COMPUTERS, GRAPHICS, & LEARNING

LLOYD P. RIEBER

Computers, Graphics, & Learning

Lloyd P. Rieber
The University of Georgia — Athens

Lloyd Rieber
6114 Nowhere Road
Hull, Georgia 30646
USA

CONTENTS

LIST OF BOXES

Preface

As the title indicates, this book is about computers, graphics, and learning, as opposed to computer graphics for learning. There is a difference. This book considers and integrates a broad spectrum of information related to the instructional design of visual information for learning and how the computer supports this process. Another way to understand the distinction is to first consider the importance of the three topics independently and then how they relate to each other. The title lists the topics in order from least to most important, so we must start with the last topic — learning — and work our way forward. This is also the order that must be considered when making design decisions involving visualization techniques.

The learning process takes center stage and clearly dominates the other two topics throughout this book. Although the learning process is fascinating in and of itself, this book also guides and directs the construction of environments that nurture and enhance learning, often referred to simply as instruction. This book is written for the professionals who design and develop these environments in both formal and informal settings. These individuals are usually referred to as instructional designers and/or instructional developers. Many carry this title as the formal result of graduate-level training; others find such a role thrust upon them, perhaps unexpectedly. For this reason, this book is relevant to anyone concerned with or involved in designing graphics for instruction.

This book has a more specific mission beyond general instructional design: to exploit the potential of visualization techniques to enhance and improve learning. Graphics long have been a common part of all instructional strategies. Many of the most valuable principles of how visuals can help learning have been identified apart from computer applications. Therefore, designers have much to gain from applying the general theory and research related to visuals, memory, and learning to instructional design. Considering these knowledge bases becomes even more important when one understands that all graphics are not appropriate for all learning outcomes. Indeed, inappropriate uses of graphics can actually thwart well-intentioned instructional design. We resist the tendency to believe that efforts to apply computers and graphics to learning “break all the rules” of the available theory and research (even if it turns out later to be true in some places). Despite the attention to theory and research, we are careful to remember that our overriding goal throughout this book is application — to apply what we know about visualization and learning to instructional design and development.

We finally come to the role of the computer. There is no question that the computer offers unprecedented graphical power for all designers, instructional and otherwise. The range, power, and number of graphical tools for desktop computers are increasing at an astonishing rate. Some of these tools are meant to increase the productivity of traditional print-based materials. Some, such as animation packages, increase the productivity of traditional video-based materials. Other tools, such as those that provide learners with “real-time, on-line” interaction, offer potentially new learning environments that would not be possible without computer technology. All this often creates a sense of urgency among designers and developers to know and incorporate the latest graphical tools in their courseware. However, we need to continually remind ourselves that “a power saw does not a carpenter make.” There is a need to exploit the graphical power of computers for learning but not fall prey to the idea that using the latest technology is a substitute for good design. The danger of confusing good design with the mere use of the “latest and greatest” technology is particularly potent in the computer arena. The computer is considered here as but an arsenal, albeit an important and powerful one, of resources to facilitate learning by and through visualization techniques. As instructional designers, we need to stay in control of our “tools of the trade.”

Finally, it is important to recognize that this book is not intended to teach you how to use a computer or to create computer graphics. Some attention is given to the development of computer graphics, but only to serve as an organizer to help you understand the range of desktop computer graphics applications. It is hoped the principles of this book remain relevant and useful as the computer industry improves and expands desktop computer graphics technology and, perhaps more importantly, as your ability and knowledge of how to develop computer graphics for instructional purposes grows. Although design and development are interdependent processes, our concern and attention is first and foremost on design. No formal training or background in psychology, instructional design, or computer graphics is considered prerequisite to reading this book, as all topics are written at an introductory level. However, this book is intended for graduate-level students.

OVERVIEW, SCOPE, AND ORGANIZATION OF THIS BOOK

Chapters 1, 2, and 3 provide a broad overview of instructional computer graphics. Chapter 1 provides a general introduction and describes the rationale and philosophy on which the book is based. Chapter 2 describes the three most common types of graphics in instruction and the range of learning outcomes in which these graphics can be applied. It then presents a brief overview of the most common instructional applications of graphics. Chapter 3 discusses the development of computer graphics. The purpose of this chapter is to compare and contrast general production procedures and techniques. However, this chapter is not meant to provide an exhaustive summary of the “how to's” of producing computer graphics. Production techniques in both “command-based” versus “GUI-based” graphics and authoring applications are presented.

Chapters 4, 5, and 6 present an overview of the status of instructional visual research. Chapter 4 provides an introduction to psychological foundations that sometimes support using graphics in instruction and other times warn against it. Included in this chapter are discussions on visual perception, visual cognition, and theories on storing visual information in short-term and long-term memory. This chapter also briefly describes some of the implications of learning theory on instructional graphic design. Chapter 5 summarizes the large pool of research dealing with static graphics, and chapter 6 summarizes the relatively scant pool dealing with animated graphics. Many of the research studies discussed in chapter 6 were conducted by the author, and so the discussions are presented firsthand.

Attention is turned to designing graphics in CBI in chapters 7 and 8. Chapter 7 summarizes the major aspects of CBI design in the context of when and how static and animated graphics should be integrated within CBI lessons. Chapter 8 deals specifically with the design of highly interactive visually based lesson activities, such as computer-based simulations. This chapter uses the microworld paradigm that has been suggested by the constructivist perspective on learning. Also discussed in this chapter is the concept of “virtual reality.”

Finally, chapter 9 considers other sources of visuals, such as video, by presenting a brief overview of multimedia, which consists of integrated learning systems that join computers and peripherals — such as videodisc and videotape players. These systems let learners experience a full range of sensory stimulation, including sound. Specific video applications of multimedia are better known as interactive video.

Design takes center stage in this book. The scant coverage of production does not mean it is unimportant, but rather underscores the perspective that design must drive production (although design is and should be influenced by production capabilities). This text also does not consider the design of visuals that are associated strictly with printed text, at least not as separate topics. Examples of issues not covered include how to select text type, text size, and text orientation (such as page justification). However, discussions of how graphics interrelate with text are relevant and are addressed here to some degree. The design of instructional text is represented well elsewhere (see, for example, Hartley, 1987; Hooper & Hannafin, 1986; Jonassen, 1982, 1984).

Finally, there is the issue of delivering instructional graphics produced by the computer, which can take one of at least two perspectives. The first and most obvious is delivery by computer, including (but not limited to) CBI applications. The other involves using the computer as the principal graphics design and development tool, then transferring the graphics to other delivery sources or “platforms” such as print-based materials, film, and video. Both delivery perspectives figure prominently in the book.

Each chapter begins with a brief overview, followed by instructional objectives for both the literal comprehension of the text and subsequent application of the principles described in the chapter. Readers are expected to complete the application objectives in a learning context that includes many other resources besides this book, such as training and guidance in instructional design and the use of relevant computer hardware and software. Together, the overview and objectives are meant as an orientation and guide. Each chapter is carefully organized and sectioned. Use the outline generated by the headings and subheadings as an additional learning guide. In addition to illustrative material, most chapters also have information boxes that contain separate and complete discussions and activities related to the main text. Chapters end with reviews which are not meant to substitute for actually reading the chapters but should help your understanding by emphasizing each chapter's main points. The reviews are also meant as a quick way to refresh your memory of the chapters without rereading.

ACKNOWLEDGMENTS

The ideas represented in this book are the product of many years of play and work in educational computing, beginning for me as a public school teacher in New Mexico. I thank all the students I have worked with so far — graduate students, as well as those I knew as grade school students — for all they have taught me. The formal idea for this book was a result of teaching a course on instructional computer graphics in the Educational Technology Program at Texas A&M University. I am especially grateful to those students who patiently labored through and provided invaluable feedback on early drafts of this text. Many thanks go to them for their understanding as I struggled to put my ideas into written form. I especially thank Evelyn Wells for graciously allowing her work on color principles to be included in this book. I also thank Ronald Zellner and William Kealy, colleagues of mine at Texas A&M University, for sharing their expertise in this area with me. Many of the ideas represented here started as a result of "water-cooler" conversations with them on the topics of learning theory, instructional design, and visualization. Special thanks also go to Mary Boyce and her students at the University of Oklahoma for their comments on an early draft. I am very grateful to the reviewers for their excellent comments and suggestions.

I also thank the excellent professionals at Brown and Benchmark Publishers, especially Paul Tavenner and Michelle Campbell, for their support and service.

Finally and needless to say, I am grateful to my wife, Patricia, and my children, Rebecca and Thomas, for their understanding and sacrifices. This book would not have been possible without their support. I owe them a large debt, of which Pat has been careful to keep an accurate record!

CHAPTER 1

Introduction

OVERVIEW

This chapter presents the rationale and philosophy of this book. Its general premise is that there is a strong need to guide instructional graphic design efforts because of the important role of visual communication and the increasing availability and ease of computer graphic tools. The chapter also provides examples of the power of visual communication in and out of education. Desktop computers are moving toward the use of graphical user interfaces (GUIs), a trend that appears to be gaining momentum. GUIs, in combination with systems that effectively blend and encourage graphics throughout software applications, will likely spur the use of graphics in educational courseware. Ironically, because of these and other forces, instructional designers are liable to lose sight of their original goals for graphics. For this reason, designers are encouraged to reflect on what motivates their decisions to incorporate graphics in instructional materials. Definitions of important terms are also presented.

OBJECTIVES

Comprehension

After reading this chapter, you should be able to:

1. Discuss the role of visuals in communication.
2. Define GUI and summarize ways it may influence instructional design.
3. List the three approaches to instructional design and describe their strengths and weaknesses.
4. Define what is meant by technocentric design.
5. Describe two or more examples of ways visualization can aid human problem solving.

Application

After reading this chapter, you should be able to:

1. List motives for using graphics in instruction.
2. Classify these motives as stemming from either instructional design or technocentric design.

PURPOSE OF THIS BOOK

This book is about the design of graphics for instructional purposes in the computer age. Without question, graphics are a popular part of computer-based instruction (CBI), and many leaders in educational computing advocate their use. Even a casual review of commercially available educational software demonstrates the high frequency and intensity of graphics. Unfortunately, graphics are often used to impress rather than to teach. Sound and graphics always have been seductive features of new technologies, but these (and other examples of educational “glitter”) are often considered — usually erroneously — as indications of effective design. Unfortunately, many educators continue to evaluate the design of instructional materials on the computer based on the machine's pyrotechnics rather than on a synthesis of the learning goals, demands of the task, and the needs of the learner. This may be because most educators are still unfamiliar with computer technology (Rieber & Welliver, 1989) and are easily impressed.

Regardless of their effectiveness, visuals (of which graphics are a subset, as described in the next section) remain a staple in most instructional strategies, and, therein lies the problem. Attempts at applying visuals in instructional materials are usually haphazard; sensitivity about what does and does not work comes slowly over time. Although we arguably do not know enough about how to use visuals effectively to communicate or to promote learning, one thing is certain: We need to better apply what we do know. This is even more important given the power of the budding relationship between computers and visuals. Whether that power will be used appropriately in education remains to be seen. But many people are dedicated to the idea that harnessing this power represents an important advance in human communication in general and learning in particular.

This book has been written for a wide range of people who share a need to learn more about designing static and animated graphics with microcomputers for instructional or training situations. People new to computers and graphics or new to instructional design should find this book helpful. Thus, this book is appropriate for educators needing information about computer graphic applications, or computer specialists and artists who find themselves designing computer graphics for training situations. As a result, many issues and topics dealing with computers, graphics, and learning have been necessarily distilled to present only the most essential and pertinent information. Additional references are provided to help direct and guide further studies in each of these separate directions.

Some Definitions

The meanings of the terms **visual**, **graphic**, **image**, and **picture** greatly overlap and are often used synonymously. Strictly speaking, computer visuals refer to all possible computer output, including text. Instructional computer graphics are considered a subset of computer visuals and involves the display of **nonverbal information**, or information that is conveyed spatially. Included in this definition are the range of computer-generated pictures, with **pictures** being defined as graphics that share some physical resemblance to an actual person, place, or thing. The quality of these types of graphics ranges from near-photographic to crude line drawings. Also included is the spectrum of nonrepresentational graphics, including, but not limited to, charts, diagrams, and schematics.

Besides its general meaning, the term **visualization** also describes an interdisciplinary field of study in which computer graphics techniques are used to display images that convey a wide range of information. In this sense, visualization differs from computer graphics in that visualization stresses the information that is conveyed in the resulting image (Brown & Cunningham, 1990). However, for simplicity's sake, **graphics** typically will be used throughout this book to denote all visual information conveyed through nontextual ways.

The terms **computer-assisted instruction** (CAI) and **computer-based instruction** (CBI) are also often used synonymously. However, there are distinctions between these terms. CBI usually refers to instructional systems that are completely computer-based. Instructional delivery, testing, remediation, etc., are all presented and managed by computer. On the other hand, CAI usually refers to supplemental or adjunct uses of the computer to support a larger instructional system, such as a traditional classroom (Hannafin & Peck, 1988). CAI includes traditional, structured (deductive) approaches such as those associated with tutorial and drill-and-practice software, but also the informal, discovery (inductive) approaches associated with computer games and simulations. This book, again for simplicity's sake, generally uses the term CBI to include all instructional applications of computers.

The First Principle of Instructional Graphics

The first principle of the design of instructional graphics is this: There are times when pictures can aid learning, times when pictures do not aid learning but do no harm, and times when pictures do not aid learning and are distracting.

This principle is an important one, however obvious it may seem. It speaks to an underlying philosophy of instructional graphics and instructional technology. It is important to understand this philosophy at the start, because it will guide every attempt at designing graphics when the purpose is to aid, enhance, or support learning. This book does not advocate or oppose the use of graphics. Instead, it supports the position that graphics, like most strategies and techniques, have their place in instruction. The problem is understanding when and how to design effective graphics, as well as when to avoid them altogether.

Similarly, it is important to note that this book does not advocate the misguided idea that computers can solve all educational problems. Rather, the book presents the “computer-as-tool” perspective, which looks for ways to capitalize on the strengths of computer technology for graphic designs as they support learning and instructional goals. Only the *instruction* that is delivered by or through the computer has the potential to influence learning, not the *computer* itself. A piano can be the medium by which the works of Mozart are brought to life, as well as the medium upon which a 3-year old pounds away. A delivery truck can bring a home fresh milk, eggs, or bread, as well as sweets and soda. Similarly, a computer can deliver instructional noises or inspirations, junk food or a well-balanced meal. Again, this concept may seem obvious, but educators have been prone to the misconception that the newest innovation will inspire achievement and academic success just by having students come in physical contact with it (Clark, 1983; Kozma, 1991; Simonsen, Clark, Kulik, Tennyson, & Winn, 1987). Even today, schools proudly report to their PTAs the number of computers they own but rarely explain to what uses the computers are being assigned. But then again, parents and school boards rarely ask. Asking how many computers are in a school is just like asking how many pencil sharpeners, chalkboards, or desks there are. These latter questions seem ridiculous because we assume that schools will have enough of these tools and resources to meet the educational needs of students. Why should it be any different with computers? Maybe it is because we are still not sure what needs computers serve in education, although we seem convinced that schools should have them.

THE IMPORTANCE OF VISUAL COMMUNICATION

Our sense of vision represents our richest source of information of the world (Sekuler & Blake, 1985). The partial or complete loss of sight is one of the most difficult impairments to overcome. Enormous amounts of information are transmitted visually. (see Footnote 1) Consider the sources of the information you depend on each day. As a student, consider your class materials, your notes, and the strategies you use to study. As a professional, consider how ideas are expressed and conveyed within your professional circles. In your personal life, consider how various media influence you and what part visuals play. Think about how you use (or are used by) the visuals in television, magazines, newspapers, and product catalogs. In virtually every case, visuals of some sort and variety are the main vehicle of expression and communication. Consider how influential visuals such as facial gestures and other body movements (usually referred to as **nonverbal communication**) are in face- to-face conversations and social interactions. As you become better at evaluating the ways in which visuals inform and influence people, you will not only begin to understand the uses, abuses, and misuses of visuals in communication, but you will also appreciate more the human visual processing system.

There are, in fact, few examples in which visuals of some sort do not play a role in daily communication. Telephones and radio are probably the most notable exceptions. Yet, even in these cases our minds make up visually where these media leave off. The motto for radio advertising — “Say you saw it on the radio!” — sums up our ability to conjure up images with even the slightest prompting. (See Footnote 2.) Though the root of this ability is considered innate, it is nurtured as we mature. Research indicates that adults are better than children at mental imaging and are also more likely to spontaneously form internal images (Pressley, 1977).

Visual skills are particularly important in many problem-solving situations. One such skill is the ability to quickly see patterns in information represented visually. A classic example of how pattern recognition led to an important discovery occurred during the cholera epidemic in London in the mid-1800s. Dr. John Snow plotted the location of each cholera death and each water pump in a central London neighborhood, as shown in Figure 1.1. The obvious clustering of deaths around the Broad Street pump provided strong evidence that there was a relationship between the cholera deaths and drinking water from this particular well (Tufte, 1983). This evidence was further strengthened when Snow visited the families of 10 other cholera victims who did not live near the Broad Street pump and discovered that five of these regularly sent for this pump's water because they preferred the taste, and three more of these families had children who attended schools near this pump. Even though at the time it was not proven that a contaminated water supply was the cause of the epidemic, Snow's map convinced authorities to remove the handle from the Broad Street pump, and within days the epidemic in this neighborhood ended (Wainer, 1992).

FIGURE 1.1
The famous dot map of Dr. John Snow plotting the cholera deaths in London in relation to neighborhood water pumps showing convincing evidence that the Broad Street pump was contaminated.

Another interesting example involved efforts in World War II to better armor combat aircraft. One strategy was to plot all the bullet holes on aircraft returning from combat on a crude picture and then add extra armor to other planes *everywhere else*. The logic was simply that since all the planes probably had been hit uniformly during battle, those that did not return must have been hit in the vital places *not* marked on the picture (Wainer, 1992).

Given this attention to visual modes of communication, it should not be inferred that other channels are unimportant or should be ignored. In social interactions, speaking and listening are the dominant methods of communication. In computer-enhanced training situations, the expense (in terms of hardware and/or memory) of voice integration — including both speech output and voice recognition — has prevented effective integration of aural communication channels in existing software and computer systems. The status of computer speech and voice integration changes almost daily and inevitably will change current notions about how people and computers should interact (see Box 1.1). Even when the language capabilities of computers become as natural as everyday human conversation, visual channels will flourish and remain a dominant influence in the presentation and interaction of ideas in computer environments.

Graphics in Education

The use of graphics in education has a long history. The use of illustrations in books written in English, especially those intended for children, was commonplace by about 1840 (Slythe, 1970). After that time, the use of illustrations in children's books has been especially extensive, elaborate, and artistic (Feaver, 1977). A wide variety of graphics — from photographs, pictures, and cartoons, to charts, maps, diagrams, and outlines — is common today in most teaching strategies. The use of graphics in instruction seems to make sense — it holds a certain degree of face validity. The cliché that a picture is worth a thousand words seems consistent with educational practice. However, research has shown that the relationship between the intent and results of graphics in education is often jumbled (Samuels, 1970).

There is a tendency to use armchair methods of deciding when, where, and how to incorporate graphics in instructional and training strategies and materials. This can lead to unexpected results. Research is just beginning to demonstrate conditions under which static and animated graphics are generally effective, as well as those where graphics serve no purpose or, worse, do harm (Levin, Anglin, & Carney, 1987; Levin & Lesgold, 1978; Rieber, 1990a). For example, consider the cultural symbolism of the owl. Most people from western cultures treat this wise old creature of the forest with affection, although an owl often represents an evil omen for many Native Americans. Classroom teachers should carefully consider the impact of such innocent graphics on all their students. Like most issues in education, graphics represents a qualitative, not quantitative, issue. It is not simply a question of how many graphics are used that determines their effectiveness. The interaction between instruction and learning is complex and does not lend itself to many generalizations. Answers to questions about how best to employ instructional graphics are similarly elusive and evasive.

Instructionally, the role of graphics in computer environments covers a lot of ground. The computer can be used for traditional applications, such as graphics that present static informational images or text that helps someone to understand a concept or principle. Much of the instructional visual research over the past 40 years has pertained to applications such as these. Although most of this research has been in noncomputer contexts, it is still quite relevant. However, the computer offers many more instructional applications than just presentation. One of the most exciting, yet uncharted, areas involves computer microworlds based on computer animation (Rieber, 1992).

Box 1.1

The Stuff Dreams Are Made Of

Probably the best conceptions of future human/computer interaction come from science fiction. The various Star Trek books, television shows, and movies are among the better examples known to people who are not necessarily "trekkies" and who may not be general lovers of science fiction. In the Star Trek stories, the computer becomes almost completely transparent and completely networked throughout the starship *Enterprise* and *Star Fleet*.

The more recent Star Trek television series titled *Star Trek: The Next Generation* has produced some very tantalizing images of how computer technology can help in human problem solving. In this series, the current Enterprise has many improvements and enhancements. Among the most notable is the ability of the computer to produce true holographic, or completely three-dimensional, images.

In one episode, Captain Picard is sitting at his desk contemplating a strange planetary system recently visited by the Enterprise. The orbit of one of the planets had a strange "wobble" in it and was unlike anything yet encountered. The captain is shown studying the phenomenon with the help of a three-dimensional holographic model floating peacefully above his desk. The scene shows the captain just sitting at his desk in his "ready room" studying the image with an understandably puzzled look on his face as his first officer, Commander Riker, enters. The importance of representing the planetary system in a three-dimensional model becomes even more evident as the captain gets up and discusses the model with Riker from a different perspective in the room. The 3-D model makes it possible to walk around and see it from every side. It clearly could not be understood very well with text or even with two dimensional graphics (although, of course, that's exactly how it was represented to the television viewer). Only a true three-dimensional "hologram" was able to convey the information well enough for the captain's purposes. The other message given to the television viewer at the end of this scene was that the captain was apparently able to quickly and easily "program" the computer to make this model for him.

However, the most amazing application of computer-generated holograms on this new Enterprise is the "holodec" where crew members can go and experience anything from computer-generated California beaches, medieval forests, and snow-covered mountains to smoke-filled nightclub rooms in the past, present, or future. In the holodec the computer can create any three-dimensional environment imaginable, complete with casts of characters. These environments cannot be distinguished from their real-life counterparts. Science fiction to be sure, but this is exactly the dream and goal of developers of cyberspace, also known as "virtual realities."

Everyday Uses of Graphics and Visual Images

Some of the most stunning examples of graphics that communicate come from outside of education. The popular media — such as television, newspapers, and magazines — have long abandoned any real restraint when it comes to using visuals. True, most visuals are used primarily to capture the viewer's or reader's attention for just a few precious seconds. Often, though, the visuals are intended to enhance a memory function by influencing people to remember one product over all others. Consider the many variations of the popular beer commercial that all end in a frustrated “I meant a Bud light!” Pictures go through our minds of all the wrong “lights” that the hapless people seem to uncover. Because they are novel and amusing, the pictures are easily remembered. We automatically associate the pictures with the product name because we have been subjected to countless rehearsals of the two. (See Footnote 3) So, we are likely to remember this one company's product first if we are out shopping for beer. The success of this commercial is a prime example of using pictures as a powerful mnemonic device — the images and the product name are forever associated or cemented together. You couldn't forget them if you tried. (It also demonstrates the power of applying some simple behavioral principles.)

Some of the best examples of using full-motion video to demonstrate procedural knowledge come from television toy commercials. Advertisers not only must capture the attention and interest of children (no small feat), but show them how to have fun with the toy, albeit in exaggerated and contrived ways. At their best, these commercials unravel the complex nature of a toy in as few as 15 seconds. Particularly good examples are all the varieties of commercials that tout toy robots that transform into cars, planes, and tanks. These commercials demonstrate a tremendous amount of information in a very short amount of time. Although this does not suggest that educators should become advertisers, there is still a great deal to learn from the techniques that successful advertisers use to visually communicate their ideas.

WHY COMPUTER GRAPHICS?

Computer learning environments pose particularly exciting and demanding situations for visual communication. The range and diversity of visualization that computers offer are unprecedented. The last 10 years have demonstrated marked increases in sophistication in the graphics produced and displayed on computers. The success of desktop microcomputer systems integrating graphical user interfaces (GUIs), such as the Macintosh computer, can be largely attributed to the dramatic rethinking of how people should interact with computers. The principal reason to highlight the computer in the design and development of instructional graphics is the computer's increasing range, versatility, and flexibility of graphic design. There is almost no graphic design need that the computer cannot serve. In addition, the design of computer graphics is no longer limited to delivery on computer platforms. The unprecedented spread of desktop publishing is a prime example of the computer as a design and production tool, though the delivery platform is paper.

Many believe the Macintosh computer survived and flourished (unlike its predecessor, the Lisa) because it carved its niche in desktop publishing. (Some suggest it invented it.)

Several factors contributed to this. The most important was that the Macintosh effectively combined text and graphics — the Macintosh was the first popular microcomputer to truly adopt (but not invent a GUI. (See Footnote 4) The marriage of text and graphics was inherent in the Macintosh's operating system as well as in its application software. Up to that time, computer text and graphics existed separately, with great effort needed to merge the two. The other major factor was the advent of laser printers. Fast, camera-ready quality printing, albeit in black and white, suddenly became available at the click of a button. Not only could professionals bypass the time-consuming and expensive step of having their materials professionally typeset, but they could now experiment with alternative designs quickly and easily. In this way, design and development became one process. Ideas could be laid out on the electronic and printed page in a "what you see is what you get" (WYSIWYG) format. Compared with conventional typeset quality publishing methods, the decreased turnaround time for the feedback/revision cycle using computer desktop publishing was staggering. GUI concepts have extended into the area of desktop presentations, where organized presentations are designed and developed by computer, then transferred to delivery platforms such as overhead transparencies, slides, or video.

Both desktop publishing and desktop presentations made the Macintosh computer a business success and offer similar potential to "Mac-like" products (such as Microsoft Windows) because there was an eager market for these innovations. The business world needs to communicate ideas and strategies to clients and consumers in effective and influential ways. Those in the business of education have communication needs that reach much further.

At times, the computer's role looms larger than just economy, efficiency, versatility, or flexibility of production or time. Computer visualization has become an important problem-solving tool for people. Probably the best example of this collaboration between people and computers is the relatively new science of Chaos (see Box 1.2). The computational power of computers has permitted people access to ideas that previously were off-limits because of the tremendous calculation demands. Instructionally speaking, the number of creative strategies and applications that simply would not be possible or practical without computer technology is increasing. Most cases in point involve computer animation.

Certainly, the use and history of animation in film easily predates modern computer technology. Many media, such as film and video, can present animated sequences. However, the computer is rapidly becoming an important tool in modern animation studios, including Disney. The field of visualization uses computer animation as a central tool for studying problems and issues in architecture, medicine, and fine arts. In each of these cases, the computer is used as an important production tool for creating animation sequences that are normally transferred to film or video for delivery. Rarely do these areas need or use the computer for delivery. In many cases, due to the complexity of the production, the computer might take up to 30 minutes to create each individual frame of the animation, shutting out the possibility of real-time animation. However, many tasks that can be animated in real-time present exciting possibilities in education.

Box 1.2

Play the Chaos Game

Computers are often viewed as the epitome of a mechanized, inhuman world. People often express fear that computers will create an "Orwellian" world in which personal freedom and expression will be eliminated. However, a contrasting view perceives computers as our liberator by performing the many tedious tasks heretofore done by humans, thereby allowing people to more fully realize their growth potential. This idea is simply to let machines and people do what they do best. Nowhere is this more striking than in the science of Chaos (Gleick, 1987).

Chaos, as the name implies, is a new science that explores events that are seemingly erratic, random, and haphazard. However, this science has begun to understand that many phenomena that seem to be random events on the surface may actually have a hidden order lurking below. Perhaps the most intriguing aspect of the science of Chaos is that these phenomena do not simply occur at the molecular or atomic level, but at the everyday level as well. Examples of chaotic phenomena are the turbulent patterns in mountain streams, ribbons of smoke from a cigarette, flags waving in the breeze, and the dripping of a leaky faucet.

More accurately, Chaos is a study of nonlinear systems. We are all familiar with linear systems where one variable changes predictably to changes of another variable, such as the relationship between acceleration and velocity. However, nonlinear systems are far more complex. Solving problems in nonlinear systems has usually been tried with linear models. Since the fifteenth or sixteenth century the goal of science has been to fully understand the laws that govern the universe. It was believed that the problem was not that we were incapable of understanding, just that we were not able to collect enough information to complete our mathematical models. This was the idea behind Newton's "clockwork" universe.

A good example of this is weather forecasting. The hope has always been that if enough data were collected at enough locations worldwide, we would be able to predict the weather with reasonable accuracy days, weeks, or even months into the future. Obviously, one can not collect all the data available, so as much information at as many geographic points as possible is sampled. The goal has been, for example, to create an elaborate linear model involving hundreds of components such as the noon temperature in Paris, the monthly precipitation of Florida, the humidity of Moscow, etc. The hope was that if enough information was collected, we might finally be able to accurately predict the weather tomorrow in New York City. Unfortunately, weather patterns are nonlinear systems, and Ed Lorenz, a research meteorologist at the Massachusetts Institute of Technology, discovered that the use of linear models to help solve problems in nonlinear systems is little more than wishful thinking.

In 1960, Lorenz constructed a model of a weather system on his computer. Of course, the measurements he entered into his computer could never be perfect but only approximations. He discovered that the smallest errors, even those resulting from rounding to the nth decimal, would create vastly different outcomes in his "toy" weather system. This phenomenon became known affectionately as the "butterfly effect," based on the idea that even the stirring of a butterfly's wings in Peking can be the difference between a sunny or stormy day a month later in New York. The notion of long-range weather forecasting based on linear models was doomed. The general principle learned was that even the smallest differences in input would cause tremendous differences in output when a linear model was used to study a nonlinear problem.

So what is the connection to computer graphics? Many of the hidden patterns and elements in the raw data of nonlinear systems start to become evident when they are represented graphically. Because people are generally good at pattern recognition, researchers in Chaos theory frequently have the computer construct pictures from the data. A good example of this is fractal geometry, where a particular shape is repeated infinitely within itself. The mathematician Benoit Mandelbrot is generally given credit as the originator of formal fractal geometry.

A clever way of experiencing this partnership between computer visualization and human problem solving is called the "Chaos Game," devised by Michael Barnsley, a mathematician at the Georgia Institute of Technology. There are many varieties of the Chaos Game, but the following example draws a figure commonly known as the Sierpinski Gasket.

To play the game, you need the following materials: (a) a sheet of paper; (b) a pencil; (c) a ruler; and (d) a game die. Here are the game rules:

Draw three dots on the paper (such as those that form an equilateral triangle). We will call these dots our GAME POINTS.
Mark the first GAME POINT "1,2"; the second GAME POINT "3,4"; and the third GAME POINT "5,6."
Pick another point at random on the sheet of paper. We will call this the STARTING POINT.
Throw the game die. The resulting number identifies, at random, one of the three GAME POINTS in step 2.
Draw a dot at mid-point of the STARTING POINT and the randomly chosen GAME POINT. This mid-point dot becomes your new STARTING POINT.
Repeat steps 4 and 5 for thousands of trials.

What would you expect to get when you are through? Most unsuspecting people expect a random arrangement of dots that either fill the paper or the area within the three game points. Instead, what is produced is a chaotic system, in the sense that what appears to be random and unorganized at first glance can contain an amazingly complex and ordered pattern. Of course, most people are not willing to invest the time or energy to play the game carefully for the thousands of trials necessary to see the result. So, let's let the computer do the part we hate — number crunching — and we will do the part we are good at — interpretation. The following program, written in HyperCard's HyperTalk on the Macintosh, plays the game for us for as long as we care to let it run:

```
on mouseUp
  global gp,gx,gy,bx,by,x,y
  —set up game board
  choose brush tool
  set brush to 7
  click at 50,300
  click at 256,75
  click at 432,300

  —choose first “starting point” at random
  set brush to 28
  put the random of 512 into x
  put the random of 342 into y
  click at x,y
  put x into gx
  put y into gy

  —play the game
  Repeat until the mouseclick
    —choose one of the three “game points” at random
    put the random of 3 into gp
    if gp = 1 then
      put 50 into bx
      put 300 into by
    end if
    if gp = 2 then
      put 256 into bx
      put 75 into by
    end if
    if gp = 3 then
      put 432 into bx
      put 300 into by
    end if
    —draw the midpoint of the “starting point” and “game point”
    put (bx+gx)/2 into x
    put (by+gy)/2 into y
    put round (x) into x
    put round (y) into y
    click at x,y
    —make the midpoint the new “start point” and repeat
    put x into gx
    put y into gy
  end repeat

  choose browse tool
end mouseUp
```

The program produces something quite unexpected. Turn to Figure 1.5 to see the result.

Real-time animation occurs when the computer is able to display graphic frames in a quick enough succession to produce the illusion of motion. Real-time animation permits computer applications such as video games and simulations. We are all familiar with the idea of a flight simulator, where the screen display changes depending on whether we are pushing the plane's control stick forward to dive toward the ground or pulling back to climb higher into the sky (either by pressing certain keys or moving a joystick or mouse). Simulations, both of real and imaginary things (like most dungeons-and-dragons-style video games), represent microworlds of realities or fantasies where a user goes to experience something firsthand. For example, real-time computer animation can make a simulated journey into space become a real experience. Here students might learn for themselves what it would feel like not to be bounded by gravity or friction.

By combining technologies, the illusion of leaving the real world and stepping into another computer-generated one can be multiplied many times over. Virtual realities can be created by fitting a person with headsets mounted with video screens. The individual also wears a special data glove that sends commands (like forward, back, and stop) to a computer via hand signals. This results in a convincing illusion, such as moving about on an uncharted planet or strolling through the lobby of a proposed skyscraper. Even though current technology can only produce crude graphics for these imaginary trips, the illusion remains strong because the headset does not allow any competing visual stimuli to enter the person's field of vision. (This is discussed further in chapter 8.)

Advancements in the Production of Computer Graphics

Computer hardware and software manufacturers appear dedicated to the idea that computer displays and resulting human interactions should be graphically based. The ability to produce and integrate graphics into instructional courseware should become easier as systems become more visually oriented. This, of course, poses the ironic problem of what to do with this graphical computer power. Instructional designers face tough challenges. Too many choices create the temptation to design instruction with as many features special to the computer as possible, often in the attempt to justify the use of the computer as a design and delivery system. It seems that this temptation is becoming harder to resist as systems that afford even greater graphical power at reduced development costs become available.

The trend of GUIs in many microcomputers has led to increased availability of graphics for CBI authors. Today, the compatibility of object-oriented software applications on GUI systems permits the almost-casual availability of graphics in CBI development. Textual and graphic objects can be shared among applications via concepts such as cutting and pasting, clipboards, and scrapbooks. Users can quickly create firsthand graphics and then import or paste these into CBI lessons. Commercially produced graphics — electronic clip art — provide designers with a ready supply of second-hand graphics. Beyond all these are video digitizers and optical scanners that permit real objects and print-based graphics, respectively, to be digitized and imported into courseware (copyrights notwithstanding). The quality of scanners continues to rise as the cost continues to fall. Although scanning and video capture technologies are still far from perfect, they are sure to be, by far, among the most flexible and cost-effective approaches to graphic design.

```
5000 HOME
5010 REM *********************
5015 PRINT CHR$(20)"3":REM SET TYPEFACE TO SYMBOL
5020 WL=2:WR=37:WT=2:WB=22: GOSUB 900: REM SET WINDOW
5030 GOSUB 500: REM DRAW SCREEN FRAME
5040 VTAB 6:HTAB 12: PRINT "SPACE SHUTTLE":VTAB 8:HTAB
      14:PRINT "CD"
5050 VTAB 12:HTAB 11:PRINT "FLIGHT LESSON #1"
5060 SB=3:GOSUB 600:REM SPACE BAR
5070 REM **********************PAGE 1
5080 HOME
5090 PRINT CHR$(20)"1":REM SET TYPEFACE TO STANDARD
6000 A$="The purpose of this mini-course is to introduce you
      to how and why objects move.":gosub 1000
```

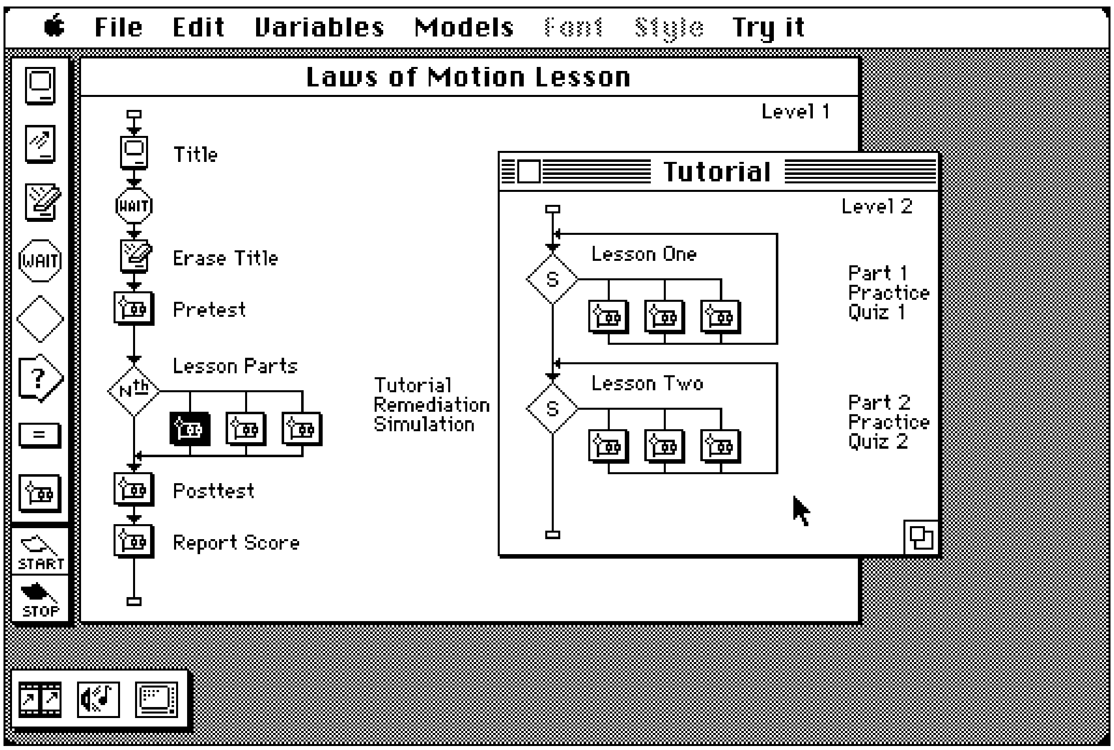

FIGURE 1.2
CBI authoring approaches are shifting from traditional command line environments (above) to ones that employ graphical elements (below).

More and more authoring approaches are also using graphics to aid the design process itself. Figure 1.2 compares a traditional “command” authoring environment with an object-oriented one. (See Footnote 5) Lessons are designed and developed by manipulating a variety of icons, each representing a particular instructional function, on flowchart-like displays.

Advances in computer animation production are particularly poignant examples of how far computer technology has come in such a short time. Traditional programming approaches to animation are tedious and labor-intensive. Real-time animation of even a modestly complex object used to require advanced programming knowledge. Improvements in producing computer animation have occurred steadily. For example, the advent of high-level programming languages, such as BASIC or PILOT, removed the need for hobbyists to learn an assembler language or machine code to create convincing animated displays. Still, nonprogrammers had to master shape tables and programming commands based on abstract mathematical coordinate systems.

GUI systems make the mathematical processing of animation almost completely transparent. Many systems offer advanced features, such as the ability to interpolate an animated path between a set of given end points. These systems also offer authoring advances in **data-driven animation**, which is the animation of objects based on constantly changing program values, such as student input. Data-driven animation is at the heart of visually based computer simulations, such as flight simulators. Animation has become a very popular feature in CBI applications, although, just like other graphics, the instructional goal that animated visuals serve is often not clearly defined.

QUESTIONING THE MOTIVE TO USE GRAPHICS IN INSTRUCTION

The use of graphics is particularly strong in training environments and situations that emphasize materials-centered instruction, of which CBI is an example. Materials-centered instructional environments depend on media other than the teacher, trainer, or workshop leader for the primary presentation of instruction. Apart from computer applications, materials-centered instruction has been a dominant influence in instruction and training since World War II, especially in the private sector and in the military (Reiser, 1987). Traditional examples of media include textual materials (books, workbooks, worksheets), videos and films, slides and filmstrips, learning centers, and overhead transparencies, to name just a few. The field of educational technology is usually associated with its contributions in refining design techniques used to produce materials-centered instruction. This field has witnessed a healthy resurgence as the era of microcomputers has matured. Most educators need to be reminded that one of the earliest examples of materials-centered instruction — the book — also remains among the earliest forms of distance learning. The ability to document the knowledge base of an expert (or group of experts) and then replicate and distribute this source of information to others cheaply and efficiently remains the hallmark of Gutenberg's invention.

Replication remains the principal appeal of materials-centered instruction. Cheap and efficient replication means that even big investments of instructional design and development efforts can be economically productive. The outlays of this investment are work and knowledge (Bunderson & Inouye, 1987). The computer epitomizes easy replication because once courseware is developed, the cost of duplicating the effort by copying disks is almost trivial. Each duplicated disk represents the total replication of perhaps years of work and knowledge. This inexpensive disk can then be shared commercially (or otherwise) with instructional consumers' such as teachers, trainers, and students. No distance is too great between the developer and the consumer so long as the consumer has the means and resources to use the product. In all likelihood, graphics will continue to be found throughout replicated computer courseware. The question begging to be asked, then, is, what is it we are replicating, exactly? In every case, the answer has two parts — a product and a process. Both need to be evaluated carefully. Let's examine the product first.

The instructional design of a given medium is biased to delivering only certain instructional attributes or stimuli. Audiocassettes deliver aural stimuli, such as the spoken word, music, and other sounds, generally in a predetermined order. Overheads can present only static visuals, such as words and pictures. Other media offer mixtures of attributes. Slide/tape projectors can offer static visuals and sound. Video (and film) can offer static and dynamic visuals and sound. Computers offer the delivery of static and dynamic visuals (animation) in linear and nonlinear formats with high-quality sound becoming increasingly available. Link the computer to other media, such as videodisc players, and you have examples of multimedia in which the number of instructional attributes is almost unlimited.

Of course, computers do a poor job of delivering attributes unique to the human medium, also known as the teacher. Anticipating a wide range of responses, especially voice inputs, always has been a stumbling block for CBI design. In addition, computers can't react to a student's puzzled or frustrated look, and computers cannot determine the right kind of encouragement to offer when a student's self-esteem needs a boost.

Each medium demands a unique set of production skills that usually emphasize special attributes of that medium. Most developers find particular delight in producing packages that highlight a medium's attributes.

The history of filmmaking is a good example of how production characteristics of a medium mature over time (Papert, 1980). Early filmmakers began by mimicking what was then status quo: a camera, placed on a pedestal, photographed a play acted out on a stage. It is remarkable how far the medium has come during this century, considering all the production effects that are now part of standard filmmaking: zooming, panning, lighting effects, color, flashbacks, and so on. The next time you watch a movie or television show, do this simple test of the sophistication of its production: count the number of different camera angles and the duration of each. Most viewers are surprised to learn that five seconds rarely pass before the angle changes. **Transparency** is a hallmark design characteristic of most media. Well-produced film and video productions make it easy for us to become personally involved in the story or drama, and less likely to remember that the actors are really just performing in front of a large production crew. Transparency also means that a person's attention can be devoted to the message and not distracted by the medium — an especially important notion for instruction.

The point is that production skills are always biased to the attributes particular to each medium. If we evaluate only the product side of instructional media, then we are judging quality only on the basis of production. It is easy to forget that there is also a process that is being replicated. This is unfortunate, because the *process* should determine the product. In instructional media applications, this process is called **instructional design**. As certain instructional media, such as the computer, increase in popularity, the importance of "putting the horse in front of the cart" also increases. Guiding the design of graphics in educational computing is especially important as the graphical power of computers increases. The temptation to incorporate a wide array of graphics, simply given the awesome ability to do so, can be overwhelming. Guidance is needed to ensure that instructional processes direct instructional product development.

INSTRUCTIONAL DESIGN VERSUS TECHNOCENTRIC DESIGN

Designing instruction is a formidable task. The development and refinement of instructional design has been slow, but continual (see Reigeluth, 1983a, for a review). There are essentially three approaches: empirical, artistic, and analytic.

The first is called an **empirical approach** because it is based largely on trial and error. You begin with your best idea, try it out, carefully observe what works and what doesn't, and then make adjustments for the next attempt. In instruction, this translates into a system in which very rough products are field-tested with students. Usually, the starting point of this system is just one's best guess, with little or no supporting rationale. The chances of success for the first few attempts are slight. During each trial, the designer tries to observe key points at which failures occur. Because a pure empirical approach does not use any particular model to guide the design effort, few improvements come from any one trial. However, the design is slowly and progressively shaped to achieve the desired goal. This iterative process is often the strategy used by content experts given teaching or training responsibilities for the first time (such as new university professors in the hard sciences, and, to a lesser extent, beginning public school teachers).

With the second approach, called here **artistic design**, the tendency is to look at instruction as an art or craft that takes years to hone and perfect. From this perspective, each instructional design project is like a painter's latest masterpiece. Similarly, teachers often begin to build up a repertoire of teaching strategies and ideas based on years of experience. A master teacher is seen as someone who is somehow able to go into a room all alone and emerge with an instructional plan that everyone trusts. These people may be excellent models and perhaps good mentors to novices but are largely unable to explain what precisely they do or how they do it.

The empirical and artistic approaches are quite inappropriate in the design of materials-centered instruction. Pure empirical approaches based on a large number of trials are expensive and waste time, and the likelihood of finding a good ready supply of instructional artists is very slim.

This leads to the third approach, here called the **analytic approach**. An analytic approach uses a systematic and systemic plan to guide the decision-making process of instructional design. (See Footnote 6) The analytic approach is best known as the **systems approach** to instructional design, on which scores of models are based (see Andrews & Goodson, 1980, and Gustafson & Powell, 1991, for examples and reviews). The analytic approach tries to keep all the variables in the design process properly balanced and views each in the overall context and perspective of the task at hand. The analytic approach also incorporates the empirical and artistic, but controls for their weaknesses and potential excesses. The methods of the analytic approach are analogous to the direction provided by an orchestra's conductor. Without proper guidance and control, many musicians soon would be competing for recognition and control of the musical score. The conductor ensures that one interpretation of the selection is followed but allows the artistry of each member to contribute to the overall effort. Other analogies liken an analytic approach to the recipe of a master chef, a road map on a cross-country trip, or the blueprint for the construction of a house (Reigeluth, 1983b). However, the cab driver's advice to the lost violinist that the best way to get to Carnegie Hall is to “practice, practice, practice” is still sound. The analytic approach, through the processes of formative evaluation or rapid prototyping, also takes the position that only so much can be anticipated. Regular and systematic field testing of the materials is an integral part of the analytic approach.

There are scores of instructional design models based on the analytic approach. An example of a model meant for CBI design is shown in Figure 1.3 (adapted from Burke, 1982). However, the application of these models is frequently prone to many pitfalls and misconceptions, even by seasoned and well-intentioned instructional design veterans. The most deadly is the tendency to overmechanize the process by believing that successful instruction will result so long as each step (procedure, or flow chart box, and so on) is followed. The other tendency is to ignore the potentially important contributions of the empirical and artistic approaches. Creativity and innovation should not be considered to be inconsistent with an analytic approach, nor should an analytic approach become too far removed from the classroom or other training environments.

Another common pitfall is the tendency to put on design “blinders” and ignore everything but a predetermined goal. This pitfall likens learning to a conveyor belt on which the student inches slowly toward mastery of a handful of objectives. No attention is given to other healthy pursuits that may occur en route, nor is the student given any credit or responsibility to take on other explorations than those charted by the designer. This pitfall is largely an artifact of the dominant behavioral influence in the history of instructional design. This perspective tends to describe learning as just a consequence of being there. This approach can work well when the learning goals are narrowly defined, such as fact learning, but it does not hold up as well when goals lean more toward problem solving (Clark, 1984a). The contrast to this is instructional design based on a cognitive approach, which, simply stated, looks at design from the learner's point of view (Case & Bereiter, 1984).

The Burke Model

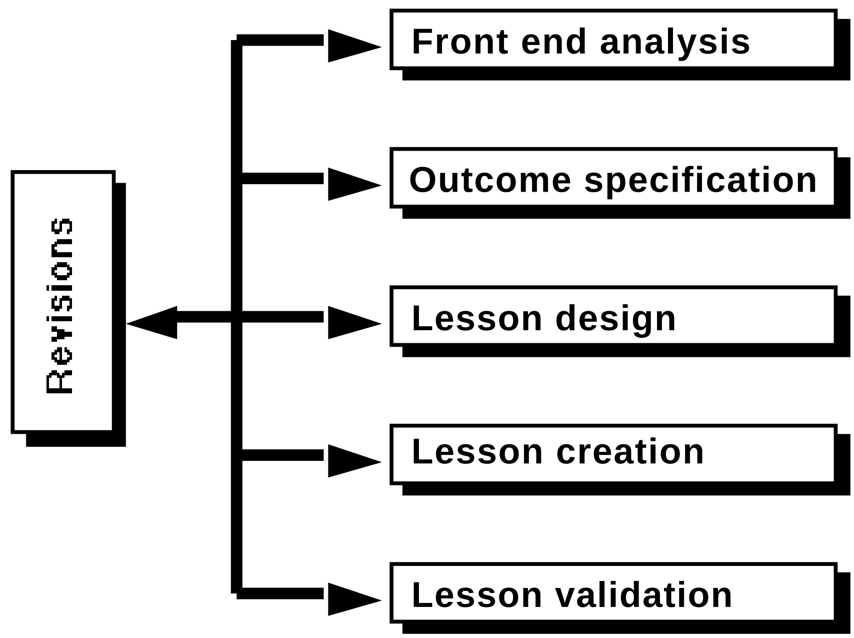

FIGURE 1.3
An example of an instructional design model applicable to CBI.

At least one other approach competes with instructional design. This approach goes in a completely different direction and is very prevalent in materials-centered instruction across all media, including computers. It is also very contagious and seems so right to instructional designers who are confronted with tough design decisions. This approach does not put the learner at the center of the process; neither does it actually put instruction at the center. This approach, which we will call “technocentric” design” (Papert, 1987), lets the technology dictate decision making, as summarized in Figure 1.4. (Usually, technology is equated with the products of instruction, such as computers, overheads, and videocassette recorders. Try to remember that technology can refer to products *and* processes.) The misuse of graphics has been especially prone to this approach. As the word technocentric suggests, a certain technology is put at the center of the process and all subsequent design decisions are based on their relationship to that technology. At first glance, it is hard to find fault with this approach. But there are many dangers inherent in this approach. For example, the decision to include high-resolution graphics, color, pull-down menus, etc., is largely based on whether or not the computer has the particular capability rather than if the feature is really necessary for learning. Often, designers and consumers of educational computing unconsciously fall into technocentric traps. We might call them “technoromantics” because they are honestly infatuated with their machines and believe that good instruction always incorporates all of their machine's capabilities (Ragan, 1989).

Technocentric design

A hardware-centered approach with evaluation of instruction based on how well the capabilities of the computer are used.

Instructional design

A learner-centered approach with evaluation of instruction based on the goals or objectives of the lesson, the needs of the learner, and the nature of the task.

FIGURE 1.4
Instructional design versus "technocentric" design.

A technocentric designer would criticize CBI, for example, that contains no graphics or animation. A technocentric attitude encourages the use of all special features, instead of questioning whether such features are relevant to the lesson goals or distract a learner's attention. Technocentric designers ask questions like “Can it be done on the computer?” instead of questions like “Who are the learners?” and “What are my instructional goals?”

Obviously, the premise taken here is that one must be very careful *not* to take a technocentric view in designing instruction. Unfortunately, many people who design CBI for a living tend to think first about the computer, not the learner, when they start a new project. As they build their instruction, every aspect of the design is related to the computer. Instructional design, by contrast, puts *final* media selection at the end of the design phase and right before final development (Gagné, Briggs, & Wager, 1992). The starting points are always the learner and the instructional objectives. These lead to the design of instructional strategies and, in turn, to the selection of the most appropriate instructional medium (or media) to deliver the instruction.

Media decisions should not be made until other instructional decisions have been made. For example, you may decide for one reason or another that real motion is important in communicating the idea or that demonstration is necessary in order to model a certain procedure. After instructional strategies and tasks are determined, you can then look to media that offer appropriate attributes to deliver them to the student. At first you consider only the *ideal* alternatives, but soon the decision-making process must also include compromises, which may lead you away from the ideal media to ones that are available or practical. Instructional design with graphics involves many considerations beyond the instructional ones. Economic considerations can override good design intentions. Constraints to the eventual installation of a program also affect media decisions. For example, you may be able to develop a course on computer, but not be able to adequately deliver it by computer given lack of hardware for large numbers of students. Instructional design makes you weigh all these issues carefully at all levels. This book's approach to design follows an analytic approach that is flexible enough to take advantage of the strengths of the empirical and artistic approaches.

REVIEW

- There are times when graphics can aid learning from instructional materials, times when they are detrimental to learning, and times when they do neither harm or good. Designers of instructional computer graphics must acquire the wisdom to know the difference.
- Regardless of their effectiveness, graphics (and other visuals) are an integral part of most teaching strategies.
- The rise in graphical user interfaces (GUIs) has permitted true integration of graphics and text in microcomputer systems. This has led to advancements in the production of computer graphics and has increased the ease with which graphics can be incorporated into instructional materials.
- Inexpensive replication of both products and processes is a major advantage of materials-centered instruction (of which CBI is an example).
- Technocentric design, featuring elements centered on computer capabilities, poses serious threats to effective materials-centered instruction.
- Instructional design is based on an analytic approach that combines empirical and artistic elements.

NOTES

1. A reminder that *visuals*, as the term is used here, refers to all visual stimuli, including both text and graphics.
2. Think back to a time when you listened to a master storyteller for an example of seeing a story in your mind.
3. Do this activity some weekend: Begin counting how many times you see a certain commercial. If you are an avid television watcher, keep the count going for the week. You'll be amazed how much “rehearsal” time you are inadvertently spending on this one product. Advertisers also use another popular memorization ploy that associates the product name with an easily remembered musical jingle. Visuals and music act as memory pointers that quickly spread to associated “links,” such as product names.
4. That distinction is usually credited to the designers of the Xerox Star, though the ideas represented in GUIs predate even it.

5. Strictly speaking, the example in the figure is not a true object-oriented programming environment, but it is close enough to well represent the concept.
6. *Systematic* refers to activities and processes that are organized and often procedural. In short, systematic means that one has a plan and is following it. *Systemic* means "system-like" and views a complex entity as a system that is comprised of many dynamic and interconnected components or subsystems. A systemic or systems approach describes, explains, and predicts behavior and change in a complex system, such as an educational or instructional system. A car could be viewed as a system, being comprised of separate, yet interdependent, subsystems, such as fuel, electrical, exhaust, brake, and power systems.

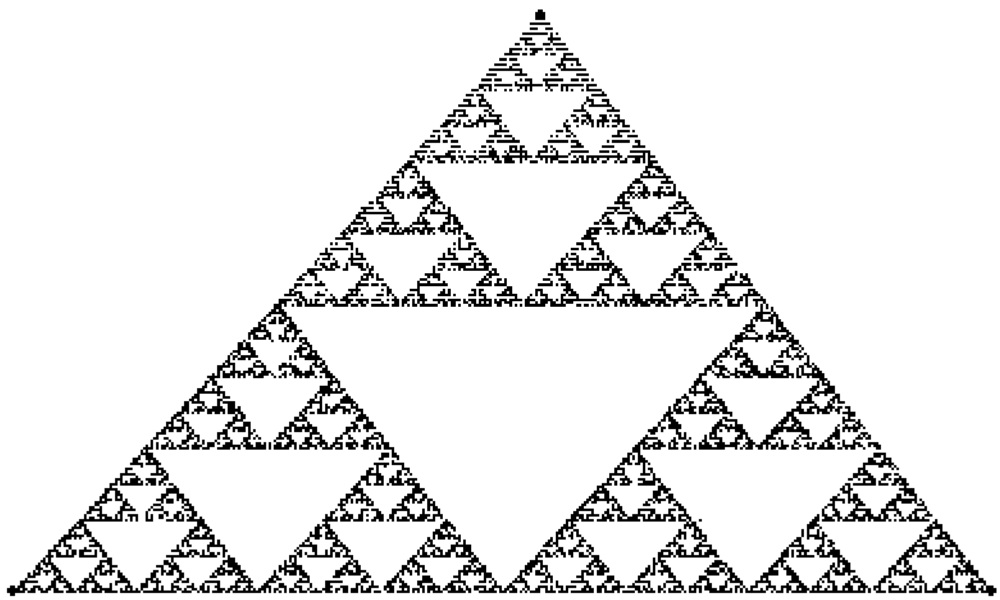

FIGURE 1.5
Here is the result of Box 1.2 — the Sierpinski Gasket. The pattern repeats itself to infinity.

CHAPTER 2

An Overview of Graphics in Instruction

OVERVIEW

This chapter presents and discusses a brief overview of the three major groups or types of graphics. It is important for instructional designers to understand these groups because they represent the types of graphics most often used in instruction. Discussed next are issues in applying these graphic types to instruction. The first step in this process is identifying the desired learning outcome. The chapter presents an overview of possible learning outcomes, called domains of learning. Finally, five instructional applications of graphics are presented according to the instructional function that each serves. In order to be effective, the instructional function of any graphic must match the particular needs of the lesson.

OBJECTIVES

Comprehension

After reading this chapter, you should be able to:

1. List the three major groups of graphics most commonly found in instruction and give a definitive example of each.
2. List the five domains of learning as defined by Robert Gagné.
3. Describe the role of graphics in facilitating learning for each domain.
4. List the three instructional applications of graphics that serve cognitive functions of learning and describe their relationship to the events of instruction.
5. List the two applications of graphics that serve affective functions.
6. Describe instructional situations where graphics are not appropriate and may distract the learner from lesson goals.

Application

After reading this chapter, you should be able to:

1. Classify graphics in given instructional materials according to the characteristics of one or more of the three graphic groups.
2. Evaluate the effectiveness of graphics in given instructional materials, based on the function that each graphic serves.
3. Evaluate the extent to which any given graphic interferes with lesson goals, such as by distraction.
4. Generate at least five examples each of effective uses of graphics from instructional materials and the popular media.

Though prescription is usually the desired goal in instructional design, description is a necessary starting point. For this reason, this chapter provides an important beginning in describing how graphics are used in instruction. First, this chapter will describe an overview of the groups or types of graphics typically used in instruction. These groups are fundamental to understanding the concepts in the remaining chapters.

The design of all effective instructional materials, including graphics, starts by defining the goals of the lesson and the nature of the learning tasks and materials. For this reason, the second part of this chapter will describe the range of learning outcomes that instruction usually addresses. Finally, this information is used to generate an informal guide to instructional applications of graphics. This guide not only provides a simple way to describe the role of graphics in instruction, but it can be used for prescriptive purposes such as those presented later in the book.

THE THREE TYPES OF INSTRUCTIONAL GRAPHICS

Given the popularity and flexibility of graphics in instruction, a way is needed to make sense out of how they can be used to improve instructional materials. First, there is a need to describe the types of graphics commonly used in instruction. Second, there is a need to describe the various functions of each type when applied in an instructional or training setting. We will use a simple classification system that describes the types of visuals commonly used in instruction. These categories describe, in general, how graphics convey information and meaning, but do not speak directly to how they can be applied in instruction. Applying these graphics types to instruction is a separate issue and will be addressed later. The three types of graphics are classified as representational, analogical, and arbitrary (Alesandrini, 1984), as shown in Figure 2.1.

Representational Graphics

Representational graphics share a physical resemblance with the object they are supposed to represent. For example, a passage of text explaining the purpose and operation of a submarine probably would be accompanied by a picture of a submarine. Representational visuals range somewhere between highly realistic and abstract.

The most common examples of realistic representational visuals are photographs or richly detailed colored drawings, the latter of which are currently the highest quality images that can be generated on microcomputers. Multimedia systems present opportunities to incorporate near-photographic images, such as composite video images taken from videodisc or videotape players, or from computers with adequate memory. Although many would argue that the quality of these video images is much lower than photographs, the issue of representational integrity is largely a function of the context. For example, although most microcomputers could represent a realistic enough submarine for most purposes, the same quality would hardly suffice for an art lesson in which fine details of the Mona Lisa are featured and discussed. Actual photographic images can be made available in multimedia systems that integrate slide/tape projectors (Pauline & Hannafin, 1987).

Types of instructional graphics

Representational

Representational graphics share a physical resemblance with an object or concept. These are further classified based on the degree of realism from highly concrete (photographs) to highly abstract (line drawings).

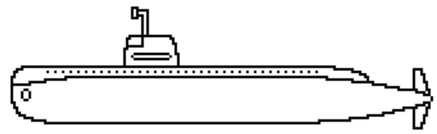

A submarine is an ocean-going vessel capable of sailing under water.

Analogical

Analogical graphics show something else and imply a similarity. It is crucial that the learner understand the analogy.

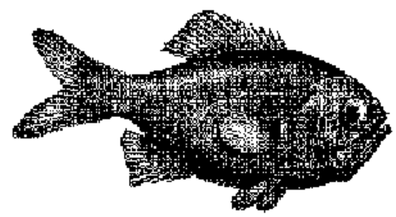

A submarine is like a fish in that it can "swim" under water.

Arbitrary

Arbitrary graphics share no physical similarity with the things they represent, but illustrate logical or conceptual relationships using a variety of visual and spatial means.

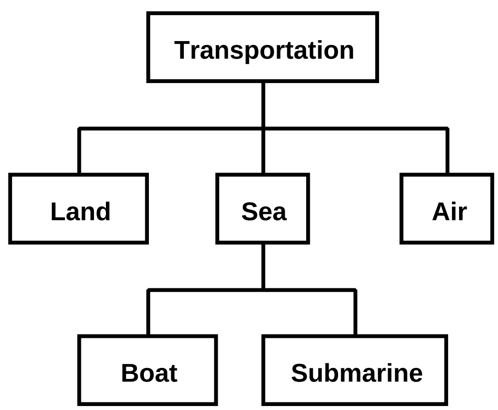

FIGURE 2.1
The Three Types of Instructional Graphics

An example of an abstract representational visual is a line drawing. These also range in quality from richly detailed to rudimentary drawings. For example, Figure 2.2 shows an example of a passage explaining the use and function of an astronaut's space suit. While it is clearly a line drawing, it was produced from a photographic original. On the other hand,

Figure 2.1 shows a rather crude drawing of a submarine. This primitive drawing still captures the most salient features of a submarine. In fact, the lack of interesting details and background makes it easier to focus on the essential characteristics of a submarine and far less likely to get confused or distracted by extraneous details. For these reasons, simple line drawings are often considered better learning aids than realistic visuals, especially when the lesson is externally paced, such as in films and video (Dwyer, 1978). The issue of realism will be discussed in more detail in chapters 5 and 7.

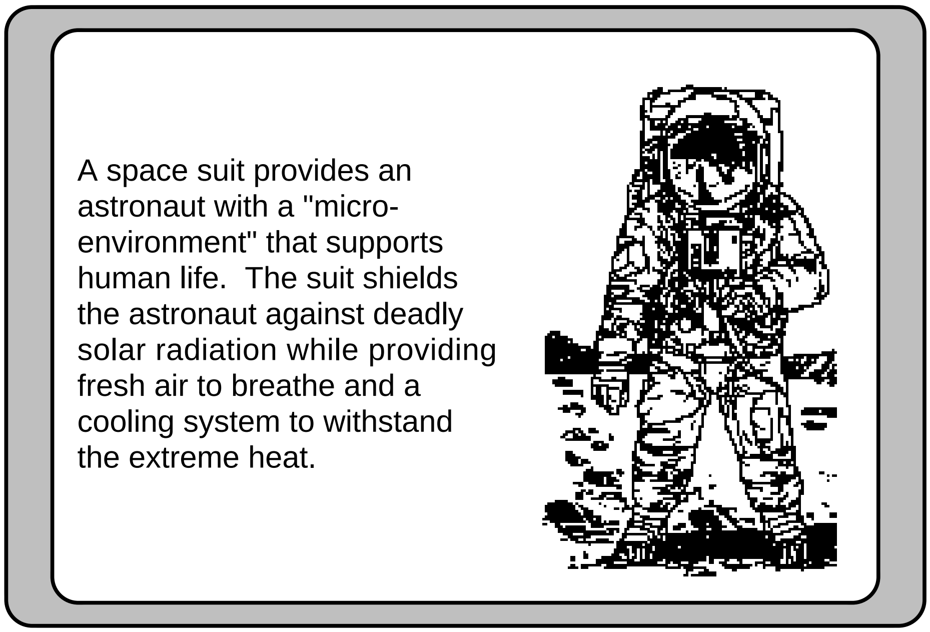

FIGURE 2.2
Snapshot of a CBI lesson using a presentation graphic consisting of a representational line drawing.

Analogical Graphics

The range of representational visuals is probably the most common type of illustration used in instructional materials today, including computer environments. However, presenting students with an accurate representation of something may not always be the best learning tool. One such example is when students have absolutely no prior knowledge of the concept. Instructional research indicates that analogies may be effective instructional strategies in such instances (Curtis & Reigeluth, 1984; Halpern, Hansen, & Riefer, 1990). For example, if students do not understand the idea that a submarine is able to dive under water, it might be more appropriate to first suggest that a submarine is analogous to a fish so students understand this characteristic. However, a better analogy would be a dolphin because it, like a submarine, must surface occasionally for air, or better yet, a whale, because of its size. Of course, a submarine is not a dolphin or a whale, so learners must understand that the analogy is being used only to represent similarities. Differences do exist, and it is important that students understand the analogy's limits.

Educational psychologists often describe learning as a process that goes from the known to the unknown (Reigeluth & Curtis, 1987). An analogy can act as a familiar “building block” on which a new concept is constructed (Tennyson & Cocchiarella, 1986). Of course, if the student does not understand the content of the analogy, then its use is meaningless and confusing. Worse yet, students may form misconceptions from an inadequate understanding of how the analogy and target system are alike and different (Zook & Di Vesta, 1991). The usefulness of the analogy, therefore, is largely dependent on the learner's prior knowledge. Graphics can help learner’s see the necessary associations between parts of the analogy. An example of a not so subtle analogical graphic is shown in Figure 2.3. The organization that paid for this ad obviously believes that America's dependency on foreign oil is a big mistake and is like a bomb ready to go off. Whether or not you agree with this position does not detract from the obvious message that is being communicated with this graphic.

FIGURE 2.3
Example of an analogical graphic.

Arbitrary Graphics

Arbitrary graphics offer visual clues, but do not share any physical resemblances to the concept being explained. In a sense, this category acts as a "catch-all" for any graphic that does not offer any resemblance of real or imaginary objects, but yet contains visual or spatial characteristics that convey meaning. Examples range from the use of spatial orientations of text, such as outlines, to flowcharts, bar charts, and line graphs.

All information can be represented as existing on a continuum. At one end are the most concrete representations — real objects. Nearby are highly realistic representational pictures. At the other end are spoken and written words that represent the most abstract form of communication. In the center of this continuum would be arbitrary graphics.

Charts and graphs are probably the most common types of arbitrary graphics (Winn, 1987). Charts refer to tables or information contained in table-like formats. Examples include taxonomies, such as the classification of animal groups, language families, or baseball teams (such as shown in Figure 2.4). The purpose of a chart is to organize and display information by one or more categories or fields. All of the information in a chart is discrete (categorical) data.

