning. This of course opens the door to pursue computer tutors that take the tutoring process to the next level. We now turn to the world of Intelligent Tutoring Systems (ITS).

Intelligent Tutoring Systems

The history of ITS has always had a cross fertilization with the psychological sciences in addition to implementing intelligent computational mechanisms. Available ITS vary with respect to their adopting cognitive architectures and strategies of human tutors, as opposed to rational models that have minimal links to psychological research and theory. The ITS under discussion in this chapter are those that have moderate or strong links to psychological mechanisms. Regardless of the ITS-psychology alignment, however, those who work in the ITS world are inspired by the notion that learning will improve by implementing powerful intelligent algorithms that adapt to the learner at a fine-grained level and that instantiate complex principles of learning.

We assume that ITS environments are a generation beyond conventional computer-based training (CBT, see Moreno, Volume 3). Many CBT systems also adapt to individual learners, but they do so at a more course-grain level with simple learning principles. In a prototypical CBT system, the learner (a) studies material presented in a lesson, (b) gets tested with a multiple choice test or another objective test, (c) gets feedback on the test performance, (d) re-studies the

material if the performance in c is below threshold, and (e) progresses to a new topic if performance exceeds threshold. The order of topics presented and tested typically follows a predetermined order, such as ordering on complexity (simple to complex) or ordering on prerequisites (Gagné, 1985). The materials presented in a lesson can vary in CBT, including organized text with figures, tables, and diagrams (essentially books on the web), multimedia, problems to solve, example problems with solutions worked out, and other classes of learning objects. CBT has been investigated extensively for decades and has an associated mature technology. Meta-analyses show effect sizes of 0.39 sigma compared to classrooms (Dodds & Fletcher, 2004), whereas Mayer's (2009) analyses of multimedia have substantially higher sigma. The amount of time that learners spend studying the material in CBT has a 0.35 correlation with learning performance (Taraban, Rynearson, & Stalcup, 2001) and can be optimized by contingencies that distribute practice (Pashler et al., 2007).

These CBT systems are an important class of learning environments that can serve as tutors. However, the next generation of ITS went a giant step further that enhanced the adaptability, grain-size, and power of computerized learning environments. ITS address the two major principles of cognitive research in human tutoring, namely tracking knowledge and adaptive instruction. The processes of tracking knowledge (called user modeling) and adaptively responding to the learner both incorporate computational models in artificial intelligence and cognitive science, such as production systems, case-based reasoning, Bayes networks, theorem proving, and constraint satisfaction algorithms. More will be said about these models later in this chapter. There will be no attempt here to sharply divide systems that are CBTs versus ITS (Doignon & Falmagne, 1999; O'Neil & Perez, 2003), but one dimension that is useful is the space of possible interactions that can be achieved with the two classes of systems. For ITS,

virtually every tutorial interaction is unique and the space of possible interactions encounters a combinatorial explosion. An ITS attempts to fill in very specific learning deficits, to correct very specific misconceptions, and implement dynamic sequencing and navigation. For CBT, interaction histories can be identical for multiple students and the interaction space is finite, if not small. Nevertheless, the distinction between CBT and ITS is not of central concern to this chapter other than to say that the distinction is blurry and that both classes of learning environments appear intelligent to the learners (Fletcher, 2001; Hu & Graesser, 2004).

Successful systems have been developed for mathematically well-formed topics, including algebra, geometry, programming languages (the *Cognitive Tutors*: Anderson et al., 1995; Koedinger et al., 1997; Ritter et al., 2007, *ALEKS*: Doignon & Falmagne, 1999), physics (*Andes, Atlas,* and *Why/Atlas*: VanLehn et al., 2002; VanLehn et al., 2007), electronics (*SHERLOCK*: Lesgold, Lajoie, Bunzo, & Eggan, 1992), and information technology (*KERMIT*: Mitrovic, Martin, & Suraweera, , 2007). These systems show impressive learning gains (1.00 sigma, approximately, Corbett, 2001; Dodds & Fletcher, 2004) compared with suitable control conditions, particularly for deeper levels of comprehension. School systems are adopting ITS at an increasing pace even though they are initially expensive to build. After an ITS is finished being built and tested, its costs of use and overhead are quite small compared with human tutors.

Recent ITS environments handle knowledge domains that have a stronger verbal foundation as opposed to mathematics and precise analytical reasoning. It is a substantial challenge to promote learning in knowledge domains that emphasize verbal reasoning, conceptual models, approximation, imprecise abstraction, and diverse rhetorical structures (Bazerman, 1988). The *Intelligent Essay Assessor* (Landauer, Laham, & Foltz, 2000; Landauer, 2007) and *e-Rater* (Burstein, 2003) grade essays on science, history, and other topics as reliably

as experts of English composition. *AutoTutor* (Graesser, Chipman, Haynes, & Olney, 2005; Graesser, Jeon, & Dufty, 2008; Graesser, Lu et al., 2004) and Why-Atlas (VanLehn et al., 2002) help college students learn about computer literacy, physics, and critical thinking skills by holding conversations in natural language. AutoTutor shows learning gains of approximately 0.80 sigma compared with reading a textbook for an equivalent amount of time (Graesser, Lu et al., 2004; VanLehn et al., 2007). These systems automatically analyze language and discourse by incorporating recent advances in computational linguistics (Jurafsky & Martin, 2008) and information retrieval, notably latent semantic analysis (Landauer, McNamara, Dennis, & Kintsch, 2007; Millis et al., 2004).

At this point we turn to some recent ITS environments that have been tested on thousands of students and have proven effective in helping students learn. These systems also have architectures that are motivated by scientific principles of learning that were developed in the psychological sciences. The systems include the cognitive tutors, constraint-based tutors, casebased tutors, and tutors with animated conversational agents.

Although these systems appear to be quite different, they are similar with respect to having a 2-level structure with an outer and an inner loop (VanLehn, 2006). The outer loop involves the selection of problems, the judgment of mastery of a problem, and other more global aspects of the tutorial interaction. The inner loop consists of covering individual steps within a problem at a more micro-level, e.g. a solution step in a mathematics problem. The tracking of knowledge and adaptivity are necessary at both the outer loop and the inner loop in a bona fide ITS.

Cognitive Tutor

One of the salient success stories of transferring science to useful technology is captured in the Cognitive $Tutor_{TM}$, a class of tutoring systems built by researchers at Carnegie Mellon

University and produced by Carnegie Learning, Inc. The Cognitive Tutor is built on careful research grounded in cognitive theory and extensive real-world trials, and it has realized the ultimate goal of improvements over the status quo of classroom teaching. Its widespread implementation has drawn interest to both its psychological mechanisms and to its efficacy.

The Cognitive Tutor instruction is based on a cognitive model developed by Anderson (1990), called Adaptive Control of Thought (ACT, or ACT-R in its updated form). In the Cognitive Tutor paradigm, the tutor has a psychological model capable of solving the problems presented to the student. This is accomplished through production rules that not only model correct steps in the solution path but also "buggy" rules with incorrect specifications that end up leading the student to deviate from the correct path. A production rule is an "IF [state S], THEN [action A]" expression that designates that action A is likely to be performed when a state S is perceived and/or in working memory. An action may be a cognitive act (such as a search for information or making a decision) or a behavior (such as pressing a button, or adding two numbers). A state S may refer to a perceived visual or auditory event, a goal, or an idea in working memory.

Over the course of problem solving, production rules in long-term memory are activated and capture the student's knowledge and mastery of the skill. The correspondence between the student's behavior and production rule firing is known as *model-tracing*. The student is performing well when there is a synchrony between the expected firing of correct production rules and the student's actions. However, sometimes the student has incorrect buggy rules that lead to errors; these are also anticipated by the Cognitive Tutor and drive the selection corrective feedback and hints to remediate the errors. *Knowledge-tracing* assigns a probability of student mastery to each correct production rule based on student behavior. It also assigns a probability to

the buggy rules that require correction. Thus, the Cognitive Tutor's cognitive model is well suited to modeling procedural tasks and skills that precisely define actions (via model tracing) that are aligned understanding (via knowledge tracing). When the system has decided that enough of the skill requirements have been met for mastery, the tutor and student move on to a new section. By incorporating both model tracing and knowledge tracing, it is possible to give the student feedback on what they are doing right versus wrong as well as what to do next (such as hinting and help).

