

Intelligent Tutoring Systems: Using AI to Improve Training Performance and ROI

By [James Ong](#) and [Sowmya Ramachandran](#)

Learning may ultimately be its own reward, but many companies recognized long ago that learning also is critical to their business success. So they spend billions each year on training. Companies teach their employees internal processes and customer-facing services. They show customers how to operate and maintain their products, and they help their partners learn to sell, install, and service their products.

Corporate training is itself big business. According to the American Society for Training & Development (ASTD), annual spending on employee training alone represents about two percent of total payroll dollars in the U.S. And like virtually every aspect of business, training has embraced technology, in hopes of increasing the effectiveness, efficiency and overall ROI of corporate education. ASTD reports that, in the largest companies, employee training time spent with “learning technologies” is increasing – over 10 percent in 2001 – while time in the classroom is decreasing.

But to what effect? The sad truth is that most training methods and technologies generally produce, at best, “trained novices.” That is, they *introduce* facts and concepts to students, present them with relatively simple questions to test this new knowledge, and provide them with a few opportunities to practice using this knowledge in exercises or scenarios. However, becoming truly proficient requires *extensive practice* solving realistically-complex problems in a wide range situations, combined with *coaching and feedback* from managers, more experienced peers, or other types of experts.

In many professions, extensive (and expensive) practice-based training is the norm. For example, doctors first serve as interns, and astronauts train for years before lift-off. During peacetime, military commanders practice tactical decision-making, command, and control in live and virtual training exercises to prepare for battle. At most companies, however, cost constraints limit the amount of training available, so employees really hone their skills on-the-job. Yet, placing trained novices on-the-job is very expensive in its own way because sub-optimal performance incurs costs such as inefficient use of

people’s time and company resources, lower product or service quality, and lost business opportunities. For many occupations, it may take years on-the-job to acquire the breadth and depth of experience needed to achieve high performance.

Current learning technologies can help create trained novices more efficiently, but they are really not up to the job of creating true experts. For example, multimedia computer-based training (CBT) systems are good at presenting information and then testing factual recall using multiple choice or fill-in-the-blank questions. However, traditional CBT systems are incapable of providing *intelligent*, individualized coaching, performance assessment, and feedback students need to acquire broad and deep expertise.

Although web-based training (WBT) reduces logistical difficulties related to distributing and installing learning software, its pedagogical limitations are just as pronounced. When using web-based learning systems, students typically rely on asynchronous communication to get help from instructors and resolve their problems. When faced with a problem, the student must wait for help from an online instructor, who, depending on the instructor-to-student ratio, may take awhile to respond.

Computer-based simulations let students learn and experience the effects of different actions in a variety of situations. Although training simulations were pioneered by the aviation industry and the military, they are being adopted by more and more businesses to help students learn diverse subjects such as equipment and software operations and maintenance, business and analytical skills, and interpersonal skills.

Today, many of these simulations are relatively simple and prompt the student to select from just a few choices at each step in the simulation. As these simulations become more sophisticated, presenting many possible actions to the student and modeling myriad cause-and-effect relationships, it becomes more and more difficult for the student to determine exactly what they did well or not so well that led to

Intelligent Tutoring Systems – Using AI to Improve Training Performance and ROI

their success (or failure) during the scenario. Even if they succeeded in achieving the objectives of the scenario or exercise, it is unlikely that everything they did was correct or optimal, so what actions or decisions could have been improved, and what underlying gaps exist in their knowledge and skills?

Instructors can help students learn much more from simulated exercises by providing coaching during simulated exercises, assessing each student's performance, and providing them with helpful comments and feedback. Ideally, an instructor could be dedicated to each student to provide maximum support, because individualized instruction has been shown to be very effective.

About 20 years ago, research by Prof. Benjamin Bloom and others demonstrated that students who receive one-on-one instruction perform two standard deviations better than students in traditional classrooms. That is, the average tutored student performed as well as the top 2 percent of those receiving classroom instruction. However, in most cases, it is much too expensive to dedicate one instructor for each student. The challenge, then, is how to encode in software the subject matter expertise and the teaching skills of a company's best instructors or mentors to provide the benefits of intelligent, one-on-one instruction cost-effectively.

The intelligent tutoring system approach

Imagine that each learner in a classroom or WBT setting has a personal training assistant who pays attention to the participant's learning needs, assesses and diagnoses problems, and provides assistance as needed. The assistant could perform many of the routine instructional interventions and alert the instructor of learning problems that are too difficult for it. By taking on basic assistance tasks, the assistant would free the instructor to concentrate on training issues that require greater expertise.

Providing a personal training assistant for each learner is beyond the training budgets of most organizations. However, a *virtual* training assistant that captures the subject matter and teaching expertise of experienced trainers provides a captivating new option. The concept, known as intelligent tutoring systems (ITS) or intelligent computer-aided instruction (ICAI), has been pursued for more than three decades by researchers in education, psychology, and artificial intelligence.

