CAMBRIDGE

Cambridge University Press
978-0-521-76903-7 - Adaptive Technologies for Training and Education
Edited by Paula J. Durlach and Alan M. Lesgold
Excerpt
More information

Introduction

This book addresses the use of technology to provide training adapted to the individual needs of different trainees. The use of technology for training has become commonplace, as a way to increase training effectiveness at reduced cost. Pressures on time and resources have impelled organizations responsible for the education and training of large numbers of people to adopt technology-based training. Initially, technological approaches replicated classroom methods (mass instruction) and generally provided either no tuning of instruction to individual student needs, simple branching schemes, or mastery approaches in which instruction essentially was repeated until a mastery test was passed.

But as the use of computer-based instructional materials has grown, so has disenchantment with one-size-fits-all passive learning solutions and the inadequate mastery-assessment methods that often accompany them. More advanced technology-based learning environments can provide tailored and personal interaction opportunities and allow students to learn by doing, supported by feedback. Still, they might not meet the needs of a varied target audience because they typically fail to integrate prerequisite or remedial instruction, which may be required for a heterogeneous user population, and they are still primarily one-size-fits-all solutions, unable to adapt their behaviors to learner backgrounds.

The need for more adaptive training and instruction arises from today's fast pace of change, both in societal complexity and in the content that needs to be learned. Historically, a great deal of training and education was personalized. The elite had private tutors, and craftsmen learned through apprenticeship. Group approaches to education and training often originated from institutional needs (e.g., religious or military) to maintain their ranks and carry out their agendas.

Less specialized education for wider audiences was aimed mainly at enculturation, to make sure that everyone shared a common body of everyday knowledge and expectations, social values, and religion. In these cases, a relatively uniform approach to training or teaching worked pretty well because there was so much shared everyday experience that the instructor could assume a lot about what the learner already knew.

The environment for instruction and training in the developed world today is quite different from the homogeneous, closed, small-scale society or institution. Society is multicultural, and education and training are undertaken by students with varied experiences and motivations. This makes it much harder to know what prior knowledge or level of motivation any individual student might have. Inaccurate instructor knowledge of each student's current level of knowledge and skill can in turn can lead to large variations in instructional effectiveness and rates of mastery. One way to cope has been to stream students into different class tracks, where the pace of progress through the material (and the ultimate end state) differs across tracks. Although this still leaves the logistic problem of multiple tracks going on at once, it represents an improvement because groups are relatively homogeneous in the characteristics that impact learning. Today, however, when students come from so many different backgrounds, learning must adapt not only to learning pace, but also to differing underlying knowledge and differing prior experience – the ontological base of different students can easily be quite different.

Within institutions such as the military, there is the additional challenge that training needs can emerge on short notice, as the result of changes in world events and affairs, as well the emergence of new technologies. Personnel without recent deployment experience may require different training from those with firsthand knowledge of a particular situation or piece of equipment. Such heterogeneity may be difficult for an instructor to cope with. Moreover, the knowledge may be required by personnel not participating currently in formal education (e.g., already deployed personnel). Technology-based training has the potential to reach a wider audience, but it needs to be appropriately tailored to the learner's prior knowledge.

Developers of learning and training systems today realize that it can be beneficial for a system to encode information about each student and to use that information to adapt the training regimen to better support mastery on an individual-by-individual basis. The information about each student – the student model – typically includes some subset of the student's experiences, knowledge, skills, and abilities in the domain of instruction, measured prior to and/or during instruction. It can also include information about student traits (e.g., spatial ability) and states (e.g., confusion) that are relevant to setting an adaptive pedagogical strategy. Various parts of this information occur and suggest adaptations on different time scales. Parts of the student model may change even within an exercise, whereas other parts may endure for much longer periods, even years.

Training system behaviors adapted on the basis of the student model can include sequencing of content, selection of content, and the type and timing of coaching and feedback provided during the learning experience. For a system to adapt effectively (i.e., better promote mastery than a one-size-fits-all version), it needs both good rules about how and what to adapt (pedagogical model) and accurate data in the student model to guide these rules. The accuracy requirement implies the need for valid and sensitive measures of student knowledge, skills, abilities, traits, and states. In addition, effective adaptation requires a thorough understanding and systematization of the instructional content itself, as well as significant insights about the knowledge structures possessed by domain experts and the cognitive processes they use to solve problems in the domain.