As an example, consider a student trying to solve the problem below.

The student would have to have the declarative knowledge component "6 + 8 = 14" stored in memory in addition to making sure that the production rules are in place to retrieve this fact. Production rules help determine the manner in which student behavior is interpreted and also the knowledge students should gain as part of the learning process. In our example problem adopted from Anderson and Gluck (2001), two of the relevant production rules would be:

IF the goal is to add the numbers in a column & n1 + n2 are in the column [Information visually perceived] THEN retrieve the sum of n1 and n2 [Retrieve from long-term memory] IF the goal is to put a sum in a column & n1 + n2 = n3 [Information retrieved from LTM] THEN set as a subgoal to write n3 in that column. [Perform behavior]

The student will perform very well and get positive feedback if the model tracing facilities reveal they have the right goals in working memory, they perceive the correct information on the visual display, and retrieve the correct information in working memory. However, there are potential breakdowns from inattentive students or buggy rules. The student may fail to perceive the correct information in the display (n1 + n2) in the column via the first production rule), fail to

have declarative information stored in long-term memory (not knowing that 6 + 8 = 14), or fail to post the correct goals in working memory. Such failures are manifested in their behavior and probabilistically update the procedural rules and declarative knowledge in the knowledge tracing.

According to ACT-R, an important part of cognitive skill consists of a large set of production rules that accumulate in long-term memory, that get activated by the contents of working memory, and that are dynamically composed in sequence when a particular problem is solved. An important part of learning is accessing and mastering these production rules in addition to factual or declarative knowledge. Behaviors performed by the student reflect the production rules and declarative knowledge implicit in the task, so the system can reconstruct what knowledge the student has already mastered versus has yet to learn (Anderson & Gluck, 2001). Using this student model, ACT-R can target those knowledge components that are missing or in error, and can select problems that specifically address those faulty components. As noted previously, such a fine grained student model is uncommon in human tutoring and thus represents a giant step forward.

The fine grained nature of the student model dictates that Cognitive Tutors have a large number of skills and productions for the student to learn. In four of their curricula (Bridge to Algebra, Algebra 1, Geometry or Algebra 2), there are 2,400 skills (or collection of knowledge components) for the student to learn (Ritter et al., 2009). This is a very large space of detailed content, far beyond school standards, the knowledge of human teachers, and conventional CBT. The goal of Cognitive Tutors is to scaffold the correct methods for solving problems with the student until they become automatized after multiple problems in multiple contexts. Since any given task is made up of a combination of procedural and declarative knowledge components, some important cognitive achievements are to chunk many productions into a larger single

production unit (Lee & Anderson, 2001) and to "proceduralize" the retrieval of some forms of declarative knowledge in order to speed up and strengthen its availability, thereby making the right facts and procedures highly accessible during problem solving (Ritter et al., 2007). For example, children solving addition problems can use a variety of strategies (Siegler and Shipley, 1995), including the sum strategy (counting from one) and the min strategy (counting from the larger of two addends). Proceduralization of declarative knowledge would increase the availability of the min strategy in situations like "1 million plus 1." Cognitive Tutors offer helps and hints during problem solving to aid the student when they do not know what to do next or they have buggy rules. Students can see their progress in Cognitive Tutors by looking at their skillmeter, which logs how many skills the student has acquired and depicts them in progress bars. The system gives just-in-time feedback when the student commits an error (Ritter et al., 2007). Cognitive Tutors include a mechanism where the student can solicit hints to overcome an impasse by pushing a button. More specifically, there are three levels of hints. The first level may simply remind the student of the goal, whereas the second level offers more specific help, and the third level comes close to directly offering the student the answer for a particular step in the problem solving. An example of an intermediate hint would be "As you can see in the diagram, Angles 1 and 2 are adjacent angles. How does this information help to find the measure of Angle 2?" when the student is learning about angles (Roll et al., 2006). These hints can be highly effective when used properly. However, some students attempt to "game the system," or abuse the hint function to quickly get through a lesson (Aleven, McLaren, Roll, & Koedinger, 2006; Baker, Corbett, Roll, & Koedinger, 2008). Gaming the system has been associated with lower learning outcomes for students and may be a consequence of learned helplessness.

As discussed earlier in this chapter, tests of Cognitive Tutor produce impressive learning gains in both experimental and classroom settings. Corbett (2001) tested various components of the Cognitive Tutor that teaches computer programming in order to identify those aspects of the system that produce the most learning. When a version with modeling tracing was compared with a version with no model tracing, there was an effect size of 0.75 sigma. Significant effect sizes have also been found in their investigations of the Cognitive Tutor in the school system. Their first classroom studies in Pittsburgh showed that Cognitive Tutor students excelled with an average of 0.6 sigma when compared to students in a traditional Algebra class. According to Ritter et al. (2007), standardized tests show overall effect sizes of 0.3 sigma and the *What Works Clearinghouse* investigations show an effect size of 0.4 sigma.

To be accurate, however, the results of the Cognitive Tutor are not uniformly rosy. For example, the representations and problems in which Cognitive Tutor students showed extremely high gains over traditional algebra students were experimenter-designed (Koedinger et al., 1997). A large-scale study in Miami with over 6,000 students showed that Cognitive Tutor students scored 0.22 sigma over their traditional Algebra student counterparts, but only scored 0.02 sigma better than the traditional students on the statewide standardized test (Shneyderman, 2001). It is widely acknowledged that there are challenges in scaling up any intervention and that the results will invariably be mixed.

Constraint-Based Tutors

Constraint-based modeling (CBM) is an approach first proposed by Ohlsson (Ohlsson, 1994, 1992) and later extended by Ohlsson and Mitrovic (2007). The core idea of CBM is to model the declarative structure of a good solution rather than the procedural steps leading to a good solution. Thus, CBM contrasts heavily with the model-tracing approach to student

modeling which models each step of an expert solution and sometimes ignores multiple alternate solutions. CBM instead has much conceptual similarity with declarative styles of programming, such as Prolog (Bratko, 1986), which define relationships between entities rather than procedural operations on entities. This abstract form of representations is designed to emphasize the general properties a good solution rather than the process of how it is obtained.

In CBM, the declarative structure of a good solution is composed of a set of state constraints (Ohlsson, 1994). Each state constraint is composed of a relevance condition (R) and a satisfaction condition (S). The relevance condition specifies when the constraint is relevant; only in these conditions is the state constraint meaningful. The satisfaction condition specifies whether the state constraint has been violated. A relevant and satisfied state constraint corresponds to an aspect of the solution that is correct. A relevant, but unsatisfied state constraint indicates a problem in the solution. The notion of state constraints can be formalized in several different ways, including a set of facts and rules.

Ohlsson (1994) gives a concrete example of constraint-based modeling in the domain of subtraction. Subtraction has two core concepts, each giving rise to a constraint. The first core concept is place value, meaning that the position of the digit affects its quantity, e.g. 9 in the tens place represents the quantity 90 whereas 9 in the ones column represents the quantity 9. The second core concept of subtraction is regrouping, in which the digits expressing a quantity may change without changing the value of the quantity, e.g. 90 = 9*10 + 0*1 = 8*10 + 10*1, so long as the decrement in one digit is offset by an increment in the other. The two constraints that follow from these core concepts are:

Constraint 1: Increments and corresponding decrements must occur together (otherwise the value of the numeral has changed)

Constraint 2: An increment of ten should not occur unless the digit in the position to the left is decremented by one.

The key observation is that a correct solution can never violate either of these constraints, no matter what order of operations is followed. Thus, the style of constraints is declarative rather than procedural. A large number of alternative procedures can be followed as long as they do not violate the constraints.

