Adaptive E-Learning

Valerie Shute

Educational Testing Service Princeton, NJ

Brendon Towle

Thomson NETg Naperville, IL

It has long been known that differences among individuals have an effect on learning. Dick Snow's research on aptitude–treatment interactions (ATIs) was designed to investigate and quantify these effects, and more recent research in this vein has clearly established that these effects can be quantified and predicted. Technology has now reached a point where we have the opportunity to capitalize on these effects to the benefit of learners. In this article, we review some of the demonstrated effects of ATIs, describe how ATI research naturally leads to adaptive e-learning, and describe one way in which an adaptive e-learning system might be implemented to take advantage of these effects.

It is usually and rightly esteemed an excellent thing in a teacher that he should be careful to mark diversity of gifts in those whose education he has undertaken, and to know in what direction nature inclines each one most. For in this respect there is an unbelievable variety, and types of mind are not less numerous than types of body. (Quintilian, ca. 90 A.D.)

Acknowledging the important relation between individual differences and education has a very long history. However, simply acknowledging versus systematically testing this relation are two quite different things. Together with Lee Cronbach, Dick Snow formalized this interaction, consequently revolutionizing the thinking and researching of human abilities in the 1970s. Snow's primary research agenda focused on how individual differences in aptitudes played out in different educational settings. This received worldwide attention in the classic book on aptitude–treatment interactions (ATIs; Cronbach & Snow, 1977).

Snow was steadfast in his belief that the psychology of human differences is fundamental to education. He also acknowledged that designers of policy and practice often ignore the lessons of differential psychology by trying to impose a "one-size-fits-all" solution even though individuals are different. His work sought to change that fact—to promote educational improvement for all. This quest, across the years, has been joined by scores of supporters who have been motivated by him, either directly—as students and colleagues—or indirectly—through his writings. The first author of this article was fortunate to have had both direct and indirect Snow influences for almost 2 decades. And his influence continues, currently manifest in a research and development stream called *adaptive e-learning*.

As e-learning matures as an industry and a research stream, the focus is shifting from developing infrastructures and delivering information online to improving learning and performance. The challenge of improving learning and performance largely depends on correctly identifying characteristics of a particular learner. Examples of relevant characteristics include incoming knowledge and skills, cognitive abilities, personality traits, learning styles, interests, and so on (Shute, 1994; Snow, 1989, 1994). For instruction to be maximally effective, it should capitalize on these learner characteristics when delivering content. Instruction can be further improved by including embedded assessments, delivered to the student during the course of learning. Such assessments can provide the basis for diagnosis and subsequent instruction (i.e., presenting more of the same topic, remediating the current topic, or introducing a new topic). In short, enhancing learning and performance is a function of *adapting* instruction and content to suit the learner (for an overview on this topic, see Shute, Lajoie, & Gluck, 2000).

Requests for reprints should be sent to Valerie Shute, Educational Testing Service, Rosedale Road, Princeton, NJ 08541. E-mail: vshute@ets.org

The effectiveness of e-learning may be gauged by the degree to which a learner actually acquires the relevant knowledge or skill presented online.¹ This acquisition is generally regarded as a constructive activity where the construction can assume many forms; thus e-learning environments should be flexible enough to accommodate various constructive activities (see Shute, 1994, for more on this topic). Moreover, individuals differ in how they learn as well as what they learn, and different outcomes of learning (e.g., conceptual understanding) reflect differences in general learning processes (e.g., inductive reasoning skill), specific learning processes (e.g., attention allocation), and incoming knowledge and skill. The definition of learning used here is that learning is a process of constructing relations. These relations become more complex, and at the same time more automatic, with increased experience and practice. Accurate evaluation of e-learning is a nontrivial task because it requires both correct measurement of learning processes and outcomes.

Technology has now advanced to the point where we can begin to implement laboratory-based adaptive instructional techniques on the Internet (e.g., differential sequencing of knowledge and skill learning objects depending on learners' perceived needs). That is, concurrent advances in cognitive science, psychometrics, and technology are beginning to make it possible to assess higher level skills (Hambleton, 1996; Mislevy, Steinberg, & Almond, 1999) and to do so effectively and efficiently. Furthermore, in contrast with paper-and-pencil multiple-choice tests, new assessments for complex cognitive skills involve embedding assessments directly within interactive, problem-solving, or open-ended tasks (e.g., Bennett & Persky, 2002; Minstrell, 2000; Mislevy, Steinberg, Almond, Haertel, & Penuel, 2001). This information can then be used to build and enhance models of cognition and learning that can further inform and guide the assessment and instructional design processes.

The story we convey in this article illustrates the natural evolution of Snow's research, embodied in adaptive e-learning environments. We begin laying the foundation by describing ATI research. This is followed by an overview of specific components that go into adaptive e-learning systems. We conclude with our thoughts on the future of education and training.

APTITUDE-TREATMENT INTERACTION RESEARCH

Theoretical Basis

Snow approached his research with the combined perspectives of a differential psychologist, experimental psychologist, personality psychologist, and cognitive scientist. Of particular interest to Snow was the match between individual differences and learning environments and how variation in learning environments elicited or suited different patterns of aptitudes. Such relations are called *aptitude–treatment interactions* (ATIs), where *aptitude* is defined in the broadest sense as a person's incoming knowledge, skills, and personality traits. *Treatment* refers to the condition or environment that supports learning (Cronbach & Snow, 1977).

