2007

Building Artificially Intelligent Learning Games

Richard Van Eck
University of North Dakota, richard.vaneck@und.edu

Chapter XIV

Building Artificially Intelligent Learning Games

Richard Van Eck, University of North Dakota, USA

"The biggest thing limiting games in education in my view is the lack of good artificial intelligence to generate good and believable conversations and interactions…We need games with expert systems built into characters and the interactions players can engage in with the environment. We need our best artificial tutoring systems built inside games, as well…Then we will get games where the line between education and entertainment is truly erased."—James Gee, 2003

Abstract

The idea of digital game-based learning (DGBL) is gaining acceptance among researchers, game designers, educators, parents, and students alike. Building new educational games that meet educational goals without sacrificing what makes games engaging remains largely unrealized, however. If we are to build the next generation of learning games, we must

recognize that while digital games might be new, the theory and technologies we need to create DGBL has been evolving in multiple disciplines for the last 30 years. This chapter will describe an approach, based on theories and technologies in education, instructional design, artificial intelligence, and cognitive psychology, that will help us build intelligent learning games (ILGs).

Introduction

The learning potential of games has been discussed in the popular press and academic journals since at least the mid-60s with the advent of simulation games in the social sciences. Yet games and learning have also always been viewed by many with a healthy dose of skepticism. One of the reasons for this has always been the dichotomization of playversus work, in which play is seen as frivolous entertainment and therefore the opposite of work and learning. This popular belief has begun to change, however, in part thanks to the efforts of scholars and researchers who have studied games and learning and published in the mainstream press (e.g., Gee, 2004; Johnson, 2005; Prensky, 2000; Reiber, 1996). Some 200 academics interested in developing and using games for learning have attended at the Game Developers Conference each year since 2002, and hundreds of academics are conducting game studies, designing games, and/or finding ways to integrate commercial games into the classroom (Foreman, 2004). This has been in part spurred by the tremendous growth in the games industry, which is currently estimated to be a $10 billion industry (eSchool News, 2005). This, of course, is in turn driven by the growing number of people who are playing games, and they 'are not all net gen-ers. The Entertainment Software Association (ESA) reports that 75% of heads of household play computer games, and that 62% of game players are over 18 with a mean age of 30. This increase in the game industry and number of games has, most recently, led to an increase in the number of colleges offering game design programs, which will further break down barriers to the acceptance of games and learning.

But even as games become more mainstream and the idea of games as a learning medium gains acceptance, the promise of learning games remains largely unrealized. Although the edutainment industry (initial attempts at learning games) has grown in sales over the years, it has not revolutionized learning nor experienced the explosive growth originally predicted. The combination of the adaptive and tireless nature of computer-based instruction with both entertainment and authentic problem solving should have produced a host of games that teach all learners at their own pace. So where are these games?

One reason for the dearth of these games may be that the dominant paradigms in education and the gaming industry are too different to allow for good synergies. The world of education is focused on providing the best path for learners to get from novice to expert in different domains. Content is thus privileged over experience. The game world, in contrast, is focused on providing a rewarding, interactive experience. Content is secondary to experience and is willingly sacrificed for game play when and where needed. In the cases of edutainment titles, these worlds often clash, with educators developing content (often linear, hierarchical, and instructivist) without regard to experience, and game developers building interactive environments (often non-linear and player-driven) without regard to the content or instruction. It is this culture clash that has led to titles in which game play is interrupted by long bouts of reading and drill and practice, and/or where game play is used as a reward for slogging through such instruction. In these edutainment titles, the game and the content are rarely if ever integrated. Seymour Papert (1998) refers to these as Shavian Reversals, which is a term from genetics indicating an offspring that has inherited the worst characteristics of both parents. As expected, these titles have rarely been financially successful, making game companies leery of anything that smacks of education. Game developers often believe that "whenever you add an instructional designer, they suck the fun out" of the game (Prensky, 2004).

While there has been some progress made through initiatives like the Serious Games initiative, the games-to-teach project at MIT, and the Education Arcade, which focus on games that teach content in the context in which it is demonstrated (e.g., Carnegie Mellon's *HazMat* project, Chris Dede's *River City* project, Education Arcade's *Revolution* history game, and Muzzy Lane's commercial game *The Calm and The Storm*), blending instructional content and games remains a significant challenge for the field.

Part of the reason for this is that the field is too young to have many established research methods and theoretical models for game design, let alone instructional games (e.g., Pearce, 2004; Prensky, 2001; Smith & Mann, 2002). What we need is to establish new models for developing learning games that account for the strengths of both the educational and game worlds. To do this, we must recognize that while games may be a new phenomenon, the tools and theory we need to forge these new models exists already in multiple fields and domains; we have just not yet examined them to see what each can contribute to the new field of digital game-based learning (DGBL). The reason for the failures of many edutainment titles was that the model of instruction, direct instruction, was not compatible with the game environments. But there are theories and instructional strategies in education and other fields that are compatible with (and indeed, used by) games. We need to examine games for their underlying strengths and weaknesses, and look to the fields of education, psychology, and to theories of narrative and storytelling to find compatible methods of instruction and learning.

For example, instructional designers have recognized for years that different types of learning require different instructional strategies and approaches. Gagné's seminal book, the *Conditions of Learning*, first published in 1965, distinguishes between five types or varieties of learning: motor skills, attitudes, cognitive strategies, verbal information, and intellectual skills. He further breaks intellectual skills into five sub-types: problem solving, rules, defined concepts, concrete concepts, and discriminations (presented in order of complexity from most to least). All of these varieties of learning require different types of instructional events and strategies. While this may seem to be common sense, teachers who fail to make this distinction have a tendency to treat all instruction the same way, and to use the same activities and strategies for all types of learning.

We face a similar situation now with games; there is a tendency to discuss all games as if they were the same. In fact the several different game genres (see Table 1), each with its own strategies and approaches, require different approaches as well. It follows that we must understand how these different game taxonomies and their attendant strategies align with learning taxonomies and strategies so that we can begin to match learning and games without sacrificing playability or learning. Table 1 provides an example taxonomy of game types (based on Bates, 2002) and learning taxonomies presented together to facilitate exploring

Table 1. Matrix of Game and Learning Taxonomies (Note: table co-authored with Joanne Gikas, 2003)

Taxonomy of Games	*Explanation of Genre*	*Gagne's Intellectual Skills*	*Bloom's Taxonomy*
Action	*Keeps the player moving and involved at all times. Primary skills are eye/hand coordination and quick reflexes. Deep thinking is generally not required.*	*Defined Concepts & Below*	*Application & Below*
Role Playing	*Revolves around characters, story and combat and takes place in large, expansive worlds. Usually collaborative, often online.*	*Problem Solving & Below*	*Evaluation & Below*
Adventure	*Story based on exploration and puzzle solving where the player is the protagonist. Player must determine best path through storyline and obstacles on own or with others.*	*Problem Solving & Below*	*Evaluation & Below*
Strategy	*Emphasize strategy and theory, often in recreations of historical or other human events.*	*Problem Solving & Below*	*Evaluation & Below*
Simulations	*Simulation of processes, events, or phenomenon. Emphasis on realistic representation.*	*Problem Solving & Below*	*Evaluation & Below*
Sports	*Allows players to play simulated sports activity.*	*Defined Concepts & Below*	*Application & Below*
Fighting Games	*Players engage in combat individually or in teams. Story is present but ancillary to fighting skills.*	*Defined Concepts & Below*	*Application & Below*
Casual	*Games for the "new gamers" easy to learn, not difficult to master.*	*Defined Concepts & Below*	*Application & Below*

where synergy exists at the intersection of both fields. These kinds of tools are only possible if we draw upon multiple fields such as instructional design, education, and games.

This chapter will focus on an approach for developing DGBL that maximizes the potentials for both learning and entertainment by drawing upon established fields including psychology (artificial intelligence, pedagogical agents, and intelligent tutoring systems), English (narrative and storytelling), and education (instructional design). Blending such disparate fields and approaches with games is not simply a matter of combination. Any instructional content we hope to integrate with games must be compatible with the underlying pedagogy and assumptions of games, or the learning will remain a separate construct and ruin the game experience. So we must also examine games for the pedagogical principles they employ in the learning that naturally takes place.

In this chapter, I will begin by outlining four principles of learning inherent in digital games:

- Play theory, cycles of learning, and engagement;
- Problem-based learning; situated cognition and learning;
- Question-asking, cognitive disequilibrium, and
- Scaffolding.

I will then discuss how pedagogical agents and intelligent tutoring systems, modified by the four game principles and by narrative, can be combined for use in designing intelligent learning games (ILGs).

Assumptions

What Is The Purpose of This Chapter?

This chapter is not a prescriptive method for developing DGBL, nor is it intended to provide specific guidance for integrating games into the curriculum. It is a preliminary outline of the ideas and approaches from which a model could be developed for developing intelligent learning games. It is written assuming little background knowledge and in a manner that is accessible to anyone. Those who are involved (or who wish to become so) with the design and development of serious games will hopefully find it most relevant.

How Is This Chapter Organized?

The order in which I address the different theories and research areas that underlie the proposed model for intelligent learning games should not be taken as an indication of the relative importance of each area. Rather, topics are presented in the order in which they address pertinent questions in the design of serious games. This progression reflects the natural progression of questions and answers (at least mine, and hopefully the reader's) one goes through when considering how best to design truly engaging and effective learning games. In this sense, topics are addressed in terms of their importance in answering the most salient questions during each stage of this process. As with most such endeavors, the answers to each question inevitably lead to other questions, which must also be addressed.

What Do I Mean By Games?

I postulated earlier that not all games are alike, and castigated those use "game" as a universal term. I can hardly proceed with this chapter, then, without clarifying what I mean by "games." First, I am speaking only of computer or console games. Within the world of computer and console games, this chapter builds a model for designing DGBL by focusing on adventure games with multiple characters. Adventure games have, in my opinion, the greatest potential for addressing all levels of the learning taxonomy.

Adventure games are situated in environments that are generally immersive, allow (and even require) exploration, are driven by narrative and story, and often require hypothesis formulation, testing, revision, and re-testing. These kinds of strategies are conducive both to adventure game play and to problem solving. Adventure games thus address all levels of the learning taxonomy, and in particular focus on the highest levels. Problem solving skills (also sometimes referred to by educators and parents as critical thinking skills) are among the most highly desired goals in education, but they are typically among the most difficult to address in any instructional medium. Adventure games are well aligned with existing pedagogical theories such as situated cognition and learning (e.g., Brown, Collins, & Duguid, 1989; Lave & Wenger, 1991), anchored instruction (e.g., Bransford et al., 1990; CGTV, 1990, 1993, 1996) and discovery-based learning (e.g., Bruner, 1960), all of which have been shown to promote problem-solving skills.

The use of multiple characters is key as well, whether those characters are driven by humans or artificial intelligence, as they allow for the social nature of learning and working with others that more closely reflects how we demonstrate knowledge in the "real world." In contrast to the way learning appears in many classrooms, in the real world we rarely work independently on problems. Rather, we work with others either formally in teams or informally as colleagues because knowledge and skills are distributed rather than concentrated in one person or position.

While today's learners, as digital natives (Prensky, 2001), may learn differently from previous generations and thus expect and prefer different strategies and approaches, this by itself is not enough to justify the expense and difficulty of developing high-quality DGBL. The ratio of learning to effort/expense must be favorable to justify the use of games in learning. Adventure games, because they embody problem-based learning, have the potential to not only promote problem solving and critical thinking but to do so while teaching content at all other levels, which given the difficulty of addressing these goals, makes their use justified. However, precisely because of their complexity, they are resource intensive to develop, and thus require more planning and effort. This makes the need for theoretical frameworks and development models most critical for this type of DGBL.