Given all these requirements, the dearth of adaptive instructional technology in practical use is hardly surprising. Typically it has had high development costs and required multiple generations of tryout and refinement over years (for further discussion see Chapter 15 in this volume, by Lesgold). Only a few intelligent adaptive systems have made it past the development stage because of this. Mostly, they have been well-defined and relatively static educational domains (i.e., where the content remains stable), such as algebra and geometry, because this is where all the needed elements most readily coalesce (expert cognitive models, problems sets with known psychometric properties, systematized content, and step-based problems with objectively correct or incorrect solutions). The other real applications have been in areas where no deep training existed before and much was needed, as in the language and cultural training described by Johnson and colleagues in Chapter 14 of this volume.

The success of the extant applications and the findings that one-on-one human tutoring provides an advantage over classroom instruction suggest that adaptive training technology should produce superior learning outcomes compared with nonadaptive technology. The purpose of this volume is to provide an overview of the latest advancements in adaptive training technology and to provide a basis for determining what further advancements would be required to make this approach more amenable to wider practical usage.

Part I, "Adaptive Training Technology," provides the reader with a foundational understanding of adaptive training technology. In Chapter 1, Shute and Zapata-Rivera provide an overview of adaptive training technology: why to adapt, how to adapt, and what to adapt. In Chapter 2, VanLehn and Chi present a case study of how adaptive technology can produce accelerated learning. In Chapter 3, Brusilovsky reviews the application of adaptive techniques to Web-based systems, specifically adaptive educational hypermedia.

Part II, "Student Modeling Beyond Content Mastery," presents work focusing on student modeling. The title of the section was chosen to acknowledge the core role that students' domain knowledge and skill play in a student model, but also to suggest that other elements of student-associated data may contribute to adaptive strategies. Aleven, Roll, and Koedinger in Chapter 4 and Conati in Chapter 5, respectively, discuss their research on using adaptive training technology to foster metacognitive learning skills, such as help seeking and self-explanation. In Chapter 6, D'Mello and Graesser describe the interaction of learning and affect, the measurement of affect during learning, and the design of affect-sensitive intelligent tutors (see also Litman's Chapter 13). Concluding the section, the potential benefits and challenges involved in creating persistent student models – long-term models that the learner takes along from one training application to the next – are explored by Kay and Kummerfeld in Chapter 7.

As previously mentioned, adaptive intelligent training technology (e.g., intelligent tutors) has been successfully applied in well-defined domains and step-based problem solving. Many domains and fields requiring training are not so well-defined, nor characterized by step-based problems. Part III, "Experiential Learning and Ill-Defined Domains," turns the focus toward how adaptive training technology might handle domains that are less well-defined and/or where the employment of a less structured interface (compared with step-based problem solving) bestows a greater latitude of activities, and thus greater challenges for interpreting student state from overt behavior. In Chapter 8, Gonzalez describes models of dynamic decision making where there is uncertainty about when events will occur and alternative options unfold over time. She describes her research investigating the decision-making process and provides implications for the training of decision making for dynamic environments. Lynch, Ashley, Pinkwart, and Aleven in Chapter 9 explore what exactly it means for a domain or problem to be ill-defined, and the implications for education and training strategies. In Chapter 10, Lane and Wray discuss experiential training designed to promote the acquisition of social and intercultural skills and describe a framework for adapting the behavior of virtual humans to support this type of learning. Similarly, Flynn in Chapter 12 is concerned with the behavior of virtual characters, and describes how semantic-web technology could be harnessed to create virtual humans for training, based on "person ontologies" and "action agents."

Mangos, Campbell, Lineberry, and Bolton in Chapter 11 lay out the challenges of designing pedagogically sound experiential training and focus on methods to orchestrate the interplay of student assessment and content selection, which adaptive experiential training requires.

There are multiple ways in which the ability to use language supports training and education. This is obviously the case when the goal is to teach reading, writing, language skills, and tasks that are inherently language based (e.g., negotiation or legal argumentation). Additionally, language can play a crucial role in a teacher's assessment of student mastery and stimulation of student cognitive processing. Part IV, "Natural Language Processing for Training," includes two chapters providing examples of both of these functions. In Chapter 13, Litman discusses the role of natural language processing during interactive tutorial dialogues, not just for purposes of content understanding, but also for monitoring student affect. Johnson, Friedland, Watson, and Surface bring the issues of experiential learning and language processing together in Chapter 14, describing language and culture training systems aimed at adult learners for whom foreign language and intercultural competency must be learned to conduct their work.