The goal of ATI research is to provide information about learner characteristics that can be used to select the best learning environment for a particular student to optimize learning outcome. Although hundreds of studies have been explicitly conducted to look for ATIs (especially during the 1960s and 1970s), it has been very difficult to empirically verify learner by treatment interactions. In the aforementioned 1977 book on this topic, Cronbach and Snow (1977) provided an excellent review of ATI research being carried out during that time. It is obvious from reading their book now that a major problem with those older ATI studies concerned data "noisiness." That is, experimental data obtained from classroom studies contained plentiful, extraneous, uncontrolled variables, such as differing teacher personalities, instructional materials, and classroom environments. Recently, however, there has been renewed interest in examining the ATI issue using computers as controlled learning environments (e.g., Maki & Maki, 2002; Sternberg, 1999). We now summarize one study where each part of an ATI is described and measured in a controlled manner.

Empirical Support

Individuals come to any new learning task with (often large) differences in prior knowledge and skill, learning style, motivation, cultural background, and so on. These qualities affect what is learned in an instructional setting. For this example, we focus on a particular learning style measure described in Shute (1993), namely exploratory behavior, as evidenced during interaction with an intelligent tutoring system (ITS) instructing principles of electricity. During the course of learning, a student's main goal was to solve progressively more difficult problems with direct current circuits presented by the ITS. However, at any given time, he or she was free to do other things in the environment, such as read definitions of concepts, take measurements on the online circuit, or change component values (e.g., voltage source, resistors). All explorations were optional and self-initiated. To quantify an individual's exploratory behavior, proportions were created-the ratio of time spent engaged in exploratory behavior divided by the total time on the tutor. This was necessary to control for differential tutor completion times, which ranged from 5 to 21 hr.

In addition to aptitude, the learning environment (i.e., treatment condition) may also influence learning outcome.

¹Technically and generally, e-learning can refer to any learning activity that largely involves technology for its presentation. We more narrowly define e-learning as that taking place in front of a computer that is connected to the Internet.

One way learning environments differ is in the amount of student control supported during the learning process. This can be viewed as a continuum ranging from minimal (e.g., rote or didactic environments) to almost complete control (e.g., discovery environments). Two opposing perspectives, representing the ends of this continuum, have arisen in response to the issue of optimal learning environment for instructional systems. One approach is to develop a straightforward learning environment that directly provides information to the learner; the other requires the learner to derive concepts and rules on his or her own. The disparity between positions becomes more complicated because the issue is not just which is the better learning environment, but rather, which is the better environment for what type or types of persons-an ATI issue. That issue motivated this controlled experiment, and the design of two environments representing the ends of this control continuum-rule application and rule induction.

These two environments were created from an ITS that teaches basic principles of electricity as a complex but controlled learning task. They differ only in the feedback delivered to the student. For instance, in the rule-application environment, feedback clearly states the variables and their relations for a given problem. This was communicated in the form of a rule, such as "The principle involved in this kind of problem is that current before a resistor is equal to the current after a resistor in a parallel net." Learners then proceeded to apply the rule in the solution of related problems. In the rule-induction environment, the tutor provides feedback, which identifies the relevant variables in the problem, but the learner must induce the relations among those variables. For instance, the computer might give the following comment, "What you need to know to solve this type of problem is how current behaves, both before and after a resistor, in a parallel net." Learners in the rule-induction condition, therefore, generate their own interpretations of the functional relations among variables comprising the different rules.

Four posttests measured a range of knowledge and skills acquired from the tutor, and all tests were administered online after a person finished the tutor.² The first test measured declarative knowledge of electrical components and devices and consisted of both true or false and multiple-choice questions. The second posttest measured conceptual understanding of Ohm's and Kirchhoff's laws. No computations were required, and all questions related to various circuits. The third posttest measured procedural skill acquisition. Computations were required in the solution of problems. The student would have to know the correct formula (e.g., voltage = current \times resistance), fill in the proper numbers, and solve. Finally, the fourth test measured a person's ability to generalize knowledge and skills beyond what was taught by the tutor. These problems required

computations. The experiment involved over 300 paid participants, mostly men, all high school graduates, 18 to 28 years old, and with no prior electronics instruction or training. On their ar-

107

rival, they were randomly assigned to one of the two environments (rule-application vs. rule-induction), and both versions permitted learners to engage in the optional, exploratory behaviors described previously. Exploratory behaviors were monitored by the computer and later quantified for post hoc ATI analysis. Although we hypothesized that the inductive environment would support (if not actively promote) the use of exploratory behaviors, results showed no differences among environments for exploratory behavior. Within each environment, however, there were wide individual differences on the exploratory learning dimension.

Learning outcome was defined as the percentage of correct scores on the four tests, combined into a single "outcome" factor score. An interaction was hypothesized-that learners evidencing greater exploratory behavior would learn better if they had been assigned to the inductive environment, and less exploratory learners would benefit from the more structured application environment. Results supported this hypothesis, showing a significant ATI (see Figure 1).