Finally, adventure games require discourse and narrative as part of the game world. Adventure computer games are narrative-based problem-solving activities in which the storyline drives the actions of the player and the movement through the game through a continuous cycle of hypothesis formation, action, and feedback. Narrative and storytelling, I will argue, is the oldest form of instruction, and therefore one of the most powerful instructional strategies available to us. Narrative theory also provides guidance for how we can blend the "narratives" of games and learning technologies like pedagogical agents (PAs) and intelligent tutoring systems (ITSs). This represents the kind of intersection and alignment of powerful learning and game play strategies that we must explore if we are to create great DGBL instead of Shavian Reversals.

Pedagogical Principles in Games

Principle 1: Games Employ Play Theory, Cycles of Learning, and Engagement

Games are effective not because of what they are, but because of what they embody. One could argue that play is the dominant feature of games (e.g., Pearce, 2004). I will argue later that games, when viewed as narrative texts, also illuminate the concept of play and learning. Researchers like Crawford (1982), Gee (2004), Lepper and Chabay (1985), Papert (1998), and Rieber (1996) point out that play is a primary socialization and learning mechanism common to all human cultures and many animal species. Play theory says that play is the most effective instructional technique. This conclusion is based largely on the observation that we learn more in the first years of life than we do in any other corresponding time in our lives (Lepper & Chabay, 1985). The play of young animals as they are growing up is the means by which the most important life skills are naturally learned. "Games are thus the most ancient and time-honored vehicle for education" (Crawford, 1982, Chapter 2). Only mammals and birds engage in play, indicating that the role of play in fostering higher learning is critical (Crawford, 1982). Rieber (1996) says research in "anthropology, psychology, and education indicates that play is an important mediator for learning and socialization throughout life" (p. 44) and that, "Having children play games to learn is simply asking them to do what comes naturally… However, playing a game successfully can require extensive critical thinking and problem-solving skills" (p. 52).

The problem, according to play theorists, is that at some point in our development, play is replaced by work, which may account for poor motivation in schools today. "Work is respectable, play is not" (Rieber, 1996, p. 43), and so our school and work lives are dominated by work instead of play. Proponents of play theory, in contrast, say play and work can be synonymous when work is its own reward (Rieber, 1996).

According to Brian Sutton-Smith (1997), there are seven kinds of play:

1. Play as progress: purpose of play is to learn something useful.
2. Play as fate: gambling and games of chance.
3. Play as power: winners and losers.
4. Play as identity: when play serves to confirm the identity and power of those playing.
5. Play as the imaginary: improvisation, imagination.
6. Play as self: solitary play activities like solitaire.
7. Play as frivolous: the intrinsic worth of the experience is of primary concern (pp. 9-11).

So we can see that play is just as complex a concept as are games or learning. It seems logical to expect, then, that different kinds of play will support different kinds of learning and be appropriate for different learners. Computer adventure games seem to most closely reflect type five, although there are certainly elements of types two through four as well. Type one might also be said to come into play with DGBL since games provide constant feedback regarding progress, and because players are focused on making progress within the game. Key to this aspect of play, however, is that what is being learned in the game is that which is required to "solve" the game; one does not *usually* pursue a game to learn something. The question of how different types of play interact with learning outcomes and individual differences is an important one for DGBL designers to answer.

One of the key aspects of play that makes it so effective is interaction. It is not possible to be passive during play; play always requires some form of input or response on the part of each person. Play in its most free-form sense (e.g., kids in a backyard) appears to be unconstrained, but in fact is guided by rules and goals. These rules may not be fixed ahead of time, and indeed may change frequently during play, but they demand and constrain actions on the part of each player; anyone who does not "play by the rules" will suffer consequences (in the game, socially, or both). Likewise, there is a constant cycle of action and reaction that occurs in play, although the pace and frequency may vary. We take turns in board games, at bat, or on offense and defense. We roll dice, spin spinners, act, and respond to others who act. Thus play requires active participation by all involved, both physical (during your turn) and mental (during other players' turns). Players are always either acting or preparing to act. Likewise, every act results in feedback, usually contiguous to the action. This constant cycle of action, feedback, and reaction according to the constraints of the rules is in large part what underlies the effectiveness of the learning process (play) in digital games.

Not surprisingly, the active interaction cycle also helps account for the high engagement in digital games. We know that learners who are engaged (e.g., those who formulate and ask questions, make predictions, practice and demonstrate what they are learning, incorporate feedback, monitor their learning) tend to learn the most. The characteristics, or events, that occur during this kind of engaged learning are key to designing instruction that is effective and engaging. Gagné (2005), after examining the literature on cognitive psychology, learning, and education in 1968, derived nine principles, or events, of instruction that unify the external events of instruction with the internal events of information processing:

1. Gaining attention
2. Informing the learner of the objective
3. Stimulating recall of prerequisite learned capabilities
4. Presenting the stimulus material
5. Providing learning guidance
6. Eliciting performance
7. Providing feedback about performance correctness
8. Assessing the performance
9. Enhancing retention and transfer

These events remain one of the most significant contributions to instructional design today, and are widely used to ensure effective instruction. Each one of these events is designed to promote the learner's active engagement and metacognition. In particular, events 4 through 7 mirror the cyclic process of active interaction described in play, and in fact make up the bulk of activities in effective instruction.

Many who look at Gagné's nine events for the first time immediately dismiss them (and all instructional design) as inappropriate for the design of learning games. The perception is that the nine events might run like this:

1. Hey! Listen up, students.
2. You are going to learn the following information during this lesson (insert objectives here).
3. Remember what we studied last week?
4. (Insert long-winded content here).
5. You can remember this best if you make up a story about it for yourself.
6. What did I just say about (content from 4 here)?
7. That's (right or wrong), and here's why. You should study the things you missed. Read Chapter 15.
8. Take this test.
9. Any time you come across (content here), you should remember that it shares these common principles (insert principles here), and so you will always be able to apply (content here) when you learn about other (content here).

Obviously, this would be stultifying dull, either as a game or formal instruction. But this is a misinterpretation that needs to be corrected. These events are not a prescriptive list of actions to be employed one at a time, one right after the other, or just once during instruction. In other words, these are not strategies; they are events. Any good designer of instruction attends to these events, but does so through a variety of strategies. Each of Gagne's events can be achieved in much more subtle, recursive, and engaging ways. In fact, *every single one of these events* occurs naturally in games. Table 2 presents the nine events again along with typical examples for each event in games.

So the first principle of learning in games is that all learning must adhere to the tenets of play, and be subservient to the game play.

Principle 2: Games Employ Problem-Based Learning

Problem-based learning is "an instructional approach that organizes the curriculum around loosely structured problems that students attempt to solve by using knowledge and skills from several disciplines or subject areas" (Collins & O'Brien, 2003, pp. 282-283). Problem-based learning is effective in promoting greater comprehension because it serves as a vehicle for discovery of concepts and rules, as well as a means of learning how to think about and apply that knowledge to solve the kinds of problems learners are likely to encounter later (Delisle, 1997). Problem solving is also the highest level of learning we can strive for. Gagné, Wager, Golas, and Keller (2005) discuss five different varieties of learning: verbal information (facts, labels and name, propositions), attitudes (a person's affective stance toward something, measured by their choices), motor skills (coordinated physical movements), cognitive strategies (strategies for enhancing and monitoring one's own learning process), and intellectual skills (those skills used for solving problems and for learning rules and concepts).

Each of the varieties of learning is distinct, and requires its own kinds of instructional events, strategies, and approaches. For example, learning to hit a baseball is both a motor skill (the mechanics of the swing) and intellectual skill (knowing the rules for positioning the bat, when to start the swing, using different swings for different pitches, and knowing what to do at different ball and strike counts). The motor skill portions can be best taught through guided practice and modeling. The intellectual skills parts can be taught out of context, without modeling, using traditional forms of direct instruction, guided discovery, and so

Table 2. Gagné's Nine Events of Instruction and Examples from Games

Instructional Event	*Examples from Games*
Gain attention	*Motion, cut scenes, noise, music, character speech, health meters, attacks, death*
Inform of objective	*Documentation for the game, introductory movies, cut scenes, character speech, obstacles that limit movement or interaction*
Stimulating recall	*Environmental cues (e.g., in Laura Croft: Tomb Raider, ledges that look like those trained on in the earlier tutorial), obstacles (search for solutions involves recalling solutions and events from earlier in the game)*
Present stimulus	*All of the above (characters, environment, objects, puzzles and obstacles, conversation) arranged according to goals of game*
Provide guidance	*Cut scenes, non-player character (NPC) or player character (PC) speech, hint books, cheats and walkthroughs, friends, partial solutions to puzzles (pressing on the wall makes it rumble, but it does not open). Also, much comes from the learner themselves as they process what has occurred in the game, but the arrangement of the actors and objects in the environment and the structure of the story itself also provide implicit guidance*
Elicit performance	*Players cannot progress through the game without demonstrating what they know or think they know—all knowledge is demonstrated within the confines of the game narrative and structure.*
Provide feedback	*Character speech, sounds, motion, etc., Player gets past the obstacle or achieves the goal, or does not. Every action has immediate feedback, even if that feedback is that nothing happens.*
Assess performance	*Movement through the game IS assessment. Nothing is learned that is not also demonstrated.*
Enhance retention	*Things learned early in games are brought back in different, often more complex forms later. Players know that what they learn will be relevant in the short and long term.*

forth. One cannot hit a baseball without both aspects, however, and one cannot learn either aspect effectively without using the appropriate instructional strategies.

In addition to separating learning into these five varieties, Gagné et al. (2005) describe how intellectual skills themselves are comprised of different kinds of skills (see Table 3), at the top of which lies problem solving. Each level of intellectual skill requires as pre-requisites, the skill below it. Thus to solve a problem, one must be able to combine rules learned previously to form more complex rules that can solve a novel problem. For instance, a customs official in India must determine whether to let people into the country according to many different criteria, including citizenship, passports, visas, lists from Interpol, and so forth. Each person may meet different criteria to varying degrees, and the agent must decide based on dozens of rules whether that person meets the criteria for admission to the country. In order to solve novel customs entry problems using all of the different rules, the agent must have mastered the defined and concrete concepts that comprise the building blocks of those rules. For instance, it is not possible to know the rule that, "U.S. citizens must have a passport in order to enter other countries" (one of the many rules the agent must use to solve the problem of entry) without also knowing that "a citizen is anyone who was born or nationalized in the country in question" (defined concept) and what a U.S. passport looks like (concrete concept). Finally, to be able to know what a U.S. passport looks like, the agent must also be able to tell the difference between the many things that determine what a U.S. passport is, including at its most basic level, the ability to tell whether blue differs from other colors (a discrimination). This illustrates how problem-based learning of necessity requires the full range of intellectual skills in learning.

If you refer back to Table 1, you will see that many games reach the problem-solving level. One could argue that every game in fact is a form of problem solving, and varies from other games primarily in the complexity and type of problem, and the requisite solution strategies required to win.

According to Jonassen (2002), there are two critical attributes of a problem. First, all problems have at their heart some goal, and the fact that we don't know how to get to the goal without generating new knowledge makes it a problem (Jonassen refers to this first

Table 3: Gagné's Taxonomy of Intellectual Skills

Problem Solving
Requires the generation of new rules based on a combination of several prerequisite
Rules
Which in turn require the mastery of different
Defined Concepts
and
Concrete Concepts
Which require the ability to make
Discriminations

attribute as the unknown). Second, there is some kind of value to the problem-solver that is inherent in finding a solution to the problem. Likewise, problem solving is also comprised of two stages, in which we formulate a representation of the problem for ourselves (called the problem space), and then we work within that representation to change and modify it in an attempt to find the solution (the unknown).