There are two proposed advantages to CBM over an ITS with student models such as those with the model-tracing and buggy production rules in the Cognitive Tutor (Ohlsson, 1994). First, CBM is allegedly able to account for a wider array of student behavior because it can accommodate greater deviations from the correct solution path. This advantage stems from the basic property of CBM that solution paths are not directly handcrafted and modeled; instead, a large family of solution paths is evaluated with respect to their fitting the designated constraints. The second major advantage of CBM is that it takes less effort for the researcher to create student models using CBM than traditional methods. Again, this advantage arises because solution paths in CBM are not crafted by hand. Because solution paths are not explicit in CBM, feedback is triggered by violation of satisfaction constraints. Moreover, the content of the feedback is not tied directly to a solution step but rather to a global constraint in the solution space. The implications of constraint-based student models have been explored in the literature, leading to a more elaborate and nuanced understanding of the tradeoffs between CBM and other modeling approaches.

Several CBM tutoring systems have been built by Mitrovic and colleagues with encouraging results (Mitrovic, Martin, & Suraweera, 2007; Mitrovic & Ohlsson, 1999; Mitrovic, Martin, & Mayo, 2002; Suraweera & Mitrovic, 2004). Particularly noteworthy are those that

support learning how to design computer databases and how to retrieve information through Structured Query Language (SQL). These systems have been incorporated into Addison-Wesley's *Database Place*, accessible by anyone who has bought a database textbook from Addison-Wesley (Mitrovic, Martin, & Suraweera, 2007). *KERMIT* (Suraweera & Mitrovic, 2004) is an entity-relationship tutor that focuses on database design. Constraints in KERMIT represent database design principles. Suraweera and Mitrovic (2004) give a very abstract example: If a composite multi-valued attribute is part of the ideal solution, then all corresponding entities in the student solution must have a matching composite multi-valued attribute. Note that this constraint follows a familiar "is-a" relation that allows the attribute to be instantiated in multiple ways. In a two hour pretest-posttest study with randomized assignment, students using KERMIT had significantly higher learning gains than students from the same class who used KERMIT without constraint-based feedback, with an effect size of 0.63.

CBM has been contrasted with the model-tracing architecture of the Cognitive Tutor. Indeed, this has been one of the fundamental debates in the ITS literature in recent years. Properties of the two models have been contrasted by Mitrovic, Koedinger, and Martin (2003) whereas Kodaganallur, Weitz, and Rosenthal (2006) has conducted a systematic investigation of two complete systems in the domain of hypothesis testing. According to Kodaganallur et al., when compared with model-tracing, CBM accounted for a narrower array of student behavior, required libraries to handle student misconceptions and errors, was unable to give procedural remediation, was incapable of giving fine-grained feedback, and was likely to give incorrect feedback to proper solutions. However, a number of methodological problems with this analysis were noted by Mitrovic and Ohlsson (2007). It suffices to say that the debate continues on the relative strengths and liabilities of these two architectures, but it is beyond the scope of this

chapter to clarify or resolve the debate.

Case-Based Reasoning

ITS with Case-based reasoning (CBR) are inspired by cognitive theories in psychology, education, and computer science that emphasize the importance of specific cases, exemplars, or scenarios in constraining and guiding our reasoning (Leake, 1996; Ross, 1987; Schank, 1999).

CBR can accommodate verbal reasoning and ill-defined problem domains whereas the Cognitive Tutor and constraint-based models are reserved for mathematics and semantically well-formed domains.

There are two basic premises in CBR that are differentially emphasized in the literature. The first basic premise is that memory is organized by cases that consist of a problem description, its solution, and associated outcomes (Watson & Marir, 1994). Accordingly, problem solving makes use of previously encountered cases rather than by proceeding from first principles. Although there is wide variation in the field as to the adaptation and use of cases, implemented systems generally follow four steps (Aamodt & Plaza, 1994):

RETRIEVE the most similar case(s), as indexed in the memory system

REUSE the case(s) to attempt to solve the current problem

REVISE the proposed solution to the current problem if necessary, and

RETAIN the new solution as a part of a new case that is stored in the memory system.

The second basic premise of CBR is that memory is dynamically organized around cases (Schank, 1999). That is, the outcome of the four steps above can not only cause cases to be reindexed according to an existing scheme, but can also drive the indexing scheme itself to change. Therefore, learning in the CBR paradigm goes beyond adding new cases after successful REUSE and beyond adding new cases after failed REUSE and successful REVISION. Success and

failure can indeed drive the creation of new cases, but they can also drive the way in all existing cases are organized and thus their future use (Leake, 1996).

Like ACT-R, CBR has been used as an underlying theory for the development of learning environments (Kolodner et al., 2003; Kolodner, Cox, & Gonzalez-Calero, 2005; Schank, Fano, Bell, & Jona, 1994). As a theory of learning, CBR implies that human learners engage in case-based analogical reasoning as they solve problems and learn by solving the problems. One might expect CBR learning environments to proceed in a similar fashion to model-tracing and CBM tutors: implement a tutor capable of solving the problem, allow the student to solve the problem, and then provide feedback based on the differences between the student's solution and the tutors. However, an automated ITS analogous to model-tracing and constraint-based modeling has not yet been implemented. Instead, CBR learning environments give learners the resources to implement CBR on their own. That is, designers of these CBR systems present learners with activities designed to promote CBR activities: identifying problems, retrieving cases, adapting solutions, predicting outcomes, evaluating outcomes, and updating a case library.

There are two CBR paradigms that proceed in this fashion of facilitating design by human learners. The first is exemplified by two environments called *Goal-Based Scenarios* (Schank, Fano, Bell, & Jona, 1994) and *Learning by Design* (Kolodner, Cox, & Gonzalez-Calero, 2005), both of which are highly integrated with classroom activities. Goal-Based Scenarios use simulated worlds as a context for learning (Schank et al., 1994). For example *Broadcast News* puts students in the scenario of having to create their own news program. Cases are news sources from the previous day, and these are RETRIEVED via tasks students perform to establish the social issues in each story. Experts are available to answer questions and provide feedback, thereby helping the students complete the REUSE and REVISE phases.

In contrast, Learning by Design frames the learning task as a design task (Kolodner et al., 2003). An example of this is designing a parachute. The Learning by Design addresses this task in 6 phases, with learning occurring in groups: clarifying the question, making a hypothesis, designing an investigation (using existing cases as input, REUSE), conducting the investigation (REUSE), analyzing results (REVISE), and group presentation. Kolodner et al. (2005) have informally reviewed educational outcomes in this paradigm and draw two conclusions. The first is that the reviewed CBR classes have significantly larger simple learning gains (posttest – pretest) compared to control classrooms on a number of measures, although effect sizes have not been reported. The second conclusion is that CBR classes have greater skill in collaborating and in scientific reasoning than matched peers. However, these conclusions need to considered with caution until rigorous tests of these claims are conducted in future research.

The second CBR paradigm involves support from a computer environment for one-on-one learning. Perhaps the best known work of this type is the law-based learning environment by Aleven and colleagues (Aleven, 2003; Ashley & Brüninghaus, 2009). The most current system, CATO, is a CBR system for legal argumentation, i.e. the arguments attorneys make using past legal cases. CATO uses its case library and domain background knowledge to organize multicase arguments, reason about significant differences between cases, and determine which cases are most relevant to the current situation. Students using CATO practice two types of tasks, theory-testing and legal argumentation, both of which rely heavily on CATO's case library. Theory-testing requires the student to predict a ruling on a hypothetical case by first forming a hypothesis, retrieving relevant cases from CATO (RETRIEVE), and then evaluating their hypothesis in light of the retrieved cases (REUSE, REVISE). Legal argumentation requires the student to write legal arguments for both the defendant and plaintiff on a hypothetical case.

Students first study the hypothetical case and then retrieve relevant cases from CATO (RETRIEVE). Next the students study example arguments that CATO generates dynamically based on the selected cases, in a kind of multi-case REUSE/REVISE. Students iteratively use this dynamic generation capability to explore the outcomes of combining different sets of cases, successively refining their arguments until they are complete (RETRIEVE/REUSE/REVISE). For argumentation, learning with CATO was not significantly different than learning from a well-respected human instructor when matched for time and content, in both cases in a law school setting (Aleven, 2003). The fact the there was no difference from a human instructor can be viewed as a success in the sense that human resources are expensive compared with an automated computer system that is finished being developed.