Implications

What are the ramifications of these findings? If we were able to map the variety of aptitudes or trait complexes (e.g., Ackerman, 1996, 2003) to associated learning environments or sequences of instructional components, we would be able to adapt and hence customize instruction for any given learner. For instance, Ackerman (1996; Ackerman & Heggstad, 1997) compiled ability-interest, ability-personality, and interest-personality correlates to support his more general process, personality, interests, and knowledge the-



FIGURE 1 Exploratory behavior by learning environment.

²We also used four matched pretests that served as covariates in the subsequent data analysis.

ory. Analysis has refined four cross-attribute (ability-interest-personality) trait complexes: social, clerical-conventional, science-math, and intellectual-cultural. The psychological importance of these trait complexes is similar to Snow's (1991) aptitude complexes (i.e., ability and interest constellations for classifying educational treatments). In general, the result, subject to empirical verification, would undoubtedly be to enhance the effectiveness, efficiency, and enjoyment of the learning experience. This is the premise underlying adaptive e-learning, which we now discuss.

ADAPTIVE E-LEARNING

The key idea to keep in mind is that the true power of educational technology comes not from replicating things that can be done in other ways, but when it is used to do things that couldn't be done without it. (Thornburg, as cited in National Association of State Boards of Education Study Group [NASBE], 2001)

Examples of e-learning on the Internet today are, too often, little more than lecture notes and some associated links posted in HTML format. However, as noted in the previous quote, the true power of e-learning comes from the exploitation of the wide range of capabilities that technologies afford. One of the most obvious is to provide assessments and instructional content that adapt to learners' needs or desires. This would comprise an online, real-time application of ATI research. Other effective technology uses include providing simulations of dynamic events, opportunities for extra practice on emergent skills, as well as the presentation of multimedia options.

The goal of adaptive e-learning is aligned with exemplary instruction: delivering the right content, to the right person, at the proper time, in the most appropriate way—*any time, any place, any path, any pace* (NASBE, 2001). We now focus our attention on what would be needed to accomplish this lofty goal in an e-learning context. Following are the necessary ingredients in an e-learning system. Each model is briefly introduced, then followed by a more in-depth description per module.

Components of E-Learning

The *content model* houses domain-related bits of knowledge and skill, as well as their associated structure or interdependencies. This may be thought of as a *knowledge map* of what is to be instructed and assessed, and it is intended to capture and prescribe important aspects of course content, including instructions for authors on how to design content for the model. A content model provides the basis for assessment, diagnosis, instruction, and remediation. In relation to the ATI study described earlier, the content model may be likened to the hierarchical array of knowledge and skills associated with Ohm's and Kirchhoff's laws.

The *learner model* represents the individual's knowledge and progress in relation to the knowledge map, and it may include other characteristics of the learner as a learner. As such, it captures important aspects of a learner for purposes of individualizing instruction. This includes assessment measures that determine where a learner stands on those aspects.

The *instructional model* manages the presentation of material and ascertains (if not ensures) learner mastery by monitoring the student model in relation to the content model, addressing discrepancies in a principled manner, and prescribing an optimal learning path for that particular learner. Information in this model provides the basis for deciding how to present content to a given learner and when and how to intervene.

Finally, the *adaptive engine* integrates and uses information obtained from the preceding models to drive presentation of adaptive learning content.

Content Model

The requirements for any content model fall into two categories: requirements of the delivery system and requirements of the learning content that is to be delivered. On the delivery side of the equation, we need a system that is content independent, robust, flexible, and scalable. *Content independent* means that the system can serve any content that is designed within the content requirements detailed later in this article. *Robust* means that it lives on the Internet and should be capable of delivering instruction to multiple users concurrently. *Flexible* implies adaptivity, requiring different types and sequences of content. And *scalable* means the system can adjust to increased demands, such as accommodating more components, more users, and so on.

On the content side of the equation, the content must be composed in such a way that the delivery system can adapt it to the needs of the particular learner. Each content aggregation will need to be composed of predictable pieces, so that the delivery system can know what to expect. This means that all of the content served by this delivery system will have to be built to the same specification. Issues such as grain size will vary, depending on the purpose or use of the content.

Learning objects. Fortunately, the need for content to be built to the same specifications dovetails nicely with current industry research and development associated with learning objects (e.g., see IEEE LTSC, 2003; IMS Global Learning Consortium, 2001). What are learning objects? Like Lego blocks, learning objects (LOs) are small, reusable components—video demonstrations, tutorials, procedures, sto-

ries, assessments, simulations, or case studies. However, rather than use them to build castles, they are used to build larger collections of learning material. LOs may be selectively applied, either alone or in combination, by computer software, learning facilitators, or learners themselves, to meet individual needs for learning or performance support. For example, all of Thomson NETg's currently delivered educational content takes the form of LOs that have been assembled into courses, where each LO addresses a coherent subset of the educational objectives of the course.