One of the keys to the effectiveness of problem solving as an instructional strategy is the process of goal setting. To solve a problem, one is initially given a goal, which is then internalized. In order to achieve that goal, the learner must formulate sub-goals related to identifying what is known and unknown, strategies for acquiring what is unknown, a process for testing and revising hypotheses, and so forth. Immediate and short-term goals such as these promote more *effective* learning and cognitive development in general, as well as self-efficacy *about* learning (Bandura, 1997).

One of the most popular strategies for solving problems is to make "moves" that appear to reduce the gap between where we are now (the current problem space) and where we want to be (the goal state). The strategy works in some problems, but is counter-productive at times in others (e.g., the tower of Hanoi, http://www.mazeworks.com/hanoi/, and missionaries and cannibals http://www.plastelina.net/games/game2.html), where moves that appear to increase the gap between the problem state and the goal state must be made first in order to solve the problem. It is for this and other reasons that problem solving is considered domain specific and cannot be taught as abstract rules or principles, instead requiring multiple practice opportunities in multiple domains (Larkin, 1989). Games can be a great vehicle for this repeated, multidisciplinary problem-solving practice.

Games always have a goal, which some might argue makes them problem solving by default. From the marketing and promotion of games to the documentation that accompanies them, the goal of the game is made the focus. Consider the box panel description of a recent Game of the Year Award Winner, *The Longest Journey*:

Between science and magic, between order and chaos, between Stark and Arcadia, there is an ancient balance. For thousands of years, this balance has weighted the scales of the cosmos evenly, ensuring harmony between the twin worlds. But now, in an age of great turmoil, chaos threatens to turn the scales and bring our most terrifying dreams to life. The Guardian of the balance has abandoned his throne... the armies of the Vanguard are advancing... a storm is coming... and the fate of the worlds is in the hands of one person: April, a Shifter. April's future is shrouded in a veil of mystery, and the journey ahead is treacherous and winding. A journey not only through twin worlds, but into her own heart and soul.

There is clearly an unknown here, which is how the player character April will achieve this balance between science and magic, order and chaos, Stark and Arcadia. Just as clearly, we have no way at this time of knowing really what any of this means, nor of any strategies we might already possess to achieve the goal. Despite this, many people (10, 873 in the first month of its release, and more than 50,000 the next year) felt that this was a problem whose solution held value for them.

Like problem solving in other venues, playing a game requires us to formulate a problem space for both the overall goal of the game (e.g., to help April maintain this balance) and the subordinate problems along the way (often numbering in the hundreds for adventure games). Two key points are clear. First, just about everything one does in an adventure game is problem solving—there is very little "down" time. Second, we rarely have any of the prerequisite knowledge needed to solve the problem. This is the strength of both games and of problem-based learning. The problem (and a game is a complex problem made up of multiple problems) itself guides the learning, and serves as the impetus and vehicle for learning all of the subordinate intellectual skills (rules, concepts, and discriminations).

Consider Gagné et al.'s (2005) description of the conditions of learning needed for problem solving. Learner "performance requires the invention and use of a complex rule to achieve the solution of a problem novel to the individual. When the higher order rule has been generated, it should also be possible for the learner to demonstrate its use in other physically different but formally similar situations" (p. 73). Game players will immediately recognize this as part of what one does during a game.

For example, consider the following scenario from the game *Mysterious Island*. I am stranded on an island with only a few items, among them, a satellite phone with a built-in encyclopedia. Unfortunately, it is out of power and I do not have access to electricity. I know that my phone needs power to work and that an outlet is needed to charge it (prerequisite rules). I also know that I have an inventory with some items in it and, if I have played any game before, know that things I have found will be useful in some way during the game (cognitive strategy and a rule). Later, in a laboratory in a cave, I find instructions on how to build a battery out of common objects, some of which I have already located on the island. Rather than recharging the battery I have, I have found a way to replace that battery. I combine the required objects in my inventory by dragging and dropping them onto each other to build a battery that provides minimal power (enough to activate the encyclopedia on my phone, which will help me solve other problems in the game). I have combined several rules in the game, some of which I knew (phones and batteries) and some of which I had to learn (how to make a battery, alternate ways to power my phone, and how to combine things in inventory). These rules have helped me formulate a new complex rule: information can be found (on the island or in my phone) that can help guide me as I combine useless things into things that will help me solve problems. This new rule will help me later in the game (many times). For instance, I have to heal a monkey I find by researching common antibiotics, combining items I find on the island to make them, and administering them to the monkey who gets better (a reward of its own) and later helps me get past other puzzles.

These examples are typical of the literally hundreds of instances of problem solving in many complex adventure games. So according to the performance requirement outlined by Gagné et al., performance in the game is indicative of problem solving.

Gagne et al (2005) then go on to list six external conditions, or characteristics of the instruction, that are required of the students to support problem-solving in traditional instruction: 1) be confronted with novel (new to the learner) problems for which they have the requisite rules; 2) apply problem solving strategies (there is no direct instruction involved); 3) receive feedback; 4) be encouraged to reflect on the solution; 5) be provided with practice on similar problems to encourage transfer; and 6) engage in problem-solving that is facilitated by collaborative group work. If games were instructional applications of problem solving, we would then expect these attributes to be present in some form or fashion in the game as well.

From the *Mysterious Island* scenario described earlier, one can see that conditions one, two, three, and five are all present. Condition one because I was presented with a novel problem, although I do not have the requisite rules (yet) for solving it. Condition two is satisfied because the game does not provide direct instruction for solving the problem (although it does provide direct instruction for a rule, building a battery, that will be requisite for solving the problem). Condition two lists several other requirements as part of the condition. Among these are that instructors should make sure that students should monitor their own progress, identify ineffective strategies and irrelevant rules, select appropriate rules, and should provide just enough assistance and guidance for the learner to accomplish things they could not do autonomously. All of these things occur during game play, although whether the learner is responsible themselves or the game provides them varies.

Condition three is satisfied because my phone works or does not, my battery parts assemble or do not, and so forth. Condition four is satisfied because knowing that what I learn and demonstrate on this problem (and what I have learned already) may be useful for the solution of other problems, I am encouraged to reflect on my actions. The game requires reflection on the part of the player, and if they are unable or unwilling to do it, they will not learn as much, nor succeed as quickly than if they do reflect and learn during play. Condition six—collaborative problem solving—is not typically present in games, although the massively multiplayer online role playing games (MMORPGs) embody the concept, as we'll discover later in the discussion of pedagogical agents and intelligent tutoring systems.

In addition to the instructional benefits of problem-based learning, DGBL that embodies problem solving is cost effective. Building truly intelligent learning games will require significant human and financial resources, both time and money, which is most justifiable if the results address the highest levels of learning, often the hardest and thus least frequently addressed in public education. Further, problem solving, because it requires all the lower level intellectual skills as well, is necessarily complex and offers the greatest potential for integrating large amounts of content. In other words, it has the best educational cost to benefit ratio because we can use it to address large parts of the curriculum.

So our second pedagogical principle of learning in games is that learning is problem-based, and involves solving problems that are used to structure additional learning of prerequisite skills.

Principle 3: Games Embody Situated Cognition

Another reason that games are effective is because the learning takes place within a meaningful context; what you must learn is directly related to the environment in which you learn it, and is thus not only relevant but also applied and practiced within that context. Most research has shown that congruence of learning and performance contexts promotes performance (Anderson, 1995), and changes in context decrease what is recalled from prior sessions (Bower, 1981, 1987; Clark, Milberg, & Ross, 1983; Smith, Glenberg, & Bjork, 1978).

Context effects can be partially explained by a theory called "encoding specificity." According to Begg and White (1985) congruence between the encoding and recall contexts will improve recall. These context effects have been demonstrated using a variety of modalities, including textual, verbal, visual, and emotional. Learning in games is "encoded" within the context of the problem(s) being solved in the game.

Just as real-world contexts for mood can impact performance, real world and/or meaningful contexts for learning can improve performance as well. Gildea, Miller, and Wrutenberg (1990) found that embedding newly defined words in a story line and supplementing them with visual context information enhanced learning the words. The story context produced good performance on vocabulary learning, although not as good as the story context with pictures, which in turn was surpassed by those who received the story context, pictures, and solicited help in the form of illustrative sentences. These and similar results have led researchers like Griffin (1995) to call for the embedding of material in a complex environment (like games) to make it more meaningful, which will in turn lead to improved learning and performance.

These learning context principles are at the heart of what researchers like Brown et al. (1989), and later Bransford et al. (1990) and the Cognition and Technology Group at Vanderbilt (1992a, 1992b, 1992c) identify as situated cognition. Situated learning, the practical application of situated cognition to formal learning, arises out of a movement in cognitive studies in the 1970s that began to study human cognition in the contexts in which they naturally occur (Cohen & Siegel, 1991; Graesser & Magliano, 1991). Situated learning holds that learning is effective to the degree that it is embedded in a meaningful context (e.g., Choi, 1995; Choi & Hannafin, 1995).

Games make perhaps the best use of both context and situated cognition of all other instructional media outside of actual apprenticeship. They do this within the context of Gagne's nine-events of instruction as well, employing each of them regularly, recursively, and contiguously to the action. The goal of the game (the unknown) drives every aspect of the game, from the nature of each puzzle or obstacle to the learner actions and the constraints thereof. Almost no learning takes place out of the context of the game; no learning is unrelated to what is currently going on in the game; and no learning advances you through the game unless it is demonstrated. Even when game players seek assistance from friends and hint books, the learning is contextualized to the game world.

Because the context of the game world and the learning that takes place are perfectly aligned, it follows that any instruction we attempt to embed in a game must do the same, which thus becomes our third principle: games embody situated cognition.

Principle 4: Games Encourage Question Asking Through Cognitive Disequilibrium and Scaffolding

Researchers have claimed that tutoring can increase learning performance by as much as two standard deviations over traditional instruction, a phenomenon often referred to as Bloom's 2-sigma effect (Bloom, 1984). The reason for this, many say, is that the instruction is individualized (content, strategies, media, and pace of learning are customized according to the student's knowledge, skills, abilities, and preferences). Individualizing instruction is much more practical when teaching under a one-to-one (as with tutoring) than it is under a 25-to-1 (as with traditional classroom-based instruction) student-teacher ratio.

Another reason tutoring is so effective is that it promotes question asking (Graesser & Person, 1994), which is critical to the learning process. Unfortunately, in most classrooms, question asking is rarely done (Otero & Graesser, 2001). The typical student asks only six to eight questions per hour (Graesser, Wiemer-Hastings, K., Wiemer-Hastings, P., Kreuz, & Tutoring Research Group, 1999), for example, and most of those are shallow (e.g., Graesser & Person, 1994). According to Otero and Graesser (2001), the research not only shows that question asking is key to comprehension, problem solving, reasoning, and other cognitive activities, it also shows that students who are trained to ask good questions become better learners. In fact, one finds that question asking in one form or another is a part of most effective learning strategies. For instance, problem-based learning requires that students formulate questions as part of the process. The problem is presented, and students are expected to formulate their own questions to guide their acquisition of the knowledge needed to solve that problem (Delisle, 1997).

Questions are also key to the process of self-regulation and metacognition (mental self-reflection, awareness, and regulation) in learning. The best learners are those who are constantly making predictions and asking themselves questions about what they are learning (e.g., Do I understand that? How does that relate to what I already know? What does this mean? How can I best understand and remember this?). Asking such questions activates prior knowledge structures, or schemas, and promotes better learning. Because schemas are in effect networks of propositions and declarative knowledge, asking questions promotes more complex and refined schemas because questions help emphasize, refine, and build the relationships between and among propositions. Research has shown that one of the differences between experts and novices is the depth and complexity of the schemas and mental models they hold in a given subject (e.g., Bransford, Brown, & Cocking, 1999).