A "cognitive map" is an interesting example of a chart that has much support from research as a learning tool. Cognitive maps are part of an instructional technique called spatial mapping (Holley & Dansereau, 1984). The purpose of cognitive maps is to show graphically the relationships and hierarchies of related ideas and concepts. Figure 2.1 depicts a simple example of a cognitive map that shows how two concepts — submarine and transportation — are related. Each fact or concept is called a node and is connected to other nodes by links that indicate the relationship between the nodes. Often, these links are then labeled further to clarify the relationships between the connected nodes. Research has shown that these graphics tend to be most useful when the *student* constructs the map or when the map is constructed in front of the student, usually during the explanation of the ideas, rather than just providing a completed map to a student to study.

Similarly, graphs also logically represent information along one or more dimensions, but the main purpose of graphs is to show relationships among the variables in the graph, as shown in Figure 2.5. The most common types of graphs are line graphs and bar graphs, although many other types abound, such as pie graphs, scatterplots, etc. Another difference between charts and graphs is that at least one of the variables in a graph usually will be continuous. Continuous data contain an infinite number of points along a continuum. Height or weight are continuous variables. Someone's height may be reported as six feet, one inch, but this is just for convenience because height can never be measured exactly.

National League	
East	**West**
Chicago Cubs	Atlanta Braves
Florida Marlins	Cincinnati Reds
Montreal Expos	Colorado Rockies
New York Mets	Houston Astros
Philadelphia Phillies	Los Angeles Dodgers
Pittsburgh Pirates	San Diego Padres
St. Louis Cardinals	San Francisco Giants

American League	
East	**West**
Baltimore Orioles	California Angels
Boston Red Sox	Chicago White Sox
Cleveland Indians	Kansas City Royals
Detroit Tigers	Minnesota Twins
Milwaukee Brewers	Oakland A's
New York Yankees	Seattle Mariners
Toronto Blue Jays	Texas Rangers

FIGURE 2.4
Example of presenting categorical information in a table.

The way space is used in a chart or graph to form sequences and patterns is very important. Research has shown that more rapid problem solving results from diagrams in which conceptual relationships are shown spatially, rather than by text (Win, Li, & Schill, 1991). In charts, the sequence of information is usually not a critical feature. For example, there are many ways to sequence the various names and hometowns of major league baseball teams. It doesn't really matter if the American League or National League is listed first or second. The teams are listed alphabetically in Figure 2.4, but changing this order does not change how information in the chart is conveyed.

The pattern of a chart is typically conveyed through row or column headings. The baseball chart is informational because of the primary and secondary groupings: a) American and National; and b) East and West. Also, the proximity of items to one another in a chart may also convey information. A chart that describes an animal family, such as marsupials, would show how much one group is related to another by how close the groups are located on the chart along one dimension.

The sequence of a graph is crucial to understanding the information it contains. For example, the usefulness of a graph that describes average monthly temperatures, such as those shown in Figure 2.5, would be seriously curtailed if it were arranged alphabetically by month instead of chronologically. “Reading” the graph is easier when the graph displays information in a natural sequence. Also, the purpose of a graph is usually to compare information across parts of the graph, such as which times of the year are the hottest or the coldest.

This also speaks to the importance of the pattern of information displayed in a graph. Consider Figure 2.5 containing three separate line graphs: one graph showing the monthly temperatures for both Houston and Pittsburgh, and then one superimposing the two graphs. Graphs such as these are meaningful if they convey trends and comparisons quickly at a glance. When superimposed, the line graphs quickly allow the reader to compare the climates of the two cities.

An effective and popular graph type is the **time-series plot**, where one axis is tied to some chronological variable, such as seconds, minutes, or years (Tufte, 1983). Scientists often use time-series plots to show how large and complicated data sets change over time. A simple example of the resulting motion of a bicycle's pedal as it turns while the bicycle moves forward at different speeds is shown in Figure 2.6. Of course, computer animation provides many opportunities for improving time-series plots, since the actual dynamics of the display over time could be shown and potentially controlled.

One of the most influential figures on how to visually display quantitative information has been Edward Tufte (1983). His most fundamental principle of statistical graphics is simply “above all else show the data” (Tufte, 1983, p. 92). Yet, it is amazing how often this simple principle is violated, sometimes unintentionally and sometimes deliberately to distort the data (such as for political motives). For this reason, Tufte defines the “lie factor of graphs” as the size of the effect shown in the graph divided by the actual size of the effect in the data. A lie factor of 1 denotes no lie, but ±.05 constitutes a substantial distortion of the data. Tufte (1983) also admonishes designers of graphs to keep “chartjunk,” nonessential graphical decoration, to a minimum. Tufte feels that “the best designs are *intriguing* and *curiosity-providing*, drawing the viewer into the wonder of the data. . .” (p. 121).

Combining Characteristics of the Three Types of Graphics

It should be noted that graphics are frequently constructed to contain characteristics of two or more of the three graphic types. Representational and arbitrary graphics are often mixed, such as the use of arrows and labels superimposed on a drawing. **Pict-o-graphs** (or **isotypes**), another popular type of graph (especially in magazines and newspapers), overlap characteristics of representational and arbitrary graphics, as shown in Figure 2.7.

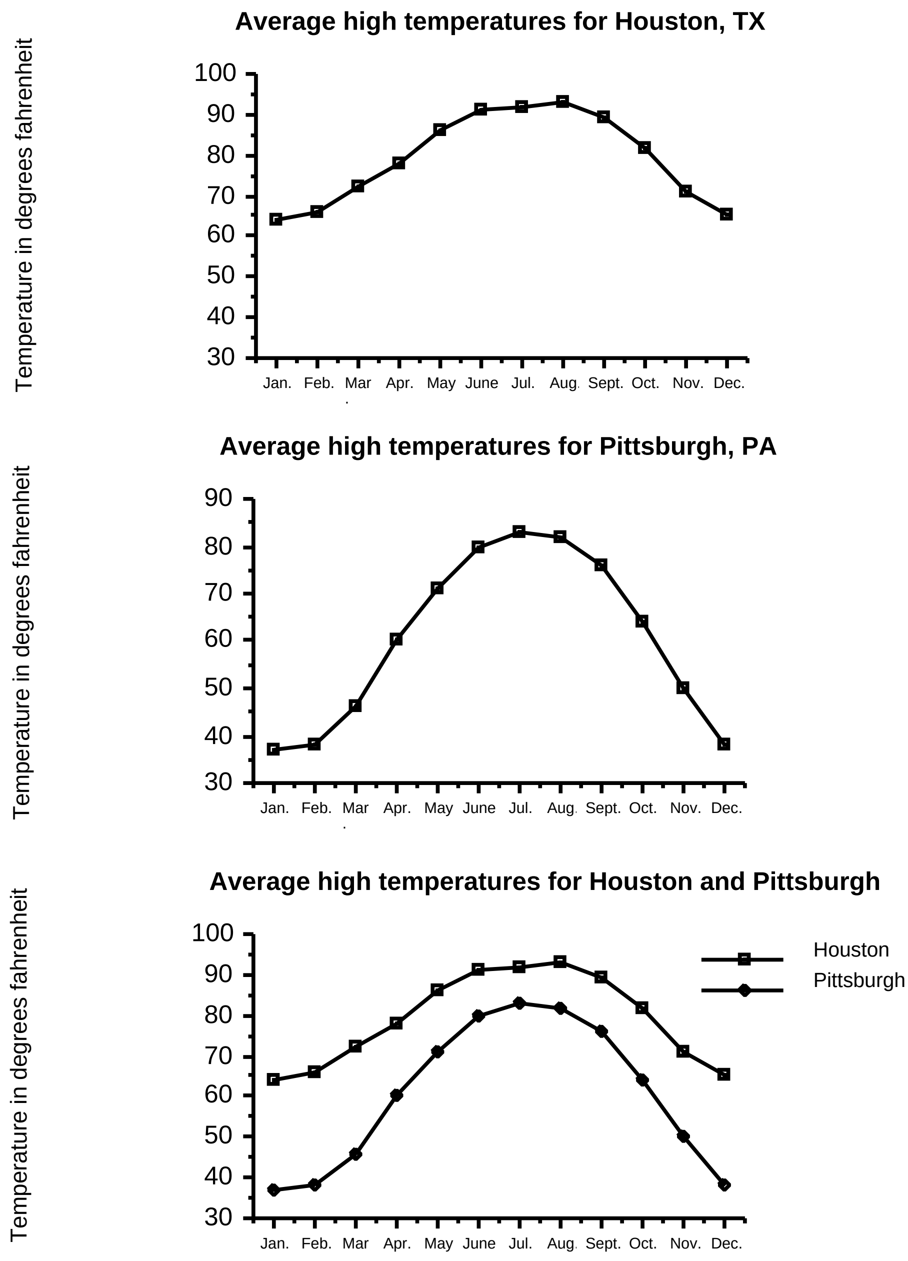

FIGURE 2.5
Examples of line graphs.

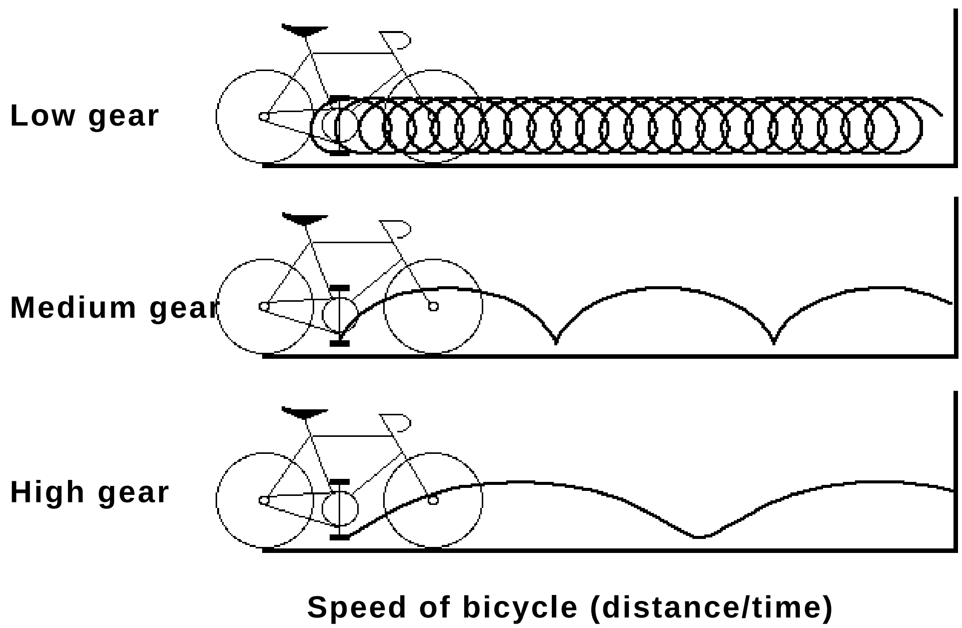

FIGURE 2.6
Three time-series plots showing the path that a bicycle's pedal follows as the bicycle moves forward, given different gear ratios.

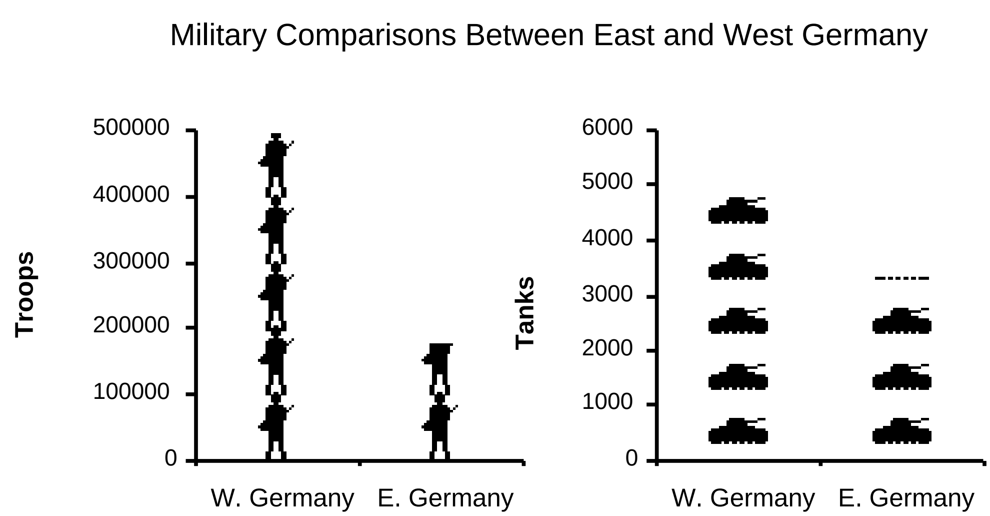

FIGURE 2.7
Examples of "pict-o-graphs."

The overlay of representational and arbitrary graphics onto geographical maps is one of the oldest mixtures of graphical forms. One of the most striking examples is the map drawn by the French engineer Charles Joseph Minard in 1861 to show the tremendous losses of Napolean's army during his Russian Campaign of 1812. The map, shown in Figure 2.8, is best described by Tufte (1983):

> Beginning at the left on the Polish-Russian border near the Niemen River, the thick band shows the size of the army (422,000 men) as it invaded Russia in June 1812. The width of the band indicates the size of the army at each place on the map. In September, the army reached Moscow, which was by then sacked and deserted, with 100,000 men. The path of Napoleon's retreat from Moscow is depicted by the darker, lower band, which is linked to a temperature scale and dates at the bottom of the chart. It was a bitterly cold winter, and many froze on the march out of Russia. As the graphic shows, the crossing of the Berezina River was a disaster, and the army finally struggled back into Poland with only 10,000 men remaining. Also shown are the movements of auxiliary troops, as they sought to protect the rear and the flank of the advancing army. Minard's graphic tells a rich, coherent story with its multivariate data, far more enlightening than just a single number bouncing along over time. *Six* variables are plotted: the size of the army, its location on a two-dimensional surface, direction of the army's movement, and temperature on various dates during the retreat from Moscow (p. 40).

MATCHING GRAPHICS WITH LEARNING GOALS

An understanding of the three graphic types is prerequisite to an understanding of how they can be used in instruction. As one might guess, there are a myriad of specific ways that each of these graphics can be used in any one instructional situation. The next section begins the discussion of the important issues surrounding *instructional* applications of graphics, such as *should* one or more graphics be included in an instructional design and, if so, what role or function should those graphics serve. A necessary first step in the design of effective instructional materials and, subsequently, instructional graphics, is the determination of the lesson goals and objectives.

Instructional Objectives

The most common vehicle for describing learning goals in any one lesson are instructional objectives, also known as performance objectives (Briggs & Wager, 1981). The purpose of instructional objectives is to describe as clearly and precisely as possible what the learner should be able to do *at the completion of the lesson*. If constructed properly, objectives not only serve as an appropriate guide in the design of instructional strategies and materials, but also indicate appropriate methods of evaluating whether the objectives have been met.

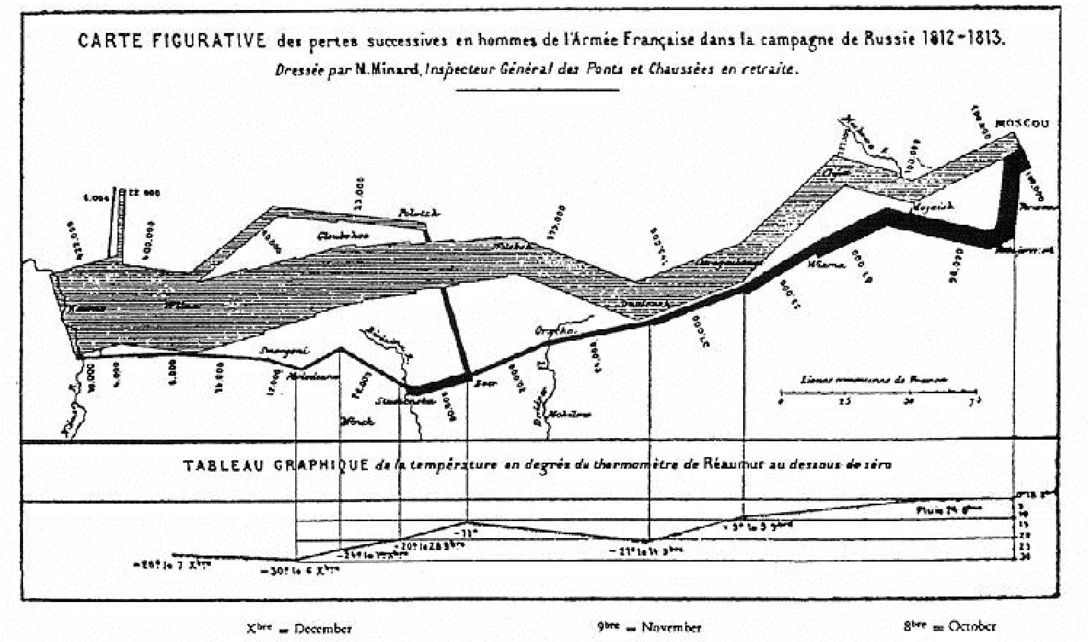

FIGURE 2.8
Minard's map of Napolean's russian campaign of 1812.
Edward R. Tufte, *The visual display of quantitative information* (Cheshire, Connecticut: Graphics Press, 1983).

There are many recipes for how to write objectives, but the ABCD model is one of the simplest. This model, as the name implies, has the following four parts: A — a detailed description of the *Audience* or learner; B — a clear and unambiguous description of the *Behavior* or skill that the learner should be able to do at the end of the instruction (usually containing a carefully chosen action verb); C — the *Conditions* under which the behavior will take place, such as tools permitted in performing the behavior or any time restrictions; and D — the *Degree* to which the performance must be accomplished, such as complete accuracy (100%).

The effectiveness of instructional graphics largely depends on the nature of the learning task (as described by the behavior) as it interacts with the profile (aptitude and interests) of the learner. The behavior also should reflect the learning domain being emphasized in the lesson. A well-written objective not only helps guide instructional design, but also the type and function of any appropriate graphic.

Domains of Learning

An understanding of general learning theory is essential in designing effective visual displays. Although a thorough overview of learning theory and its applications to instructional graphics design will be presented in chapter 4, an initial understanding of the range of possible learning outcomes will be discussed next because it is a vital first step. Therefore, the purpose of this section is to introduce the range of learning outcomes in the context of the role that graphics may serve in facilitating those outcomes.

There is a wide range of learning outcomes. Probably the most well-known description of learning and knowledge is that provided by Benjamin Bloom (1956). Though somewhat dated, Bloom's original taxonomy of domains of learning is still considered as the standard against which current perspectives are compared. Bloom divided learning and knowledge into three domains: cognitive, affective, and psychomotor. The psychomotor domain involves the learning of physical tasks that require eye-hand-mind coordination. The affective domain largely comprises a person's attitudes and value systems. The cognitive domain concerns the learning of facts and skills and, for better or worse, comprises the “lion's share” of mainstream educational activities.

Robert Gagné (1985) has refined and extended Bloom's original descriptions to include five domains (as shown in Figure 2.9). He divides the cognitive domain into the three separate domains of verbal information, intellectual skills, and cognitive strategies, while keeping the affective and psychomotor domains essentially the same. Just as Bloom's model provides the most detail for the cognitive domain, Gagné's model is most complete for verbal information and intellectual skills. Gagné's model will be used here because of its wide acceptance and application in educational technology. In addition, his model has been widely applied to instructional design (see the section elaborating on Gagné's events of instruction later in this chapter). These events comprise a micro-model of instructional design that describes and prescribes the instructional “milestones” of each lesson. Each of the five domains requires special considerations for the design of graphics. As a field, educational technology has been primarily interested in the cognitive domain, as it encompasses the bulk of instructional questions and problems. For this reason, verbal information and intellectual skills will be the two domains emphasized throughout this book.

Verbal Information Domain

Verbal information involves the learning of factual material and includes verbatim learning, nonverbatim learning, and substance learning. Verbatim learning is learning by rote, such as memorizing a poem word for word. Nonverbatim learning is the memorization of isolated facts, but in a learner's own words. An example is “Columbus discovered America in 1492” (although a Native American might dispute this “fact”). Substance learning involves the summarization of an instructional passage (such as that presented by text, video, or lecture) in a student's own words without requiring interpretation or application.

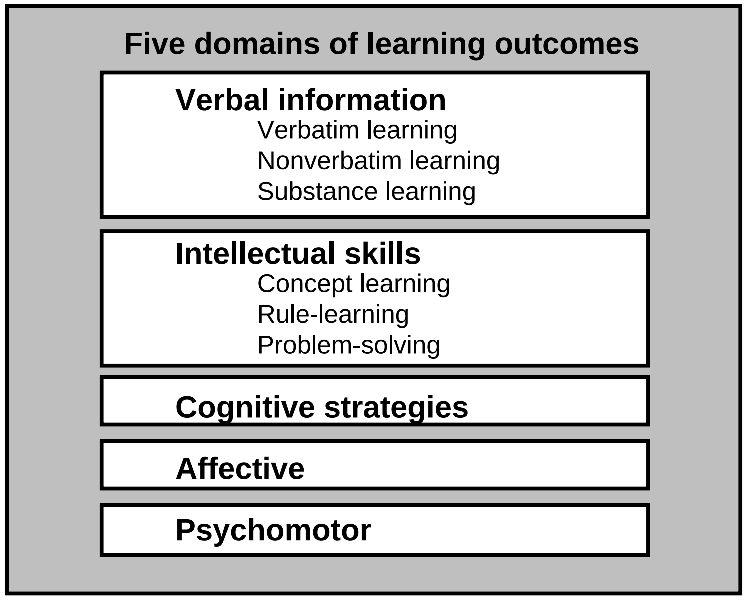

FIGURE 2.9
The five domains of learning.

An example of how images are used to help a person remember a fact is perhaps best illustrated with a television commercial for margarine popular several years ago. In the commercial, a person spreads the margarine on a piece of toast and takes a bite. Trumpets immediately sound, and a crown appears on the person's head. The intent of the commercial is to associate the image of a crown with the product name (Imperial, in this case). When the viewer goes shopping and has to decide which brand of margarine to buy, the amusing commercial and the image of a crown may be recalled. This, in turn, elicits the product's name. (In contrast, few people remember the exact brand name of another margarine based on the slogan “It's not nice to fool Mother Nature!”) (See Footnote 1.)

Using visuals to help remember isolated facts and details is an old, but successful, strategy (Bower, 1970; Carney, Levin, & Morrison, 1988). Visual mnemonics is at the heart of many “memory improvement” systems. In these systems, participants learn ways to quickly contrive and associate some visual image to a fact to be remembered, such as a person's name. The weirder and wilder the image, the stronger the memory trace will be. For example, try this little experiment for remembering the Spanish word for duck — pato (pronounced pah’ toe). Take 30 seconds and visualize a duck with a pot on its head like a hat. Think about how “pot on duck” resembles “pato duck” while visualizing this image. As a test, write “Spanish word for duck” on a slip of paper and place it in your pocket. When you find the paper hours or days later, see if you can remember the word. For most people, this little activity makes them remember the Spanish word for duck forever!

Other examples of visual mnemonics include the many pegword systems (e.g., one is a gun, two is a shoe, etc.) and the method of loci (Just & Carpenter, 1987). The method of loci is a classic strategy often associated with public speakers. The trick is to associate parts of a speech with a mental and visual “tour” of a place you know well, such as your house. You then remember the speech by going on a mental "walk" through your house in your mind. This technique was often the same one used by traveling minstrels or poets hundreds of years ago, many known for their prodigious memories. Although visual mnemonics are proven strategies for recall tasks and other fact learning, there is also some evidence of their utility for higher-level learning as well (Levin & Levin, 1990).

Visual mnemonics serve a transformation function (Levin, Anglin, & Carney, 1987) to directly impact and influence a student's associative memory. Such visuals are believed to be more memorable because of the way they target the most critical information to be remembered. Three components of transformational pictures help to explain this effect (known as the "three Rs" of associative memory techniques): the visuals *recode* the critical information into a more concrete form, *relate* it to a well-organized context, which subsequently helps a student to *retrieve* the information later.

Intellectual Skills Domain

Intellectual skills comprise a hierarchy of skills, each considered to be prerequisite to the other, beginning with concepts, then rules or principles, and, finally, problem-solving.

Concepts. Concept learning entails cognitive classification systems. Concepts are frequently classified as concrete or abstract (Tennyson & Park, 1980). For example, most nouns represent concepts, each falling on the continuum between concrete or abstract. Understanding the concept "chair" means that you can, for example, pick out all the chairs from a group of chairs and tables. To do this, you must understand the distinguishing attributes of chairs and tables — what makes a chair a chair and a table a table. Sometimes the context or situation can make even strange objects become part of the concept family. For example, a log taken from a stack of firewood can assume "chair" status if the wood cutter wants to take a small break and sit down.

Because there are many varieties and possible examples of any one concrete concept, individuals usually construct their own personal prototype for concrete concepts. For example, what image first comes to your mind when the concept "bird" is suggested? This image is your personal prototype for a bird. Chances are it was something similar to a robin, sparrow, or cardinal. These usually best represent the essence of "bird" for most people. Few people immediately associate an ostrich or penguin, perhaps because the inability to fly or the ability to swim do not match most people's attribute lists for birds. Some words, like cardinal, belong to several concept families, such as birds, religion, and sports (especially for people living in St. Louis). The particular concept family that gets triggered depends on the context (E. Gagné, 1985). Mental prototypes help us to organize world knowledge, although such prototypes can also explain the tendency of people to form stereotypes. Much research suggests that pictures can help in remembering concrete concepts (Paivio, 1986).

Abstract concepts are much harder to represent because they have no tangible form. Examples include the concepts of justice, freedom, honesty, and family. Designers often use a visual strategy where one concrete concept shows a snapshot of the abstract concept, such as a tree to represent the environment. Some abstract concepts hold strong cultural meanings for people. The concept of justice is often represented in American culture with the image of a blind-folded woman holding a set of scales, as shown in Figure 2.10. This image tries to communicate a concrete image of what is meant by justice, although it is certainly just one of thousands of possible representations (another common one is a judge's gavel). Of course, the danger of using such an image is that it may oversimplify the concept or bias the learner away from the breadth or range of examples that the concept actually represents. Analogies can be effective ways to teach abstract concepts (Newby & Stepich, 1987).

FIGURE 2.10
In many western cultures, this graphic is prototypical of the abstract concept of “justice” for many people.

How would you illustrate the familiar concept of “education”? Figure 2.11 also uses an analogical graphic for this purpose. This graphic is compelling because it allows many interpretations. The graphic causes one to pause and reflect, rather than providing only one narrow meaning. For example, what do you think the flame represents? To some, it might symbolize knowledge or enlightenment. To others, it might represent the learner. Do you see the hands holding out the flame to others (i.e., sharing), protecting the flame, or nurturing or helping the flame? In all cases, the graphic is an analogy or metaphor for an extremely abstract concept.

Rule Learning and Problem Solving. Rule learning and problem solving are examples of higher-order learning. Rules, also known as principles, comprise the learning of “if/then” situations and relationships. It is easier to precisely define rules in some content areas, such as mathematics and science, than in others. Examples of rule using in math would be the rules of addition and how to reduce fractions to lowest terms. In science, the application of any scientific formula, such as Newton's second law, would be an example of a rule. Of course, rules apply to all content areas, including social situations, such as deciding when to shake hands with someone. Research has shown that some visuals, such as schematics, can be used to help children learn mathematics (Fuson & Willis, 1989; Willis & Fuson, 1988).

FIGURE 2.11
How would you interpret this analogical graphic of the abstract concept "education"?

Problem solving is very controversial and is not easily defined. Gagné has operationally defined problem solving as the application of two or more rules at the right time and in the right sequence. Problem solving here consists of first isolating or defining the problem and then devising a solution based on rule selection, followed by the decision of when to apply what rule. Problem solving can become very complex very quickly, even using this simple model. There are hundreds of everyday examples. Consider what happens when you are at the grocery store trying to decide what brand of coffee to buy. You have decided to buy one of two brands, depending on which is a better bargain, but unfortunately each brand comes in a different size container. In order to determine the better buy, you have to select and apply the correct rules of how to calculate the unit cost of each brand.

Gagné's definition of problem solving as described above suggests a hierarchical nature of learning in the intellectual skills domain. Problem solving is seen as largely a function of how well all relevant and subordinate rules have been mastered and how well the many rules are associated. Consequently, mastering any one rule requires adequate understanding of the concepts that comprise it. This hierarchy of learning obviously imposes constraints on instructional design. There are competing theories of how people solve problems, such as those viewing the process in a holistic way or those dealing with mental heuristics (e.g., Polya, 1957). Also, many inductive learning theories suggest that it is possible for people to induce rules and concepts when put into problem-solving situations unprepared (Bruner, 1966).

Psychomotor Domain

The psychomotor domain involves the learning of motor skills that require eye-hand coordination, such as typing, riding a bike, and sharpening a pencil. Driving a car is a good example of a psychomotor task as a new driver tries to learn eye/hand, eye/foot coordination to the point of automaticity in order to make the car move and respond according to moment-to-moment demands.

Demonstration coupled with lots of practice remains an effective instructional strategy for the psychomotor domain because most motor skills involve the mastery of physical tasks that are procedural in nature. Media that possess motion, such as films, computer animation, or a real person, are logical choices for the delivery of instructional materials in this domain. For example, the military has long used films to train recruits to do tasks such as how to take apart and put back together weapons, such as rifles and machines guns (Spangenberg, 1973). Computer animation permits the visualization of the many stages of a task over time in concrete ways.

Affective Domain

The affective domain is best thought of as a person's attitudes, beliefs and value systems (Keller, 1983). "Choose" is the key action word for describing behaviors in the affective domain. Attitudes are often reflected in the free-choice patterns of people (Maehr, 1976). Most graphics used in magazine, newspaper, and television commercials deal with the affective domain. Some are very blatant, especially those targeted for certain subgroups, such as those using football or basketball players as role models to promote a product to teenage males. Other graphics present very pleasant, appealing, or highly interesting images (often with implied or expressive sexual connotations), which try to capture a person's attention for a few seconds. Still other visuals may try to associate a certain mood or feeling, such as power or success, with the product. Very few visuals actually provide consumers with accurate product information. (See Footnote 20 Billboards, particularly in highly populated urban areas, are notorious for tailoring their messages to the general profile of people living in the neighborhood, to the point of being stereotypic.

Cognitive Strategies

Cognitive strategies deal with personal mental activities that govern and control other mental operations. Gagné has called these executive control functions. Cognitive strategies originate with each individual. For example, think about what study strategies you use and why you use them. Many students simply read and reread text in order to remember it instead of taking the time and effort to learn more effective and efficient study strategies. Much of the literature dealing with metacognition (thinking about thinking) refers to cognitive strategies (Flavell, 1985). Cognitive strategies probably represent the least understood domain of learning.

Inter-Domain Relationships

The fact that the cognitive domain is stressed in most educational research literature should not suggest that the other domains are unimportant. It is just that researchers thus far have spent more time studying the cognitive domain. In practice, the domains strongly interact (Gagné, Briggs, & Wager, 1992). Crossovers between verbal information and intellectual skills are the most common ones discussed in the literature. People usually will have a need for factual information throughout the learning of intellectual skills. An example would be the experience of following a recipe that unexpectedly calls for grams instead of ounces half-way through the recipe. Factual information, the conversion formula, is needed while performing the procedural task (cooking).

Crossovers between the cognitive and affective domains also have been studied to a degree. For example, we know that a student's attitude for learning can strongly influence the time and intensity invested in a task. The crossover between the affective and cognitive domains is particularly important as it primarily refers to motivation and locus of control. Although research frequently indicates gender differences in learning about science, math, and computers, for example, no real evidence supports a psychological basis for the differences; rather, the differences can be attributed to environmental and cultural influences. Any feeling of inadequacy will certainly make one less likely to want to attempt to participate in a certain subject or task.

There are many crossovers between the cognitive and psychomotor domains. For example, there is no natural reason why anyone would come to a stop at a red light while driving a car. This information is part of the cognitive domain and must be related to the psychomotor skill of bringing the car to a controlled stop smoothly in anticipation of a red light.

The design of instructional graphics is strongly influenced by the interrelationships and interdependency of the five domains. For example, graphics meant to motivate (affective domain) should not interfere with other learning tasks in the cognitive or psychomotor domains. Frequently, designers lose sight of their original goals when deciding on the number and nature of the graphics they wish to include. The distinctions among the five domains must be maintained when designing graphics. Before you can begin to consider how graphics can enhance learning, you must first understand the importance of clearly identifying the desired learning outcomes and then choosing to design a graphic so that it supports these outcomes.

A GUIDE TO THE INSTRUCTIONAL FUNCTIONS OF GRAPHICS

Understanding the most common types of instructional graphics and how they are applied in the various learning domains is an important first step. Of course, simply describing the types of visuals says nothing of their uses and functions in instruction. This section presents an informal guide to help you choose the right type of graphic for the right job. This guide is primarily intended for CBI design, but it applies to other media as well.

The five applications of instructional graphics described in this section are **cosmetic**, **motivation**, **attention-gaining**, **presentation**, and **practice**. It is important that you understand the instructional philosophy from which these applications originate. With the exception of the first one, cosmetic, these applications represent major groups of instructional strategies. There are many ways for designers to integrate appropriate graphics in each group.

Characteristics of Successful Instruction

Think back over your many years of experience in education, whether as a teacher or a student, and try to think of one particular time when the instruction really worked. What was it about it that seemed to make learning click?

The search for the essential components of “good” instruction has a very long history, and there are many models and opinions of what actual components are involved. The hope has long been to reduce good instruction to a fundamental group of principles that could be easily replicated. Activities in the social sciences, however, are never that clear-cut. Still, the search for fundamental characteristics of good instruction is a worthwhile endeavor. There are some things on which most professional educators agree. For example, motivation ranks high on the list as an adaptation of an old adage points out: “You can send me to school, but you can't make me think!” It is easy to agree that motivation is important, but difficult to agree on what makes instruction motivating.

Many instructional models are based on behavioral philosophies where learning is viewed as an “input/output” activity — good instruction goes in and learning comes out. More recent ideas, based on cognitive psychology, recognize the role of what goes in between the input (stimulus) and output (response) as the most important element — student thought processes (Clark, 1984a; Gagné & Dick, 1983; Gagné & Glaser, 1987; Hannafin & Rieber, 1989a). One longstanding model that has been adapted to fit current theories of learning is called the *events of instruction*, also provided by Robert Gagné (1985) (Hannafin & Rieber, 1989b). The model has nine events, or “milestones,” as shown in Figure 2.12. Instruction needs to consider, though not necessarily incorporate, all of the events.

Most people outside of education usually think of instruction only in terms of the presentation of information, or event 4. It is easy to think of instruction as the “pouring” of information into a learner's head. However, event 4 includes the careful and deliberate selection, organization, and presentation of content. But it is not enough to simply present information to students. Good presentations must be coupled with careful guidance of what is being presented, as suggested by event 5. For example, students should recognize and distinguish among major and minor points and among relevant, incidental, and trivial information. Good instruction assures that this occurs. For this reason, events 4 and 5 are grouped together for our purposes as *presentation*.

The allure of this model is that it is simple and generic. But, as cautioned in chapter 1, don't let the simplicity of the model mislead you into mechanizing the process it represents. One of the most important premises of this model is that it views purposeful learning as a combination of external and internal conditions. The internal conditions are represented by student thought processes, the external by the instructional environment in which the learner is placed. This premise is based on an information-processing model of learning where thought processes influence how information from the environment is perceived, understood, and potentially stored in memory. Perhaps the largest determinant of all this is what the student already knows (prior knowledge) (Ausubel, 1968).

Orientation

1. Gain the learner's attention.
2. Inform the learner of the goals of the lesson and what to expect from the lesson.
3. Make the learner recall any prerequisite information that is important to the current lesson or that the current lesson builds upon.

Presentation

4. Present the lesson information.
5. Guide the learner as the lesson information is presented.

Practice

6. Provide opportunities for the learner to interact with the lesson.
7. Provide the learner with informational feedback based on these interactions.

Testing

8. Test the learner in reliable and valid ways on the predetermined learning outcomes.

Retention and transfer

9. Throughout the lesson, consider how to help the learner remember and apply the lesson information in similar and dissimilar contexts (learning transfer) in the near and distant future.

FIGURE 2.12

The events of instruction.

Notice how these events have been grouped in Figure 2.12. Rather than describing the events separately, it is useful to consider how events in each group interact within the group and then how one group influences other groups. To understand this, try to relate your personal experiences of "good" instruction with the discussion that follows.

While the importance of events 4 and 5 may be rather obvious to most people, students must be properly prepared for these events. It is important to "set the table" properly before "sitting down to eat." This is the general purpose of events 1, 2, and 3. These first three events act as an *orientation* to prepare the learner for what the following events have to offer. Event 1 makes the deliberate effort to gain and hold the learner's attention. Event 2 sets up learner expectancies, which are extremely important because they help learners to be selective as they learn. Event 2 gives learners a sense of what they should be doing or looking for during the lesson. This helps them to monitor their own learning in order to know when to go over material a second or third time, or to stop and ask questions. Event 2 also helps students to understand lesson procedures to prepare them for intense instructional “sprints” or to settle back and pace themselves for an instructional “marathon.” The importance of event 3 is easy to overlook. Event 3 is based on the philosophy that there is very little, if anything, worth learning that is not related, directly or indirectly, to other knowledge. Very little meaningful learning exists in a vacuum. So, if the current lesson relates to something important that was previously learned, instruction must assure that learners actively recall that prior information or knowledge into their working memories, so that they can actively relate the old information to the new. Again, the point is that instruction must not leave this to chance. Therefore, event 3 demands careful consideration. Taken together, these first three events do an important job of preparing the learner for the “meat” of the lesson.

Events 1 through 5 can be viewed as a “one-way street” going from the instruction to the learner. However, learning requires “transactions” between the learner and the instruction (Merrill, Li, & Jones, 1990b). This view sees learning as a process where the lesson information goes on many “round trips” between the instruction and the learner. For this reason, event 6 deliberately requires the learner to become an active agent in the learning process. Attention to event 6 assures that instruction will be very interactive. Giving the learner a chance to respond and interact with the lesson material is only worthwhile if the learner is then given additional information about the degree to which responses were appropriate. This is known as feedback and is identified in event 7. Feedback is an extremely important and potent instructional component and has two qualities that frequently overlap. First, feedback informs a learner to the degree of “rightness” and “wrongness” of a response to reinforce the making of more correct answers in the future. The application of feedback as reinforcement is a pillar of behavioral learning theory. The second quality of feedback, usually considered the more important of the two, is the *information* that it provides (Kulhavy, 1977). Every time a learner interacts with the lesson, there is a “window” of opportunity for feedback to provide pertinent and relevant information based on the learner's response. This window is probably widest when the learner's confidence in the answer is high, but the answer is wrong. An extreme case would be studying all night, thinking you answered a question well the next day, but then finding out your answer was wrong. You would understandably want to know why. Good informational feedback can be crucial at these times (as well as those occasions that are less dramatic). We will group events 6 and 7 together as *practice*.

Event 8 simply recognizes that there are times when assessment of learning is necessary. The purpose of testing, as it is defined here, is to judge the quality of the instruction as objectively as possible. Whereas the purpose of practice is to improve learning, event 8 is meant to assess just what learning has occurred. The major goal of event 8, therefore, is instructional accountability, although testing can and should serve other purposes as well, such as increasing motivation and providing more feedback.

Event 9, although shown last in the model, is certainly not least. In fact, event 9 is arguably the most important event of all because it describes the overall purpose of the model and perhaps most instruction as well. Event 9 should constantly be in the designer's mind because it serves as a reminder that the purpose of instruction is not only to remember what we have learned soon after we have learned it, but also later in a variety of contexts. Event 9 encompasses three important learning issues: retrieval, durability, and transfer (Clark & Voogel, 1985; Di Vesta & Rieber, 1987). It is not enough to just remember something when asked; students also should be able to retrieve it long after the lesson has ended and in situations that may not resemble the context in which it was learned. Students should not only be able to answer math questions on a worksheet, for example, but should also be able to use the information on the next shopping trip to the mall. Event 9 must be continually considered because the ability to retrieve information is thought to be largely dependent on the way in which it was initially encoded into memory. Event 9 also completes the discussion started with event 3. Just as the current lesson is related to other lessons that came before, so too will it relate to those that follow. Again, this is accomplished by considering event 9 throughout the design and implementation of the lesson. Event 9 provides much guidance and requires much vigilance.

Lastly, you probably noted that few details were given in the above discussion about particular strategies useful for accomplishing each event. This was intentional. The starting point for deciding *how* to apply these events is the identification of the learning outcome as defined by the instructional objectives. Selection of the relevant events and strategies for each chosen event depends in large measure on the nature of the learning outcome, the content, and the learners. It must be restated that although all these events must be *considered* each and every time instruction is designed, not all must necessarily be used. For example, drill and practice software would not need all of the events to accomplish its objectives. Many authors have tried to define particular instructional strategies appropriate to each event (e.g., Gagné, Wager, & Rojas, 1981; Joyce & Weil, 1980), but that is beyond our scope and purpose here. Often, one particular strategy can be useful across many strategies. A good example is questioning (Hamaker, 1986). It can help to gain and focus attention, and it can help to recall prerequisites, guide learning, and, of course, provide practice. It is also probably the most popular test of learning.

Though frequently debated (e.g., Gagné & Merrill, 1988), the identification of particular strategies is also seen by some as part of the "art" of instructional design. Our task in the next section will be to consider how graphics can be a viable part of these groups of instructional events.

Five Instructional Applications of Graphics

The following five instructional applications of graphics are offered as an informal guide to the ways graphics can be used in instruction: cosmetic, motivation, attention-gaining, presentation, and practice. These five applications come as a direct result of the discussions of learning outcomes and the events of instruction. Their purpose is to describe instructional situations in which all three graphic types can be applied. While they are listed here for the purpose of describing the role of graphics in instruction, they will be used throughout the rest of the book for the purposes of design (prescription) and evaluation. Even though these are listed in discrete fashion, their functions frequently overlap, making it possible for the instructional intent and result of any one graphic to be classified across more than one application. These applications are presented as easy-to-remember guideposts for the various uses of graphics in and out of CBI. The rationale for needing these guideposts is that it is very likely that the instructional intent of a graphic can be entirely different from the instructional result when designers make decisions to include graphics based on misinformation, misinterpreted information, or no information.

Three of the applications — attention-gaining, presentation, and practice — serve cognitive functions and two of the applications — cosmetic and motivation — serve affective functions (as shown in Figure 2.13). These functional categories should help you to design and evaluate instructional graphics based on whether the intent of a graphic is to contribute to learning or to the affective appeal of a lesson. All instructional graphic designers should ask themselves this all-important question each time they begin a project: “What function is my graphic going to serve in this lesson?

A guide to using graphics in instruction	
Instructional applications	**Function**
Cosmetic **Motivation**	**Affective**
Attention-gaining **Presentation** **Practice**	**Cognitive**

FIGURE 2.13
Five instructional applications of graphics.

Affective Functions

The purpose of cosmetic graphics and motivational graphics is to enhance the affective appeal of a lesson. Affective applications are designed to improve a student's attitude toward learning or to increase the incentive of a student to participate in the lesson.

Cosmetic Graphics. Graphics are often used for purely cosmetic reasons. In a sense, it is a misnomer to call this an instructional function, because, by definition, no direct learning benefits are expected from cosmetic graphics. The purpose of a cosmetic graphic is to merely add to the polish or decoration of a package to make a program more attractive or aesthetically pleasing (Levin, Anglin, & Carney, 1987). There are too many examples to list them all, but a few common cosmetic graphics are fancy screen borders, some uses of color, and the use of special effects (like animation at the start of a program to display a product's title and publisher). Cosmetic graphics often add a certain level of completeness or sophistication to a package. This may promote the feeling among students that the instruction is important, whether or not this is true.

At their best, cosmetic graphics help maintain student interest and perhaps regain student attention and would heavily overlap the attention-gaining and motivational functions described next. At their worst, cosmetic graphics distract student attention from other important material. An example of a cosmetic graphic is shown in Figure 2.14. Here the graphic is included in a lesson on the history of sports. Notice that the graphic has nothing to do directly with the lesson text, but merely adds visual appeal to the frame. Unfortunately, students may get the impression that the graphic is directly relevant to the text and thereby might spend time looking for learning clues in the graphic. When this happens, the graphic poses the risk of distracting the student from the intended lesson goals. Distraction is the Nemesis of instructional graphic design (e.g., Willows, 1978).

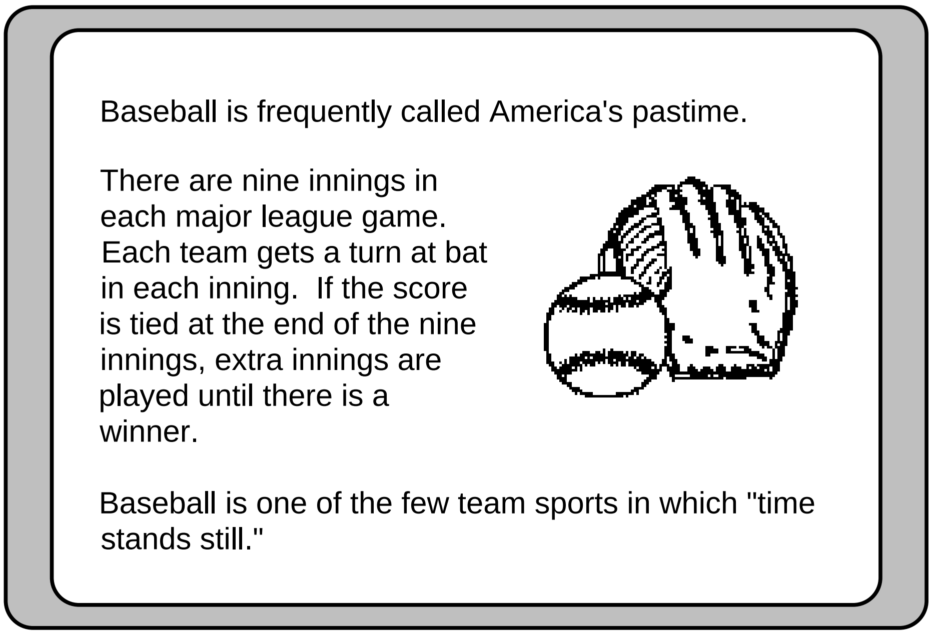

FIGURE 2.14
Snapshot of a CBI lesson using a cosmetic graphic.

Distraction effects pose threats to learning because of the severe processing limitations of short-term memory (this will be discussed in more detail in chapter 4). Therefore, anything that offers the potential of distracting students' attention from the lesson goals must be carefully evaluated. The haphazard use of cosmetic graphics is an example of where good intent can lead to unfortunate outcomes. Steps must be taken to assure that learners will not be misled into perceiving some underlying instructional value of a cosmetic graphic. The frequency and position of cosmetic graphics should be strictly controlled.

Motivational Graphics. Graphics are often incorporated into instruction to raise the general motivational level of a lesson. Much of the motivating appeal of graphics is due to novelty. Unfortunately, novelty effects are temporary, gradually disappearing over time (Clark, 1983). A good example of failing to recognize novelty effects is the early history of microcomputers in the classroom. Many believed that students naturally learned more from microcomputers because they wanted to work on them and because it was so easy to keep them on task. However, comparative reviews of media research favoring the computer over traditional instructional media were often found to be based on novelty effects (Clark, 1985). There is nothing wrong with taking advantage of novelty effects so long as one understands that the opportunity to enhance learning solely because of novelty is short-lived. As students become more familiar with computers, the prospect of interacting with one becomes less and less exciting, and hence the novelty effect disappears. The inherent instructional design of the materials delivered by computer is all that's left to influence the learner. But that is the way it should have been from the beginning.

Using graphics to arouse general curiosity and interest is seen by many as a very superficial way to increase motivation. There are deeper ways to maintain attention and interest beyond the simple provision of interesting graphics. For example, if the nature of the learning task is satisfying, relevant, and challenging, students are more likely to participate in meaningful ways (Keller & Suzuki, 1988; Kinzie, 1990; Lepper, 1985; Malone, 1981). Hence, their time on-task is not only increased, but the quality of this learning time is enhanced as well.

Professional educators frequently argue about whether instruction should contain entertainment-like qualities. We all probably agree with the two ends of this debate. Learning certainly demands effort and hard work, but instruction does not need to be boring and dull. Instruction is certainly a serious business, but it need not be grim. At what point, we must ask ourselves, do we feel that instruction is *responsible* to entertain students? Graphics can be used as one strategy to maintain motivational appeal by constantly refreshing the lesson's level of novelty and curiosity. However, the power of computer graphics as a long-term motivational tool designed to increase student perseverance does not have much empirical support (e.g., Surber & Leeder, 1988).

Cognitive Functions

Graphics that serve cognitive functions are designed to directly enhance the ability of students to learn from instructional materials. These graphics should be designed to achieve, or help achieve, one or more of the events of instruction.

Attention-Gaining Graphics. Of the three orienting events of instruction, graphics are used more often, by far, for attention-gaining — and for good reason. There are many sources of stimuli that compete for a person's attention in and out of the classroom. Many, if not most of these sources are probably far more interesting than the instruction itself. Usually, these competing stimuli come from the student's environment, such as a buzzing light, sniffling nose, screeching chairs, music or laughter from down the hallway, a growling stomach, or an attractive member of the opposite sex sitting in the next row. Competing stimuli also can come from within the student's own mind, such as personal concerns like a home crisis or just general daydreaming.

For these reasons and many more, attention-gaining is an important initial event of instruction (Gagné, 1985). Attention-gaining applications are obvious, practical, and rational uses of graphics. For example, animation can be an effective way of arousing and maintaining a learner's attention during CBI, as depicted in Figure 2.15. In this example, an animated space shuttle flies across the computer screen. The purpose of the animation in this example is not to teach something, but only to attract and focus the student's attention onto the computer screen. Hopefully, this attention will be maintained long enough to capture the student's interest in the learning material on the screen. As with cosmetic applications, graphics that purposely serve to gain attention should not subsequently distract attention from other important and salient lesson features.

Other examples include interesting special effects for transitions between instructional frames or lesson parts. Special screen washes, moving symbols or characters (cartoon or text), animated prompts, such as arrows that direct attention to key words, paragraphs, graphics, or other screen items are still other examples of animated attention-gaining devices. In addition, animated figures offer contrast to a static background, thus bringing the animated figure to prominence and allowing important lesson information to be amplified or emphasized (Hannafin & Peck, 1988). Interesting graphics contained throughout a CBI lesson can help maintain a student's attention. One reason that graphics seem to work is that they offer a degree of novelty. Attention is naturally drawn to what is new and different. Remember, however, the temporary nature of novelty effects.

Among the qualities of static graphics that increase the level of student interest is moderate to heavy richness of detail (Dwyer, 1978; Fleming, 1987). Some people may notice a contradiction with this and the principle, discussed in the last section, indicating that learning often results from representational graphics containing relatively low levels of realism. But there is a big difference between using a graphic to capture the attention of someone versus using that graphic to teach something. This is just one of many examples of the “form follows function” principle, where the type and design of a graphic must be determined by the function that the graphic is supposed to serve.

Presentation Graphics. Graphics are frequently used to teach. This application represents the main body of reported research discussed in chapters 5 and 6 (Alesandrini, 1984; Alesandrini, 1987). Graphics can be used with or without accompanying text to demonstrate or elaborate a lesson concept, rule, or procedure. The processing partnership between visual (e.g., static or animated graphics) and verbal (textual) information is the foundation of several theories of long-term memory (Bower, 1972; Paivio, 1979, 1983, 1986) and is discussed in more detail in chapter 4. The use of static and animated graphics as a presentation device has been called a “learning-by-viewing approach” to instructional graphics (Reed, 1985). Most instructional applications of this type use representational graphics to directly visually depict critical information and are probably the most common way pictures are used to help students learn from text (Levin, Anglin, & Carney, 1987).

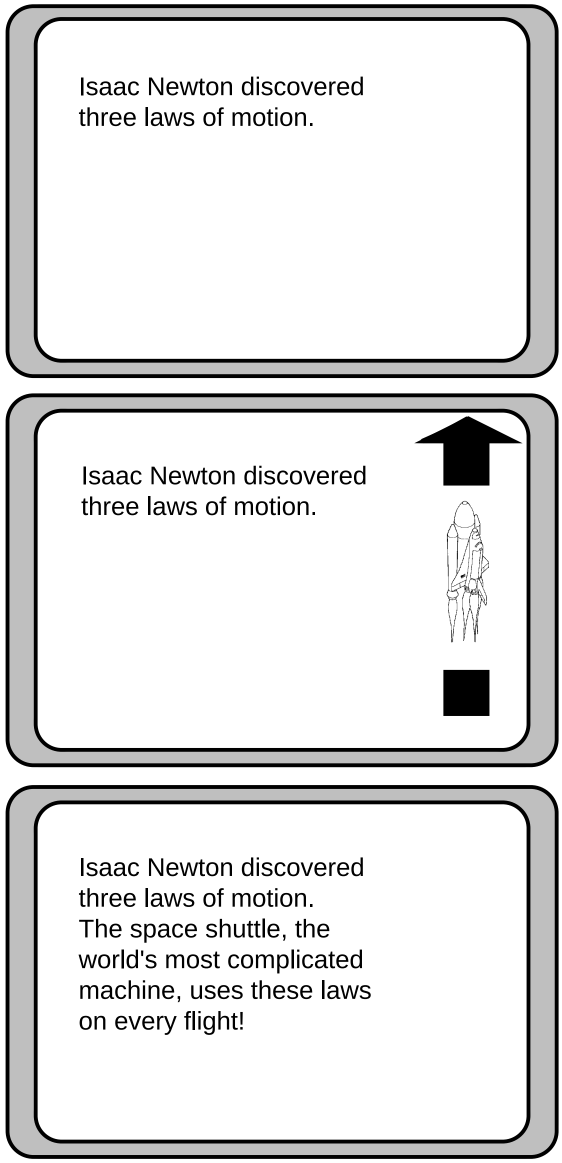

FIGURE 2.15
Snapshot of a CBI lesson using animation as an attention-gaining device.

Another example of a presentation graphic is illustrated in Figure 2.2 in a lesson describing the functions of an astronaut's space suit. Again, the graphic and the text share a special relationship. The graphic provides a visual elaboration of the information contained in the text. However, this graphic could be greatly improved with the use of labels to highlight each component of the suit. A CBI lesson could easily be designed to precisely indicate each part of the suit, either on separate frames or interactively on one frame.

An example of using animated graphics as a presentation tool is shown in Figure 2.16. Here, animated graphics demonstrate an application of the laws of motion. Animated presentations can also aid a student's conceptual understanding of interrelated lesson variables. In this way, presentation graphics help students to interpret difficult-to-understand information. The use of visual analogies, such as a graphic of a mechanical pump to help describe the difference between systolic and diastolic blood pressure (Levin, Anglin, & Carney, 1987), is a common strategy. Another good example of how graphics can help a student's conceptual understanding is the program shown in Figure 2.17. The goal is for students to try to create or replicate a given “product” (Kosel & Fish, 1983). The product starts out in raw material form as a square wafer. The raw material is sent through any number of combinations of punch, stripe, and rotation “machines” chosen by the student. These machines successively alter the product into its final shape. Animation helps students understand how the changes to the product occur in intermediate phases and how the order of the changes affects the final outcome. Children, or other novices, would be expected to benefit from these kinds of animated displays, since they probably would have difficulty in visualizing abstract relationships on their own.

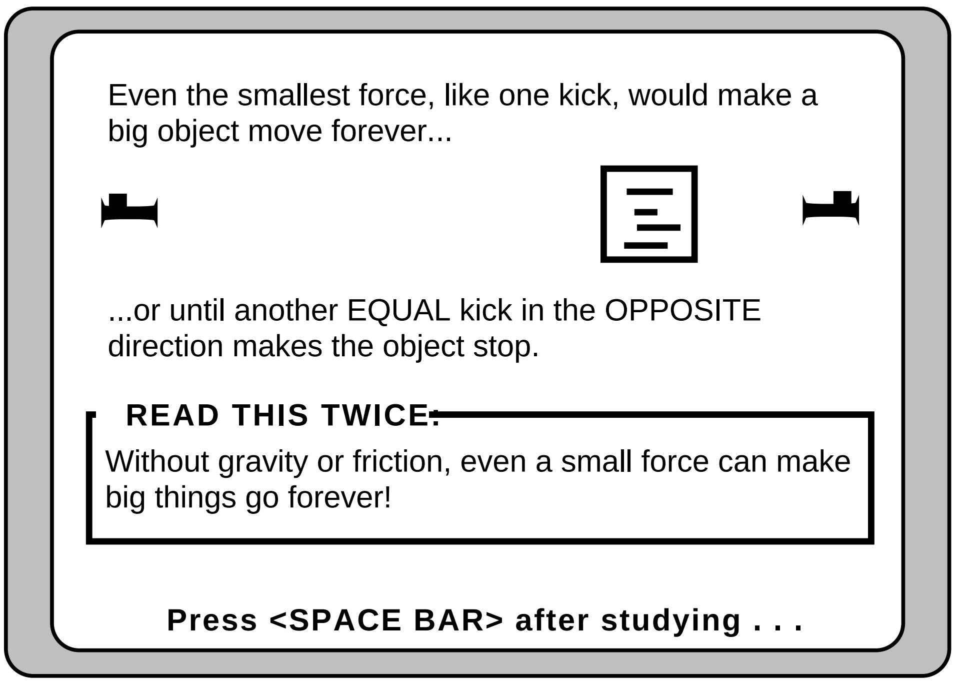

FIGURE 2.16
Snapshot of a CBI lesson using an animated presentation graphic. The block moves slowly across the screen after being kicked.

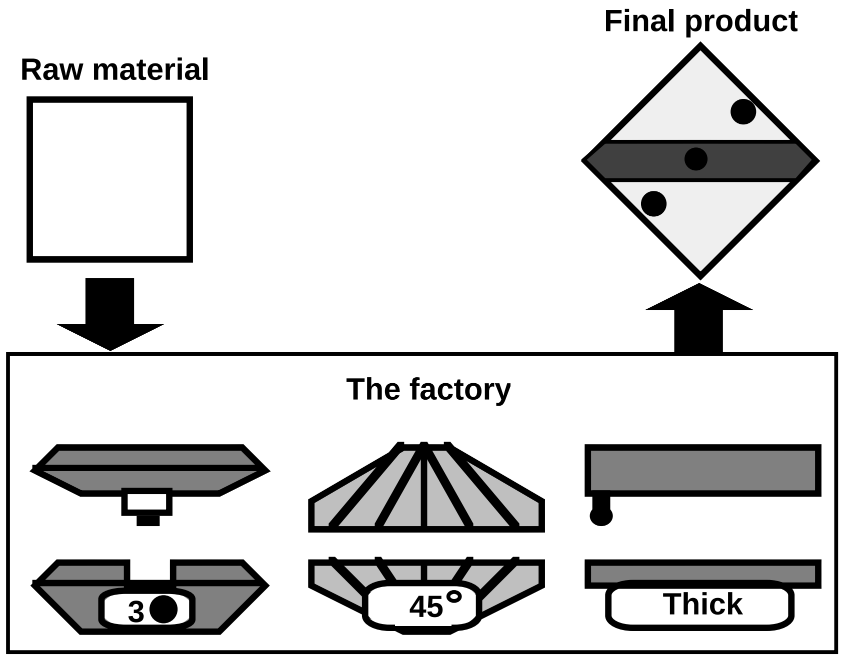

FIGURE 2.17
Using computer animation, the raw material get transformed into the final product by successively going through the three machines in the “factory”.

Representational graphics are an effective presentation strategy when combined with text. The graphics help learners focus their attention on the explanative information in the text (Mayer, 1989). Graphics also help learners form visual mental models of the materials explained by the text. Mayer and Gallini (1990) suggest that visuals are useful presentation strategies when they satisfy four conditions: 1) the text is *potentially* understandable by students; 2) the visuals are designed and evaluated in terms of learner understanding; 3) the visuals are used to *explain* information provided by text; and 4) students have little or no previous experience with the content.

Finally, presentation graphics can also serve an organizational function (Levin, Anglin, & Carney, 1987) to help make relationships between ideas more apparent. The most common examples of these are “how-to-do-it” graphics that show a set of step-by-step procedures in visual form. Examples include how to assemble a household device or how to perform an emergency medical procedure, such as cardiopulmonary resuscitation (CPR). Such procedural applications of graphics are very relevant for many psychomotor tasks, as shown in Figure 2.18.

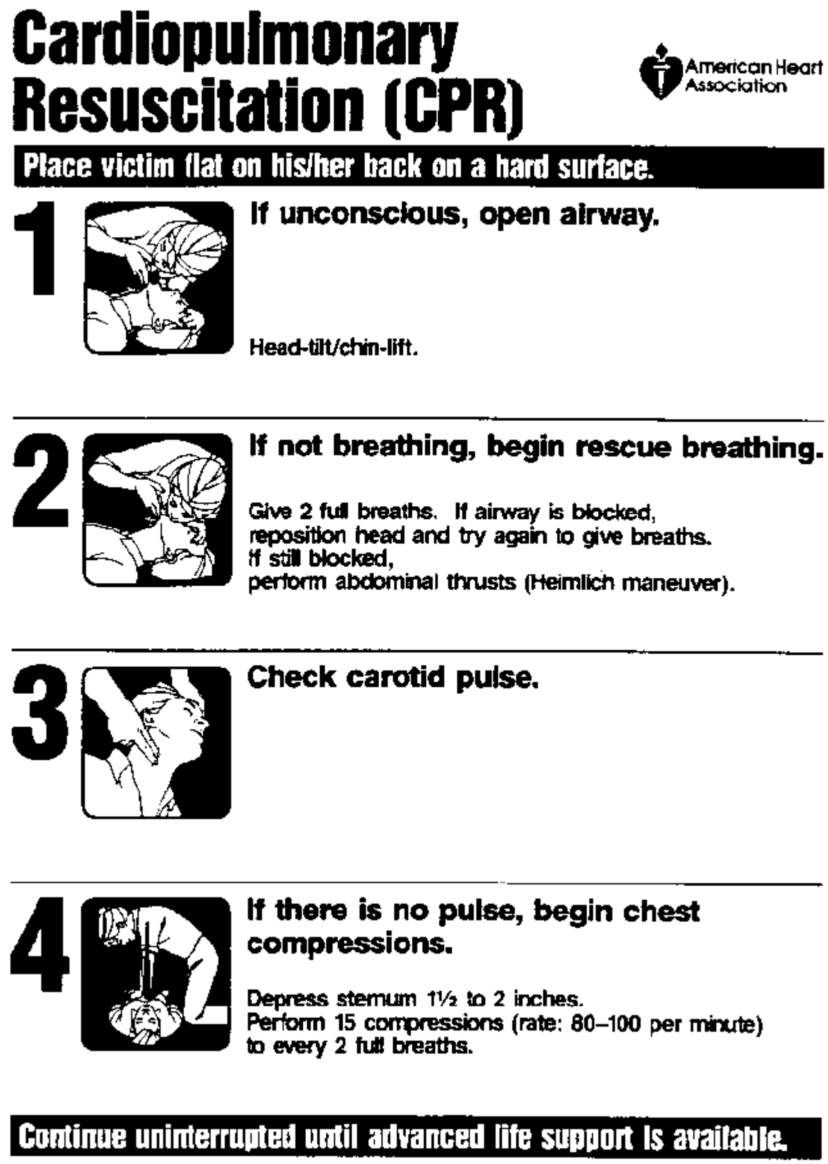

FIGURE 2.18

An example of a procedural graphic that illustrates a step-by-step sequence of tasks. Reproduced with permission. "Cardiopulmonary Resuscitation CPR Wall Chart," 1986 ©1986 American Heart Association.

Graphics in Practice Activities. Graphics can be very useful in practice activities. Graphics can act as visual feedback to students as they interact with lesson ideas and concepts. This application of graphics is particularly suited to the computer medium, such as those involving visually based simulations. Real-time animated graphics in interactive learning displays are also known as “interactive dynamics” (Brown, 1983). Real-time animated graphics change continuously over time, depending on student input. Students learn in these highly interactive visual environments by discovery and informal hypothesis-testing. Graphics act as instantaneous feedback. Brown (1983) called this application of animation “learning by doing.” Examples include graphic, real-time simulations, such as piloting an airplane as shown in Figure 2.19, interacting with a Newtonian particle in a gravity-free/frictionless environment (diSessa, 1982; Rieber, 1990b; White, 1984), and graphic programs where students learn musical concepts (Lamb, 1982). Other examples include graphic programming procedures in LOGO, where students drive an animated “turtle” (Papert, 1980). However, students must be able to perceive differences in the graphic feedback, an ability that novices especially have a difficult time attaining (Brown, 1983; Cohen, 1988; White, 1984). Interactive dynamics should be structured to offset this deficiency (White, 1984; Rieber, 1989) or to augment such interactions with coaching or other prompts (Reed, 1985). Learning from interactive dynamics appears very contextually bound. This use of animation is not easily replicated with media other than the computer.

Graphics are commonly used in more traditional practice activities in CBI, such as question and answer. Often, the role of graphics is merely to reinforce correct responses, such as displaying a happy face for right answers. The danger is that attractive and interesting graphics may actually reinforce wrong responses or other behaviors. Some computer chess games, for example, visually personify the chess pieces. When one piece “takes” another, some programs actually show the “execution” of the captured piece. A person who finds such visuals motivating might *want* to lose the game just to witness the graphical results.

A simple guessing game illustrates the value of different types of graphic feedback. Most people, at some time in their lives, have played a guessing game where player 1 thinks of a number from 1 to 100 and player 2 tries to guess what it is. The only clues given to player 2 are whether the guesses are too high or too low. By considering each clue, player 2 should be able to quickly narrow down the range in which the mystery number falls, and then finally pinpoint the exact number. The players reverse roles, and the one who took the fewest number of guesses is the winner of the round.

It is easy to construct this mystery number game on the computer, where the computer assumes the role of player 1. The computer chooses a number at random and tallies the number of guesses until the mystery number is discovered. For each incorrect guess, the computer displays the clue “too high” or “too low.” A slight modification to the game can be made to involve visuals. The simplest modification would be displaying the prompt to be spatially congruent to the clue, such as displaying “too high” at the top of the screen and “too low” at the bottom of the screen. However, a more intriguing modification would be to display the clues relative to the mystery number along a vertical or horizontal axis. In this way, the student would not only see in which direction the guess is wrong, but by how much.

FIGURE 2.19
Snapshot of a computer flight simulator as the user tries to land the airplane. The graphics change continuously in real-time depending on the user input of the various controls.

In another variety of the game, the clues change from “too high” or “too low” to “hot” or “cold,” as shown in Figure 2.20. Very inaccurate guesses are “freezing,” but get “hotter and hotter” as the student's accuracy improves. The purpose of the graphics in each of these examples is to provide visual feedback to students based on their guesses. Secondarily, the graphics help make the game more entertaining as well. The graphic feedback in Figure 2.20 can be easily improved further by adding some cultural conventions. The graphic could be rotated to match the convention that thermometers are usually oriented vertically with hot always toward the top. This is an obvious place for color as well because red is a standard in western cultures for hot and blue for cold.