The last group of chapters in Part V address various "Culminations" for adaptive training development. In Chapter 15, Lesgold presents lessons learned during the development of five generations of intelligent coached apprenticeship systems. In Chapter 16, Levchuk, Shebilske, and Freeman present the challenges of designing adaptive technology for team training. In Chapter 17, Bienkowski describes and advocates the Design-Based Research approach to the study of technology-based learning. This approach draws on engineering, software, and industrial design practices, such as agile design rapid prototyping, participatory design, and user-centered design. Finally, in Chapter 18, Durlach describes discussions held by contributors to this book (and others) regarding the current state and the future of adaptive training technology. Four topics were targeted for discussion: student models, pedagogical models for experiential learning, training efficiency, and military training applications.

Cambridge University Press
978-0-521-76903-7 - Adaptive Technologies for Training and Education
Edited by Paula J. Durlach and Alan M. Lesgold
Excerpt
More information

Part I

ADAPTIVE TRAINING TECHNOLOGY

Cambridge University Press
978-0-521-76903-7 - Adaptive Technologies for Training and Education
Edited by Paula J. Durlach and Alan M. Lesgold
Excerpt
More information

CHAPTER 1

Adaptive Educational Systems

Valerie J. Shute and Diego Zapata-Rivera

Introduction

Adaptive educational systems monitor important learner characteristics and make appropriate adjustments to the instructional milieu to support and enhance learning. The goal of adaptive educational systems, in the context of this chapter, is to create an instructionally sound and flexible environment that supports learning for students with a range of abilities, disabilities, interests, backgrounds, and other characteristics. The challenge of accomplishing this goal depends largely on accurately identifying characteristics of a particular learner or group of learners – such as type and level of knowledge, skills, personality traits, affective states – and then determining how to leverage the information to improve student learning (Conati, 2002; Park & Lee, 2004; Shute et al., 2000; Snow, 1989, 1994).

We present a general evidence-based framework for analyzing adaptive learning technologies. We then describe experts' thoughts on: (1) the variables to be taken into account when implementing an adaptive learning system (i.e., *what* to adapt) and (2) the best technologies and methods to accomplish adaptive goals (i.e., *how* to adapt). We conclude with a summary of key challenges and future applications of adaptive learning technologies. These challenges include: (1) obtaining useful and accurate learner information on which to base adaptive decisions, (2) maximizing benefits to the learner while minimizing costs associated with adaptive technologies, (3) addressing issues of learner control and privacy, and (4) figuring out the bandwidth problem, which has to do with the amount of relevant learner data that can be acquired at any time.

Rationale for Adapting Content

The attractiveness of adaptive technologies derives from the wide range of capabilities that these technologies afford. One capability involves the real-time delivery of assessments and instructional content that adapt to learners' needs and preferences. Other technology interventions include simulations of dynamic events, extra practice opportunities on emergent skills, and alternative multimedia options, particularly those that allow greater access to individuals with disabilities. We now provide evidence that supports the importance of adapting content to students to improve learning. These arguments concern individual and group differences among students.

Differences in Incoming Knowledge, Skills, and Abilities

The first reason for adapting content to the learner has to do with general individual differences in relation to incoming knowledge and skills among students. These differences are real, often large, and powerful; however, our educational system's traditional approach to teaching is not working well in relation to the diverse population of students in U.S. schools today (Shute, 2007). Many have argued that incoming knowledge is the *single* most important determinant of subsequent learning (Alexander & Judy, 1988; Glaser, 1984; Tobias, 1994). Thus, it makes sense to assess students' incoming knowledge and skills to provide a sound starting point for teaching. A second reason to adapt content to learners has to do with differences among learners in terms of relevant abilities and disabilities. This addresses issues of equity and accessibility. To illustrate, a student with visual disabilities will have great difficulty acquiring visually presented material, regardless of prior knowledge and skill in the subject area. Student abilities and disabilities can usually be readily identified and content adapted to accommodate the disability or leverage an ability to support learning (Shute et al., 2005).