Conversational Agents

As discussed earlier in the context of human tutoring, conversational interaction may potentially explain why tutoring is so impressive in facilitating learning (Graesser, Person, & Magliano, 1995). Conversation, dialogue, and social interaction indeed have a long history in the field of education as an explanatory foundation for pedagogy in classroom talk and apprenticeship training models (Resnick, 2010; Resnick, Michaels, & O'Connor, in press; Rogoff & Garner, 1984; Vygotsky, 1978). This has presented a challenge to the ITS community. Can an ITS be designed to simulate tutoring with natural language in social contexts?

Animated conversational agents have met this challenge in some of the recent advanced learning environments (Atkinson, 2002; Baylor & Kim, 2005; Gholson et al., 2009; Graesser, Chipman, Haynes, & Olney, 2005; Hu & Graesser, 2004; Johnson, Rickel, & Lester, 2000; McNamara, Levinstein, & Boonthum, 2004; Moreno, Volume 3; Moreno & Mayer, 2007; Olney,

2009; Reeves & Nass, 1996). These agents interact with students and help them learn by either modelling good pedagogy or by holding a conversation. The agents may take on different roles: mentors, tutors, peers, players in multiparty games, or avatars in the virtual worlds. The students communicate with the agents through speech, keyboard, gesture, touch panel screen, or conventional input channels. In turn, the agents express themselves with speech, facial expression, gesture, posture, and other embodied actions. Intelligent agents with speech recognition essentially hold a face-to-face, mixed-initiative dialogue with the student, just as humans do (Cole et al., 2003; Graesser, Jackson, & McDaniel, 2007; Gratch et al., 2002; Johnson & Beal, 2005). Single agents model individuals with different knowledge, personalities, physical features, and styles. Ensembles of agents model social interaction.

AutoTutor is an intelligent tutoring system that helps students learn through tutorial dialogue in natural language (Graesser, Dufty, & Jeon, 2008; Graesser, Hu, & McNamara, 2005; Graesser et al., 2005). AutoTutor's dialogues are organized around difficult questions and problems that require reasoning and explanations in the answers. For example, below are two example challenging questions from two AutoTutor domains: Newtonian physics and computer literacy.

PHYSICS QUESTION: If a lightweight car and a massive truck have a head-on collision, upon which vehicle is the impact force greater? Which vehicle undergoes the greater change in its motion, and why?

COMPUTER LITERACY QUESTION: When you turn on the computer, how is the operating system first activated and loaded into RAM?

These questions require the learner to construct approximately 3-7 sentences in an ideal answer and to exhibit reasoning in natural language. These are hardly the fill-in-the-blank questions or

multiple-choice questions that many associate with learning technologies on computers. It takes a conversation to answer each one of these challenging questions. The dialogue for one of these challenging questions typically requires 50-100 conversational turns between AutoTutor and the student.

The structure of the dialogue in AutoTutor attempts to simulate that of human tutors, as was discussed earlier. More specifically, AutoTutor implements the three conversational structures: (a) a 5-step dialogue frame, (b) expectation and misconception-tailored dialogue, and (c) conversational turn management. These three levels can be automated and produce respectable tutorial dialogue. So far AutoTutor has been developed for computer literacy, physics, and scientific methods, but there is an open source version, called GnuTutor (Olney, 2009), that allows software developers to develop conversational dialogues for virtually any subject matter.

As an example, consider the following dialogue with AutoTutor on the example physics problem.

STUDENT: The lightweight vehicle experiences the greater impact and the greater change in motion because the mass of the lighter vehicle is less, therefore having less inertia to resist the change in motion. The momentum of the truck is greater than that of the car.

TUTOR: Right! What does Newton's third law of motion say about the forces that bodies exert on each other? [Positive feedback followed by a hint.]

STUDENT: The forces must be equal in magnitude and opposite in reaction.

TUTOR: Yes, correct! Can you add to that? [Positive feedback followed by a pump for more information.]

STUDENT: The forces must be equal and opposite reactions acting on opposite bodies.

TUTOR: Ya! We've almost covered it, but there's one part I think is still missing. I am sure you can get this. During an interaction, each body exerts an equal but opposite force on the other _____? [positive feedback followed by a prompt for a specific word.]

STUDENT: body

TUTOR: Yes, right! on the other body.

On each turn AutoTutor gives feedback to the student on the previous turn and also advances the dialogue with a pump, hint, or prompt for more information. AutoTutor attempts to get the student to do the talking rather than merely delivering the correct answer. Moreover, the hints and prompts are intelligently selected so that the student fills in missing words and ideas (what we call *pattern completion*). AutoTutor also corrects misconceptions expressed by the student and answers some classes of questions that the student may ask (such as definitions of terms, "What does inertia mean?").

AutoTutor can keep the dialogue on track because it is always comparing what the student says to anticipated input (i.e., the expectations and misconceptions in the curriculum script). For example, in the above question, AutoTutor has the following correct expectation: "According to Newton's third law of motion, if one body exerts a force on the other body then the other body must exert an equal and opposite force on the first body." Pattern matching operations and pattern completion mechanisms drive the comparison between what the student expresses and each expectation. These matching and completion operations are based on latent semantic analysis (Landauer et al., 2007) and symbolic interpretation algorithms (Rus & Graesser, 2006) that are beyond the scope of this article to address. We emphasize that AutoTutor cannot interpret student contributions that have no matches to content in the curriculum script. Novel

contributions of the students will not be considered close matches to the expectations or misconceptions so they will receive low scores in covering the expected content. This of course limits true mixed-initiative dialogue. That is, AutoTutor cannot explore the topic changes and tangents of students as the students introduce them. However, analyses of human tutoring (Chi et al., 2004; Graesser et al., 1995) has revealed that (a) human tutors rarely tolerate true mixed-initiative dialogue with students changing topics that steer the conversation off course and (b) most students rarely change topics, rarely ask questions, and rarely take the initiative to grab the conversational floor. Instead, it is the tutor that takes the lead and drives the dialogue. AutoTutor and human tutors are very similar in these respects.

The learning gains of AutoTutor have been evaluated in over 20 experiments conducted during the last 12 years. Assessments of AutoTutor on learning gains have shown effect sizes of approximately 0.8 standard deviation units in the areas of computer literacy (Graesser et al., 2004) and Newtonian physics (VanLehn, Graesser et al., 2007). AutoTutor's learning gains have varied between 0 and 2.1 sigma (a mean of 0.8), depending on the learning performance measure, the comparison condition, the subject matter, and the version of AutoTutor.

The agents described above interact with students one-to-one. Learning environments can also have pairs of agents interact with the student as a trialogue and larger ensembles of agents that exhibit ideal learning strategies and social interactions. It is extraordinarily difficult to train teachers and tutors to apply specific pedagogical techniques, especially when the techniques clash with the pragmatic constraints and habits of everyday conversation. However, pedagogical agents can be designed to have such precise forms of interaction.

As an example, iSTART (Interactive Strategy Trainer for Active Reading and Thinking) is an automated strategy trainer that helps students become better readers by constructing self-

explanations of the text (McNamara et al., 2004). The construction of self-explanations during reading is known to facilitate deep comprehension (Chi et al., 1994; Pressley & Afflerbach, 1995), especially when there is some context-sensitive feedback on the explanations that get produced (Palincsar & Brown, 1984). The iSTART interventions teach readers to self-explain using five reading strategies: monitoring comprehension (i.e., recognizing comprehension failures and the need for remedial strategies), paraphrasing explicit text, making bridging inferences between the current sentence and prior text, making predictions about the subsequent text, and *elaborating* the text with links to what the reader already knows. Groups of animated conversational agents scaffold these strategies in three phases of training. In an *Introduction* Module, a trio of animated agents (an instructor and two students) collaboratively describes selfexplanation strategies amongst themselves. In a *Demonstration Module*, two Microsoft Agent characters (Merlin and Genie) demonstrate the use of self-explanation in the context of a science passage and the student identifies the strategies being used. In a final *Practice* phase, Merlin coaches and provides feedback to the student one-to-one while the student practices selfexplanation reading strategies. For each sentence in a text, Merlin reads the sentence and asks the student to self-explain it by typing a self-explanation. Merlin gives feedback and asks the student to modify unsatisfactory self-explanations.