Today, prototype and operational ITS systems provide practice-based instruction to support corporate training, K-12 and college education, and military training. Indeed, the technology is now ready for prime time.

The goal of ITS is to provide the benefits of one-on-one instruction automatically and cost effectively. Like training simulations, ITS enables participants to practice their skills by carrying out tasks within highly interactive learning environments. However, ITS goes beyond training simulations by answering user questions and providing individualized guidance. Unlike other computer-based training technologies, ITS systems assess each learner's actions within these interactive environments and develop a model of their knowledge, skills, and expertise. Based on the learner model, ITSs tailor instructional strategies, in terms of both the content and style, and provide explanations, hints, examples, demonstrations, and practice problems as needed.

Research on prototype systems indicate that students taught by ITSs generally learn faster and translate the learning into improved performance better than classroom-trained participants. At Carnegie Mellon University, for example, researchers developed an intelligent tutoring system called the LISP Tutor in the mid-1980s that taught computer programming skills to college students. In one controlled experiment, students who used the ITS scored 43 percent higher on the final exam than a control group that received traditional instruction. When given complex programming problems, the control group required 30 percent more time to solve these problems, compared to the ITS students. Another ITS, known as Sherlock, was developed in the early 1990s to train Air Force personnel on jet aircraft troubleshooting procedures. Learners taught using Sherlock performed significantly better than the control group and, after 20 hours of instruction, performed as well as technicians with four years of on-the-job experience.

Since those early implementations, ITSs have been developed for a widening variety of training applications. Here are some notable examples:

- An ITS developed by Alan M. Lesgold and colleagues at the University of Pittsburgh on behalf of a multinational semiconductor firm trains technicians to repair complex semiconductor chip-manufacturing equipment.




Figure 1 - Navy tactical tutoring system simulator and report card

- Carnegie Learning developed a suite of ITS-based "cognitive tutors" in secondary-level math subjects. The systems, based on earlier research carried out by John Anderson and Ken Koedinger at Carnegie Mellon University, were tested in a Pittsburgh public high school, where students showed 50- to 100-percent improvement in problem solving and use of equations, tables, and graphs.
- [Stottler Henke Associates](#) developed an [ITS for the U.S. Navy to teach Tactical Action Officers](#) using simulated scenarios and automated evaluation of the student's actions. Using this system, learners receive 10 times more hands-on tactical decision-making experience than before. This ITS is used as part of classroom instruction and as a stand-alone system aboard Navy ships, where learners use it for self-study.

How intelligent tutoring systems work

Many traditional instructional methods present learners with facts and concepts followed by test questions. These methods are effective in exposing people to large amounts of information and testing their recall. However, they often instill "inert knowledge" that learners can recall but may not apply correctly when needed. By contrast, ITS systems use simulations and other highly interactive learning environments that require people to apply their knowledge and skills. These active, situated learning environments help them retain and apply knowledge and skills more effectively in operational settings.

In order to provide hints, guidance, and instructional feedback to learners, ITS systems typically rely on three types of knowledge, organized into separate software modules:

- The "expert model" represents subject matter expertise and provides the ITS with knowledge of what it's teaching.

Intelligent Tutoring Systems – Using AI to Improve Training Performance and ROI

- The "student model" represents the student's knowledge, skills, and other attributes that affect how the student should be taught. This model lets the ITS know who it's teaching.
- The "instructor model" enables the ITS to know how to teach, by encoding instructional strategies used by the tutoring system.

Here's how each of these components works. An **expert model** is a computer representation of a domain expert's subject matter knowledge and problem-solving ability. This knowledge enables the ITS to compare the learner's actions and selections with those of an expert in order to evaluate what the user does and doesn't know.

A variety of artificial intelligence (AI) techniques are used to capture how a problem can be solved. For example, some ITS systems capture subject matter expertise in rules. That enables the tutoring system to generate problems on the fly, combine and apply rules to solve the problems, assess each learner's understanding by comparing the software's reasoning with theirs, and demonstrate the software's solutions to the participant's. Though this approach yields a powerful tutoring system, developing an expert system that provides comprehensive coverage of the subject material is difficult and expensive.

A common alternative to embedding expert rules is to supply much of the knowledge needed to support tutoring in each scenario definition. For example, scenario-based tutoring systems enable the course developer to create templates that specify allowable sequences of actions and states. This method avoids the need to encode the ability to solve all possible problems within an expert system. Instead, it requires only the ability to specify the far fewer number of actions that are appropriate in each scenario. Which technique is appropriate depends on the nature of the task and its underlying knowledge and skills.