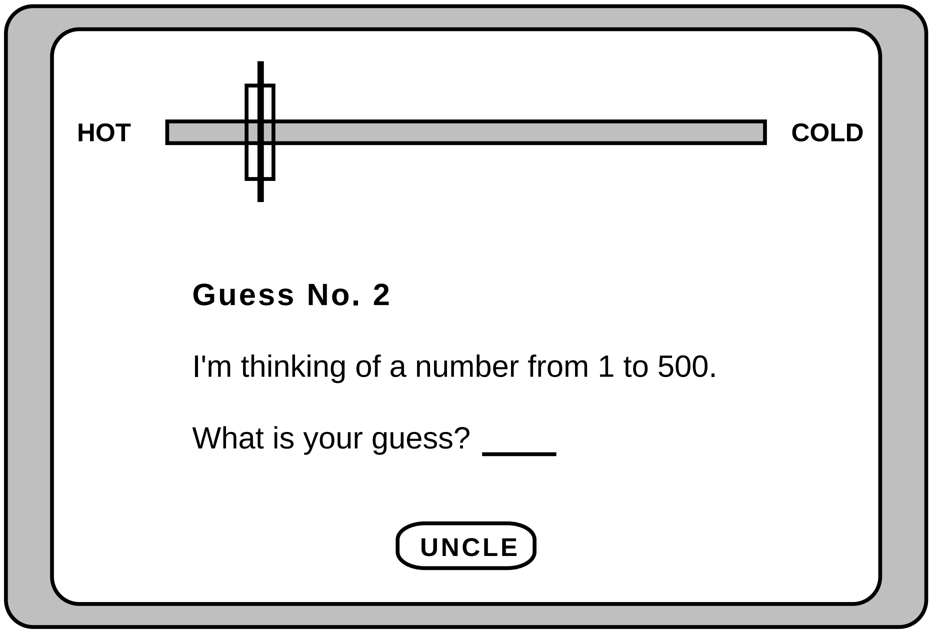

FIGURE 2.20
An example of poviding visual feedback.

REVIEW

- Representational, analogical, and arbitrary represent the three major groups or types of graphics under which most instructional graphics can be classified.
- Each of the three types of graphics convey information in different ways
- Instructional materials, such as those involving graphics, should be selected or designed to fulfill the instructional objectives.
- Graphics are typically used to serve three cognitive functions: attention-gaining, presentation, and practice.
- Graphics are typically used to serve one of two affective functions: cosmetic and motivation.
- Graphics should be included in instruction only on the basis of the instructional function that each serves.
- When graphics are used to increase student motivation or interest, care must be taken to assure that the graphics do not interfere with any cognitive functions served by other lesson components. This interference is known as a distraction effect.

NOTES

1. This example was originally described by Kathryn Alesandrini in her address at the annual meeting of the Association of the Development of Computer-based Instructional Systems (ADCIS) in Oakland, CA, in 1987.

2. It is easy to question whether some visuals, like those describing the “lift and cut” process of an electric razor, are really disguising these moods or feelings under the pretense of product information by implying the association of “sophistication” or “state-of-the-art technology” with the product.

CHAPTER 3

Developing Instructional Computer Graphics on Microcomputers

OVERVIEW

This chapter presents a conceptual overview for the production of computer graphics for instruction. The goal is not to teach the use of any one particular graphics application, but rather to provide an organizer for the different approaches commonly found on microcomputer systems and to give a sense of what features and effects currently exist. The development of static and animated computer graphics are considered separately. The chapter also introduces the concept of "second-hand" graphics, including those produced from scanning print-based pictures or capturing video snapshots and then converting the images into digital form for later use on the computer.

OBJECTIVES

Comprehension

After reading this chapter, you should be able to:

1. Describe the relationship between instructional design and instructional development.
2. List the three graphic primitives.
3. Explain the difference between raster and vector graphics displays
4. Describe differences in producing static and animated graphics displays using command-based and GUI-based approaches.
5. List some of the features common in GUI-based graphics applications.
6. Explain some of the procedures involved in scanning and digitizing analog pictures.
7. Describe the difference between fixed-path and data-driven animation.
8. Describe the differences and implications between graphics stored algorithmically as computer programs, **paint** files (bitmaps), **drawings** (object-oriented files), and generic **pict** files.

Application

After reading this chapter, you should be able to:

1. Produce simple static and animated graphics using graphics commands from a programming language.
2. Produce simple static graphics using a GUI-based graphics application.
3. Produce a simple fixed-path animation sequence using command-based and GUI-based graphics approaches.

4. Classify a given graphics application as a command-based, GUI-based, or scanning/digitized approach.

This chapter deals with issues surrounding the development of computer graphics for instruction. The differences and relationships between design and development are analogous to those between the blueprint and construction of a house (Reigeluth, 1983b). Design proposes instruction and describes its specifications, often in great levels of detail. Development concerns the actual production or "construction" of the design. Changes to the design are easy and relatively inexpensive. Changes to the development can be costly and time-consuming. This is not to suggest that design and development are mutually exclusive. In fact, some media, like computers, often permit an instructional materials production cycle where design and development are intertwined. This approach, called rapid prototyping, is discussed in more detail in chapter 7. At the very least, designers must consider the resources, conditions, and constraints of development.

Instructional designers, like their counterparts in architecture, must consider many tradeoffs and compromises throughout the design and development phases. Instructional designers must carefully consider how decisions affect groups of instructional variables, just as the decision to use two-by-four-inch lumber, two-by-six-inch lumber, or fabricated metal incurs tradeoffs among cost, strength, and installation time.

Instructional designers and developers often talk about the relationship between instructional effectiveness and efficiency. No instructional design can ever hope to be perfect in every respect. Consider the situation of learners achieving 85%, instead of 95%, of the "goals" of the instruction. A designer must ask whether it is worth the time and cost to revise the instruction to attain the extra 10%. Obviously, the answer is based on the context and will be different, for example, for the training of medical personnel on emergency room procedures versus elementary instruction on art appreciation. Of course, saving time and money is meaningless if the instruction, like the house, falls apart after it is built. This is a little like buying a pair of pants that are the wrong size just because they are on sale. The hope is to maximize effectiveness while minimizing cost, design time, development time, and instructional time (i.e., the time required by a learner to complete the instruction).

The purpose of this chapter is not to provide instruction on the "how to's" of computer graphics applications. Teaching how to use even a small number of specific commercial graphics applications is not a goal of this book. The rapid rate at which new commercial graphic software packages are introduced, combined with the fickle nature of developers and users, would make this chapter obsolete before the ink dried. Instead, the goal of this chapter is to provide a brief *conceptual* overview of past, present, and (hopefully) future approaches to developing computer graphics. It is important to note that this chapter will focus entirely on graphics produced from microcomputer systems.

Since this is a book concerned with the *design* of computer graphics for instruction, one might question why development is even considered at all. There are several reasons. As already mentioned, design and development are not independent of one another. Knowledge and sensitivity about development issues influence design decisions. For example, although design may call for an animated sequence, knowledge about the capability and cost (in terms of money *and* time) of the particular microcomputer system will influence how elaborate the sequence can be, as well as the possibility of changing the design altogether, such as to a sequence of multiple static graphics. A final reason to consider development is simply to gain some sensitivity and appreciation for the rapid advancement in graphics production on microcomputers.

HARDWARE SYSTEMS: TYPES OF COMPUTER GRAPHICS DISPLAYS

Regardless of how a graphic is produced on a computer, the end result will be either a computer display of the graphic or a computer file that stores information about the graphic, or both (Artwick, 1985). There are three fundamental structures, known as **graphic primitives**, that act as the building blocks for all computer graphics (Pokorny & Gerald, 1989). The first is the picture element, or **pixel**, which is simply a single point of light on the computer screen. The next two elements are the **line** and the **polygon**. All computer-generated images can be created from these three graphic primitives.

Essentially there are two major kinds of computer graphics display systems that can produce pixels, lines, and polygons: raster graphics displays and vector graphics displays (Conrac Corporation, 1985). Each uses a fundamentally different hardware approach to controlling the scan rate and pattern of the cathode-ray tube (CRT) to display the graphic. Vector graphics displays, as the name implies, use vectors to define lines, which, in turn, comprise polygons. A vector is a mathematical entity comprised of two or more elements. A line, for example, is defined on a vector display in terms of its magnitude (e.g., length) and direction. Diagonal lines on vector graphics displays are true diagonals and do not suffer from the "jaggies" associated with the more common raster graphics displays. However, virtually all desktop computer systems use raster graphic displays.

Raster graphics are formed by a pattern of pixels on the computer screen. A single graphic consists of a matrix of on or off pixels (or "0" or "1," either of which defines a "bit" of information on a computer). For this reason, raster graphics are sometimes referred to as bit-mapped displays. One can more easily understand bit-mapped graphics by imagining that the computer screen is a matrix of tiny light bulbs. In order to draw the letter "H," the computer must be told which light bulbs (pixels) should be off or on, such as shown in Figure 3.1. Information related to shades of gray or color could also be stored in relation to each pixel.

In general, displaying a black and white bit-mapped graphic consumes the same amount of computer memory, regardless of the simplicity or complexity of the graphic display, because it takes the same amount of memory to store information about each pixel whether it is off or on. (See Footnote 1) Scanning and digitizing processes transform analog images, such as line drawings and photographs, into the digital form of bit-mapped graphics. An everyday example of graphics produced by the combination of dots, apart from computers, are photographs in newspapers. The continuous tones of shading in a photograph must be broken down into tiny dots of varying size and intensity. The resulting image, called a halftone, is printed in the newspaper as a reproduction of this configuration of dots. From a distance, the human eye perceives continuous shades of gray from the discrete collection of dots.

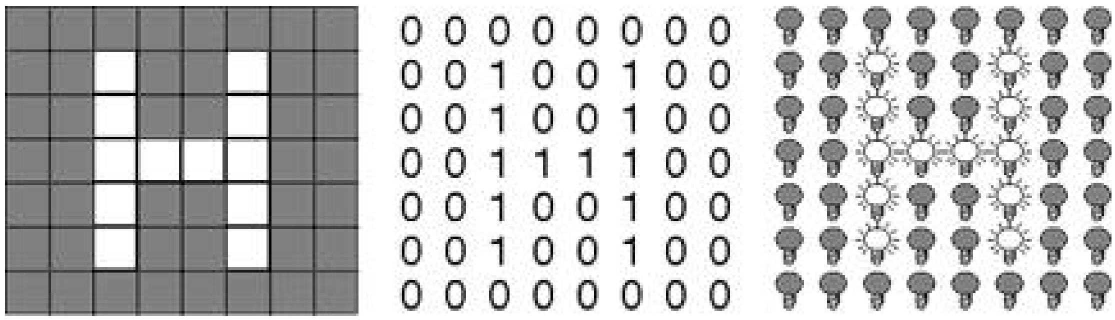

FIGURE 3.1
Raster graphics are formed by a pattern of "pixels" on the computer screen. A "bit" map of a graphic, such as the letter "H," is analogous to a pattern of lit light bulbs, where "1" means the light bulb is on and "0" means the light bulb is off.

The clarity and sharpness of a bit-mapped graphic, known as **resolution**, depends on the number of pixels contained in a certain display area. Low-resolution displays offer crude graphic representations because the smallest point of light that can be manipulated is quite large. Increasing the resolution means increasing the number of rows and columns in the graphic matrix to produce a greater number of smaller and smaller pixels in a given area, such as that shown in Figure 3.2. High resolution is a function of the level of detail produced by the pixel size and is therefore relative to the capability of the computer hardware.

A diagonal line presents problems on raster displays because of the row and column orientation of the pixels. Diagonals often have a jagged look resembling a staircase, such as the exaggerated example in Figure 3.3. This effect is minimized as the resolution of the display is increased. A high-end software technique, called antialiasing, also can be used to minimize the jagged effect by averaging the shading of pixels adjacent to the diagonal.

PRODUCING STATIC COMPUTER GRAPHICS

This section presents a brief conceptual overview of the software approaches to produce computer graphics on microcomputer systems and how these graphics can be electronically stored for future use.

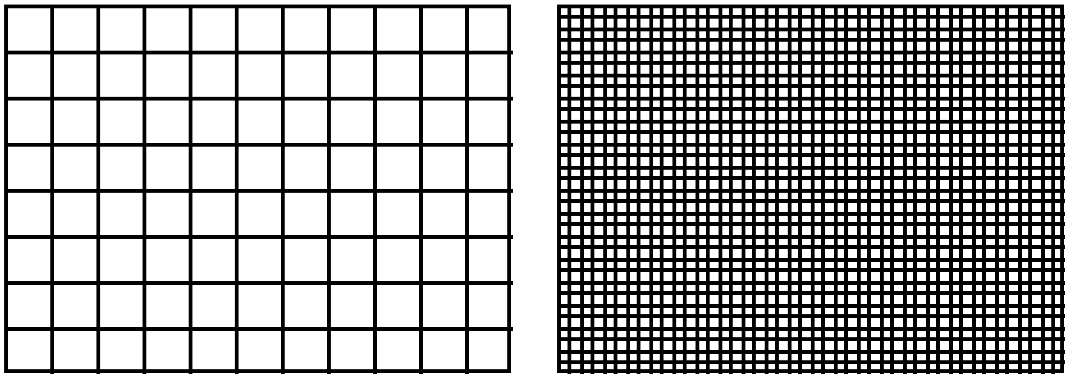

FIGURE 3.2
A comparison of low-resolution (left) and high-resolution (right) display screens.

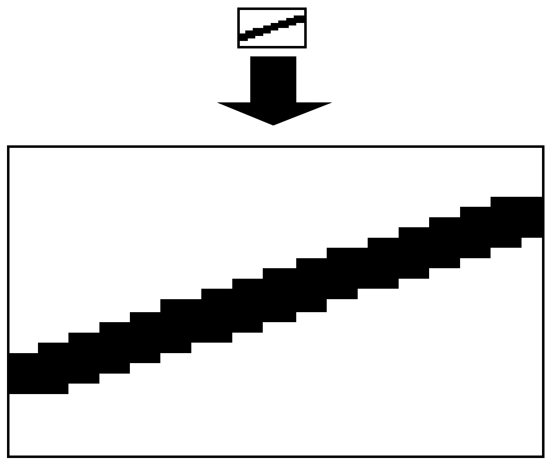

FIGURE 3.3
Diagonals displayed on raster systems are prone to the "jaggies." In this example, a diagonal is enlarged many times to show the way individual pixels are "staircased." The greater the screen resolution, the less the distortion.

One can produce computer graphics on microcomputer systems in essentially one of three ways:

1. Command-based approach
2. GUI-based approach
3. Use of "second-hand" graphics (clip art, scanned/digitized images, etc.)

The *command-based approach* involves algorithmic processes for defining a graphic, such as the writing of programming code using special graphics commands particular to the programming language (e.g., PASCAL, C, BASIC, LOGO).

The *GUI-based approach* is based on the graphical user interface discussed in chapter 1 and involves graphic tools such as "pencils," "brushes," "fill buckets," "box makers," etc. GUI-based approaches most commonly use input devices such as a mouse, light pens, and graphic tablets, although some of the earliest GUI-based approaches used the keyboard. The GUI-based approach is now the status quo on most microcomputer systems. Some authoring environments, both old (e.g., PILOT) and new (e.g., HyperCard, Authorware), offer a combination of command-based and GUI-based approaches.

Second-hand graphics are copied, not created, and include clip art and all of the scanning and digitizing technology (both hardware and software). For example, graphics can be drawn on paper, converted into digital form with a scanner, and then "imported" into one of many computer applications. This approach also includes the electronic "capture" of video and photographic images.

Each of these three approaches will be briefly described, after an overview of various formats in which graphics can be stored on disk.

Overview of Graphic File Formats

Command-based, GUI-based, and scanned/digitized graphics can be stored in a variety of formats on a computer disk (floppy, hard, optical, or compact), as listed in Table 3.1. The format of a stored graphic image directly affects the way it can be used or revised later. For example, when computer graphics are stored as bit-mapped images, only the "on/off" pixel pattern is saved. Bit-mapped files are commonly known in some systems as **paint** files. **TIFF** (Tagged Image File Format) files also store bit-map graphics, but can include additional grayscale and color information.

Other formats allow a graphic to be stored on disk as a collection of one or more individually defined and editable "objects." (See Footnote 2) These files are sometimes referred to as **drawings**. Instead of storing the actual graphic as a bit map, the visual attributes of the graphic are stored as a group of mathematically defined objects. In this way, the graphic is simply redrawn by the computer every time it is retrieved from disk to random-access memory (RAM). Almost all of the latest graphics packages that use the GUI-based approach store the graphic in this way. Most graphics packages save the drawing in a format that is specific to the software. In addition, some graphical programs allow object-oriented drawings to be converted to bit-mapped paint files. Once converted to a bit map, individual information for each of the drawing's objects is lost.

Table 3.1 Ways to Save or Store Computer Graphics

- Graphic file formats
  Paint file — graphic stored as bitmap
  Drawing — graphic stored as a group of one or more editable objects
  Pict file — generic graphic file format for both bitmap and object-oriented graphics

- Computer program — a "verbal description" of the graphic using a set of executable graphic commands; the graphic is redrawn each time the program is run

Several generic formats have been created to allow a graphic to be stored and imported into a wide range of applications. One common format, called a **pict** (for picture) file, can store both bit-mapped graphics or object-oriented drawings. Most major brand-name graphics software packages can both read and save pict files, allowing for easy swapping of graphics from one application to another. Many word processors, desktop publishing and presentation packages, and authoring packages are *only* able to import pict files.

There are a variety of specific file formats available, depending on what hardware and software are used. Readers can take both warning and solace in knowing that the issue of multiple-storage formats often confuses and bewilders novice computer users, even though it is not really that complicated. In order for a graphic to be used in an application different than the one in which it was created is analogous to home plumbing problems, such as trying to connect 1/2-inch and 5/8-inch pipes. In order to get a graphic from "here to there," some kind of "adapter" must be used as an in-between step, such as first converting one specific drawing format to a pict file before importing it into another application. This is a good example of an idea more easily understood by actually working with a computer than reading about it in a book.

Computer graphics also may be stored as a computer program. This is not considered a graphic file format because the graphic itself is not stored, only the *idea* of the graphic as represented in the all verbal form of the computer program. Most common programming languages on microcomputers have graphic commands and functions. The next section elaborates on this idea.

Command-Based Approaches to Producing Static Computer Graphics

Long gone are the days when computer company executives argued about whether to include upper- *and* lowercase letters on their computer terminal displays. When microcomputers became readily available in the late 1970s and early 1980s, the main way to produce graphics was to master one of several programming environments. True hackers learned low-level programming languages, such as assembler and machine code. However, the opportunity to access low- and high-resolution graphics "pages" through graphic commands in high-level programming languages, such as BASIC, LOGO, and PASCAL, made graphics a popular and easy (relative to machine code) context for programming projects. Command-based approaches are typically based on a Cartesian coordinate system. However, another creative, though less well-known system, called turtle geometry, also has been used.

Cartesian Coordinate System

Anyone who has played "Bingo" or "Battleship" is already well acquainted with the idea of dividing a flat space into a grid system in such a way as to precisely define a specific location on a two-dimensional surface. A system based on Cartesian coordinates (named after French mathematician René Descartes) is simply a formal mathematical approach for doing the same thing. Figure 3.4 shows a typical example of such a system, which would be recognized by any first-year high school algebra student. In a two-dimensional plane, a set of Cartesian coordinates consists of two numbers, each separated by a comma. The first number refers to the horizontal (or X) axis, and the second number refers to the vertical (or Y) axis. A traditional Cartesian system is separated into four polar (i.e., positive or negative) quadrants, depending on values of each of the coordinates. Most computers usually only use the positive/positive quadrant, although it is often modified such that it is "flipped," making the point of origin (0,0) reside in the top left corner of the screen, as illustrated in Figure 3.5. Users simply need to remember that the screen is divided into a matrix of rows and columns starting with point 0,0 and extending as far as the resolution of the particular system permits (e.g., 512, 342 in the case of the standard Macintosh screen).

Drawing screen graphics is really just a matter of playing "connect the dots" through a variety of graphical programming commands. Although programming languages vary in their names of specific commands, the fundamental functions of these commands are very similar across languages. In order to draw a box, for example, one must draw a series of four lines from, say, point 50,50 to 100,50 to 100,100 to 50,100, and back again to 50,50. (Figure 3.6 shows two small programs that accomplish this task.) Of course, some programming languages will have more commands covering a wider range of graphic functions than others. For example, many languages have special commands that allow simple objects, such as boxes and ovals, to be drawn more quickly by defining the objects in terms of its diagonal, as shown in Figure 3.7.

The graphic could be saved either as the program consisting of the series of graphic commands or as a bit-mapped image. If saved as a program, the visual representation of the graphic itself is not stored, only the algorithm, which, when run, produces the graphic. In terms of computer memory, the size of the stored file depends directly on the length of the program, not on the complexity of the graphic image. Our box example would only require the storage of a few simple lines. Saving it as a bit-mapped image, on the other hand, saves only the visual representation currently appearing on the computer screen. This bit-mapped image can be recalled later, but the computer has no way of knowing how it was created. The computer merely retrieves and reconstructs the matrix of on/off pixels. The memory to store a graphic as a bit-mapped image is the same whether the graphic was actually produced by one line or 1,000 lines of programming code.

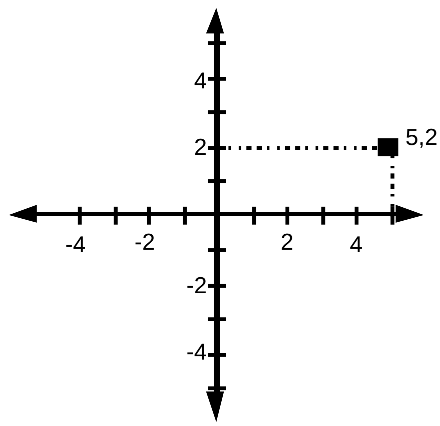

FIGURE 3.4
An example of a Cartesian coordinate system.

A more interesting example of "connect the dots" strategies concerns graphics that contain curves, such as circles, ellipses, and arcs. As already mentioned, many languages include "oval makers" to produce the oval represented within the perimeter of an imaginary rectangle. Lacking such a functional tool, we are left with the task of defining our own circle as a collection of connected dots. Since a circle is mathematically defined as an infinite number of points equidistant from its center, compromising is essential. Rather than draw a true circle, a polygon can be constructed to represent a circle — the more sides, the better the representation. We can either manually decide which pixels will be connected or we can use a mathematical model of a circle and let the computer compute the dots. Box 3.1 demonstrates two examples of the latter approach, using modifications of the traditional formula for defining a circle with Cartesian coordinates and another based on trigonometric functions. This is a simple example of using a pure mathematical model for driving the production of computer graphics and is essentially the same concept used in the most sophisticated computer-assisted design (CAD) systems.

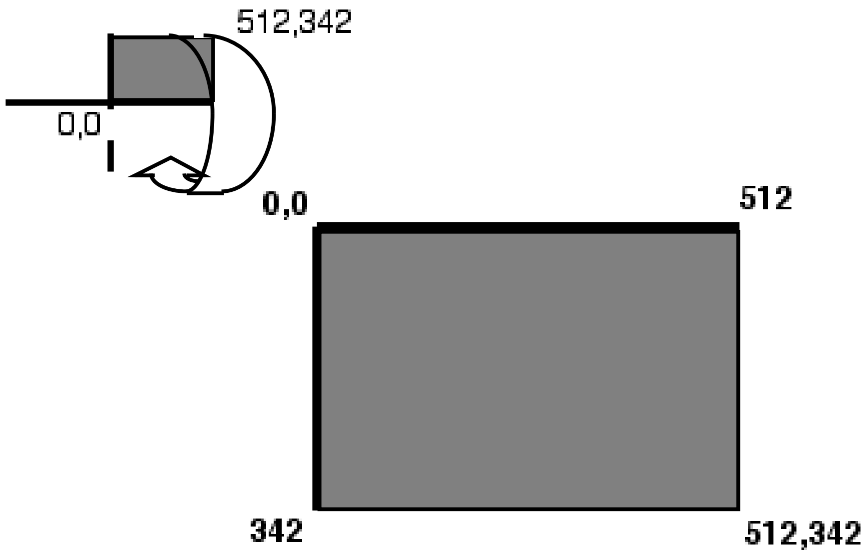

FIGURE 3.5
The coordinate system as used on the standard Apple Macintosh screen. This system is based on vertically "flipping" the traditional positive quadrant.

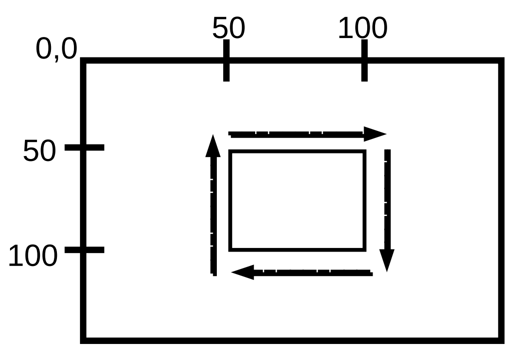

BASIC
HPLOT 50,50 TO 100,50
HPLOT 100,50 TO 100,100
HPLOT 100,100 TO 50,100
HPLOT 50,100 TO 50,50
HyperTalk
CHOOSE LINE TOOL
DRAG FROM 50,50 TO 100,50
DRAG FROM 100,50 TO 100,100
DRAG FROM 100,100 TO 50,100
DRAG FROM 50,100 TO 50,50

FIGURE 3.6
Two programs that draw a box by connecting lines from each of the box's corners.

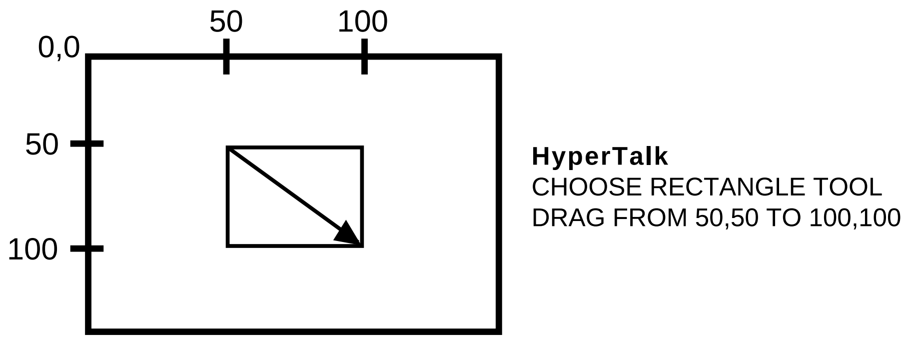

FIGURE 3.7
A program that draws a box as defined by its diagonal.

Turtle Graphics

Although command-based approaches to producing computer graphics are usually based on the Cartesian coordinate system, there is one other notable and unique way to do it. This approach, called **turtle graphics**, was initially developed for use with the LOGO programming language and was founded on learning principles associated with the *process* of creating the graphic, and not on the product itself (i.e., the resulting graphic image) (Abelson & diSessa, 1981; Lockard, Abrams, & Many, 1990; Papert, 1980). (See Footnote 3) In other words, the goal was to find a way for people of varying ages and ability to think and communicate about geometry without using the rather cryptic (and often meaningless) method associated with Cartesian systems. The LOGO language capitalized on the use of graphics to allow users, even young children, to gain access to powerful ideas associated with mathematics and computers. LOGO is often misinterpreted as a "toy" language, but in reality it is a sophisticated procedural programming language. Turtle graphics is but one of many "microworlds" users can explore in LOGO.

As the name implies, users create graphics by manipulating a graphic object, called a turtle, on the computer screen. As users "drive" the turtle, it leaves a trail. The turtle has vector-like qualities in that it has two characteristics — position and heading. Turtle graphics is a fundamentally different approach to mathematics when compared to the more common and traditional approach based on Cartesian coordinate systems. Cartesian systems define a figure, such as a circle, by its relative position to a set of points outside of the figure, such as a perpendicular axes, and an Euclidean system defines it in relation to one inside point, its center. Turtle graphics, on the other hand, defines the figure in relation to the relative position of the turtle on the figure itself. For this reason, turtle geometry is based on differential or "intrinsic" mathematics. Whereas Cartesian systems define graphics on the basis of fixed, absolute points, turtle graphics are drawn by commands that are relative to each other. The movement of one turtle graphic command is always in relation to its position and heading immediately before the command's execution. For example, the command FORWARD 50 will draw a line in whatever direction the turtle is pointing.

Some of the most common turtle graphic commands, called **primitives**, are shown in Figure 3.8. In order to draw a box, for example, a series of FORWARD and RIGHT commands must be executed, as shown in Figure 3.9. Some very interesting and powerful mathematical ideas can be expressed and explored with just this small list of commands. To draw a circle, the idea of "move a little, turn a little" is repeated until the turtle has made a complete "round trip" and arrives back at its original position with its original heading, also shown in Figure 3.9.

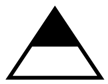

Turtle geometry

Some common commands:

FORWARD	PENUP
BACK	PENDOWN
RIGHT	HIDETURTLE
LEFT	SHOWTURTLE

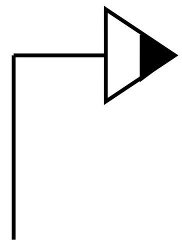

FORWARD 100
RT 90
FORWARD 50

FIGURE 3.8
Some common commands of turtle geometry.

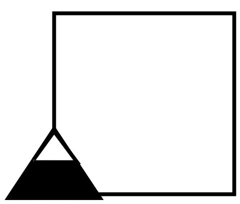

FORWARD 100
RIGHT 90
FORWARD 100
RIGHT 90
FORWARD 100
RIGHT 90
FORWARD 100
RIGHT 90
and
REPEAT 4 [FORWARD 100 RIGHT 50]

REPEAT 36 [FORWARD 5 RIGHT 10]

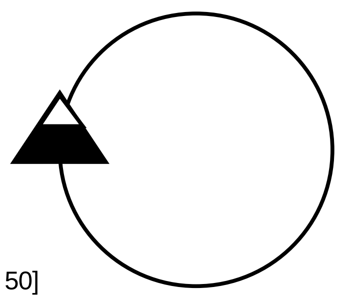

FIGURE 3.9
Two sample turtle geometry programs that draw a square and one that draws a circle.

Box 3.1

Drawing Circles the Hard Way

Here are two examples in which circles are drawn mathematically. The first example is based on the traditional circle formula:

$$(x-h)^2 + (y-k)^2=r^2$$

where h,k are the coordinates of the center of the circle, r is the radius, and x,y are the coordinates of any one position on the circle itself.

The second way defines the x,y position on the circle using the trigonometric functions of sine and cosine.

Both examples are presented using HyperTalk, the language of HyperCard on the Apple Macintosh (also known as scripting). The scripting on the following page corresponds to each of the two "buttons" shown on the "card" below:

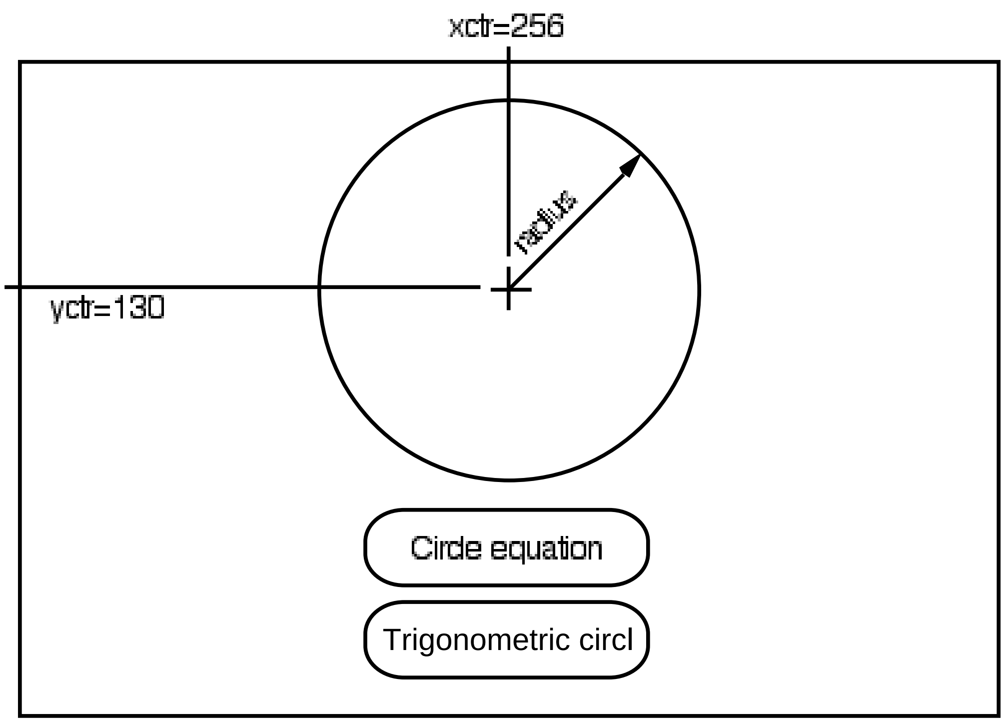

One positive consequence of this approach, assuming you are successful, is that you will *really* understand what the mathematics of the formulas mean. Chapter 8 will discuss this issue of empowering students with tools, such as the computer, to help them understand the process of mathematics and science.

Here is a legend for the major variables in each of the two scripts that follow:

xctr— horizontal coordinate of the center of the circle
yctr— vertical coordinate of the center of the circle
rad — the radius of the circle
x — horizontal coordinate of any one position on the circle
y — vertical coordinate of any one position on the circle

Script of card button "circle equation:"

```
on mouseUp
  choose line tool
  global xctr,yctr,rad,a,b
  put 256 into xctr
  put 130 into yctr
  put 100 into rad
  put 156 into x
  put xctr-rad into a
  put yctr into b
  repeat until x>355
    put (sqrt((rad^2)-((x-xctr)^2)))-yctr into y
    put the abs of round (y) into y
    drag from a,b to x,y
    put x into a
    put y into b
    add 5 to x
  end repeat
  repeat until x<156
    put (sqrt((rad^2)-((x-xctr)^2)))+yctr into y
    put the abs of round (y) into y
    drag from a,b to x,y
    put x into a
    put y into b
    subtract 5 from x
  end repeat
end mouseUp
```

```
Script of card button "Trigonometric Circle:"

on mouseUp
  global xctr,yctr,rad,i,a,b
  choose line tool
  put 256 into xctr
  put 130 into yctr
  put 100 into rad
```

```
    put 0 into i
    put xctr+rad into a
    put yctr into b
    repeat until i>6.3
      put (round(rad*cos(i)+xctr)) into x
      put (round((rad)*(sin(i))+yctr)) into y
      drag from a,b to x,y
      put x into a
      put y into b
      add .1 to i
    end repeat
  end mouseUp
```

GUI-Based Approaches to Producing Static Computer Graphics

Programming a sequence of commands, even in turtle geometry, is a very abstract way to draw a picture. A much more concrete method, based on the idea of the graphical user interface (GUI), comes closer to the everyday experience of actually sketching a picture with paper and pencil. GUI graphics applications have a variety of graphics tools, functions, and effects. Selecting these tools, functions, and effects, and then using them to draw a graphic is done with one of any number of input devices. GUI-based approaches work fundamentally the same way, irrespective of which input device is used. The most direct approach is using a light pen to actually draw on the computer screen. Pressure-sensitive graphic tablets or sketchpads also can be used. The user draws on the tablet with a blunt stylus. The motion of the stylus on the tablet is mirrored on the computer screen. The feel of these electronic sketchpads is less natural than light pens, and it usually takes awhile to develop the necessary eye-hand coordination. In between the light pen and graphics tablets on the "feel scale" is the mouse — a hand-held device with one or more buttons that mirror the motion of the user's hand. Mouse users have their own vocabulary as they point to and manipulate screen objects, such as "aim and click," "double-click," and "click, hold, and drag."

Two main types of GUI-based graphics packages are available, painting packages and drawing packages, which are named closely after the way the graphics are stored. Paint packages are analogous to painting or printing directly on a sheet of paper with a pencil or pen. (See Footnote 4) After the graphic has been painted, it cannot be edited or modified. Instead, you have to use an "eraser" to correct mistakes and make changes or "cut" out entire sections of the graphic. As one might guess, graphics produced by painting packages can only be saved as a bit map.

Drawing packages, on the other hand, allow a graphic to be composed of one or more objects, each of which can be continually edited and modified. Figure 3.10 shows a screen snapshot of a graphic package called MacDraw II for the Macintosh. The features found in MacDraw II are typical of drawing packages. The menu bar across the top of the screen designates categories of graphic and text effects, as well as file functions such as saving and printing. Just below the menu bar and title line are some of the many patterns that can be used to "fill" any screen object created (including simple objects like straight lines). A palette of graphic tools is shown on the left edge of the screen. Horizontal and vertical rulers mark the dimensions of the drawing page. The slide bars on the right and bottom edges of the screen let the user move the drawing page around the screen (necessary because the computer screen can only show a small portion of the entire drawing page at any one time; most packages use 8 1/2-by-11-inch paper as the standard size, but this can be greatly expanded).

The screen arrow, controlled by the user via the input device (like the mouse), is used to select any tool, function, or effect, as well as for drawing. It is common for the screen arrow to change shapes to reflect its particular function at any time. For example, as an arrow, it represents a selection tool. When freehand drawing, it may look like a pencil or a paint brush. As this example shows, a GUI-based approach uses many graphical symbols to represent the tools and functions, such as the box tool, oval tool, and line tool. Some of the other symbols are less obvious. The capital letter "A" represents a text editor to generate and modify text objects. The pair of "mountains" at the extreme bottom-left edge of the screen either enlarge or reduce the view of the drawing page.

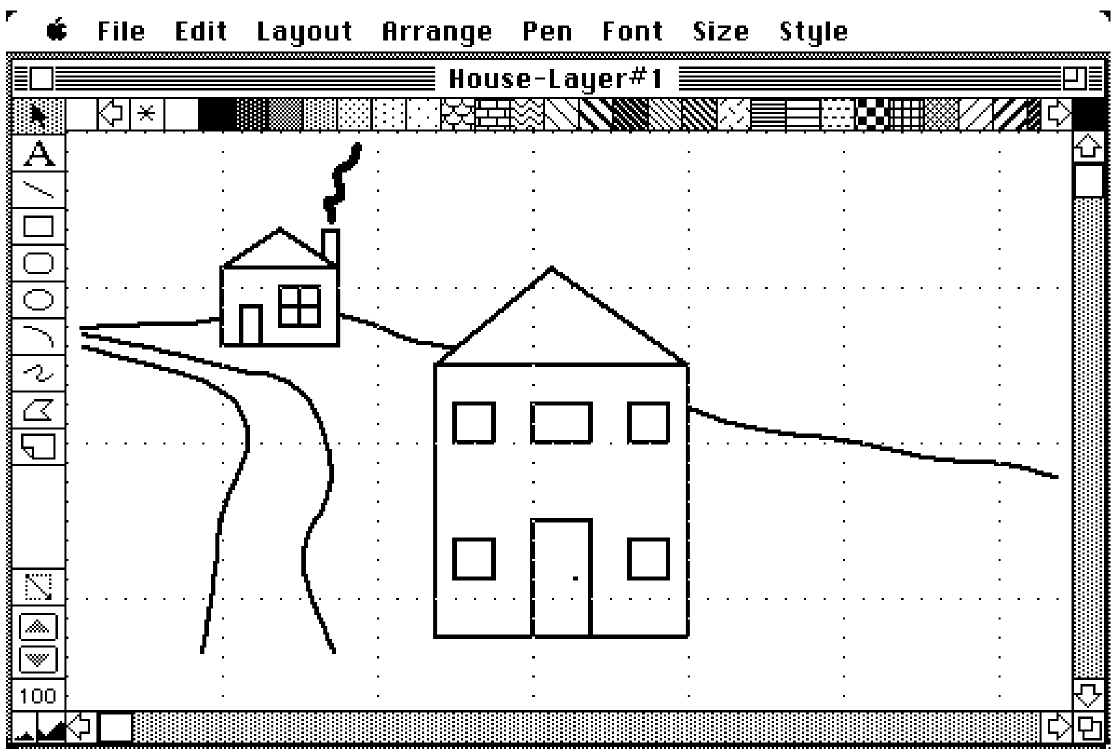

FIGURE 3.10
An example of a typical graphics package based on a graphical user interface (GUI).

Just as an artist working with traditional drawing materials, a computer graphic is produced in a GUI-based approach by alternately selecting and using tools, functions, effects, and other features (such as color). Most GUI packages are object-oriented, meaning that as objects are drawn they retain their separate identities. This allows them to be moved, edited, and copied. Hence, any one drawing is comprised of a collection of individual objects.

Examples of Typical Functions and Effects

There are too many graphic tools, functions, and effects across the many applications currently available on the market to possibly describe them all. However, the next section describes a core set of functions and effects common to many graphics packages.

Grouping and Ungrouping Objects. Even a simple graphic like a house is made up of multiple individual objects. The walls, door, window, roof, chimney, smoke, roadway, and pasture in Figure 3.10 are all separate objects. It is therefore much more convenient to *group* the separate objects into meaningful sets, such as all of the objects that comprise each house (e.g., door, window, roof, etc.). The grouped object then can be manipulated in the same way as any other object. When needed, the house can be *ungrouped* at any time.

Object Arrangement. When objects overlap, it is often necessary to define which object should be "on top" or "in front of" the other. This is analogous to cutting out figures from construction paper and laying them down on top of each other. Most GUI systems allow any object to be moved progressively "backward" or "forward." The arrangement of objects is particularly important when each closed object has been filled with a pattern or color. When no pattern or color has been chosen, the objects appear transparent, or as simple line drawings.

Alignment. Freehand drawing on a computer is a tough task. It is extremely difficult to precisely control most input devices, like the mouse, no matter how steady your hand. In order to provide greater accuracy in drawing objects, most packages allow objects to "snap" to imaginary grid lines in small increments, such as one-eighth-inch increments. This feature is similar to the kinds of control and accuracy necessary in computer-aided design. This feature usually can be turned on and off at will.

Rotation. Once an object is created, its orientation on the screen usually can be changed. Almost all systems allow an object to be "flipped" vertically or horizontally, but rotation features are particularly useful. Most packages allow an object to be rotated freely, as well as constraining the rotation to increments of 15, 30, 45, or 90 degrees. An example of rotating an object is shown in Figure 3.11.

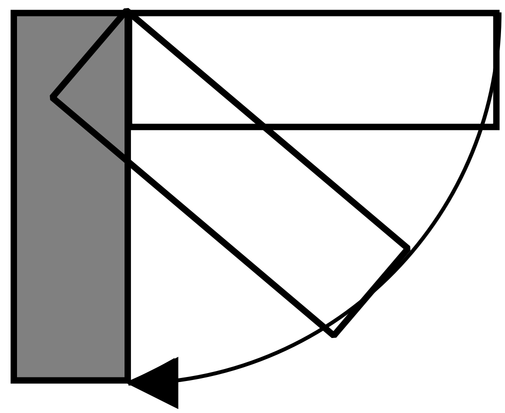

FIGURE 3.11
An example of rotating an object 90 degrees.

Layering. Complicated graphics can consist of hundreds of objects. Many packages allow a complex graphic to be constructed in layers, such that once a layer is defined, it becomes part of the background. The user cannot manipulate objects except those on the currently active layer. This helps the user organize the graphic and helps prevent accidentally selecting the wrong object. Layering is analogous to constructing one graphic on several plates of glass, each stacked on top of one another, as shown in Figure 3.12. Any one "plate," or layer, can be drawn on at a time. Once drawn, the layers can be stacked in any order. Just like grouping, the objects on layers above will cover objects on layers below. The simplest example would be a graphic consisting of two layers, where the one below acts as the background.

Line Smoothing. Freehand drawing of smooth, rolling curves on a computer is very difficult. When using paper and pencil, one can use a variety of sketching techniques and tools, like plastic templates, to make the task easier and to improve quality. These techniques just do not work well on a computer. Many packages, however, allow irregular curves to be constructed as a series of straight lines comprised of "valleys" and "peaks," which are then "smoothed" over by the computer, as shown in Figure 3.13. The low and high spots of the curve become "handles" with which to modify the curve.

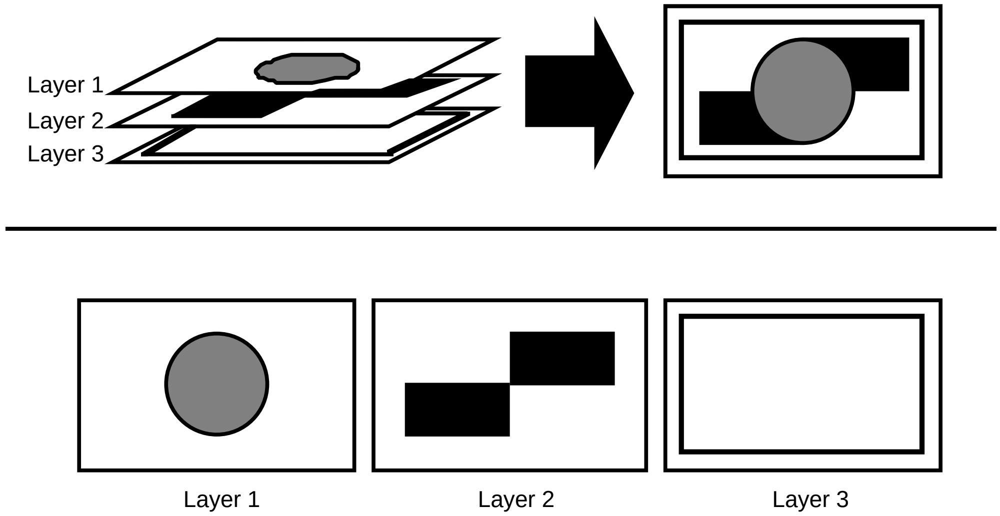

FIGURE 3.12
The concept of layering. In this case, one graphic is comprised of three separate layers, each of which can be edited. The layers can be rearranged in any order. Any one layer can be deleted and more layers can be added, if necessary.

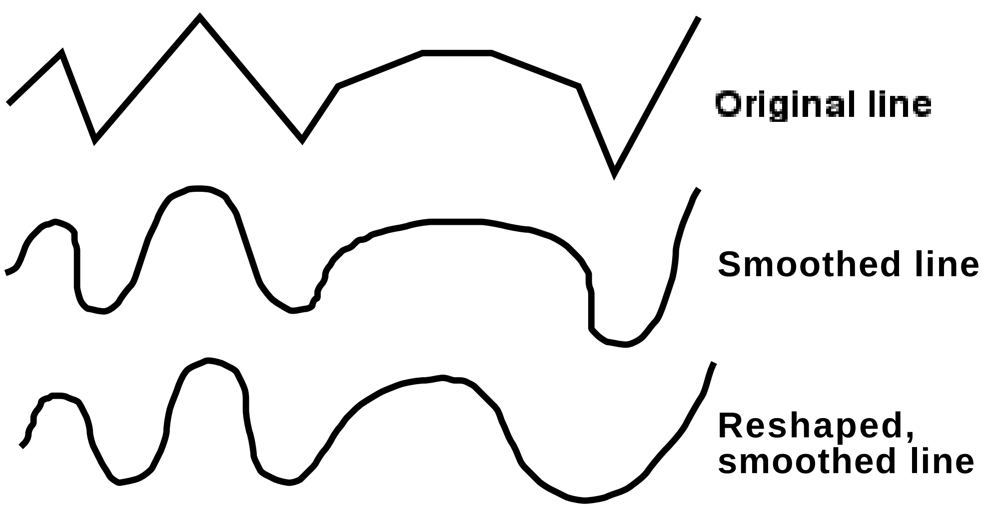

FIGURE 3.13
The concept of line smoothing.

Graphing. The capability to create spatial representations of categorical and numerical information, such as line graphs, bar graphs, and pie graphs, represents a distinct set of graphics applications. These and other graphing functions are usually provided in separate graphing software packages that construct graphs based on raw data entered by the user, such as that shown in Figure 3.14. However, graphing functions are becoming a more popular feature of many commercial spreadsheets.

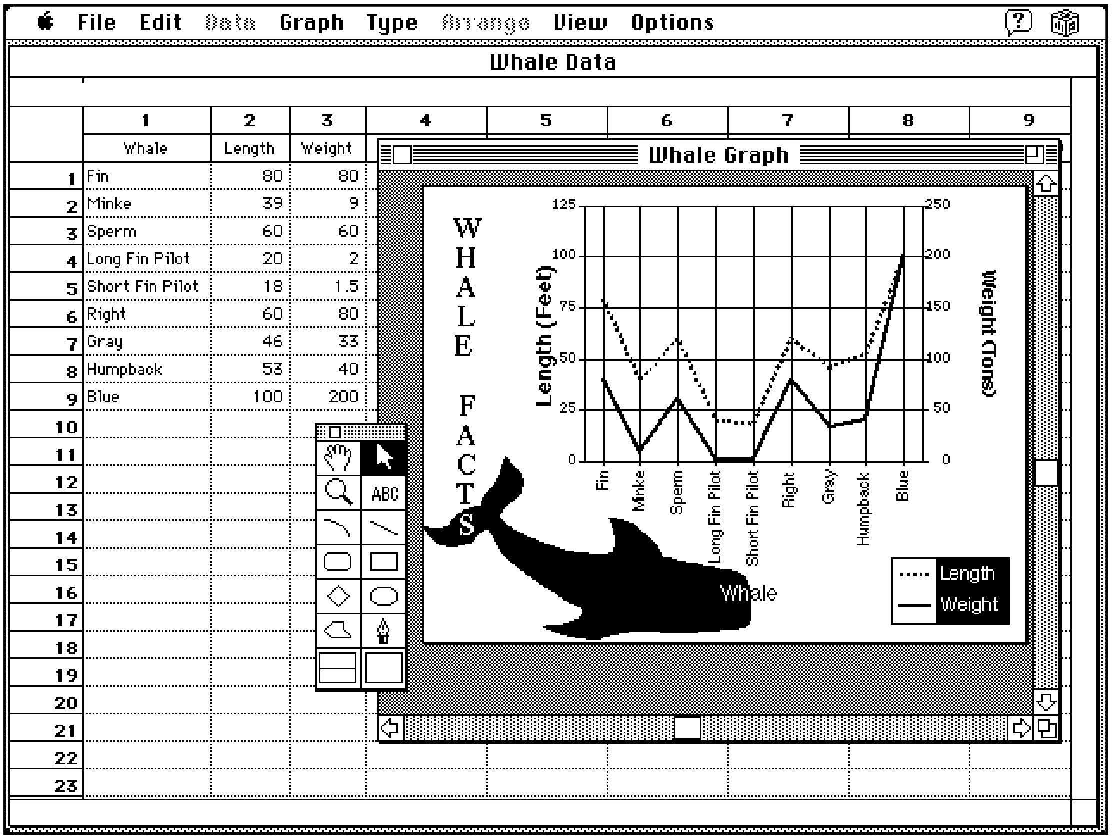

FIGURE 3.14
An example of a graphing package. The user enters raw data and the software subsequently constructs one of many possible graphs.

Second-Hand Computer Graphics: Clip Art, Scanning, and Digitizing

Despite the many features and effects that graphics applications now provide and the many more they likely will provide in the future, there always will exist two inherent user limitations related to creating an original computer graphic — talent and time. Professionals increasingly turn to two alternative methods to get high-quality computer graphics in their materials. Neither method demands much talent because, instead of creating an original graphic from scratch, you either find and borrow a graphic drawn by someone else or take a photograph and convert the picture to digital form. We will refer to these as "second-hand" graphics to distinguish them from the graphics a user draws from scratch.

The most popular form of second-hand graphics are called *clip art* files. The idea is simple: Hire computer graphics artists to draw a collection of graphics using common graphics applications. The files can be sold to users who can load, use, and edit the files as if they had drawn the graphics themselves, as shown in Figure 3.15. The idea of clip art is not new; it has been used for many years in the printing industry. Most arts and craft stores sell print-based clip art that can be cut with scissors and used in newsletters and other publications. Although many companies produce and sell clip art, computer user groups frequently swap graphics files among members. Commercially produced clip art is usually sold with the understanding that the user is given the right to reproduce the art work freely in whatever work it is needed. Unfortunately, copyrighted graphics are also frequently shared in this manner among users; reproducing copyrighted material without permission is, of course, illegal.

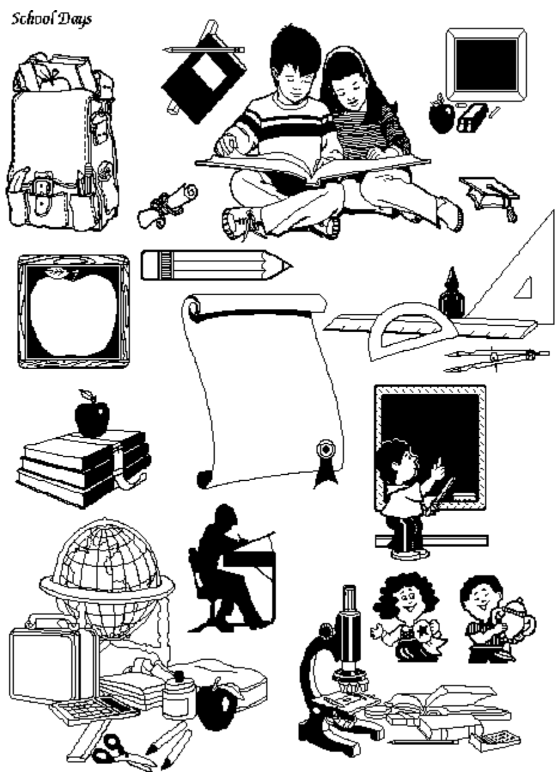

FIGURE 3.15
An example of "clip art."

Second-hand graphics also can be created by converting analog images, such as print-based or video pictures, to digital form. Optical scanners are probably the most common device used for this purpose. A typical hardware configuration is shown in Figure 3.16. Optical scanners work in much the same way as paper copiers. The document to be scanned, typically called the original, is usually placed face down on a glass plate. A bright light is pulled across underneath the original to detect variations in the amount of light reflected back, called **reflective density**.

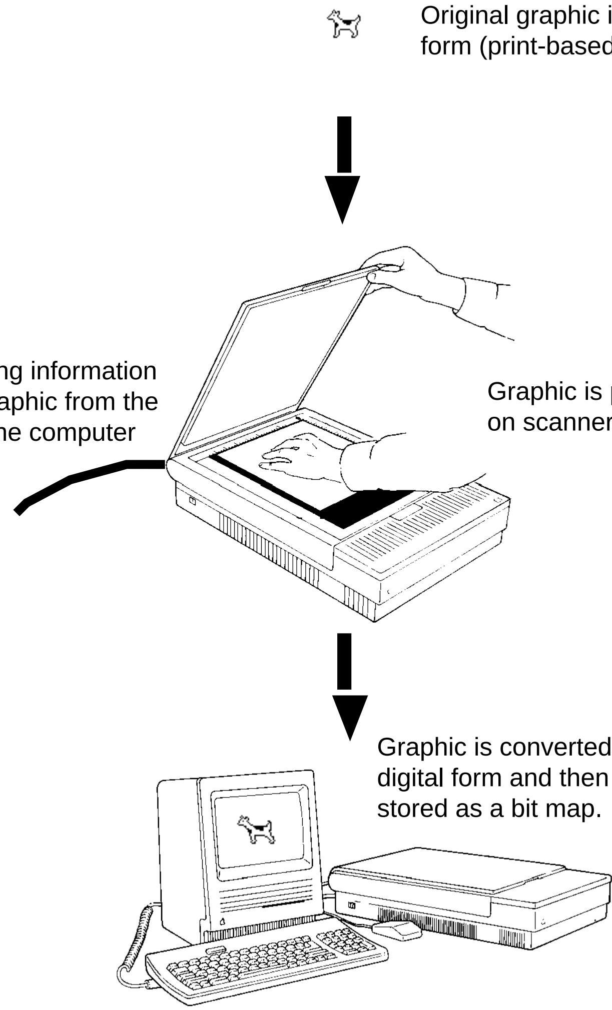

FIGURE 3.16
A computer interfaced with an optical scanner.

Scanners vary in the amount of information that is sent back to the computer for each point scanned in the original, and this information is used to define one of several **composition types**. The simplest, known as **line art**, is when each scanned point is recorded as either black or white. Most scanners allow for the handling of shades of gray. Some use **halftone** patterns, similar to that used in a newspaper photo. Others use **grayscale** settings, where continuous shades of gray are approximated. Finally, the most sophisticated scanners can record color.

Various software features allow the user to change various scanning settings. The **threshold** setting determines whether a specific dot on the original is recorded as black or white. Other common settings include the ability to change the *brightness*, or the degree of overall whiteness of the image, and the *contrast*, or the relative difference between black and white. At the highest contrast settings, black and white are emphasized and few gray shades are left. At the lowest contrast settings, the scanner emphasizes the middle gray shades, leaving little which is pure white or black.

Obviously the more information a scanner records about an image, the more computer memory is needed to store the file. Fairly simple graphics about the size of a typical computer screen can take as little as 5 to 10 kilobytes for simple line art drawings. Grayscale images, on the other hand, must record an exact shade of gray for each scanned dot. For example, a scanner that records one of 16 shades of gray for each dot must use four bits of memory for every scanned point in the original. At the extreme end, it is not uncommon for a scanned color image to contain up to one megabyte (approximately 1 million bytes) of memory. Regardless of the brand and the features, the issues of economy and processing ability related to computer memory become very important to understand.

There are other examples of devices that allow images or objects to be digitized, including specially designed or adapted photographic or video equipment that takes digital snapshots of real objects. Scanning and digitizing technology is advancing at a tremendous rate and is worth watching over the coming years.

PRODUCING ANIMATED COMPUTER GRAPHICS

Similar to static graphics, animated graphics can be produced either by a command-based or GUI-based approach. Producing animated displays with command-based approaches are really just extensions of the techniques discussed in relation to static graphics. On the other hand, GUI-based approaches can vary greatly from one animation package to another. We will discuss various approaches, starting with some simple, yet fundamental ideas, and then proceed to other, more sophisticated, approaches. It is useful to understand development issues of animated displays in terms of two animation designs: fixed-path and data-driven.

Fixed-path animation is analogous to choreographing a movie sequence. The same exact animation is supposed to happen the same way, in the same place, at the same time, each and every time the sequence is executed. Fixed-path animation, therefore, is a good technique when a design calls for a specific presentation of an animated sequence. We will consider some of the fundamental programming techniques in creating computer animation. However, many GUI animation packages offer the ability to record the real-time motion of a screen object while a user simply moves it around the screen. The software can then play back the animation just as it was "performed." The software does all the dirty work for the developer, such as storing and processing all of the mathematical operations actually responsible for the animation to take place.

In contrast, motion and direction of screen objects in data-driven animation do not vary according to the actual movement of the human hand, but by some data source. Although fixed-path animation also can be created by a data source and then "captured" or "recorded,"(See Footnote 5) we will define the data in data-driven animation as that generated by the student during the instructional sequence. Visually based simulations, such as flight simulators and video games, are good examples of what we will call data-driven animation. By our definition, animation is produced in *real-time*, or in the actual time that the user watches the display. In this way, the animation acts as visual feedback to students as they interact with the simulation moment to moment. Obviously, there is no way to anticipate when or if a particular student will "dive" or "climb." Instead, a mathematical model of the physical environment being simulated must be programmed into the computer in such a way as to refresh the graphics realistically in order to create the illusion that the student is actually controlling the "plane." Whereas fixed-path animation does only one thing, data-driven animation, theoretically, can produce an infinite number of displays with a finite amount of information. Both fixed-path and data-driven animation can be manipulated in one, two, or three dimensions on most microcomputer systems. For simplicity, the next several sections will deal exclusively with one or two dimensions.

Command-Based Approaches to Fixed-Path Animation

Animation is an illusion that tricks a person into seeing something that really is not there (the psychology behind this trick is explained more fully in chapter 4). The trick to inducing the perception of a moving object on the computer screen involves creating a series of carefully timed "draw, erase, move, draw" sequences. In order for convincing real-time animation to be produced, the computer must be able to complete about 16 of these sequences in one second. The mathematical model is essentially the same behind both command-based and GUI-based approaches. The difference is simply that in a command-based approach the user must actually program the mathematics of the algorithm into the computer.

An annotated, illustrated example of a simple graphics program written in BASIC to produce a fixed-path animation is shown in Box 3.2. The object that is being animated is a single point of light. The example is presented as a progression of some fundamental animation concepts. Even if you know nothing about programming, you should be able to follow its logic. The result of the program is a ball bouncing back and forth on the screen. Obviously, this example can not be presented well given the static medium of a book. But it is hoped that by reading and following the example, you will get a sense of the animation principles at work. In order to really understand the principles, however, you should read Box 3.2 while trying out the example on a computer.

Of course, a single point of light is not a very interesting screen object to manipulate. More sophisticated shapes, such as arrows, planes, boats, animals, or space ships, can be moved in much the same way. However, the perception of animation will be lost if the shape takes too long to be drawn and erased before it is moved to a new position and drawn again. Command-based approaches on microcomputers allow complex objects to be coded into a

shape table, or a precisely defined memory location that stores information about one or more shapes in the form of instructions called **plotting vectors**. While the details of how to do this are beyond our scope, the point to be remembered is that once a shape table is correctly defined, each shape in it can be manipulated as easily as the single dot of light discussed in Box 3.2. Some systems combine command-based and GUI-based approaches, such as HyperCard, and allow most screen objects, such as buttons, fields, and "lassoed" screen areas, to be treated as shapes and moved in a similar mathematical way as the dot in Box 3.2.

Box 3.2

Follow the Bouncing Ball

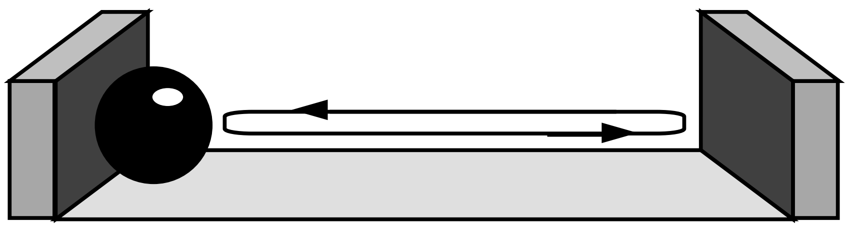

In this example, a simple computer program, written in AppleSoft BASIC for the Apple II, animates a “ball” (a single point of light) bouncing back and forth on the computer screen. The purpose of this program is to show one real example of applying the “draw, erase, move, draw” idea in a command-based approach to create a fixed-path animation sequence.

Here is a simple computer program in which the ball is animated left to right across the computer screen:

```
 100 GR
1000 REM ANIMATE LEFT TO RIGHT
1040 LET H=0
1060 LET V=20
1080 COLOR=2
1100 PLOT H,V
1120 COLOR=0
1140 PLOT H,V
1160 LET H=H+1
1180 GOTO 1080
```

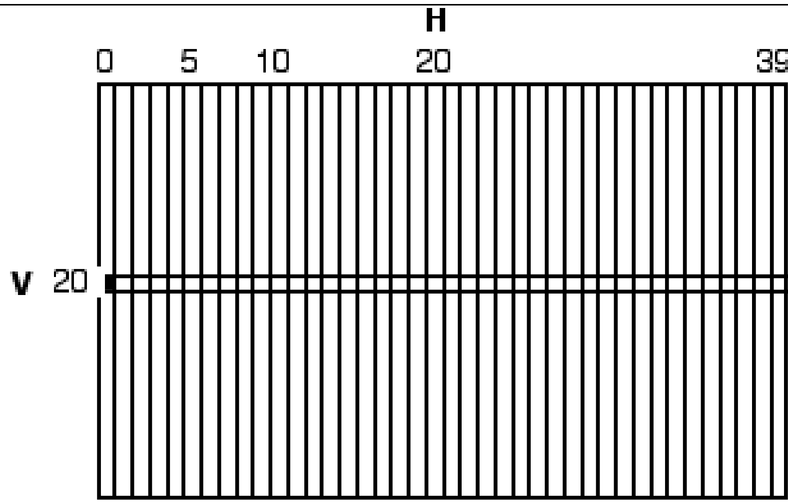

When the ball goes just beyond the right edge of the screen, it triggers an error because it exceeds the window limit of "39." Let's walk through each of the lines to discuss how each contributes to the final animated sequence and also to discuss some problems which exist in the program.

Line 100 tells the computer to call up a fresh low-resolution graphics page. Line 1000 is simply a "remark" or "comment" line. Line 1040 sets a variable called H to zero and Line 1060 sets another variable called V to 20. The variable H will be used to define the position of the ball on the Horizontal axis and V will do the same on the Vertical axis. Note that since the ball will only be moving back and forth, the variable H will change, however V will not. Line 1080 chooses a color for the ball: two is the code for blue. Line 1100 finally plots a point at the screen location H,V which translates into a dot at the intersection of 0 across and 20 down. We have therefore completed the first stage of our "draw, erase, move, draw" animation model.

The purpose of the next two lines is to erase the ball. Line 1120 chooses another color for drawing: zero is the code for black. Line 1140 again draws a ball at the same screen location H,V. However, since the ball is drawn in black the ball disappears because the background color of the screen is also black. Technically speaking, the ball was not erased, it was simply "painted over" in the same color as the background, so it vanishes. This little "trick" completes the second stage of our "draw, erase, move, draw" animation model.

Line 1160 performs the mathematical calculation necessary to identify another screen location, in this case, the cell immediately to the right of the first. The command "LET H=H+1" loosely translates "make H what it was before plus 1." Since 0+1=1, H becomes 1. Mathematically incrementing the H variable simply tells the computer to "aim" at a different screen location which fulfills the third stage of our "draw, erase, move, draw" animation model. Line 1180 tells the computer to immediately branch to line 1080 and continue working from there. Line 1080 switches the drawing color back to blue. Line 1100 again plots a point at the screen location H,V. However, this time, H is now 1, so a blue ball is drawn at the intersection of 1 across and 20 down. This completes the first of many “draw, erase, move, draw” sequences necessary to move the ball across the screen.

The program is now involved in a loop and will continue executing the loop until we tell it to stop (by pressing Control-C), the computer loses power, or something else unforeseen happens (like an error). Continuing the program’s logic, line 1120 again changes the drawing color back to black. Line 1140 again draws a black ball over the blue ball which causes the ball to again disappear. Line 1160 again adds 1 to H, making it 2 and Line 1180 again branches the program back to 1080 starting the whole process over again. The program continues in the loop which causes it to continually draw, erase, move, and draw the ball over and over going from the left edge of the screen to the right.

However, there are two major problems. One problem is perceptual and the other is technical. The technical problem is that this little program “crashes” upon reaching the right edge of the screen because there is no horizontal screen location at 40. The result is a rather rude and cryptic error message like “Illegal Quantity Error.” However, this technical problem is easily dealt with, so we will deal with the perceptual problem first. When the program is actually run, the ball does not *smoothly* travel from left to right. Instead, it often appears as though it is “skipping” sporadically from left to right. The problem is not a computer malfunction, in fact, the computer is working *too* well. Even low-end microcomputer systems can process and execute this little program so fast that the blue ball is erased before the eye has time to actually perceive it. In order to give the eye a chance, we need to add a small delay to the program in the form of Lines 1110 and 1115:

```
 100 GR
1000 REM ANIMATE LEFT TO RIGHT
1040 LET H=0
1060 LET V=20
1080 COLOR=2
1100 PLOT H,V
1110 FOR D=1 TO 50
1115 NEXT D
1120 COLOR=0
1140 PLOT H,V
1160 LET H=H+1
1180 GOTO 1080
```

Lines 1110 and 1115, in essence, give the computer “busy work” to perform, counting from 1 to 50 in this case, *while the blue ball is displayed on the screen*. This delays the program long enough at this crucial point to give the human eye a “long,” clear view of the ball. These two lines will make a dramatic perceptual difference. The ball will now go *smoothly* from left to right when the program is run.