Studies have evaluated the impact of iSTART on both reading strategies and comprehension for thousands of students in K12 and college (McNamara, Best, O'Reilly, & Ozuru, 2006). The three-phase iSTART training (approximately 3 hours) has been compared with a control condition that didactically trains students on self-explanation, but without any vicariously modeling and any feedback via the agents. After training, the participants are asked to self-explain a transfer text (e.g., on heart disease) and are subsequently given comprehension

tests. The results have revealed that strategies and comprehension are facilitated by iSTART, with impressive effect sizes (0.4 to 1.4 sigma) for strategy use and for comprehension.

Future Directions

This chapter has made the case that tutoring by humans and computers provide powerful learning environments that incorporate a variety of theories and principles of learning. Many of the computational architectures are quantitatively precise, complete, and formally well-specified, which are virtues for helping students learn about mathematics, computer programming, engineering, physics, and other domains in the STEM areas. The Cognitive Tutor and constraint-based models are excellent examples of such ITSs. However, the ITS that involve case-based reasoning and conversational agents are designed to handle verbal reasoning, conversations, and qualitative knowledge domains and skills. Therefore the ITS enterprise is prepared to handle a diverse curriculum. Nevertheless, there are still a large number of unanswered fundamental questions that need attention in future research. This section identifies some directions for further inquiry from the standpoint of ITS environments.

Computer tutors allow researchers to have more control over the tutoring process than human tutors can provide. Human tutors are limited by strategies and habits that are so weathered into their cognition and social interaction, over a lifetime of experiences, that it is an uphill challenge to train them to implement specific pedagogical routines. The reliability of an ITS lies in a sharp contrast. This opens up the possibility of new programs of research that systematically compare different versions of an ITS and different types of ITS. All ITS have multiple modules, such as the knowledge base, the student's ability and mastery profile, decision rules that select problems, scaffolding strategies, help systems, feedback, media on the human-computer interface, and so on. Which of these components are responsible for any learning

gains of the ITS? It is possible to systematically manipulate the quality or presence of each component in *lesion* studies that manipulate the presence or absence of different tutoring components. The number of conditions in manipulation studies of course grows with the number of components. If there are 6 major components, with each level varying in 2 levels of quality, then there would be $2^6 = 64$ conditions in a factorial design. That would require nearly 2000 students in a between-subjects design with 30 students randomly assigned to each of the 64 conditions. If a variable has an impact on learning in a curvilinear fashion, then we would need three levels, resulting in $3^6 = 729$ conditions and nearly 22,000 students. This of course seems impractical, although thousands of students receive training on some ITS environments on the web (Heffernan, Koedinger, & Razzaq, 2008). The alternative would be to selectively focus on one or two modules at a time.

Comparisons need to be made between different computer tutors that handle the same subject matter. Algebra, for example, can be trained with the Cognitive Tutors, ALEKS, constraint-based models, and perhaps even case-based learning environments. Which of these provides the most effective learning for different populations of learners? It may be that there are aptitude-treatment interactions and, therefore, no clear winner. Eventually we would sort out which tutoring architecture works best for each population of learners.

The subject matters with verbal reasoning, as opposed to mathematical computation, need a different ITS architecture. The conversational agents are expected to play an important role in these topics that require verbal reasoning. Questions remain on how effective the conversational agents are compared to the more conventional graphical user interfaces. Is it best to make the interface resemble a face-to-face conversation with a human? Or does such anthropomorphic realism present a distraction from the subject matter? If the animated agent does resemble a

human, what is the ideal personality of the agent and to what extent should it appear intelligent? Should a computer agent claim it understands the student, has empathy, and recognizes the student is frustrated? Or should it not pretend to have such human elements? What are the pragmatic ground rules of a computer agent who wants to bond with a human learner? These questions are currently unanswered.

One of the provocative tests in the future will pit human versus machine as tutors. Most people place their bets on the human tutors under the assumption that they will be more sensitive to the student's profile and be more creatively adaptive in guiding the student. However, the detailed analyses of human tutoring challenge such assumptions in light of the many illusions that humans have about communication and the modest pedagogical strategies in their repertoire. Computers may do a better job in cracking the illusions of communication, in inducing student knowledge states, and in implementing complex intelligent tutoring strategies. A plausible case could easily be made for betting on the computer over the human tutor. Perhaps the ideal computer tutor emulates humans in some ways and complex non-human computations in other ways. Comparisons between human and computer tutors need to be made in a manner that equilibrates the conditions on content, time on task, and other extraneous variables that are secondary to pedagogy. As data roll in from these needed empirical studies, we suspect that there will be unpredictable and counterintuitive discoveries. If an ITS ends up being an improvement over the human tutor, we would not view the implications to be shocking, cold, or inhumane. A human tutor helps the student with the experiences and perspective of a single human, whereas an ITS potentially captures the wisdom of thousands.

References

- Aamodt, A., & Plaza, E. (1994). Case-based reasoning: Foundational issues, methodological variations, and system approaches. *AI Communications*, 7, 39-59.
- Aleven, V. (2003). Using background knowledge in case-based legal reasoning: a computational model and an intelligent learning environment. *Artificial Intelligence*, *150*, 183-237.
- Aleven, V., McLaren, Roll, I., & Koedinger, K. (2006). Toward meta-cognitive tutoring: A model of help seeking with a cognitive tutor. *International Journal of Artificial Intelligence in Education*, 16, 101-128.
- Anderson, J. R. (1990). The adaptive character of thought. Hillsdale, NJ: Erlbaum.
- Anderson, J.R. (2007). *How can the human mind occur in the physical universe?* Oxford: Oxford University Press.
- Anderson, J. R., Corbett, A. T., Koedinger, K. R., & Pelletier, R. (1995). Cognitive tutors: Lessons learned. *Journal of the Learning Sciences*, 4, 167-207.
- Anderson, J. R. & Gluck, K. (2001). What role do cognitive architectures play in intelligent tutoring systems? In D. Klahr & S. M. Carver (Eds.) *Cognition & Instruction: Twenty-five years of progress* (pp. 227-262). Hillsdale, NJ: Erlbaum.
- Ashley, K.D., & Brüninghaus, S. (2009). Automatically classifying case texts and predicting outcomes. *Artificial Intelligence and Law*, 17, 125-165.
- Atkinson, R. (2002). Optimizing learning from examples using animated pedagogical agents.

 **Journal of Educational Psychology*, 94, 416 427.

Baker, R.S., Corbett, A.T., Roll, I., Koedinger, K.R. (2008) Developing a generalizable detector of when students game the system *User Modeling and User-Adapted Interaction*, 18, 287-314.

- Baker, R.S., D'Mello, S.K., Rodrigo, M.T., & Graesser, A.C. (2010). Better to be frustrated than bored: The incidence, persistence, and impact of learners' cognitive-affective states during interactions with three different computer-based learning environments. *International Journal of Human-Computer Studies*, 68, 223-241.
- Baylor, A. L., & Kim, Y. (2005). Simulating instructional roles through pedagogical agents.

 International Journal of Artificial Intelligence in Education, 15, 5–115.
- Bazerman, C. (1988). Shaping written knowledge. Madison, WI: University of Wisconsin Press.
- Bloom, B. S. (1984). The 2 sigma problem: The search for methods of group instruction as effective as one-to-one tutoring. *Educational Researcher*, *13*, 4-16.
- Bransford, J. D., Brown, A. L., & Cocking, R. R. (Eds.). (2000). *How people learn* (expanded ed.). Washington, D.C.: National Academy Press.
- Bratko, I. (1986). *Prolog programming for artificial intelligence*. Wokingham, England: Addison Wesley Publishing Company.
- Burstein, J. (2003). The E-rater scoring engine: Automated essay scoring with natural language processing. In M. D. Shermis & J. C. Burstein (Eds.), *Automated essay scoring: A cross-disciplinary perspective*, (pp. 113-122). Mahwah, NJ: Erlbaum.
- Carlson, S. (1985). The ethical appropriateness of subject-matter tutoring for learning disabled adolescents. *Learning Disability Quarterly*, 8, 310-314.
- Chi, M. T. H., de Leeuw, N., Chiu, M., & LaVancher, C. (1994). Eliciting self-explanations improves understanding. *Cognitive Science*, *18*, 439-477.
- Chi, M. T. H., Roy, M., & Hausmann, R. G. M. (2008) Observing tutorial dialogues collaboratively: Insights about human tutoring effectiveness from vicarious learning. *Cognitive Science*, 32, 301-341.