Differences in Demographic and Sociocultural Variables

Another reason to adapt content to learners relates to demographic and sociocultural differences among students, which can affect learning outcomes and ultimately achievement (Conchas, 2006; Desimone, 1999; Fan & Chen, 2001). For example, training on a foreign language may contain different content depending on whether the learner is a child or an adult.

Differences in Affective Variables

In addition to cognitive, physical, and sociocultural differences, students' affective states fluctuate both within and across individuals. Some of these states – such as frustration, boredom, motivation, and confidence – may influence learning (Conati, 2002; Craig et al., 2004; D'Mello & Graesser, Chapter 6 in this volume; Ekman, 2003; Kapoor & Picard, 2002; Litman & Forbes-Riley, 2004; Picard, 1997; Qu et al., 2005).

In summary, there are a number of compelling reasons to adapt content to learners. We now provide context and coherence for adaptive technologies by way of a general evidence-based, four-process model. This model has been extended from (1) a simpler two-process model that lies at the heart of adaptive technology (diagnosis and prescription) and (2) a process model to support assessment (Mislevy et al., 2003).

Four-Process Adaptive Cycle

The success of any adaptive technology to promote learning requires accurate *diagnosis* of learner characteristics (e.g., knowledge, skill, motivation, persistence). The collection of learner information can then be used as the basis for the *prescription* of optimal content, such as hints, explanations, hypertext links, practice problems, encouragement, and metacognitive support. Our framework involves a *four-process cycle* connecting the learner to appropriate educational materials and resources (e.g., other learners, learning objects, applications, and pedagogical agents) through the use of a learner model (LM) (see Figure 1.1).[1] The components

[1] The terms "student model" and "learner model" are used interchangeably in this chapter. They are abbreviated as either SM or LM. Because this chapter focuses on the educational functions of adaptive systems, we limit our modeling discussion to the context of students or learners rather than more broadly defined users.

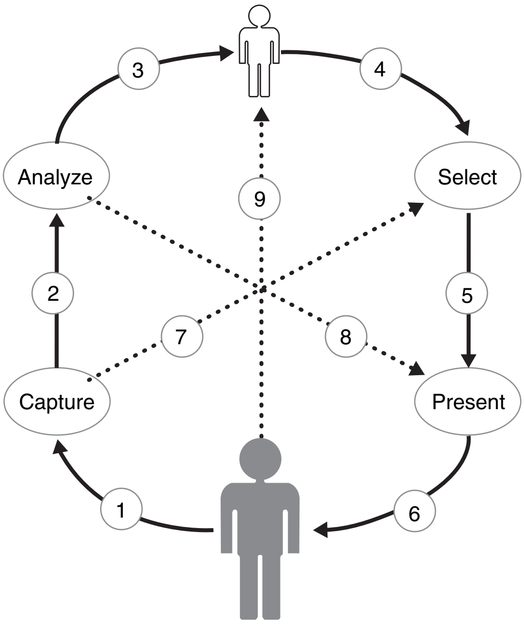

Figure 1.1. Four-process adaptive cycle. The larger human icon represents the student. The smaller human icon represents the student model.

Source: From "Adaptive technologies," by V. J. Shute and D. Zapata-Rivera, 2007, in J. M. Spector, D. Merrill, J. van Merriënboer, & M. Driscoll (Eds.), *Handbook of research on educational communications and technology* (3rd Ed.) (pp. 277–294). New York: Lawrence Erlbaum Associates, Taylor & Francis Group. Copyright © 2007 by the Taylor & Francis Group; reprinted by permission of the publisher.

of this four-process cycle include capture, analyze, select, and present.

CAPTURE

This process entails gathering information about the learner as the learner interacts with the environment (depicted in Figure 1.1 by the larger human figure). Relevant information can include cognitive data (e.g., solution to a given problem) as well as non-cognitive aspects of the learner (e.g., engagement). This information is used to update internal models maintained by the system.