Related to the idea of schemas and mental models are Piaget's theories of assimilation and accommodation. Piaget believed that when confronted with new world knowledge, the learner has two options. The learner can integrate that knowledge into what is already known and believed (a schema for how the world works in relation to the new knowledge). Piaget referred to this as assimilation, and it requires the least effort of the two approaches. One can think of this as walking down a shopping aisle at the grocery store and filling your cart with bread, milk, eggs, and canned goods (new world knowledge). Each item is easily placed in the cart (your mental model, schema, or propositional network). With assimilation, we attempt to fit new, often complex information into existing slots or categories by simplifying or re-conceptualizing it to match what we already know. Imagine, however, that you then find an elephant. The cart cannot possibly hold the elephant, and the cart will have to change drastically for you to be able to accommodate the elephant (the new knowledge). Accommodation requires much more effort than assimilation, including replacing or reconstructing existing ideas. With accommodation, we must modify our existing model of the world to accommodate new information that does not fit into an existing slot or category. Accommodation occurs in order to deal with the problem of holding contradictory beliefs, a state that Piaget referred to as cognitive disequilibrium. Piaget believed that intellectual maturation was dependent upon cycles of assimilation and accommodation, and that cognitive disequilibrium was the key to the process.

Piaget was concerned primarily with children's development, and knowledge construction for children necessarily involves a great deal of accommodation. This is because their existing models are ill formed and weak, driven as they are by limited experience with the world. As we mature, we spend more time in assimilation than accommodation, because our mature mental models are more robust, having been tested and modified over many years and many encounters with new information. One might erroneously conclude that accommodation and cognitive disequilibrium are most relevant in the early years of development, and less so during our adult years.

However, it is important to recognize that just as children have weak models due to lack of exposure in general, adults as novices in a new domain of knowledge or experience share the same characteristics. Thus one can think of development both in terms of maturation and in terms of expertise in a given domain. The design of any learning endeavor, including DGBL, in which novices are moving toward expertise by necessity, must attend to cognitive disequilibrium, accommodation, and assimilation.

Question asking and assimilation and accommodation are also closely related to metacognition and self-regulation in learning. The more the learner is responsible for in the learning process, the deeper and more efficient their learning will be. This is the same principle behind what Vygotsky (1962, 1978) called the zone of proximal development and scaffolding. The essence of these concepts is that there are some activities (physical or mental) that are completely within the ability of the learner to accomplish on their own. Alternatively, there are some things that are completely beyond the ability of the learner to accomplish regardless of whatever assistance might be given them. Somewhere in between these two extremes lie tasks that the learner is capable of doing with some assistance, often from adults or more competent peers, or in the case of games, from role-playing characters, hints, and responsive interactive features of the game space. Vygotsky believed that this area, or zone, was where learning occurred. More importantly, he believed that the most effective learning would occur when students were in the upper ranges of this zone of proximal development (ZPD). In other words, educators should strive to provide just *enough* assistance and guidance to allow the learner to progress, but no more. This supports the concept that by maximizing the role of the learner in acquiring knowledge, we maximize learning. Vygotsky called this minimal assistance the "scaffolding" for its metaphoric relation to the role it plays in building and supporting learning.

Scaffolding also theoretically promotes cognitive self-efficacy, the learner's perceived competence in learning within specific domains. Because learners in the ZPD are continually accomplishing things that would normally be beyond their abilities, their perception of their own abilities in regards to the learning will also improve (Bandura, 1997). This is important because cognitive self-efficacy (beliefs about one's competence and ability) improves learning *independently* of metacognitive abilities (the ability to monitor their own learning and select and use appropriate cognitive strategies) (Bandura, 1997). When consistent and accurate feedback is available to increasingly confident learners, they will regularly set goals near the upper limit of their zone of ability in an effort to monitor and ensure their own progress (Bandura, 1991; Bandura & Cervone, 1983, as cited in Bandura, 1997).

Games embody the processes of cognitive disequilibrium, accommodation, and scaffolding. Interacting with a game requires a constant cycle of asking questions, and forming, testing, and revising hypotheses. This cyclic learning process happens rapidly and frequently during the game, with immediate feedback. The extent to which games foil expectations (create cognitive disequilibrium) without exceeding the capacity of the player to succeed (going beyond the zone of proximal development), determines to a large extent whether they are engaging.

Games also often provide hints to help advance the game, either directly (through characters or game hints) or indirectly (the wall rumbles but does not move when pressed, indicating you have found a secret door, but it is locked). Hints serve as scaffolding for new learning, and have been shown to facilitate the active construction of knowledge (e.g., Graesser, Person, & Magliano, 1995; Lepper, Aspinwall, Mumme, & Chabay, 1990; Merrill, Reiser, Ranney, & Trafton, 1992). Far from interrupting the game, hints are sought out by players, who populate message boards and Web sites with questions, and who purchase not solution books, but hint books. Players instinctively seek out the least amount of help necessary to advance them through the game, intuitively implementing a scaffolding approach to keep themselves in the ZPD.

Making this guidance and scaffolding process a more regulated feature of the game would lessen the extra steps players must now take outside the game environment to get assistance. Indeed, many games provide these kinds of resources through lists of questions the player can choose from when speaking with a character in the game, or through agents in the game who can provide limited hints.

The Problem of Integration

We have seen that traditional approaches to direct instruction are incompatible with games. The result of ignoring this incompatibility during the early years of edutainment development led to what Papert (1998) called Shavian Reversals, which in turn led many to conclude that games and education cannot be merged. However, we now also see four principles of learning in games that are tied to established educational and instructional learning theories. These principles can provide relevant guidance as a new generation of DGBL designers strives to find ways to develop new games for learning.

There are several questions that we must answer in our quest to integrate instructional events, strategies and subject area content into games, some driven by the principles outlined earlier in this chapter, and some that are driven by practical reality. These questions will guide and inform the balance of this chapter:

1. What mechanisms exist in other fields that can be used to present content within a game in a way that is compatible with the game and game principles?
2. What mechanisms exist in other fields that can support the principles of scaffolding, question asking, and problem solving?
3. How must these mechanisms be modified according to the principles outlined here and other theories or approaches?
4. How, assuming we can answer the first three, can we make sure that intelligent learning games are extensible to multiple problems and domains, and ensure that any content expert can generate content for these games without "sucking the fun out" of them?

These questions are addressed in the following sections. The answer to question one, I believe, lies in pedagogical agents. The answer to question two lies in the field of artificial intelligence and Intelligent Tutoring Systems. The answer to question three lies in the principles I have outlined earlier and in the area of narrative in games and learning. The answer to question four lies in the use of a form of electronic performance support system (EPSS) called authoring tools.

Pedagogical Agents

One of the problems we face when attempting to bring learning content into a game is that the content is often at odds with the game world, storyline, or structure of the game. While it may be realistic to say that a character in a game who does not know something can go to the library or return to headquarters for training, such strategies are assumed to take place off camera. If we instead force the player to wade through electronic texts or a tutorial, we have interrupted the game, or what Csikszentmihalyi (1990) calls flow—the optimal learning state that occurs when learners (or game players) are immersed in an activity to the extent that they lose track of time and the outside world. Flow is the ultimate goal of any game designer and educator, and anything that interrupts it should be avoided. Often, designers will use PDAs and tablets as tools within the game context to help deliver information. Such devices partially satisfy Principle 3 (situated cognition and learning), but they cannot be used extensively without violating Principle 1 (play theory, a continuous cycle of learner input and feedback, engagement), and making it difficult to address Principle 4 (question asking, cognitive disequilibrium, and scaffolding), at least in the same manner that games do.

One partial solution to this problem lies in an area of study in cognitive psychology and instructional design called pedagogical agents. Pedagogical agents are typically used in computer-based instructional environments where learners interact with a computer-based character to get advice, feedback, or instruction. Pedagogical agents can look like humans, animals, inanimate objects, or fantastic creatures (e.g., genies or space aliens). The increasing research on the use of pedagogical agents in learning environments (e.g., Baylor, 2000; Baylor & Kim, 2005; Baylor & Ryu, 2003; Graesser, VanLehn, Jordan, Rose, & Harter, 2001; Johnson, 2004; Lester, Converse, Kahler, Barlow, Stone, & Bhoga, 1997; Moreno, 2004; Moreno, Mayer, Spires, & Lester, 2001) presents a rich resource to draw on for the use of agents in games. Perhaps the most well known (though not most effective) agent is "Clippy," the paperclip character in *Microsoft Word*. The idea is that he/she pops up with information that might be helpful to the person working with a word-processing document at the time. Clippy takes on a persona that is tied to the context in which he/she appears. A paperclip, after all, is a highly recognized part of an office that is relevant to printed documents. By embedding help systems in pedagogical agents, we provide social aspects of learning when humans cannot be present, and we do so without violating the environmental context. These animated pedagogical agents can take on different roles, or persona, in computer-based instruction, including assistant, pedagogical expert/mentor, learning companion (Baylor, 2001), and motivators (Baylor & Kim, 2005), making them ideal candidates for these same roles in intelligent learning games (ILGs).

Likewise, then, a pedagogical agent in a game allows us to provide the instructional content within a game, without having to resort to decidedly non-game-like methods such as stopping to read a manual, or using a built-in tutorial. Actions like those just described do not fool game players—they know they have been asked to leave the game world for that of the classroom. Agents in games can provide the necessary content in the context of any of several different roles or personas within the game including as co-investigators, mentors, team members, or peers with content knowledge.

However, using pedagogical agents merely trades one problem for another; how do we best use agents to deliver content and provide learning guidance in a game? Agents are compatible with the situated cognition nature of games (Principle 3), but though ostensibly part of the game play are not in and of themselves part of the constant cycle of learner input and feedback (Principle 1), nor are they by themselves part of the problem solving, question asking, cognitive disequilibrium, or scaffolding of learning that occurs in games (Principles 2 and 4). It does no good to simply replace the PDA in a game with a human tutor who delivers the content in a tutorial approach, which would clearly not be part of the game, and might be much more irritating. Research has shown that people treat pedagogical agents the same way they treat humans in terms of their expectations and responses (Reeves & Nass, 1996). As a player, I am not expecting an online tutorial to act like a human, so while it may be annoying, I am at least getting what I expect. A character who comes up to me as part of the game, however, and instead of answering questions in a forthright manner begins to quiz me, evades my direct questions, and provides long didactic statements is likely to receive the same response I would give to a co-worker who did the same thing.

Just as importantly, an agent who behaved this way would clearly not be a part of the game world, regardless of the role he or she held in the game. Every character in a game (including the player) is guided entirely by the goals and sub-goals in the game. Non-player characters (NPCs), characters controlled entirely by the game, never step "out of character"; all their actions take place within the context of their roles and the narrative structure of the game. The player, likewise, is guided entirely by her or his desire to solve the game, again within the narrative structure of the game.

Pedagogical agents in and of themselves do not entirely solve the problem of content integration and instructional guidance in games, because most current agents (there are none in games as of this writing) utilize the same instructional approaches that doom edutainment titles—dull, boring monologues. What is needed to bring agents into games is an approach that more closely reflects the natural exchange of information among characters and players in the game world. When only one character or entity in a game has the content knowledge, the learning will of necessity interrupt the game flow as the player interacts with that one person. And while much can be done to keep the style of instruction thematically connected to the game in such cases, there is no getting around the fact that you have just spent 20 minutes (or more!) talking to someone to learn something to move on in a game.