However, remember that our goal was to have the ball bounce back and forth. So far, it travels only left to right. We also have to deal with the technical problem of the computer crashing when the ball goes over the right edge of the screen. The following changes take care of both issues:

```
 100 GR
1000 REM ANIMATE LEFT TO RIGHT
1040 LET H=0
1060 LET V=20
1080 COLOR=2
1100 PLOT H,V
1110 FOR D=1 TO 50
1115 NEXT D
1120 COLOR=0
1140 PLOT H,V
1160 LET H=H+1
1170 IF H > 39 THEN GOTO 2000
1180 GOTO 1080
2000 REM ANIMATE RIGHT TO LEFT
2040 LET H=39
2060 LET V=20
2080 COLOR=2
2100 PLOT H,V
2110 FOR D=1 TO 50
2115 NEXT D
2120 COLOR=0
2140 PLOT H,V
2160 LET H=H=1
2170 IF H < 0 THEN GOTO 1000
2180 GOTO 2080
```

The first thing you will probably notice is that the program is about twice as long as it was before. This is because it was copied and pasted below itself, starting with line 2000. These duplicated lines are identical to their counterparts from 1000-1180 with one simple, yet crucial exception. Line 2160, instead of adding 1 to H, subtracts 1. Adding makes the ball go from left to right, and subtracting makes the ball go from right to left (in other words, it “bounces”). Line 1170 and its counterpart at line 2170 were also added. These IF/THEN lines tell the computer to branch between the two parts of the program when the ball is about to go either too far to the right (i.e. H>39) or too far to the left (i.e. H<0). The result is a program which animates a little ball bouncing back and forth on the computer screen.

GUI-Based Approaches to Fixed-Path Animation

The "draw, erase, move, draw" model also guides the development of fixed-path animation in GUI-based approaches. However, most GUI-based approaches do not use a mathematical model to produce fixed-path animation at the user interface level, even though mathematics may still underlie the animation sequence. Before discussing some examples of software designed expressly for the production of fixed-path animation, we will consider the simplest and oldest approach to animation — frame-by-frame animation.

Frame-by-Frame Animation

Frame-by-frame animation, as the name implies, is when a set of successive frames is constructed so that when shown in rapid succession at just the right rate, one or more objects appear to move. This is the same technique that one would use to create paper and pencil animation where the object to be animated is drawn at slight variations from page to page. Animation is produced by flipping through the pages at just the right speed with your thumb. Frame-by-frame animation is the oldest form of animation and is the same technique used by the most sophisticated examples of film animation, such as Disney, whether drawn by hand or by computer. This is also the technique used in the labor-intensive "claymation" films, where real objects, such as bendable figurines, are moved slightly and photographed one frame at a time. (See Footnote 6)

Frame-by-frame animation requires that every detail between one frame and the next be exactly the same — except for the object or objects being animated. These objects are drawn in small, but discrete, variations from the preceding frame. If the animation is shown at 30 frames per second, a standard video display rate, each frame shows the object in increments lasting one-thirtieth of a second. Reducing the animation to a rate of 10 frames per second, will, of course, be cheaper and quicker to produce, but the quality of the animation will suffer accordingly. Anyone who has seen and compared a classic Disney cartoon with typical Saturday morning television knows this quality difference firsthand. Even though the frames are discrete, the human perception system will fill in the gaps and perceive continuous movement (see chapter 4).

GUI computers make frame-by-frame animation readily available through their copy and paste features. In HyperCard on the Macintosh, for example, once the first frame, or card, is created, it can be copied and pasted as the second card, which, in turn, can be edited to change the appearance or position of the animated object. This and all succeeding frames can be copied, pasted, and revised as shown in Figure 3.17. When the entire sequence is shown at the proper rate, animation is produced.

Figure 3.18 shows a typical way some animation software packages accomplish frame-by-frame animation. In this particular case, the software allows for single objects to be animated "in place." Figure 3.18 shows a simple example to create the illusion of a spinning helicopter rotor. The software allows for any single frame to be created pixel by pixel and then copied over into the next frame for minute editing necessary to create the illusion.

Animating a person running, a horse galloping, or clock hands revolving would be other examples of animation where this approach would be useful.

FIGURE 3.17
An example of frame-by-frame animation.

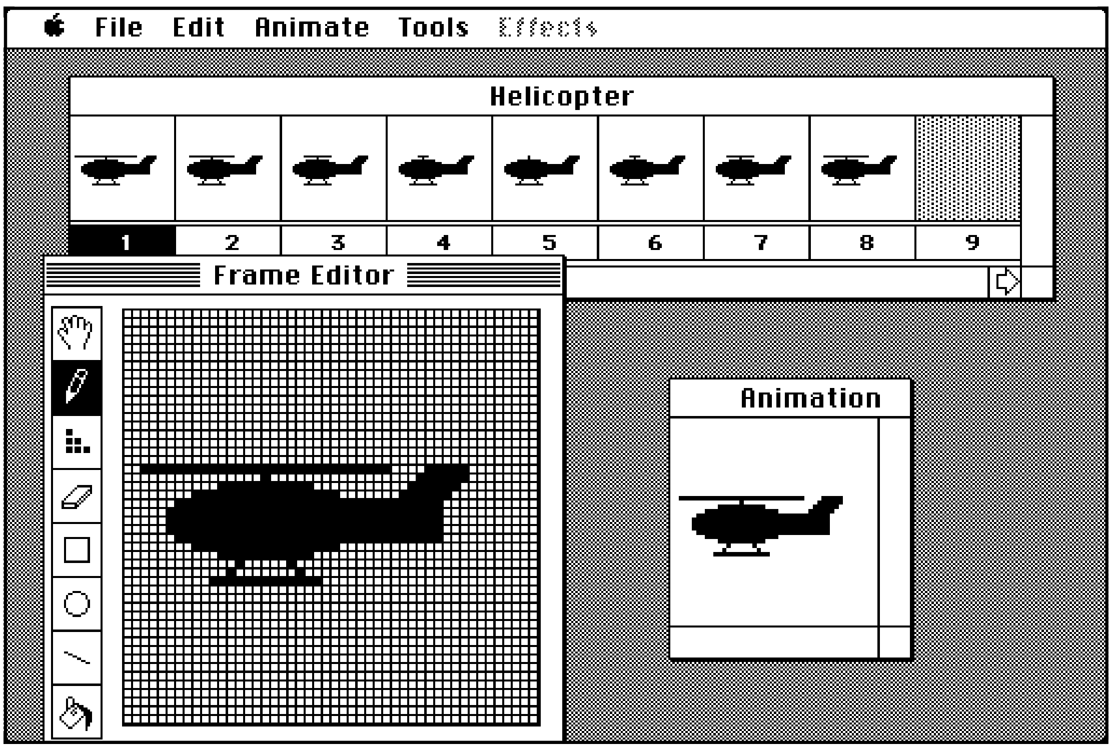

FIGURE 3.18
An example of a frame-by-frame animation package. When shown in rapid succession, the helicopter's blade will appear to spin.

Of course, it becomes a surprisingly complex task to keep track of even a few animated objects at a time. For this reason, some computer software packages separate the task of manipulating *individual* animated objects from the assembly of all the objects into one completed animation sequence. Several packages use the analogy of a movie production to do this. (See Footnote 7) Individual animated objects become the cast of characters. As shown in Figure 3.19, the entire animated production is called the "score" and consists horizontally of the individual frames and vertically of "tracks" on which cast members perform. Any one frame combines any number of animated objects, each on its own separate track. There are also separate tracks available for sound and special effects. When finished, the score can be rewound or fast-forwarded to any position and played, meaning that frames are then displayed one by one at a specified rate.

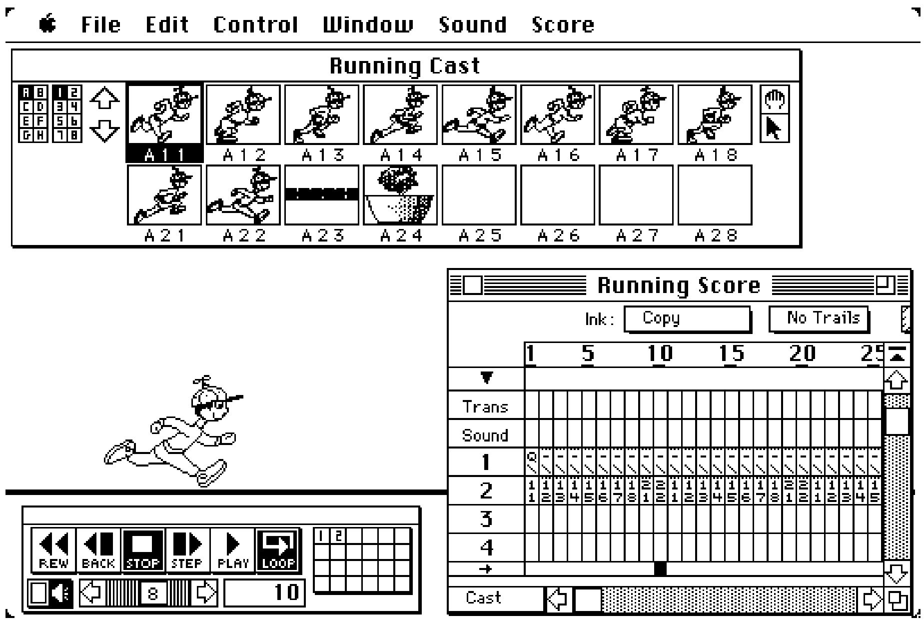

FIGURE 3.19
A snapshot of a sophisticated animation package that uses the analogy of a stage production.

Other GUI-Based Approaches to Producing Fixed-Path Animation

There are a variety of other ways to produce fixed-path animation. Some applications allow a user to move an object around the screen freehand while recording the motion in real-time. The software then automatically converts the information into frame-by-frame animation. Of course, if the user's hand shakes a little in the middle, it will be necessary to edit those particular frames. This is usually a more difficult task than first imagined. Other approaches allow for the path of the object to be defined first and edited, rather than the object itself.

The simplest methods for defining an animated path are those that animate an object along a straight line, as shown in Figure 3.20. The entire animated sequence is defined by the object to be animated, the two end points of the imaginary line along which the object moves, and the time (sometimes defined as speed) that the object takes to complete its "run." The object's initial screen position usually defines one end point. The object is then physically grabbed and moved to the other end point with the computer recording the time taken to get there. The sequence then can be fine-tuned until the desired effect is reached.

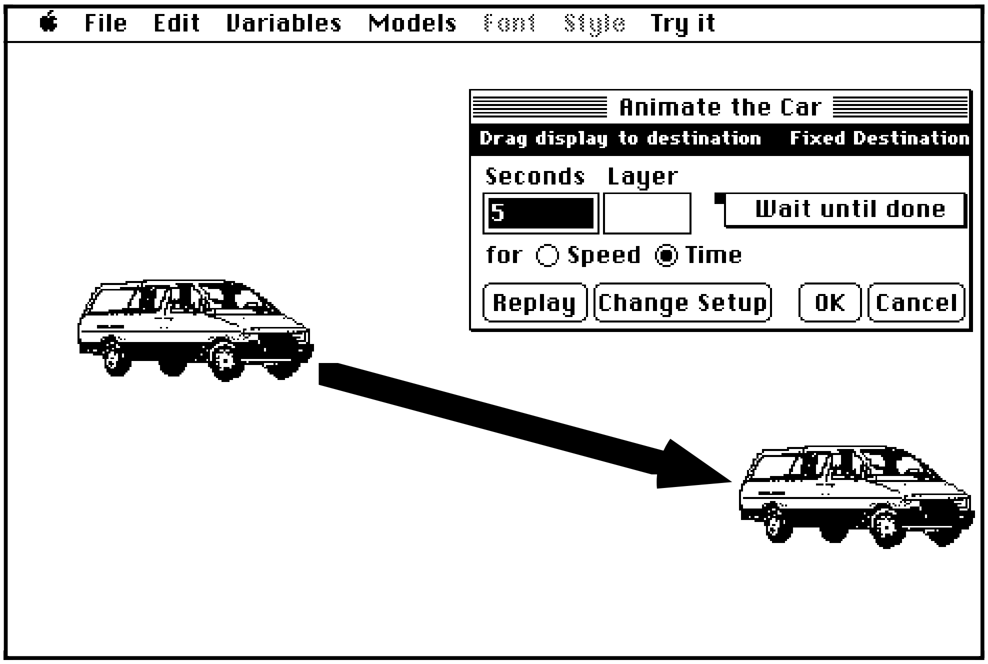

FIGURE 3.20
An example of "fixed destination" animation. All the user has to do is drag the object to be animated from the starting point to the ending point. The computer extrapolates all of the points in-between.

The next level of sophistication is animation along a crooked path, analogous to a piece of string. Again, the sequence has a starting point and an ending point. However, there are also one or more points along the way that "pull" on the string, as shown in Figure 3.21. Points can be added or deleted to the string. The points also can be defined to make the string go along a straight or curved line between the points.

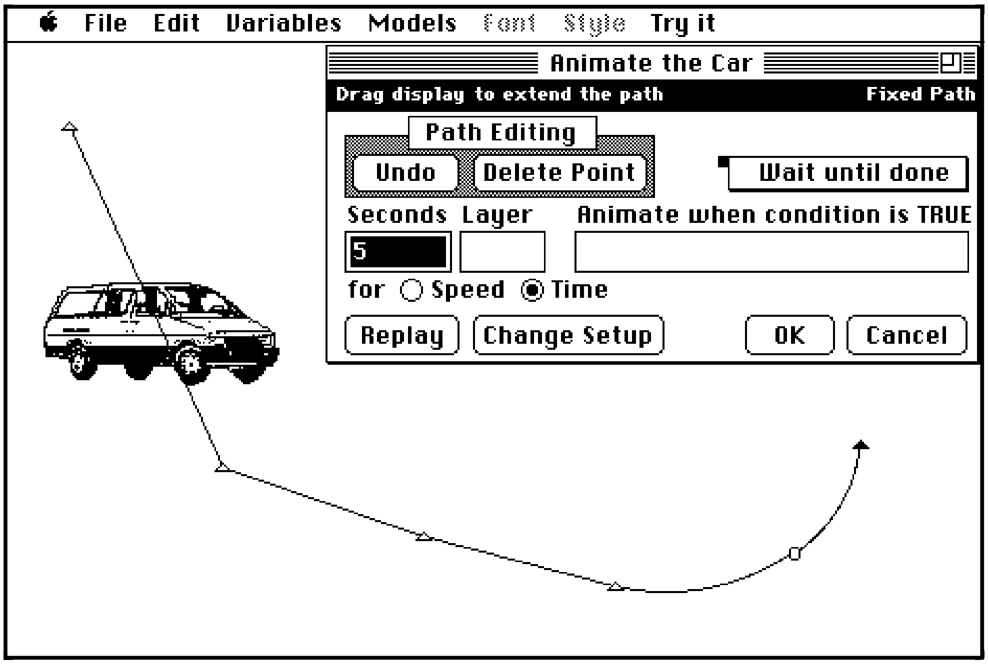

FIGURE 3.21
An example of fixed path animation. The animated path of the object can be continually extended and reshaped.

Data-Driven Animation

Command-based and GUI-based approaches to data-driven animation are quite similar. Each requires methods and techniques to keep track of one or more variables that control the animation of one or more screen objects. Even though many systems may capitalize on GUI features, all data-driven animation methods are essentially command-based approaches.

Recall the example of fixed-path animation using a command-based approach from Box 3.2. Rather than have the computer decide on the values of H and V, we can let them be determined by user interaction. We need some method of allowing the user to provide *input* to the computer system. For example, most computer keyboards have four arrow keys, one for each direction (i.e., left, right, up, and down). The computer can be programmed to move a screen object, such as the ball, in the direction of whatever arrow key is pressed or clicked on. Recall from our "bouncing ball" example that the variable H controlled horizontal motion and the variable V controlled vertical motion. The program can be easily modified to add 1 to H when the right arrow is pressed, or subtract 1 from H if the left arrow is pressed. Similarly, the program can add 1 to V when the down arrow is pressed and can subtract 1 from V when the up arrow is pressed. This simple program lets the user move, or drive, the ball anywhere on the screen.

With a little creativity, the program can be incorporated into a variety of instructional contexts: teaching how to follow map directions by adding a street background and having the user drive a car between points (as shown in Figure 3.22); or teaching latitude and longitude by switching to a nautical map and changing the object to a boat; having the object leave a trail, the program becomes an electronic Etch-A-Sketch; etc.

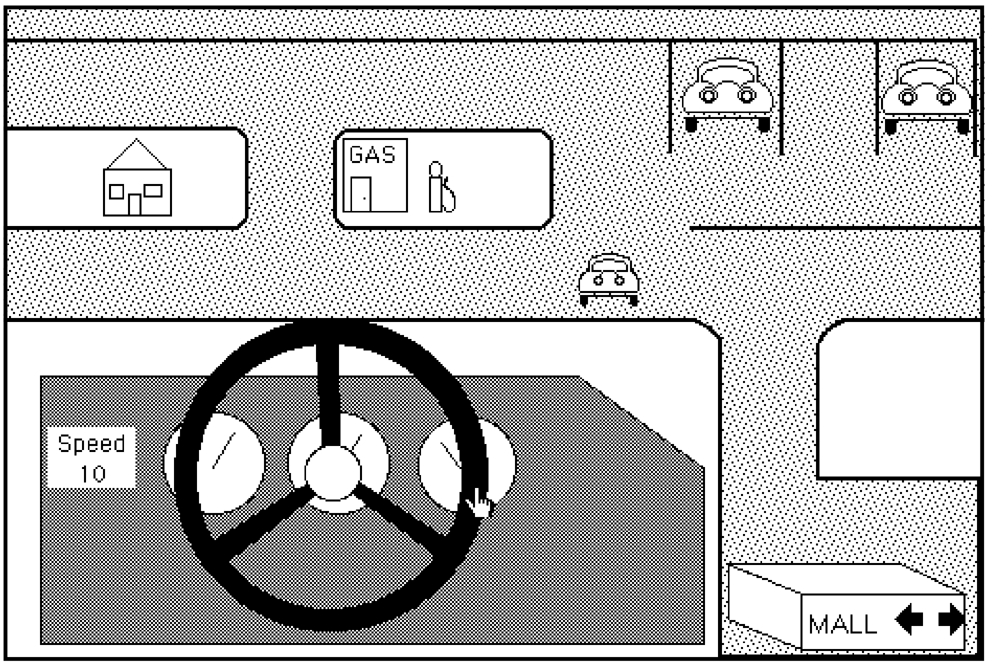

FIGURE 3.22
An example of data-driven animation. The user drives the car around the town just by pointing the hand icon on the steering wheel.

THE INSTRUCTIONAL DELIVERY OF COMPUTER GRAPHICS

Regardless of how static or animated graphics may be produced, the ultimate instructional question is how the computer graphics will be implemented or delivered in an instructional system. The most obvious delivery platform for computer graphics is the computer. This includes all instructional materials delivered by computer, such as computer-assisted instruction and computer-controlled presentations.

However, it is not always feasible to deliver the materials by computer, either because of economy or practicality. Probably the next most obvious delivery platform is print-based material. Computer animation presents unique concerns. Data-driven animation, by definition, must be delivered by computer because the student input must be processed moment to moment. Fixed-path animation, however, offers delivery alternatives. A practical solution to delivering fixed-path animation is to transfer the animation to videotape in order to take advantage of the widespread availability of videocassette players. However, several compromises must be made. Unless the user is willing to manually control the player, the program will be externally paced, an issue that must be included in the original design. Other problems are related to the quality of the video image, depending on the computer system used to produce the animation in the first place.

Computer systems and video systems vary widely in their methods of composing and transmitting a video signal. Video manufacturers typically use one of three signal systems to define the scan line: NTSC, PAL, and SECAM. NTSC, named after the American National Television Standards Committee that established it, is the most common in the United States. PAL, or Phase Alternation Line, is the system adopted throughout most of the United Kingdom, Europe, the Middle East, Africa, Australia, and South America. SECAM, short for séquential couleur á memoire, was developed in France. Computer manufacturers vary even more widely in the way images are transmitted, and few use a standard video signal. (See Footnote 8) Computer screens, as already mentioned, are based on the pixel, not the scan line. The lack of compatibility between computer and video displays poses problems when trying to transpose a signal from one platform to the other. It is usually necessary to use a special device to synchronize the computer and video signals, such as a "gen-lock" that locks the pulses generated on separate signals to a common beat. Standardization between computer and video displays is being recognized as a very real need, especially as computer- generated graphics are becoming more common in video productions (Parsloe, 1983).

REVIEW

- Generally speaking, instructional design should drive instructional development. However, each process provides feedback and guidance to the other.
- The pixel, the line, and the polygon are known as graphic primitives from which all other graphics can be constructed.
- There are two main types of computer graphics displays. Raster graphics displays, the most common on microcomputers, divide the computer screen into a matrix of pixels. Vector graphics displays draw lines in terms of their length and direction but are far less common.
- There are three main approaches to producing computer graphics. The command-based approach involves writing computer programs that draw graphics only when the programs are run. The GUI-based approach capitalizes on graphical user interfaces to allow a more natural method for drawing graphics, such as selecting

and drawing with one of many graphical tools and effects. Finally, analog pictures can also be scanned and digitized into the computer. Depending on the approach, computer graphics can be stored as bit-mapped images, computer programs, or mathematically definable objects.

- There are two types of computer animation. Fixed-path animation repeats the same animated sequence each time, as in frame-by-frame animation. Data-driven animation is produced when the motion and direction of screen objects depend on a user's input.

NOTES

1. This was especially true on some of the first microcomputers, like the Apple II, in which the memory associated with a graphics "page" was always constant. However, applications on more recent desktop computers, like the Macintosh, are able to store bit-mapped graphics more economically by saving only the actual area of the graphic display. In this way, a small graphic, like an inch-square box, will require less memory than a large graphic. Graphics can be stored on disk in a variety of ways, including bit maps. The next section discusses software issues associated with how graphics are produced and stored.
2. Designing graphics as a collection of objects is a direct application of object-oriented programming systems (OOPS) and is really an extension of procedural programming languages, such as SMALLTALK, LISP, LOGO, and Pascal. Object-oriented programming has had a strong influence on the field of artificial intelligence.
3. In fact, it is a mistake to think that the purpose of LOGO has anything to do with creating graphics. LOGO is a programming language and not a graphics application. LOGO is concerned with how people learn and how the computer can provide a rich source of problem-solving tools. A brief background of LOGO and its learning philosophy will be discussed in chapter 8.
4. About the only common GUI-based graphics application on the Macintosh that has no object-oriented features is MacPaint. MacPaint deserves special notice because it was only one of two software packages available when the Macintosh was first introduced — mainly because it and a text editor (MacWrite) were given away with the machine. In MacPaint, although graphic tools are selected and used with the GUI-based approach, once the graphic is drawn it is only interpreted by the computer as a bit map. The only way to change the graphic is to edit it pixel by pixel or to "erase" it and start over. MacPaint is still a commonly used graphics application. Special credit goes to its creator, Bill Atkinson. The best testimonial to its design was how it was quickly copied by other producers of graphics applications.
5. Many fixed-path animated sequences are generated by data. The command-based approach using BASIC programming in the next section would be a simple case in point. Other, more sophisticated examples would include defining some complex chain of events using a mathematical model, and then allowing the mathematics to define the animated sequence frame by frame. Particularly good examples of the important role that mathematics plays in producing fixed-path animation include

visual displays of chaotic systems, such as a flag waving in the breeze or the motion of a water-filled balloon as it bounces on the floor. These are examples of chaotic systems and are based on nonlinear mathematics. Because of the time needed to calculate the moment-to-moment positions of these chaotic systems, few, if any, can be produced with real-time animation.

6. This process is so labor intensive, that when it was in production, the original *Flintstones* cartoon series was among the most expensive television shows of its time.
7. This analogy is the one used and promoted in Macromedia Director, currently one of the most advanced animation packages available on microcomputer systems. It is important to note that the terms "advanced" and "sophisticated" refer to the software, not necessarily to the skills needed by the user. In fact, the "holy grail" of software design is to maximize both the software simplicity and capabilities. The past history of software on all computer systems has usually meant that the most powerful software was also the hardest to learn and use. I take the position that there is no reason not to expect powerful software that is also easy to use. Using a good analogy, such as producing a "movie" with a "cast" and "score," is one strategy to design software that's easy to learn. Analogies are a powerful instructional tool in general, but their use has been a particularly successful strategy for computer hardware and software design. This issue will be discussed in more detail in chapter 7.
8. One of the few microcomputers to use the standard NTSC video signal is the Amiga, and, for this reason, it is a favorite among videographers.

CHAPTER 4

Psychological Foundations of Instructional Graphics

OVERVIEW

This chapter reviews several major learning theories related to using graphics in instruction. While not pretending to be a substitute for a more thorough description of these psychological foundations, this chapter should provide a substantive review of major points and issues that should be considered and understood when designing graphical displays for instruction. Behavioral and cognitive learning theories are reviewed and compared, as are theories related to visual perception. Some of the topics include perception, attention, memory, and motivation. Particular attention is given to the theory underlying computer animated displays.

OBJECTIVES

Comprehension

After reading this chapter, you should be able to:

1. Describe the major features of behavioral and cognitive learning theories.
2. Define visual perception and visual cognition.
3. Summarize and compare major theories of how visual information is stored in memory, such as dual coding theory and propositional theories.
4. Describe the phenomena of apparent motion and how it relates to computer animation.

Application

After reading this chapter, you should be able to:

1. Generate a list of instructional graphic design principles derived from behavioral learning theory.
2. Generate a list of instructional graphic design principles derived from cognitive learning theory.
3. Given a computer display containing static or animated graphics, generate a hypothesis related to its effectiveness as an aid to learning.

The purpose of this chapter is to provide an overview of some of the psychological foundations related to learning from graphic displays. This chapter is solely theoretical in nature. It is easy to get the feeling sometimes that theory just gets in the way of instructional design. However, if the goal of instructional design is to affect learning, it seems reasonable that knowledge about learning and cognition should help in making appropriate instructional design decisions. This chapter will present some general themes related to psychological processes and try to relate these to the design of instructional graphic displays. A firm theoretical grounding in the psychological foundations should help to explain and predict some of the conditions under which graphics (static and animated) support learning, as well as those conditions under which graphics do not support or are detrimental to learning. These principles extend to all instructional environments involving graphics, including multimedia.

LEARNING THEORY: A PRIMER

Learning theory has had considerable influence on instructional practice in general, and computer-based instruction in particular. Main ideas of the two dominant classes of learning theories — behavioral and cognitive — will be briefly presented. Each makes qualitatively different assumptions about how people learn and remember.

Behavioral Learning Theory

The design of computer-based instruction has gone through an evolutionary cycle. One of the strongest influences on instructional practice in America in this century has been behavioral learning theory and its applications, such as programmed instruction. Although a definite conceptual shift toward cognitive views of learning has occurred in instructional technology, behavioral designs are likely to continue to dominate instructional practice.

Behaviorism is founded on the formation and strength of **stimulus-response (S-R) associations** (Gropper, 1983). An instructional stimulus is presented, such as a screen containing computer text or graphics, that prompts the learner to respond in an overt, observable manner, such as by typing an answer. The relationship between the stimulus and the response is strengthened through the use of reinforcement.

Operationally, the *stimulus* can be defined as either the material to be learned or the instructional event that leads to the learner's initial response. One of the basic goals of behavioral methods is to attain a measure of control or predictability of given instructional stimuli. The repeated use of particular stimuli helps to establish predictable control of learner responses, placing student responses under a form of *stimulus control*. Initially, for example, learners have no innate reason to press the <SPACEBAR> on a computer keyboard or to aim and click on a screen icon, but such responses can be readily elicited, thereby placing the desired response under stimulus control. Through systematic stimulus control, responses can be shaped through the presentation of directions or the repeated presentation of the same stimulus requiring the same response, paired with appropriate reinforcement for the desired response.

The *response* is the learner's overt behavior made in response to the instructional stimuli. The learner's response is the only recognized behavioral link to the instructional stimulus. Therefore, it is crucial that a causal link between stimulus and response be established in order to evaluate the effectiveness of the instructional stimulus. It is important that the learner clearly understand what response is required in order for the S-R bond to be formed. The response must be judged as appropriate or inappropriate in clear, objective terms, and this information is then conveyed to the learner as feedback. Instructionally, responses are elicited through the presentation of instructional activities and shaped through the presentation of systematically controlled activities, responses, and response consequences.

After a response is made and judged, a follow-up stimulus is presented to the student as a consequence to the response. This learner-activated stimulus, or **reinforcement**, is applied systematically to strengthen desired responses and is chosen conditionally on the desirability of the response. A key element of reinforcement is the principle of contiguity, where reinforcement is given closely in time with the response to be strengthened (Gagné & Glaser, 1987). Reinforcement for learner responses can be positive or negative. Both positive and negative reinforcements are designed to increase desired responses (Kazdin, 1980). Correct answers typically receive positive reinforcement in the form of statements such as “Good job!” or “Super!”; incorrect answers receive negative reinforcement such as “You missed that one” or “Your answer is incorrect.” Presumably, learners seek to avoid negative reinforcement by striving for appropriate or correct responses during response opportunities that follow. In the absence of the overt reinforcement available during instruction, behaviors are thought to be maintained by the intrinsic rewards of success.

Whether or not reinforcement actually is “reinforcing” to a student is determined by trial and error. Obviously, things that are strongly reinforcing to one person may be totally ineffective for another. A child who is a fan of superheroes, such as Superman or Batman, should find such graphics as reinforcing for correct responses and would be expected to seek to make more correct responses on the condition that similar graphics will follow.

Instructional designers often fall prey to forgetting that care should be taken to reinforce only correct responses. Fanciful graphics, used to add cosmetic or affective appeal to the learner, can be harmful when used as feedback to incorrect responses. The worst-case scenario is when the consequence of a wrong response is actually more reinforcing than that provided for a correct one. A good example is from an actual CAI lesson on weights and measures, where a student is asked to convert gallons to quarts. If the student responds correctly, for example, that four quarts equal one gallon, appropriate (though boring) praise is given, such as “That's correct.” However, if the student answers three quarts, the well-intentioned lesson shows, via animation, a gallon jug pouring its contents into the three awaiting quart containers with the last quart being spilled onto the floor. If the student finds a graphic of spilled milk more reinforcing than the praise, you can wager that more incorrect responses will follow.

Reinforcement and feedback are commonly misinterpreted as being synonymous. Though they may be used similarly in practice, there are several important distinctions between them. Feedback generally includes information related to the accuracy of a response, with the purpose to guide the student to make correct answers (Kulhavy, 1977; Schimmel, 1988). Feedback can be instructive without necessarily increasing desired responses. Reinforcement, on the other hand, can increase the probability of desired responses without necessary ties to the substantive requirements of a response. For example, information that smoking causes cancer may not actually lead a spouse to stop; however, always leaving the room when he or she lights up, even without explanation, may lead to decreased smoking. It is possible, therefore, to provide informational feedback that does not necessarily increase the probability of desired responses; likewise, it is possible to provide reinforcement that offers little informational feedback. Feedback that is reinforcing yet informational satisfies both behavioral and cognitive descriptions.

Though both feedback and reinforcement are continuously applied during typical CBI lessons, this practice may be both unnecessary and inappropriate. Reinforcement instead can be provided intermittently where only selected occurrences of a response are reinforced. The purposeful application of reinforcement schedules is important, since one goal of instruction is to decrease the amount of sustained reinforcement needed to elicit and maintain desired responses (Reynolds, 1968). Intermittent reinforcing is believed to be much more powerful than continuous reinforcement. A good example is a gambler at a slot machine. Winning only occasionally provides strong incentive to linger a little longer.

In order for instruction to be effective in most settings, the learner must respond to stimuli that are not necessarily identical to one another. In most cases, it is impossible to anticipate and teach all variations under which instructional stimuli might be encountered. It is essential that learners first be able to *discriminate* among classes of stimuli and then *generalize* relevant stimulus features and attributes to include appropriate alternatives. For example, though the short vowel word “cat” may be presented via individual worksheets, the student is expected to read “cat” from books, pet food boxes, and computer displays. Ambiguous cues can force the student to incorrectly discriminate or generalize. An illiterate shopper may choose a lemonade brand based on the picture of a lemon on the label, only to find out later at home that dish soap was purchased by accident.

During instruction, numerous stimuli are presented. Part of the learner’s task, therefore, is to selectively identify relevant stimuli from irrelevant stimuli. Graphics are just one source of visual stimuli competing for the student's attention. What determines which stimuli will be attended to and which will be ignored? The answer involves the issue of **selective attention**, a phenomenon that has been studied under both behavioral and cognitive frameworks. From the behavioral point of view, when two or more stimuli are provided, the learner will select the one that most easily results in the correct response. This is known as the “principle of least effort” (Underwood, 1963). A typical behavioral task would be to provide a student with a card displaying both the word and picture of an object, such as a cat. If the student's task is to identify the object, the principle of least effort predicts that the learner's attention will focus predominantly on the picture instead of the word in order to achieve a correct response (Samuels, 1967). In learning how to read, of course, the goal would be to shift stimulus control from the picture to the word.

Several behavioral techniques can strengthen desired S-R associations, such as showing the printed word and saying the word. **Cueing** helps learners to interpret complex stimuli by providing contextual “hints.” Shoppers who buy dish soap instead of lemonade would benefit from paying attention to which supermarket aisle they are in, for example. **Prompting** refers to supporting instructional features that amplify critical stimuli features, such as important words or concepts (Hannafin & Peck, 1988). In effect, prompts emphasize relevant stimuli by explicitly directing learners to relevant aspects of the lesson. In practice, however, stimulus-response associations should not be dependent upon prompts that will be unavailable under actual performance conditions. Therefore, prompts should be progressively *faded* from instruction to ensure that responses can be elicited by appropriate stimuli alone. Responses are then *shaped*, both individually and in complex chains, to meet predetermined response requirements.

The principle of least effort requires instructional designers to be wary of providing any stimuli, such as pictures, that may compete for the learner's attention. If the picture is perceived as providing the information necessary to respond correctly, it will dominate which S-R associations are formed and which are neglected. If a student learns beginning reading skills with the use of pictures, for example, the strength of the S-R association of the picture and the task may be stronger than other S-R bonds. Therefore, if a picture is present, a learner may defer to the picture by default. From a behavioral point of view, therefore, graphics can be potent stimuli resulting in both appropriate and inappropriate learning. Graphics can provide the foundation for strong S-R associations or they can be disruptive or cause interference (Willows, 1978). Research has shown that learners vary in their susceptibility to interference by graphics. For example, learners with poor reading skills seem particularly vulnerable (Samuels, 1967).

Although S-R associations are the basic unit of analysis in the behavioral model, few tasks involve isolated stimuli or address only simple responses for simple associations. In practice, individual associations are linked together to produce networks of interdependent S-R *chains*. Successful learning almost always requires complex sets of related S-R events to collectively guide responses. The learning of complex tasks and problem solving is explained through chaining, a point frequently challenged by cognitivists.

Strict behavioral applications to learning frequently come under criticism. For example, behavioral designs, though usually considered quite effective for lower-level learning such as that associated with verbal information (see chapter 2), tend to be insufficient for higher-level learning such as intellectual skills (which include problem-solving).

Cognitive Learning Theory

In contrast to focusing on strengthening S-R bonds, cognitive orientations to learning consider the actual thought processes occurring in between the stimulus and the response as the most important aspects to learning. The emphasis is on how a learner selects, perceives, processes, encodes, and retrieves information from memory (Di Vesta, 1987).

The Information-Processing System

There are many aspects to cognitive psychology. However, almost all recognize some model of human cognition based on **information processing**. Information processing models are just that — models. They are not meant to describe learning in any physiological way. Instead, these models provide a computer analogy to help understand the learning process by suggesting that the processing of information by humans is like that of a computer. Information processing provides a vocabulary for describing mental events, as well as explaining why learning does or does not occur. A standard criticism of information-processing theories is that they generally focus on semantic learning and do not take into account social and emotional aspects of learning. However, this criticism is more related to the application of the models, rather than the models themselves. There is no reason why an information-processing model cannot be extended to serve these functions as well. It is also important to recognize that the S-R model of behaviorism is subsumed in information processing. Of interest, however, is what happens in between the stimulus and the response; in other words, what goes on between one's ears.

Information-processing models describe learning as a series of knowledge transformations, starting with the input of information (stimulus) from the environment, and ending with either an output (response) or the storage of the information in memory, or both (Dodd & White, 1980; R. Gagné, 1985; E. Gagné, 1985). These transformations require a progression of nonobservable, mental steps for learning to occur.

Figure 4.1 shows a standard form of the information-processing model. First, the learner must filter the large amounts of stimulation coming from the environment. This information is transformed by the senses into neural information, where it exists briefly in the **sensory register** (Sperling, 1960). However, the amount of information available to a learner at any one moment in time can be enormous. The human response to this bombardment of information is known as **selective perception**, which is the process where only a small portion of the incoming stimuli will be given any consideration at all. Successfully ignoring all stimuli except the most pertinent and relevant is an extremely important ability, without which life would be unbearable. Because of selective perception, the bulky sensation of a winter coat soon fades away. Even the most subtle stimuli, such as the stray floating particles in the vitreous humor of the eye, called nuscae volatantes (Latin for “flying gnats”) or more commonly “floaters,” would soon overwhelm your attentional processes if your perceptual system did not filter them out.

The main instructional issue related to selective perception is attention. Current theories of attention have evolved from earlier ones based on mental “filters” (Broadbent, 1971) and the limited capacity of many mental operations (Norman & Bobrow, 1975). Attention involves cognitive decisions related to which information will be attended to, given the fact that the environment contains far more information that any one person can handle at any given time. Information-processing theories describe attention as consisting of a general sequence of stages, some of which are made subconsciously, such as sensory input, and others consciously, such as selective attention as introduced in the last section. Selective attention, therefore, implies the element of intentionality in focusing on one set of information while blocking other incoming information. In this way, selective attention serves a gate-keeping function. These perceptual and attentional processes are believed to be influenced by both the intensity of the input and expectations based on prior knowledge — these are known as bottom-up and top-down processing, respectively (Anderson, 1980). Research has shown that people are often predisposed to select information based on physical characteristics (e.g., color and motion), as well as information that is novel or unique (see Dodd & White, 1980, for a review). Visual or aural intensity, externally provided cues, prompting, and organization of the material are among the many presentation factors that can facilitate attention.

Research has shown that attention is naturally drawn to the novel or unique (Fleming, 1987). Obviously, the use of graphics can be an important strategy for influencing attention, but only if the graphics are used deliberately in novel or unique ways. Graphics in a lesson saturated with illustrations would, of course, soon lose their attention-gaining capability, as novelty effects eventually wear off. Attention-gaining graphics should be used judiciously in order to optimize their effectiveness.

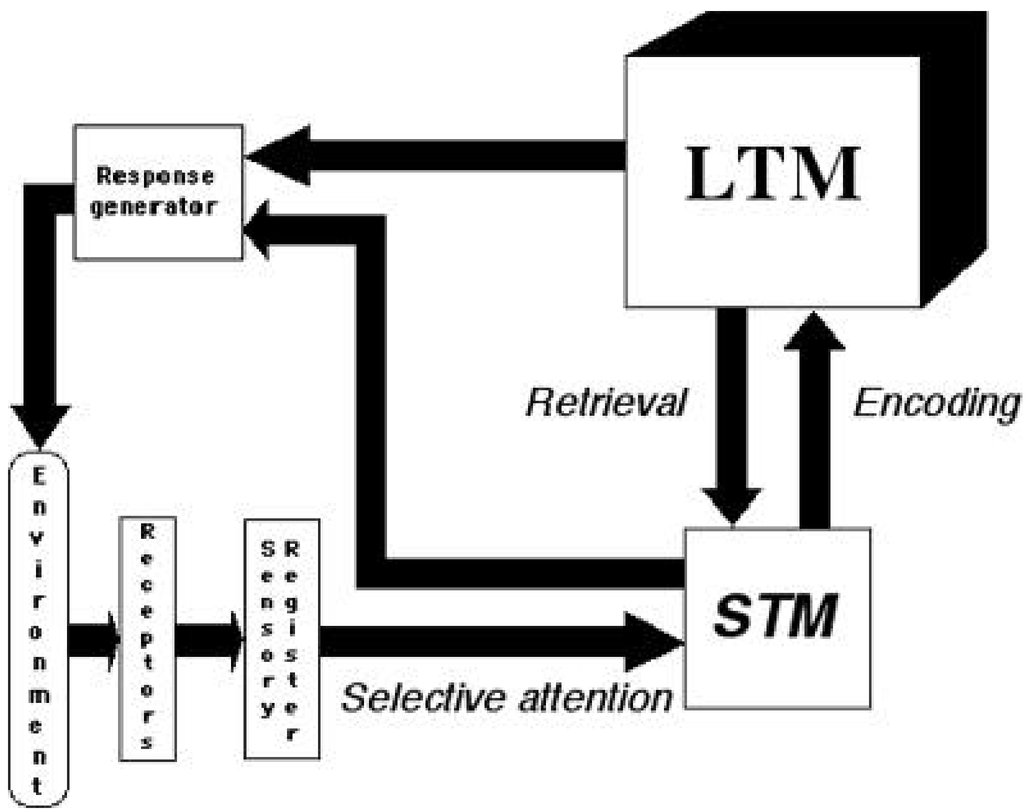

FIGURE 4.1
The information-processing system.

Selected information is stored temporarily in **short-term memory** (STM), or **working memory**. STM acts essentially as a buffer containing information units to be acted upon. It does not, by itself, serve as a permanent storage location, but instead acts as a “broker” for the selective exchange of information from instruction, prior knowledge, and long-term memory. STM is analogous to a computer's random access memory (RAM). Unlike computers, however, STM has a severely limited capacity. Although the limits are debatable, STM is generally considered to hold five to nine informational units. Each unit corresponds roughly to one idea or expression (Miller, 1956). Some researchers have proposed strategies, such as mnemonics or meaningful "chunking" of information, to increase the functional capacity of STM (Chase & Ericsson, 1981). For example, memory for the phone number 355-1224 would be facilitated by noticing the mathematical coincidence that 24 is two times 12 in the final four digits. This would reduce the burden on STM from the seven original informational units (roughly one unit per digit) down to about five.

Again, graphics offer the *potential* for an efficient means for coping with the precious and limited processing capabilities of STM. Similar to the principle of least effort, students are likely to revert to a given graphic because they are apt to perceive the amount of invested mental effort (Salomon, 1983) as being less for the graphic than for the surrounding text, even though this may not be the case. If this would happen, the graphic would be serving to subvert the usefulness of the text by serving as a distraction. On the other hand, graphic organizers can much more efficiently represent information when compared to presenting pure verbal information, such as the spoken or written word.

If information is to be learned, it must be transferred from STM to **long-term memory** (LTM) — a permanent mental storage location analogous to computer disks or magnetic tape. In order for this transfer to be successful, the information must be coded. The retrievability of knowledge afterward is directly related to the manner in which information is coded. From an instructional perspective, however, it is important that successful strategies are applied to strengthen the coding of relevant content, since poorly coded information can be discarded at any point prior to successful storage in LTM (E. Gagné, 1985).

Encoding and Retrieving Information to and from Long-Term Memory

Although controversial, many cognitive psychologists believe that once information is successfully stored, or encoded, in LTM, it is never lost. Subsequent problems in remembering or recalling encoded information is believed to be a matter of deficient retrieval strategies, rather than simply of forgetting (i.e., storage deterioration). Successful retrieval of information from LTM, therefore, is dependent on both the quality of initial encoding into LTM and the methods governing retrieval — both those supplied by the lesson and those triggered solely by individual learners.

Beyond the processing system, and central to the likelihood of retrieval, is the manner in which information is stored within LTM. Knowledge is believed to be stored in LTM in a variety of mental representations, including propositions, productions, and images. A **proposition** is the smallest, single information unit, corresponding generally to an idea. Declarative knowledge consists of propositions. It is useful to consider propositions as simple idea units, rather than as actual words or sentences (Wanner, 1968). **Productions**, on the other hand, represent procedural knowledge. Declarative knowledge can be thought of as knowing *about* something, whereas procedural knowledge is knowing *how* to do something. Productions are conditionally based action sequences that are executed under highly specific conditions, similar to an *if-then* condition programmed into a computer. They are linked together in production systems where one production leads directly into the next production. A simple example of a node-link representation involving propositions and productions is shown in Figure 4.2. Functionally, productions and production systems provide a parallel to the S-R associations and S-R chains aspects of behaviorism, though their conceptual roots and underlying assumptions are quite dissimilar.

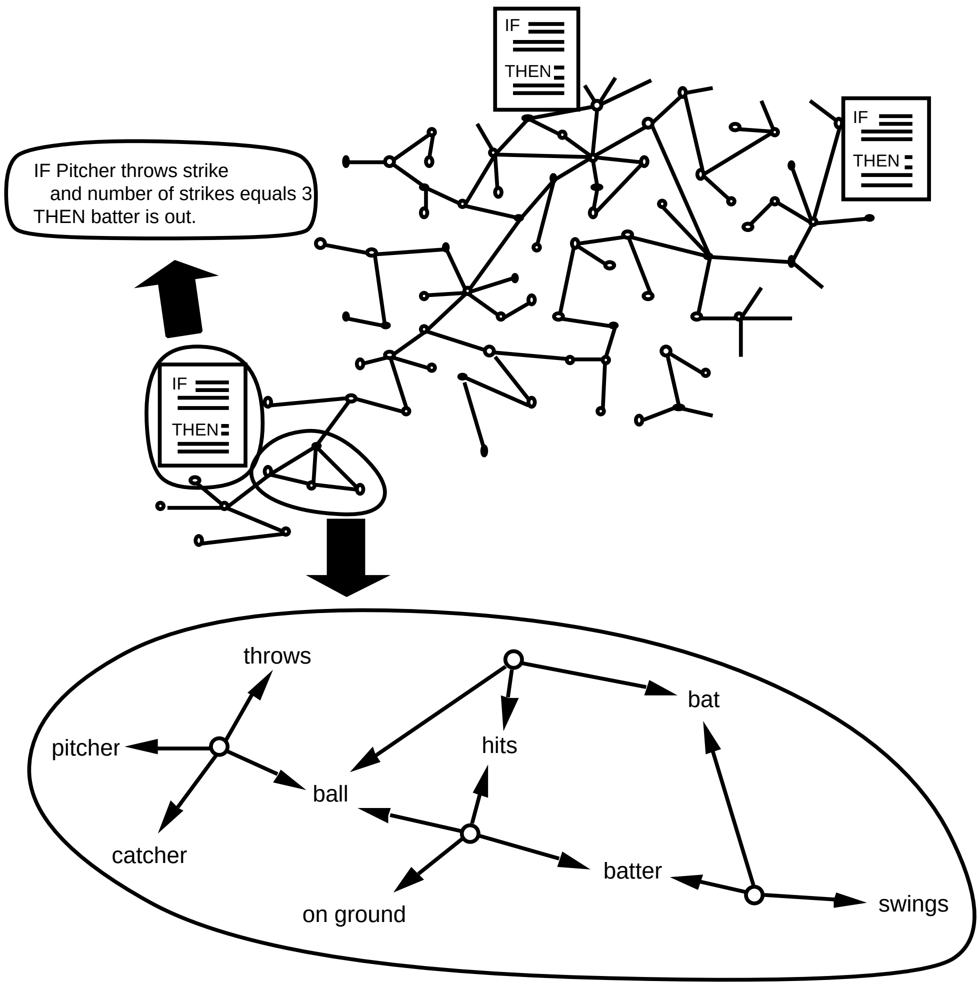

FIGURE 4.2
An example of a propositional network.

Schema is another psychological construct, meant to be more metaphorical than physiological, that represents an individual's entire organized knowledge network (Norman, 1982), as well as the representation of that knowledge (Rumelhart & Ortony, 1977). Schema theory has been used to describe interrelationships among prior knowledge structures. Just as a play requires characters with defined roles and associations, schemata has certain variables and slots that must be filled for meaning to exist (Rumelhart & Ortony, 1977).

Each individual develops and refines a multitude of individual schemata; each schemata is instantiated (triggered, then enacted), in order to understand or make sense of complex situations. Furthermore, schemata systematically access one another. Examples of common schema are the “scripts” associated with going out to eat at a restaurant, riding a city bus, and going out on a date.

Schema theory refers to two main types of mental processing: top-down and bottom-up. Bottom-up mental processing begins with isolated facts that eventually instantiate a schemata. Top-down processing begins with a schemata leading a person to search for appropriate information to satisfy or complete the schemata. An everyday example of top-down processing is seeing a person walking away from you down and across the street in a way that looks familiar. If the walk triggers the schemata of “my good friend Joe,” you might yell across the street to the person in a very friendly way, just to find out when the person turns toward you that he is a total stranger (making you feel quite foolish as a result). In this case, top-down processing took over inappropriately. Analogies can be one way to capitalize on top-down processing in instruction. When students are given instruction in an area or domain in which they are not familiar, analogies can provide a bridge between where the students are and where the instruction is going to take them, assuming that the students understand the analogy (Newby & Stepich, 1987). As described in chapter 2, graphics can be a very appropriate basis for presenting analogies to students.

Three processes are crucial to the encoding and retrieving processes (E. Gagné, 1985): **elaboration, organization,** and **spread of activation**. *Elaboration* is a process whereby supporting information is added to the information being learned. Elaboration occurs when an individual uses knowledge already stored in LTM to enhance, extend, or modify new information while in STM, as well as during subsequent transfer into LTM. In the propositional network model, elaborations provide links between previously stored propositions and new information. Effective elaborations serve to combine or link related propositions to stimulate retrieval of the learning context; ineffective elaborations do not (Reder, 1982; see also Wang, 1983). Several simple, but effective elaboration strategies have proven effective in improving retrieval, such as, again, the use of analogies (Hayes & Tierney, 1982). Linden and Wittrock (1981) reported that simply reminding learners to elaborate during learning helped to facilitate retrieval.

Whereas elaboration affects the storage of information, *spread of activation* increases linkages among related propositions and nodal links in the propositional network. It is a process whereby a given active proposition passes activation along to related propositions. Isolated information in LTM is not easily activated, since few paths to other propositions exist. Certain information is difficult for individuals to recall because retrieval is dependent on a limited number of very specific prompts or cues. In order for activation to spread, direct links must be established between the propositions. Several researchers have advocated the use of a type of visual networking of information in textual materials, also known as spatial mapping, as an instructional technique (see studies by Chi & Koeske, 1983; Holley & Dansereau, 1984; and Novak & Gowin, 1984).

In addition to elaboration and spread of activation, successful learning requires dividing and organizing information into subsets. People attempt to organize information automatically and spontaneously (Reitman & Reuter, 1980), though success and efficiency vary widely for retrieval purposes. *Organization*, the intentional shaping of information into meaningful parts, plays a key role in effective retrieval of learned materials (Frase, 1973; Meyer, 1975; Thorndyke, 1977). Effective organization may provide additional pathways among network nodes in a manner similar to elaboration. Organization might also help overcome some of the inherent memory limitations of STM. The organization of information into subsets might also help provide *pointers* in STM that help individuals manage large amounts of interrelated information in LTM (E. Gagné, 1985; Glynn & Di Vesta, 1977).

The Role of Prior Knowledge

Prior knowledge is paramount to all aspects of cognitive psychology (see, for example, Di Vesta, 1987; Mayer, 1979; Shuell, 1986; Tobias, 1987). Its significance is probably best summarized by David Ausubel (1968, epigraph): “If I had to reduce all of educational psychology to just one principle, I would say this: The most important single factor influencing learning is what the learner already knows. Ascertain this and teach him accordingly.” Ausubel (1963) maintained that the availability, organization, and strength of existing supporting cognitive structures are the foremost factors governing the meaningfulness of newly learned material, as well as the ease and efficiency of acquisition and retention.

For cognitive psychologists, prior knowledge also provides the potential for supporting schema related to forthcoming instruction, improved capacity for comprehension monitoring (Garhart & Hannafin, 1986), and individual lesson choices (Gay, 1987), and the capacity for elaboration and meaningful learning. The availability of related prior knowledge permits the learner, to a degree, to uniquely define information needs. When supporting knowledge exists, learners gain the capacity to compare and contrast to-be-learned instruction within existing knowledge, providing uniquely relevant elaboration unavailable to learners with limited prior knowledge. Consequently, lesson knowledge generally will be encoded more meaningfully and retrieved more successfully by learners with high versus low prior knowledge.

VISUAL COGNITION

Visual cognition includes all the mental processes involved in the perception of and memory for visual information (Pinker, 1984). Perception is the process of selectively attending to and scanning a given stimulus, interpreting significant details or cues, and, finally, perceiving some general meaning (Levie, 1987). Memory for visual information involves the cognitive processes of storing and recalling information from visual stimuli. It is difficult to pinpoint where perception ends and cognition begins. For this reason, we will deal with the issues of perception and memory independently.

Visual Perception

Visual perception is the process of being able to selectively attend to and then subsequently perceive some meaning from a visual display. All the senses are involved in perception, although the visual sense is usually stressed in most perceptual theories. Most people usually think of visual perception in terms of human physiology, or the “mechanics of seeing,” such as how the eye receives visual stimuli, and converts and transmits this information as an electrochemical signal along the optic nerve until it reaches the visual cortex of the brain, as that shown in Figure 4.3. However, visual perception is “more than meets the eye.” Certainly, the physiology that accounts for perception is remarkable; however, it's not the issue here. Instead, we are concerned with what happens to the information once it reaches the processing centers of the brain. Visual perception is far from an objective process and instead is based on all our previous knowledge and experiences. The use of prior knowledge to guide perception is known as knowledge-guided perceptual analysis, better known as top-down processing. Visual perception is not like taking a photograph with a camera.

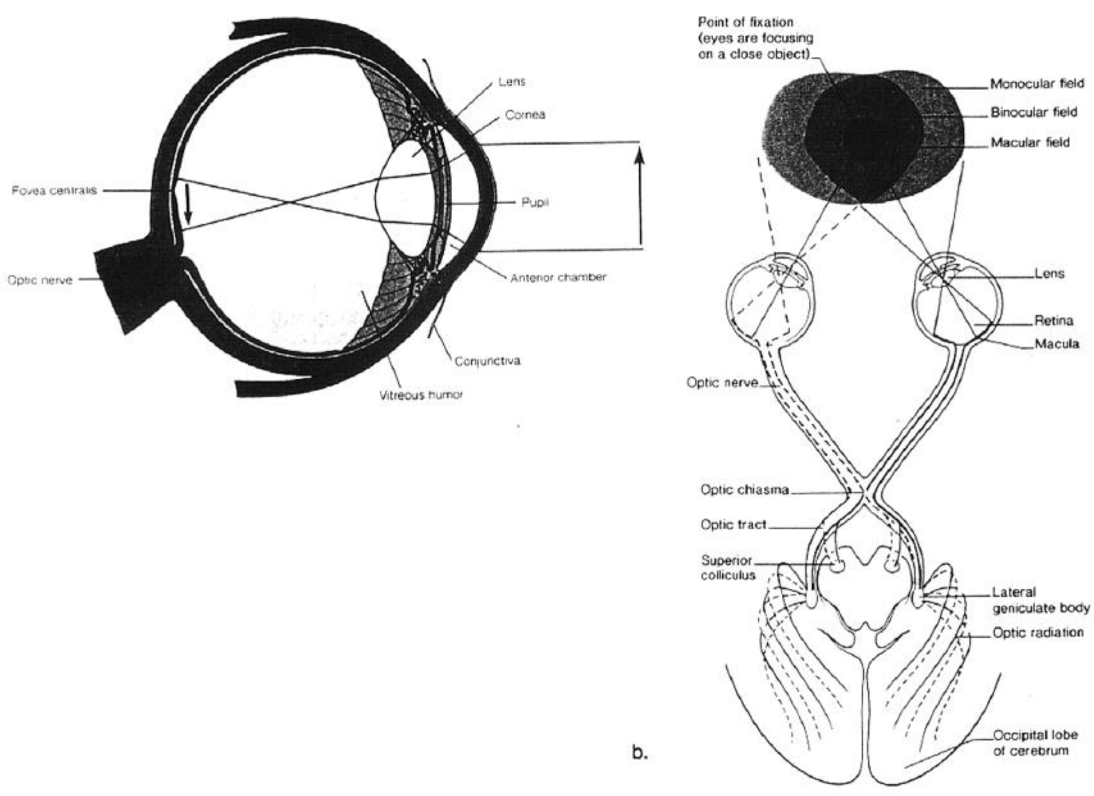

FIGURE 4.3
Some of the physiological mechanics of perception: the left eye and the visual centers of the brain (both views are from the top).

Visual perception is largely concerned with visually recognizing shapes and patterns of objects directly in our visual field. There are several traditional theories of pattern recognition, such as template matching and feature models. All of these traditional approaches assert that knowledge about the regularities of the world is used to limit the number of possible recognizable shapes from which the perceptual system can choose. Gestalt psychologists from the 1920s, such as Max Wertheimer, Kurt Koffka, and Wolfgang Köhler, were among the first to be interested in visual cognition. Whereas one might define perception on the basis of the individual elements (structuralism) in the visual display, Gestalt psychology defined a series of perceptual principles based on global characteristics. These principles are still useful in studying pattern recognition, where the total is more than the sum of the parts.

Four gestalt principles are still particularly relevant to designing instructional visuals. The first principle, **closure**, is based on the idea that humans naturally look for meaning. This principle accounts for the phenomenon of seeing dead presidents in fluffy white clouds.

Study Figure 4.4 for a moment and try to find a recognizable object. At first, all you might see is a random scattering of ink blotches. Here's a hint if you are having trouble seeing something further: look for an animal. Still having trouble? Look for a dog. Once people see the figure, they instantly recognize not only a dog, but a Dalmatian. Once you achieve recognition of this figure, begin to reflect on the elaborated meaning you begin to attach to it, such as a dog eating something, or retrieving a thrown stick.

FIGURE 4.4
An example of the principle of closure. Do you see a smattering of ink blotches, or something more meaningful? Read the preceding page for some clues.

Similarly, look at Figure 4.5. Again, what do you see? Most people instantly see a young girl of about 20 with her head turned away. Look again and you might see something completely different. Instead of a young girl, there is an old woman. The young woman's left cheek becomes the old woman's big nose. The interesting point from a perceptual point of view is that it is impossible to see both at the same time. Your perceptual system will not allow it. Even though you may switch rapidly between the two meanings, you cannot accept both simultaneously. Again, reflect on the meanings you might be tempted to attach to either meaning. Perhaps thoughts quickly went through your mind of the young girl going hurriedly off to meet her fiancé, or the old woman waiting in line for groceries. Your mind wants to attach meaning to visual displays.

FIGURE 4.5
Another example of the principle of closure. There are two distinct meaningful images here. Can you see both of them?

The other three gestalt rules for the organization of visual information are the principles of **proximity, similarity,** and **continuity**. The principle of proximity states that objects physically closer to one another will be perceived as being grouped together in some meaningful way. Look at the dots in Figure 4.6. One quickly judges that there are five columns of five dots each, simply because of the spatial distances. The principle of similarity simply states that similar objects also will be grouped together in a meaningful way. Look at the way the circles appear to be "surrounding" the one "lonely" cross in Figure 4.7. Figure 4.8 illustrates the principle of continuity: The line is perceived as unbroken and continuous because the mind looks for unity in objects.

People perceive meaning from animated visuals when they are tricked into seeing something that really is not there. Certain perceptual factors help to explain this phenomenon (Rieber & Kini, 1991).

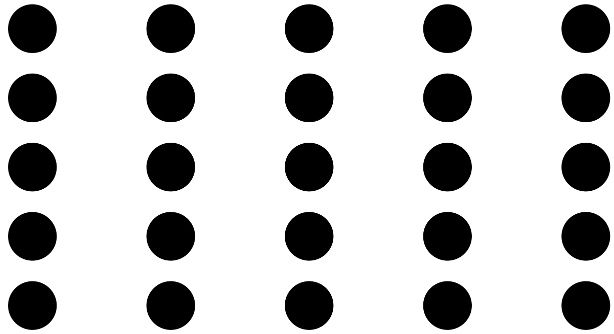

FIGURE 4.6
An example of the principle of proximity. The physical closeness of some of the dots make natural groupings of five columns.

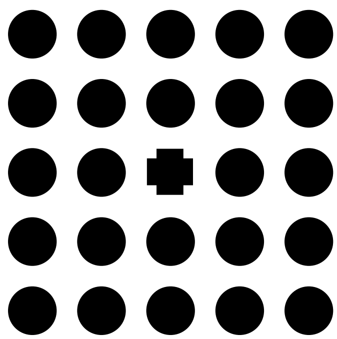

FIGURE 4.7
An example of the principle of similarity. The cross appears to be "surrounded."

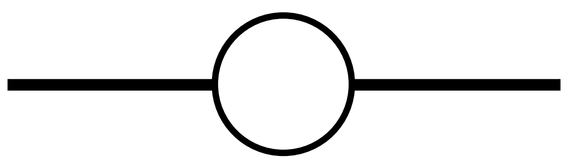

FIGURE 4.8
An example of the principle of continuity. The two line segments are perceived as one continuous line.

Perceptual Factors Related to Animation

Motion perception research has a long history. Animation is an example of **apparent motion**, or the phenomenon of seeing motion when there is actually no physical movement of an object in the visual field (Ramachandran & Anstis, 1986). In contrast, a person perceives **real motion** when a moving object actually triggers visual detecting neurons (Schouten, 1967). Apparent motion results when two or more static objects, separated by a carefully determined distance, are alternately presented to the observer over time. Even though there is no actual motion of the image on the eye's retina, the visual system perceives motion by combining this discretely presented information into a smooth and continuous set. When the conditions are just right, the mind fills in the gaps between the frames, resulting in the *perception* of continuous motion, even though it is only being confronted with a rapid series of *still* images. You see examples of apparent motion everyday, such as when moving arrows created by carefully timed neon lights try to attract your attention to a particular store.

The most intensively studied version of apparent motion is known as **stroboscopic** motion, or the motion perceived when two lights are presented at different times and different locations (also called beta movement) (Kaufman, 1974; Schiffman, 1976). Stroboscopic motion has a critical threshold of about 16 frames per second in order for it to be *perceived* as smooth and continuous; anything less results in choppy or jumpy displays. Even the most inexpensive computer systems available today easily match this critical rate.

The **phi phenomenon** is closely related to stroboscopic motion (Schiffman, 1976). This illusion of motion is produced when stationary lights are turned on and off in sequence. Examples include the use of carefully sequenced lights to create dynamic visuals around billboards, theater marquees, scoreboards in sports arenas, and other “Las Vegas-like” displays. The phi phenomenon also accounts for any animation on computer displays produced by the coordinated switching of pixels. Several perceptual theories help to explain both stroboscopic motion and the phi phenomenon. (It is interesting to note that other examples of apparent motion have been identified, such as the delta phenomenon, produced when the brightness of the stimuli is varied [William, 1981].)

There are many factors that determine the nature of apparent motion. However, three are particularly relevant to computer animation: (1) the time between projection of the separate displays; (2) the light intensity of the displays; and (3) the spatial distance between each of the displays. For example, no motion will be perceived when two lights are alternately presented at too slow a rate. Instead, one simply sees *two* lights being switched on and off. If the lights are alternated at too fast a rate, then the two points are perceived simultaneously and, again, no motion will be perceived. However, if the two lights are alternated at just the right speed, the perception that a *single* object is moving back and forth will be induced. The light intensity of the two points and the space between them must also be just right.

Although each of the three factors described above must be considered individually, “Korte's law” states that apparent motion can only result when these three factors are properly synchronized (Korte, 1915, as cited in Kaufman, 1974; and Carterette & Friedman, 1975). If one factor is held constant, then the other two factors will vary proportionally (though not necessarily directly). For example, if light intensity is held constant, the distance between the displays must vary proportionally to the amount of time between each display for apparent motion to be produced. So, as the displays are moved closer together, the time between the projection of the displays must also decrease. Many other factors, such as spatial orientation, depth, color, size, shape, and texture, also impact the perception of apparent motion, but the influence of each of these factors is considered minimal.

In order for a group of intermittently displayed objects to be perceived as *one* object in continuous motion, the visual system must trigger a psychological process called **correspondence** (Mack, Klein, Hill, & Palumbo, 1989; Ramachandran & Anstis, 1986; Ullman, 1979). In this process, the brain imposes *one* organized and meaningful pattern to the separate images. For example, an animated scene of a person riding a bicycle across the computer screen may actually consist of 30 separate frames. Although the visual system "sees" all 30 frames, an individual *perceives* only one object in motion. The correspondence process cognitively "assembles" the separate images into one meaningful set; motion becomes the "glue" in this assembly. Explaining how the visual system achieves correspondence detection is controversial in visual cognition research. Perceptual psychologists have shifted from the view that apparent motion is the result of a single psychological process over the years to a two-process theory where two distinctly qualitative processes are at work: long-range and short-range apparent motion (Braddick, 1974, as cited in Petersik, 1989; Julesz, 1971). Motion is perceived using either or both of these processes. (Research indicates that there is competition between the two processes when each is equally stimulated. Researchers have not yet adequately determined the conditions necessary to trigger one or both processes [Pantle & Picciano, 1976].)

The short-range process of perceiving motion relies on the brain solving a complex mathematical matching game. First, a display stimulates the retina and the image is transformed into an array of tiny points of varying brightness. Second, this image is converted into an electrical signal and is carried along the optic nerve to the brain. Third, the brain compares each point with corresponding points in each successive display and determines that *one* set of matched points composing a *single* object has changed its position. The way the brain achieves this complex mental computation is not yet understood (Petersik, 1989).

The long-range process of perceiving motion provides a very different explanation. It holds that the visual system uses strategies that limit the number of matches the brain needs to consider in order to avoid the need for a complex point-to-point comparison (Ramachandran & Anstis, 1986). Instead, it is believed that the visual system uses a number of special strategies evolved over thousands of years. Given the regularities of the physical world, the visual system extracts the most relevant features from a complex display (such as clusters of dots rather than individual dots) and then searches for those features in the successive images. The features could be edges, short outlines, blotches of brightness and darkness, or texture. The visual system limits motion perception to events consistent with universal physical laws. Long-range motion perception is based on the assumption (usually at a subconscious level) that the physical world is organized and predictable, not chaotic and random. Such assumptions lead the visual system to perceive one kind of motion in "preference" to other kinds. For example, linear motion is perceived over abrupt changes.

The visual system also anticipates the need for objects to overlap at times; therefore, the unseen objects still will be perceived to exist (Ramachandran & Anstis, 1986). This is similar to the principle of continuity described earlier.

In summary, apparent motion explains how the illusion of motion is produced from a wide assortment of media, such as motion pictures (including film, videotape, laser disc, and television), theater marquees, scoreboards, billboards, and animation produced by computer. Smooth and continuous motion perception is largely dependent upon the rate of picture presentation. The next time you watch a movie, try to remember that your attention to the plot, script, and acting would not be possible without the projector or player first displaying the individual photographic frames at the proper rate, followed by your visual system fusing each successive frame into one continuous and integrated set.

Memory Considerations for Visual Information

Active visualization in short-term memory is an accepted, though not universal, phenomenon. Strong evidence comes from classic studies where subjects were asked to mentally rotate letters and other three-dimensional objects (Pinker, 1984), such as the example described in Figure 4.9. This research has shown convincingly that the time it takes people to determine whether a given rotated letter is normal or mirror-reversed increases depending on how far it deviates from an upright position (e.g., Cooper & Shepard, 1973). Similarly, determining whether 2 three-dimensional objects have the same shape has been shown to be a linear function of how different their orientations are represented. A simple example is how people determine moderate to complex compass directions, such as in the activity in Box 4.1. Consider how you use visualization techniques to construct meaning for phrases such as “Northeast.”

In contrast to visualization in STM, the question of if and how visually based information is stored in and retrieved from long-term memory is more controversial. Retrieving information from long-term memory to produce internal visual images in short-term memory can be described as “the process of remembering or reasoning about shapes or objects that are not currently before us but must be retrieved from memory or constructed from a description” (Pinker, 1984, p. 3). There is considerable evidence that, in general, memory is greater for pictures than for words. People are particularly adept at remembering certain kinds of visual information. Many studies show that people's recognition memory for pictures is extraordinary. For example, Shepard (1967) showed 612 different pictures to people. When tested immediately afterward, subjects correctly recalled more than 98% of the pictures. Even a week later, subjects recalled more than 85% of the pictures. These kinds of results were repeatedly found in similar tests of recognition for visual information, such as photographs and faces (e.g., Nickerson, 1965; Standing, Conezio, & Haber, 1970).