In current industry practice, LOs are assembled into courses before the time that they are delivered. The arrangement of these LOs to achieve instructional goals is done during this assembly. These collections can be specified using the Sharable Content Object Reference Model (SCORM; Advanced Distributed Learning, 2001) specification for defining courses. Among other things, this specification defines a way to define the structure of a course in a simple, easy-to-author manner. Because SCORM is agnostic to the nature of the content, we can easily define the collections described previously.

Although current practice is to assemble the LOs before the collection is delivered, there is no reason why the LOs could not be assembled into a structure that would allow an intelligent adaptive system to reassemble them on the fly to meet the needs of the particular learner. The rest of this section will describe how an LO collection could be specified such that an adaptive system could present it according to the needs of the learner. The basic idea involves dividing the LO collection such that it contains subcollections, each of which teaches a particular skill or knowledge type. Each of the subcollections thus contains all of the instructional components necessary to teach that skill. Using this hierarchy, the adaptive system can first decide what needs to be taught and then decide how to teach it, as we describe in the following pages.

Knowledge structures. The purpose of establishing a knowledge structure as part of the content model in any e-learning system is that it allows for dependency relations to be established. These, in turn, provide the basis for the following: assessment (What's the current status of a particular topic or LO?), cognitive diagnosis (What's the source of the problem, if any?), and instruction or remediation (Which LOs need to be taught next to fix a problem area or present a new topic?).

Each element (or node) in the knowledge structure may be classified in terms of different types of knowledge, skill, or ability. Some example knowledge types include

• *Basic knowledge* (BK): This includes definitions, examples, supplemental links (jpgs, avis, wavs), formulas, and so on, and addresses the *What* part of content.

• *Procedural knowledge* (PK): This defines step-by-step information, relations among steps, subprocedures, and so on, and addresses the *How* part of content.

• *Conceptual knowledge* (CK): This refers to relational information among concepts and the explicit connections with BK and PK elements, draws all into a "big picture" and addresses the *Why* part of content.

Restricting a node in the knowledge structure (again, note that each node has an associated collection of LOs) to a single knowledge type helps ensure the course is broken down to an appropriate grain size, by limiting the scope of what can be in any single node. This restriction also suggests different strategies for the authoring of instruction and assessment: Different knowledge types require different strategies. We suggest the following guidelines (for more on this topic, see Shute, 1995): BK instruction should involve the introduction of new definitions and formulas in a straightforward, didactic manner, whereas BK assessment relates to measuring the learner's ability to recognize or produce some formula, basic definition, rule, and so on. PK instruction should occur within the context of experiential environments where the learner can practice doing the skill or procedure (problem solving), whereas PK assessment relates to the learner's ability to actually accomplish some procedure or apply a rule, not simply recognize those things. Finally, CK instruction typically occurs after the learner has been presented with relevant base information (BK-PK), and then the big picture may be presented, either literally or via well-designed analogies, case studies, and so on. CK assessment refers to a learner being able to transfer BK and PK to novel areas, explain a system or phenomenon, predict some outcome, or strategize. The outcome tests described earlier in relation to the electricity tutor study exemplify each of these outcome types.

A simplified network of nodes and associated knowledge types is shown in Figure 2. Each node has an associated collection of LOs that teach or assess a certain component of a concept or skill. In summary, we posit different knowledge types associated with their own special way of being instructed and assessed. So now the questions are: How do we optimally assess and diagnose different outcome types, and what happens after diagnosis? Before answering those questions, we now present the learner model—the repository of information concerning the learner's current status in relation to the various LOs (i.e., domain-related proficiencies).

Learner Model

The learner model contains information that comes from assessments and ensuing inferences of proficiencies. That information is then used by the system to decide what to do next. In the context of adaptive e-learning, this decision relates to customizing and hence optimizing the learning experience. Obviously, a critical component here is the validity and reliability of the assessment. One idea is to employ what is called the evidence-centered design approach (e.g., Mislevy, Steinberg, & Almond, 1999) to assessment. This al-



FIGURE 2 Sample knowledge hierarchy.

lows an instructional designer (or whomever) to (a) define the claims to be made about the students (i.e., the knowledge, skills, abilities, and other traits to be measured), (b) delineate what constitutes valid evidence of the claim (i.e., student performance data demonstrating varying levels of proficiency), and (c) create assessment tasks that will elicit that evidence. Evidence is what ties the assessment tasks back to the proficiencies, and the entire process is theory driven, as opposed to a more common data- or item-driven manner.

Assessing the learner. The first issue concerns just what is to be assessed. There are actually two aspects of learners that have implications for adaptation: (a) *domain-dependent information*—this refers to knowledge assessment via pretest and performance data to allow the system to initialize a learner model in relation to content and LOs, eliminate those already "known," and focus instruction or assessment (or both) on weak areas; and (b) *domain-independent information*—this relates to learner profile data (e.g., cognitive abilities or personality traits) and allows the system to pick and serve optimal LO sequences and formats.