So it is not enough for a pedagogical agent to be merely thematically tied to the game (e.g., a trainer within the game world); the agent must be a character who is engaged in advancing the goals and story of the game world according to the same motivations and constraints as other characters in the game. They must present any instructional content as contiguously as possible to the events in the game that require the application of that content. What we need to integrate agents into games, then, is a pedagogical approach that is compatible with this constant cycle of player action and game feedback, the "conversation" between player and game.

There is an established pedagogical technology that can address Principles 1, 2, and 4, and because it has already been combined with pedagogical agents, potentially can address Principle 3 as well. Intelligent tutoring systems (ITSs) have been around for more than 30 years, and engage the learner in problem-solving tutoring conversations with frequent student contributions and feedback. Moreover, they do this through natural language processing (so the learner can say anything they want and the system will respond appropriately) and recently, with pedagogical agents acting as the tutors.

Intelligent Tutoring Systems

ITSs are computer-based applications that engage learners in a tutoring dialog to help them construct knowledge in a given domain. The system attempts to get the student to articulate knowledge about a given topic, and provides feedback and structured guidance (e.g., hints, prompts, etc.) through conversational "turns." ITSs have been around for more than 30 years, but have essentially remained unchanged from the three features proposed by Hartley and Sleeman (1973): the Expert Model, Student Model, and Tutor. The expert model is the component of the tutoring system that contains all that an expert in the domain knows. This component can be thought of as a recognized expert in the field or a definitive textbook on the subject.

The expert model represents not only the content of the tutoring system, but also the model for how that content organized. The dual representation comes from the fact that experts not only know a lot about the given content, they *organize* what they know in ways that are efficient. Learning theorists refer to these organizations as schemas, mental models, or propositional networks. One way to think about the role of organization of knowledge is to think about a spreadsheet with multiple columns of data. One could sort the data by any column, and the information would in each case be equivalent, but the ordering would allow different patterns and meanings to become more visible.

Experts have very well-developed schemas and mental models for what they know, developed through assimilation and accommodation over a long period of experience with the domain, with multiple connections and relationships between concepts. Learning theorists suggest that it is therefore not enough for a learner to know only *what* an expert knows, but to also know it in the same *way* that an expert knows it. In other words, we want the student's mental model of the domain to eventually approximate the mental model of an expert.

The expert model thus approximates both the content knowledge and the structure or organization of that knowledge. The tutoring system then uses the expert model as a source of knowledge and structure for that information. The goal of the system is to reduce the disparity between the expert model and what the learner knows. Of course, the system does not have direct access to what the student knows, and so it must develop a representation during the tutoring session by tracking what the student says that is correct, incorrect, or irrelevant, during the tutoring conversation. This representation is called the student model. Each time the student articulates something, the system compares what was said to what it knows about the structure and content of the domain (the expert model) and determines how closely the two are aligned. It then modifies the student model to reflect its best guess about what the student knows, and selects the best pedagogical response that it believes will reduce the gap between the student model and the expert model, while maintaining the principle of scaffolding.

The means by which the system makes these pedagogical decisions is called the tutor. Through a series of different pedagogical and conversation moves, the tutor provides feedback to the student based on what it gets from the expert model and on the underlying pedagogical approach used by the system. Tutors generally operate under the assumption that the best learning occurs when the student contributes the most to the tutoring dialog (similar to our earlier discussion of self-regulation, efficacy, problem-solving, and scaffolding). Accordingly, the tutor tries to get the learner to articulate as much as possible without any assistance. As the learner hits the limit of his or her knowledge and either is unable to contribute anything further or begins to make errors, the tutor provides corrective feedback, but again just enough to get the learner to be able to contribute more to the tutoring conversation.

There are many examples of ITSs (e.g., PACT, Andes, Atlas, Why Tutor, Why2, LISP, Smithtown, Sherlock, Stat Lady, Geometry Tutor, *AutoTutor*) that have been shown to be effective in teaching computer literacy (Graesser et al., 1999), algebra, geometry, computer languages (Anderson, Boyle, & Reiser, 1995; Bonar & Cunningham, 1988; Koedinger, Anderson, Hadley, & Mark, 1997, Schofield & Evans-Rhodes, 1989), physics (Gertner & VanLehn, 2000; Graesser et al., 1999; VanLehn, 1996), and remediating misconceptions (Stevens & Collins, 1977). In all cases, the tutors result in learning gains, reduction of instructional time, or both. Research shows that human tutors result in a .4 to 2.3 SD gain in learning over traditional classroom instruction (Graesser et al., 2001). ITSs, on the other hand, result in a .3 to 1.0 SD increase in learning (Corbett, Anderson, Graesser, Koedinger, & VanLehn, 1999). So while ITSs are not as good as humans are in all cases, they produce reliable and significant learning gains when they are used.

As mentioned earlier, ITSs are effective in part because they embody Vygotsky's concept of scaffolding and the zone of proximal development (Principle 4) and in part because they require constant input from the learners and promote a frequent cycle of response and feedback (Principle 1). One other reason that tutoring is effective is that it employs problem-based learning (Principle 2). ITSs specify complex problems that require application and synthesis of lower-level skill and facts. Rather than presenting information in direct instruction fashion, the system presents a complex problem that the student and system work together to solve, with the system pushing the student to do as much of the work as possible, and providing guidance when and if needed.

To illustrate how ITSs manage this process, it may be helpful to examine a typical problem scenario used by an actual ITS, *AutoTutor*. This has the added advantage of illustrating an ITS that has already integrated a human-like pedagogical agent (in this case, a 3D talking head), illustrating that integrating an ITS into a game is not *much* more complicated than integrating a pedagogical agent into a game.

AutoTutor was developed by the Institute for Intelligent Systems' Tutoring Research Group at the University of Memphis. *AutoTutor* uses a talking head agent to engage in a tutoring dialog with the learner. The agent uses gestures, facial expression, and synthesized speech. In a typical tutoring episode, *AutoTutor* presents a tutoring problem to the learner. The expert model has been defined for the system by an expert in the given domain (e.g., physics). A typical problem might be as follows:

"Suppose a runner is running in a straight line at constant speed, and the runner throws a pumpkin straight up. Where will the pumpkin land? Explain your answer."

The learner responds by typing his or her response, although the capability for speech recognition exists as well. As soon as the learner responds in any way, several things happen in the background. While the specifics of this process are very complex, they can be characterized adequately by discussing them in more general terms (see Graesser, Person, & Harter, 2000; Graesser et al., 2001 for more detailed descriptions). *AutoTutor* examines the student answer and compares it to expected good and bad answers that have been pre-specified by an expert. The goal is for the student to articulate an "ideal answer" that encompasses not just the correct response, but also an elaboration of the logic of that response, drawing upon relevant concepts, principles, and rules of, in this case, physics. However, this comparison is not a simple matching process, but a sophisticated set of computations. *AutoTutor* first takes the learners response and runs it through a language module (syntactic parser) to decompose the utterance into its component elements and to categorize it as an assertion, question, metacognitive comment, and so forth. Latent semantic analysis (LSA) is then used to compare the student response to a whole host of possible responses using sophisticated mathematical computations to determine the likelihood of a match. Based on the type of response returned by the parser and the LSA result, *AutoTutor* does one of several things. If the response is a question, *AutoTutor* either provides the answer from a corpus of knowledge, for example, the electronic text of a physics textbook, (if the student prompted the agent with a domain related question), indicates yes or no (in response to a question that is right or wrong), or provides the one of several other possible responses based on the learner's request (repeats the questions, clarifies a point) or the appropriate tutoring move (prompts the learner to elaborate, provides a hint, etc.).

As the learner articulate parts of the answer, *AutoTutor* updates the student model accordingly, and adjusts the dialog to address areas the student has not yet articulated. In this process, *AutoTutor* uses a series of hints, prompts, and assertions to get the learner to contribute as much as possible (scaffolding). This behavior is based on existing research on tutoring and the dialog moves that are made by novice and expert tutors.

So ITSs provide a pedagogical approach that is compatible with learning in games because they promote the continuous cycle of student response and feedback (Principle 1), embody problem-based learning (Principle 2), emphasize learner contributions and promote question asking (Principle 4), and employ naturalistic conversation and dialog, all of which games do naturally. Because ITSs have already been combined with human-like agents (e.g., *AutoTutor*), and because many people support blending ITSs with agents in games (e.g., Laird, 1999) and immersive environments (e.g., Ravenscroft & Matheson, 2002; Rickel, Marsella, Gratch, Hill, Traum, & Swartout, 2002; Shute & Psotka, 1996), ITSs represent an ideal way of bringing content and instructional guidance into a game according to Principle 3 as well. One can easily imagine the architecture running in the background to govern the dialog of the characters within the game during relevant portions of a game, and lying dormant during portions of the game where the game AI is dominant.

But once again we find that we have traded old questions for new. The models of dialog in tutoring are only marginally more naturalistic than are didactic tutorials or lectures. Their purpose is still quite obviously to teach content, which as a goal outside the game in which they reside, will always feel like a separate endeavor. Players would quickly become irritated with an agent who continually responded with suggestions for them to think more about the material, or who gave them hints and prompts instead of direct information relevant to solving the problem at hand. These are effective moves and strategies in learning environments, but are incompatible as formulated for game environments and the characters within them. Also, tutoring dialogs are designed to take place in one-to-one environments, but games require the learner to interact with multiple characters and in many places—rarely does the player stay in one place or interact with one person for any length of time. No matter how natural or integrated an agent is in a game, if the learner must interact with that character for long periods of time, it will be difficult to maintain the integrity of the game.

What is needed is a way to re-conceptualize the dialog structure to account for the game environment. One way to do this is to study how this kind of tutoring dialog could be adapted to include multiple agents (all connected to the ITS) in the game, and to determine how traditional tutoring dialog structure could be modified to take place during shorter interchanges, and in a more distributed (among agents and within the game environment itself) fashion. To understand how we might begin to do this, we can study how narrative and discourse are structured in both ITSs and in games to see where they align and how they could potentially be modified in ITSs to allow for a more seamless integration in intelligent learning games.

Narrative

"…If you're going to tell a story, you must not only master what storytellers in other media already know, you must also learn to adapt that knowledge to games" (Bates, 2001, on the prevalence of story elements in modern games).

There are three reasons we should attend to narrative in our design of intelligent learning games. First, narrative is among the most effective learning strategies and may in fact be at the heart of all effective learning. Second, narrative is arguably the dominant feature of games when they are conceptualized as texts that are read and co-constructed through player interaction. Third, narrative psychology studies how people makes sense of inanimate and animate objects. Games, and the characters within them (including agent/ITSs) straddle the world between animate and inanimate objects, and can thus be informed by attention to narrative psychology principles.

Narrative as a Learning Strategy

Narrative is among the oldest forms of learning, predating the written word. Oral histories and traditions were the only means of education for the approximately 7,000 years prior to the first written language by the Sumarians in 3,200 B.C., and the written language preserved much of this tradition in narrative form from then on. Narrative still drives much of what we do outside of learning environments as well, with books, television, and film remaining among the most popular forms of information organization and sharing. The prevalence of narrative as a strategy for learning makes it perhaps the most powerful and accessible instructional strategy available. Narrative itself is simply a structure for organizing information (by theme, chronology, etc.). As such, it is a powerful means of organizing knowledge while preserving the relationships between ideas and concepts; conditions well suited to higher-order learning such as problem solving and the attendant processes of assimilation and accommodation.

We have already seen how problem-based learning, one of the most effective means of learning, involves the generation of questions and predictions (themselves effective learning and metacognitive strategies) and a running narrative of proposed strategies for acquiring new information, assessment of progress (e.g., dead ends versus effective strategies), and reflection on the solution once achieved.