Two competing theories about this “picture superiority” effect have been offered. The first, called **dual coding theory**, proposes that long-term memory consists of two distinct, though interdependent codes, one verbally based and the other visually based (Paivio, 1990, 1991; Clark & Paivio, 1991). The second theory suggests that *all* information is represented by a single propositional coding model. Since neither theory is physiological in nature, the arguments for and against each theory are useful only in finding a reasonable model that helps us to better understand thought processes and make better instructional decisions. Useful perspectives can be gained from considering both dual coding and propositional theories. Of the two theories, however, dual coding has considerable empirical support (Anderson, 1978) and is the theory favored here. The next two sections describe dual coding theory and contrast it to the alternate propositional hypothesis.

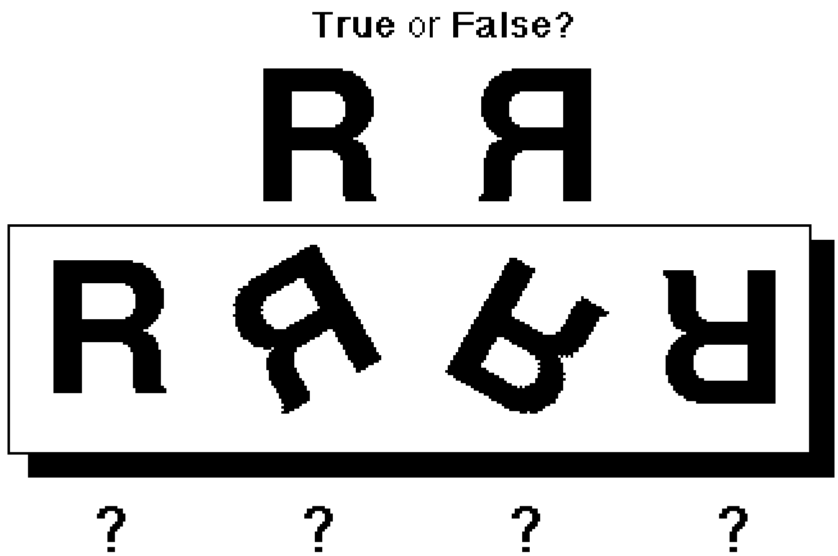

FIGURE 4.9
Here is an example of the mental rotation research that provides strong evidence for visualization in short-term memory. The activity simply asks subjects to respond "true" if the letter shown is in its normal position, although rotated, or "false" if the letter is shown as "mirror-reversed." The letters that are rotated more from the original position consistently take people longer to respond, suggesting that people must mentally rotate the figure back to the upright position before being able to respond with 100% confidence.

Box 4.1

"You Are Here": Visualizing in Short-Term Memory

Here's a little map activity which requires you to actively visualize in short-term memory. Below is a map of Texas with letters of the alphabet scattered about. See how long it takes you to answer each of the following questions under the condition that you *must* be correct each time. As you answer the questions, take note of the visualization techniques you use to get the right answer. Also consider which questions took you the longest and shortest time to answer and why.

1. Which letter is southwest of Austin?
2. Which letter is furthest south in the state?
3. Which letter is northeast of Dallas?
4. Which letter is southeast of San Antonio?
5. Which letter is in the northeast corner of the state?
6. Which letter is in the southeast corner of the panhandle?
7. Which letter is furthest west and in the state?
8. Which letter is east of Houston?
9. Which letter is west of Dallas?
10. Which letter is directly north of San Antonio?
11. Which letter is northeast of Brownsville?

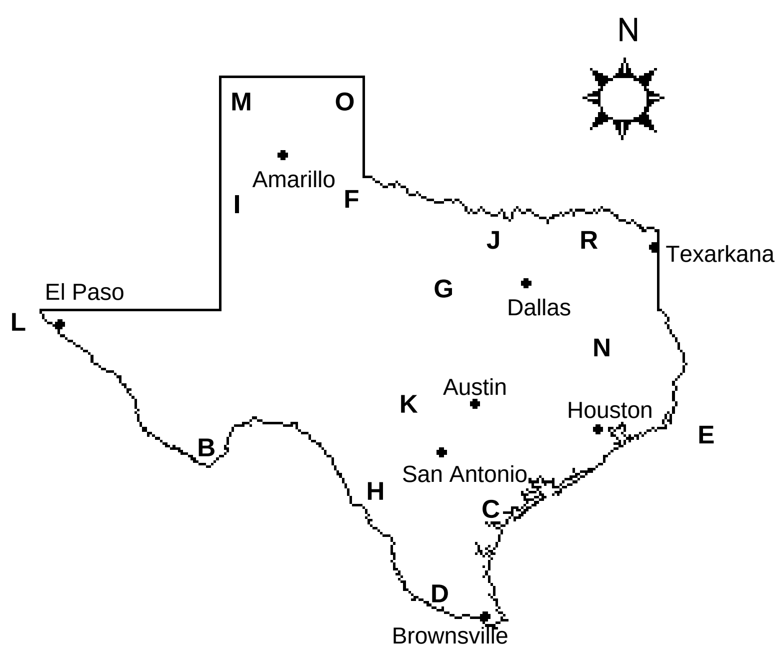

An Overview of Dual Coding Theory

Dual coding theory is a complete set of assumptions and hypotheses about how information is stored in memory (Sadoski, Paivio, & Goetz, 1991). Dual coding theory suggests that memory consists of two separate and distinct mental representations, or codes — one verbal and one nonverbal. The verbal system is “language-like” in that it specializes in linguistic activities associated with words, sentences, and so on. Although the nonverbal system includes memory for all nonverbal phenomenon, including such things as emotional reactions, this system is most easily thought of as a code for images and other “picture-like” representations (although it would be inaccurate to think of this as pictures stored in the head). For simplicity and clarity, we will refer to the nonverbal elements of dual coding as the visual system.

Dual coding supports the idea that knowledge is represented on a concreteness-abstractness continuum and that human cognition is predisposed to storing mental representations in one of two forms corresponding to the ends of the continuum. At one end are the visually based representations in which knowledge is stored in concrete and nonarbitrary ways. For example, the image of an airplane necessarily resembles characteristics of the real thing. At the other end are the verbal, or semantic, representations in which knowledge is stored in discrete and arbitrary ways. There is no natural reason that the word “airplane” must be used to represent the real object, for example. The most fundamental memory units are called **logogens** in the verbal system and **imagens** in the visual system.

Both the verbal and visual subsystems have unique properties. Logogens are stored in the verbal system as discrete elements, resembling words and sentences, whereas imagens are stored as continuous units in the visual system having an “all-in-oneness” quality. This is similar to the difference between digital and analog information. Although the two coding systems are assumed to be structurally and functionally separate, they can become interconnected. Informational units in one system can cue or trigger elements stored in the other.

As shown in Figure 4.10, dual coding predicts that three distinct levels of processing can occur within and between the verbal and visual systems: representational, referential, and associative. Representational processing describes the connections between incoming stimuli and either the verbal or visual system. Verbal stimuli directly activate verbal memory codes, whereas visual stimuli activate visual memory codes. For example, hearing the word “dog” first activates the verbal system, but seeing a picture of a dog directly activates the visual system.

Referential processing is the building of connections *between* the verbal and visual systems. Hearing or reading the word “dog” will stimulate the appropriate logogen in the verbal system. Subsequently forming a mental image of a dog, perhaps your own pet, implies that the verbal system has directly activated the imagen corresponding to your pet. The other direction is also possible, such as a person who is asked to name a given picture of a dog and says “collie.” An important assumption of referential processing is that these connections between the verbal and visual system are not necessarily one to one, but can be one to many. Hence, showing a picture of a dog may invoke many different verbal responses, such as “animal,” “dog,” “collie,” or “Rover.”

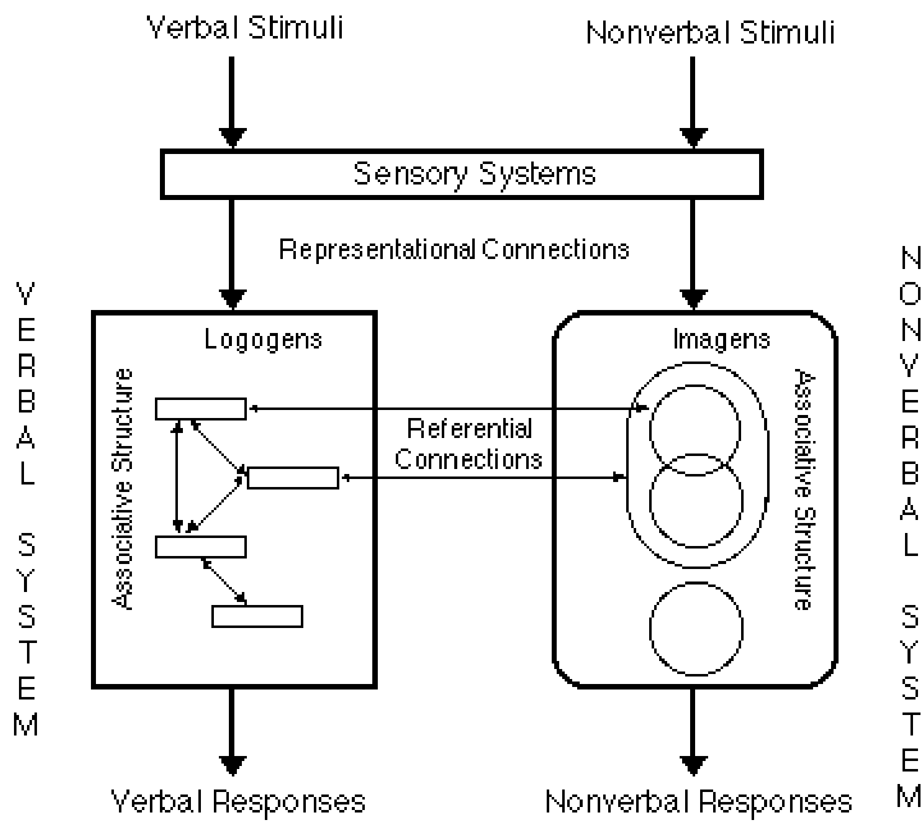

FIGURE 4.10
A dual coding model for memory and cognition.

However, the example of a single picture evoking many different verbal labels can be explained two ways. First, the different responses may be the result of multiple links between the single imagen of the visual system and the many logogens of each verbal representation. Second, it is possible that the image was linked *only* to the logogen corresponding to “my pet Rover,” which, in turn, invoked a search strategy within the verbal system resulting in something like “Rover is a collie, which is also a dog, which is also an animal.”

Associative processing refers to the activation of informational units *within* either of the systems. However, processing in the verbal system is believed to be sequential or linear, whereas processing in the visual system is thought to be parallel or synchronous. Both imply hierarchical organizations, but access from one logogen to another versus one imagen to another is qualitatively different. If you form a mental image of the refrigerator in your kitchen for example, you can then decide to “look” left or right, up or down in your mind.

Mental scanning can be accessed easily or quickly, regardless of which direction you choose to take. However, recalling the middle line from the "Pledge of Allegiance" requires a linear or sequential search from beginning to end implying a very different storage mechanism for verbal information.

Dual coding theory predicts that pictures and words provided to students will activate each of these coding systems differently. The superiority of pictures for memory tasks is explained by dual coding on the basis of two important assumptions (Kobayashi, 1986). The first is that the two codes produce additive effects. This means that if some piece of information is coded both visually *and* verbally, the probability of retrieval is doubled. The second assumption is that the ways in which pictures and words activate the two codes are different. Pictures are believed to far more likely to be stored both visually and verbally. Words, on the other hand, are less likely to be stored visually. For example, if a picture of a bus is shown to someone, dual coding theory says the picture provides adequate cueing to the visual memory trace and the individual is *very* likely to also add semantic labels. Thus, the picture is being stored in long-term memory two times, once visually and once verbally. On the other hand, if just the word "bus" is shown to someone, the verbal code may be activated, but the visual code may not be activated unless the person actively forms and processes an internal image of the bus. Information that is dually coded is twice as likely to be retrieved when needed because if one memory trace is lost (either verbal or visual) the other is still available. When it comes to memory, two codes are better than one.

Instructional designers should be most interested in ways to increase the likelihood that information will be dual encoded in long-term memory. Information encoded in both verbal and visual forms with strong and flexible links between the codes should enhance retention, retrieval, and transfer. Dual coding is more likely to occur when the content lends itself to imaging (Paivio & Csapo, 1973). Concrete concepts, like "tree" or "house," are good examples of information that readily produce internal images in most people. Concrete concepts are easier for people to visualize simply because they refer to tangible objects that have a physical form. Conversely, people do not automatically form internal images for abstract concepts, like "patriotism" or "kindness." In these cases, it is often useful to provide the learner with a prototype image that communicates the most important characteristics or attributes of the concept (Klausmeier, 1990), such as two people shaking hands to represent "friendship." This prototypical image is frequently analogical to the concept (Newby & Stepich, 1987), like the example discussed in chapter 2 of a blindfolded woman holding a set of scales to represent justice. Research shows that words, sentences, and paragraphs that are easy for people to form internal mental images of are generally recalled better than those that are difficult to imagine. Dual coding theory provides a plausible explanation for this empirical evidence. It is also useful to note that it is generally believed that the primary code for concrete concepts is visual, that the primary code for abstract concepts is verbal, and that concrete concepts are learned before abstract concepts.

Arguments Against Dual Coding Theory

In contrast to dual coding theory, propositional theories suggest that *all* information can adequately be stored in long-term memory in semantic or verbal form, similar to the idea of a propositional network described earlier. Therefore, the assertion that a second code is not needed discounts dual coding on the principle of **parsimony** — the idea that all things being equal, a simpler model should be preferred to one that is more complex.

Propositional theories suggest a process where visual information is transformed into a semantic form. Incoming visually based information from the environment is converted into propositions as the information is passed from short-term to long-term memory. When retrieved, the propositions are transformed back into visual information, as shown in Figure 4.11. This is analogous to how a computer stores information about a graphic in memory on a disk. Information about a graphic appearing on the computer screen (i.e., STM) is converted to digital form for storage on a disk (i.e., LTM). This digital information must later be retrieved from disk, processed, and redrawn on the computer screen for the information to reappear as a picture.

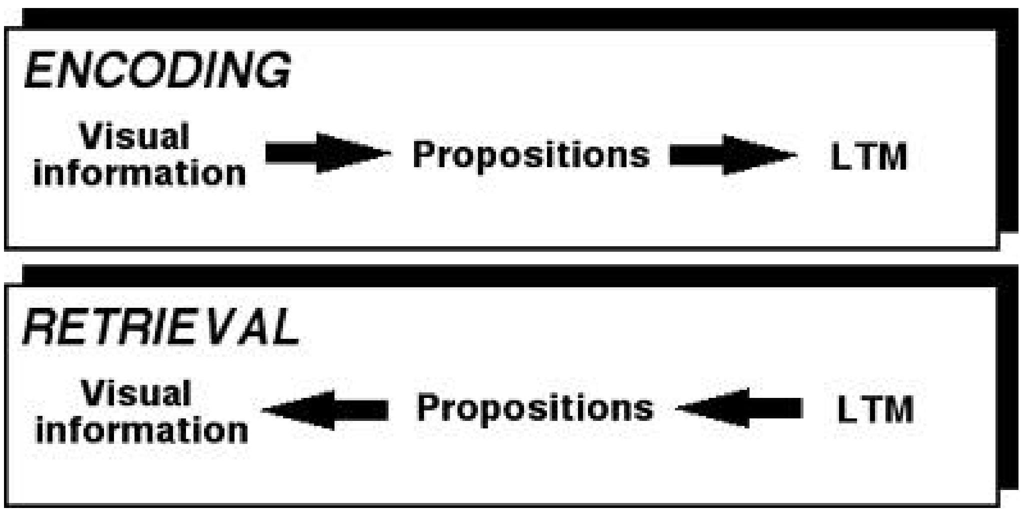

FIGURE 4.11
Proponents of a propositional or semantic model of encoding of visual information in long-term memory suggest that visuals are converted to propositional form as they are encoded, or "passed," from short-term to long-term memory and reconstructed again into visual form when later retrieved back into short-term memory.

Introspection data of people reporting seeing “pictures in their heads” suggests that they are processing the information in short-term memory. Most propositional theorists do not argue against imagery in short-term memory; their arguments are only related to the way information is stored in long-term memory. Propositionists explain empirical evidence of the superiority of pictures over words on the basis of increased elaboration. That is, people provided with pictures just naturally spend more time and effort processing pictures. People process and rehearse pictures more fully than words and sentences. This rehearsal results in more propositional information, as well as more durable traces between the propositions stored in long-term memory, when visual representations are provided than when information is given only in verbal form.

Proponents of a pure propositional theory contend that simple everyday examples show the inadequacy of a dual coding approach. For example, people are usually unable to remember simple facts about objects that they come in contact with every day. Does Lincoln face left or right on a penny? If stored visually, it should be a matter of simply recalling the image from long-term memory and looking for this one detail. Propositionists would say that the image of a penny is not actually stored in memory, only bits of information that, when reconstructed, form the image. Since the proposition related to the “direction of Lincoln's profile” is not a salient feature for most people, it is usually a toss-up between left or right. As another example, visualize the Washington Monument. Now, how many windows are there at the top on any one side? If stored visually, it should just be a matter of counting. Of the many bits of information stored for “penny” and “Washington Monument,” there is no reason why either of these should be rehearsed; therefore, they are not stored. Upon reconstruction, the informational “holes” become obvious only when pinpointed.

Proponents of dual coding meet this challenge by suggesting that a proposition-only theory soon collapses under its own weight. The amount of information contained in even simple images, such as the square shown in Figure 4.12, and the subsequent mental processing necessary to adequately relate the information make pure propositional models impractical. Contrary to a pure propositional model, proponents of the dual coding model suggest that long-term memory is predisposed to verbal *and* visual information. Dual coding advocates suggest that proposition-only models soon collapse under their own weight. For example, the amount of pure propositional information contained in even the simplest visuals, such as this square, is staggering. What, then, would be all of the individual propositions necessary to define even one face? Although the debate over which model more adequately represents actual human cognition may never be resolved, dual coding appears to provide instructional designers with a useful theoretical framework for designing and developing instructional visuals.

Memory for Animated Visuals

Animation, like any picture, should aid recall when it illustrates highly visual facts, concepts, or principles. However, the difference between animated graphics and static graphics for memory tasks is not as clear. Animated graphics are probably better at communicating ideas involving changes over time because of their ability to show motion. Animation should help reduce the level of abstraction for many temporal concepts and principles. For example, the motion of an animated car traveling from New York to Washington, accompanied by a display of the miles traveled and gallons of gasoline consumed, should help reduce the abstraction level for how to compute the average “miles per gallon.” Learners would have to consciously work to connect the visual “snapshots” represented by static visuals for such problems. A common strategy is the use of abstract symbols, such as arrows and dotted lines, in the hope that these represent or suggest the motion attribute to learners. In contrast, animation triggers the automatic ability of the visual system to induce apparent motion, thus freeing short-term memory for other tasks.

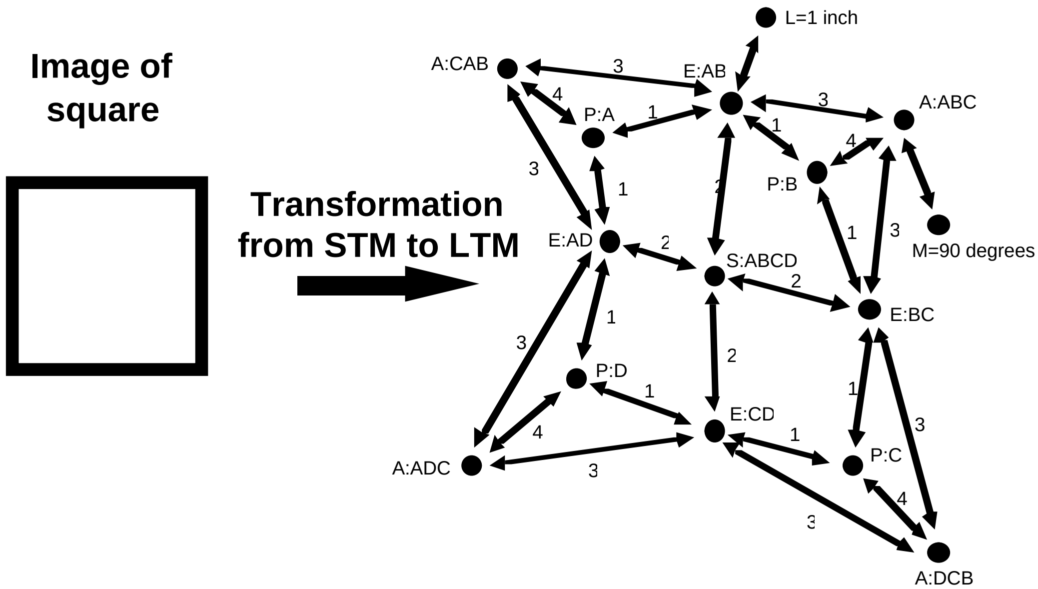

FIGURE 4.12
Contrary to a pure propositional model, proponents of the dual-coding model suggests that long-term memory is predisposed to verbal *and* visual information. Dual-coding advocates suggest that proposition-only models soon collapse under their own weight. For example, the amount of pure propositional information contained in even the simplest visuals, such as this square, is staggering. What, then, would be all of the individual propositions necessary to define even one face? This propositional network for a square was adapted from Larkin, McDermott, Simon, & Simon (1980), where "nodes represent corners (P), edges (E), angles (A), and the surfaces (S). Links connect corners with edges (1), edges with the surface (2), angles with edges (3), and angles with corners (4). Descriptors can be linked to nodes, as shown for the length (L) of Edge AB, and the magnitude (M) of Angle ABC" (p. 1337).

Many important concepts and principles not only change over time, but also change in a certain direction. The direction in which an object is moving is defined as its **trajectory** (Klein, 1987). For example, many concepts and principles in physical science demand that learners understand not only that an object is moving at a certain speed, but also that it is moving in a certain direction. An example is velocity, which is defined as the speed and direction of a moving object. The motion and trajectory of an object can be represented by both verbal and visual means. Film, video, or computer animation can provide visual representations and verbal representations can be conveyed by words such as up, down, fast, or slow. Dual coding theory suggests that the learning of facts, concepts, or principles involving motion and/or trajectory should be facilitated by instruction that presents appropriate *combinations* of visual and verbal representations of these attributes due to increased likelihood of redundant encoding.

Static visuals would be sufficient for tasks that only require learners to visualize information. However, if a task demands that learners understand changes over time or in a certain direction, then static visuals can only hope to prompt learners to mentally construct these attributes on their own. (See Footnote 1) However, animation makes this cognitive task more concrete and spontaneous by providing the motion and trajectory attributes directly to the learner. This would reduce the processing demands on short-term memory and should increase the potential for successful and accurate encoding into long-term memory. Just as dual coding theory would predict, preliminary research has shown that animation displayed with accompanying narration produces greater retention and recall than when either are presented separately or when verbal descriptions are presented before or after the animation (Mayer & Anderson, 1991).

MOTIVATION

Up to this point, the theoretical nature of this chapter may be setting a rather dull tone to human learning. In reality, human learning is an amazingly complex and dynamic interchange of events that current theory sorely falls short of capturing. The cognitive orientations discussed so far hopefully speak to the role of the individual in determining whether learning will occur. The most well-articulated, well-organized, and well-managed instruction will not have a chance to be effective unless it takes into account all the social and motivational factors within which instruction takes place (Weiner, 1990). What motivates an individual to initiate and complete a task? Interpretations have been offered from many points of view. We will again look at this issue from behavioral and cognitive perspectives.

Given the behaviorist's general lack of interest in nonobservable aspects of learning, it may surprise many that motivation plays an important role in traditional behavioral models. Although motivation can be conceptualized in a number of ways, perhaps the most common view of motivation for behaviorists is related to the strength of the reinforcement stimuli — the stronger the reinforcement, the stronger the motivation to respond. As motivation increases, heightened levels of arousal become evident. Learners are thought to seek positive reinforcement through producing desired responses during instruction. If the presumed positive reinforcers are viewed as marginal, neutral, or negative by learners, then the motivation to seek the reinforcement by producing a response decreases. Traditional school and training situations abound with the use of **extrinsic** motivators, such as stars, report cards, or paychecks. As previously mentioned, graphics are frequently used as extrinsic motivators, such as when a graphic, like Superman, appears as a reward to a correct response. However, care must be taken not to “turn play into work,” as research suggests that the well-intentioned use of extrinsic motivators, such as grades, can destroy the natural appeal of an activity for some children (Condry, 1977; Greene & Lepper, 1974; Lepper, Greene, & Nisbett, 1973).

Although cognitive psychologists do not discount the role and reality of extrinsic motivation, most look at motivation from a different perspective. For example, we all experience times when we choose to complete an activity not because of the promise of some external reward, but because the activity itself is satisfying and enjoyable. **Intrinsic** motivation refers to times when a certain activity is its own reward. When activities are intrinsically motivating, people demonstrate continuing motivation by choosing to participate even after external pressures to do so are removed (Deci, 1975, 1985; Kinzie & Sullivan, 1989; Maehr, 1976).

As the motivation literature suggests, the design of activities in which learners demonstrate commitment and perseverance in the thoughtful completion of a task depends on the degree to which the activity is perceived as relevant and its completion as personally satisfying (Keller & Suzuki, 1988; Lepper, 1985). By definition, a meaningful learning context is an intensely personal affair. The goal in education, however, is to discover contexts that have a wide appeal to learners of varying interests and aptitudes. LOGO (see chapters 3 and 8), for example, seems to attract the attention of children through the use of interactive computer graphics to produce interesting visual designs.

Ordinary people differ in their explanations of why they succeed or fail at a task. **Attribution theory** suggests that people interpret their ability to succeed as caused or controlled by several attributes. For example, some people see themselves as in control of their success, whereas others do not and instead believe that external forces control their destiny (Rotter, 1954; Weiner, 1979). One's perception of control over one's success can lead to different patterns of time, effort, and attention to a given task. Another potent attribute is one's perception of how stable the cause of success is over time, whether it is temporary or permanent. Stability is evidenced by questions such as “I was successful today, but will I be able to do it tomorrow?” Obviously, instruction should lead people to believe that they control the frequency and stability of their own success (Milheim & Martin, 1991).

Probably the most applied instructional motivation model is the ARCS model by John Keller (1983, 1988), named after its four components: Attention, Relevancy, Confidence, and Satisfaction. As previously discussed, getting and sustaining a student's attention is a prerequisite cognitive task for any and all subsequent learning. Tasks must be individually relevant to one's needs and expectations. For some people, relevancy may be future-oriented, such as perceiving that a skill will help them in some way in the future; for example, getting a good job or a promotion. For others, relevancy is more present-oriented, meaning that a task will be considered relevant if it satisfies an immediate need, such as providing social rewards with friends or pertinent information to an important question. People generally need to feel confident that they are likely to succeed at the task at hand. It is not that the task need be or should be easy, but success should be within their grasp. People generally have a need to maintain a positive self-image. Feelings of impending success can trigger such an attitude. Lastly, people will generally seek to maintain participation in activities they perceive as interesting and relevant. In other words, we find such activities satisfying. One of the main characteristics of satisfying experiences is that people generally will continue to participate once external pressure to do so has ceased.

Similar to the ARCS model, Malone (1981) has suggested a framework of intrinsically motivating instruction based on *challenge, curiosity*, and *fantasy*. Tasks need to be designed to be optimally challenging — not too easy or too difficult. But perhaps most important, the tasks should elicit feelings of competence, or self-efficacy, as students solve problems they perceive as relevant and important. This enhances one's self-concept and leads to a feeling of control over one's own success (Weiner, 1979, 1985). Similarly, a person's curiosity is usually piqued when an activity is viewed as novel or moderately complex. Curiosity is also usually increased by activities that offer an element of surprise. This occurs when the expected and actual outcomes of an activity are different or incongruent, a phenomenon that Berlyne (1965) termed “conceptual conflict.” Again, however, both challenge and curiosity produced by a conceptual conflict must be optimally maintained to be effective. A task perceived by students as too easy quickly loses appeal, and a task perceived as too demanding is avoided. Likewise, a conceptual conflict between expected and actual task outcomes can make a learner seek to resolve the conflict, but can quickly lead to frustration if the conflict is too confusing or bewildering. Norman (1978) has termed optimal levels of conceptual conflict as “critical confusion.” Fantasy entails providing students with a meaningful context for learning that easily triggers their imaginations. In addition, young students easily transfer such fantasy contexts to play activities. Chapter 8 applies this model of motivation to the design of simulations, games, and microworlds.

Graphics offer the potential to increase the challenge and curiosity of a task, as well as encouraging students to be creative and use their imaginations. Recall from chapter 2 that motivation was one of the instructional applications of graphics. This can be interpreted and applied as including both extrinsic and intrinsic motivation. However, the other instructional applications of graphics also can be understood in relation to motivation as well. Cosmetic and attention-gaining applications are more closely related to extrinsic motivation, akin to the phenomena of “wanting” to go through the instruction because of all the “pretty pictures.” On the other hand, graphics for presentation and practice are more closely related to intrinsic motivation by helping to create an intrinsically interesting learning environment. Certainly, the use of graphics as visual feedback during practice activities, including simulations, can provide an intensely satisfying and challenging learning environment. For example, graphics can easily trigger the illusion of going on a “safari” to find “c” words or getting students to use their estimation skills to “save a whale.”

Finally, intrinsic motivation is just one of the characteristics of **self-regulated learning**, defined as “individuals assuming personal responsibility and control for their own acquisition of knowledge and skill” (Zimmerman, 1990, p. 3). In addition, self-regulated students are metacognitively and behaviorally active (Zimmerman, 1990). Metacognitive attributes involve the student's attempt at the planning, goal-setting, and organization of learning in tandem with self-monitoring and self-evaluation (Borkowski, Carr, Rellinger, & Pressley, 1990). These attributes subsequently lead students to take appropriate actions associated with their own learning, such as the selection, structuring, and creation of environments that will best suit their learning styles and needs.

The advantages of self-regulated learning are obvious. Students become not only more active in the learning process, but also assume responsibility for it. The implication is that students do not simply participate in a given lesson, but actually help design it. At issue is what instructional designers should consider in helping to nurture the self-regulation process. Certainly, learning environments should be designed with a "self-oriented feedback loop" to provide a rich and continual stream of feedback to students to help them establish and maintain goal-setting and goal-monitoring (Zimmerman, 1989). Schunk (1990) referred to students' deliberate attempts to attend to and evaluate their behavior in relation to their goals as self- observation and self-judgment.

REVIEW

- Both behavioral and cognitive learning theories suggest times when graphics can aid learning and times when they might interfere with learning.
- The behavioral principle of least effort suggests that learners may divert their attention to a given graphic even if it is inappropriate for them to do so.
- Cognitive-based information-processing theories suggest that graphics, when appropriately designed, may be useful in attending to, encoding, and retrieving lesson information.
- Propositional theory and dual coding theory are two of the prevalent theories of memory for visual information
- Dual coding theory offers support for the memory of pictures by suggesting that memory has both verbal and visual memory stores.
- Animation is an illusion and occurs when people are tricked into seeing something that really is not there.
- Theories of motivation suggest that students learn best when the learning environment is relevant to their needs and interests, when students are confident of their abilities to participate in a given task, and when that participation is perceived to be a satisfying experience.
- Graphics offer the potential to increase the extrinsic and intrinsic motivation of a learning environment.
- Graphics can increase the intrinsic motivation of a learning environment by piquing students' curiosity, optimizing the challenge of an activity, and by encouraging students to use their imaginations.

NOTES

1. Often, merely prompting learners to internally image can be sufficient to produce learning effects. Studies involving static visuals have indicated that adults are much more likely to spontaneously form internal images than children, suggesting maturation effects. This issue is discussed in more detail in chapter 5.

CHAPTER 5

Review of Instructional Visual Research: Static Visuals

OVERVIEW

This chapter presents a review of research dealing with the effect of pictures and other graphics on learning. It concentrates on research of *static* graphics, whereas the next chapter deals with *animated* graphics. A simple research evaluation model and instructional visual research taxonomy are first discussed to help guide readers through some of the issues related to interpreting educational research. An overview and summary of major published reviews is presented chronologically in order to provide a historical context for the research literature. In general, research suggests that pictures do not always help learning, and, at times, may actually be counterproductive. The chapter concludes with a series of general instructional design principles to help guide practitioners on the basis of available research evidence.

OBJECTIVES

Comprehension

After reading this chapter, you should be able to:

1. Describe the difference between internal and external validity as it relates to research on instructional visuals.
2. List four of the major classes of variables to be considered when interpreting educational research.
3. Describe how the four major classes of variables interrelate or are interdependent.
4. Identify areas in which substantial research on instructional visuals has been conducted, as well as areas in which there is little systematic research.
5. Describe some of the conditions or circumstances under which pictures should or should not help the learning process.

Application

After reading this chapter, you should be able to:

1. Critically read and interpret instructional visual research.
2. Critique the decisions by published reviewers to include or exclude individual studies in their reviews.
3. Derive a set of instructional design principles resulting from research on instructional visuals.

4. Apply this set of instructional design principles in the design and development of instructional materials.

Given the widespread use of illustrations and other types of graphics in instruction, one would think that there would be definitive research literature to either support or dispel their usefulness. Although the use of pictures as an instructional aid has been a very popular research issue, the literature is far from definitive and, at first glance, can even appear contradictory. For example, researchers studying the effects of pictures in prose learning prior to 1970 concluded that pictures often did not aid children's learning and was even distractive at times (Braun, 1969; Samuels, 1967, 1970). Research conducted since 1970 has been more supportive of instructional visuals, not because students are somehow different now, but because we have a better sense of the conditions under which visuals work. This implies that not only do a set of conditions exist, but that pictures will not, and should not, help learning in every instance.

It is easy to be misled into thinking that relatively simple questions like “Do visuals help people learn?” or “Will color make visuals more effective?” should have equally simple answers. An analogous question would be “Is a hammer a good tool to use?” The answer is sure, sometimes, but it all depends. It's great for hammering and pounding nails, but pretty lousy for cutting hair. Researchers in recent years have done a much better job of unraveling and explaining what the effectiveness of instructional visuals depends on. The purpose of this chapter, and the one that follows, is to present a simplified overview of instructional visual research.

This chapter concentrates on the large pool of research conducted on static visual research, a pool which is largely media independent (i.e., not more or less related to one instructional medium, like the computer, than another). The next chapter focuses on the relatively small group of studies dealing with animation, particularly that associated with computers. Both chapters deal exclusively with experimental research that strives to find some causal relationships between instructional variables, such as graphics, and learning. In contrast is correlational research that only looks for relationships between variables, but cannot tell whether one variable (e.g., a picture) caused a change in the other (e.g., learning). Some studies try to bridge the two by studying the relationship between, for example, learner aptitudes and instructional variables (Cronbach, 1957; Snow, 1977). This is known as aptitude-by-treatment interactions (ATI) and has been used to some degree, in research on instructional visuals (Dwyer, 1978). A third branch of research methodology, best known as naturalistic inquiry, is virtually without representation in research on pictures in instruction. Whereas experimental and correlational research is largely quantitative in nature, methodology based on naturalistic inquiry is predominately qualitative in nature. (See Footnote 1)

There have been many recent reviews of research on static instructional visuals. Rather than just offer another review, this chapter will present a “meta-review,” or review of reviews. Furthermore, this review will be presented chronologically, in the hope of providing an historical context. Before beginning, it is important that we have a way of understanding and interpreting the research. To this end, the next section will present a simple model useful in evaluating and interpreting research. This evaluation model, though useful for understanding any educational research, is particularly helpful in knowing when research is either a friend or a foe to instructional designs that incorporate graphics.

INTERPRETING RESULTS OF INSTRUCTIONAL VISUAL RESEARCH

Readers are cautioned at the onset to be suspicious of any experimental research result in the social sciences, for example, educational research. Unlike much research in the hard sciences, control of all extraneous variables in the social sciences, except that which you are actually studying (called the independent variable), is not only extremely difficult, but actually can be criticized for jeopardizing the generalizability (or usefulness to the field in general) of the results. This dilemma is summarized by the distinction between control and maintenance of internal and external validity.

For example, if one wants to study a question related to the effectiveness of illustrations, then internal validity calls for only that variable, visuals, to be allowed to vary between groups in the study: for example, one group is given visuals and another is not. A high control over the study's internal validity, or control of the conditions that actually influence the experiment, makes the researcher confident that if one group learns more than the other, it must be the visuals that are “causing” the difference. However, this type of learning is usually criticized as being too artificial or “sterile” to be generalized to *real* training situations in which learning occurs in the midst of many mediating factors. Generally, the higher the internal validity, the lower the external validity. Conversely, studies that try out variations of instructional variables in “the real world” under “real” conditions unfortunately risk the criticism that the lack of internal controls makes interpretation of any resulting differences difficult or impossible. The issue of balancing internal and external validity can never be totally resolved. However, this is an issue that must always be in the mind of anyone who intends to interpret research results (see Borg & Gall, 1989, for a discussion of sources of validity in relation to educational research).

There are, however, many more issues that must be taken into consideration when trying to interpret educational research. Learning involves a dynamic interplay of many variables, most of which vary even during the learning event itself. A simple model of the many interdependent instructional variables consists of four elements (Jenkins, 1978): the learner; the learning activities; the learning materials; and the testing environment. Though this model is simple, it illustrates a complex set of relationships among general groups of factors that must always be kept in mind when exploring questions about learning, understanding, and remembering. In other words, understanding research requires that each of the four variables be recognized and interpreted in light of the other three. This model likens educational research to a game of pick-up sticks. One cannot manipulate one of the variables without affecting the other three. Figure 5.1 visually represents the dynamic interplay of these variables into a model useful for the evaluation of educational research.

When reading and interpreting research results, many questions should arise based on the interplay of this model. For example, consider a study that reports significant results in favor of visuals. (See Footnote 2) If children were used as subjects, an important question for the practitioner might be whether the results generalize to an adult population. Similarly, if a study reported significant results related to a test of factual recall, do the results generalize to problem-solving situations? If the study reported results with the use of simple line drawings, do the results generalize to color photographs? If the study used an individualized learning approach where learners had control over the pacing of the lesson frames, do the results generalize to group instruction where the visuals are externally paced (such as a video)? It is easy and tempting to believe that a study's results will completely generalize to your situation, especially when the results seem to support your position. Obviously, no study can ever completely match an applied circumstance, but it is vital that the main issues and procedures of a study dovetail with one's purposes and needs sufficiently before research results can be used to support or refute one's design decisions.

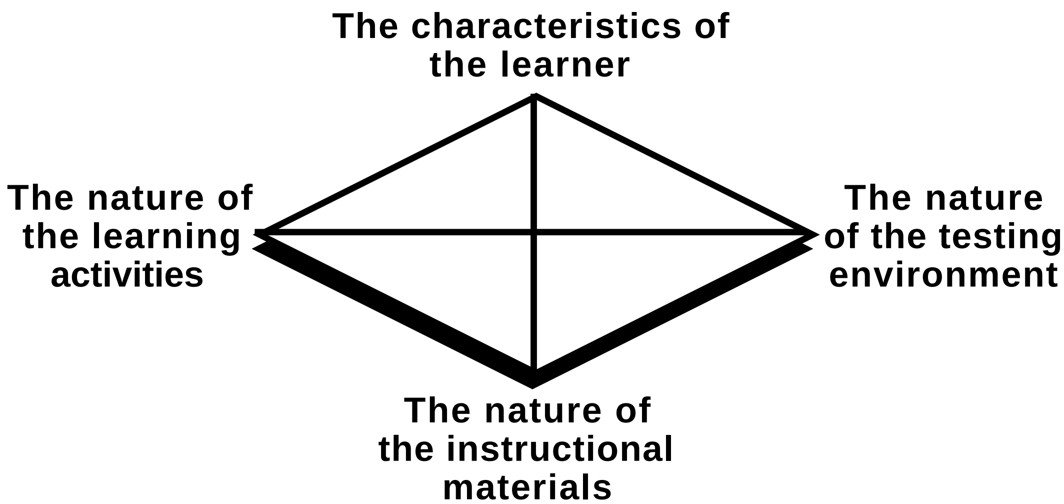

FIGURE 5.1
Some criteria useful in the interpretation of educational research. The four points of this model are interdependent; that is, you cannot change one without affecting the others. Therefore, each point must be intepreted in light of the other three.

A useful analogy for understanding and interpreting research on instructional visuals is the construction of a brick wall. Any one research study represents but one brick. A collection of unrelated studies does not begin to answer a question like “Do visuals aid learning?” any more than a pile of bricks defines a wall. Similarly, the results of any one study must be viewed in relation to the rest of the literature, just as any one brick's contribution to the wall can only be understood when looking at the entire structure. Some studies, like some bricks, become the foundation for others. Other studies, such as replications of previous work, act as functional facades, supporting and confirming what we already know. Obviously, no one study can answer all questions, even if all of the questions are known. At its best, educational research is an organized collection of studies, each of which connects to the ones that precede and follow it.

Guided by the evaluation model in Figure 5.1, we can construct our own “brick wall” of research on instructional visuals in the form of a matrix or taxonomy. For our purposes, the types of graphics (i.e., representational, analogical, arbitrary) relate directly to learning materials. The instructional functions (i.e., cosmetic, motivation, attention-gaining, presentation, practice) and the domain of learning in which each is applied (i.e., verbal information, intellectual skills, cognitive strategies, affective, psychomotor) relate to the learning activities. For example, the crossing of graphic types and learning outcomes results in a matrix with 15 cells or “bricks,” as illustrated in Figure 5.2. By adding the five instructional functions discussed in chapter 2 (three cognitive and two affective) and the most obvious psychomotor function (i.e., demonstration of procedural, or “how to,” tasks), the matrix becomes three-dimensional, as shown in Figure 5.3. Legitimate research can and should be conducted in any one of the 36 cells. The matrix is further complicated by considering the remaining two issues suggested by the evaluation model *within each cell* (or brick) — the nature of the learner and testing components, as well as other issues, such as Gagné's events of instruction. This matrix, in essence, creates a rather thorough taxonomy for general goals related to research of instructional visuals. This research taxonomy helps to clarify where research has been done and to what instructional designs the research applies. The taxonomy also guides work remaining to be done.

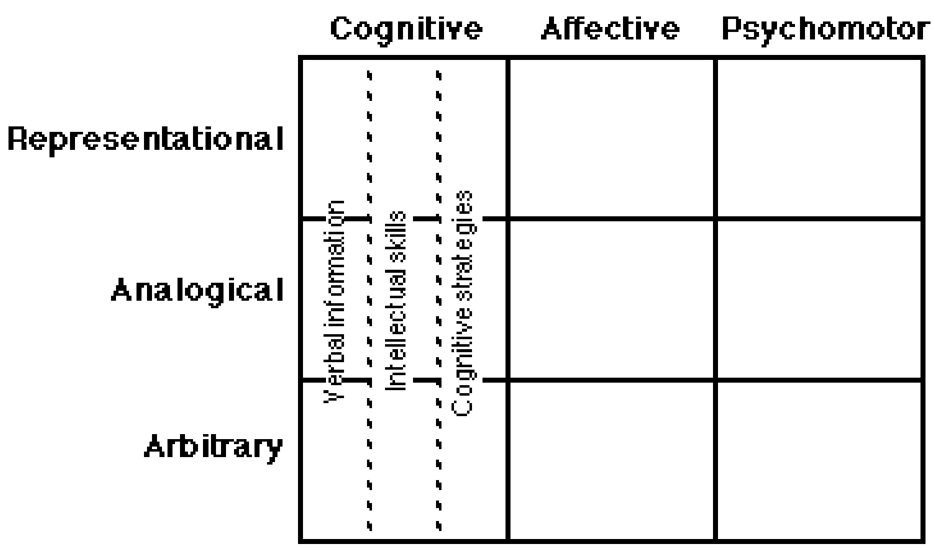

FIGURE 5.2
Simply crossing the various domains of learning with the illustration types describes some of the areas in which research has been conducted.

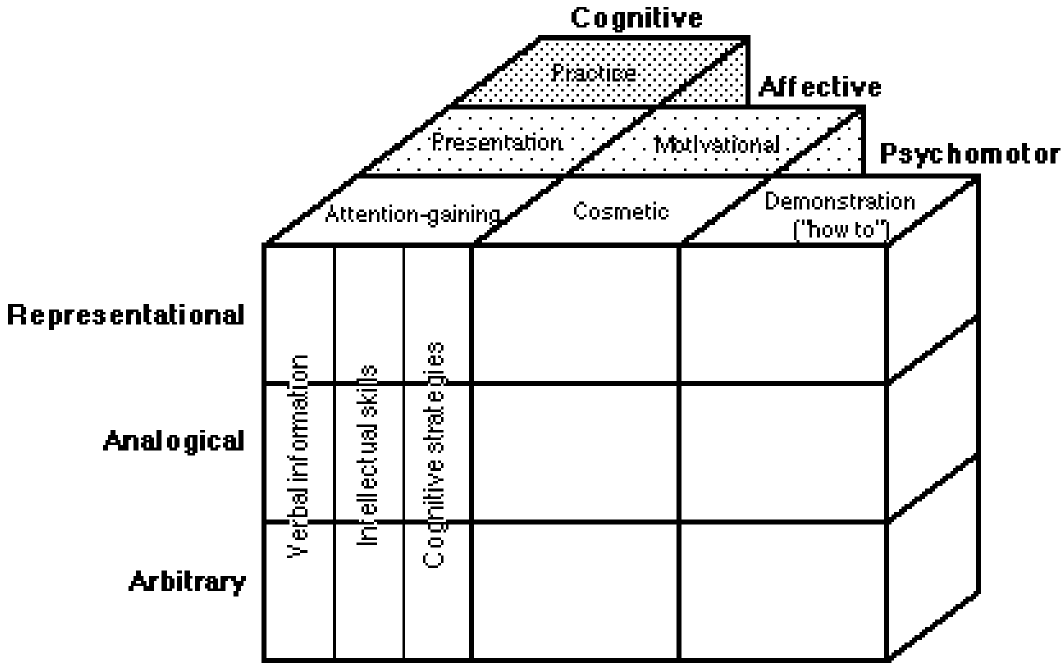

FIGURE 5.3
Adding the consideration of instructional functions to the matrix in Figure 5.2 provides a fairly thorough research taxonomy to help interpret existing research and to show where more research needs to be done.

OVERVIEW OF STATIC VISUAL RESEARCH

The question is, then, what do we know about visuals within areas suggested by the taxonomy of Figure 5.3? Simply put, we know a lot about a little. The lion's share of the research literature deals with the use of static visuals to supplement the presentation of textual information in reading, commonly known as prose learning. The remainder of this chapter summarizes this pool of research. The sequence of instructional visual research and its subsequent interpretation unfolds in an interesting and not so predictable way. For this reason, this overview will chronologically present a series of the most serious and notable reviews available. Considered together, these reviews demonstrate the frustration of early and mostly negative results of using visuals to teach, followed by more positive results. The latest reviews point to a greater understanding of how instructional visuals may contribute to learning when part of global instructional strategies. Again, this chapter focuses only on static visual research. The role of animated visuals is discussed in the next chapter, especially in relationship to computer-based instruction.

Distraction Effects of Pictures: Review by S. Jay Samuels, 1970

One of the earliest serious reviews of instructional visual research was published by Samuels (1970) and began with the following sentence: “If fish were to become scientists, the last thing they might discover would be water” (p. 397). Samuels chastised researchers for failing to investigate “the ubiquitous use of illustrations in books for beginning reading instruction” (p. 397). This admonishment was directed at the failure of researchers to study one of the most widely adopted and touted instructional aids in use — pictures. Samuels reviewed studies on the basis of picture effects on three learning outcomes: learning to read, comprehension, and attitudes.

The first category of experiments — learning to read — studied the use of pictures in teaching simple vocabulary to children. Most of these studies were strictly behavioral in nature and used a pair-association task for the learning activity. This usually involved presenting a series of printed words with or without accompanying pictures of the words. In almost all of the studies Samuels reviewed, there was usually either no difference between the picture and no-picture groups, or students performed better with no accompanying pictures. It was believed that children's attention was naturally drawn to the stimulus that most easily resulted in the correct answer during the learning trial (i.e., the picture), as suggested by the principle of least effort (See Footnote 3) (Underwood, 1963). Therefore, students were usually unable to recall the word when tested only with the written word. Samuels concluded that pictures generally interfere with young children's learning of sight vocabulary. Usually, children classified as poor readers were most susceptible to this interference effect.

The last two categories of studies reviewed by Samuels, comprehension and attitudes, unfortunately posed many problems, making interpretation ambiguous at best. First, too few studies were represented in each case — two in the case of attitudes. Second, the studies that were represented appeared prone to confounding on many counts. For example, many of the comprehension studies tested for memory, not comprehension. The quality of the design of the studies is also easily questioned. The Samuels' review, therefore, largely left unanswered the questions of whether pictures affect reading comprehension and attitudes.

The value of the Samuels' review relates to its evidence of the potential distracting nature of visuals. Samuels concluded that students, usually those with below-average reading skills, had difficulty shifting their attention from a picture to a written word because the picture required less effort. Anyone who has ever tried to assemble a children's bicycle late on Christmas Eve knows of the temptations to skip the written directions and defer to the accompanying pictures. However, it is possible to overlook relevant information in the picture or misunderstand relevant information.

A classic demonstration of the distractive potential of pictures is a pair of studies conducted by Willows (1978). Second- and third-grade children were presented with words superimposed on either related or unrelated pictures. A no-picture group acted as a control. Results showed that the children read more slowly and with less accuracy in the presence of pictures than without pictures. Furthermore, unrelated pictures produced more interference than related pictures. Willows concluded that “the children either consciously or automatically and unconsciously attempted to use the pictures as clues to the meanings of the words printed on or near them” (p. 261). Similar to Samuels, Willows found that younger, less-skilled readers were more susceptible to this interference effect and also attributed the result to the principle of least effort. Willows further suggested that the results from his word-picture association experiments generalized to reading prose or narrative, and discussed the distraction effects in relation to designing books for children:

> In most books used for reading instruction, the illustrations which accompany a story are complex and include representations of many components of the text. If a child comes to a word he already knows, then the pictures in the periphery are superfluous and probably distracting. If he does not know a word and looks to the picture for a clue to its meaning, he may well be misled by those aspects of the picture which are not closely related to the meaning of the particular word he is trying to decode. (p. 261)

Describing the Conditions Under Which Pictures Facilitate Learning

Research conducted since the Samuels' review points to the positive effect of pictures in prose learning. Recent findings suggest that pictures can exert strong positive influences on learning, given certain conditions. Further evidence suggests that children's dependence on pictures decreases with age. As children grow older, they become better able to produce their own internal images (Pressley, 1977).

For example, Guttmann, Levin, and Pressley (1977) read stories to kindergarten, second-grade, and third-grade children and presented either sets of pictures that fully illustrated the content of the stories (imposed pictorial conditions), sets of pictures that illustrated the story but omitted all information contained in the questions (partial-picture conditions), or no pictures. Subjects in the partial-picture condition were explicitly instructed to form mental images. Children in the control condition neither saw pictures nor were instructed to internally construct them. Results showed that even kindergarten children learned more when presented pictures illustrating the story; however, the partial-picture group showed no difference from the control group. Second-grade children were more able to use the incomplete pictures and seemed to be able to construct an internal image if given a partial picture, but incomplete pictures still did not produce significant differences in learning when compared to that of the control group. Third graders in all three experimental conditions, on the other hand, significantly outperformed control subjects. This study suggests the developmental importance of imagery ability. Imagery abilities and skills, like other cognitive processes, probably develop over time. Other studies support this contention (Shimron, 1975; Lesgold, Levin, Shimron, and Guttmann, 1975).

Other research, although supporting the claim that children depend less on outside images as they grow older, demonstrated that pictures can decrease the difficulty of prose material for older children. Levin and Divine-Hawkins (1974) demonstrated that fourth-grade children do not automatically construct images, although they are capable of doing so. This finding has led to many examples of successful training of subjects to form mental images (see, for example, Lesgold, McCormick, & Golinkoff, 1975; Pressley, 1976).

Review by Joel Levin and Alan Lesgold, 1978

Since the late 1970s, many reviews of instructional visual research have been published that point to positive effects of pictures on learning. One of the earliest was a review by Levin and Lesgold (1978), which described "abundant empirical evidence to document the positive value of pictures" (p. 233). They reported that consistent learning gains were to be expected by the addition of pictures *when the following five ground rules were followed*: (1) the prose passages were presented orally; (2) the subjects of the experiments were children; (3) the passages were fictional narratives; (4) the pictures overlapped the story content; and (5) the learning was tested by factual recall.

The purpose of employing these ground rules was to try to sift through the hundreds of studies available for some consistent set of conditions that could be examined, interpreted, and debated. These ground rules also made it possible to exclude studies that, although published, were poorly designed or confounded in some way. Results of studies such as these are not only meaningless, but also distort one's interpretation of the total pool of published research when these studies are included in a review. Based on their review, the ground rules provide a strong supportable set of conditions, however small, under which picture facilitation effects could be expected.

Even if you disagree with Levin and Lesgold's ground rules, at least you know what you are disagreeing with. For example, the first ground rule — that passages should be presented orally — makes generalization to other reading contexts arguable. However, the rationale of Levin and Lesgold to include this ground rule is simply to make certain that the subjects comprehended the context that the illustrations were supposed to support. This rule removed the possibility that any one subject's "inability to perform well on a prose-learning or comprehension test [could] be attributed to the subject's inability to read" (p. 234).

The importance of this issue cannot be understated when the goal is to interpret the effectiveness of illustrations on learning. Almost all of the available studies involve a research design in which the illustrations are supposed to *support* a passage of text. If subjects cannot read the text, then obviously there will be "no significant differences" (NSD) in varying the type of supporting illustration. The temptation is to conclude, incorrectly, that the illustrations had no effect. If subjects did not know how to read, then the results provide absolutely *no* information about the effectiveness of the illustrations. In fact, any time a study reports NSD, ineffectiveness of the instructional treatments is one of perhaps hundreds of rival hypotheses. Unfortunately, in all too many cases, the most probable cause of NSD is a problem with the research design. Levin and Lesgold go further by strongly suggesting that pictures will be of similar benefit when learners, when capable, are *reading* for comprehension.

By limiting their review only to studies involving children, as per the second ground rule, Levin and Lesgold remove the possibility that maturation effects may confound their interpretation. Similarly, the third ground rule, limiting the studies to those involving only fictional passages, greatly reduced the chance that subjects had prior knowledge of the information and therefore performed well on a test because they already knew the answers before the research began. This ground rule helps to “separate out what is *learned* from what is already known” (Levin & Lesgold, 1978, p. 235).

The fourth ground rule, that the pictures must overlap the story content, has remained one of the most fundamental principles of instructional graphic design. As discussed in chapter 2, pictures can serve a wide variety of purposes and functions, including decoration. It is important to distinguish these uses of pictures from those that try to illustrate a concrete feature of a passage. Levin and Lesgold's review provides evidence that pictures offer the potential to facilitate learning when the pictures are, indeed, congruent with the content to be learned.

The fifth and final ground rule is the most limiting for most practitioners. Levin and Lesgold limit their review only to studies in which factual recall was aptly measured. It would be inappropriate to generalize this finding to other types of learning, such as problem solving. Again, researchers frequently have been guilty of not only failing to precisely define the type of learning being studied in an experiment, but also failing to validate whether or not the testing procedures actually measured it. Due to the widely inconsistent definitions and measurements of learning in the available studies, Levin and Lesgold (1978) felt forced to limit their review to “factual information recall tapped by short-answer (generally “Wh”) questions” (p. 236).

One example of the supportive research cited by Levin and Lesgold is a study by Guttmann, Levin, and Pressley (1977). In their study, children correctly responded to about 80% of the short-answer questions when pictures accompanied oral narratives, whereas children in a no-picture condition only answered about 57% correctly.

Taken as a group (you can not accept one condition without accepting the others), the ground rules from the Levin and Lesgold (1978) review provided the first meaningful set of guidelines to help practitioners predict when instructional visuals would be worth using. Further research continues to support these ground rules (see review by Levie & Lentz, 1982, for example).

Research Conducted and Reviewed by Francis Dwyer

One of the most prolific instructional visual researchers to date has been Francis Dwyer. Since 1965, Dwyer and his colleagues have conducted research with more than 8,000 high school students and 40,000 college students (Dwyer, 1972, 1978, 1987; Canelos, 1987). However, the most unique aspect to Dwyer's research is its systematic approach to investigating the use of visuals in instruction. The instructional materials of every study involved a 2,000-word script describing the parts, locations, and functions of the human heart. True to instructional design principles, instructional visuals were only added to illustrate parts of the script found difficult to learn. A total of 37 such critical areas were augmented with visuals (p. 52, Dwyer, 1978).

The types of visuals used in any one study depended on the research question being asked. Generally, the studies investigated the effects of representational pictures varying from highly detailed color photos to simple line drawings as an aid to learning the content contained in the script. Also studied were the effects of various lesson strategies on visualized instructions, such as using moving arrows and inserted questions as cueing strategies. A variety of treatment combinations were studied, such as incorporating the programmed materials into print-based materials, slide/tape materials with audio, and computer-based materials. This research effort, known as the Program of Systematic Evaluation (PSE), has produced more than 150 published research studies (Canelos, 1987).

In a way, Dwyer's research findings act as a testimonial to all of instructional visual research because they repeatedly show that visuals are not equally effective across learning situations. Effectiveness of all instructional strategies, such as visuals, depends on a wide array of factors, such as those already suggested by the research evaluation model in Figure 5.1. No true general principles were found. Some of the visuals were found to be effective some of the time under some of the conditions.

The most consistent results found by Dwyer were related to the amount of realism in the visuals. The results suggest that pictures facilitate learning for adults under certain conditions. For example, people need sufficient time to scan and interpret visuals with highly realistic details. Richly detailed visuals require the learner to attend to and systematically scan the visual in search of essential learning cues. If insufficient time is given, students may actually choose to ignore the visuals and attend to the more familiar, printed text. When lessons are externally paced, the most effective visuals are those that usually contain relatively small amounts of visual detail.

Dwyer's finding of the curvilinear relationship between learning and realism of instructional visuals (too little or too much realism adversely affects learning) is probably the hallmark of his research, and it disputed early realism theories (e.g., Carpenter, 1953; Dale, 1946; Morris, 1946) that suggested that the more information provided, the more learning would occur — hence, more realistic pictures should invoke more learning. Dwyer's research indicates that when learners are confronted with visuals containing too much information and too little time they frequently choose to either ignore visuals or attend to wrong or inappropriate information in the visual. This phenomenon, where an optimal amount of arousal (not too little and not too much) is necessary for acceptable levels of performance, is consistent with many aspects of human nature, such as motivation, and is classically known as the Yerkes-Dodson law (1908).

Part of the value of Dwyer's research is to remind designers to be skeptical of any general, across-the-board recommendations to include visuals in instruction. Dwyer repeatedly reminds designers to weigh the effectiveness of materials against the efficiency and economy of producing them. If the instructional treatments containing visuals are no more effective than text alone, Dwyer (1978) encourages that visuals be omitted

Review by W. Howard Levie, 1987

Levie (1987) provides among the broadest reviews of picture research. (See Footnote 4) Levie reviewed available research in four areas: picture perception, memory for pictures, learning and cognition, and affective responses to pictures. Again, this review shows that much is known about too few of the "bricks" discussed in Figure 5.3. Levie suggests that "an aerial view of the picture research literature would look like a group of small tropical islands with only a few connecting bridges in between" (1987, p. 26).

The first two areas of inquiry, picture perception and memory for pictures, are essentially the same as those discussed in chapter 4. Research on recognition memory for pictures constitutes the largest pool of research on a single topic. This research is that associated with showing subjects hundreds of pictures and testing them: "Have you seen this picture before?" This research has obvious applications in advertising, for example, for issues such as product recognition. Among the important variables associated with a person remembering whether or not an image had been seen previously are meaningfulness and distinctiveness of the image, the ability to distinguish foreground objects from background graphics, complexity, color, movement, and the degree to which all objects in a picture have an obvious relationship to each other.

In considering memory for pictures, Levie, however, stresses the dual role of pictures: they are objects themselves, and, at the same time, they act as substitutes or surrogates for other objects. We need to remind ourselves that the interpretation of pictures is culturally based. For example, western cultures are inundated with information about the world represented in picture or iconic form. We tend to take for granted the degree of visual literacy of people brought up in those cultures versus other cultures. Some of these visual skills are probably innate, such as the ability to recognize and identify an object displayed in a picture. However, other visual literacy skills, such as decoding the meaning of abstract or complex visuals, require refined visual decoding skills.

The nature of what constitutes realism is also a very debatable issue. Again, we tend to take for granted the subjectiveness of picture recognition. For some, realism is somehow measured against the likeness of the object the picture is supposed to represent. For example, most people know that the object represented in a black and white photograph also contains color cues that are not presented. However, a second prominent view, one stressed in chapter 8, is that each of us constructs our own reality. In describing these two points of view, Levie (1987) writes that on one hand " . . . realism is basically a function of the degree of resemblance between the information provided by two types of optic rays" (p. 2), whereas, on the other hand, constructivists "point out that judgments of pictorial realism are influenced heavily by our learned preconceptions of how such a picture *ought* to look" (p. 4). This is just another reminder that perception, as discussed in chapter 4, is not like "taking a photograph" with our eyes.

The third area of research reviewed by Levie, learning and cognition, represents research on tasks more closely associated with school learning or training activities. This is the largest pool of research directly applicable to instructional design issues. All of the research reviewed and conducted by Samuels (1970), Levin and Lesgold (1978), and Dwyer (1972, 1978, 1987) fit here. Most of this research concerns the learning of verbal information, or facts. There is also ample research on the use of visual mnemonics as an aid to remembering information. For example, studies have shown that visual mnemonic techniques, such as the pegword system (see chapter 2), are generally ineffective in young children who are not able to internally image (Pressley & Levin, 1978). However, instructional techniques that provide a visual representation of the mnemonic have proven to be effective. Examples have included picture mnemonics that help children learn the names of minerals, state capitals, and U.S. presidents.

Little research has been done on using pictures in concept acquisition, and the research that has been conducted has not been systematic. As previously discussed, analogical pictures offer potential to help people understand abstract concepts (Tennyson & Park, 1980). Analogies that are visually based seem to facilitate learning by helping to make abstract information more concrete and salient (Mayer, 1983; Rigney & Lutz, 1975).

Levie notes that very little research is available on the role of pictures in higher-order thinking, such as problem-solving. However, there is some evidence to suggest that visual thinking plays an important part in skills associated with syllogistic reasoning (e.g., “If Tim is older than Mary and Mary is older than Brian, who is older, Tim or Brian?”). This research is not well-substantiated or corroborated; however, it is intriguing and seems to match many people's accounts of experiences of gaining insight while problem-solving. There are many stories of famous philosophers and scientists arriving at ingenious solutions to a problem by having the solution come to them in a vision. Einstein is said to have thought of the theory of relativity by imagining what it would be like to ride on a beam of light (Burke, 1985).

In contrast to research in the cognitive domain, little research is available on the affective effects of pictures. While there is little research to suggest that pictures necessarily promote stronger emotional feelings than words in the long-term, there is convincing evidence that pictures offer much more emotional impact in the short-term. One might refer to this as “shock value,” and research suggests that pictures can arouse strong feelings in less than a second (Cupchik & Berlyne, 1979, as cited in Levie, 1987). Verbal forms of similar information, such as text, require comprehension and interpretation that can temper the emotional impact more than when presented visually. It has been suggested, for example, that television images of people being brutalized during the civil rights period of the 1960s helped sway public opinion and shape public action. There is certainly a difference between calmly reading an account of protesters being attacked by police dogs and actually seeing the animal's ferocity directed toward a 9- or 10-year-old girl or boy. Emotional reactions to films also have been documented (e.g., Lagerspetz & Englom, 1979 as cited in Levie, 1987).

Similarly, attitudes can be affected by pictorial information. Television news stories have been viewed as more accurate and believable when film shot on location is shown (Ryan, 1975). Also, research shows that people show preferences for types of visual material. For example, research shows that children often prefer realistic pictures. Other preferences include color, complexity, and ambiguity. However, one's aesthetic preferences are not necessarily related to times when the goal is *learn* something from the visual material. Dwyer (1972) has noted that “aesthetically pleasing visuals may be deceptive in their instructional value” (p. 90). Again, this serves as an important reminder from previous discussions that graphics can serve different functions, such as affective or cognitive. It is important that the intent of the graphics be defined. There is nothing wrong with adding visuals simply for aesthetic appeal, so long as these graphics do not undermine other instructional goals, such as those in the cognitive domain.

Review by Joel Levin, Gary Anglin, and Russell Carney, 1987

It should be evident by now that by far the most extensive body of literature related to the influence of pictures on learning has been that associated with prose learning. Prose learning essentially is learning from text. In contrast to the research cited earlier (e.g., Samuels, 1970), which studied the role of pictures in learning how to read, the emphasis in prose learning is on the ability to comprehend and remember information contained in the text, or "reading to learn" (Levin, Anglin, & Carney, 1987, p. 51).

Levin, Anglin, and Carney (1987) conducted a meta-analysis of pictures in prose studies. A meta-analysis is a statistical analysis that takes and groups the effect sizes of the independent variable in question (i.e., pictures) across studies. Using this pool of data, inferences about the independent variable can be made from the total group of studies. While open to criticism, meta-analyses offer the best procedure to date for objectively analyzing across studies. The result of their meta-analysis has yielded what they call the "10 commandments of picture facilitation." These commandments are presented here with "secular" translations to act as a concise set of guiding principles resulting from picture effects in prose learning (see Table 5.1). Together, these commandments and principles also act as a suitable summary and action plan for the research presented in this chapter.

> *Commandment 1*: "Pictures shalt be judiciously applied to text, to remember it wholly" (p. 73). *Instructional design principle*: Since the available research is largely based on the effects of pictures on learning from prose, this research and any principles derived from it can only be consistently applied to similar situations.

This first commandment acts as an organizer for the rest. It serves as a reminder of where and how the commandments were derived. All of the remaining commandments and principles point to times when pictures offer the potential to facilitate learning *from prose*, such as text or narrative passages. This commandment also reminds us to not apply and generalize this group of research to a set of questions or problems that are inherently different. In other words, these principles are only *directly* useful if your goal is to design pictures and graphics to assist in remembering textual information presented in narrative fashion. If you are designing pictures or graphics for other tasks or functions, then you must look elsewhere for guiding principles. Unfortunately, most instructional designers deal with questions in which there is no substantive research literature. In those cases, a decision must be made whether or not research on the effects of pictures on learning from prose materials can be generalized to fit those circumstances. A pragmatic view is to cautiously accept this research pool until better and more direct information becomes available.

Table 5.1. Some instructional design principles derived from research on using pictures a an aid to learning from narrative text.
1. Since the available research is largely based on the effects of pictures on learning from prose, this research and any principles derived from it can only be consistently applied to similar situations. 2. Pictures should be congruent and relevant to the information presented in the text. 3. Pictures should be congruent or relevant to the text, otherwise they may be distractive and interfere with learning. 4. Pictures will not be necessary when the text already prompts the learner to spontaneously form internal images. 5. When pictures are prepared as an aid to learning from text, be sure that the learner can read and understand the text. 6. Pictures should be prepared to clearly represent the content which is to be remembered, and additional "dressing up" of the picture should be avoided. 7. The preceding commandments and related principles listed refer only to representational or analogical graphics, and not to arbitrary graphics, such as graphs and charts. 8. Pictures should be designed to perform their appropriate instructional function based on the needs of the learner, the instructional objectives of the task, and the instructional materials actually used.