Chi, M.T.H., Siler, S.A. & Jeong, H. (2004). Can tutors monitor students' understanding accurately? *Cognition and Instruction*, 22, 363-387.

- Chi, M.T.H., Siler, S., Yamauchi, T., Jeong, H. & Hausmann, R. (2001). Learning from human tutoring. *Cognitive Science*, 25, 471-534.
- Clifford, M.M. (1991). Risk taking: Theoretical, empirical and educational considerations. *Educational Psychologist*, 26, 263 297.
- Cohen, J. (1992). A power primer. Psychological Bulletin, 112, 155–159.
- Cohen, P. A., Kulik, J. A., & Kulik, C. C. (1982). Educational outcomes of tutoring: A metaanalysis of findings. *American Educational Research Journal*, 19, 237-248.
- Cole, R., van Vuuren, S., Pellom, B., Hacioglu, K., Ma, J., & Movellan, J., (2003). Perceptive animated interfaces: First steps toward a new paradigm for human computer interaction. *Proceedings of the IEEE*, 91, 1391-1405.
- Collins, A., Warnock, E. H., Aeillo, N., Miller, M. L. (1975). Reasoning from incomplete knowledge. In D. G. Bobrow A. Collins (Eds.), *Representation and understanding* (pp. 453-494). New York: Academic.
- Conley, M., Kerner, M., & Reynolds, J. (2005). Not a question of "should," but a question of "how:" Integrating literacy knowledge and practice into secondary teacher preparation through tutoring in urban middle schools. *Action in Teacher Education*, themed issue *Every Teacher a Teacher of Reading*, 27, 22 32.
- Corbett, A.T. (2001). Cognitive computer tutors: Solving the two-sigma problem. In M.Bauer, P.J. <u>Gmytrasiewicz</u>, and <u>J. Vassileva</u> (Eds.), *User Modeling: Proceedings of the Eighth International Conference*, (pp. 137-147). Heidelberg, Germany: Springer.

Craig, S.D., Graesser, A. C., Sullins, J., & Gholson, B. (2004). Affect and learning: An exploratory look into the role of affect in learning. *Journal of Educational Media*, 29, 241-250.

- Csikszentmihalyi, M. (1990). Flow: The psychology of optimal experience. New York: Harper-Row.
- D'Mello, S. K., Craig, S.D., Witherspoon, A. W., McDaniel, B. T., and Graesser, A. C. (2008).

 Automatic Detection of Learner's Affect from Conversational Cues. *User Modeling and User-Adapted Interaction*, 18, 45-80.
- D'Mello, S., & Graesser, A.C. (in press). Multimodal semi-automated affect detection from conversational cues, gross body language, and facial features. *User Modeling and User-adapted Interaction*.
- D'Mello, S.K., Picard, R., & Graesser, A.C. (2007). Toward an affect-sensitive AutoTutor. IEEE Intelligent Systems, 22, 53-61.
- Deci, E. L., & Ryan, R. M. (2002). The paradox of achievement: The harder you push, the worse it gets. In J. Aronson (Ed.), *Improving academic achievement: Impact of psychological factors on education* (pp. 61-87). Orlando, FL: Academic Press.
- Dodds, P. V. W., & Fletcher, J. D. (2004) Opportunities for new "smart" learning environments enabled by next genera-tion web capabilities. *Journal of Education Multimedia and Hypermedia*, 13, 391-404.
- Doignon, J.P. & Falmagne, J. C. (1999). Knowledge spaces. Berlin, Germany: Springer.
- Dweck, C. S. (2002). Messages that motivate: How praise molds students' beliefs, motivation, and performance (in surprising ways). In J. Aronson (Ed.), *Improving academic*

- achievement: Impact of psychological factors on education (pp. 61-87). Orlando, FL: Academic Press.
- Fletcher, J. D. (2001) What do sharable instructional objects have to do with intelligent tutoring systems, and vice versa? *International Journal of Cognitive Ergonomics*, 5, 317-333.
- Gagne, R. M. (1985). *The conditions of learning and theory of instruction* (4th ed.). New York: Holt, Rinehart, & Winston.
- Gee, J.P. (2003). What video games have to teach us about language and literacy. New York:

 Macmillan.
- Gholson, B., Witherspoon, A., Morgan, B., Brittingham, J. K., Coles, R., Graesser, A. C., Sullins, J., & Craig, S. D. (2009). Exploring the deep-level reasoning questions effect during vicarious learning among eighth to eleventh graders in the domains of computer literacy and Newtonian physics. *Instructional Science*, *37*, 487-493.
- Graesser, A. C., Chipman, P., Haynes, B., & Olney, A. (2005). AutoTutor: An intelligent tutoring system with mixed-initiative dialogue. *IEEE Transactions on Education*, 48, 612-618.
- Graesser, A. C., D'Mello, S. K, Cade, W. L. (2009). Instruction based on tutoring. In R.E. Mayer and P.A. Alexander (Eds.), *Handbook of Research on Learning and Instruction*. New York: Routledge Press.
- Graesser, A. C., D'Mello, S. K., & Person, N., (2009). Meta-knowledge in tutoring. In D. J. Hacker, J. Dunlosky, & A. C. Graesser (Eds.). *Handbook of Metacognition in Education* (pp. 361-382). New York: Taylor & Francis.
- Graesser, A.C., Hu, X., & McNamara, D.S. (2005). Computerized learning environments that incoporate research in discourse psychology, cognitive science, and computational

- linguistics. In A.F. Healy (Ed.), Experimental Cognitive Psychology and its Applications: Festschrift in Honor of Lyle Bourne, Walter Kintsch, and Thomas Landauer (pp. 183–194). Washington, D.C.: American Psychological Association.
- Graesser, A.C., Jackson, G.T., & McDaniel, B. (2007). AutoTutor holds conversations with learners that are responsive to their cognitive and emotional states. *Educational Technology*, 47, 19-22.
- Graesser, A. C., Jeon, M., & Dufty, D. (2008). Agent technologies designed to facilitate interactive knowledge construction. *Discourse Processes*, 45, 298–322.
- Graesser, A.C., Lu, S., Jackson, G.T., Mitchell, H., Ventura, M., Olney, A., & Louwerse, M.M. (2004). AutoTutor: A tutor with dialogue in natural language. *Behavioral Research Methods, Instruments, and Computers*, *36*, 180-193.
- Graesser, A. C., & Person, N. K. (1994). Question asking during tutoring. *American Educational Research Journal*, *31*, 104-137.
- Graesser, A. C., Person, N. K., & Magliano, J. P. (1995). Collaborative dialogue patterns in naturalistic one-to-one tutoring. *Applied Cognitive Psychology*, *9*, 1-28.
- Gratch, J., Rickel, J., Andre, E., Cassell, J., Petajan, E., & Badler, N. (2002). Creating interactive virtual humans: Some assembly required. *IEEE Intelligent Systems*, 17, 54 63.
- Hacker, D.J., & Graesser, A.C. (2007). The role of dialogue in reciprocal teaching and naturalistic tutoring. In R. Horowitz (Ed.), *Talk about text: How speech and writing interact in school learning* (pp. 253-270). Mahwah, NJ: Erlbaum.
- Heffernan, N.T., Koedinger, K.R., & Razzaq, L.(2008) Expanding the model-tracing architecture:

 A 3rd generation intelligent tutor for Algebra symbolization. *The International Journal of Artificial Intelligence in Education*, 18, 153-178.

Hock, M., Pulvers, K., Deshler, D., & Schumaker, J. (2001). The effects of an after-school tutoring program on the academic performance of at-risk students and students with learning disabilities. *Remedial and Special Education*, 22, 172-186.