We have also seen how self-efficacy influences performance. Perhaps even more than with problem solving, self-efficacy is closely related to narrative. It is, in effect, the story we tell ourselves about our abilities and place in the world of knowledge (I am a good learner; I am a poor learner; I cannot do math). As such, it is a part of the larger narrative we construct for ourselves about who we are in the world, what our place is, how we have gotten to where we are, and where we hope to go. Narrative is at the heart of all knowledge, not just as a learning strategy, but also as a *way* of knowing. For example, what is science but the story of how the universe and we came to be, backed up by observation and experimental data? A lack of attention to narrative is what accounts for our inability to relate to current agents; "observers have difficulty understanding them narratively" (Sengers, 2004).

Narrative in Games

David Braben, president of Frontier Developments and creator of the seminal game *Elite* and the popular *Rollercoaster Tycoon 3*, believes that storytelling and emotions are the biggest missing element in games, and will be what game designers will be working on in the future (Braben, 2002). Phoebe Sengers (2004) argues that the "juice" missing in current agent technology (what makes them "soul-less") is narrative. These sentiments are representative of the importance game designers and educators place on the use of narrative in games and DGBL (for more on this topic, see a recent collection of essays on the subject called *First Person: New Media as Story, Performance, and Game*, edited by Noah Wardrip-Fruin and Pat Harrigan, and published in 2004 by MIT press).

There is a seamless dialog, or discourse (albeit non-verbal), that exists between learner and game, in which the learner poses hundreds, if not thousands, of questions, hypotheses, and statements about the game (that is probably a door; something on the other side is worth seeing; there may be somebody around that corner; I bet I can pick up something I need here, back at the shop). As the player acts on and tests these questions and beliefs, the game immediately and consistently provides answers in the form of feedback (this does not move when you push it; it is locked; there was somebody around the corner and he just shot you; you now have a pry bar that can be used here). It is impossible for the game to progress without constant action on the part of the player.

Contrast this with traditional direct instructional methods like the lecture, textbook reading, watching video, and so forth, which are often characterized by long stretches of content "delivery" without response and feedback from learners. This is not to say that direct instruction does NOT require learner interaction. In fact, of course, any instructional designer will build in opportunities to elicit responses from the learners and provide feedback, and indeed the most effective learning (direct instruction or not) occurs when the learner is contributing questions, restating knowledge, and getting feedback. However, when there are large amounts of "information" that are part of the instruction, the opportunities for student contributions are limited and generally not spontaneous (i.e., the instructor elicits student responses directly).

In contrast, games automatically demand player input and do so frequently. Perhaps more importantly, most student contributions are self-initiated, rather than elicited. This is much more closely aligned with the principles of constructivism. Constructivist views of learning hold that we create meaning (knowledge) through our interactions with ideas, content, people, and so forth. The meaning derived through this interaction will differ from person to person based on that person's individual characteristics (e.g., their prior knowledge, beliefs, attitudes, sensory abilities, etc.) and the characteristics of the content or ideas or people they are interacting with. Thus, knowledge is not static, but fluid and highly dependent on contexts and individual differences. With games, then, we can conceptualize the experience as being co-constructed by the player and the game. James Gee discussed this co-construction of the game world and story in an online colloquy in 2003: "A game is an intricately designed world that encourages certain sorts of actions, values, and interactions. At the same time, the player co-designs the game's world by the actions and decisions the player takes. The player brings the world alive and in open-ended games every player ends up with a different world and having played a different game." Other game researchers echo this interpretation of games as a co-constructed narrative (e.g., Sengers, 2004). Anything that interrupts this conversation, if you will, will interrupt the game flow.

Adventure games rely perhaps the most heavily of all game genres on the use of traditional conceptions of narrative: narrative in the sense of the back-story of the game, the actual story generated through computations based on the interaction of player and game, and the running narrative players create for themselves as they consider how they will approach the game at any given moment. Just as narrative is essential to adventure games, narrative is also a powerful instructional strategy. All learning is, in one sense, a running narrative. A learner's journey from novice to expert in a domain is a kind of narrative; the running metacognitive commentary learners maintain about what they know, what they are learning, and how the things they learn relate to each other is a kind of narrative; the assessments we require of learners are a narrative both of the right answer and the rationale for it.

There are some in the gaming community who question whether interactivity and narrative are compatible goals in games (Bates, 2001; Jenkins, 2004), arguing that the needs of learner control, flexibility of game flow, and immediacy of response do not allow for the traditional components of a linear narrative. But narrative is much more complex than a simple linear definition implies. Games, as a new form of media, are also a new form of narrative, as Henry Jenkins (2004) argues in his excellent treatment of this issue in "Game Design as Narrative Architecture."

"One gets rid of narrative as a framework for thinking about games only at one's own risk" (2004). Designers should think less of themselves as storytellers than as narrative architects: designers of narrative story spaces and environments. The game serves as a frame for a story that is co-authored by the interaction of player and game. "One can imagine games taking their place within a larger narrative system with story information communicated through books, film, television, comics, and other media, each doing what it does best" (Jenkins, 2004). It takes only a small stretch to imagine instructional content as one of those other channels, even within the game environment itself. Jenkins refers to this as the concept of embedded narrative, where "comprehension is an active process by which viewers assemble and make hypotheses about likely narrative developments on the basis of information drawn from textual cues and clues." Instructional content as "embedded narrative" has obvious application to problem solving and learning in games, as does his assertion that such stories are not chronological and linear, but more a body of information distributed across the game environment, as might be the case with the interplay of agents, content, game space, and learning described earlier. In fact, some argue (e.g., Pearce, 2004) that game narrative is by definition incomplete: "It must be in order to leave room for the player to bring it to fruition." Here again we find further evidence for the requirement of learner/player participation in game, which we have already seen is a critical element in problem-based learning and perhaps all instructional methodologies.

Narrative Psychology, Socially Situated Agents, & Games

Narrative is critical also in helping us determine how to integrate agent-based ITSs into games. Narrative tools like plot structure, character development, dialog, and conversation will help us ensure that agent-ITSs are subservient to the game problems and goals, are realistic and compatible with other NPCs, and interact with other NPCs and the player in appropriate and believable ways. The field of narrative psychology accounts for how we make sense of animate objects in the world. Inanimate objects we understand in terms of cause and effect rules, whereas an animate object "is made comprehensible, not by figuring out its physical laws, but by structuring it into narrative or 'stories'" (Sengers, 1999, page 3). Because human-like agents and NPCs in games straddle both the animate and inanimate world they will be understood by game players as animate characters and woven into the story learners create through their interaction with the game. Accordingly, we have to attend to making these characters as human and natural to the story and game space as possible, rather than the "fragmented, depersonalized, lifeless, and incomprehensible" pedagogical agents that currently exist (Sengers, 2004, page 95).

So what will socially situated agent/ITSs look like in games? We can find part of the answer in an examination of problem solving in the real world. In our everyday lives, we rarely interact with just one person when seeking assistance. To be sure, we will seek out someone with expertise in the area for which we need assistance, but that person will differ by area. Accordingly, we should not have just one agent/ITS in the game who provides all assistance. Knowledge should be distributed among multiple agents in the game, just as it is in the real world. Additionally, although we seek out different experts when needed, this is not the only reason we seek those people out. We interact with people differently depending on the context and nature of our daily lives; sometimes as friends, colleagues, lunch mates, superiors, significant others, and so forth. Our human-like agents should have as many roles as those we work with on a day-to-day basis, and interact with us accordingly. Their roles should change with the social context they are in at the time game players interact with them.

Finally, the nature of problem solving is often a collaborative effort involving our seeking help and advice from many sources, both to gather the information needed as well as to help with the problem-solving strategies. Sometimes the people we interact with have the knowledge we need; sometimes they have strategies for getting that knowledge. This too should guide our design of agent/ITSs in games. We should seek assistance from multiple characters and they should seek assistance from us; sometimes we will have information that is needed, and sometimes they will have it. This is not to say that all characters in the game should be intelligent or even on our side; just that our instruction should come from multiple sources, and in a form that is consistent with human interactions and naturalistic problem solving. Sengers (2004) calls this "socially situated AI," and makes a convincing argument for how we can design more natural agents by attending to several factors, the foremost of which is that the social and culture environment in which the agent is situated. By extension, the agent then should be "socially" aware of the game environment, the other agents and NPCs, and the player, through her or his interactions with the game.

Authoring Tools

Authoring tools are the final key to implementing the process of designing intelligent learning games. Because game designers are generally unfamiliar with the educational aspects outlined here, and the educators are generally unfamiliar with game design, we need ways for non-game designers (e.g., content experts and instructional designers) to generate content for games so that we are not asking game designers to be instructors, nor instructional designers to be game developers. And we need to find ways to make this process scalable and portable, so that when appropriate, multiple learning games can be created and modified by modifying the content and connecting it to the game. What is needed is authoring tools that support both designers and educators as they develop DGBL, tools that embed the knowledge from both sides in their very structure.

The *Greenwood Dictionary of Education* defines an authoring tool as, "A software application designed for use by a non-computer expert to create computer programs… Authoring tools are designed to be used by individuals without substantial programming knowledge or skills" (2003).

Authoring tools are a kind of expert system, and as such are closely connected with the evolution of artificial intelligence as a field. Expert systems were initially conceptualized as computation tools that represent an expert's knowledge in a given domain. As such, they are similar to the expert module of an ITS. However, their purpose as expert systems was to support, augment, or in some cases replace human decision-making (Shute, 1985). Authoring tools grew out of a desire to automate the *creation* of expert systems, which initially were very difficult to create. The idea was to create tools that would allow a subject matter expert (SME) with little or no knowledge of the programming requirements for building a system (or even how a system worked) to generate the content needed by an expert system through simplified computer interface.

ITS developers faced a similar problem recently, as we now do with intelligent learning games. Ideally, a game should act as a shell, which can become a number of intelligent learning games in multiple domains through the creation and specification of new problems and content areas. This is necessary both for reasons of cost of development, and because we need to be able to have experts generate content without having to know how to program games. The authoring tools should translate what the expert knows into something that will be compatible with the game and intelligent agent architecture we embed in it. If not, intelligent learning games will not only be too expensive to develop, but will engender an expertise/content bottleneck that precludes their widespread application. This approach maximizes both what game designers know about games and what experts and educators know about instruction, without asking either to develop expertise in the other domain.

It is important to realize that the authoring tool serves to translate what an expert knows (content knowledge) into the pedagogical framework that has been established for bringing content into a game through agent/ITSs. Neither the subject matter expert nor the game designers must know anything about each other's domains. In order to understand what authoring tools are and how they will impact DGBL as outlined in this chapter, it may help to see how a development team faced and solved a similar problem in the development of *AutoTutor* (the ITS described earlier).

One of the chief challenges in the ITS enterprise has been building authoring tools for developing new content. For example, it initially took a little under three months to develop an *AutoTutor* version on a new topic. This was far too slow a process and relied on too limited a resource pool (those with knowledge of the inner working of *AutoTutor*) for widespread adoptions and diffusion of this learning technology. We needed to find a way to increase the number of people who could generate the content needed by learning technologies like *AutoTutor*. This is the same resource and expertise bottleneck we will face with building ILGs.

In the case of *AutoTutor*, the curriculum scripts (content) were written for the benefit of the tutoring system, not for the end user, and their development was driven by computational needs rather than human factors. The result was that they were complex and difficult to understand for those who were not privy to their development and the relation of their components to the various components of the tutoring system itself. Faced, as we were then, with the prospect of dissemination to educators in all domains and educational levels, we had to decide now how to map what the system needed in the form of these scripts to what SMEs know about content and instructional preferences. This is the same chasm we must bridge between the worlds of game design and the educational and psychological principles and applications outlined in this chapter; how will we support those with expertise on one side or the other as they develop ILGs?