> *Commandment 2*: "Pictures shalt honor the text" (p. 73). *Instructional design principle*: Pictures should be congruent and relevant to the information presented in the text.

This is among the most substantiated principles for the design of instructional graphics and is the same one promoted in the review by Levin and Lesgold (1978). Pictures must be congruent and relevant to the information or message contained in the prose material in order for the pictures to make a difference in learning. The most potent learning effects have been from the use of pictures that act as visual mnemonics. Further research shows that these picture effects are durable over time (Anglin, 1986, 1987).

> *Commandment 3*: "Pictures shalt not bear false fitness to the text" (p. 73). *Instructional design principle*: If pictures are not congruent or relevant to the text, they may be distractive and interfere with learning.

In a real sense, this commandment is the complement to Commandment 2. As shown by research by Samuels (1967) and Willows (1978), pictures offer great potential for distraction. Obviously, when pictures are present and do *not* offer any instructional benefit, the risk that the learner will be distracted from the intended message is great enough to reconsider if the picture or graphic is necessary. This principle should be seriously considered by instructional designers and developers who seek to add graphics to increase the affective or motivational appeal of their materials. The use of cosmetic or decorative graphics offers the potential for two types of disruptive situations. The first is simply that appealing, cosmetic graphics may distract a learner's attention from the message. Given the fact that short-term memory is limited and temporary, any stimuli that offer the potential to occupy any cognitive processing should be carefully weighed. Second, a learner may actually be misled into believing that a graphic meant as decoration by the designer *is* somehow relevant to the message or instructional activity. This latter situation is particularly dangerous, as it opens the door for potential misconceptions.

> *Commandment 4*: “Pictures shalt not be used in the presence of `heavenly' bodies of prose” (p. 74).

> *Commandment 5*: “Pictures shalt not be used with text cravin' for images” (p. 74).
> *Instructional design principle*: Pictures will not be necessary when the text already prompts the learner to spontaneously form internal images.

These two commandments are presented together because they refer to situations in which pictures would not be expected to result in any greater learning than would be expected from the text alone. Commandment 4 refers to times when learners are already learning as much as possible without using additional aids, such as pictures. This is known in research circles as “ceiling effects.” For example, if a learner already demonstrates mastery of a concept given one instructional method or strategy, then any additional strategy or aid will not result in more learning because there is no room for improvement. In regard to picture research, reaching a ceiling level of learning from text alone means either that the visualization of the content is not necessary to permit learning or that the text itself does a good enough job of cueing the learner to internally form an adequate image in short-term-memory.

Commandment 5 elaborates on the phenomenon, experienced by almost all people, in which the text or verbal message is so concrete and vivid that we are able to conjure up a suitable image in our minds ourselves. Probably the best examples of this phenomenon are from literature in which skillful writers “paint pictures” with their words. An example of this is contained in Box 5.1. Obviously, the ability of some words, whether written or spoken, to evoke mental images depends on a reader's or listener's background and prior knowledge. Of course, the ability to tell or write a story so that it paints pictures in a wide range of people's minds is a distinguishing characteristic of *master* writers and storytellers.

Box 5.1

Seeing A Story With Words Alone

To think that you must actually draw a picture for people before they see something is often unnecessary (and perhaps presumptuous). Some of the best pictures are “drawn” with words. We all have our favorite examples, whether they be from writers or storytellers. One of my favorites is from John Steinbeck’s (1939) *The Grapes of Wrath* which takes place during the “dust bowl” era of the 1930’s. (I lived for a time along Route 66 in New Mexico, a road which figures prominently in the book.) Read the following passage and reflect on the images you "see" in your mind:

Outside, the seated man stood up and looked over the cowl of the truck and watched the restaurant for a moment. Then he settled back on the running board, pulled a sack of tobacco and a book of papers from his side pocket. He rolled his cigarettes slowly and perfectly, studied it, smoothed it. At last he lighted it and pushed the burning match into the dust at his feet. The sun cut into the shade of the truck as noon approached.

In the restaurant the truck driver paid his bill and put his two nickels’ change in the slot machine. The whirling cylinders gave him no score. “They fix ‘em so you can’t win nothing, “ he said to the waitress.

And she replied, “Guy took the jackpot not two hours ago. Three-eighty he got. How soon you gonna be back by?”

He held the screen door a little open. “Week-ten days,” he said. “Got to make a run to Tulsa, an’ I never get back soon as I think.”

She said crossly, “Don’t let the flies in. Either go out or come in.”

“So long,” he said, and pushed his way out. The screen door banged behind him. He stood in the sun, peeling the wrapper from a piece of gum. He was a heavy man, broad in the shoulders, thick in the stomach. He face was red and his blue eyes long and slitted from having squinted always at sharp light. He wore army trousers and high laced boots. Holding the stick of gum in front of his lips he called through the screen, “Well, don’t do nothing you don’t want me to hear about.” The waitress was turned toward a mirror on the back wall. She grunted a reply. The truck driver gnawed down the stick of gum slowly, opening his jaws and lips wide with each bite. He shaped the gum in his mouth, rolled it under his tongue while he walked to the big red truck.

The hitch-hiker stood up and looked across through the windows. “Could ya give me a lift, mister?”

The driver looked quickly back at the restaurant for a second. “Didn’t you see the *No Riders* sticker on the win’ shield?”

“Sure — I seen it. But sometimes a guy’ll be a good guy even if some rich bastard makes him carry a sticker.”

The driver, getting slowly into the truck, considered the parts of this answer. If he refused now, not only was he not a good guy, but he was forced to carry a sticker, and was not allowed to have company. If he took in the hitch-hiker he was automatically a good guy and also he was not one whom any rich bastard could kick around. He knew he was being trapped, but he couldn’t see a way out. And he wanted to be a good guy. He glanced again at the restaurant. “Scrunch down on the running board till we get around the bend,” he said.

The hitch-hiker flopped down out of sight and clung to the door handle. The motor roared up for a moment, the gears clicked in, and the great truck moved away, first gear, second gear, third gear, and then a high whining pick-up and fourth gear. Under the clinging man the highway blurred dizzily by. It was a mile to the first turn in the road, then the truck slowed down. The hitch-hiker stood up, eased the door open, and slipped into the seat. The driver looked over at him, slitting his eyes, and he chewed as though thoughts and impressions were being sorted and arranged by his jaws before they were finally filed away in his brain. His eyes began at the new cap, moved down the news clothes to the new shoes. The hitch-hiker squirmed his back against the seat in comfort, took off his cap, and swabbed his sweating forehead and chin with it. “Thanks, buddy,” he said. “My dogs was pooped out.”

Of course, if you've seen the film version starring Henry Fonda, your mental images may be clouded with this scene from the movie. In a sense, watching a movie before reading the book robs you of the experience of creating your own visual interpretation of the author's words. Instead, you are really watching another person's interpretation of the work, such as the film's director. One might argue, for example, that the symbolism of Melville's Captain Ahab in Moby Dick is too quickly replaced with the face of Gregory Peck, and, as a result, much of the symbolism is removed as well. The role and use of film versions of classic literature is an interesting issue to debate.

> *Commandment 6*: “Pictures shalt not be prepared in vain” (p. 74).
> *Commandment 7*: “Pictures shalt be faithfully created from generation to generation.” *Instructional design principle*: When pictures are prepared as an aid to learning from text, be sure that the learner can read and understand the text.

The ability to read is an obvious prerequisite to experiments that investigate any aids to text comprehension. Unfortunately, the literature is filled with cases when this seemingly simple rule has been broken. Commandment 6 refers to the fact that pictures produced as an aid to learning from text cannot be effective unless learners have the prerequisite reading ability, such as word decoding and recognition, literal and inferential comprehension, and comprehension monitoring (E. Gagné, 1985). In these cases, the fact that pictures do not help “should be neither disturbing nor surprising” (Levin, Anglin, & Carney, 1987, p. 74).

Commandment 7 refers to the same issue, but is couched in the context of students being able to form internal images from text. Obviously, if students are not able to read, they consequently will not spontaneously form internal images of verbal information contained in the text. Other “sins” associated with commandment 7 include assuming that all learners are equally able to form internally images and that learners will be able to correctly decide which parts of the text should be used for internal imaging purposes. In other words, learners often have a difficult time distinguishing between important and unimportant text. As already mentioned, research has shown that not all people, such as young children, are appropriately able to internally image.

> *Commandment 8*: “Pictures shalt not be adulterated” (p. 75). *Instructional design principle*: Pictures should be prepared to clearly represent the content that is to be remembered; additional “dressing up” of the picture should be avoided.

This commandment refers to specific graphic design decisions of the picture itself. For example, consider an instructional design that calls for a picture of a rose to supplement a text on gardening. Should a single rose be included, a bouquet, or the whole garden? Should line art or a photograph be included? Should color be included in the picture? The advice here is to avoid adding visual characteristics to graphics that do not directly support the clarity of the information to be learned. This commandment is supported largely by Dwyer's (1978, 1987) research on the realism of instructional visuals discussed earlier. Although countering instructional wisdom at times, this research suggests that too much information can often lead to poorer learning performance because learners often cannot extract the most relevant and salient information from a complex graphic.

> *Commandment 9*: “Pictures shalt be appreciated for the art they art” (p. 75). *Instructional design principle*: The preceding commandments and related principles listed refer only to representational or analogical graphics, and not to arbitrary graphics, such as graphs and charts.

It would be a serious mistake to generalize research results taken from one class of pictures to pictures that serve totally different purposes and functions. It is important to remember that the previous commandments and related principles were generated from research where pictures were meant to bear some physical resemblance to objects, concepts, or processes represented directly or analogically to that included in a text.

> *Commandment 10*: “Pictures shalt be made to perform their appropriate function” (p. 76). *Instructional design principle*: Pictures should be designed to perform their appropriate instructional functions based on the needs of the learner, the instructional objectives of the task, and the instructional materials actually used.

This commandment acts as a final summarization of the research on pictures as an aid to learning from prose. The amount of additional learning that can be expected by adding a picture is largely a function of the type of picture selected. The resulting instructional design principle summarizes this entire chapter, and in so doing, goes well beyond the commandment offered by Levin, Anglin, and Carney (1987). This principle suggests that the function of any particular graphic must be consistent with its intent. This principle also suggests that care must be taken to assure that the intent of a graphic does not subvert or undermine the function or effectiveness of other instructional elements, which may or may not be related to the graphic. For example, this principle accepts the premise that graphics can be added for affective considerations, such as those that serve cosmetic or motivational functions, but only under the condition that such graphics do not distract a learner's attention from any relevant information or tasks.

A FINAL WORD

The apparent contradictions concerning the effectiveness of pictures in reading require cautious interpretation. It is clear that there are contexts where pictures do not facilitate learning due to distraction effects and the inability of some readers to shift attention from pictures to text. However, ample contexts (the ground rules used by Levin and Lesgold

[1978] for example) exist where pictures appear very useful in facilitating reading achievement. Dominant conclusions drawn from this research are: (1) pictures are superior to words for memory tasks; (2) adding pictures (external or internal) to prose learning facilitates learning, assuming that the pictures are congruent to the learning task; (3) children up to about the age of 9 or 10 rely more heavily on externally provided pictures than do older children; (4) children do not automatically or spontaneously form mental images when reading.

This chapter has tried to briefly present the research to date on the effectiveness of static visuals by "reviewing the reviews." This, however, is not a substitute for reading the reviews themselves and the primary sources of the research — that is, the original published studies discussed in the reviews. Anytime one reads an author's interpretation of another's work, the resulting message becomes necessarily biased by the experiences and background of the reviewer. Therefore, it is important to remember that this chapter is but a beginning step in understanding the research on instructional uses of static visuals. (See Footnote 5)

Some of the next steps include reading these and other reviews, and as many of the original research reports cited as possible (and even those that were not cited). You should be in a position to list all the concerns and issues that influenced the reviewers' decision to include or exclude a particular study and also to decide if you agree with their decision. Finally, you must make your own interpretation and consider how the available research affects your instructional design decisions.

REVIEW

- Early research on the use of pictures as an aid to learning suggested negative effects, such as pictures as distractors to learners.
- Most of the available research deals with pictures as an aid to learning from printed text in narrative form, also known as prose learning.
- More recent research suggests conditions that need to be considered in order for pictures to help learning.
- One of the most fundamental and longstanding principles of instructional graphic design is that pictures should overlap or be congruent with the instructional content.
- A variety of prerequisite skills may be necessary before a learner is able to benefit from the inclusion of a well-designed instructional graphic, including the ability to read and to interpret moderate to complex visuals.
- Research suggests a curvilinear relationship between the realism of instructional visuals and learning. Too much or too little detail may be detrimental to learning, especially when the materials are externally paced, such as in videos.
- Learners usually prefer that aesthetically pleasing graphics be included in instructional materials, but these graphics usually do not directly contribute to learning per se. Special care should be taken to assure that such graphics do not interfere with other graphics intended to serve instructional functions in the cognitive domain.

NOTES

1. The naturalistic inquiry methodology likens the researcher to an educational anthropologist who becomes part of the situation being studied, rather than trying to view it objectively from the “outside looking in.” It is expected that a rich blend of methodologies will begin to flourish in the future. In fact, the position here is that good research begins with good questions and that different research methodologies are but different tools to be used to answer those questions. Good research should involve and blend quantitative and qualitative procedures as necessary. Unfortunately, the quantitative and qualitative research methodologies are often viewed in opposition to one another (see Lincoln & Guba, 1985, for examples of the arguments, and Neuman, 1989, for specific applications to computer-based instruction). Pitting one methodology against another is counterproductive. Readers are guided to Salomon, 1991, for a good summary of the points at issue, as well as a discussion of how the two methodologies should be effectively combined and used in order for each methodology to complement the other.
2. The term “significance,” when used in a statistical context, simply refers to the probability that the results of an experiment occurred by chance. Before conducting an experiment, a researcher sets the probability level at which the hypothesis that there is no difference between groups (known as the null hypothesis) will be rejected. Usually a probability level (p) of about 5 out of 100 is chosen (i.e., $p<.05$). If experimental groups differ enough to reject the null, the researcher is simply saying, in essence, the odds of this difference occurring solely by chance is no greater than 5 in 100. Therefore, rejecting the null hypothesis is the same as accepting the belief that the results did not occur by chance, but because of some other reason. If the study is well-designed, the researcher will *infer* that the determining factor was the factor being studied, such as illustrations. If the experimental groups are not different enough to reject the null, then the researcher will report “no significant differences,” or NSD, between the group. But is a statistically *significant* difference between experimental groups an educationally *important* result? The answer is maybe yes, maybe no — it depends on the study. Critical issues are the research questions being asked, the number of subjects in the study, the reliability and validity of the materials used, and the magnitude of the difference. The difference between significance and importance must be determined by the reader.
3. See chapter 4 for a more complete explanation of the principle of least effort.
4. Although all of the reviews included in this chapter should be required follow-up reading, Levie's (1987) review is probably the best “future reference material” of the bunch. For example, Levie provides a highly detailed and organized bibliography for each of the four areas of research he reviewed. Unfortunately, this review (as well as the one that follows by Levin, Anglin, & Carney, 1987) is published in a fairly obscure (and expensive) source and therefore may not be readily accessible.
5. Reading any *review* of research is, at best, a “second-hand” account of the research. So, since this chapter consisted of a “review of reviews,” it presented, in a sense, a “third-hand” account of the research literature. This does not mean that the original, “firsthand” research reports were not consulted. (However, the analysis of the

original research that went into this chapter was, I'm sure, much different than that which went into the reviews by Samuels, Levin, Lesgold, Dwyer, Levie, Anglin, and Carney.) Again, the purpose of this chapter was to interpret and synthesize these published reviews into a readable form for people new to education or research. Therefore, it is not fair to say that this chapter truly presents a “second-hand” account. Again, let me repeat my strong encouragement that you should read the reviews discussed in the chapter and as much of the original research as possible, then make your own conclusions.

CHAPTER 6

Review of Instructional Visual Research: Animated Visuals

OVERVIEW

This chapter reviews research on using animated visuals in instruction; it complements and continues the previous chapter. Special considerations in the interpretation of animated visual research are examined. Unlike research on static visuals, a very limited number of animation studies have been conducted so far. Research on two main instructional applications of animation are reviewed: presentation strategies and visual feedback from simulations. In addition, an overview of the goals and philosophy behind an ongoing research agenda is also discussed.

OBJECTIVES

Comprehension

After reading this chapter, you should be able to:

1. List areas in which research on the effects of animation on learning have been conducted.
2. List the two characteristics of animation that distinguish it from static visuals.
3. Describe the importance of these two characteristics in interpreting instructional animation research.
4. Describe the conditions and range under which research indicates that animation may be an effective presentation strategy.
5. Describe factors that can undermine the effectiveness of seemingly well-designed animated displays.
6. Describe problems associated with studying animation when used as part of visually based practice activities, such as simulations and games.

Application

After reading this chapter, you should be able to:

1. Critically read and interpret instructional animation research.
2. Classify a given learning outcome in a content area as appropriate or inappropriate for the design of an animated display.
3. Derive a set of design guidelines from the scant pool of available research on instructional animation.
4. Apply these design guidelines in the design and development of instructional materials involving animation.

Animation is a popular and favorite effect among producers of computer-based instruction. Reviews of available courseware for education and training demonstrate its popularity (e.g., Sales, Tsai, & MacLeod, 1991). Animation is most commonly used for cosmetic purposes, with the intent of impressing rather than teaching. All too often, animation is added to CBI without serious concern for its true instructional purpose or impact. While it is easy to be critical of designers and developers, they share responsibility with consumers who tend to evaluate instructional effectiveness based on the number of frills that a package contains. This can create a cycle where the market, rather than learning needs, drives instructional design. Equally accountable are researchers who have yet to provide adequate guidance to either group. It is hoped that an attitude of shared responsibility will begin to prevail in more fully exploiting the potential of animated visuals in learning environments.

There is nothing inherently wrong with using animation for cosmetic applications, as long as it does not interfere with other lesson functions. It is tempting to say that animation scattered throughout a lesson should add to the lesson's motivational appeal. While there is a certain amount of face validity to this position, no hard evidence supports it. If animated visuals do, in fact, increase a lesson's extrinsic motivational appeal, such appeal would be based on novelty and be only temporarily effective.

Given the proliferation and power of animation production packages on microcomputers, it is expected that animation will be used even more rampantly in the future. GUI systems are increasing the ability of nonprogrammers to author dazzling graphic displays (see chapter 3). From a production perspective, animation can be dramatic and spectacular. But does this imply that learning effects should be equally impressive? Many designers and educators probably think so. Unfortunately, the research does not support such a position. It is easy to be “seduced” into believing that special effects, such as sound and graphics, *must* be very helpful to learners. Research, however, indicates much more complex conclusions.

The purpose of this chapter is to review the relatively small pool of instructional animation research (compared to that available on static visuals). This chapter summarizes and extends an earlier review of animation (see Rieber, 1990a). Similar to research on static visuals, early results in instructional animation research were generally negative, but also prone to confounding on many counts. Recent work has begun mapping out some of the conditions under which animation can effectively aid learning. Two main areas of research have been conducted on animation in instruction: (1) as a way of presenting information; and (2) as visual feedback in practice strategies, such as simulations. Given the available research, it seems clear that animation exerts a relatively subtle influence on learning and that many factors can further undermine this effectiveness. Despite recent advances in applying learning and instructional theory to CBI design (see Jonassen, 1988a), we know surprisingly little about some of the computer's most fundamental presentation and interactive components.

Although there are many additional applications of animation beyond presentation and practice, virtually no research is available on them. For example, it could be argued that animation provides a good way to gain the attention of a student and also to cue a student to attend to the most critical features of a screen display. As already discussed, attention-gaining is an important initial event of instruction (R. Gagné, 1985). Certainly, animation affords many practical methods of gaining and cueing attention, such as special effects during transitions between screens, (see Footnote 1) moving icons or characters (including cartoons and text), and animated prompts (such as arrows that direct attention to keywords, paragraphs, graphics, or other screen items or features). Animated objects contrast with a static background, thus increasing their prominence. (See Footnote 2)

The most direct application of animation in instruction is using it to present lesson content. Animation, with or without accompanying text, offers many opportunities for presenting or elaborating facts, concepts, and principles. As discussed in chapter 4, the processing partnership between visual and verbal information is well-established theoretically (e.g., Paivio's dual-coding theory). One could describe these instructional uses of animation as "learning-by-viewing" approaches (Reed, 1985, p. 297-298). Much of the research discussed in this chapter studied animation in this way.

Although not as cleanly definable as when used in a presentation strategy, animation is also frequently used in a wide array of interactive activities. The goal of these activities is usually to practice a recently learned skill or to acquire a new skill. These can range from highly-structured to discovery-based activities and approaches. In questioning strategies, animation is often used as visual reinforcement to student answers (i.e., a "pretty picture" as a reward for getting the right answer). As warned in chapter 4, this type of feedback should reinforce only correct answers. Any "pretty picture" given for a wrong answer should never be more reinforcing than that given for a right answer. It is surprising how often this simple rule is broken.

On the other end of the continuum are highly interactive visual displays, such as simulations, in which animation is presented as a continual stream of *informational* visual feedback reflecting moment-to-moment changes based on student input. Animated displays used in this way are sometimes called **interactive dynamics** (Brown, 1983). Simulations can be used in structured learning experiences or for open-ended, discovery-based learning. The distinction between these approaches is discussed later in this chapter, as well as in chapter 8. These examples of animation in instruction can be termed "learning-by-doing" approaches (Brown, 1983).

Some examples of successful animation used in practice activities include flight simulators, physical science simulators (diSessa, 1982; White, 1984), and programs in which students learn about musical concepts (Lamb, 1982). An entirely different set of examples comes from programming languages based on graphics, such as LOGO, in which students drive an animated turtle around the screen to draw pictures (see chapter 3 and 8) (Papert, 1980; Abelson & diSessa, 1981). As discussed further in chapter 8, real-time visual feedback during a simulation demonstrates a use of animation not easily replicated with media other than computers. On the other hand, animated presentations can be generated and displayed on various media other than computer, such as film and video.

SOME IMPORTANT CONSIDERATIONS IN THE INTERPRETATION OF ANIMATION RESEARCH

Chapter 5 discussed many issues to be considered when interpreting the results (or lack thereof) of educational research in general, and static visual research in particular. Before any graphic offers the potential for increased learning, *a need for external aids to visualization must be established*. For example, before evaluating the effectiveness of a picture, reviewers (Dwyer, 1978; Levin, Anglin, & Carney, 1987) have stressed the importance of first determining whether a textual passage *alone* elicits adequate internal imaging by students. If students adequately image *internally*, then, obviously, the inclusion of *external* visuals will probably not result in any additional learning gains. Although one could argue that adding such visuals may do no harm, which may be true, there is always the potential that unnecessary visuals may distract. Even if the text does not sufficiently induce appropriate (and necessary) mental imaging, visuals must be congruent, relevant, and consistent with the information presented in the text in order to be effective (Levin & Lesgold, 1978).

Lessons learned from static visual research are believed to be relevant for animated visuals, as well. Although animated visuals can be viewed as a subset of instructional visuals, to what extent does the research on static visuals extend to animated visuals? Put another way, what distinguishes learning from static versus animated visuals? Chapter 4 described two perceptual characteristics unique to animated visuals. The most obvious is the motion attribute. By definition, animation provides the illusion of movement. However, chapter 4 described another, more subtle, distinguishing characteristic — trajectory, or the path of travel by the animated (or moving) object (Klein, 1987). These two additional perceptual factors must be considered in designing and interpreting animation research.

Predicting a learning effect on the basis of externally provided animated visuals depends on two things. First, animated visuals, like static visuals, must pass the test of a “need for external visualization,” as described above. Second, the learning of the content must depend on understanding either changes to an object over time (i.e., motion) or changes in the direction in which the object is moving (i.e., trajectory), or both. If there is no case for this second requirement, then there would be no reason why animated visuals would aid learning anymore than static visuals. In fact, a case could be made that the additional (and unnecessary) characteristics of motion and trajectory could be distracting in some way to the learner. It would also be reasonable to expect stronger learning effects when *both* motion and trajectory attributes are essential to understanding a certain fact, concept, or procedure, or in solving a problem.

Unfortunately, several early reports of animation research failed to meet these requirements. In fact, two studies frequently cited as proof of the ineffectiveness of animation in instruction fall into this category. The first contained serious methodological problems in the study's overall design (Moore, Nawrocki, & Simutis, 1979). Subjects in all treatments groups were required to answer review questions after each of four lesson parts. Subjects could not proceed through the lesson until they achieved at least 85% accuracy on these review questions. Obviously, this meant that by the time subjects reached the post-test, *all* would achieve at least the 85% performance level. Not surprisingly, there were no significant differences between treatment groups on the post-test because of this artificially induced ceiling effect (i.e., all students learned the maximum amount regardless of treatment).

There were also several serious problems in the design and execution of a second frequently cited study (King, 1975). The materials were not sufficiently difficult, which probably also resulted in ceiling effects. The test materials were also heavily weighted to measure verbal kinds of information and thus may not have been sensitive enough to parts of the lesson demanding active visualization on the part of the students. Finally, the actual graphics used were very crude.

In addition, both of these studies used an adult population and neither provided any evidence to indicate that visuals of any type were needed to learn the material. Accepting these studies as conclusive evidence for the inability of animation to promote learning is akin to basing all design considerations on picture effects solely on the studies reviewed by Samuels (1970) (see chapter 5).

Finally, animation researchers need to design their studies to control for static versus animated visuals. Otherwise, inferences cannot be made about animation directly. For example, a very early study found positive effects for animation for factual recall, but only in comparison to a verbal presentation of the materials (Rigney & Lutz [Alesandrini], 1975). Due to the lack of proper controls, there is no evidence to suggest that the animation was any more effective than just static graphics illustrating the same material.

OVERVIEW OF AN INSTRUCTIONAL ANIMATION RESEARCH AGENDA

Most of my research to date has been devoted to understanding how animation can be used to influence learning. I will share my research on animation in this chapter, but I will leave many of the design implications of this research for discussion in chapter 8. Believing that the dramatic visual effects of animation should beget equally dramatic learning effects was a trap I fell into early in my research. Fortunately, I persevered and have made progress toward mapping out conditions that both facilitate and undermine animation's effectiveness. The history of my research, albeit short, tells an important story.

My instructional animation research has had a number of simple goals:

1. Study animation in the context of varied instructional strategies and expect interactions.
2. Study animation across age groups given the assumption that people's ability to learn from animation, like any visual, changes over time (maturation effects).
3. Study higher-level learning outcomes.
4. Study animation in a suitable content area (Newton's laws of motion).

Based on the simple philosophy that learning is a complex set of events, the first goal was to study a variety of instructional strategies. I anticipated that different types of instructional strategies used in different combinations should interact (i.e., vary depending on the combination used). For my work, I have used Robert Gagné's (1985) events of instruction as the guiding set of lesson design principles (see chapter 2). Interactions should be expected among the five groups, or "families," of lesson components (orientation, presentation, practice, testing, retention and transfer). As a beginning point, I have focused on the two groups that generally account for most of the variance in instruction — presentation and practice strategies.

Given earlier research indicating that people differ in their ability to form and use images as they grow older (i.e., maturation effects) (Pressley, 1977), a second goal was to study animation using very different age groups. So far, I have used students in the early to middle grades (grades four to six) and young adults (college undergraduates). Researchers all too often overlook maturation effects. What works for adults may not work for children. Results of several studies that concluded that animation made no significant difference in learning may be attributable to this point.

A third goal of my research was to study higher-level learning outcomes, such as rule-learning and problem solving. Although lower-level learning is important and, at times, prerequisite to higher-level learning, I felt there was a need for visualization research to go beyond investigating issues involving only recognition and recall tasks.

A fourth goal was to study the effects of animation in a content area not only suitable for specific visualization issues, but also generally relevant and important to people of all ages. For this reason, I chose Newton's laws of motion. Recall the two distinguishing perceptual attributes of animation discussed in the previous section: motion and trajectory. *Both* attributes are crucial to understanding many aspects of laws of motion. The motion attribute is obvious enough for understanding laws of *motion*. Understanding whether an object should be moving or not, and at what speed, is a fundamental aspect of the laws of motion. However, it is usually not enough to simply understand whether or not something is moving and at what speed. Often, an object will move, say, from left to right or right to left, depending on the forces acting on it. These issues relate to the attribute of trajectory. In fact, motion and trajectory help define many physical principles. The best-known two are probably acceleration and velocity, which are defined on the basis of speed and direction of travel.

For the purposes of presentation, laws of motion also offered an almost unlimited assortment of examples to be presented and explained, as well as contexts in which the material could be applied. Even more important, the laws of motion afford a rich variety of practice environments, including, of course, traditional question-and-answer scenarios, but also games and simulations. A broad research literature is available from science education on how people learn the laws of motion, as well as where they have difficulty. Some of this research is known as misconceptions in science (see Eylon & Linn, 1988, and Perkins & Simmons, 1988, for reviews).

Research materials were developed according to standard principles of instructional design and visualization research procedures suggested by Dwyer (1978). A basic narrative script was developed and field-tested with children and adults. Areas in the script needing additional lesson support were determined based on the level of difficulty experienced by subjects. These pilot studies also initially established acceptable estimates of the content's validity and reliability. Although the materials vary somewhat depending on the specific goals of the study, in general, the materials have the following four lesson objectives in common:

Students will be able to:

1. Apply the rule that, without any outside forces, an object in motion will remain in motion and an object at rest will remain at rest (Newton's first law of motion).
2. Apply the rule that, in one-dimensional space, when an object at rest is put into motion by a force, an equal force applied in the opposite direction is needed to stop the object.
3. Apply the rule that, in one-dimensional space, when an object is acted on by unequal forces, the final speed and trajectory of the object is the result of the sum of the forces acting on the object from both sides.
4. Apply the rule that, in two-dimensional space, forces occurring at right angles to each other act independently on the object.

All four objectives denote application learning, or rule-learning (R. Gagné, 1985). Objectives 2, 3, and 4 are specific applications of Newton's second law (i.e., force equals mass times acceleration, or $f = ma$). In general, these four objectives corresponded to separate parts of the lesson script. While other objectives — such as applications of acceleration and velocity and varying the mass of an object — were studied at times, these four objectives remain the core focus of the learning materials.

Besides testing for the standard kinds of performance on these objectives, other data were also collected to provide a broader perspective of the influence that the animated visuals may be having on a learner. For example, cognitive perspectives on learning, as discussed in chapter 4, expect other kinds of differences in learning than only performance. Information-processing models predict that the time necessary to both encode information from short-term memory (STM) to long-term memory (LTM), and subsequently retrieve information from LTM to STM, will vary depending on a variety of factors, such as the learning outcome (e.g., fact learning versus problem solving), the level at which mental processing is occurring (i.e., shallow versus deep), and the nature of the learning and testing activities (e.g., multiple-choice questions, games, essays, simulations).

Performance data alone often do not adequately reveal information about the fluency of encoding and retrieval, whereas latency data (the time needed by a student to complete an activity) offer more information, albeit indirect, about mental processing. Although using latency data in isolation may be criticized as biased or misleading (e.g., Siegler, 1989), latency data interpreted in combination with performance data should provide a much more complete understanding of learning than using either data alone. For this reason, most of my studies have included latency data in some or all of the analyses. Some studies recorded the time needed by students to view individual frames of instruction from the tutorial, whereas other studies recorded the time needed by students to answer individual questions. In addition, other data, such as student introspections and opinions, were also often collected to see if qualitative aspects of the students' experiences matched the results of the quantitative analyses.

The presentation strategies consisted of presenting the material in a traditional CBI tutorial, generally involving four parts. The first part typically reviewed fundamental background vocabulary and concepts, such as mass, weight, and force. This part also formally presented Newton's first law of motion. The second part generally introduced Newton's second law of motion and specific applications given equal forces in opposite directions in one-dimensional space. The third part extended applications of Newton's second law to include the effects of *unequal* forces in one-dimensional space. The fourth part generally extended all of the concepts and principles covered up to that point, but in the context of *two*-dimensional space. Therefore, the tutorial was generally hierarchical in nature, that is, increasing in difficulty and making information and skills in early parts prerequisite to understanding later parts. Two basic versions of the script were produced: one for adults and one for children. In general, the script provided about 40 frames of instruction with animation embedded in about 40% of these frames.

To date, two basic types of practice strategies have been produced. The first includes traditional questions, the second includes visually based simulations. In either case, suitable activities were developed to provide students with the opportunity to practice the skills taught in the respective part. (See Footnote 3) Examples of both presentation and practice activities will be illustrated in the rest of this chapter.

Learning a Valuable Lesson Early On

Although I am not partial to reliving unsuccessful experiences, I feel there is much to gain from learning from one's errors and misconceptions. One of my earliest studies in instructional animation research is a case in point. As already mentioned, I became susceptible to the position that animation *should* be a potent influence on learning — so potent, that even a little should make a difference. With this assumption in mind, I designed a study in which animation was integrated into an orienting strategy presented before each of the four lesson parts (Rieber & Hannafin, 1988). In essence, the orienting strategy presented the main principle for the ideas in the lesson part in generalized form. An example of one of the four orienting frames is shown in Figure 6.1. I expected this one frame to provide students with an "anchor" for understanding the rest of the ideas in the respective lesson part. This is, incidentally, the underlying principle behind all orienting strategies (Ausubel, 1978; Hannafin & Hughes, 1986).

In a way, I expected each of the four animated orienting frames to act as a "magic vitamin pill" that would make the rest of the instruction all the more meaningful or "nutritious." Needless to say, the study resulted in no significant differences. Although I knew that orienting activities were generally weak instructional variables in general, I thought that animation would make a powerful enough "ingredient" to turn the orientations into an effective strategy. In addition, the fourth-, fifth-, and sixth-grade students found the content very difficult to understand, which may have reduced even further the effects of the animation. This experience made me realize for the first time that the effects of animation on learning, if any, are far more subtle than I had first imagined.

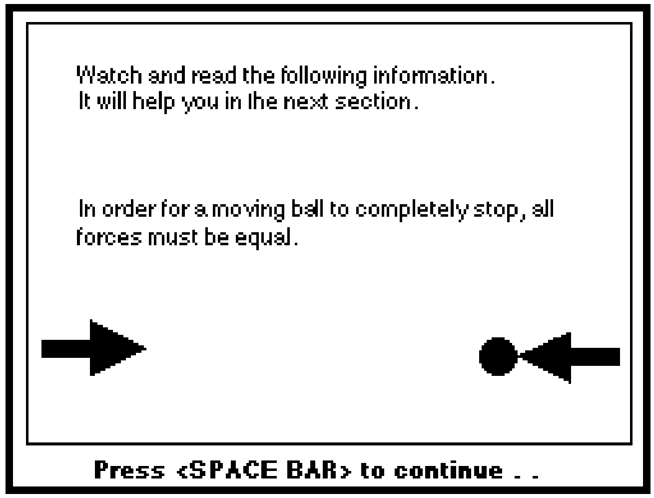

FIGURE 6.1
An example of using animation (ineffectively as part of an orienting presentation strategy. This frame was designed to introduce students to the main principle of the lesson part that followed.

REVIEW OF ANIMATION IN COMPUTER-BASED INSTRUCTION

This section reviews evidence that is beginning to demonstrate the conditions under which animated instruction may be effective. For several reasons, this review will be confined to computer applications of animation. Some of the research, such as that based on visual feedback from simulations, is largely medium-dependent, that is, there is really no other medium other than the computer in which comparable instructional activities can be designed and studied. However, this is certainly not the case with the presentation of information that “moves,” because other media, such as film and video, can be used to develop and deliver such presentations. However, the majority of research that explored the “motion attribute” from film research is dated and its generalizability is quite limited.

For example, several studies demonstrated that learning increased when using full-motion film, but primarily for procedural tasks in the psychomotor domain, such as taking a machine gun apart and putting it back together (Spangenberg, 1973). However, these results were eliminated when the instructional materials using still pictures were improved, meaning that poor instructional design explained the differences in the results more than the simply whether or not “motion” was used. (See Footnote 4) Other examples of this film research is quite useful, such as that suggesting that people with low spatial aptitude benefit from instruction containing motion sequences (Blake, 1977). There is also evidence that people's recognition memory is greater for full-motion video scenes than for static video scenes, perhaps due to the innate human abilities of perceiving movement of objects that have evolved from living in a dynamic world (Goldstein, Chance, Hoisington, & Buescher, 1982). For a thorough, yet concise, overview of instructional research of motion in films and for a general review of television research, see Chu and Schramm (1979).

In comparison to the large amount of research on the effects of static visuals (even though much of this is focused on a small number of topical areas, as discussed in chapter 5), little research is available on animated visuals, and much of what is may be confounded, as discussed earlier. In my earlier review (Rieber, 1990a), I suggested three recommendations or conclusions be drawn from the research. Although more research is now available, these three recommendations are still valid and will be used to organize the rest of this chapter. Table 6.1 summarizes the studies discussed in this section.

> **"*Recommendation 1*: Animation should be incorporated only when its attributes are congruent to the learning task" (Rieber, 1990a, p. 79).**

This recommendation is obviously an extension of the design principles extrapolated from the static visual research. However, this recommendation must be interpreted on the basis of the motion and trajectory attributes discussed previously. In order for animation to have an effect on learning above and beyond that associated with a static visual, not only must a need for visualization be present, but there also must be a need to conceptualize changes to an object over time (motion) and/or in a certain direction (trajectory).

Results from studies have been mixed. One study on teaching functions of the heart found no significant differences in comparing text, text plus still visuals, and text plus animated visuals (J. Caraballo, 1985). However, the study never presented evidence that additional visual support was needed by the adult subjects to learn the material. A subsequent study also found no differences in similar treatment groups in teaching how to compute the area of a polygon, even though care was taken to validate a need for external visualization through prior field tests (A. Caraballo, 1985). However, it turned out that the animation that actually was produced did not specifically teach the mathematical rules, but only indirectly showed relationships between various geometric shapes. For example, the program demonstrated, through animation, how two identical triangles could be combined to form a parallelogram. Generally, most college seniors and graduate students would know or remember these relationships with little prompting. Thus, the addition of animated presentations of these relationships probably had little effect on learning. Also, since both studies used an adult population, the subjects may have already been able to form *internal* images of the content, thereby reducing any benefit of the animation.

Having the benefit of this earlier research, I designed a study to investigate the presentation aspects of animation on learning about laws of motion (Rieber, 1989). I carefully field-tested the materials, as previously mentioned, to establish both a need for visualization and to specifically pinpoint locations where visuals should be added. I also used a younger

Table 6.1. Summary of Empirical Research on the Instructional Effects of Animation (listed alphabetically by author).

Study	Subjects No.	Age	Content	Learning Outcome	Results Regarding Animation
Baek & Layne, 1988	119	High School	Mathematics (ave. speed)	rules	Anim>Static> Text
Comments: Students' attention was focused on information contained in the animation.					
Caraballo, A., 1985	109	Adult	Computation of area of geometric shapes	facts, concepts, rules	NSD
Comments: Animation was used an aid to conceptual understanding, not as an elaboration of the lesson content.					
Caraballo, J., 1985	109	Adult	Physiology of human heart	facts, comprehen-sion	NSD
Comments:No pilot studies conducted to determine a need for external visualization.					
Collins, Adams, & Pew, 1978	18	Adult	Geography of South America	facts	Interactive Map> Labeled Map> Unlabeled Map
Comments: Animation for attention-gaining within an interactive graphic (blinking dots).					
King, 1975	45	Adult	Mathematics (sine ratio)	rules	NSD
Comments: Possible confounding due to an easy learning task, verbally-heavy tests which may not have been sensitive toward visualization tasks, and crude graphics.					
Mayer & Anderson, 1991	102	Adult	How a bicycle pump works	problem solving	Animation with narration> Animation only= Narration only= Control
Comments: Study involved three separate experiments testing various predictions of dual coding theory. Summary above shows general results of the final experiment. Results of the first two experiments showed that animation <u>with</u> narration supports dual coding more than narrations <u>before</u> animation.					
Moore, Nawrocki, & Simutis, 1979	90	Adult	Psychophysi-ology of audition	facts, rules	NSD
Comments: The study was confounded due to instructional design. A review was given to all students after each of four lesson parts if they did not achieve at least 85% master on the respective part. This induced an artificial ceiling effect.					

Study	Subjects No.	Age	Content	Learning Outcome	Results Regarding Animation
Reed, 1985	180	Adult	Algebra word problems	rules, problem-solving	Mixed according to word problem type

Comments: Lessons were iteratively improved over the course of four separate experiments. Animation was effective when replacing rather than supplementing verbal information.

Study	Subjects No.	Age	Content	Learning Outcome	Results Regarding Animation
Rieber & Hannafin, 1988	111	Children (elem.)	Newton's laws of motion	rules, problem-solving	NSD

Comments: There was not enough variance in the lesson treatments. Animation not powerful enough effect, especially within an orienting activity. Students also found the material very difficult to understand.

Study	Subjects No.	Age	Content	Learning Outcome	Results Regarding Animation
Rieber, 1989	192	Children (elem.)	Newton's laws of motion	facts, rules	NSD

Comments: The study may have been confounded due to the difficulty of the content. Students may not have been attending to the animation appropriately as evidenced by latency data of the time spent by students processing each frame of instruction.

Study	Subjects No.	Age	Content	Learning Outcome	Results Regarding Animation
Rieber, 1990b	119	Children (elem.)	Newton's laws of motion	rules, problem-solving	Animation>Static= None; Interaction between visuals & practice

Comments: Animation most effective as a presentation strategy when practice support was moderate. The structured simulation was effective as a practice strategy.

Study	Subjects No.	Age	Content	Learning Outcome	Results Regarding Animation
Rieber, 1991a	70	Children (elem.)	Newton's laws of motion	rules, problem-solving (intentional and incidental)	Animation>Static

Comments: Evidence that the visually-based structured simulation was intrinsically motivating for students. Students also learned a rule incidentally from an animated display, but they also became prone to a scientific misconception.

Study	Subjects No.	Age	Content	Learning Outcome	Results Regarding Animation
Rieber, in press	39	Children	Newton's laws of motion	rules	"Chunked" Anim.> Static

Comments: Animation was an effective presentation strategy, but only when screens were presented in parts, or "chunks," to aid students in selectively attending to information in the animated visual.

Study	Subjects No.	Age	Content	Learning Outcome	Results Regarding Animation
Rieber, Boyce, & Alkindi, 1991	127	Adult	Acceleration & Velocity	rules (near and far transfer)	NSD on orienting activity; Simulation results mixed

Comments: A visually-based simulation was ineffective as an orienting activity, but effective as a practice strategy for near transfer tasks only. Feedback from subjects indicated that they found the content quite demanding. They also seemed uncomfortable with the simulation in that they seemed to expect more structure.

Study	Subjects No.	Age	Content	Learning Outcome	Results Regarding Animation
Rieber, Boyce, & Assad, 1990	141	Adult	Newton's laws of motion	rules, problem-solving	NSD on learning; Animation>Static> None on response latency
Comments: Although no differences were found on performance measures, animated presentations may have aided organization and retrieval as evidenced by latency data on posttest. Structured simulation generally effective as practice strategy.					
Rieber & Parmley, 1992	160	Adult	Newton's laws of Motion	rules	Structured Sim= All Tutorial Groups> Unstructured Sim & Test Only Groups
Comments: Subjects were able to inductively learn from a structured simulation, but not an unstructured simulation. Subjects' response confidence lower without access to traditional tutorial.					
Rigney & Lutz [Alesandrini], 1975	40	Adult	Science: How a battery works	facts, concepts, rules	Pictorial group> verbal group
Comments: Since there was no control for the use of static versus animated graphics, effectiveness of animation can not be inferred from the results.					

Notes: NSD—No Significant Differences

population (fourth, fifth, and sixth graders). Despite the foresight, no significant differences were found.

However, two findings provided evidence of a "smoking gun" that may have prevented the animation from doing its job. First, the post-test scores of all the students clearly demonstrated that they found the material exceedingly difficult. (See footnote 5) Second, latency data on the time taken by students to view the animated presentations indicated that something peculiar was going on. Students actually spent significantly *less* time viewing frames containing animation, such as the one illustrated in Figure 6.2. The computer began recording the viewing, or processing, time of students only when the prompt to press the space bar to go on was presented. Students had to wait until the computer finished presenting text and any animated sequences before the prompt was given. I was suspicious that students were using for other tasks the time taken by the computer to execute the animation sequences. In other words, although the computer was presenting a *potentially* useful animated sequence from which they might learn, students probably ignored the animated sequence and used the time for other things, such as reading the screen text. Therefore, by the time the prompt to press the space bar was displayed, they were ready to move on. The difficulty of the lesson combined with insufficient cueing to the animated sequence could have been more than sufficient to confound the study.

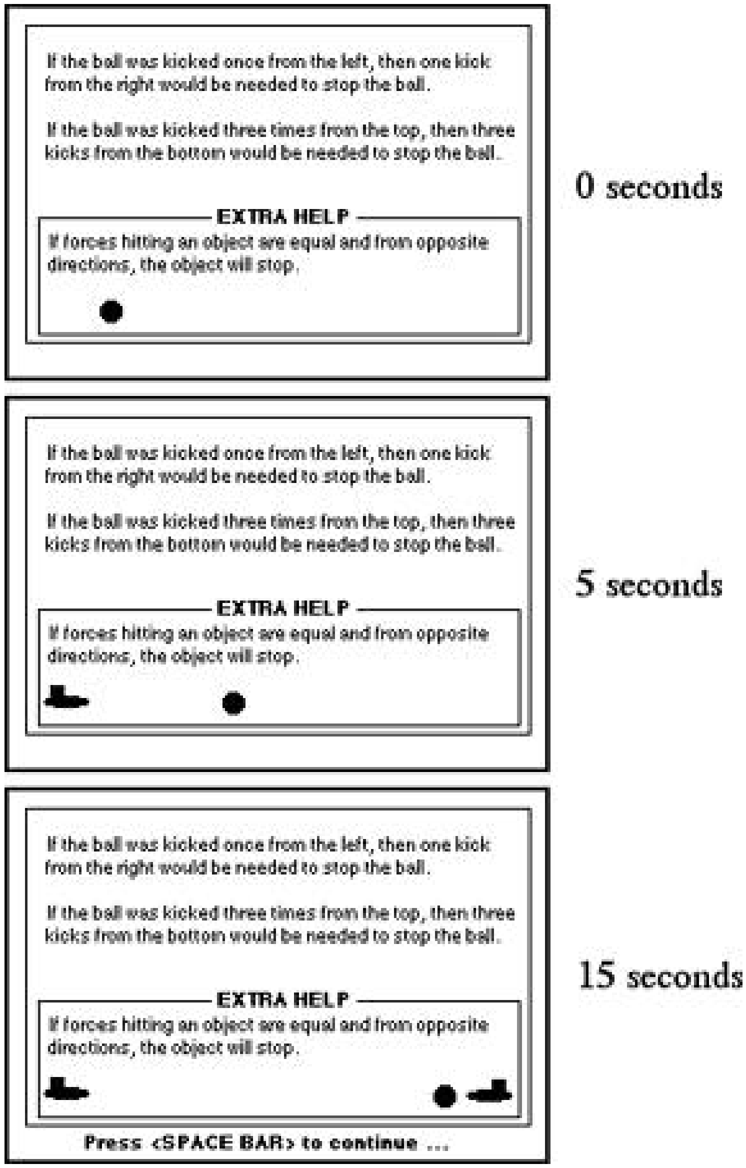

FIGURE 6.2

A time lapse sequence showing the us eof animation to visually elaborate a lesson principle.

Rieber, L.P. (1989). The effects of computer animated elaboration strategies and practice on factual and application learning in an elementary science lesson. Journal of Educational Computing Research, Vol. 5, Issue 4, p. 431-444. (© Baywood Publishing Company)

I designed a follow-up study to improve the materials based on this feedback (Rieber, 1990b). I changed the instructional design on two counts: I greatly simplified the materials and also added a special cueing strategy, as illustrated in Figure 6.3. The cueing strategy simply made it easier for students to pay attention to the animation and reduced the temptation to use the time of the animated sequence for other tasks. The strategy called for each presentation frame to be broken down into three, four, or five parts, or “chunks.” Rather than viewing one *screenful* of information at a time, students viewed a chunk of screen information at a time. Students pressed the space bar when ready to view the next chunk. Presumably, students would better attend to the animated sequence because they had no reason to do anything else — they would have already had sufficient time to read everything else on the screen.

Results of this study (Rieber, 1990b), in contrast to all the other animation studies conducted to date, clearly showed that students receiving animated graphic presentations learned more than students receiving static graphics or no graphics. There was one additional qualifier, however: this result was only found when students also received some sort of practice. (Practice was an additional factor studied, as will be discussed more in the next section.) This suggested that animation was effective, but only in the context of *full lesson support*. These results showed, finally, a situation where animation was a modestly effective presentation strategy.

This study was replicated again, but this time using an adult sample to see if the results would generalize to an older population (Rieber, Boyce, & Assad, 1990). No differences were found among the treatment conditions on the post-test measures. However, subjects' response times on the post-test indicated that those who received the animated presentations took significantly less time to answer the questions. This suggested that the animated presentations may have encouraged mental organization of the material as it was being learned. Increased mental organization of the content should result in faster, more confident, responses. This was exactly the pattern in the latency data of the post-test — students receiving animated presentations needed less time to reconstruct the information as they answered the test questions. The implication is that although the adult subjects were sufficiently able to *internally* image, allowing all groups to achieve similar performance levels, the externally provided animated displays nevertheless aided the learning process, even though the performance measure was unable to detect such differences. Open-ended comments by students after the study matched this hypothesis. Students given animated presentations commented about their value, whereas students given static graphic commented that “examples of moving balls and kicks were needed” (Rieber, Boyce, & Assad, 1990, p. 50). Students given all-text versions commented that “pictures and graphics were needed” (Rieber, Boyce, & Assad, 1990, p. 50).

A more recent study shows considerable evidence of both the range and limitations of adults learning from animated presentations (Mayer & Anderson, 1991). Students were taught how a bicycle pump works. In three separate experiments, some students watched only an animation of the principles, others heard a narration of the same information but without pictures, and others saw both the animation and heard the narration either together or with the narration coming before the animation. Students given the animation *along with* the

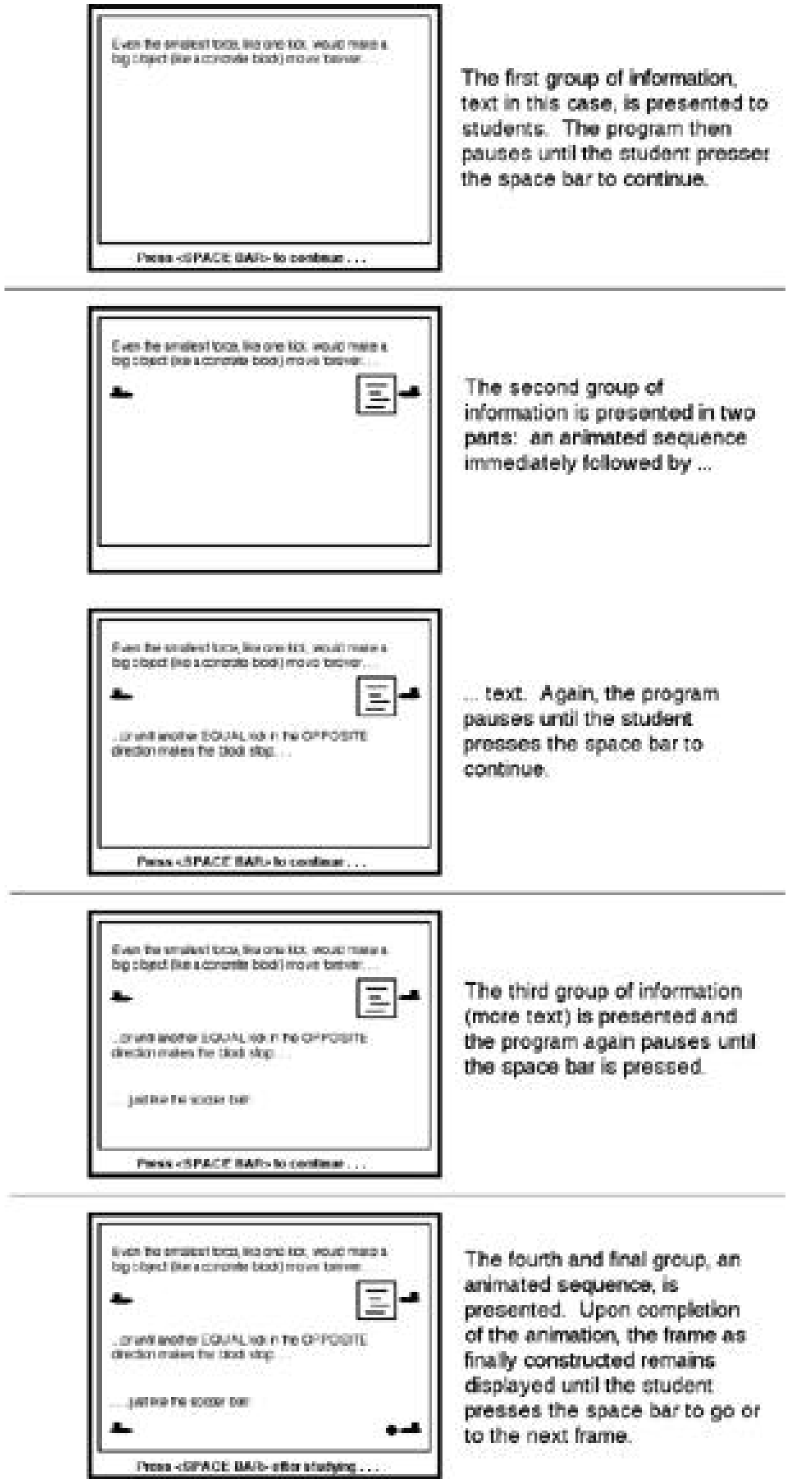

FIGURE 6.3

An illustration of a special cueing strategy which "chunked" each frame of instruction into verbal and visual parts.

Rieber, L.P. (1991) Effects of Visual Grouping Strategies of Computer-Animated Presentations on Selective Attention in Science. *Educational Technology Research & Development*, Vol. *39*, Issue 4, p. 5-15 (© Associations for Educational Communication & Technology.

narration significantly outperformed students who either in isolation watched the animation or heard the narration or who heard the narration right before seeing the animation on the problem-solving tasks. Even more important, the animation without the verbal description was completely ineffective, as students in this treatment compared equally with students provided no instruction at all. Consistent with Paivio's dual coding theory described in chapter 4, learning from animation, like any visual, is best when paired with appropriate verbal support because of the increase to both representational and referential encoding.

This series of studies has begun to shed some light on some of the conditions necessary for animated presentations to aid learning. As mentioned at the onset, the demands of the learning task must match the three attributes of animation (visualization, motion, and trajectory) in order for learning to occur. However, this is a necessary, but not *sufficient* condition for learning. Other factors can undermine the effectiveness of animation. *Some* of the factors indicated by this research include: exceedingly demanding learning tasks, poor instructional design, and the inability of students to focus on or attend to the information contained in the animated display. This final intervening factor suggests the next recommendation.

> **"*Recommendation 2*: Evidence suggests that when learners are novices in the content area, they may not know how to attend to relevant cues or details provided by animation" (Rieber, 1990a, p. 82).**

Based on the previous research, there seemed to be evidence that the "chunking" strategy shown in Figure 6.3 helped students focus on the animated sequence. However, the purpose of the Rieber (1990b) study was not to investigate the effects of this particular cueing strategy. Instead, the strategy was used throughout the instructional design of all the treatments, even those containing static graphics or no graphics. For this reason, I designed a study to directly test the hypothesis that this strategy did, in fact, account for greater selective attention of the animated information on the part of the students.

The study compared two versions of two visual treatment groups (static and animated visuals) (Rieber, 1991a). One version presented one screenful of information at a time, as was used in the Rieber (1989) study. This method can probably be considered the traditional approach in CBI design. The second method used the chunking strategy from the Rieber (1990b) study (see Figure 6.3). Results showed that students in the animated *grouped* condition performed significantly better on the post-test than students in either of the two static visual treatments (grouped or ungrouped). Post-test scores of the students in the animated *ungrouped* condition were not significantly different than any of the other three conditions (Rieber, 1991a). This study provided good preliminary evidence that the animated presentations would only be more effective than static visuals when students are properly cued to the information contained in the animated sequence. In the study by Mayer and Anderson (1991), animation presented without any verbal support was completely ineffective, indicating that students were either unable to appropriately focus on or to understand the most important visual parts of the presentation.

The implications of this second recommendation are easy to overlook when designing animated visual displays. Designers and developers forget that they become content experts of the materials they produce. Information contained in an animated sequence, though wonderfully obvious to them, may be totally overlooked by students. Even if students appear to be attending to the surface-level features of an animated display, they still may be unable to draw out, or “read,” the information contained in the animation. Why don't students pick up on the information in a seemingly well-produced animated display better or more frequently? The answer may lie in the fact that students are probably not accustomed or trained in interpreting animated information, perhaps because much of the animation they view is meant to appeal to their affective domain, such as video cartoons. This research indicates that students must be sufficiently cued and guided in order to take advantage of the potential learning effects of animation. I believe that the “chunking” strategy used in my research is only one of many possible cueing strategies that should prove effective.

A complex study that investigated the use of graphics to teach algebra word problems (Reed, 1985) suggested that students who are beginners in an area may have great difficulty perceiving differences from animation when only required to *view* the displays. The study involved a series of four iterative experiments (meaning that the results of each experiment were used to improve the design of the next). The animated displays were only effective when paired with an interactive strategy that forced students to attend to critical features of the animated display. A replication of this study (Baek & Layne, 1988) provided additional evidence that students need external cueing in order to learn from animated displays.

White (1984) has described instances when students' misconceptions of a content may interfere with their perceptions of what is actually happening in an animated display. For example, if your personal “theories” of physical science tell you that an object should be moving at a constant rate, you will probably misinterpret or ignore motion cues of an object that is actually *accelerating* at a slow rate. Again, in contrast, the designer or expert may see these differences as obvious. As discussed in chapter 4, perception is a function of prior knowledge or experiences (described as top-down processing in chapter 4).

> **“*Recommendation 3*: Animation's greatest contributions to CBI may lie in interactive graphic applications” (Rieber, 1990a, p. 82).**

This final recommendation from my earlier review (Rieber, 1990a) represents a Pandora's box of issues and applications of animation in instruction. The recommendation speaks to interactive activities in which animation plays an important role. The most obvious examples are visually based simulations, such as flight simulators, where animation is used to represent visual feedback from the artificial world modeled by the computer. In these applications, one cannot study animation per se, but only the activity within which the animation is contained. It is therefore virtually impossible to cleanly extract the effects of animation because its effects are contextually bound to the activity. For consistency and simplicity, I have repeatedly referred to such highly interactive activities as practice strategies.

This is an area in which I have just begun to do systematic research. As a first step, I have taken the position that it is more useful to compare design philosophies than small variations within a single activity. For example, in some of my work I have compared designs based on behavioral orientations, such as questioning strategies, to those based on cognitive orientations, such as visually based simulations. Although the “behavioral versus cognitive” label may be an oversimplification, it nonetheless remains a useful distinction. As discussed in chapter 4, each philosophy makes vastly different assumptions about human learning. My goal is not to divide the positions further, but merely to resolve and better understand differences in their applications to instructional design.

Relevant and sustained student interactivity is one of the most critical features of instructional design (R. Gagné, 1985; Gagné, Briggs, & Wager, 1992; Jonassen, 1988b). Successful practice strategies, such as questioning techniques, have a long history, especially for lower-level learning such as recall (Anderson & Biddle, 1975; Hamaker, 1986). Practice enhances learning in these situations by increasing overt attention to and rehearsal of relevant lesson information, combined with positive reinforcement and informational feedback (Kulhavy, 1977; Schimmel, 1988; Wager & Wager, 1985). Traditional questioning strategies have been successfully applied to CBI (e.g., Hannafin, Phillips, & Tripp, 1986), but they tend to quickly become monotonous or boring. In addition, they rely heavily on extrinsic motivation, such as rewards and reinforcers, to be successful.

Practice strategies that promote higher levels of learning demand different design assumptions (Salisbury, 1988). Learning is promoted by presenting problems or conflicts that encourage a student to use novel and original strategies, such as hypothesis-testing, to derive solutions. Cognitive psychology suggests many factors that need to be considered in designing practice strategies. These include, but are not limited to, meaningful contexts based on a learner's prior knowledge and experiences, issues related to comprehension monitoring, and intrinsic motivation (Craik & Lockhart, 1972; Keller & Suzuki, 1988; Lepper, 1985; Malone, 1981; Ross 1984).

I and my colleagues have studied visually based simulations, structured in various ways, to teach Newton's laws of motion. In most studies, students were given varied control over an animated “starship.” Students manipulated the direction and frequency of forces acting on the starship, as shown in Figure 6.4. In the experiments, the simulation was used as a strategy to practice the information and skills learned in an accompanying tutorial. These simulation activities were compared to the more traditional questioning activity, as well as to a “no practice” control.

When studied with children, those given the simulation outperformed students in the no-practice control whereas those given the questioning technique did not (Rieber, 1990b). In addition, a follow-up study provided strong preliminary evidence about the intrinsic motivating appeal of the activities (Rieber, 1991b). Students overwhelmingly chose to return to the simulation when given total freedom of choice at the end of the experiment. The questioning activity and a word find puzzle were among the students' other choices. The puzzle activity was deliberately meant to be strong competition for the simulation. Not only do children of this age group (fourth graders) traditionally find such puzzles a lot of fun, but the experiment allowed them to use the puzzle as a general “escape” from any “school-like” features they may have associated with the computer materials. Measuring intrinsic motivation is tricky business, but such free-choice methods have a credible history. Choosing to participate in an activity when all external pressure to do so has been removed is generally known as continuing motivation (Maehr, 1976; Kinzie & Sullivan, 1989).

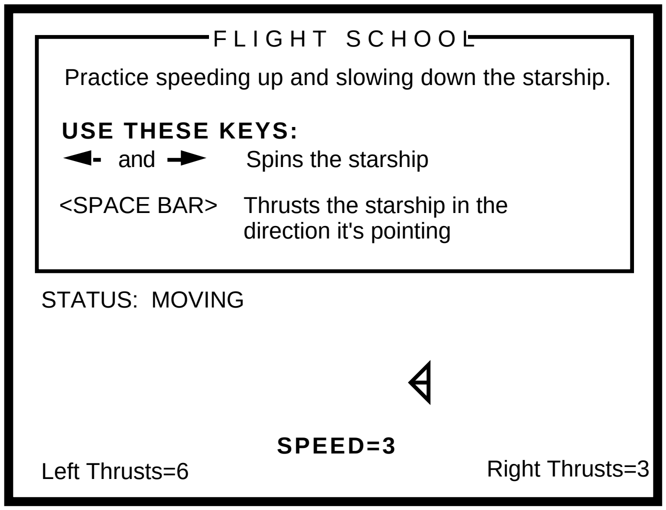

FIGURE 6.4
A snapshot of the screen during an episode of a structured simulation used as a practice strategy. The “starship” is under student control. This simulation is structured in various ways. For example, the starship spins in 180 degree increments, resulting in one-dimensional motion.
Rieber, L.P. (1990) Using Computer Animated Graphics in Science Instruction with Children. Journal of Educational Psychology, Vol. 82, No. 1, p. 135-140. © 1990 by the American Psychological Association. Reprinted by permission of the publisher.

When studied with adults (Rieber, Boyce, & Assad, 1990), subjects performed equally well given either the simulation or the questions — both groups outperformed the no-practice control. However, the simulation group took significantly less time to answer post-test questions than either the question or no-practice groups. Again, this latency data suggests that the simulation activities may have aided students' organization of the material, resulting in a decrease in retrieval time.

Research on Inductive Learning

My most recent work has been studying how people learn from the simulations with and without the use of accompanying tutorials. Again, my purpose is to compare design philosophies within CBI, which in this case is the difference between deductive and inductive learning. Deductive learning involves the most traditional approaches to education, such as presenting the rule for a concept (e.g., evergreen trees) with lots of examples and nonexamples, and providing plenty of practice. Inductive approaches involve "discovering" general rules or concepts through constant and varied interaction with specific cases (e.g., repeatedly walking through the forest until you notice different kinds of trees). The way children (and adults) learn complex "dungeon and dragon" video games are good examples of inductive learning.

A good illustration of the difference between deductive and inductive approaches is learning how to swim. A deductive approach would be the Red Cross method of carefully teaching the prerequisite skills to people step by step, such as breathing, kicking, and arm techniques, first in shallow water and then in deep. Deductive methods determine learning goals and the *best* way to achieve them in advance. Of course, deductive methods can also be boring and routine, and teaching strategies often begin to resemble one another.

The inductive method is similar to throwing someone into the deep end and seeing what happens. If the person learns how to swim, the skill is probably learned for life because of the meaningfulness of the experience. Of course, there is also the danger of drowning on your first attempt or surviving just to be afraid of water for the rest of your life!

In daily life, people use a combination of deductive and inductive strategies. Learning how to do simple plumbing or electrical jobs around the house are good examples. If you are confident, you may just take a device apart to see if you can find what is wrong. Other times, you may consult a "how-to" book for step-by-step instructions.

Research conducted so far with adults on learning inductively from computer- and visually based simulations has been mixed. For example, a "stepping stone" study looked at the use of visually based simulation activities as orienting and practice activities combined with tutorials on learning about acceleration and velocity (Rieber, Boyce, & Alkindi, 1991). Examples of the simulation activities are shown in Figure 6.5. In general, the activities were ineffective as an orientation for later learning experiences. Similar to previous studies, the simulation activity was generally useful as a practice activity, but not under all conditions. For example, the activity was effective as practice when subjects were tested on near transfer tasks (i.e., questions that closely matched the context in which learning occurred). However, the effect disappeared when tested on the same content using novel contexts (i.e., far transfer). However, the study was partially confounded by the complexity of the material. In post-experiment surveys, students stressed that they found the principles of acceleration and velocity very difficult to learn. The survey data also seemed to indicate that students were uncomfortable with the open-ended simulation activities — they expected more structure.

This study, as well as many informal experiences, has given me some indication that adults are generally very uncomfortable with open-ended, discovery-based activities, at least when they perceive the learning environment to be formal or "school-like" (such as in the case of participating in a research study). Adult educators have long echoed similar messages (Seaman & Fellenz, 1989). A recent study directly compared deductive versus inductive

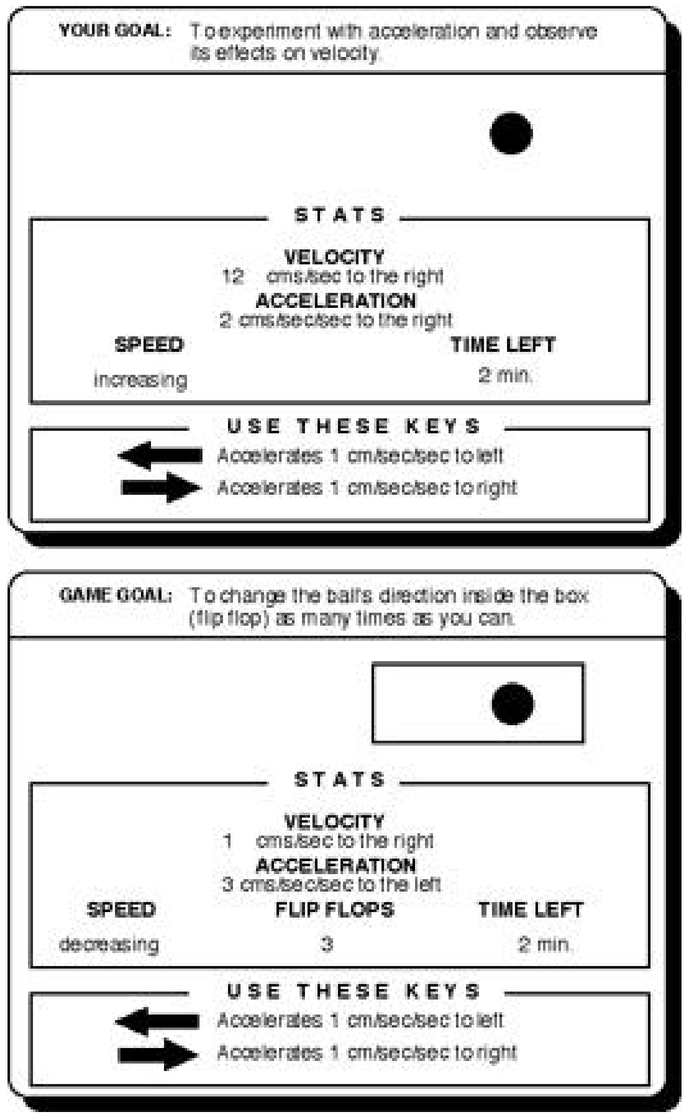

FIGURE 6.5

Two examples of simulations using game-like features that allow students to interact with the principles of velocity and acceleration.