- Hock, M., Schumaker, J., & Deshler, D. (1995). Training strategic tutors to enhance learner independence. *Journal of Developmental Education*, 19, 18-26.
- Hu, X., & Graesser, A.C. (2004). Human Use Regulatory Affairs Advisor (HURAA): Learning about research ethics with intelligent learning modules. *Behavioral Research Methods*, *Instruments, and Computers*, 36, 241-249.
- Invernizzi, M., Rosemary, C., Juel, C., & Richards, H. (1997). At-risk readers and community volunteers: A three year perspective. *Scientific Studies of Reading*, *1*, 277-300.
- Johnson, W.L., & Beal, C. (2005). Iterative evaluation of a large-scale, intelligent game for learning language. In C. Looi, G. McCalla, B. Bredeweg, & J. Breuker (Eds.), *Artificial Intelligence in Education* (pp. 290-297). Amsterdam: IOS Press.
- Johnson, D.W., & Johnson, R.T. (1992). Implementing cooperative learning. *Contemporary Education*, 63, 173–180.
- Johnson, W. L., Rickel, J. W., & Lester, J. C. (2000). Animated pedagogical agents: Face-to-face interaction in interactive learning environments. *International Journal of Artificial Intelligence in Education*, 11, 47-78.
- Jurafsky, D., & Martin, J.H. (2008). Speech and language processing: An introduction to natural language processing, computational linguistics, and speech recognition. Upper Saddle River, NJ: Prentice-Hall.
- Kodaganallur, V., Weitz, R. R., & Rosenthal, D. (2006). An assessment of constraint-based

 Tutors: A response to Mitrovic and Ohlsson's critique of "A comparison of model-tracing

- and constraint-based intelligent tutoring paradigms". *International Journal of Artificial Intelligence in Education*, *16*, 291-321.
- Koedinger, K. R., Anderson, J. R., Hadley, W. H., & Mark, M. (1997). Intelligent tutoring goes to school in the big city. *International Journal of Artificial Intelligence in Education*, 8, 30-43.
- Kolodner, J., Camp, P., Crismond, D., Fasse, B., Gray, J., & Holbrook, J. (2003). Problem-based learning meets case-based reasoning in the middle-school science classroom: Putting learning by design into practice. *Journal of the Learning Sciences*, *12*, 495-547.
- Kolodner, J., Cox, M., & Gonzalez-Calero, P. (2005). Case-based reasoning-inspired approaches to education. *The Knowledge Engineering Review*, 20, 299-303.
- Landauer, T. K. (2007). LSA as a theory of meaning. In T. Landauer, D. McNamara, D. Simon, & W. Kintsch (Eds.) *Handbook of Latent Semantic Analysis* (pp. 3-34). Mahwah, NJ: Erlbaum.
- Landauer, T. K., Laham, D., & Foltz, P. W. (2000). The Intelligent Essay Assessor. *IEEE Intelligent Systems*, 15, 27-31.
- Landauer, T., McNamara, D. S., Dennis, S., Kintsch, W. (2007) (Eds.), *Handbook of latent semantic analysis*. Mahwah, NJ: Erlbaum
- Leake, D. (1996). CBR in context: The present and future. In D. Leake (ed.), *Case-bBased reasoning: Experiences, lessons, and future directions* (pp. 3-30). Menlo Park, CA: AAAI Press/MIT Press.
- Lee, F. J. & Anderson, J. R. (2001). Does learning of a complex task have to be complex? A study in learning decomposition. *Cognitive Psychology*, 42, 267-316.

Lepper, M. R., Drake, M., & O'Donnell-Johnson, T. M. (1997). Scaffolding techniques of expert human tutors. In K. Hogan & M. Pressley (Eds), *Scaffolding student learning: Instructional approaches and issues* (pp. 108-144). New York: Brookline Books.

- Lepper, M. R., & Henderlong, J. (2000). Turning "play" into "work" and "work" into "play": 25 years of research on intrinsic versus extrinsic motivation. In C. Sansone & J.

 M.Harackiewicz (Eds.), *Intrinsic and extrinsic motivation: The search for optimal motivation and performance* (pp. 257-307). San Diego, CA: Academic Press.
- Lesgold, A., Lajoie, S. P., Bunzo, M., & Eggan, G. (1992). SHERLOCK: A coached practice environment for an electronics trouble-shooting job. In J. H. Larkin & R. W. Chabay (Eds.), Computer assisted instruction and intelligent tutoring systems: Shared goals and complementary approaches (pp. 201–238). Hillsdale, NJ: Erlbaum.
- Linnenbrink, E. A. & Pintrich, P. (2002). The role of motivational beliefs in conceptual change. In M. Limon & L. Mason (Eds.), *Reconsidering conceptual change: Issues in theory and practice* (pp. 115-135). Dordretch, Netherlands: Kluwer Academic Publishers.
- Litman, D.J., Rose, C.P., Forbes-Riley, K., VanLehn, K., Bhembe, D., & Silliman, S. (2006).

 Spoken versus typed human and computer dialogue tutoring. *International Journal of Artificial Intelligence in Education*, 16, 145-170.
- Mathes, P. G., & Fuchs, L. S. (1994). Peer tutoring in reading for students with mild disabilities:

 A best evidence synthesis. *School Psychology Review*, 23, 59-80.
- Mayer, R. E. (2009). Multimedia learning (2nd ed). New York: Cambridge University Press.
- McArthur, D., Stasz, C., & Zmuidzinas, M. (1990). Tutoring techniques in algebra. *Cognition and Instruction*, 7, 197-244.

McNamara, D. S., Levinstein, I. B., & Boonthum, C. (2004). iSTART: Interactive strategy training for active reading and thinking. *Behavioral Research Methods, Instruments, and Computers*, 36, 222 - 233.

- McNamara, D. S., O'Reilly, T. P., Best, R. M., & Ozuru, Y. (2006). Improving adolescent students' reading comprehension with iSTART. *Journal of Educational Computing Research*, 34, 147-171.
- Meier, J., & Invernizzi, M. (2001). Book Buddies in the Bronx: A model for America Reads. *Journal for the Education Placement of Students Placed At-risk*, 6, 391-333.
- Merrill, D. C., Reiser, B. J., Merrill, S. K., & Landes, S. (1995). Tutoring: Guided learning by doing. *Cognition and Instruction*, *13*, 315-372.
- Meyer, D. K., & Turner, J. C. (2006). Re-conceptualizing emotion and motivation to learn in classroom contexts. *Educational Psychology Review*, *18*, 377-390.
- Millis, K., Kim, H. J., Todaro, S. Magliano, J. P., Wiemer-Hastings, K., & McNamara, D. S. (2004). Identifying reading strategies using latent semantic analysis: Comparing semantic benchmarks. *Behavior Research Methods, Instruments, & Computers, 36*, 213-221.
- Mitrovic, A., Koedinger, K., & Martin, B. (2003). A comparative analysis of cognitive tutoring and constraint-based modeling. In P. Brusilovsky, A. Corbett, F. de Rosis (Eds.), *Proceedings of the Ninth International Conference on User Modeling* (pp. 313-322). Berlin,
 Germany: Springer-Verlag.
- Mitrovic, A., Martin, B., & Mayo, M. (2002). Using evaluation to shape ITS design: Results and experiences with SQL-Tutor. *User Modeling and User-Adapted Interaction*, *12*, 243-279.
- Mitrovic, A., Martin, B., & Suraweera, P. (2007). Intelligent tutors for all: The constraint-based approach. *IEEE Intelligent Systems*, 22, 38-45.

Mitrovic, A., & Ohlsson, S. (1999). Evaluation of a constraint-based tutor for a database language. *International Journal on Artificial Intelligence in Education*, 10, 238-256.

- Moreno, R., & Mayer, R. E. (2007). Interactive multimodal learning environments. *Educational Psychology Review*, 19, 309-326.
- O'Neil, H. & Perez, R. (Eds.). (2003). Web-based learning: Theory, Research and Practice.

 Mahwah NJ: Lawrence Erlbaum Associates.
- Ohlsson, S. (1992). Constraint-based student modelling. *International Journal of Artificial Intelligence in Education*, 3, 429-447.
- Ohlsson, S. (1994). Constraint-based student modeling. In J. E. Greer & G. McCalla (Eds.), Student modelling: The key to individualized knowledge-based instruction (pp. 167-190). Berlin, Germany: Springer.
- Ohlsson, S., & Mitrovic, A. (2007). Fidelity and efficiency of knowledge representations for intelligent tutoring systems. *Technology, Instruction, Cognition and Learning (TICL)*, 5, 101-132.
- Olney, A.M. (2009). GnuTutor: An open source intelligent tutoring system based on AutoTutor.