One way to build this bridge between what an expert in one domain knows and what the learning technology needs lies in Electronic Performance Support Systems (EPSS). This term was first coined by Gloria Gery (1991) in her book by the same name. She defined an EPSS as,

An integrated electronic environment that is available to and easily accessible by each employee and is structured to provide immediate, individualized online access to the full range of information, software, guidance, advice and assistance, data, images, tools, and assessment and monitoring systems to permit job performance with minimal support and intervention by others.

EPSSs are usually knowledge management or decision-making support systems, but their definition and purpose has grown over time. Bill Miller (1996) broadens their definition:

"An electronic performance support system is any computer software program or component that improves employee performance by either reducing the complexity or number of steps required to perform a task (process simplification), providing the performance information an employee needs to perform a task, or providing a decision support system that enables an employee to identify the action that is appropriate for a particular set of conditions.

The *AutoTutor Script Authoring Tool* (*ASAT*) was the first prototype of such a tool for use with *AutoTutor*. This system adopts a coaching methodology and language that is familiar to subject matter experts to help guide them as they develop the type of content needed by the system. The pedagogy and vagaries of the system (i.e., tutoring, hints, prompts, and scripting tags) are hidden, and are instead embedded in the tool itself. The end result is that an expert in some subject area is supported by the tool as they develop content appropriate for the system. Once we have established models for blending the different principles, approaches, and technologies outlined in this chapter, we can embed those models into authoring tools to support those who would develop the content for the new generation of digital learning games.

This also yields the solution to the problem of scalability, since more people can develop the content for the games, more quickly. Perhaps more importantly, authoring tools also open the door for generating a single game environment with dozens of scenarios for learning without having to redesign the entire game each time our learning goals change, although significant challenges face us in accomplishing this.

Summary

This chapter began by examining games for existing principles that can help guide the integration of instructional content with games. Through this process, the chapter identified the following four principles:

1. Games employ play theory, cycles of learning, & engagement
2. Games employ problem-based learning
3. Games embody situated cognition & learning
4. Games encourage question-asking through cognitive disequilibrium and scaffolding

Any attempt to build content into games that hopes to avoid Shavian Reversals must adhere to these principles. From there, I discussed research, theory, and technologies from other disciplines (artificial intelligence, psychology, English, and instructional design) that I believe can help us build intelligent learning games. The structure of this discussion was guided by the following questions:

1. What mechanisms exist in other fields that can be used to present content within a game in a way that is compatible with the game?
2. What mechanisms exist in other fields that can support the principles of scaffolding, question-asking, and problem solving?
3. How must these mechanisms be modified according to the principles outlined here and other theories or approaches?
4. How, assuming we can answer the first three, can we make sure that intelligent learning games are extensible to multiple problems and domains, and ensure that any content expert can generate content for these games without "sucking the fun out" of them?

The answer to each of these questions, is that pedagogical agents can serve as vehicles for instructional guidance and co-construction of knowledge within the game, adhering as they potentially do to the situated nature of agents in the context of the game world, but that they must be guided by a pedagogy that reflects the problem-solving and "conversational" nature of player/learner and game interaction. Intelligent tutoring systems, by virtue of their use of tutoring dialogs, problem-solving, continuous learner-system interaction, and emphasis on self-regulation and scaffolding, can provide this pedagogical approach. Because ITSs have already been combined with pedagogical agent technology (e.g., *AutoTutor*), models already exist for how the two technologies can be combined for use in intelligent learning games.

Neither ITSs nor pedagogical agents are yet suitable for integration into games, however, because they ignore important aspects of play and social/culturally situated agents. The field of narrative, and in particular narrative psychology provides clues to how these systems must be modified, including socially situated AI (Sengers, 2004), and dialogs that are suitable for the distributed kinds of problem-solving strategies we employ in the "real" world.

Finally, authoring tools allow us to constrain the input of subject matter experts to fit the new pedagogical model for intelligent learning games created through the approaches outlined here. This allows game designers to design games, SMEs to design content, and authoring tools to design the learning. This process maximizes quality, minimizes costs, and ensures more widespread application of intelligent learning games across domains and populations.

There is a big difference, of course, between describing the process of blending all these technologies and actually getting them to work together. None of these approaches are perfectly suited for integration with games without modification. We need to find ways to modify ITSs so that the tutoring is done in a distributed manner, with conversational turns occurring between player-agent, agent-agent, and perhaps even multiple agents within each conversational turn. We need to explore discourse models that are appropriate to team and group interaction, rather than the typical tutoring dyad of tutor and student, and to examine how question asking and problem solving are done in group situations. These models will in turn drive and constrain the design of the ITSs and the agents.

Finally, our authoring tools will need to reflect the needs of all of these elements as well as the manner in which the agents and content are connected to the game engines (no small feat itself). By looking to these established models and approaches rather than reinventing the wheel, however, we are well on the road to developing a new model for integrating learning into games that draws from decades of proven techniques in learning while staying true to the power and nature of games. While the map is not the journey, it is nonetheless the pre-requisite to planning and taking that journey.

References

Abdullah, M. H. (2001). *Self-directed learning*. ERIC Digest (ED459458). ERIC Clearinghouse on Reading English and Communications. Retrieved December 18, 2005, from http://ericdigests.org/2002-3/self.htm

Anderson, J. R. (1995). *Cognitive psychology and its implications* (4th ed.). New York: W. H. Freeman.

Anderson, J. R., Boyle, C. B., & Reiser, B. J. (1985). Intelligent tutoring systems. *Science, 228*, 456-462.

Bandura, A. (1997). *Self-efficacy: The exercise of control*. New York: W. H. Freeman and Company.

Bates, B. (2001). Story: Writing skills for game developers. Presentation at the Game Developers Conference, San Jose, CA, March 20-24. Retrieved October 5, from http://www.gdconf.com/archives/2001/index.htm

Bates, B. (2002). Game design: The art and business of creating games. Indianapolis, IN: Prima Tech.

Baylor, A. (2000). Beyond butlers: Intelligent agents as mentors. *Journal of Educational Computing Research, 22*(4), 373-382.

Baylor, A. (2001). Permutations of control: Cognitive consideration for agent-based learning environments. *Journal of Interactive Learning Research, 12*(4), 403-425.

Baylor, A. L., & Kim, Y. (2005). Simulating instructional roles through pedagogical agents. *International Journal of Artificial Intelligence in Education, 15*(1).

Baylor, A., & Ryu, J. (2003). Does the presence of image and animation enhance pedagogical agent persona? *Journal of Educational Computing Research*, *28*(4), 373-395.

Begg, I., & White, P. (1985). Encoding specificity in interpersonal communication. *Canadian Journal of Psychology, 39*(1), 70-87.

Berlyne, D. E. (1960). *Conflict, arousal and curiosity*. New York: McGraw-Hill.

Bloom, B. S. (1984). The 2-sigma problem: The search for methods of group instruction as effective as one-to-one tutoring. *Educational Researcher, 13*(6), 4-16.

Bonar, J., & Cunningham, R. (1988). Bridge: An intelligent tutor for thinking about programming. In J. Self (Ed.), *Artificial intelligence and human learning*. London: Chapman and Hall Computing.

Bower, G. H. (1981). Mood and memory. *American Psychologist, 36*(2), 129-148.

Bower, G. H. (1987). Commentary on mood and memory. *Behavior Research Therapy, 25*(6), 443-455.

Braben, D. (2002). Another five years from now: Future technologies. Game Developers Conference, San Jose, CA, March 21-23, 2002.

Bransford, J. D., Brown, A. L., & Cocking, R. R. (1999). *How people learn: Brain, mind, experience, and school*. Washington, D.C.: National Academy Press.

Bransford, J. D. et al. (1990). Anchored instruction: Why we need it and how technology can help. In D. Nix & R. Spiro (Eds.), *Cognition, education and multimedia*. Hillsdale, NJ: Erlbaum Associates.

Brown, J. S., Collins, A., & Duguid, P. (1989). Situated cognition and the culture of learning. *Educational Researcher, 18*, 32-42.

Bruner, J. S. (1960). *The process of education*. Cambridge, MA: Harvard University Press.

Choi, J. I. (1995). The effects of contextualization and complexity of situation on mathematics problem-solving and attitudes. (Doctoral dissertation, Florida State University, 1995). *Dissertation Abstracts International, 56*(10), 3884A. (UMI Microform No. 9605031)

Choi, J. I., & Hannafin, M. (1995). Situated cognition and learning environments: Roles, structures, and implications for design. *Educational Technology Research and Development, 43*(2), 53-69.

Clark, M. S., Milberg, S., & Ross, J. (1983). Arousal cues arousal-related material in memory: Implications for understanding effects of mood on memory. *Journal of Verbal Learning and Verbal Behavior, 22*, 633-649.

Cognition and Technology Group at Vanderbilt. (1990). Anchored instruction and its relationship to situated cognition. *Educational Researcher, 10*(6), 2-10.

Cognition and Technology Group at Vanderbilt. (1993). Anchored instruction and situated cognition revisited. *Educational Technology, 33*(3), 52-70.

Cognition and Technology Group at Vanderbilt. (1996). Multimedia environments for enhancing learning in mathematics. In S. Vosniadou, E. De Corte, R. Glaser, & H. Mandl (Eds.), *International perspectives on the design of technology-supported learning environments* (pp. 285-305). Mahwah, NJ: Erlbaum.

Cohen, R., & Siegel, A. W. (1991). A context for context: Toward an analysis of context and development. In R. Cohen & A. W. Siegel (Eds.), *Context and development* (pp. 3-23). Hillsdale, NJ: Erlbaum.

Collins, J. W., & O'Brien, N. P. (Eds). (2003). *Greenwood dictionary of education*. Westport, CT: Greenwood.

Corbett, A., Anderson, J., Graesser, A., Koedinger, K., & VanLehn, K. (1999). Third generation computer tutors: Learn from or ignore human tutors? In *Proceedings of the 1999 Conference of Computer-Human Interaction*, (pp. 85-86). New York: Association of Computing Machinery.

Crawford, C. (1982). The art of computer game design. Out-of-print book. Retrieved December 20, 2005, from http://www.vancouver.wsu.edu/fac/peabody//game-book/Coverpage.html

Csikszentmihalyi, M. (1990). *Flow*: *The psychology of optimum experience*. New York: Harper Perennial.

Decker, K. A., & Ware, H. W. (2001). *Elementary teacher planning time: Teacher use; parent perception*. Paper presented at the Annual Meeting of the American Educational Research Association, Seattle, WA. April 10-14, 2001). Eric Document Reproduction Services (ED46324). Retrieved from http://www.eric.ed.gov/ERICDocs/data/ericdocs2/content_storage_01/0000000b/80/0d/d4/aa.pdf

Dillon, J. T. (1988). *Questioning and teaching: A manual of practice*. London: Croom Helm.

Duffrin, E. (no date). *Direct instruction making waves*. Retrieved December 15, 2005, from http://www.catalyst-chicago.org/arch/09-96/096main.htm#Critique

eSchool News. (2005). $10B gaming field inspires new curricula. September 30, 2005. Retrieved October 1, 2005, from http://www.eschoolnews.com/news/showStoryts.cfm?ArticleID=5896

Festinger, L. (1957). *A theory of cognitive dissonance*. Stanford, CA: Stanford University.

Foreman, J. (2004). Video game studies and the emerging instructional revolution. *Innovate Journal of Online Education, (1)*, 1.

Gagne, R. M., Wager, W. W., Golas, K. C., & Keller, J. M. (2005). *Principles of instructional design (5th ed.).* Belmont, CA: Wadsworth/Thomson Learning.

Geary, G. (1991). *Electronic performance support systems*. Gery Associates.