Rieber, L.P., Boyce, M., & Alkindi, M. (1991). The effects of computer-based interactive visuals as orienting and practice activities on retrieval tasks in science. International Journal of Instructional Media, Vol. 18, No. 1, p. 1-17.

strategies, again using adults as subjects (Rieber & Parmley, 1992). Subjects given a *structured* simulation activity *without* a tutorial performed as well on performance measures as any of the conditions that included a tutorial. However, subjects given an *unstructured* simulation performed no better than subjects given no instruction at all. Figure 6.6 illustrates an example of a “structured” simulation, and Figure 6.7 illustrates an “unstructured” simulation. A simulation was defined as structured when it had a clear goal and guided students through stages of the skill to be learned. Unstructured simulations were much less goal-oriented and fully immersed subjects in all of the physical principles of the lesson.

Subjects were also asked to rate their response confidence as they answered the post-test questions; that is, how confident, on a scale of 1 to 5, they were that they were answering correctly. The most confident students in the experiment were those in any of the conditions containing a tutorial. However, the next confident were the subjects given the structured simulation/no tutorial treatment. Therefore, even though this group *performed* as well as the tutorial groups, they did not feel as confident in their learning. The lowest in confidence were subjects in the unstructured simulation and, not surprisingly, students who were given the test without any instruction.

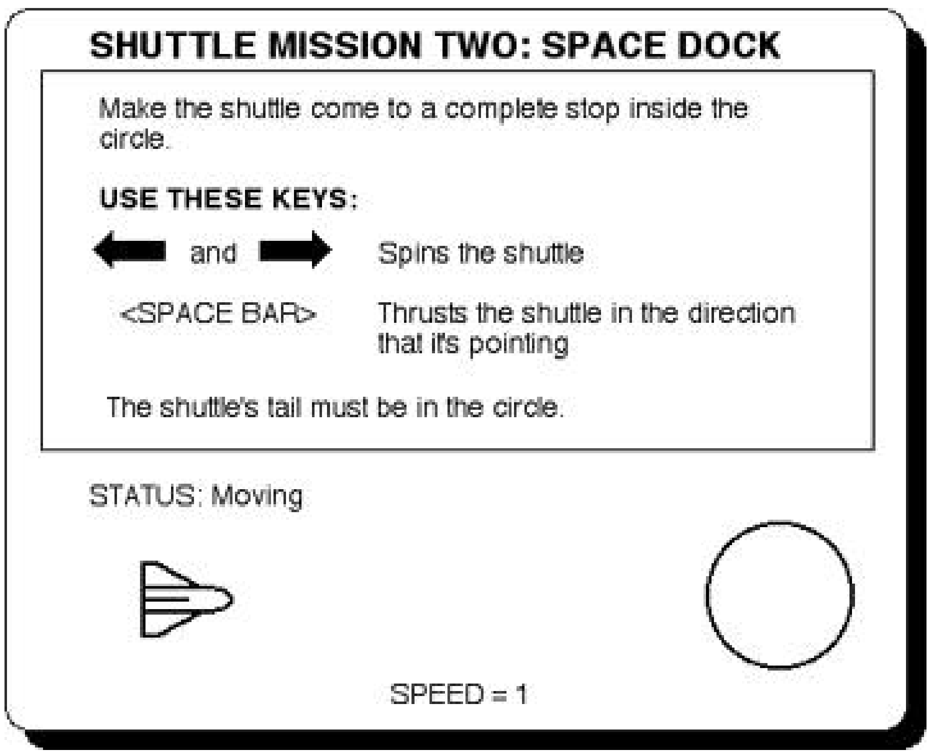

FIGURE 6.6

An example of using a structured simulation activity to inductively learn about laws of motion. When structured in ways like this, strudents were able to learn as much with or without formal tutorials.

Rieber, L.P. (1991). Animation, incidental learning, and continuing motivation. *Journal of Educational Psychology*, *83*(3), 318-328. Copyright 1991 by the American Psychological Association. Reprinted by permission of the publisher.

There is also a wealth of related research in this area, such as that dealing specifically with the design of simulations (Alessi & Trollip, 1985, 1991; Atkinson & Burton, 1991; Orbach, 1979; Reigeluth & Schwartz, 1989). In fact, I designed the simulation activities to follow closely the earlier development work by diSessa (1982) and White (1984), which is based on using simulations as “microworlds” for learning physics. The constructs of simulations and microworlds are the basis for the discussion in chapter 8.

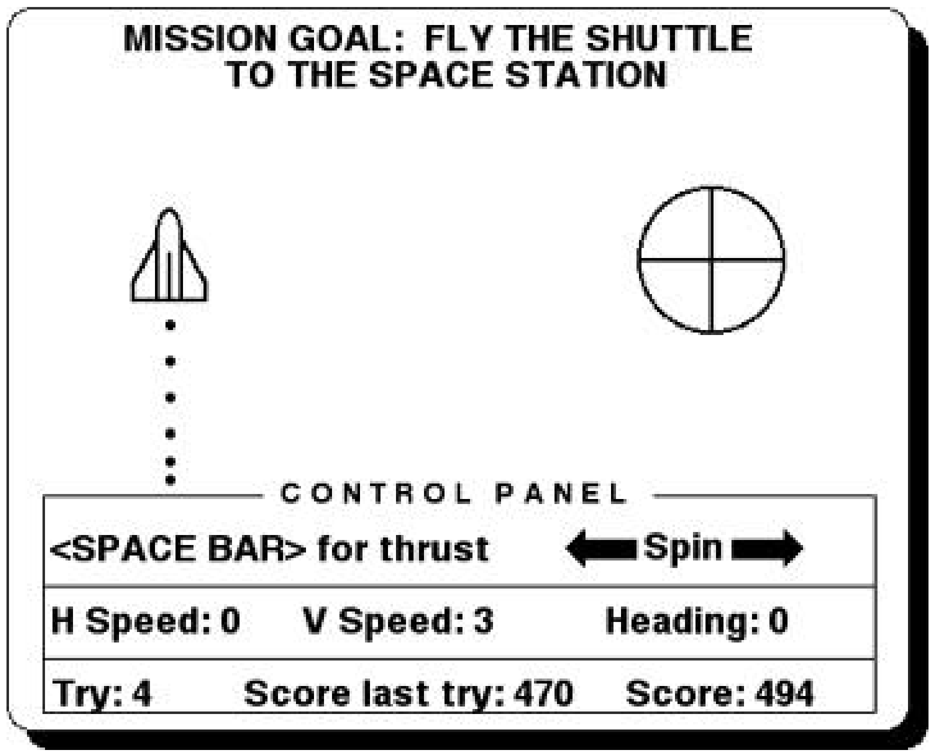

FIGURE 6.7
An example of using an *unstructured* simulation activity to inductively learn about laws of motion. Although minimal goals were established, adult subjects were not receptive to such open-ended activities. The “verdict” on whether this pattern will be found with children is still open and being researched.

Research on Learning Incidental Information from an Animated Display

An area of research somewhat related to inductive learning is incidental learning, or learning that occurs without deliberate attempt by the instruction or teacher (Klauer, 1984; Lane, 1980). Traditional instructional design, upon which most of CBI is based, has been primarily concerned with *intentional* learning, or that specified by carefully chosen and predetermined instructional objectives. Proponents of incidental learning accept the premise that a wide variety of learning is continually in progress, only some of which is anticipated.

Research has usually indicated tradeoffs between intentional and incidental learning; that is, increases to one kind of learning usually means decreases to the other.

In an early study (Rieber, 1990b), I had some preliminary evidence that students were extracting information from an animated presentation other than what was intentionally being taught. I decided to test the hypothesis that students might be able to learn incidentally from animation and to see if there were any consequences to intentional learning (Rieber, 1991b). Students were given a tutorial on a simple application of Newton's second law, where the acceleration of an object with constant mass varies depending on the size of the force that is applied to it. The larger the force, the larger the initial acceleration and the faster the ball ultimately goes. This application was an *intentional* learning outcome.

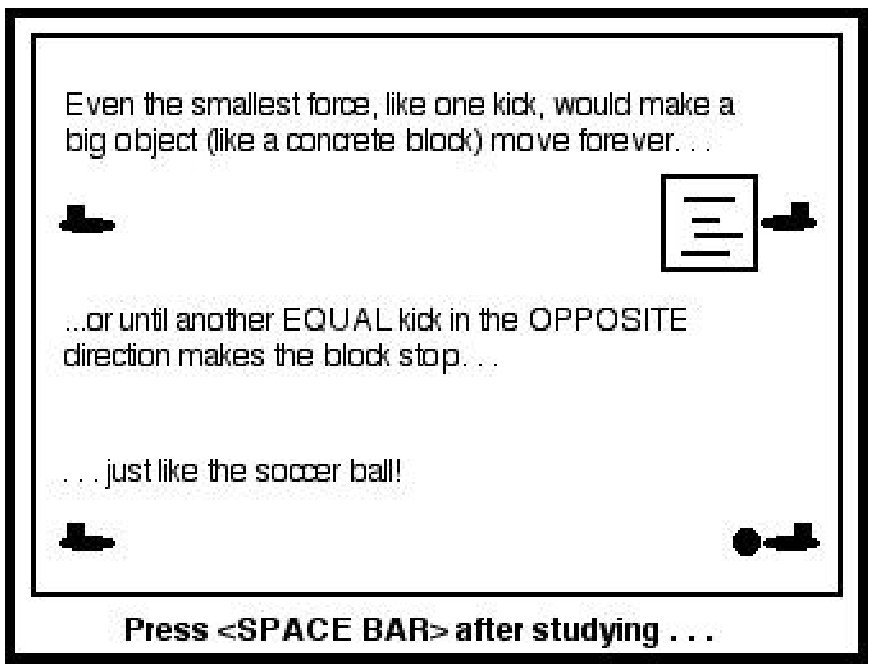

FIGURE 6.8
An example of using animation to present incidental information to students through the motion attribute.
Rieber, L.P. (1991). Animation, incidental learning, and continuing motivation. *Journal of Educational Psychology*, *83*(3), 318-328. Copyright 1991 by the American Psychological Association. Reprinted by permission of the publisher.

However, students given animated presentations were also exposed to another application of Newton's second law, where the *mass* of an object varies but the force remains constant. Through animation, students were shown the consequences of what happens when you apply the same size force to objects of different mass, such as a concrete block and a soccer ball, as illustrated in Figure 6.8. This application was an incidental learning outcome, meaning that no formal attempt was made to actually teach the application — it just happened to be part of the animation. In fact, the purpose of the animated sequence was to promote the intentional learning goal.

Results were quite startling. The fourth graders given the animated sequences successfully extracted this incidental information and applied it in appropriate ways. Furthermore, there seemed to be no obvious decrease in their intentional learning. However, there were consequences to this "extra" learning. Not only were the students able to apply this incidental information to appropriate contexts, they also applied it to *inappropriate* contexts. They incorrectly used the information to help solve problems dealing with the concept of gravity, whereas students only given static graphics were not prone to such interpretations. This "good news, bad news" story shows that students are constantly processing information in a variety of ways. Sometimes, their interpretations are constructive and relate to a set of larger goals; other times they may be building misconceptions.

Some Final Comments about Animation Research

As you can see, the available research on the effects of animation on learning is quite small. That fact has influenced me to be more systematic in my research agenda in order to efficiently investigate a wide range of issues. Other work, though largely nonexperimental, has been done on the instructional effectiveness of animated graphic displays. For example, Margaret Withrow (1978, 1979) and her associates have successfully used computer animation for languaging activities with hearing-impaired students. Other work includes motion perception research (e.g., Proffitt & Kaiser, 1986; see also the history of apparent motion research in chapter 4), testing (Hale, Okey, Shaw, & Burns, 1985), learning geography (Collins, Adams, & Pew, 1978), and understanding three-dimensional orthographic drawings (Zavotka, 1987).

Given the limited research, designers and developers should cautiously and prudently interpret and apply the research results. It is hoped that much more research will be forthcoming, especially for presentation issues, as the visualization community begins to apply its high-end computer graphics systems to instructional issues. A whole range of questions related to two versus three dimensions, texture, color, and lighting and shading, remain largely unexplored (see Tufte, 1990 for discussions on designing graphics that escape from "flat land," for example). The ending from my earlier review of animation is still very relevant here: "CBI designers are faced with a curious dilemma. They must resist incorporating special effects, like animation, when no rationale exists, yet must try to educe creative and innovative applications from the computer medium" (Rieber, 1990a, p. 84).

REVIEW

- Despite the popularity of animation among CBI designers and developers, little research is available on its effectiveness.

- Although animation can be a dramatic visual effect, research indicates that animation's effects on learning are quite subtle.
- Early animation research was heavily prone to confounding.
- In order for animation to be effective, there must be a need for external visualization of changes to an object over time (motion attribute) and/or in a certain direction (trajectory attribute).
- Children and adults vary in the degree to which they benefit from animated displays.
- Learners may need to be carefully cued to information contained in an animated display.
- Young children seem able to extract information incidentally from animated displays, although they may form misconceptions without proper guidance.
- Animation, as continuous visual feedback, is an important part of visually based simulations, although the role that animation plays in such activities cannot be isolated and studied apart from the activity itself.
- Research indicates that visually based simulations can be effective practice strategies, as compared to traditional questioning activities.
- Visually based simulations have shown to be intrinsically motivating for children in intermediate grades.
- Early research on using visually based simulations as inductive learning strategies indicates that adults are frequently uncomfortable with open-ended, discovery-based activities, especially when they perceive the learning environment to be formal or “school-like” in other ways.

NOTES

1. Many authoring products offer a wide range of transition effects between frames, such as “slide right,” “slide left,” “jaws” open and close, “Venetian blinds,” “barn door” open and close, etc.
2. In some early investigations of the use of motion, Dwyer (1978) examined using arrows as a cueing strategy (though not animated arrows). His results were inconsistent. The arrows were effective at times, but this result did not generalize across treatment groups and measures. For example, arrows seem to cue students to relevant information when used with line drawings, but not with realistic visuals (Dwyer, 1969, as cited in Dwyer, 1978). However, when students were prompted in advance of the kind of learning to expect, moving arrows helped them learn from realistic visuals, but not line drawings (Dwyer, 1977, as cited in Dwyer, 1978).
3. I have further developed the various research materials into a separate software package called *Space Shuttle Commander*. Too many times, I feel, researchers “preach” about what designers and developers should do without fully understanding the design and development cycle. I decided to try to model many of the lessons I had learned and then make the result, however humble, available to educators. SSC is being distributed free of charge to educators through the Educational Technology Branch of NASA (Rieber, 1990c). SSC and its underlying instructional principles are discussed in detail in chapter 8.

4. There is some additional value to this study. Unfortunately, poor instruction is sometimes more frequently the rule rather than the exception. It is useful to know which features might help to "make up" for otherwise inadequate instruction.
5. In fact, none of the 12 between-subject cells scored higher than 60% on the post-test. The study had, incidentally, a complex design involving three between-subjects factors and two within-subjects factors. Consult the original research report for more details. Given the general NSD findings, this all became rather academic (I mean this literally, as this study was used for my doctoral dissertation).

CHAPTER 7

Designing Graphics for Computer-Based Instruction: Basic Principles

OVERVIEW

This chapter synthesizes information contained in preceding chapters and applies it to the instructional design of computer-based instruction. One premise is that graphics can only be designed in the context of the entire instructional system. Therefore, this chapter carefully examines the relationship between instructional design and development from both traditional and alternative views. In particular, this chapter compares and contrasts formative evaluation with rapid prototyping techniques, which appears to be aptly suited to computer-based instruction. This chapter also introduces the concepts of screen design, frame protocol, and procedural protocol, and discusses the related areas of human factors and the roles of color and realism in visual displays. The chapter concludes with instructional graphic design recommendations for four of the five instructional applications introduced in chapter 2 — cosmetic, motivational, attention-gaining, and presentation. The fifth application, practice, is the topic of the next chapter.

OBJECTIVES

Comprehension

After reading this chapter, you should be able to:

1. Define and describe formative evaluation.
2. Define rapid prototyping, and summarize its fundamental principles and assumptions.
3. Define modularity and plasticity, and describe their implications in the instructional development of various media.
4. Compare and contrast rapid prototyping with formative evaluation.
5. Define frame protocol, and describe how and when a screen should be divided into functional zones.
6. Define distribution of emphasis, and describe two strategies for amplifying the most relevant screen information.
7. Define procedural protocol, and describe its importance in designing effective human/computer interfaces in instructional materials.
8. Discuss the role of color and realism as instructional variables.

APPLICATION

After reading this chapter, you should be able to:

1. Conduct rapid prototyping procedures to design and develop instructional materials that appropriately incorporate visualization techniques.
2. Design computer displays that demonstrate effective and consistent frame protocol.
3. Design computer-based instruction with effective and consistent procedural protocols.
4. Design computer displays that effectively incorporate basic principles of graphic design, color, and realism.
5. Design effective and functional instructional computer graphics for cosmetic, motivational, attention-gaining, and presentation applications.

Up to this point, this book has laid the groundwork for the role of graphics in instructional design. Previous chapters have explored and considered knowledge bases most relevant to the design of instructional graphics — learning theory, instructional theory, instructional design, and research of instructional visuals (static and animated). Issues surrounding computer graphics production also have been considered. This book has stressed the importance of designing visual materials in relation to an entire instructional system. That is, instructional design must try to take into account all of the instructional variables inherent in a learning environment, rather than considering one or more in isolation. Our intention has not been so much to consider the design of *visuals* for instruction, but to approach visualization techniques as important tools for designing all instructional strategies.

The purpose of the next two chapters is to consider the design of instructional materials, given the powerful arsenal of computer visualization techniques. These are *not* meant to be chapters on graphic design, although such knowledge or skills would obviously complement the ideas contained here. Specific design recommendations are presented in the next two chapters to help apply graphics in instructional design.

COMPUTER GRAPHICS AND INSTRUCTIONAL DESIGN

Instructional design is complex and involves a dynamic blend of subjective and objective processes. Instructional designers must appropriately identify and maintain an effective balance of hundreds of variables (“juggle” is a good metaphor). They do not simply and objectively apply a series of stoic principles culminating in effective instruction, nor do they merely use intuition and guesswork. The best instructional designers blend analytical, empirical, and artistic approaches in ways even they probably do not sufficiently understand.

Instructional designers need to recognize their personal philosophies of learning and instruction, because these philosophies ultimately account for the instructional products they produce. Consider, for example, how one's philosophy of how people learn impacts one's view of instructional design. For example, compare the differences between behavioral and cognitive philosophies (discussed in chapter 4). Since each orientation makes radically different assumptions about how people learn, instructional designs based on each philosophy will necessarily be different. One's philosophy of learning “bubbles up” to influence the resulting instructional design, as shown in Figure 7.1.

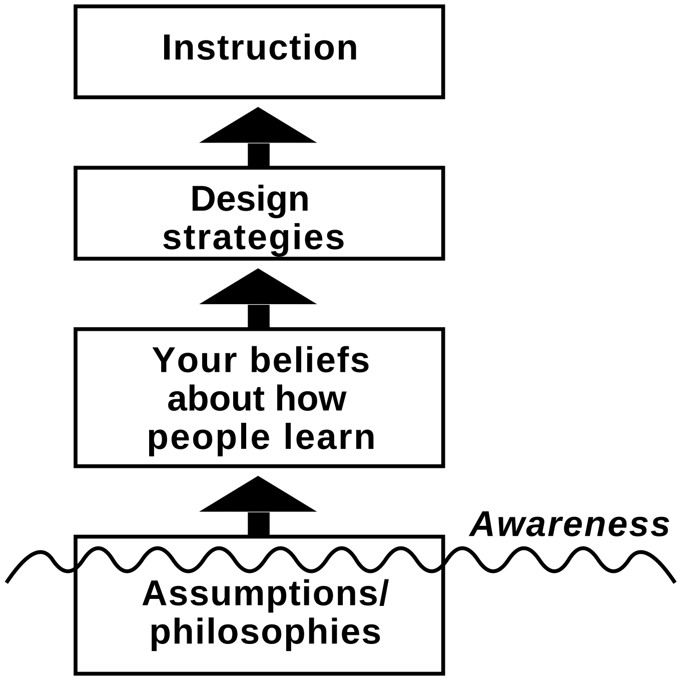

FIGURE 7.1

An illustration of the cause/effect relationship between the assumptions and belief systems of the instructional designers and the instruction that results. A designer's fundamental assumptions and philosophies, even those below the awareness level, affect beliefs about how people learn. These affect instructional design and, in turn, the final instructional products which follow.

In the next section, we will briefly revisit traditional views on instructional design and compare these to possible alternatives. Whether you see more differences than similarities will largely depend on your personal philosophy and understanding of instructional design. The importance and relationship of these issues to graphics is simply that the role and design of instructional graphics are necessarily linked to the overall design philosophy under which a design is operating — the designer must be fully aware of this philosophy, as well as other related assumptions, and, potentially, biases.

Traditional ISD

Chapter 2 presented an overview of the fundamental principles associated with traditional instructional systems development (ISD). The focus here is on a few key elements of lesson design. This section concentrates on the results of traditional ISD in the "real world," even though its intent and purpose may be more flexible in theory. As a reminder, these principles will be discussed in the preparation of instructional materials, as opposed to the process a teacher goes through in preparing a lesson for delivery in a regular classroom. Although aspects of the processes are similar for both, materials-centered instruction goes further and considers the selection and development of instructional media. This last condition (or assumption) is an important distinction, which also makes the role of graphics particularly relevant.

Traditional views of instructional design at the lesson level begin with identifying lesson objectives. Although the intent of traditional ISD is to continually remain open to revision, there is a tendency to avoid questioning the appropriateness of the lesson objectives after they have been identified.

The next step is to begin the design phase. A lesson plan or "prescription" is drafted and revised until a lesson **script** is written. The script becomes the starting point for the development phase. After the first drafts of the instructional materials are produced, the process of **formative evaluation** begins. The purpose of formative evaluation is to improve the materials through a structured series of field tests. In contrast, the purpose of *summative evaluation* is to provide information on the *final* effectiveness of the materials without intending to use the information for revision (Gagné, Briggs, & Wager, 1992).

Usually, the field tests begin with one-on-one trials with individuals most representative of the target population. The designer/developer observes these individuals' early reactions to the materials. Only the major weaknesses of the design can be identified at this stage. The formative evaluation process usually proceeds to small group tryouts and then to environments with conditions and constraints that designers believe closely resemble those of the final installation. All along the way, feedback from each field test is used to successively revise the design from which new materials are produced.

Obviously, the costs of producing instructional materials varies widely depending on the medium chosen. The development of video materials, especially those involving studio productions, can be very expensive due to the costs and demands of equipment and technical personnel. Before setting one foot in the studio, it is usually necessary to have all aspects of the production finely detailed and orchestrated, down to every word in the script and every aspect of every camera angle. There is little tolerance for creative "fudging" with the production at this stage. Given studio production costs, it is important to get it right the first time. But how can one be sure that the instructional design is appropriate and will work? It would be impractical to assume that the production can simply be repeated until the instruction is appropriate.

One common strategy is for instructional designers to use **prototypes** of the materials early in the formative evaluation process. That is, even though video may have been selected as the most appropriate instructional medium, nonvideo materials can be developed first to test various aspects of the instructional design. For example, early versions of the materials may be tried out in print-based form, such as with the use of printed texts and graphics, so that by the time the first studio session is conducted there is confidence in the design.

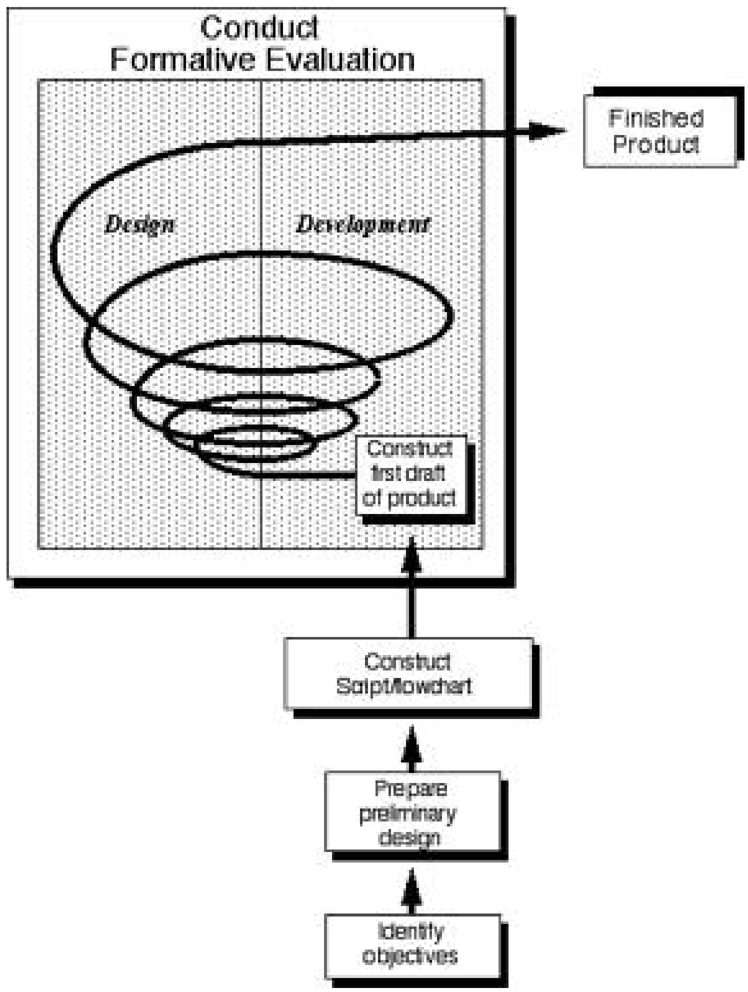

FIGURE 7.2

An illustration of the cyclic nature of formative evaluation as it is usually practiced within traditional ISD. The flow lines represent feedback loops from one stage of the process to the next. Once formative evaluation begins, design and development provide a continual spiral of feedback as early drafts of materials are field-tested and revised resulting in a finished product.

A characteristic of formative evaluation in traditional ISD is that design and development remain separate processes, though each is expected to provide critical feedback to the other. This revision cycle spirals from the earliest rough drafts to the final instructional product in media form, as illustrated in Figure 7.2. Traditional ISD frequently assumes and expects that designers and developers will not necessarily be the same people. A strength of ISD is that it allows for effective management and communication of many people working together with limited resources. Still, the process of formative evaluation in traditional ISD is prone to the weakness that some designs may be committed to prematurely once development work has begun, even though traditional ISD speaks to an openness and willingness to revise a design throughout the process. It might be argued that such weaknesses, though unfortunate, may be a necessary part of the process, given certain media (such as broadcast-quality video).

It is reasonable that certain media, such as video, simply demand a more rigid delineation between the design and development phases. But is this true for all media? What about print-based or computer materials? Word processing and desktop publishing have transformed the process of producing high-quality printed materials from raw copy. Even video offers many examples of inexpensive and portable equipment, such as camcorders and small-format editors, which offer greater availability and flexibility than studio equipment. For the rest of this chapter, we will focus on how the computer affords a much different design and development cycle than that suggested by traditional ISD. The next section describes what some perceive as an alternative approach to the formative evaluation procedures commonly discussed by traditional ISD advocates.

Rapid Prototyping

Rapid prototyping describes a design philosophy based on the idea that design, development, and implementation can never truly be separated and distinguished from one another. To some, rapid prototyping may represent mere extensions of formative evaluation of traditional ISD; for others, it represents a dramatic shift — a true design alternative. Again, your interpretation will largely depend on your current view of instructional design. (See Footnote 1) Tripp and Bichelmeyer (1990) present a compelling argument that rapid prototyping represents a philosophical shift in instructional design.

Any one instructional design is but a reflection of the designer's interpretation of the problem at a given time. Therefore, instructional design can be characterized in as many ways as there are instructional designers. To suggest that there exists any one model of instruction may be totally misleading. Design only begins to assume some meaning or value when it is implemented. For this reason, it makes little sense to try to pretend that it is possible to accurately prescribe and predict anything in advance. By definition, instructional design is relevant and important only in terms of its proven effectiveness with the intended learners in the intended setting. Given the complexity of the design process, designers are necessarily making decisions with either incomplete information or no information. Therefore, much design is necessarily based on conjecture.

On the other hand, it would be a mistake to believe that designers who practice rapid prototyping are actually “winging it.” Rapid prototyping assumes the need for any of a number of required knowledge bases — including an understanding of learning theory, instructional theory, and instructional practice — as well as profiles of the intended learners, an understanding of the content area, and consideration of any number of contextual constraints (such as the physical learning environment, motivational factors, availability and limitations of resources, and so on). A thorough understanding of traditional ISD (including its strengths and weaknesses) would also be an important knowledge base. Design associated with rapid prototyping does not begin with blind guessing, but with intelligent and informed first-generation prototypes, which may or may not develop into a final product.

There are a number of critical principles and assumptions associated with rapid prototyping that potentially serve to distinguish it from formative evaluation in traditional ISD. The first is the relationship between design and development. Rapid prototyping assumes that key aspects of design, including even a more complete understanding of the problem (and therefore, the identification of the lesson objectives), can only be determined as a consequence of constructing and testing prototypes. In this way, design and development become interdependent and intertwined and, in a sense, are really one process, though perhaps separated into individual activities and tasks to simplify managing the process. In reality, the degree to which design and development are either implemented as one or two processes is more of a continuum than a dichotomy, such as that represented in Figure 7.3. There seems to be a point somewhere on this continuum where one crosses the threshold between a traditional application of formative evaluation and rapid prototyping.

If design and development act as one process, the final definition of the lesson objectives can only be accomplished as a result of the prototyping phase. This is in contrast to traditional ISD, where the objectives are usually fixed as design and development are initiated. (In theory, it is meant to remain flexible and open, but this is usually not the case in actual application.) An example of a rapid prototyping model, adapted from Tripp and Bichelmeyer (1990), is shown in Figure 7.4. In addition, rapid prototyping offers the flexibility to test radically different designs and approaches, including rival hypotheses. Classic examples are those based on structured learning approaches, compared to letting individuals discover principles on their own.

Another critical assumption of rapid prototyping concerns the medium within which the designer is working. Not all media lend themselves to rapid prototyping techniques. Tripp and Bichelmeyer (1990) note that “rapid prototyping presupposes a design environment that makes it practical to synthesize and modify instructional artifacts quickly” (p. 38). Media must possess at least two important attributes — modularity and plasticity — in order for rapid prototyping to be practical (Tripp & Bichelmeyer, 1990). Computer tools are usually able to easily satisfy both attributes.

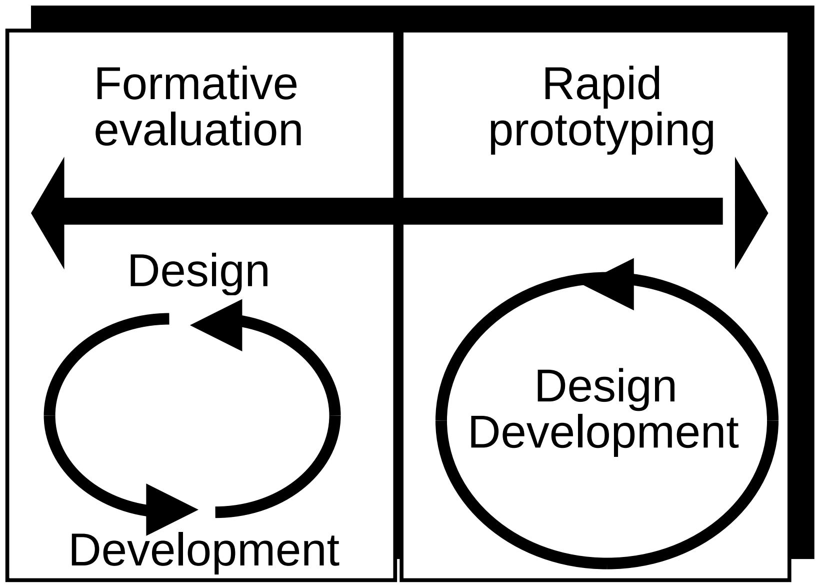

FIGURE 7.3
The relationship between instructional design and instructional development is continuous, not dichotomous. There seems to be a point where design and development act as one process. This "threshold" represents a distinguishing characteristic between formative evaluation and rapid prototyping.

Modularity refers to the ability to add, delete, or rearrange entire sections of the instruction quickly and easily. The way many people prepare presentations or lectures based on the use of overhead transparencies is a good example of modularity. Plasticity refers to the ability to make modifications to the existing prototype materials with only minor time and effort costs. Obviously, different media vary widely in regard to plasticity. Again, take video, for example. Imagine shooting a scene set in the context of the old west as the pioneers are traveling on the Oregon Trail. What would the designers need to do if they discover the next day that many of the shots inadvertently show a distant highway with an occasional car or truck? There would be no alternative but to reshoot all of the affected scenes. Most currently available computer tools offer optimal environments for rapid prototyping based on the constraints and assumptions of *both* modularity and plasticity, especially those based on graphical user interfaces (GUIs) as discussed in chapter 3. Box 7.1 presents an overview of rapid prototyping principles in a unique way — by comparing instructional design to building a paper plane.

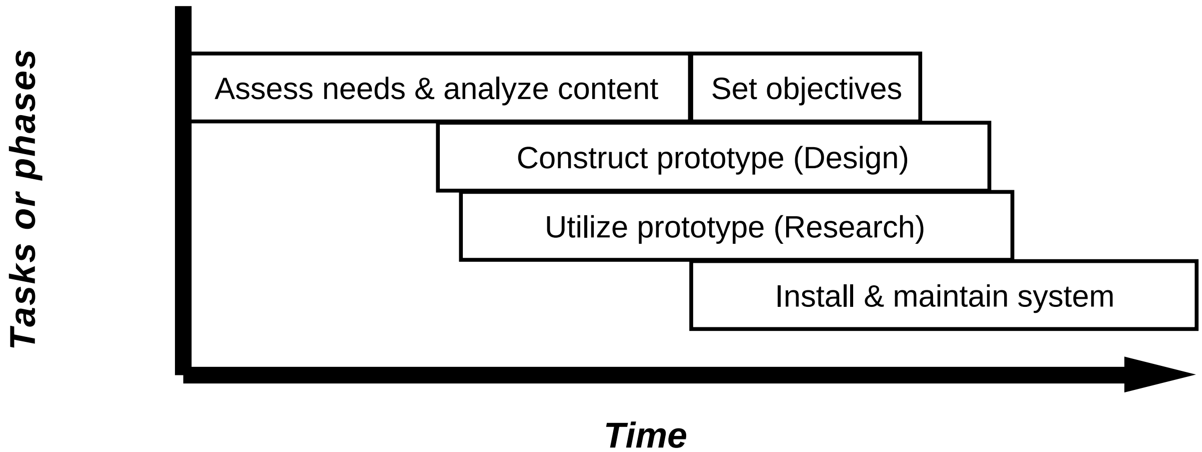

FIGURE 7.4
A rapid prototyping model applied to instructional design adapted from Tripp & Bichelmeyer (1990). A critical aspect of this model is that final identification of the lesson objective/s occurs *while* constructing and utilizing the prototype.

Box 7.1

Understanding Rapid Prototyping by Analogy: Making Paper Planes

Here's a simple example which may help you to understand the important principles and assumptions of rapid prototyping (RP). As Tripp & Bichelmeyer (1990) point out, RP has been a common tool for designers and engineers in many other fields, such as the aerospace industry. Fittingly, this example uses an everyday understanding of a complicated system — aeronautics — in the context of a universal experience (I hope) — building a paper airplane. The intent is to use this example as an analogy to *instructional* systems. If possible, find someone to do this project with you (I suggest a child). As you do it, get in the habit of expressing whatever thoughts are on your mind at the time.

Take a 8 1/2" X 11" sheet of paper and, without doing any research, make a paper airplane. Go ahead and try it out. How well does it fly? Not so good? What should you do, start over? Resist this temptation at first, and, instead, try to make some modifications to your plane that you think will help it fly better. For example, based on your personal theory of flight, fold the back edges of the wings either up or down slightly. Try flying your plane again. How much a difference does this little modification make? Continue to make additional minor modifications. Go with your feelings and intuitions.

Make modification after modification and test lots of hypotheses regarding what you think should improve the plane's design. Maybe get paper clips and pennies to see if adding weight helps or hurts. You'll probably discover many important principles of flight as well as the boundaries and limits to these principles. Spend *at least* 10 minutes in this experimental phase, still only using the first plane you constructed.

Now, stop for a moment to pause and reflect on what you just experienced and what you think you now know about the design of a paper airplane. Reflect on the meaning and importance of the design and development cycle and how important your test flights are to understanding if your design hypotheses are worth pursuing further. How many throws does it take to test one design hypothesis — one, two, five, ten? Consider the simple and straightforward feedback that each toss give you. How many total flights have you made? You probably lost count.

Take another sheet of paper and do it again. See if this next attempt begins with a better or more refined design. Are there things you will immediately do differently as you begin to test this plane? Unfortunately, you have probably been testing your plane very subjectively up to this point. That is, you know a good flight when you see it, but you probably have not, as yet, expressed in some objective way what you feel are the characteristics of a good flight. Therefore, let's give some serious thought to the testing of this plane. What criteria should you use to judge the effectiveness of the plane's design? Most people usually use distance as an important measure, perhaps followed by accuracy. Try to make your testing as objective as possible. Set up a testing environment which builds in your criteria and start collecting data. Let's focus on distance and accuracy. Clear a flight path in the room where you are working (or better yet, find an unobstructed hallway). Choose a starting line. Throw your plane and keep track of the number of times your plane lands within a certain path (to test for accuracy) and also how far your plane travels (to test for distance). Keep a careful record and study your data. Calculate some statistics. What is the average distance? What is the percent of accurate flights? Get some other people involved and have contests. See who can design the best plane.

After you've arrived at the winning design for a paper plane, do something that is both unnerving and disconcerting: take another sheet of paper, crumble it up into a ball, and see how well this design fares under the testing environment you've devised. My guess is that this "plane" does as well or better for both distance and, especially, accuracy than the planes you've designed up to this point. Why? There are several reasons. Distance and accuracy, though important, do not capture other important elements of aerodynamics related to concepts such as gliding ability or lift. You also probably did not control for the amount of force allowed for each throw. Add these to your criteria and see what happens. Some designs seem better able to "ride on" or glide in the air.

How similar are the designs of all the people involved? Unless there is an individual in the group who either is very creative or has an interest in paper planes and knows other designs due to past experiences, chances are the planes are almost identical in terms of their fundamental design. The first design most people come up with usually resembles something like this:

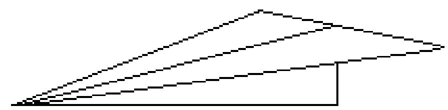

See the last page of the chapter for examples of other radical designs (insert a graphic of a "flying wing" and a "ring wing" plane on a later page). Build them, test them, and compare their results to those you've created.

What is this little activity supposed to teach us about rapid prototyping of instructional materials? First, notice that design and development were intertwined and interdependent. Did you plan out your design by, say, drawing it out on a separate sheet of paper? Of course not. Not only is it a rather silly thing to do, it's also very difficult. Try it sometime. How do you represent the procedural nature of the design? How do you represent hidden folds or the strategic tearing of paper. In fact, it is much simpler to show your design in the completed model. Similarly, it is important to consider the "gulf" between design and development of instructional materials. Even in traditional instructional approaches, design and development are expected to provide important feedback loops to the other. You are expected to learn about design through the testing of early prototypes. In other words, instruction is meant to be improved over a succession of design and development cycles. This is the role of formative evaluation at the lesson level. Unfortunately, all too often there is too great a gulf between design and development so that by the time the first draft materials are developed, there is already too great an investment in the original design. There is a high risk early on to growing complacent and accepting and committing to inferior designs. RP assumes that you will learn much, if not most, about your design only through rapid turn-over of design, development and evaluation.

Second, consider the medium in which you have been working — paper. As Tripp & Bichelmeyer (1990) note, RP assumes that the medium satisfies the attributes of modularity and plasticity. Paper allows you to fold, bend, and shape the plane in a variety of ways. It allows you to change small, but important features of the plane's form quickly and easily (remember the bending of the back of the wings).   Adding and subtracting elements, such as weight via paper clips and pennies, can be tried after the basic design is finalized. How well would other media work? Try using light and then heavy cardboard. Now try using modeling clay. You might as well throw a rock. The lesson here is that the medium must still be appropriate for the task. Paper, unlike clay, has some inherent characteristics which are appropriate for flight: light, strong, holds shape, long and flat while still rigid enough to catch and ride the air. Similarly, one cannot choose an instructional medium arbitrarily. The medium must still be appropriate for the activity. The virtues of RP, therefore, cannot be realized in every instructional medium, at least not at the same implementation level.

As you built the plane you probably felt the urge to try it out as soon as possible. Although you were probably careful as you made your folds, you probably did not want to spend more than a few minutes designing and constructing your plane. That's easy to understand, because the fun is in flying the plane, not building it. You also know from the start that the value and interest in the project is in the act of flying, not in the act of folding. Although one could make a case about the aesthetic appeal of the plane (as in the case of building a store-bought plastic model kit), most people rarely use aesthetics as a way to judge their paper plane. Similarly, the value and interest in an instructional design is in its implementation. You learn about the value of the design only by developing it and trying it out. Did you expect your plane to fly perfectly the first time you threw it? Perhaps you had high expectations at first, but you probably discovered soon that the plane would never fly as well as you first thought. Similarly, the only way to get a realistic analysis of the completeness, appropriateness, and limits of the goals of an instructional design, and subsequently, its effectiveness, is in the actual context of its use.

Consider how you arrived at your first design. The traditional model for a paper airplane is usually the only one with which most people experiment. Thereafter, it is usually very difficult to come up with truly alternative, creative designs. The same probably holds true for many instructional designers. They tend to think of instruction in only a limited number of ways. The traditional paper plane design might be analogous to a traditional application of Gagné's events of instruction. Once you have formulated your own sense of what instructional design is or should be (such as "first present information to students, and then you have them practice it"), it can be difficult to consider alternatives. Although one may make many minor adjustments (such as trying different questioning strategies), the basic fundamental design remains the same. RP allows (encourages) you to design, develop, and compare seemingly opposite designs.

The ways in which you tested the paper plane also offer very important lessons for instructional design. Feedback from testing is only valid if you are gathering appropriate kinds of information and interpreting it in appropriate ways. Distance and accuracy seem like obvious data to collect — until you compare the performance of a paper ball (especially one filled with paper clips). When testing instructional designs, one needs to be sure that the tests properly reflect the objectives. RP goes further to help designers to fine tune their initial objectives and to determine others. Too often, we only interested in performance data, like scores on a posttest, and fail to consider other sources of information, such as a student's motivation to participate and persist in an activity. It might be argued that almost any design will produce results if we somehow require (or force) students to comply, such as with external incentives such as grades, much like the paper ball being thrown as hard as possible simply to satisfy the goals of distance and accuracy. Perhaps the lesson for instructional designers is to be on a constant vigil for information which may provide useful insights to improve the instruction. In addition, it is important to remain open to consider information from a wide variety of sources. The best designs for a paper plane take advantage of the plane gliding through air with only the slightest momentum. So too, the best instructional designs usually work by intrinsically motivating the student to go as far as they can with but the slightest prompting or coaxing.

Traditional ISD versus Rapid Prototyping in the Design of Instructional Computer Graphics

The purpose of comparing formative evaluation to rapid prototyping is simply that one cannot separate decisions regarding the design of computer graphics from the instructional and learning contexts within which they will be used. All the questions related to the appropriateness of graphics, the type of graphics, and the nature of the graphical elements must be considered and answered within the context of the overall instructional design. The more aware designers are about their instructional design philosophy, the better equipped they will be to make appropriate decisions. There are some obvious advantages and opportunities when computer tools are mixed with rapid prototyping techniques. CBI designers are encouraged to use rapid prototyping to test instructional designs that use a variety of graphical techniques. In addition, the principles of rapid prototyping should be used as the basis for understanding the remaining ideas presented in this chapter. Rankin (1989) has proposed an illustration design model founded on a mixture of principles related to formative evaluation and rapid prototyping procedures.

SOME GENERAL GRAPHIC PRINCIPLES OF SOFTWARE DESIGN FOR COMPUTER-BASED INSTRUCTION

The next time you take an automobile trip on a major highway, try to imagine the amount of planning that went into all aspects of the roadway's design. Give special attention to road signs. Reflect on their purpose and importance. Considering all of the other demands on a driver, signs serve an extremely important function. When designed well, they provide critical information to the driver in plenty of time to make decisions, but not so far in advance that the information is not yet relevant. When designed well, they are taken for granted, allowing the driver to devote attention to more important matters, such as traffic patterns. When designed poorly, drivers get lost, miss exits, waste time, and get very frustrated. Worse yet, they may have accidents.

Interacting with computer software, instructional or otherwise, is similar to taking a journey on a highway. It can be easy to get confused and disoriented. Well-intentioned software designers, like their highway engineer counterparts, sometimes try to map out the course that either *will* or *can* be followed through the use of software “signs” and “markers” that convey important information about “where you are” and “how to get where you want to go." It is easy to forget that the end user will not be a computer expert, nor perhaps even interested in computers, but someone who merely wants to use the software for a personal need or interest.

The use of software should be clear and straightforward, but this is easier said than done. In one sense, computer software is simply one more thing for the user to deal with in an already complex world. Interacting with computer software is very similar to interacting with other things in the world, such as automobiles, television sets, microwave ovens, and even doors and water faucets. It is arguably even more important to design *instructional* software with a clear and easy-to-understand **interface** (that which connects the user with the system) than other kinds of software, since the purpose of instructional software is to teach (or to help someone learn) about some content or domain. Time spent just figuring out how to use the software will obviously distract and detract from the instructional value of the software.

Fortunately, work from several fields of inquiry offers guidance to software designers, especially to those of us who design instructional software. Some of these fields go by the names of software engineering, cognitive engineering, ergonomics, human factors, and computer/human interface design (Box 7.2 reviews one of the more well-written and entertaining books on the subject). The principles of rapid prototyping are very compatible with these areas, helping to ensure that the software, in final form, will be easy to understand and use. A few of the most pertinent principles and issues related to the design of CBI are discussed next.

Box 7.2

The Psychology of Everyday Things

Does your VCR still blink "0:00" because you never figured out how to set the time? Have shower controls in hotel rooms ever baffled you? Did you ever buy a major appliance because you were seduced by the myriad of features and options advertised, only to discover that it was practically impossible to actually *use* the features once you bought and installed it? Have you ever become hopelessly lost while going through a hypertext or hypermedia stack? If any of the experiences sound even vaguely familiar, then you need to read *The Psychology of Everyday Things* (POET) by Donald Norman (1988).

POET is about the application of cognitive psychology to the design of human/machine interfaces. In fact, Norman has long been an advocate of something called "cognitive engineering," a label which could be used easily here as a two-word summary of POET. POET has become a kind of cult reading for engineers and industrial designers.

Interestingly, Norman has attained pseudo-celebrity status because of POET and is frequently called upon by national network news shows to comment on design issues in the workplace (POET has also been reviewed in the popular media, such as Time magazine[2]). Indeed, many of the examples Norman cites of bad design frequently occur in business and industry (such as nuclear power plants) — not a real comforting thought. Fortunately, most of POET is optimistic in its attitude that good design can become the rule and not the exception by following some of its relatively simple principles and recommendations. However, the book is not written for other psychologists, but for designers, engineers, and especially users. Norman uses everyday and often humorous examples of where even well-intentioned design goes bad and how psychology and its related fields (such as human factors) can be brought in to help.

Much of the design advice that Norman gives can be summarized by the principles of *visibility*, *mapping*, and *feedback*. Briefly put, the key parts and actions related to successfully using an object or completing a task should be made clearly visible. Controls and functions should follow natural mapping techniques (such as a kitchen's stove controls precisely mirroring the layout of the burners). This type of design effectively puts "knowledge into the world" and helps to reduce the user's reliance on memory, past experiences, or the operator's manual (Norman frequently suggests that an obvious indicator of bad design is when a simple object, such as a water faucet, needs written directions on how to use it). Feedback is essential in all stages of the action sequence. In particular, Norman discusses the role of feedback in narrowing the "gulfs of execution and evaluation" for the user. These gulfs refer to the separation between mental and physical states. The gulf of execution is the difference between one's goals (intentions) and the allowable actions (e.g. wanting to make a slide projector go back to the previous slide and actually being able to do it). The gulf of evaluation refers the amount of effort a user must exert to figure out if one's expectations and intentions have actually been met (e.g. being on hold while trying to connect to someone via an automatic phone routing system and wondering if you've been cut off). Obviously, the best designs have very small gulfs of execution and evaluation.

Norman suggests that designers practice the principles of visibility, mapping, and feedback within the context of a good *conceptual model*, combined with appropriate application of *constraints* and *affordances*. The idea of a conceptual model comes from research on mental models (which is well explained in the book). A conceptual model is a model chosen by a designer or engineer which is meant to convey the meaning of a system in a way appropriate for a user who probably has no idea of how the system really works and probably doesn't even care. Analogies and metaphors can often be good conceptual models, such as suggesting that working with a computer is like an office desktop. Constraints and affordances describe ways to naturally limit the range of possible actions to ensure a person's success when using the object or system. An example of a constraint would be a car designed to prevent locking the keys inside by forcing a person to use the ignition key to lock the door from the outside. Affordances are natural uses of objects — buttons are for pushing, handles are for pulling, knobs are for turning, etc. Often times bad design results from allowing too many functions to be handled by too few controls (high-tech wrist watches and telephones are common culprits) or by using an object in a way which is counter to its natural intent (people pushing a door meant to be pulled is a sure sign of bad design).

What does this book have to do with educational technology or instructional design? While at first glance it may appear that POET is most relevant to educational media specialists in their efforts to help educators use often bewildering media equipment in their teaching, the book is really about how to design any complex system so that it becomes understood and usable. Although POET is specifically about the design of everyday systems, such as the telephone, a car's dashboard, or the kitchen stove, the concepts and principles also easily extend to the design of *instructional* systems. One of the most obvious applications would be in the design of instructional computer simulations. The principles of POET are relevant so long as the intent is to design an environment where a learner will be *interacting* with a system, whether that system be a kitchen stove or physics. Such interactive learning environments are, of course, frequent and common in educational media, such as computer-based instruction (CBI). On the other hand, Norman's principles are less relevant if one's instructional design is largely based on presenting explanations, such as in linear, non-interactive tutorials. The other most obvious application is the design of the software's user interface, which corresponds to the CBI principles of frame and procedural protocol.

POET is also useful as a good introduction to many concepts and principles from cognitive psychology which Norman which goes on to apply to the design of everyday objects. Since instructional designers are expected to translate educational psychology theory into educational practice, POET illustrates how to go about this application, albeit in non-instructional systems. Fortunately, it is not difficult to see how the examples in POET can be used as analogous to instructional systems, so the book's many important lessons can be very relevant to instructional designers.

A few of the psychological concepts covered in POET include the nature of memory (including connectionism and parallel distributed processing) and mental models. Norman's accounts and explanations of these concepts are written in nontechnical language, again, making the book very readable to general audiences. POET includes some other important insights for educational technologists, such as the psychology of making errors (distinguished in POET as *slips* versus *mistakes*). Closely related is the phenomenon of *learned helplessness* where someone begins to falsely blame themselves during an activity which, in turn, frequently leads to the self-fulfilling expectation that success will never be possible. Norman also describes two “deadly temptations" for designers which are directly applicable to instructional designers — *creeping featurism* and the *worshiping of false images*. These refer to the temptation by designers to include features and options merely because they are technically possible, not because they are necessarily relevant or useful to the task. This, coupled with the tendency of consumers to buy products on the basis of a "the more features the better" mentality can lead to self-perpetuating cycles of bad design. The rampant use of sound and graphics in educational media closely parallels these two temptations.

The book is also written in an interesting way. The book's main narrative text is sprinkled with italicized, editorialized accounts and examples, largely from Norman's own personal and professional experiences. This style gives the reader the sense that Norman is giving a talk on stage while walking back and forth between two microphones (kind of like "Mr. Right Brain and Mr. Left Brain"). In speaking with others who have read the book, some like this style, and others are either annoyed or confused by it. Also, I found the format of the book's headings and subheadings to be very ambiguous and confusing, making it difficult to follow the outline of the book (a strategy I like to use when reading a book). Ironically, this violates the some of the very design issues and recommendations that Norman is advocating.

Throughout the book, Norman makes one sarcastic remark frequently as his commentary to bad design — "it probably won a prize." Norman has observed that aesthetics, not usability, is usually the overriding consideration in design. For this reason, the book's last two sentences are meant as advice to we, the consumers of everyday objects: "Give mental prizes to those who practice good design: send flowers. Jeer those who don't: send weeds" (p. 217). But this advice also provides a fitting testimonial to Norman's attitude to design in general and POET in particular.

Screen Design

Unlike many other media (such as print), the computer screen acts as the only interface, or "doorway," to the rest of the software. Screen design continually walks a fine line between incompleteness and excessiveness. Despite the attempt to offer a science of screen design, much remains an art form, depending on the skill, experience, and insight of the designer (see related discussions by Falio & DeBloois, 1988; Grabinger, 1984; Heines, 1984; and Hannafin & Hooper, 1989). Again, rapid prototyping techniques offer some of the best tools to design and test the screen design of materials because of the need to consider users' reactions and interpretations.

The design of any one computer frame is defined as **frame protocol**. The design of two or more related frames, as well as how two or more lesson parts are related or bridged, is defined as **procedural protocol**.

Frame Protocol and Functional Zones

As previously stated, at any one given time, the only interface between the user, the software, and the computer is the particular screen that the user is viewing. All information related to the software's use must somehow be conveyed in this limited space. It is common for computer displays to be as small as about 35 square inches. Obviously, this space becomes “prime real estate,” and future development must be carefully considered. The user may become confused and “lost” if too little information is provided and overwhelmed if too much information is provided. Hannafin and Peck (1988) describe three typical types of instructional computer frames: instructional, question, and transitional.

Instructional frames present instructional material or content to the student. Question frames provide any assortment of interactive situations in which students practice with the material at whatever level is appropriate (such as recall, recognition, or application). Transitional frames frequently act as bridges between major parts of the instruction. Typical examples are title frames (which carry “you are here” messages) and feedback frames (which help to bridge lesson information with the student's understanding, as demonstrated in practice frames). Traditional examples of instructional, practice, and transitional frames are illustrated in Figure 7.5. It is not necessary for these frame types to remain mutually exclusive.

A fundamental principle of frame design is **step size**, which is the amount of information presented in any one frame. Step size is an important consideration, given the tendency for instructional software to teach with presentation techniques such as explanations (although some would argue that presenting long, involved explanations is usually not considered a very appropriate use of the computer and is probably better suited for print-based materials). The issue becomes how much information to present on any one screen. Some research shows that students prefer “lean” presentations, or presentations that explain using abbreviated sentences and paragraphs, which may teach as well as traditional narratives (Morrison, Ross, & O'Dell, 1988; Morrison, Ross, Schultz, & O'Dell, 1989). Too much information on any one screen will require small font sizes and can be difficult and annoying to read. Too little information on one screen will require a string of many screens, making it difficult to follow and reducing the lesson's continuity.

Frame protocol is “the consistent designation of various zones of a frame for specific uses” (Hannafin & Peck, 1988, p. 175). Although it is acceptable for the frame protocol to change within a lesson, in general frame protocol should remain consistent for related frames within any one lesson section. It is usually helpful to consider a frame as a collection of two or more zones in which certain types of information will consistently be presented. Instructional frames, for example, typically include header, informational, and directional zones, such as that shown in Figure 7.6. These zones need to be effectively communicated to the student, either by familiarization through use or by deliberately and overtly calling the user's attention to them. Once this has been accomplished, the user should be better able to focus on the instructional message, since other information, such as navigational aids and directions, becomes secondary. Although the types of zones and their formats can vary widely for instructional, question, and transitional frames, it is crucial that the rules of consistency be enforced.

Instructional frame

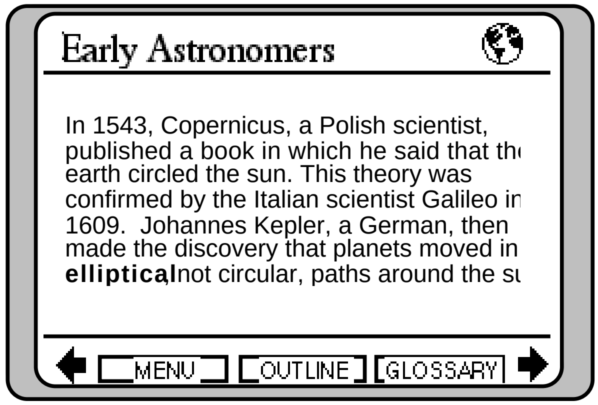

Practice frame

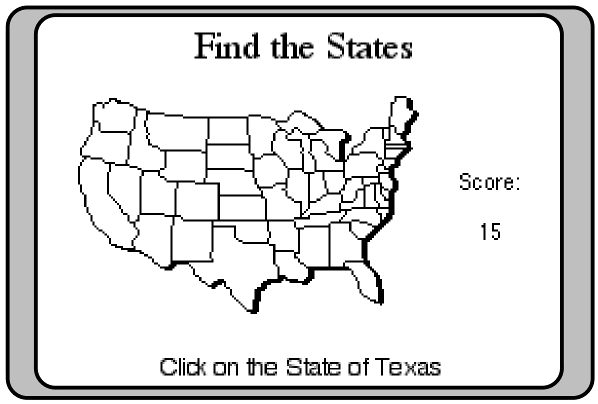

Transitional frame

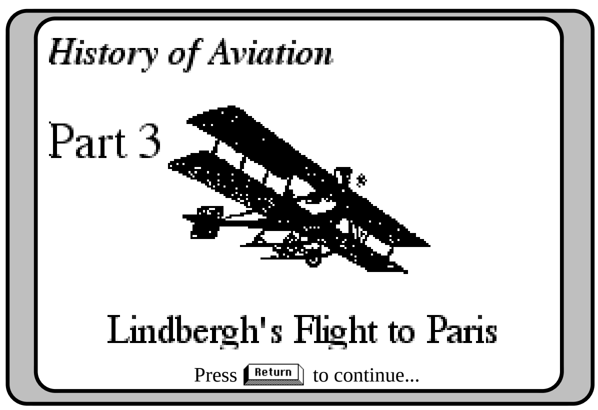

FIGURE 7.5
Examples of instructional, practice, and transitional frames in CBI.

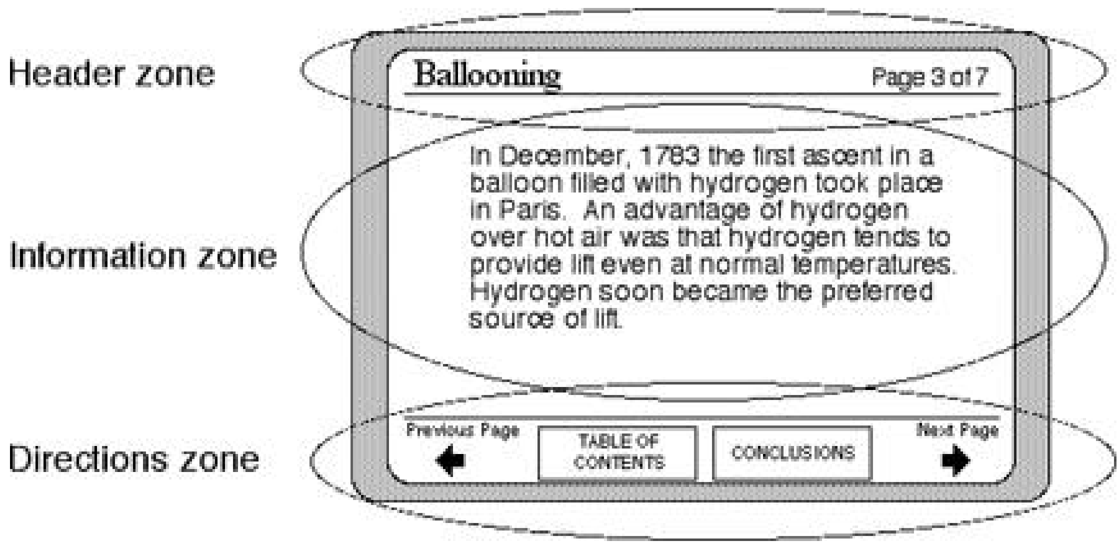

FIGURE 7.6
Examples of typical functional zones within a CBI frame. A "header zone" provides information about the relationship of the specific screen as it relates to the rest of the lesson. The "information zone" presents lesson content. The "directions zone" presents directions to the learner of what actions are required or available at this point.

Some screen zones can be further divided into subzones, such as an informational zone containing both pictures and words. Similar to print-based media, the zone can be split horizontally or vertically to present both the graphic and text, as shown in Figure 7.7. Given the interactive nature of computers, further information can be added or suppressed, via either program or student control. Additional information about the graphic or key words, such as labels, definitions, and examples, can be accessed through the use of **pop-up** windows available on command.

A recent phenomenon is software's heavy reliance on the use of graphical symbols (icons) instead of words for communicating software functions and options to users. There are obvious advantages to using icons. Icons don't take up much space and can be placed off to the side or periphery of the screen. Also, the graphical nature of icons can make it easy for a user to distinguish screen text from iconically based screen directions, hence reducing screen complexity and distractions. However, an icon's meaning must be absolutely clear to the user; often designers use very arbitrary symbols to convey complicated messages. Graphical symbols should be chosen very carefully so that even a first-time user recognizes the meaning at a glance. If no graphic can be found to satisfy this criterion, then a written label (with or without the icon) will be necessary. Designers often provide help options, which explain the function of screen symbols or options. While this seems like a good idea, help options are typically at least one level removed from the interface. Users may not even know that help is available. Of course, any icon that requires even a short paragraph of explanation is probably a clear indication that the graphic is a poor choice and should be changed.

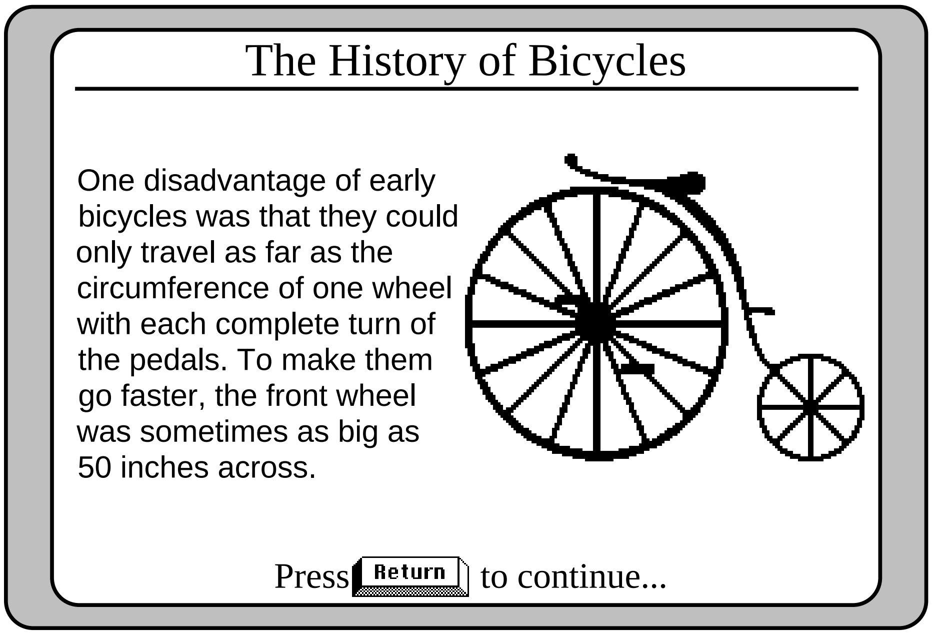

FIGURE 7.7
A frame protocol in which the "information zone" is split vertically to include both text and graphics.

Distribution of Emphasis

Whether designing instructional, question, or transitional frames, there will be a constant struggle to make the most important information jump out at the user, while having less important information remain unobtrusive, yet supportive. Therefore, there will always be a strong need to effectively direct a student's attention to the most relevant and salient information in the display. Other kinds of secondary information — for example, screen buttons for options such as returning to a main menu, or available instructional aids such as glossaries — should be, at best, in the periphery. This is especially true when graphics are used. The concept of **distribution of emphasis** refers to designing a frame so that a user is more likely to attend to the most important information and less likely to dwell on or be distracted by other information. Two approaches are typically used to increase a user's awareness of key information: cosmetic-based and information-based amplification techniques. Sometimes, as in advertising, these techniques are used in reverse to make critical information (such as the finance rate) as difficult to notice as possible, while complying with federal or state laws that require such information to be disclosed, such as that illustrated in Figure 7.8.

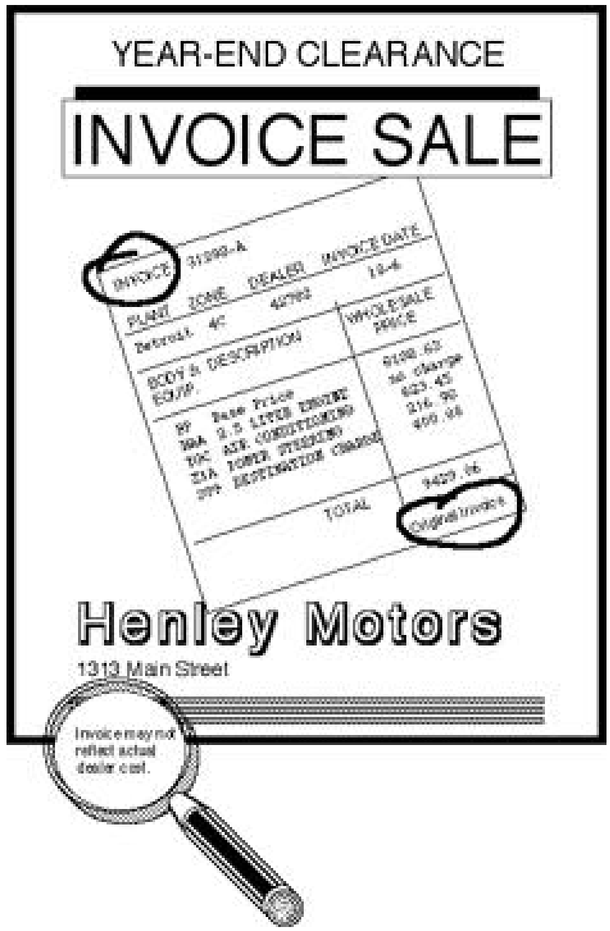

FIGURE 7.8
Does this look familiar? This illustration is based on actual automobile advertisements in local newspapers. Notice the interesting relationship between the largest and smallest typed words. The ad tries to capture your attention to the "invoice sale" in progress at this particular dealership. Most people think this means that the dealer is selling the car at their cost from the manufacturer, yet the small print indicates that "invoice may not reflect actual dealer cost." What, therefore, does "invoice" mean? Absolutely nothing.