Assessing specific learner traits indicates particular kinds of content delivery in relation to either differential sequencing of topics or providing appropriate, alternative formats and media. In relation to differential sequencing, a typical, valid instructional system design sequence consists of these steps: (a) Present some introductory material, (b) follow with the presentation of a rule or concept, (c) provide a range of illustrative examples, (d) give liberal practice opportunities, and (e) summarize and call for reflection. Thus, across topics, one general sequencing rule may be to serve easier topics before more difficult ones. And within a topic, a general rule may be the default delivery of certain LOs: introduction, body (rule or concept, followed by examples), interactivity (explorations, practice, and explicit assessments of knowledge or skill), and reflection (summary).

Alternatively, suppose you knew that an individual was high on an inductive reasoning trait. The literature (e.g., Catrambone & Holyoak, 1989) has suggested that those learners perform better with examples and practice preceding the concept. On the other hand, learners with low inductive reasoning skills perform better when the concept is presented at the outset of the instructional sequence. This has direct implications for the sequencing of elements within a given topic.

Now, consider a learner assessed as possessing low working memory capacity. One instructional prescription would be to present this individual with smaller units of instruction (Shute, 1991). And going back to our example at the beginning of this story (depicted in Figure 1), an individual possessing an exploratory learning style suggests a less structured learning experience compared to a person with a less exploratory learning style (Shute, 1993).

In Figure 3, we can see how the parts come together (modified from Quinn, 2001). The bottom row of the figure shows the typical way of serving content based on inferred gaps in a learner's knowledge structure. This is also known as *microadaptation*, reflecting differences between the learner's knowledge profile and the expert knowledge structure embodied in the content model. "Instructional rules" then determine which knowledge or skill element should be selected next (i.e., selecting from the pool of nonmastered objects). The top row shows an additional assessment, representing another way to adapt instruction based on learners' cognitive abilities, learning styles, personality, or whatever else is deemed relevant. This is also known as *macroadaptation* and provides information about how to present the selected knowledge or skill chunk.

Instructional Model

There are several general and specific guidelines for systematic approaches to instructional design, such as those described by Robert Gagné. But how do we progress from guidelines to determining which LOs should be selected, and why? To answer this question, after delineating the guidelines presented in Gagné's (1965) book entitled *The Condi*- *tions of Learning,* we describe a student modeling approach that implements some of these instructional ideas within the context of an ITS.

The following represents an abridged version of Gagné's "events of instruction" (see Kruse, 2000), along with the corresponding cognitive processes. These events provide the necessary conditions for learning and serve as the basis for designing instruction and selecting appropriate media (Gagné, Briggs, & Wager, 1992). We include them here as they offer clear, obvious guidelines for designing good e-learning environments.

- 1. Gain the learner's attention (reception).
- 2. Inform the learner of the objectives (expectancy).
- 3. Stimulate recall of prior learning (retrieval).
- 4. Present the learning stimulus (selective perception).
- 5. Provide learning guidance (semantic encoding).
- 6. Elicit appropriate performance (responding).
- 7. Provide feedback (reinforcement).
- 8. Assess the learner's performance (retrieval).
- 9. Enhance retention and transfer (generalization).

Applying Gagné's (Gagné, Briggs, & Wager, 1992) nine-step model to an e-learning program is a good way to facilitate learners' successful acquisition of the knowledge and skills presented therein. In contrast, an e-learning program that is replete with bells and whistles, or provides unlimited access to Web-based documents, is no substitute for sound instructional design. Although those types of programs might be valuable as references or entertainment, they will not maximize the effectiveness of information processing or learning.

In addition to the specific prescriptions mentioned previously, there are a few key presumptions and principles for instructional design that should be considered when designing an e-learning system. In general, these include the following: Knowledge is actively constructed, multiple representations for a concept or rule are better than a single one, problem-solving tasks should be realistic and complex, and opportunities for learners to demonstrate performance of activities promoting abstraction and reflection should be provided.

In terms of the relevant features of the adaptive system, this means several things. First, the activities presented to the student should involve the creation and manipulation of representations. If the student is expected to have a mental model that corresponds to the representation, he or she needs to be actively involved in the creation or manipulation of the representations.

Second, the content should be designed with multiple representations for a concept or a rule. This serves two purposes: It allows the adaptive engine (described later) to provide the student with the single representation that best matches the student's aptitude profile, while simultaneously giving the engine additional representations to present in the event that the student fails to master or acquire the topic the first time. These multiple representations should include different visual representations (textual vs. graphical, or different graphical representations of the same concept) as well as different styles of conceptual explanation.

Third, the student should be provided with a final learning activity that encourages reflection and integration of the knowledge learned into the body of knowledge as a whole. Finally, the system should incorporate enough support and help so that the student can spend time learning the material and not the system. This is simply to ensure that as much of the student's cognitive effort as possible goes into learning the material being presented, and not into learning the system that is doing the presenting.

How do we move from these general and specific guidelines to determining which learning object or objects should be selected, and why? One solution is to employ something



FIGURE 3 Learning management system framework including two types of assessments.

like Student Modeling Approach to Responsive Tutoring (SMART; Shute, 1995), a principled approach to student modeling. It works within an instructional system design where low-level knowledge and skill elements are identified and separated into the three main outcome types previously mentioned (i.e., BK, PK, CK). As the student moves through an instructional session, LOs (i.e., the online manifestations of the knowledge and skill elements) are served to instruct and assess. Those knowledge elements showing values below a preset mastery criterion become candidates for additional instruction, evaluation, and remediation, if necessary. Remediation is invoked when a learner fails to achieve mastery during assessment, which follows or is directly embedded within the instructional sequence.