ANALYZE

This process requires the creation and maintenance of a model of the learner in relation to the domain, typically representing information in terms of inferences on current states. That is, the computer can infer what the learner knows or can do directly from aspects of the learner's performance in the learning domain (e.g., if the learner solves a relatively difficult problem correctly, the inference is that his/her knowledge and/or skill related to the topic is likely pretty good, and if he/she solves another difficult problem correctly, the confidence in the inference that he/she knows the content well increases). In Figure 1.1, this is depicted as the smaller human figure and is often referred to as the student model or the LM.

SELECT

Information (i.e., content in the broadest sense) is selected for a particular learner according to: (1) his/her current status as represented in the student model and (2) the purpose(s) of the system (e.g., next learning object or test item). This process is often required to determine how and when to intervene.

PRESENT

Based on results from the select process, specific content is presented to the learner. This entails appropriate use of media, devices, and technologies to efficiently convey information to the learner.

This model accommodates alternative types and levels of adaptation. Table 1.1 describes some of the different possibilities, starting with a completely adaptive cycle and continuing to a nonadaptive presentation.

In general, the architecture of adaptive applications has evolved in a way that reflects the evolution of software systems architecture; for example, it is possible to find *stand-alone* adaptive applications where the complete adaptive system – including its student model – resides in a single machine. Also, adaptive applications have been implemented using a *distributed* architecture model. Some examples of distributed applications include: (1) client-server adaptive applications that make use of student modeling servers and shells (Fink & Kobsa, 2000); (2) distributed agent-based platforms (Azambuja et al., 2002; Vassileva et al., 2003); (3) hybrid approaches

Table 1.1. Scenarios Represented in the Four-Process Adaptive Cycle

Scenario	Description
A complete outer cycle, automated adaptation (1, 2, 3, 4, 5, and 6)	All processes of the cycle are exercised: capturing relevant information, analyzing it, updating the variables that are modeled in the learner model, selecting appropriate learning resources and strategies that meet the current needs of the learner, and making them available to the student in an appropriate manner. This cycle will continue until the goals of the instructional activity have been met.
Automated adaptation with user input (1, 2, 3, 4, 5, 6, and 9)	The learner is allowed to interact with the learner model. The nature of this interaction and the effects on the learner model can vary (e.g., overwriting the value of a particular variable). Allowing student input to the model may help reduce the complexity of the diagnostic and selection processes by decreasing the level of uncertainty inherent in the processes. It can also benefit the learner by increasing learner awareness and supporting self-reflection.
Diagnosis only (1, 2, and 3)	The learner is continuously monitored; information gathered is analyzed and used to update learner profiles, but not to adapt content. This may be seen as analogous to student assessment.
Short (or temporary) memory cycle (1, 7, 5, and 6)	The selection of content and educational resources is done by using the most recent information gathered from the learner (e.g., current test results and navigation commands). Adaptation is performed using information gathered from the latest interaction between learner and the system.
Short (or temporary) memory, no selection cycle (1, 2, 8, and 6)	A predefined path on the curriculum structure is followed. No learner model is maintained. This predefined path dictates which educational resources and testing materials are presented to the learner.

involving distributed agents and a student modeling server (Brusilovsky et al., 2005; Zapata-Rivera & Greer, 2004); (4) peer-to-peer architectures (Bretzke & Vassileva, 2003); and (5) service-oriented architectures (Fröschl, 2005; González et al., 2005; Kabassi & Virvou, 2003; Trella et al., 2005; Winter et al., 2005).

To illustrate how our four-process adaptive model can accommodate more distributed scenarios, Figure 1.2 depicts an extended version of our model. Agents (e.g., application, personal, and pedagogical agents) maintain a personal view of the learner using their own representation of the "four-process adaptive cycle" (see Figure 1.1). Agents share (or negotiate) personal information with other agents to accomplish goals on behalf of the learner. A common LM is maintained in a learner modeling server. The term "common learner model" refers to a subset of the LM that is common to all the agents (e.g., identification information) and other information the agents share (e.g., long-term goals and interests).

Summary of Current Adaptive Technologies

This section describes adaptive technologies currently in use and relevant to the context of this chapter. The technologies have been divided into two main sections: soft and hard technologies; this distinction may be likened to *program* versus *device* and may be used across the array of processes described in the previous section (i.e., capturing student information, analyzing it, selecting content, and presenting it). The technologies selected for inclusion in this section are those that make use of, to some extent, an LM in its formulation. Also, this listing is intended to be illustrative and not