Gee, J. P. (2003). *Video games in the classroom?* Colloquy live, The Chronicle of Higher Education, August 27, 2pm ET. Retrieved October 5, 2005, from http://chronicle.com/colloquylive/2003/08/video/

Gee, J. P. (2004). *What video games have to teach us about learning and literacy*. New York: Palgrave-MacMillan.

Gertner, A., & VanLehn, K. (2000). Andes: A Coached problem solving environment for physics. In G. Gauthier, C. Frasson, & K. VanLehn (Eds.), *Intelligent tutoring systems: 5th international conference*. Berlin: Springer.

Gildea, P. M., Miller, G. A., & Wrutenberg, C. L. (1990). Contextual enrichment by videodisc. In D. Nix & R. Spiro (Eds.), *Cognition, education, and multimedia* (pp. 1-29). Hillsdale, NJ: Erlbaum.

Graesser, A. C., & Magliano, J. P. (1991). Context and cognition. In R. Cohen & A. W. Siegel (Eds.), *Context and development* (pp. 57-76). Hillsdale, NJ: Erlbaum.

Graesser, A. C., & McMahen, C. L. (1993). Anomalous information triggers questions when adults solve quantitative problems and comprehend stories. *Journal of Educational Psychology, 85*(1), 136-151.

Graesser, A. C., & Person, N. K. (1994). Question asking during tutoring. *American Educational Research Journal, 31*, 104-137.

Graesser, A. C., Person, N., & Harter, D. (2000). Teaching tactics and dialog in *AutoTutor*. *International Journal of Artificial Intelligence in Education*.

Graesser, A. C., Person, N. K., & Magliano, J. P. (1995). Collaborative dialogue patterns in naturalistic one-on-one tutoring. *Applied Cognitive Psychology 9*(4), 495-522.

Graesser, A. C., VanLehn, K., Rose, C., Jordan, P., & Harter, D. (2001). Intelligent tutoring systems with conversational dialogue. *AI Magazine*, *22*(4), 39-50.

Graesser, A. C., Wiemer-Hastings, K., Wiemer-Hastings, P., Kreuz, R., & Tutoring Research Group. (1999). *AutoTutor*: A simulation of a human tutor. *Journal of Cognitive Systems Research*, *1*, 35-51.

Hartley, R., & Sleeman, D. H. (l973). Towards more intelligent teaching systems. *International Journal of Man-Machine Studies, 5*, 2l5-236.

Hu, X., Mathews, E., Graesser, A. C., & Susarla, S. (2002). EBOOK.EXE: A desktop authoring tool for HURAA. In M. Driscoll & T. C. Reeves (Eds.), *Proceedings of E-Learn 2002* (pp. 471-476).

Jenkins, H. (2004). Game design as narrative architecture. In N. Wardrip-Fruin & P. Harrigan (Eds.), *First person: New media as story, performance, game*. Cambridge, MA: MIT.

Johnson, S. (2005). *Everything bad is good for you*. New York: Penguin Group.

Johnson, W. L. (2004). *Motivational effects of socially intelligent pedagogical agents*. Paper presented at the annual meeting of the American Educational Research Association, San Diego, CA.

Jonassen, D. H. (2002). Integration of problem solving into instructional design. In R. A. Reiser & J. V. Dempsey (Eds.), *Trends and issues in instructional design & technology* (pp. 107-120). Upper Saddle River, NJ: Merrill Prentice Hall.

Koedinger, K. R., Anderson, J. R., Hadley, W. H., & Mark, M. A. (1997). Intelligent tutoring goes to school in the big city. *International Journal of Artificial Intelligence in Education, 8*, 30-43.

Larkin, J. H. (1989). What kind of knowledge transfers? In L. B. Resnick (Ed.), *Knowing, learning, and instruction: Essays in honor of Robert Glaser* (pp. 283-305). Hillsdale, NJ: Erlbaum.

Larochelle, M., & Bednarz, N. (1998). Constructivism and education: Beyond epistemological correctness. In M. Larochelle, N. Bednarz, & J. Garrison (Eds.), *Constructivism and education* (pp. 3-20). New York: Cambridge University.

Lave, J., & Wenger, E. (1991). *Situated learning: Legitimate peripheral participation*. Cambridge, MA: Cambridge University Press.

Lepper, M. R., Aspinwall, L. G., Mumme, D. L., & Chabay, R. W. (1990). Self-perception and social-perception processes in tutoring: Subtle social control strategies of expert tutors. In J. M. Olson & M. P. Zanna (Eds.), *Self-inference processes: The Ontario symposium (Vol. 6)*. Hillsdale, NJ: Erlbaum.

Lepper, M. R., & Chabay, R. W. (1985). Intrinsic motivation and instruction: Conflicting views on the role of motivational processes in computer-based education. *Educational Psychologist, 20*(4), 217-230.

Lester, J. C., Converse, S. A., Kahler, S. E., Barlow, S. T., Stone, B. A., & Bhoga, R. S. (1997). The Persona Effect: Affective impact of animated pedagogical agents. *Association of Computing Machinery*. Retrieved 2001, from http://www.acm.org/sigchi/chi97/proceedings/paper/j1.htm

Macías, M., & Castells, P. (2001). Authoring tool for building adaptive learning guidance systems on the Web. In *The Proceedings of the Sixth International Computer Science Conference on Active Media Technology (AMT'01)*. Hong Kong, December 18-20, 2001.

Merrill, D., Reiser, B., Ranney, M., & Trafton, J. (1992). Effective tutoring techniques: A comparison of human tutors and intelligent tutoring systems. *The Journal of Learning Sciences, 2*, 277-305.

Miller, B. (1996). EPSS: Expanding the perspective. Retrieved February 19, 2003, from http://www.pcd-innovations.com/infosite/define.htm

Moreno, R. (2004). *Agent-based methods for multimedia learning environments: What works and why?* Paper presented at the Annual Meeting of the American Educational Research Association, San Diego, CA.

Moreno, R., Mayer, R., Spires, H. A., & Lester, J. C. (2001). The case for social agency in computer-based teaching: Do students learn more deeply when they interact with animated pedagogical agents? *Cognition and Instruction*, *19*(2), 177-213.

Murray, T. (1998). Authoring knowledge-based tutors: Tools for content, instructional strategy, student model, and interface design. *Journal of the Learning Sciences, 7*, 5-64.

Murray, T., Blessing, & Ainsworth, S. (2003). *Authoring tools for advanced technology learning environments*. Norwell, MA: Kluwer.

Otero, J., & Graesser, A. C. (2001). PREG: Elements of a model of question asking. *Cognition & Instruction 19*, 143-17.

Papert, S. (1998). Does easy do it? Children, games, and learning. *Game Developer*, June, 88.

Pearce, C. (2004). Towards a game theory of game. In N. Wardrip-Fruin & P. Harrigan (Eds.), *First person: New media as story, performance, game* (pp 143-153). Cambridge: MIT.

Phillips, D. C. (1995). The good, the bad, and the ugly: The many faces of constructivism. *Educational Researcher, 24*(7), 5-12.

Prensky, M. (2000). *Digital game-based learning*. New York: McGraw-Hill.

Prensky, M. (2004). *Future predictions: Trends in virtual education. What does the future hold?* Department of Education Secretary's NCLB eLearning Summit: Increasing options through e-Learning, July 13, 2004, Orlando, FL. Retrieved October 5, 2005, from http://www.nclbtechsummits.org/summit2/presentations/5.3.Prensky.pdf

Ravenscroft, A., & Matheson, M. P. (2002). Developing and evaluating dialogue games for collaborative e-learning. *Journal of Computer Assisted Learning, 18*, 93-101.

Reeves, B., & Nass, C. (1996). *The media equation: How people treat computers, television, and new media like real people and places*. New York: Cambridge University Press.

Rickel, J., Marsella, S., Gratch, J., Hill, R., Traum, D., & Swartout, W. (2002). Toward a new generation of virtual humans for interactive experiences. *IEEE Intelligent Systems*, *17*, 32-37.

Rieber, L. P. (1996). Seriously considering play: Designing interactive learning environments based on the blending of microworlds, simulations, and games. *Educational Technology Research and Development*, *44*(2), 43-58. http://it.coe.uga.edu/~lrieber/play.html

Schofield, J. W., & Evans-Rhodes, D. (1989). Artificial intelligence in the classroom. In D. Bierman, J. Greuker, & J. Sandberg (Eds.), *Artificial intelligence and education: Synthesis and reflection* (pp. 238-243). Springfield, VA: IOS.

Sengers, P. (1999). Narrative intelligence. In K. Dautenhahn (Ed.), *Human cognition and social agent technology* (pp. 1-26). Philadelphia, PA: John Benjamins.

Sengers, P. (2004). Schizophrenia and narrative in artificial agents. In N. Wardrip-Fruin & P. Harrigan (Eds.), *First person: New media as story, performance, game* (pp. 95-116). Cambridge: MIT

Shute, V. (1985). Artificial intelligence. In T. Husen & T. Neville Postlethwaite (Eds), *The international encyclopedia of education* (pp. 333-340). Oxford, England: Pergamon.

Shute, V. J., & Psotka, J. (1996). Intelligent tutoring systems: Past, present and future. In D. Jonassen (Ed.), *Handbook of research on educational communications and technology*. New York: Simon and Schuster.

Smith, L., & Mann, S. (2002). Playing the game: A model for gameness in interactive game based learning. In *Proceedings of the 15th annual NACCQ*, Hamilton, New Zealand, July, 2002.

Smith, S. M., Glenberg, A., & Bjork, R. A. (1978). Environmental context and human memory. *Memory and Cognition, 6*(4), 342-353.

Stevens, A., & Collins, A. (1977). The goal structure of a Socratic tutor. In *Proceedings of the National ACM Conference*. New York: ACM.

Susarla, S. C., Adcock, A. B., Van Eck, R. N., & Moreno, K. N. (2003). Authoring for *AutoTutor*: Adding a new dimension to an intelligent tutoring system. In *Proceedings of the 2003 World Conference on E-Learning in Corporate, Government & Higher Education*. Phoenix, AZ, November, 2003.

Susarla, S., Adcock, A., Van Eck, R., Moreno, K., & Graesser, A. C. (2003). Development and evaluation of a lesson authoring tool for *AutoTutor*. In V. Aleven, U. Hoppe, J. Kay, R. Mizoguchi, H. Pain, F. Verdejo, & K. Yacef (Eds.), *AIED2003 Supplemental Proceedings* (pp. 378-387). Sydney, Australia: University of Sydney School of Information Technologies.

Sutton-Smith, B. (1997). *The ambiguity of play*. Cambridge, MA: Harvard University Press.

Toole, J., & Heift, T. (2002). The tutor assistant: An authoring system for a Web-based intelligent language tutor. *Computer Assisted Language Learning, 15*(4), 373-386.

VanLehn, K. (1996). Conceptual and metalearning during coached problem solving. In C. Frasson, G. Gauthier, & A. Lesgold (Eds.), *Proceedings of the Third Intelligent Tutoring Systems Conference* (pp. 29-47). Berlin: Springer-Verlag.

Vygotsky, L. S. (1962). *Thought and language*. Cambridge, MA: The M. I. T. Press.

Vygotsky, L. S. (1978). *Mind in society*: *The development of higher psychological processes*. (M. Cole, V. John-Steiner, S. Scribner, & E. Souberman, Eds.). Cambridge, MA: Harvard University Press.

Zimmerman, E. (2004). Narrative, interactivity, play, and games: Four naughty concepts in need of discipline. In N. Wardrip-Fruin & P. Harrigan (Eds.), First person: New media as story, performance, game (pp. 154-164). Cambridge, MA: MIT.