SMART includes capabilities for both micro- and macroadaptation of content to learners, mentioned earlier, and is based on a distinction originally made by Snow (1989). Basically, microadaptation relates to the domain-dependent learner model (i.e., the individual knowledge profile in Figure 3) and is the standard approach, representing emerging knowledge and skills. In this case, the computer responds to updated observations with a modified curriculum that is minutely adjusted, dependent on individual response histories during instructional sessions. Macroadaptation refers to the domain-independent learner model (i.e., the individual learner profile in Figure 3), representing an alternative approach. This involves assessing students prior to as well as during their use of the system, focusing mainly on general, long-term aptitudes (e.g., working memory capacity, inductive reasoning skill, exploratory behavior, impulsivity) and their relations to different learning needs.

An alternative approach to student modeling includes using Bayesian inference networks (BINs) to generate estimates of learner proficiencies in relation to the content (e.g., Mislevy, Almond, Yan, & Steinberg, 1999). Both the SMART and BIN approaches are intended to answer the following questions: (a) What is the learner's current mastery status of a topic, and (b) what is the nature and source of the learner's problem, if any? Typical ways of evaluating success (e.g., pass or fail, or correct solution of two consecutive problems) do not offer the degree of precision needed to go beyond assessment into cognitive diagnosis. Both SMART and BINs provide probabilistic mastery values associated with nodes or topics (regardless of grain size). With regard to the instructional decision about what should subsequently be presented, the knowledge structure, along with an indication of how well a learning objective is attained, informs the adaptive engine of the next recommended bit or bits of content to present.

Adaptive Engine

Given the content model, the learner model, and the instructional model, the fundamentals of the adaptive engine are fairly simple. The first step involves selecting the node (topic) to present, based on a diagnosis of the student's knowledge needs. The next step involves deciding which LO or LOs within that node to present, sequenced or flavored according to the characteristics and needs of the particular learner. The presentation of LOs is continued until the student has mastered the topic or node, and the node selection process is then repeated until all nodes have been mastered.

Although this is a fairly simple overview, the actual process is obviously more complicated. We will examine each part of the process (selecting a node, and then presenting the content within the node) separately.

In our solution, selecting a node, in the general case, is a fairly simple exercise; the engine can simply choose from the pool of nodes that have not been completed and whose prerequisites have been mastered. However, one additional feature of our structure is that a pretest can be generated on the fly, and assessment can incorporate a sequencing algorithm.

Recall that the LOs in each node have been categorized by their role in the educational process, and the authoring guidelines have restricted each learning object to a single role only. Because of this, for any collection of nodes, the system can create another collection of nodes that contains only the assessment tasks from the original collection. Presenting this new collection functions as a pretest. If the student passes the assessment without any presentation, he or she is presumed to have already mastered the associated content.

When the engine is presenting objects relating to a particular node, it uses a set of rules to drive the selection of individual LOs for presentation to the student. These rules examine the information contained in the student model, the student's interaction within the node so far, and the content model of each individual LO contained within the node. Using this information, the rules assign a priority to each LO within the node. Once the priority of every LO has been calculated (which occurs almost instantly), the LO with the highest priority is delivered to the student.

As an example of how this works for instructional objects, consider the following. An initial arbitrary weighting is assigned to every LO in the node. One rule states that if the student's interaction with the node is empty (i.e., the student is just beginning the node), then decrease the priority of every LO except those which fulfill the role of "introduction." This rule ensures that the default sequence provides for an introduction-type LO to be presented at the beginning of an instructional sequence. On the other hand, consider a learner who prefers a more contextualized learning experience, such as a learner characterized as very concrete and experiential. To handle that case, there is a rule that states that if the learner is "highly concrete" and "highly experiential," and if the learner's interaction with the node is empty, then increase the priority of associated assessment-task LOs. If the learner is not concrete and experiential, then the second rule has no effect; however, if he or she is, then the second rule overrides the first, and the learner sees an assessment task at the beginning of the instructional sequence.

The remaining rules work in an analogous fashion. That is, each one examines a set of conditions that has an associated instructional prescription and adjusts priorities on the appropriate LOs. Working together, all of the rules serve to provide the instructionally correct learning object for the student at every point of the student's interaction with the node.

One issue that should be addressed concerns the accuracy of the rule set: designing it such that it provides a natural and effective learning experience regardless of learner characteristics. One way to accomplish this is by using the techniques of genetic programming (GP; Koza, 1992) to improve the performance of the rule set. Research has shown that this technique is applicable to the design of rules for rule-based systems (e.g., Andre, 1994; Edmonds, Burkhardt, & Adjei, 1995; Tunstel & Jamshidi, 1996). The general idea is to treat each individual rule set as a single individual in the population of algorithms; the rule sets can then be evolved according to standard GP methods.