 *Proceedings of the AAAI Fall Symposium on Cognitive and Metacognitive Educational Systems (pp. 70-75). Menlo Park, CA: AAAI.
- Palincsar, A.S., & Brown, A.L. (1984). Reciprocal teaching of comprehension- fostering and monitoring activities. *Cognition and Instruction*, *1*, 117-175.
- Pashler, H., Bain, P. M., Bottge, B. A., Graesser, A., Koedinger, K., & McDaniel, M., (2007).

 Organizing instruction and study to improve student learning. *IES practice guide* (NCER 2007-2004). Washington, DC: National Center for Education Research.

Person, N. K., & Graesser, A. C. (1999). Evolution of discourse in cross-age tutoring. In A. M.O'Donnell and A. King (Eds.), *Cognitive perspectives on peer learning* (pp. 69-86). Mahwah, NJ: Erlbaum.

- Person, N., & Graesser, A.C. (2003). Fourteen facts about human tutoring: Food for thought for ITS developers. In V. Aleven, U. Hoppe, J. Kay, R. Mizoguchi, H. Pain, F. Verdejo, and K. Yacef (Eds.), *AIED2003 Supplemental Proceedings* (pp. 335-344). Sydney, Australia: University of Sydney School of Information Technologies.
- Person, N. K., Kreuz, R. J., Zwaan, R., & Graesser, A. C. (1995). Pragmatics and pedagogy:

 Conversational rules and politeness strategies may inhibit effective tutoring. *Cognition and Instruction*, 13, 161-188.
- Person, N., Lehman, B., & Ozbun, R. (2007). Pedagogical and motivational dialogue moves used by expert tutors. Presented at the 17th Annual Meeting of the Society for Text and Discourse. Glasgow, Scotland.
- Pressley, M., & Afflerbach, P. (1995). *Verbal protocols of reading: The nature of constructively responsive reading*. Hillsdale NJ: Erlbaum.
- Pressley, M., & McCormick, C. (1995). *Cognition, teaching and assessment*. New York: Harper Collins.
- Reeves, B. & Nass, C. (1996). *The Media Equation: how people treat computers, televisions, and new media like real people and places*. Stanford, California: University Press.
- Resnick, L.B. (in press). Nested learning systems for the thinking curriculum. *Educational Researcher*.

Resnick, L.B., Michaels, S., & O'Connor, C. (in press). How (well structured) talk builds the mind. In R. Sternberg and D. Preiss (Eds.), *From genes to context: New discoveries about learning from educational research and their applications.* New York: Springer.

- Ritter, S., Anderson, J. R., Koedinger, K. R., Corbett, A. (2007) Cognitive Tutor: Applied research in mathematics education. *Psychonomic Bulletin & Review*, *14*, 249-255.
- Ritter, G., Barnett, J., Denny, G., & Albin, G. (2009). The effectiveness of volunteer tutoring programs for elementary and middle school students: A meta-analysis. *Review of Educational Research*, 79, 3-38.
- Rogoff, B., & Gardner, W. (1984). Adult guidance of cognitive development. In B. Rogoff and J. Lave, J. (Eds.), *Everyday cognition: Its development in social context* (pp. 95-116). Cambridge, MA: Harvard University Press
- Rohrbeck, C. A., Ginsburg-Block, M., Fantuzzo, J. W., & Miller, T. R. (2003). Peer assisted learning Interventions with elementary school students: A Meta-Analytic Review. *Journal of Educational Psychology*, 95, 240-257.
- Roll, I., Aleven, V., McLaren, B.M., Ryu, E., Baker, R.S., & Koedinger, K.R. (2006). The Help Tutor: Does metacognitive feedback improve students' help-seeking actions, skills and learning? In M. Ikeda, K.D. Ashley, & T-W. Chan (Eds.), *Proceedings of the 8th International Conference on Intelligent Tutoring Systems* (pp. 360-369). Berlin: Springer.
- Ross, R.H. (1987). This is like that: The use of earlier problems and the separation of similarity effects. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 13*, 629–639.

Rus, V., & Graesser, A.C. (2006). Deeper natural language processing for evaluating student answers in intelligent tutoring systems. In the *Proceedings of the American Association of Artificial Intelligence*. (pp. 1495-1500). Menlo Park, CA: AAAI.

- Schank, R. C. (1999). Dynamic memory revisited. Cambridge University Press.
- Schank, R. C., Fano, A., Bell, B., & Jona, M. (1994). The design of goal-based scenarios. *Journal of the Learning Sciences*, 3, 305-345.
- Siegler, R. S., & Shipley, C. (1995). Variation, selection, and cognitive change. In T. Simon & G. Halford (Eds.), *Developing cognitive competence: New approaches to process modeling* (pp. 31-76). Hillsdale, NJ: Erlbaum.
- Shah, F., Evens, M.W., Michael, J., & Rovick, A. (2002). Classifying student initiatives and tutor responses in human keyboard-to keyboard tutoring sessions. *Discourse Processes*, *33*, 23-52.
- Shneyderman, A. (2001). *Evaluation of the Cognitive Tutor Algebra I Program*. Miami, FL: Miami-Dade County Public Schools Office of Evaluation and Research.
- Slavin, R.E. (1990). *Cooperative learning: Theory, research, and practice*. New Jersey: Prentice Hall.
- Slavin, R., Karweit, N., & Madden, N. (1989). *Effective programs for students at risk*. Boston: Allyn and Bacon.
- Sleeman D. & J. S. Brown. (1982). (Eds.). *Intelligent Tutoring Systems*. Orlando, Florida: Academic Press.
- Suraweera, P., & Mitrovic, A. (2004). An intelligent tutoring system for entity relationship modelling. *International Journal of Artificial Intelligence in Education*, *14*, 375-417.

Taraban, R., Rynearson, K., & Stalcup, K. (2001). Time as a variable in learning on the World Wide Web. *Behavior Research Methods, Instruments, & Computers, 33*, 217-225.

- Tollefson, J. (1997). Lab offers strategic help after school. Strategram, 1-7.
- Topping, K. (1996). The effectiveness of peer tutoring in further and higher education: A typology and review of the literature. *Higher Education*, *32*, 321-345.
- VanLehn, K. (2006) The behavior of tutoring systems. *International Journal of Artificial Intelligence in Education*, 16, 227-265.
- VanLehn, K., Graesser, A. C., Jackson, G. T., Jordan, P., Olney, A., & Rose, C. P. (2007). When are tutorial dialogues more effective than reading? *Cognitive Science*, *31*, 3-62.
- VanLehn, K., Jordan, P., Rosé, C. P., et al. (2002). The architecture of Why2-Atlas:A coach for qualitative physics essay writing. In S. A. Cerri, G. Gouarderes, & F. Paraguacu (Eds.), *Intelligent Tutoring Systems: 6th International Conference* (pp. 158-167). Berlin: Springer.
- VanLehn, K., Siler, S., Murray, C., Yamauchi, T., & Baggett, W.B. (2003). Why do only some events cause learning during human tutoring? *Cognition and Instruction*, 21, 209-249.
- Vellutino, F., Scanlon, D., Sipay, E., Small, S., Pratt, A., Chen, R., et al. (1996). Cognitive profiles of difficult-to-remediate and readily remediated poor readers: Early intervention as a vehicle for distinguishing between cognitive and experiential deficits as a basic cause of specific reading disability. *Journal of Educational Psychology*, 88, 601-638.
- Vygotsky, L.S. (1978). *Mind and society: The development of higher mental processes*. Cambridge, MA: Harvard University Press.
- Wasik, B. (1998). Volunteer programs in reading: A review. *Reading Research Quarterly*, 33, 266-292.

Wasik, B. A., & Slavin, R. E. (2003). Preventing early reading failure with one-to-one tutoring:

A best- evidence synthesis. *Reading Research Quarterly*, 28, 179-200.

- Watson, I., & Marir, F. (1994). Case-based reasoning: A review. *Knowledge Engineering Review*, 9, 327–354.
- Woolf, B. P. (2009). *Building intelligent interactive tutors*. Burlington, MA: Morgan Kaufmann Publishers.

.