The **student model** evaluates each learner's performance to determine his or her knowledge, perceptual abilities, and reasoning skills. Dr. Valerie Shute at the Educational Testing Service presents the following simple example of a hypothetical arithmetic tutoring system. Imagine that three learners, Jack, Jill, and Bill, are presented with addition problems that they answer as follows:

$\begin{array}{r} 22 \\ + 39 \\ \hline \textcircled{5}1 \end{array}$	$\begin{array}{r} \overset{1}{2}2 \\ + 39 \\ \hline \textcircled{1}61 \end{array}$	$\begin{array}{r} \overset{1}{2}2 \\ + 39 \\ \hline \textcircled{6}2 \end{array}$
$\begin{array}{r} 46 \\ + 37 \\ \hline \textcircled{7}3 \end{array}$	$\begin{array}{r} \overset{1}{4}6 \\ + 37 \\ \hline \textcircled{1}83 \end{array}$	$\begin{array}{r} \overset{1}{4}6 \\ + 37 \\ \hline \textcircled{8}5 \end{array}$
Jack	Jill	Bill

Figure 2

Diagnosing misconceptions that cause different types of math errors

Though all three students answered incorrectly, different underlying misconceptions caused each person's errors: Jack fails to carry, Jill always carries (sometimes unnecessarily), and Bill has trouble with single-digit addition. In this example, the student supplies an answer to the problem, and the tutoring system infers the student's misconceptions from this answer. By maintaining and referring to a detailed model of each student's many strengths and weaknesses, the ITS can provide carefully selected instruction and practice opportunities to make the most of the student's learning time.

In more complex domains, the tutoring system can analyze each learner's actions during simulated scenarios to estimate their knowledge and skills. For example, the Tactical Action Officer (TAO) ITS, developed for the U.S. Navy by Stottler Henke, teaches the use of tactical rules of engagement in realistic scenarios to the officers who control the ship's weapons and sensors. This system applies finite state machines, specified graphically, that look for sequences of actions and states that indicate principles the student does or doesn't understand. The simulation window in the left part of Figure 1 lets each student command the ship in simulated battles. The report card at right is displayed at the end of the scenario. It lists appropriate and inappropriate actions carried out by the student, along with associated principles. Red bulleted items describe actions the learner performed incorrectly, and green bullet items describe correct actions.

Intelligent Tutoring Systems – Using AI to Improve Training Performance and ROI

The **instructor model** encodes instructional methods that are appropriate for the target domain and the learner. Based on its knowledge of a person's skill strengths and weaknesses, participant expertise levels, and student learning styles, the instructor model selects the most appropriate instructional intervention. For example, if a student has been assessed a beginner in a particular procedure, the instructor module might show some step-by-step demonstrations of the procedure before asking the user to perform the procedure on his or her own. It may also provide feedback, explanations, and coaching as the participant performs the simulated procedure. It might even pose questions, using Socratic teaching methods, to encourage students to reflect upon their actions and reasoning. As a learner gains expertise, the instructor model may "decide" to present increasingly complex scenarios. It may also decide to take a back seat and let the person explore the simulation freely, intervening with explanations and coaching only upon request. Additionally, the instructor model may also choose topics, simulations, and examples that address the user's competence gaps.

Authoring tools reduce costs

Like multimedia authoring tools, ITS authoring tools can simplify the development of tutoring systems. One pioneer is a powerful tool developed by Douglas Towne and Allen Munro at the University of Southern California. Their systems incorporate simulation, graphics, and intelligent tutoring capabilities to teach equipment maintenance and other procedural tasks.

Other ITS authoring tools are currently being developed at academic and commercial institutions. For example, Stottler Henke developed the [SimBionic® intelligent agent toolkit](http://www.simbionic.com) (www.simbionic.com) that lets instructors and subject matter experts specify intelligent simulation behaviors quickly and easily, without programming. These behaviors specify how simulated characters, systems, and other entities respond to user actions, state conditions, and simulated events. Traditional development methods require simulation behaviors to be coded within software by programmers, so it is difficult for subject matter experts to specify these behaviors directly. By contrast, the SimBionic authoring tool presents behaviors graphically, so

experts can create, edit, and review them easily, without programming. Graphical representations can be understood by experts and by software programmers alike, so they can speak the same language, resulting in more effective collaboration.

Another ITS authoring tool developed by Stottler Henke, called the Task Tutor Toolkit™, is designed for technical training tasks. It lets trainers enhance training simulators with automated hinting and performance evaluation, without programming. Like SimBionic, Task Tutor Toolkit offers a visual, graphical interface, so subject matter experts who are not computer programmers can directly create content for the ITS.

ITS is still perceived by many as a technology of the future, but the rapid evolution of learning software and artificial intelligence is making it a viable option. ITSs for technical training is a particularly promising field, and Web-based ITS will soon add substantial improvements to that popular training medium.

So while learning will continue to be its own reward, businesses can also expect to reap greater rewards from their training investments in the future, as they embrace and deploy powerful new technologies such as ITS. The beneficiaries of this learning technology evolution will be organizations' employees, customers, partners – and their bottom lines.

For additional information about intelligent tutoring systems, please contact:

James Ong
ong@stottlerhenke.com
(650) 931-2710

Sowmya Ramachandran
sowmya@stottlerhenke.com
(650) 931-2716

For additional information about Stottler Henke's intelligent tutoring systems, please visit our web site at: www.stottlerhenke.com.