For our purposes, the interesting feature of GP is that it turns a design task (create a rule set that treats learners effectively) into a recognition task (determine how well a given rule set performs at treating learners). We believe that a large sample of learner data can be used to evaluate a learner's potential experience with a given rule set, and this can be used as the basis of the evaluation function that drives the GP approach. Further, we believe that the combination of human-designed rules with computer-driven evolution gives us a high likelihood of success and avoids many of the risks inherent in rule-based systems.

CONCLUSION

There are many reasons to pursue adaptive e-learning. The potential payoffs of designing, developing, and employing good e-learning solutions are great, and they include improved efficiency, effectiveness, and enjoyment of the learning experience. In addition to these student-centered instructional purposes, there are other potential uses as well, such as online assessments. Ideally, an assessment comprises an important event in the learning process, part of reflection and understanding of progress. In reality, assessments are used to determine placement, promotion, graduation, or retention. We advocate pursuing the ideal via online diagnostic assessments. As Snow and Jones (2001) pointed out, however, tests alone cannot enhance educational outcomes. Rather, tests can guide improvement-presuming they are valid and reliable-if they motivate adjustments to the educational system. There are clear and important roles for good e-learning programs here.

However, and as mentioned earlier, the current state of e-learning is often little more than online lectures, where educators create electronic versions of traditional printed student manuals, articles, tip sheets, and reference guides. Although these materials may be valuable and provide good resources, their conversion to the Web cannot be considered true teaching and learning. Instead of the page-turners of yesterday, we now have scrolling pages, which is really no improvement at all. Adaptive e-learning provides the opportunity to dynamically order the "pages" so that the learner sees the right material at the right time.

There are currently a handful of companies attempting to provide adaptive e-learning solutions (e.g., see LearningBrands.com, AdaptiveTutoring.com, and Learning Machines, Inc.). Further, adaptive e-learning has become a rather hot topic in the literature recently (Nokelainen, Tirri, & Silander, 2002; Kurhila, Miettinen, Sampson, Karagiannidis, & Kinshuk, 2002). However, many of these are not concerned with adaptive instruction at all; rather, they are concerned with adapting the format of the content to meet the constraints of the delivery device, or adapting the interface to the content to meet the needs of disabled learners. Of those that are concerned with adaptive instruction, most tend to base their "adaptivity" on assessments of emergent content knowledge or skill or adjustments of material based on "learner styles"—less suitable criteria than cognitive abilities for making adaptive instructional decisions.

We believe that the time is ripe to develop e-learning systems that can reliably deliver uniquely effective, efficient, and engaging learning experiences, created to meet the needs of the particular learner. The required ingredients in such a personalized learning milieu include rich descriptions of content elements and learner information, along with robust, valid mappings between learner characteristics and appropriate content. The result is adaptive e-learning, a natural extension of Snow's considerable contributions to the field of educational psychology.

ACKNOWLEDGMENT

We thank Aurora Graf, Jody Underwood, Irv Katz, and several anonymous reviewers for their helpful comments on this article.

REFERENCES

- Ackerman, P. L. (1996). A theory of adult intellectual development: Process, personality, interests, and knowledge. *Intelligence*, 22, 227–257.
- Ackerman, P. L. (2003). Aptitude complexes and trait complexes. *Educa*tional Psychologist, 38, 85–93.
- Ackerman, P. L., & Heggestad, E. D. (1997). Intelligence, personality, and interests: Evidence for overlapping traits. *Psychological Bulletin*, 121, 218–245.
- Advanced Distributed Learning. (2001). SCORM (version 1.2). Retrieved April 10, 2003, from http://www.adlnet.org/ADLDOCS/Other/SCORM_1.2_doc.zip
- Andre, D. (1994). Learning and upgrading rules for an OCR system using genetic programming. In *Proceedings of the First IEEE Conference on Evolutionary Computation* (Vol. 1, pp. 462–467). Piscataway, NJ: IEEE.

114 SHUTE AND TOWLE

- Bennett, R. E., & Persky, H. (2002). Problem solving in technology-rich environments. In Qualifications and Curriculum Authority (Ed.), Assessing gifted and talented children (pp. 19–33). London, England: Qualifications and Curriculum Authority.
- Catrambone, R., & Holyoak, K. J. (1989). Overcoming contextual limitations on problem-solving transfer. *Journal of Experimental Psychol*ogy: Learning, Memory, and Cognition, 15, 1147–1156.
- Cronbach, L. J., & Snow, R. E. (1977). Aptitudes and instructional methods: A handbook for research on interactions. New York: Irvington.
- Edmonds, A. N., Burkhardt, D., & Adjei, O. (1995). Genetic programming of fuzzy logic rules. In *Proceedings of the Second IEEE Conference on Evolutionary Computation* (Vol. 2, pp. 765–770). Piscataway, NJ: IEEE.
- Gagné, R. M. (1965). *The conditions of learning*. New York: Holt, Rinehart & Winston.
- Gagné, R. M., Briggs, L. J., & Wager, W. W. (1992). Principles of instructional design (4th ed.). Fort Worth, TX: Harcourt Brace.
- Hambleton, R. K. (1996). Advances in assessment models, methods, and practices. In D. C. Berliner & R. C. Calfee (Eds.), *Handbook of educational psychology* (pp. 889–925). New York: American Council on Education/Macmillan.
- IEEE LTSC. (2003). Learning object metadata. Retrieved April 10, 2003, from http://itsc.ieee.org/doc/wg12/LOM_1484_12_1_v1_Final_Draft.pdf
- IMS Global Learning Consortium. (2001). Learning resource metadata (version 1.2.1). Retrieved April 10, 2003, from http://www.imsglobal.org/metadata/index.cfm
- Koza, J. (1992). Genetic programming. Cambridge, MA: MIT Press.
- Kruse, K. (2000). *Web rules*. (20000. Retrieved April 10, 2003, from http://www.learningcircuits.org/feb2000/feb2000_webrules.html
- Maki, W. S., & Maki, R. H. (2002). Multimedia comprehension skill predicts differential outcomes of web-based and lecture courses. *Journal of Experimental Psychology: Applied*, 8, 85–98.
- Minstrell, J. (2000). Student thinking and related instruction: Creating a facet-based learning environment. In J. Pellegrino, L. Jones, & K. Mitchell (Eds.), *Grading the nation's report card: Research for the evaluation of NAEP*. Washington, DC: National Academy Press.
- Mislevy, R. J., Almond, R. G., Yan, D., & Steinberg, L. S. (1999). Bayes nets in educational assessment: Where do the numbers come from? In K. B. Laskey & H. Prade (Eds.), *Proceedings of the Fifteenth Conference on Uncertainty* in Artificial Intelligence (pp. 437–446). San Francisco: Kaufmann.
- Mislevy, R. J., Steinberg, L. S., & Almond, R. G. (1999). On the roles of task model variables in assessment design (CSE Tech. Rep. No. 500). Los Angeles: University of California, Center for the Study of Evaluation, Graduate School of Education & Information Studies.
- Mislevy, R. J., Steinberg, L. S., Almond, R. G., Haertel, G., & Penuel, W. (2001). Leverage points for improving educational assessment (CSE Tech. Rep. No. 534). Los Angeles: University of California, Center for Studies in Education/CRESST.

- National Association of State Boards of Education Study Group. (2001). Any time, any place, any path, any pace: Taking the lead on e-learning policy. Retrieved April 10, 2003, from http://www.nasbe.org/e_Learning.html
- Nokelainen, P., Tirri, H., Kurhila, J., Miettinen, M., & Silander, T. (2002). Optimizing and profiling users online with Bayesian probabilistic modeling. In *Proceedings of the Networked Learning 2002 Conference* (Berlin, Germany, May 2002). The Netherlands: NAISO Academic Press.
- Quinn, C. (2001). Framework for a learning management system [Slide]. Unpublished PowerPoint presentation, KnowledgePlanet.com, Emeryville, CA.
- Sampson, D., Karagiannidis, C., & Kinshuk. (2002). Personalised learning: Educational, technological and standardisation perspectives, interactive educational multimedia (on-line journal ISSN 1576–4990), Special Issue on Adaptive Educational Multimedia, 4 (invited paper), April 2002.
- Shute, V. J. (1991). Who is likely to acquire programming skills? Journal of Educational Computing Research, 7(1), 1–24.
- Shute, V. J. (1993). A comparison of learning environments: All that glitters. In S. P. Lajoie & S. J. Derry (Eds.), *Computers as cognitive tools* (pp. 47–74). Hillsdale, NJ: Lawrence Erlbaum Associates, Inc.
- Shute, V. J. (1994). Learning processes and learning outcomes. In T. Husen & T. N. Postlethwaite (Eds.), *International encyclopedia of education* (2nd ed., pp. 3315–3325). New York: Pergamon.
- Shute, V. J. (1995). SMART: Student Modeling Approach for Responsive Tutoring. User Modeling and User-Adapted Interaction, 5, 1-44.
- Shute, V. J., Lajoie, S. P., & Gluck, K. A. (2000). Individualized and group approaches to training. In S. Tobias & J. D. Fletcher (Eds.), *Training* and retraining: A handbook for business, industry, government, and the military (pp. 171–207). New York: Macmillan.
- Snow, C. E., & Jones, J. (2001, April 25). Making a silk purse. *Education Week Commentary*.
- Snow, R. E. (1989). Toward assessment of cognitive and conative structures in learning. *Educational Researcher*, 18, 8–14.
- Snow R. E. (1991). The concept of aptitude. In R. E. Snow & D. Wiley (Eds.), Improving inquiry in social science (pp. 249–284). Hillsdale, NJ: Lawrence Erlbaum Associates, Inc.
- Snow, R. E. (1994). Abilities in academic tasks. In R. J. Sternberg & R. K. Wagner (Eds.), *Mind in context: Interactionist perspectives on human intelligence* (pp. 3–37). New York: Cambridge University Press.
- Sternberg, R. J. (1999). *Thinking styles*. New York: Cambridge University Press.
- Tunstel, E., & Jamshidi, M. (1996). On genetic programming of fuzzy rule-based systems for intelligent control. *International Journal of Intelligent Automation and Soft Computing*, 2, 273–284.