Cosmetic-Based Amplification. Consider the experience of driving down the business district of a town where all the shops are competing for your attention. Some stores seem to jump out at us — we can't help but notice them. Others (usually the ones we are looking for) seem to hide away as if camouflaged. Storefronts that capture our attention with creative signs and blinking lights, whether we like it or not, are practicing their own form of cosmetic-based amplification.

In computer software, cosmetic-based amplification techniques use surface-level graphical features to attract and direct attention to specific screen parts or text. Although a multitude of techniques can be used, each intends to provide contrast between various frame parts and subparts through obvious graphical cues. Such overt cues are designed to make some information stand out and to make other information fade into the background. Graphical features, such as different fonts, different font sizes, different font styles (e.g., boldface, italics, etc.) are most commonly used to highlight key text, such as important vocabulary or important phrases. Functional zones can be highlighted with borders. Animation, such as moving words and other objects, is also a favorite cosmetic amplification technique. Other techniques include using blinking lights, words, and objects designed to capture the user's fleeting attention.

The intention of all cosmetic-based approaches is simply to promote contrast between various screen information in order to take advantage of a person's natural tendency to notice that which is different from the rest. However, if cosmetic techniques are used excessively, users will become either numb to the barrage of displays or so distracted that they may choose to ignore much of the screen information. The best software uses cosmetic-based amplification techniques carefully and prudently. Again, the rule of consistency is important — use the same types of techniques to convey the same sorts of information throughout the software. Consider the example of an on-line glossary. Software should always have the user perform the same action (e.g., click a button or keystroke) to trigger access to the glossary. The software should also use the same text format to show which words in a given paragraph are included in the glossary, such as boldface or color. Another common example is user directions, which should always be in the same screen location and be easily discriminated or separated from the rest of the information on the screen. Users will constantly be wondering “what should I do next?” or “what are my options at this point?” and will depend on a clear format to get this information.

Information-Based Amplification. In contrast to cosmetic-based amplification techniques, this class of amplification techniques uses learning strategies to amplify the most important or critical information (Hannafin & Peck, 1988). Repetition is probably the simplest technique. Repeating a key point or idea periodically throughout the lesson will help cue the user to the idea's importance or centrality. Other techniques include orienting strategies (similar to previews of what to pay attention to next) and review (summarizing the key ideas at the end of the lesson part). The old adage for preparing a speech — “Tell 'em what you are going to tell 'em, tell 'em, then tell 'em what you told 'em” — is an example of a simple repetitive strategy for making the message of the speech obvious and clear.

Information-based amplification techniques can also include interactive strategies, such as questioning, to force users to focus on the most relevant information.

Any technique that uses redundant displays of the information, though in different form, constitutes information-based amplification. An example would be a picture that repeats or duplicates textual information. All of the pictures types — representational, analogical, and arbitrary — can be used in this way. For example, information in a passage of text citing highway fatality statistics over the past five years can also be represented in a pie chart or line graph. A lesson section devoted to Abraham Lincoln's Emancipation Proclamation can be accompanied by his portrait to amplify his role in authoring the document. An arbitrary graphic can be used to show the relationship between related concepts, such as how branches of the U.S. government interact, in order to supplement and complement a given text.

Procedural Protocol

Whereas it is relatively easy to design any one frame adequately, it is much more difficult to design a series of frames in such a way that they contain overt cues to their relationship to the program as a whole. Hannafin and Peck (1988) define procedural protocol as “the consistent use of conventions for lesson procedures, obtaining student responses, indicating the availability of lesson options, and prescribing other features that affect lesson use” (p. 178). In other words, procedural protocol provides all of the **navigational aids** a user will need to successfully complete the software.

Consider, for example, how a book's thickness, size, and weight possess information about it's overall length and scope. We tend to take for granted how a reader responds and reacts to a book's tactile cues. Any one page has an implied relationship to the beginning or end of the book. The idea of turning pages is so intuitive that no user instructions are necessary. In this way, a book can convey a variety of secondary cues to the reader. On the other hand, some navigational cues require overt attention to be of value to the user, such as how to read or use a table of contents, titles, subtitles, and indices. Unlike a book, any given single frame of computer-based instruction acts as the user's total interface with the rest of the program. Procedures of how to navigate through the lesson and participate in interactive strategies, as well as take advantage of other options, must be clearly communicated to the user.

The study of human factors, also known as cognitive engineering, corresponds closely with procedural protocol. As discussed in Box 7.2, Norman (1988) suggests that the three principles of visibility, mapping, and feedback must be carefully considered in the design of any interface between humans and an object in the world.

Software should make all relevant options and functions visible and plain. When a multitude of options are available, the functions should be stratified, so that only the most relevant options are in the user's perceptual view. But there should be “doorways” to second or third layers of options. Software should organize features in meaningful ways so that users do not need to deal with all options and decisions at the same time. For example, the option to increase or decrease the size of printed output could be nested within a “print” function, which, in turn, might be further nested within a group of “file” functions. This may be less an issue in structured CBI packages where users are frequently led down only one learning path. However, in unstructured packages, such as hypertext stacks, the risk of disorientation is great, despite the flexibility of the software.

Users should be given **conceptual maps**, such as outlines, to help guide their explorations and decision-making. Users should be completely confident about where they are in a software package, where they have been, and where they may be going. The range of allowable actions should constantly be evaluated against the range of goals a user may have at any given moment. Once a user defines a goal, the amount of time and effort that must be expended in order to execute an action to meet the goal should be minimized. (Recall from Box 7.2 that Norman [1988] refers to this difference between a user's intentions and the system's allowable actions as the gulf of execution.)

Computer specialists have come to tolerate cryptic and abstract interactive strategies, such as those involving a keyboard. Other input devices, such as the mouse, light pens, touch screens, and voice recognition, are beginning to reduce the level of abstraction to which computer systems so far have been prone. A goal of “putting the red ball in the blue box” is trivial in the real world, yet can be an exercise in frustration for a similar simulated action on a computer screen. The principle of mapping is simply the degree of relationship between two things, such as a screen button of a right arrow and the function of going to the next page. Mapping is an important tool in helping to be sure users can quickly complete an action sequence, whether their goals are simple (i.e., go to the next page) or complicated (i.e., mix chemical A and chemical B to produce solution C). The range of available input devices should be considered (including, but not limited to, the keyboard). Computer systems should use “natural mapping” in their designs where the action sequence is concrete and mirrors how actions would be executed in the real world. Natural mappings are important, whether one is designing an aircraft instrument panel (Hicken, 1991) or a computer “book,” where one clicks on the right side of the screen to go to the next page or on the left side to go to the previous page.

The widespread availability of the computer mouse has made it easier to design simple user interfaces. Though not as concrete as some input devices, such as touch screens, even novice users seem able to master the “point and click” technique with only minimal experience. The most common mouse strategy is for the user to interact with the system through screen buttons. However, the mere use of a mouse and buttons does not guarantee that an appropriate procedural protocol has been designed.

Buttons seem to be the most effective as interfaces between a lesson's procedural protocol and the user when the button is based on a concrete concept (see chapter 2). The advantage of using concrete concepts as navigational aids (whether representational or analogical) is that they are memorable and can easily be depicted in graphical form, such as screen icons (as introduced in the last section). Whenever possible, use simple and intuitive representational graphics in designing screen buttons, such as a picture of a printer to convey print functions, so that users know at a glance the function of the button. Use an analogy if no direct graphical representation of the concept exists, such as a picture of a tree to denote the “branches” of an outline or menu, or a picture of a “football coach” to denote how to get help, advice, or other coaching. Computer systems frequently use concrete analogies, such as “file folders” for subdirectories, to convey abstract ideas in meaningful ways to novice users.

Feedback is an essential element of procedural protocol. Users need simple and direct feedback as they interact with the software to let them know if an intended action has been completed. Once users complete some action sequence, the system must tell the users if they have been successful. (As discussed in Box 7.2, Norman [1988] calls this gap between what a user intended to do and knowledge about what actually happened as the gulf of evaluation.)

Software should take into account a variety of learners with a variety of needs and situational conditions. Software should reward users for an exploratory attitude. Likewise, software should not penalize users for making mistakes in their attempts to navigate through the software. Navigational mistakes, when made, should be easily detected and remedied by the users.

Finally, when all other design efforts fail, the last resort is to standardize a function, such as the example of using a special key or button to back out of a function. Again, consistency is very important. Students will usually begin to pick up on well-designed procedural protocols just by use. Once a designer determines the placement and function of various screen locations (functional zones), care should be taken not to inadvertently change these and any deliberate changes should be carefully and cautiously made. Users need to be able to accurately predict what they are able to do and how to do it all along the way.

The concept of transparency, as discussed in chapter 1, is an important benchmark in evaluating both frame and procedural protocol. The better frame and procedure protocol are designed, the more likely the user will not notice them. Instead, the user can focus on the lesson's ideas or activities. This creates a type of seamless software design, such that there is no obvious distinction between using the software and knowing *how* to use the software.

Some Basic Principles of Graphic Design

Despite the resistance here to entering headlong into the world of graphic design, there are some generally accepted issues and concepts related to visual layout that are relevant at this point. Even the most specific design principles can be traced to one of four broad categories of visual design and layout: simplicity, unity, emphasis, and balance. In a sense, these principles aptly summarize much of the previous discussion of frame and procedural protocol. All four categories can be supported and described on the basis of human perception and information processing discussed in chapter 4. For example, visual layout should take into account a person's abilities and limitations in selectively perceiving the most important information in a given display. All information coming from the environment will be competing for a person's limited attention. Therefore, a display should be designed to maximize the chances that a person will notice the most relevant quickly; it should also be able to sustain user attention over a period of time. When combined with rapid prototyping procedures and a lifelong hobby of “software watching,” these principles can help guide beginners in designing effective screens.

An easy way to understand the importance of these four principles is simply to consider the effect of their opposites — complexity, disorganization, sameness, and imbalance. For example, it is difficult for most novice designers to resist the temptation to fill a given display area with as much information as possible, as though the display were like a refrigerator shelf. The simpler the design, the easier it will be for a person to quickly scan, interpret, and extract meaning from a display. The most important information should be readily discernible. This does not necessarily mean that the user is released from the responsibility of studying a frame of instruction. However, if the user is expected to take sufficient time reading a block of text, the perceived demands of the task, such as reading, should be obvious to the user, both in terms of time and effort. See Pettersson (1989) for more specific information about the graphic design of instructional visuals.

Color and Realism as Instructional Variables

There is something intuitively appealing about using color in the design of instructional materials. It is tempting to believe that learning must be a natural consequence of the strong visual sensation often provoked by color. (See Footnote 3) Therefore, it seems obvious that merely adding color to a given display should increase learning. Given the ubiquitous nature of color in the natural world, it is easy to believe that if one wants to learn about the world in some way, then color should play a role. Educators often criticize instructional materials solely because they lack color. (See Footnote 4)

Once again, what appears to be an obvious truth is complicated and clouded by the fact that color is simply one of many variables that must be considered in instructional design. For simplicity, we will discuss the use of color in instructional displays in the context of the two families of instructional functions introduced in chapter 2 — affective and cognitive. This should help us to better understand whether or not evidence, speculation, and intuition agree regarding the design of color. Simply understanding the differences between affective and cognitive functions of color and realism can be a monumental first step for designers and can help prevent creating well-intentioned displays that interfere with learning. The discussion presented here is offered as a doorway to understanding the issues that are unquestionably more complicated than their relationship to only these two functions may suggest.

The appeal of color graphics, by definition, is directly associated with affective or motivational considerations. However, most of the available research is concerned with using color in direct instruction, as opposed to motivation. Even here, research on using color for cognitive functions has largely been inconsistent and inconclusive (see Dwyer & Lamberski, 1982-83, for a review). In general, color, in and of itself, does not seem to be a critical variable for instructional tasks. At best, color serves only a secondary instructional purpose, such as cueing or directing a learner's attention to some critical feature in a display.

Color can be an effective amplification technique by helping to bring important information to prominence. Even here, it is not color per se that makes a difference, but the potential contrast that color provides (Goldsmith, 1987).

Closely related to the use of color is realism, or the degree to which a pictorial representation resembles an object in the world (Dwyer, 1978). The concept of realism exists on a continuum ranging from representations that cannot be distinguished from their real-world counterparts to the most abstract representations of objects, such as printed words. Consider a picture of a goldfish in an aquarium. The most realistic representation would confuse a person into believing that the picture *really is* a fish. All of the visual cues associated with a goldfish, such as color, texture, depth, motion, and reflected light shimmering on the water, would offer no visually perceptible differences between the representation and a real fish. The truest representation would need a high-resolution, animated holographic 3-D image. Less realistic (because some of the visual cues, such as motion and depth, would be missing) would be a high-resolution color photograph. Someone might mistake the photograph for a real fish at first, but would soon realize that it is only a picture.

As with color, it is tempting to say that effective instruction should contain a high degree of realism, without fully understanding or recognizing the intent or function of the visual. Visuals used for cosmetic or motivational functions will have entirely different design assumptions than visuals designed to teach or instruct. Recall that by instructional, we mean that the visual's purpose is usually to communicate a fact or intellectual skill (such as a concept or principle, or aid in a person's problem solving). As discussed in chapter 5, Dwyer's (1978, 1987) research suggests that visuals designed for a cognitive task that contain either too little or too much realistic detail adversely affect learning, especially when learners do not have control over the pacing of the visual presentation, such as video. Dwyer's research has shown that students may have difficulty identifying and attending to relevant information in a highly realistic visual, such as a color photograph. For example, realistic details (including color) can interfere with a student's ability to recognize and understand critical features of a visual.

However, when the intent is motivational in nature, it can be argued that color and realism are important attributes to consider. The next time you go to a bookstore, reflect on your browsing patterns. See if you are more likely to browse longer through a book that contains an assortment of interesting visuals, especially those with color photographs. Of course, browsing behavior is not instructional in nature. You are not trying to read the book for meaning when browsing. On the other hand, you definitely will not learn anything from the book *unless* you first open it up and spend time with it. Therefore, it can be suggested that color and realism might serve to affect one's choice to engage in a learning behavior, and subsequently, to choose to persist in the activity. This may be true even though accompanying visuals offer no direct instructional value. Some research has shown that student preferences and expectations for visuals that differ in the amount of realism (i.e., video vs. computer graphics) may influence the effectiveness of the technology more than direct instructional uses of these visual characteristics (Acker & Klein, 1986).

There is no reason not to use color and realism when they are used solely with the intent to increase the motivational appeal of instructional materials and when, of course, they do not undermine the effectiveness of other instructional variables. However, even this relatively bland recommendation is made cautiously, as no substantive research is available to support the argument that color and realism directly increase the motivational appeal of materials (see Surber & Leeder, 1988).

Given the decision to go ahead and use color in instructional materials, we still are left with all of the graphic design issues of how to use color to create appealing and motivating materials. Again, this book does not presume to teach graphic design, although it is, of course, related to instructional design — poor *graphic* design of materials can easily undermine what would be otherwise appropriate *instructional* design. As a beginning guide, Box 7.3 offers some good advice on the use of color. Here are some general design principles related to color:

1. Use color as an effective attention-gaining device.
2. Use color to show contrast, in order to direct or focus attention.
3. Use color to show relationships between screen information, such as using the same color for labels and the screen objects to which they correspond.
4. Use color to increase motivation, interest, and perseverance, but be careful to avoid distraction effects.

Box 7.3

Color Use Principles

Color is probably best considered a secondary graphic element in the design of instructional materials. That is, the effect of color on learning has not been shown to be particularly potent. This simply means that there should be little expectation that color, in and of itself, will lead to greater levels of learning. In fact, the greatest role of color may simply be in supporting other instructional elements, such as a cosmetic-based amplification technique to help gain a learner's attention to important screen information. Color use should not be avoided, but nor should it be used indiscriminately, as poor color choices have a high potential for distraction. Although the research is vague, there is some reason to believe that color may play an important role when designing materials with affective or emotional learning goals.

The following is a synthesis of the literature related to the use of color in graphic design, including, but not limited to instructional applications. This set of design principles was created by Evelyn Wells.

General Principles

1. *Do not use color indiscriminately*. Color can enhance the appearance and function of your project, but only if it is used properly. Incorrect color use can significantly detract from the effectiveness of your project.

2. *Consider color as an aesthetic and cognitive design tool during every stage of your design*. If you reduce color to only a cosmetic afterthought, you will miss out on all the advantages it could give you.

3. *Limit the use of different colors*. Exercise simplicity, clarity and consistency. Make sure you have a reason for each color you use, and how you use them.

4. *Do not rely on color alone*. Color is best at redundantly emphasizing information. Other design elements (text, shape, layout, fonts, etc.) should carry the bulk of the information, and work in tandem with color. Your project should still be usable if converted to a monochrome format.

5. *Try to work in the same conditions in which your project will be used in final form*. Many factors can change the appearance of colors. Even worse, the change is not uniform for all colors. Variations in ambient lighting (fluorescent, daylight, darkness) and media transfer (print, video, film) can have a drastic effect. Also be aware that screen layout dynamics will cause colored elements to change in relative size and position, which will affect their appearance a great deal.

6. *Experiment, experiment, experiment*. Do not underestimate the potential of color. Consider alternative color schemes and consult with potential users of your product and your colleagues. Examine other products (in your field, and in everyday life) and become aware of how color is used.

Cognitive Design Principles

1. *Group categorically related elements with the same color*. Emphasize relations between elements on the same screen, or on successive screens, by coloring the elements or their backgrounds with the same color. For example, in a hypermedia document, all words which link to the glossary could be colored blue, and all words which link to a video segment could be colored green. Do not use the same color for elements which are not related. Even if the elements are not viewed at the same time, they could be subconsciously linked.

2. *Use similar colors to denote relationships between elements*. Relationships, such as chapter and section hierarchies, can be represented by choosing colors that

systematically vary in a dimension such as hue or saturation. The degree of change (such as dark green for a chapter heading, medium green for a section, light green for a subsection) can indicate the strength of the relationship.

3. *Link color change to dynamic events*. Changing color (from green to yellow to orange to red, for example) can portray elapsing time or other critical levels. Changing color is also good for dynamic data visualization.

4. *When using colors for coding information, use a maximum of 5 +/- 2 colors*. Human memory for color is even less than the "magical number 7+/-2" accepted for memory of units such as digits and words. If your users must remember and recognize colors, use only 3-7, and make them as distinct and meaningful as possible.

5. *Use extremely bright and saturated colors only for special purposes*. These colors immediately draw attention, and should be used sparingly, and only for the very most important parts of the design. They are useful for error messages, urgent commands, key words, or some introductory information. For the rest of the content, give the most important elements the most contrast with the background.

6. *Use "temperature" of colors to indicate action levels or priorities*. Warm colors (red, orange, yellow) tend to advance from the image and imply action or a required response, while cool colors (purple, blue, green) tend to recede from the image and imply rest or background status.

7. *Use logic in choosing meaningful color schemes*. Color schemes based on the ROYGBIV spectrum are often naturally understandable by most users. However, you should take advantage of colors conventions which have cultural or application-specific meaning, such as using red for unpaid bills, or using pink for female children and blue for male children.

8. *Be aware of the social connotations of colors*. Colors carry meaning individually and in combinations. For instance, blue can mean masculinity, death, water, or coldness, and red, black and white are often associated with Nazi Germany. Choose colors with your audience in mind.

9. *Use the same colors for all aspects of your project*. Aid the users of your product by consistently using the same color scheme for all documentation: support materials, training, testing, and advertisement. Again, be aware of the color change when transferring across media.

Physiological Design Principles

1. *Do not use highly saturated, spectrally extreme colors simultaneously*. The focus of the eye changes according to the wavelength of the color, and spectrally extreme

colors cause frequent refocusing which may in turn cause visual fatigue and afterimages. Avoid juxtaposing colors such as red and blue or yellow and purple, or use them only in desaturated forms.

2. *Use blue for large background areas, but not for text, thin lines, or small shapes*. The eye is not sensitive to blue in the foveal (center) portion of the retina, where detailed vision occurs. This makes it hard to discriminate small blue shapes, especially for pure shades of blue. However, blue makes an ideal background color.

3. *Use red and green for central colors, but not for background areas or for small peripheral elements*. The eye is insensitive to color in the periphery, especially to pure reds and greens. If these colors must be used for elements in the periphery, use high contrast or blinking.

4. *Avoid adjacent colors which differ only in hue*. Because edges are mainly perceived through brightness gradients, adjacent colors should always differ in value as well as hue. Blue does not contribute to the perception of brightness, and edges created by a difference in blue only will appear especially indistinct.

5. *Consider the final viewing environment*. In general, use a dark background with light elements (text, etc.) for dark viewing conditions (slide presentations, etc.) and a light background with dark elements for light viewing conditions (paper, normal computer use). Contrast is the most important factor in text legibility.

6. *Increase the brightness of the display for older operators*. With age, the eyes loose much sensitivity. The overall brightness of the display should be increased, and color contrasts will need to be enhanced.

7. *Avoid single-color distinctions*. Color deficient vision (color-blindness), which occurs in about eight percent of the male Caucasian population, stems from the partial or complete dysfunction of either the red or green (and occasionally the blue) photoreceptors in the retina. Colors which vary only in their amounts of red or green will be hard to distinguish for color deficient viewers. Again, colors should vary in at least two of the three primary colors, and subtle differences should be used with caution.

FUNCTIONAL DESIGN RECOMMENDATIONS FOR INSTRUCTIONAL COMPUTER GRAPHICS

The final section of this chapter summarizes many of the issues discussed in this book related to the use of graphics in designing instructional materials. This section presents general design principles based on four of the five instructional applications of graphics introduced in chapter 2 — cosmetic, motivation, attention-gaining, and presentation. These principles are relevant to both static and animated graphics. The fifth application — practice — is the topic of the next chapter, which will address the design of highly interactive learning environments on the computer and especially the use of animated graphics as visual feedback. As discussed in chapter 2, these five applications are not mutually exclusive and frequently overlap.

There are a few general instructional graphic principles that apply to all materials. These are followed by principles specific to each of the four instructional applications.

1. *There are times when pictures can aid learning, times when pictures do not aid learning but do no harm, and times when pictures do not aid learning and are distractive.* This is the "first principle of instructional graphics," as presented and discussed in chapter 1. Its obvious message is repeated here as a reminder that graphics are not innocuous variables to be casually included in instructional design, but must be given careful thought and planning. The intent and outcome of graphics should be considered and evaluated throughout the instructional design process.

2. *Select the type of visual based on the needs of the learner, content, and the nature of the task.* The type of visual used (representational, analogical, arbitrary), as well as the instructional function it serves (cosmetic, motivation, attention-gaining, presentation, practice), should be selected and designed based on the interplay of three variables — learner, content, and task. It must be remembered that not all people learn in the same ways, nor do they all have the same interests, backgrounds, or experiences. Different content or domains (i.e., the material to be learned) demand different considerations when it comes to visuals. Also, the nature of the task, as defined by one or more strategies suggested by Gagné's events of instruction, and the delivery system (i.e., individual, small group, large group, distance learning, etc.) must be taken into account. All of these instructional variables, of which visuals only contribute to, must be congruent and consistent with one another.

3. *Graphics should not distract attention from the lesson goals or objectives.* The best guide to what should be achieved in the lesson is the lesson objectives as determined through instructional design procedures — whether those based on traditional ISD or rapid prototyping. Unless one knows what the goals are, there will be no way to know if the goals have been met.

4. *Graphics should be designed carefully to serve their appropriate function*. Graphics should be designed as an integral part of an instructional design, not as an

afterthought. Deliberate attention should be given to what type of graphic is chosen and the function it serves as it relates to the overall design. The component parts of a lesson, whether based on Gagné's events of instruction or another model, should complement and support each other. For example, once a deliberate decision is made to use a line graph to show the relationship between economic growth of two countries over a five-year period, design efforts should go toward accomplishing this goal. An effort should be made to resist using the graph for other reasons, such as cosmetic. When the graph has served its purpose, it should be removed and the lesson should proceed. It also goes without saying that graphics should not depict or promote cultural, ethnic, or gender stereotypes.

Cosmetic Graphics

5. *Be extremely cautious in the use of cosmetic graphics in the design of instructional materials.* Cosmetic graphics, by definition, do not carry any instructional value. The intent of cosmetic graphics is simply to make the materials more attractive and polished. There is inherently nothing wrong with this, and there is no reason why instructional materials should be drab and boring, but there is always the danger that frilly graphics may distract a learner's attention from the instructional message. The best advice is to design cosmetic graphics with a true motivational purpose. In other words, design graphics that at least have the primary intent of motivating students and that secondarily give the software a finished, polished, or commercial quality.

6. *Make design decisions related to the use of cosmetic graphics early in the process and include such graphics in the evaluation of the final materials.* One worst-case scenario is when cosmetic graphics are added *after* all the time and effort has been spent in designing, developing, and evaluating quality instruction, again solely to satisfy some estimate of commercial quality.” Not only is the risk high that these graphics may disrupt the effectiveness of the instruction, but the designer may never know of this influence since all evaluation procedures will have long since concluded. At the very least, summative evaluation should not occur until the materials are in final form — and final form means final form.

   Some common cosmetic graphics include the wide array of backgrounds, such as those that simulate pages in a book or note pad, television screens, or text etched on simulated marble slabs. Designers also commonly use “glitzy” transitions, such as one screen sliding in front of the previous. Rather than only consider these as cosmetic features, designers should seek to use these transitions in ways that help users understand the flow of information. For example, the screen transition of “slide left” (where the next frame is “slid” from the right of the screen toward the left edge of the screen, covering up the previous frame as it moves) could be used to simulate the natural mapping of “future pages come from the right.” Similarly, going back to a previous page should use “slide right.” If a special glossary screen is available, the transition effect “slide down” could be used when accessing it, only to use “slide up” when the user is finished with it. This helps build the concept that such extra information is “above,” waiting to be used, and is “put back” when finished. In this way, designers can use such simple effects for cosmetic, motivational, and instructional applications simultaneously.

Motivational Graphics

7. *Use graphics to increase motivation and interest, but be careful to avoid distraction effects.* Motivation is a popular reason or rationale that designers often cite to justify using graphics. As discussed in chapter 3, there are two principle types of motivation: extrinsic and intrinsic. Although there is nothing wrong with using graphics to increase the extrinsic appeal of the lesson through the use of pretty pictures, designers should seek to creatively use graphics to increase the *intrinsic* appeal of the lesson through the use of graphics (as per the next recommendation). A common application of motivational graphics, especially in instruction meant for children, is personified cartoon characters (such as “Mr. Tooth” explaining the importance of dental care). Of course, the use of graphics solely for motivation excludes their real power to communicate, to inform, and to provide feedback.

8. *Use graphics to present meaningful contexts for learning and to increase the intrinsic motivation of the learner.* Probably the best use of realistic graphics, such as photographs and especially video, is in triggering strong emotional and affective responses in people. Therefore, one of the best motivational uses of graphics may be in stimulating a learner's fantasy and imagination through powerful visual displays. Video sequences, such as a shuttle lift-off, the beating of civil rights demonstrators, or the spectacle of a passing tornado, can help create real reasons for a learner to participate and persevere in the mathematical, social, or scientific issues about to be introduced and discussed. Once again, the graphics are used for a specific purpose to complement and supplement the entire instructional system.

   One strategy that can overlap motivational, instructional, and even cosmetic functions of a given graphic is the use of photographic images of an event, including portraits, to complement the verbal accounts of the event. This use might be termed the “Life magazine effect” and refers to the intrigue that a photograph can stimulate while one is reading or listening to the event to which it is related. A photograph of Abraham Lincoln's face does much more than words to capture his essence in an article about America's Civil War. Again, designers walk a fine line between using the graphic to instill a sense of the emotional quality of how one person affects history and distracting the reader or viewer from the details of what the person actually did.

Attention-Gaining Graphics

9. *Graphics can be an effective attention-gaining device.* Surprisingly, little direct research is available on the use of graphics for attention-gaining and much of that is dated (see Dwyer, 1978). However, there seems to be consensus among developers that graphics can be used to draw students' attention to the instructional materials.

That said, the following principles related to attention should be used to guide the design of graphics for attention-gaining and attention-sustaining purposes.

10. *Attention is a highly selective and controllable process* (Fleming, 1987). As discussed in chapter 3, the limitations of selective perception and short-term memory require an individual to focus on only a limited amount of information at a time and block out other incoming stimuli from the environment. Attention-gaining graphics should be designed to pull learners back to the instruction in general and to a task in particular. The graphics (and the rest of the lesson components) will constantly be competing with other sources of information (such as other incoming stimuli from the environment and the person's prior experiences). The graphics should give students a reason to attend to the lesson information, either because of the graphic's visual appeal or the meaning that the graphic may hold for the learner. Similarly, a learner's expectations can strongly influence attention.

    There is an important relationship between attention-gaining and presentation principles associated with graphics. Even though a graphic may be used to present information, learners often do not know how or when to use the information contained in the graphic. This is a problem related much more to selective attention than how to present the information graphically. Certainly, students will not learn anything from a presentation graphic unless they first attend to it. One strategy to increase the chance that students will attend to a graphic is to provide students direct and overt directions to actively search for or use specific information in the visual. In this way, the strategy overlaps attention-gaining, presentation, and practice functions. Unlike other media, such as print-based, computers afford a variety of interactive strategies. Designers should take advantage of these.

11. *Attention is naturally drawn to what is novel or different.* An enemy to attention-gaining includes any monotonous stimuli. Screen after screen of text will soon make even the most motivated individuals lose their ability to focus on the message. Even presenting small changes occasionally will help attract and maintain attention. In CBI, transition screens between lesson parts can help to not only provide markers to students as they complete a lesson, but also help to break up even short series of presentations. Contrasting screen elements, such as animated objects, can also help attract attention.

    However, do not confuse this principle with those associated with graphics used to present or communicate information. Although a complex graphic may gain or attract a learner's attention, the learner subsequently may be quite unable to extract any meaning from the graphic. So, while the graphic may have been successful in gaining attention in general, excessive realism may distract student attention from the specific (or essential) information contained within the graphic.

Presentation Graphics

12. *Graphics should be congruent and relevant to the accompanying text, or distraction may result.* This is probably the most fundamental principle associated with using graphics in presenting information that is supported by the research on both static and animated visuals (see chapters 5 and 6). Using graphics as part of presentation strategies is the most direct use of a visual for instructional purposes. Presentation graphics should be carefully designed to convey only the information intended, and adding any further visual details should be avoided. The most salient and critical features of the graphic related to the information should be clearly distinguishable to the learner. The graphic may be used as the primary carrier of the lesson information or may be used to supplement the verbal information (whether textual or aural). Animated displays should be used to present information that changes over time. Such changes are usually operationalized through either the attribute of motion or trajectory (path of travel), or both.

13. *Students should be cued to process the information contained in a graphic in some overt way.* This principle complements the attention-gaining principle (#10) above. The most well-designed graphic will be totally useless if the learner does not first attend to the visual and then consciously use the information in the visual in some way. Most designers assume that the learner will be visually literate enough to know what to do with the visual, when this may not be the case. Part of the problem is that designers become so close to the instructional materials that they can lose touch with the learner's point of view. Rather than assume that a learner will instinctively know how to interpret the visual's information or leave it to chance, the lesson should be designed to have the learner interact with the visual in some way.

    Some interactive strategies may seem more related to practice than presentation, but this is not an important issue. A presentation strategy may simply be to rhetorically ask a learner to "look and see (from a graph) how many bushels of wheat were produced in America last year and compare it to 1960." A true interactive lesson strategy on the computer would force the learner to input the information. Interaction during the presentation of information would conform to learner guidance (Gagné's fifth event of information), since it is aiding the learner in attending to and selecting the information as it is presented. Practice strategies are more related to the rehearsal and application of information once it has been taught.

14. *Graphics are unnecessary when the text alone produces mastery.* Despite some of the advantages and appeal that the addition of external visuals may carry in instructional design, it should be remembered that a well-constructed verbal message, whether textual or aural, may sufficiently cue a learner to internally form appropriate mental images. This is the "master story-teller" phenomenon discussed elsewhere in this book. In addition, as the research indicates, the reliance on external visuals may decrease with age. Therefore, using pictures and other graphics may be less necessary with adults than with children. Of course, the graphics may still support the learning, although no additional effects on performance may result from

their use. As some of the research with animation has indicated, the graphics may be helping the encoding and retrieval processes (as evidenced by response time), even though no additional learning seems to be occurring (see chapter 6). Of course, such graphics may also provide some motivational incentive to students, even though some distinctive instructional value of the graphics may be questioned. Also, the instructional value of spontaneous internal imaging depends heavily on the context. Internal imaging may be sufficient in the comprehension and inference of a story, but may be totally insufficient in learning technical information, such as how to change a flat tire.

Practice

The use of graphics for attention-gaining and presentation purposes, though appropriate and practical, does not come close to the potential of computer-based applications of instructional graphics. Of the many strengths associated with the computer as an instructional medium, its interactive capabilities represent the richest and most exciting areas for instructional design and development. As yet, these capabilities remain largely unexplored and much of the potential remains untapped. In this book, the term “practice” is broadly defined as any strategy in which the learner interacts directly with the content or domain. The potential of graphics in practice strategies can be essentially reduced to their role as immediate visual feedback. The computer's ability to quickly process student input to provide moment-to-moment visual feedback (i.e., animation) extends beyond the utility of immediate feedback to questioning techniques and into the realm of artificial worlds where learners do not just study the material, but begin to live it. In this way, intrinsic motivation and practice become intertwined. The next chapter is devoted to this topic.

REVIEW

- Both formative evaluation and rapid prototyping imply a strong relationship between instructional design and instructional development.
- In rapid prototyping, design and development become intertwined, whereas design and development are usually considered separate processes in formative evaluation procedures from traditional ISD.
- Rapid prototyping procedures provide a very appropriate context for design decisions related to the type and nature of instructional computer graphics.
- The most effective frame and procedural protocols are unobtrusive and satisfy the principle of transparency.
- The purpose of cosmetic-based and information-based amplification techniques is to help the learner identify and distinguish the most important and salient information in a display.
- Effective procedural protocols act as interfaces between the user and the computer materials and help the user “navigate” in and around the materials.
- The relationship of color and realism to the design of effective instructional materials should be distinguished on the basis of affective and cognitive components.

- Color and realism seem to be more important as an affective consideration than as a cognitive, or instructional, consideration.
- Each of the design recommendations for the instructional graphic design of cosmetic, motivational, attention-gaining, and presentation applications should be applied and adapted to fit the specific instructional context.

NOTES

1. Interestingly, I see more similarities than differences between rapid prototyping and traditional interpretations of formative evaluation. I have only recently come to realize that rapid prototyping may be quite radical to many instructional designers who take a traditional instructional systems development (ISD) perspective. I think it is because I have largely been involved only in computer-based applications of instructional design since leaving my past role as a classroom teacher. As it will be stated frequently in this chapter, it is far easier to implement rapid prototyping procedures and philosophies with computers, given their “plasticity” and “modularity,” than with what are commonly referred to as “traditional” media, such as print-based, film, video, and photography.
2. Time magazine, July 4, 1988, pgs. 48-49.
3. Unless, of course, you are color blind. Complete color-blindness is very rare. Much more common is partial color blindness, which occurs much more frequently in men than women. People who are partially color blind, called **dichromats**, can only perceive two colors — red and green or blue and yellow — and the other colors that are a blend of the two.
4. Such as this book, I suppose. Due to a variety of reasons (not the least of which was cost), color was not an available resource for this book. This decision was made at the book's inception. On the other hand, color is not a critical design component for most of the concepts being discussed, although, I admit, it would have increased the motivational appeal of this text.

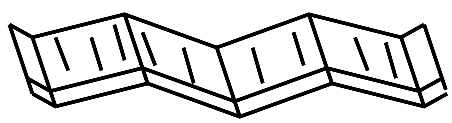

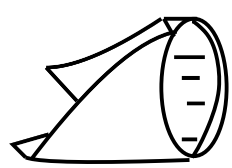

Here are some of other, more radical paper plane designs as referred to in Box 7.1. The top design is call a "flying wing" and when properly constructed is a superb glider. The paper plane below is based on the "ring wing" design by George Allison of the NASA Langley Research Center. This latter design is supposedly only half the weight of a conventional airplane, but with the same payload.

CHAPTER 8

Designing Highly Interactive Visual Learning Environments

OVERVIEW

This chapter provides recommendations for designing interactive learning environments. The type and nature of interactive strategies depend on the underlying learning philosophy. This chapter describes a philosophy of learning, called constructivism, that views learning as individual "constructions" of knowledge. This philosophy and its implications in education are compared to "instructivism," a term used to denote the other cognitive applications to instructional design considered up to this point. In constructivism, the computer is viewed as a source of rich, computational, cognitive tools with which the user can explore and *experience* many concepts and principles. These learning environments are often referred to as microworlds. Microworlds are compared to both instructional simulations and games. A series of design recommendations based on a merger of instructivist and constructivist philosophies is presented and discussed. A software package called *Space Shuttle Commander* is presented as one concrete application of these design recommendations.

OBJECTIVES

Comprehension

After reading this chapter, you should be able to:

1. Summarize the philosophy of constructivism.
2. Compare and contrast constructivism with other cognitive orientations to learning and instructional design (termed "instructivism").
3. Describe the Piagetian principle of equilibration and the enabling mechanisms of accommodation and assimilation as they relate to the learning process.
4. Summarize the goals of mental model research and integrate the concept of a conceptual model into microworld design.
5. Compare and contrast microworlds with simulations.
6. Describe some game attributes that offer the potential to increase the intrinsic motivational appeal of instruction.

Application

After reading this chapter, you should be able to:

1. Recognize and apply advantages and strengths from both instructivism and constructivism to instructional design.
2. Design highly interactive learning environments that combine characteristics of microworlds, simulations, and games.

This chapter continues the discussion begun in the previous chapter of how visualization techniques may contribute to instructional design. This chapter is devoted to the fifth instructional application of graphics first introduced in chapter 2 — practice. As we will see, the term “practice” may become either insufficient or inappropriate in capturing many of the ideas presented in this chapter. A more general and appropriate term might be “interaction,” because the focus is really on how the student *participates in* and *contributes to* the learning event. Such interactions within a learning environment would include, but not be limited to, practice strategies. For many instructional technologists, the opportunities for highly interactive learning environments that computers make possible represent the major reason for investing (both economically and intellectually) in computer technology (Hannafin, 1992).

As with the previous chapter, the focus here is on how graphics may contribute to the design of the total instructional system (or learning environment). But this chapter goes much further in stressing the most fundamental issues that influence instructional design. For this reason, this chapter will have, by far, less direct discussions of graphics than any other. Graphics are considered as but one resource for developing interactive learning environments. The goal is not to promote graphics, but to build rich and engaging environments where learners can come in contact with the most intriguing ideas that society has to offer. Graphics offer but one interesting medium with which to “paint this landscape.” We might continue this analogy by considering how the human need and talent for artistic expression and inspiration are served by many media — oil, watercolor, written words, spoken words, stone, marble, clay, etc. — as well as by many forms — realistic, impressionistic, surrealistic, functional, natural, etc. Likewise, computer-based graphical techniques offer powerful resources to help fulfill the basic needs of learning and support the talents of instructional design. Throughout this chapter, you are encouraged to consider all instructional media and strategies, but you are also reminded to carefully consider all the graphical ideas and resources discussed so far.

In addressing the issue of instructional interactions, this chapter will present another, completely different, orientation to learning than that presented so far — constructivism. The concept of constructivism represents a dramatic alternative view to instructional technology. The advice from the previous chapter that instructional designers need to recognize and confront their own philosophical beliefs about learning and instruction becomes even more crucial in this chapter. Again, your interpretation and resolution of these issues will largely depend on this philosophical introspection.

CONSTRUCTIVISM AND ITS IMPLICATIONS FOR INSTRUCTIONAL DESIGN

An historical context may be useful at this point to better understand constructivism and its implications in instructional design. At present, there are two dominant and divergent interpretations of instructional technology, and both envision a significant role for computers in learning and education. The first view is closely aligned with instructional systems development (ISD) and treats instructional applications of computers as related, at least historically, to conventional applications of other educational media. This is the view that has dominated this book thus far. The second interpretation of instructional technology, based on constructivism, considers the computer as a rich source of cognitive tools for learners — an electronic type of "Play Doh" (Rieber, in press). Let's consider the roots of these two perspectives.

The formal beginning of modern instructional technology is usually traced to the convergence of B. F. Skinner's application of behavioral learning principles to instruction, usually called **programmed instruction** (PI), and the audiovisual movement of the mid-1900s (see Reiser, 1987, and Saettler, 1990, for detailed historical overviews). Skinner was well-known for creating various teaching machines designed to deliver highly structured instructional treatments to learners. Teaching machines carefully controlled and delivered predetermined reinforcement schedules during instruction — a skill that Skinner found teachers largely unable to perform. These teaching machines were highly interactive, but also tended to be quite dull and tedious. PI, though generally effective for lower-level learning such as fact learning, was largely inappropriate for higher-level learning. Many current applications of computer-based instruction are really just extensions of the PI paradigm.

Instructional systems development (ISD), as previously defined and discussed, also has its roots in PI. Many PI principles became cornerstones of ISD. For example, the PI principle of **objective specification** was the precursor to behavioral objectives — the idea that the required learner response should be determined in advance in precise, observable terms. **Empirical testing**, the idea that successful lesson components (e.g., appropriate reinforcement, cueing, step size, etc.) could only be determined based on actual field-testing, was the forerunner to formative evaluation (Hannafin & Rieber, 1989a). The PI movement is often criticized today, especially given the popularity (and potential) of the cognitive movement. It is true that PI had serious limitations in covering the breadth of learning outcomes. It is also true that PI conformed to the behaviorist assertion that, essentially, environments control people's behaviors. However, PI remains the first true experiment in seriously attempting to apply learning theory to instructional practice. PI successfully fulfilled the criterion that defines any technology — the application of basic knowledge for a useful purpose — and for that reason PI offers many important lessons for future attempts at harnessing other technologies for instructional design.

Cognitive psychology has had a strong influence on ISD in recent years. Cognitive influences have, for the most part, successfully shifted primary attention from the instruction to the learner (Gagné & Glaser, 1987). Cognitive psychology has persuaded instructional technologists to accept the need to consider what happens in between the stimulus and response (i.e., cognitive or mental processing) as the most important part of the learning process, despite the inability to directly observe this process. At first glance, this point may seem trivial and academic — stuff that makes for good discussions in graduate school classes and nothing else. In actuality, this is a significant turning point for the field and is especially relevant for instructional designers. Cognitive models, such as the information-processing model introduced in chapter 4 (see Figure 4.1), have become the focus of instructional design. Cognitive concepts, such as mental encoding and retrieving, depth of processing, metacognition, and so on, have expanded the range of instructional ideas and have opened up new approaches for identifying and solving instructional problems.

Despite the positive influence of cognitive psychology on instructional design, the skill, task, and procedural aspects of "the model" are still largely retained. As discussed in the last chapter, instructional design is still largely based on achieving the learning objectives identified early in the process. Thus, in general, the goal of any one instructional design is to bring the learner to the point of mastering the learning objectives as efficiently and as effectively as possible. Certainly, a learner's prior knowledge, abilities or aptitudes, needs, and interests have a major influence on how the instruction is designed. However, most of the major instructional decisions, such as how content is selected, sequenced, structured, and presented is usually made on behalf of the learner. Some use the term "neo-behavioral" to define this "mingling" of behavioral and cognitive philosophies (Case & Bereiter, 1984).

The term "practice" is most appropriate in this first interpretation of instructional technology because it describes the interaction as per the events of instruction. By following presentation strategies with practice, the lesson information completes, in a sense, a cycle or "round trip" between the instructional materials and the learners — the instruction elicits a response from the learners, followed by the instruction providing the learners with appropriate informational feedback about their performance. Practice is viewed as but one part of an instructional system, and, therefore, its purpose is to complement the other instructional components (i.e., orientation strategies, presentation strategies, testing, and strategies to enhance retention and transfer).

Given the dominant role that instruction continues to play in this type of learning environment, we might coin our own "-ism" word by using the term "instructivism" to describe this interpretation of instructional technology (Rieber, 1992, in press). Instructivist models characterize learning as a progression of stages starting at the novice or beginner level in a particular domain and ending at the point where the learner becomes an expert. This characterization is similar to Gagné's concept of a learning hierarchy where lower-level learning is considered prerequisite to higher-level learning. All instructivists make the assumption that one purpose of instruction or education is to help the learner understand the "real world." Another assumption is that one group of people, such as teachers and other educators, have the authority and responsibility to make decisions about *what* should be taught and *how* it should be taught to another group of people, such as students. Of course, this means that one assumes that there *is* one objective interpretation of the world to be recognized and accepted and that certain pieces of this world knowledge are important enough for everyone in the society to learn. As we are about to see, not all educators share these views or assumptions.

The second interpretation of instructional technology is patterned after a philosophy of human learning and cognition known as constructivism (Jonassen, 1991a). Constructivists consider the major goal of education to be the creation of a rich assortment of **cognitive tools** that are made available to learners to help them explore their environments. It is then up to learners to decide for themselves what is real or true. Constructivists usually define instructional technology as the generation of computer-based tools that provide rich and engaging environments for learners to explore. These environments are frequently referred to by constructivists as **microworlds** (an idea we will revisit in depth later in this chapter) because they allow learners to participate in a set of ideas until they begin to "live" the ideas, not just study them (Dede, 1987; Papert, 1980, 1981). The next section will provide a brief overview of some of the main tenets of constructivism as they apply to learning and instruction.

Constructivism: An Overview

There is a story that someone once commented to philosopher Ludwig Wittgenstein that people living in medieval Europe before the time of Copernicus must have been pretty stupid to have believed that the sun actually circled the Earth and that common sense should have told them the opposite was true. Wittgenstein is said to have agreed, but also wondered what it would have looked like if the sun *had* been circling the Earth — the point being that it would have looked exactly the same to most people (Burke, 1985). The idea that the Earth was at the center of the universe was just as true to these people as the concept that the Earth orbits the sun is to us. Information does not become knowledge just by its telling.

It is tempting to believe that we, living today, somehow know the *real* truth about the world, that we are somehow better informed than those poor, ignorant folks who lived many years ago. Ours is the real science, right? But before you answer this question, you need to examine your beliefs, even those of supposedly objective truths from mathematics and science. How do you *really* know that the Earth goes around the sun? Just as Wittgenstein observed, our perceptions tell us something very different, yet we have come to accept another fact as being true and our perceptions as being false. All too often, we teach people something as being true without considering what this really means at the individual level. Much education is involved in telling people what to believe. However, true understanding cannot be *imposed* on someone, but instead must come about by a personal revelation (Bruner, 1990).

Actually, science offers some stunning historical examples of how differences in interpreting the world actually meant that the world *was* a different place to people. It all depended on one's point of view. Consider the idea above, proposed first by Aristotle, that the Earth is at the center of the universe and is unchanging. If you do not believe that the Earth changes, then you do not *look* for changes. Our view of science as discovering and exploring the heretofore unknown does not exist in such a world. Compare this to a Newtonian world, where the Earth circles the sun in an elliptic orbit according to certain laws of nature. Aristotle's ideas are just plain wrong in Newton's world. For hundreds of years, Newtonian physics represented the truth of the physical world. The role of science was to gather more information and search for other laws of nature — a view that persists today. But in Einstein's world of curved space and black holes, Newton's laws do not seem to be enforceable — even the behaviors of time and light can change (Hawking, 1988). If Newton is wrong, then maybe so is Einstein. Perhaps the universe really is just a grain of sand on some cosmic beach.

Constructivists believe that each of us defines the world (and ourselves) by what we know and believe (Goodman, 1984; Watzlawick, 1984). Each person perceives and interprets the world in a unique way. Instead of suggesting that knowledge can be transferred from one person to another, information from the environment is used as building blocks for individuals to construct knowledge. This construction process is believed to be a natural consequence of meaningful interaction with one's environment or culture. One's knowledge is never static, but dynamic and ever-changing.

But what constitutes *meaningful* interaction? Consider Newton's first law that states that an object at rest remains at rest and one in motion remains in motion unless acted on by some outside force. Compare two very different instructional designs for teaching this principle. First, consider a physics class where a teacher lectures about the principle to a roomful of students sitting attentively in their chairs, followed by a series of homework problems from the textbook. Next, consider a second classroom where the teacher has each student build and test a series of ramps with a variety of objects (in order to test different levels of friction). The first scenario has students interacting with information selected and interpreted by someone else. In the second scenario, students begin by interacting with the principle *itself*. The teacher's job is to facilitate, manage, or at times, guide, the students' interactions.

Is Newton's first law *for real*, or are there a series of general conclusions based on shared experiences that people can resolve among themselves? With help from the teacher, the group may form some consensus about Newton's first law, but the truth of the law rests within each individual. Interestingly, there is research indicating that physics students who learn physics given instruction similar to the first scenario can pass tests, but may actually revert to their personal view, or theory, of the world when confronted with novel physics problems to solve (Eylon & Linn, 1988). Students may know how to compute the formulas, but their conceptual understanding may not have been changed. See Box 8.1 for a follow-up discussion on this second instructional scenario using the physics of baseball.

Constructivists believe that learning is enhanced in environments that provide a rich and varied source of engaging experiences (Papert, 1988). Computer enthusiasts feel that the computer offers a powerful medium for exploring and discovering many ideas, just as a young child might explore the concepts of volume with a sandbox and mass and momentum with marbles. The computer's ability to present graphical representations is usually considered one of its most important attributes. In constructivism, quality of knowledge structures, not their quantity, is the issue. In other words, learning is not about acquiring new knowledge, but the constant reconstruction of what someone already knows (Forman &

Pufall, 1988a; Fosnot, 1989). As a person's knowledge structures are continually “revised,” there is the occasion where a new structure is formed because new information just no longer matches the available structures. As Forman and Pufall (1988b) note: “Central to constructivism is the assumption that to know is to continually reconstruct, to move from a more to a less intuitive state” (p. 240). The cognitive theories of Jean Piaget still provide among the best accounts of the constructivist view in education.

Box 8.1

How Far Can You Throw? — An "Exercise" in Constructivism

Having been born and raised on the Southside of Pittsburgh, I grew up surrounded by baseball stories. People there sometimes debate who had the best throwing arm of all time. My own personal choice is Roberto Clemente. Clemente played right field for the Pittsburgh Pirates until his tragic death in a plane crash on New Year's Eve, 1972. According to one account, which may be perhaps more legend than fact (though I have chosen to believe it), is that at old Forbes Field he once threw a baseball over 400 feet — on a fly — just in time to tag out the base runner sliding in at home plate.

How far can *you* throw a baseball? One hundred feet? Two hundred? Three? How about a mile? "Whoa!" you say, "I have a major league arm, but it isn't bionic!" The point is that no matter how far you think you can throw, you know that eventually the ball is going to come to a stop. Find the highest hill or wait for the strongest wind before you toss it but the outcome will inevitably be the same — the ball will come to a dead stop.

Little wonder that Aristotle thought that the "natural" state of an object was at rest. Objects seem to "seek" or "prefer" to be at rest. However, Isaac Newton said otherwise. His first law of motion states that "every body persists in its state of rest or of uniform motion in a straight line unless it is compelled to change that state by forces impressed on it." In other words, an object at rest will stay at rest (overlapping with Aristotle), and an object in motion will continue in motion (so much for Aristotle), unless something else comes into the picture. But all our everyday experiences lend far more support to Aristotle than Newton. So why do we believe (and teach) Newtonian, rather than Aristotelian, physics in our schools? Perhaps a better question is do we really *believe* Newton?

In fact, most of the credit for Newton's first law really belongs to Galileo (Newton was born the year in which Galileo died). Galileo proved to himself that Aristotle was wrong by the following set of experiments. Place a wooden block on a perfectly horizontal surface. Give the block a push and watch it slide a short distance until it stops. Now repeat the experiment over and over with smoother and smoother blocks and surfaces. Assuming that you keep giving the block the same size push each time, you will notice that the block goes a little further each time the smoothness of either the block or the surface is increased. Now try the experiment with a ball instead of the block.

Again, the distance traveled is increased (be sure to keep the size of the push the same each time). You have probably figured out by now that we are simply reducing the amount of friction that is being exerted against the object. A ball rolling across a frozen lake or a golf ball hit on the moon will go a long way before stopping because of the same principle (although both *will* eventually stop). Although Galileo could not eliminate friction completely, he extrapolated his findings to a "what *if* we removed all the friction" situation and concluded that the object would move forever in a straight line.

So how far could *you* throw a ball if all other forces, such as gravity and friction, were removed? According to Galileo or Newton, even the weakest toss would make the ball go forever. Do you believe this, I mean, *really* believe this? Maybe you need to prove it to yourself. Just don't throw the ball near the speed of light, then the rules change.

Influence of the Work of Jean Piaget

Educational applications of constructivism are closely associated with the learning theories of Jean Piaget (Vuyk, 1981). Although the brief account that follows may appear to many as a caricature of Piaget's theory, it should help in understanding and applying constructivism to the discussion of educational implications in later sections.

Piaget's theories can be classified in two ways — stage-dependent and stage-independent (Mayer, 1983). Most of the attention is usually given to Piaget's stage-dependent theory, which suggests that there are four stages of cognitive development that people supposedly progress through (at least potentially) in their lives — sensorimotor, preoperational, concrete operations, and formal operations. However, our attention will be devoted to the Piaget's **stage-independent** theory.

Piaget's stage-independent theory concerns two assumptions about how internal mental structures are formed (Piaget, 1952, 1970). The first is the need for **adaptation**, or the ability of an individual to survive and prosper given an ever-changing environment. The second is **organization**, which is one's need or desire for a stable or coherent world. These two processes create an internal or intrinsic conflict for people. The goals or needs of one process directly contrast those of the other. Lifelong learning requires a constant balancing between the two. Just as one struggles to achieve an organized world, the environment presents a new situation or problem. Piaget defined a process, called **equilibration**, that explains how people accomplish this balancing act. Equilibration consists of two mechanisms: **assimilation** and **accommodation**. New information from the environment is assimilated, or subsumed (or understood), under an already existing mental structure. For example, a baby who has learned to throw a tennis ball is just as likely to throw an orange or an apple the first time each is encountered. Accommodation, on the other hand, describes the process where the child builds new structures from the existing structures when the new information no longer fits. Thus, the baby soon learns that some round objects are meant to be thrown, but others are to be eaten.

Life's everyday encounters with the environment inevitably lead to one natural conflict after another, conflicts that are resolved by assimilation and accommodation. Interestingly, learning can only occur when an individual is in a state of **disequilibrium**, also known as **cognitive conflict**. When confronted with new information from the environment, a person naturally seeks to assimilate, or incorporate, this information into structures that already exist. The process of accommodation is triggered when new information no longer fits or matches the existing structures, necessitating the formation of new structures. According to Piaget, this process never ends, though the range or breadth of potential new structures that can be formed are linked to the developmental stage of the individual. But that is another story.

Educational interpretations of constructivism consist of three properties that are closely aligned with Piaget's theories: **epistemic conflict**, **self-reflection**, and **self-regulation** (Forman & Pufall, 1988b).

Epistemic conflict is really just the Piagetian process of equilibration described above. Learning is a result of trying to resolve a problem encountered in the environment that is outside the person's repertoire. Of course, the conflict may have been *artificially* induced, such as a problem presented by a teacher, but resolution of the problem can only be achieved by the individual. In the constructivist vernacular, each resolution is a construction. Just because the environment has posed a problem or conflict does not mean that the individual will choose to pursue resolution. If the problem is perceived as too easy or trivial, then the individual will not find the problem worth pursuing. If the problem is too difficult, the individual may simply choose to ignore it.

The property of self-reflection involves an individual's deliberate attempt at objectively and explicitly representing reality in response to a conflict. Arriving at a resolution or solution to the conflict involves the property of self-regulation. Cognitive structures are spontaneously restructured according to the mechanism of assimilation and accommodation. Old mental structures become more refined or comprehensive. New mental structures are formed. Once conflict and reflection trigger self-regulation, the individual acts until resolution is attained, either by explaining the new information as another, extended example of something that was already known (assimilation), or by the formation of something new (accommodation).

Microworlds

What does all of this have to do with instructional design? If one accepts the constructivist notion that knowledge is not transferred from one source to another — such as from instruction to the individual — but is personally constructed as a result of cognitive conflicts with the environment, then “instruction” is really a misnomer because individuals teach themselves. However, we will use the term “instruction” to describe the deliberate attempt to structure the environment in such a way so as to foster, nurture, or trigger the equilibration process in an area of inquiry believed relevant. Constructivists usually use the term **microworld** to describe placing learners in contact with such learning environments (Papert, 1980; Dede, 1987). Table 8.1 lists some characteristics of microworlds, as defined and explained in the sections to follow.

TABLE 8.1 Characteristics of a microworld

- A small, but complete subset of a domain.
- The simplest model of a domain that is recognizable by an expert of the domain.
- Provides an immediate "doorway" for novices to gain immediate access to a domain through experiential learning.
- Provides general, useful, and syntonic learning experiences.
- Provides learners with "objects to think with."
- Promotes problem solving through "debugging."
- Shares characteristics of an interactive "conceptual model."

Probably the most well-known computer-based application of constructivism is LOGO, a computer language that reflects and promotes Piagetian learning. LOGO was the result of a collaborative effort between the Massachusetts Institute of Technology, and Bolt, Beranek, and Newman, and was initially funded by the National Science Foundation. Many people contributed to LOGO's development, including Wally Feurzeig, Daniel Bobrow, Hal Abelson, and Andy diSessa. However, Seymour Papert is usually credited as LOGO's chief developer and spokesperson. LOGO lets learners explore many areas, including mathematics, science, and metacognition (thinking about thinking), by placing them in contact with a microworld in which these concepts are represented. A microworld, as the name suggests, is a small, but complete subset of reality to which one can go to learn about a specific domain. Personal discovery and exploration are essential ingredients of learning in a microworld (Dede, 1987; Papert, 1981).

Microworlds are among the most promising attempts at creating computer environments that foster an individual's construction (assimilation and accommodation) of knowledge. Microworlds, though a constructivist invention, offer instructional designers two key advantages. First, microworlds present learners with experiences *within* specific boundaries of a domain. Second, microworlds offer learners “stepping stones” *between* interconnected ideas within the domain by allowing rudimentary ideas to first become established and then transformed into more sophisticated aspects of the domain.

Turtle geometry, as defined and discussed in chapter 3, is one such LOGO microworld that gives learners access to geometric principles through interactive graphics (Abelson & diSessa, 1981; Lockard, Abrams, & Many, 1990; and Lukas & Lukas, 1986). Students “drive” the turtle, which leaves a trail as it goes around the screen. The turtle commands, known as primitives, express fundamental geometric ideas of space and distance. As alluded to in chapter 3, the purpose of turtle graphics is not to produce graphics, but to use graphics as the key for experiencing a set of powerful ideas that, in turn, leads to learning about mathematics and science. The rest of this section elaborates on these powerful ideas.

Successful LOGO learning experiences are founded on several key ideas, many associated with programming. For example, LOGO is a procedural language that encourages *top-down* problem solving in which a large problem can be broken down into more manageable chunks. Students can increase the turtle's vocabulary by creating new commands, or procedures. The definitions of new turtle procedures consist of LOGO primitives, as well as procedures created earlier.

However, the turtle geometry microworld best represents a constructivistic learning environment by the turtle simply being a good “object to think with” (Papert, 1980). At the heart of constructivism is a search for other good objects that learners can use to construct knowledge. Almost anything can become a good object to think with: pots, pans, mud pies, blocks, Legos, etc. Some are more flexible and generalizable to a variety of domains than others. Meaningful interaction with objects in the environment liberates and encourages the equilibrium process.

The turtle is but one example of an object to think with that is made possible through computer-based visualization. Papert contends that the computational and graphical power of the turtle makes accessible to children certain ideas from the world of mathematics and problem solving that previously were considered too formal or abstract for young learners. Two characteristics of the turtle help make this possible: the turtle as a transitional object, and the turtle as an aid to debugging. Both characteristics offer many lessons to other would-be microworld designers.

It is common for young children to begin using LOGO for self-guided learning within minutes of encountering the turtle. This is believed to be achieved by the role of the turtle as a transitional object between the children and the computer. The turtle is **body syntonic** with the child in that both share two important characteristics — a position and a heading. This simple fact has powerful learning consequences. From the start, even a young child has something in common with the turtle. This commonality immediately provides the bridge to new ideas. Papert contends that young children quickly anthropomorphize the turtle (giving it human characteristics), thus creating an **ego syntonic** relationship with the turtle. This encourages the Piagetian concept of **decentering**, in which young children begin to interpret the world from several perspectives. The anthropomorphization of the turtle also gives children a way to express mathematical ideas. Children begin to acquire the vocabulary of turtle geometry through their communications with the turtle. Since the language of the turtle is LOGO and the language of LOGO is mathematics, LOGO gives children a means to verbalize mathematics (Papert, 1980).

The second important characteristic of the turtle is as an aid to **debugging**, or the identification and correction of errors within a computer program. Whereas errors are something to be avoided in most forms of instruction, constructivists prefer the idea that errors are a natural consequence of interactions with the environment. Instead of being negative, errors are useful so long as they provide a rich source of information to help guide subsequent interactions. Successful error handling drives the way in which an individual adapts to meet other challenges from the environment. The informational feedback that errors provide is very potent, especially when a learner has a strong commitment to the action that triggered the error. Error detection is made intuitively obvious in LOGO with the turtle's role as a graphical tool. The turtle's animated graphics provide instantaneous graphic feedback. This rapid exchange between the learner and computer in the form of learner action/intention and animated feedback encourages risk-taking and hypothesis-testing. The forming and testing of hypotheses based on animated graphical feedback can be a powerful learning strategy.

Papert (1980) suggests that microworlds, like all powerful ideas, should fulfill four criteria: they should be simple, general, useful, and **syntonic**. Syntonic learning, which loosely translates as “it goes together with,” in a sense subsumes the other three criteria. Syntonic learning involves connecting new ideas to prior knowledge and engaging the learner in a never-ending pattern of going from the “known to the unknown.” Constructivists say learner control is essential in microworlds, a point that contrasts with research on learner control of direct instruction that suggests that learners are often poor judges of their own learning paths (Clark, 1982; Steinberg, 1977, 1989).

Learning within a microworld relies on a learner's natural tendency to seek equilibrium. Successful microworlds actually encourage learning conflicts in order to activate the process of equilibration, since it is believed that only through the resolution of these conflicts can learning take place. The trick is to structure the microworld so that learners have an environment in which conflict resolution is within their grasp. It is this purposeful structuring by the microworld designer that offers a link with instructional design and the other issues discussed so far in this book. Microworlds offer learners an opportunity to exercise a cognitive or intellectual skill that they would be unable or unlikely to do so on their own, either because there is no intrinsic reason to do so or because no sufficient tool is available with which to allow them to begin the experience.

THEORY INTO PRACTICE: BLENDING CONSTRUCTIVISM WITH INSTRUCTIONAL DESIGN

Although many principles of constructivism offer much potential in developing successful learning environments, it is usually difficult for people to see practical examples, given the typical constraints found in most schools and training situations. Indeed, any form of instruction, that is, some form of structured learning experience, is totally outside of extreme interpretations of constructivism. In other words, radical constructivism translates into instructional chaos. I feel that a compromise between the instructivist and constructivist “camps” *can* be reached. As a start, the next section will discuss several areas of research and development that complement the design of microworlds.

Mental Models

Mental model research closely parallels microworld design. Everyday activities require us to interact with a complex environment. It has been suggested that people form mental models of the physical world (see Gentner & Stevens, 1983, for a review). A mental model is simply an individual's conceptualization, or theory, of a specific domain or system. The purpose of mental model research is to lay out as precisely as possible how people understand a certain domain.

Students develop and use mental models to help explain and solve general classes of problems. Similar to the Piagetian idea of a mental structure, mental models are loosely organized and forever changing as new interactions with the environment suggest adaptations. So far, mental model research has focused on technical domains, like physics or electricity, simply because they are far more normative and are more easily made explicit than most other domains, such as parenting. Despite the use of such highly specific domains, theorists suggest that people form mental models of a large number of systems ranging from the kitchen stove to Newtonian mechanics (revisit Box 7.2 from chapter 7 for a discussion of ways in which people form mental models of everyday things).

Mental models serve us with both explanatory and predictive skills. Survival demands that we are able to predict everyday events with a high degree of success. The routine need to cross a street is a good example. Beyond all of the perceptual requirements (such as estimating the width of the street and the speed of oncoming cars) is a need to understand the many “street systems” that operate together. Just a few of these systems include the workings of automobiles, traffic lights, and physics. Our understanding of each system is crucial as we decide when is an appropriate time to cross the street as well as if we can casually stroll across or should attempt an Olympic sprint. Any misunderstanding of one of these systems could be as deadly as any misjudgment of distance or speed. For example, consider your mental model of an electric “walk” sign at an intersection and what it means when it begins to flash. Can you still initiate the crossing? What should you do if you're already part way across? Obviously, each interpretation can have dramatically different consequences. Mental models of everyday things usually form through interactions with the environment. However some systems, such as the physical sciences, are difficult to understand through a wide range of random interactions. Microworlds may offer a platform for people to accurately understand any number of systems.

The application of mental models to the instructional design of microworlds involves considering three things: the target system, the user's mental model of the target system, and the building of a **conceptual model** of the target system (Norman, 1983). The target system is the actual system that a learner is trying to understand. Newtonian mechanics, thermodynamics, or even a refrigerator can be examples of target systems. A user's mental model describes his or her personal understanding or theory of the target system. People use their mental models to describe and predict how the target system works. Of course, a user's mental model may not be an accurate reflection of the target system. Consider your understanding of how your home's thermostat controls the furnace. Does setting the thermostat to 90 degrees warm a chilled room any faster than setting it to 80 degrees? If you hold the *valve theory*, you would answer yes. This mental model is based on the idea that the thermostat controls a valve that lets more heat into the room. If you hold the *timer theory*, you would answer no because this theory states that when the furnace is activated, it always puts out the same amount of heat. The thermostat simply signals the furnace to turn itself off when the desired temperature is reached (Norman, 1988). Does pressing the already-lit elevator button in the lobby help ensure that the elevator will really come? Does repeated pressing of the button make the elevator come to your floor faster? Your actions are a result of your mental model for elevators (though some may also be rooted in superstitious behavior).

To help users more accurately understand a system (and subsequently to alter their mental models), a conceptual model may be designed and presented to them. Conceptual models act as both bridges between the target system and a user's model and anchors upon which a user's model can grow and develop (Mayer, 1989). Conceptual models are usually invented by teachers, designers, or engineers. A microworld is largely synonymous with an *interactive* conceptual model. It embodies the simplest working model of a system in which an individual can begin to understand the target system. A conceptual model can often be metaphorical to the target system, such as suggesting that a computer system is like a "desktop." In such cases, conceptual models, like microworlds, offer a temporary doorway to a set of larger ideas. For example, Papert (1980) has recounted the way in which his fascination with gears as a young boy offered him a beginning conceptual model of mathematical ratios and proportions. For Papert, gears became a personal microworld that helped make the many abstract mathematical ideas more concrete for him.

In order for an interactive conceptual model to truly become a microworld, one more condition must be met — students must find the experience personally satisfying and rewarding. Designing a microworld in such a way so that students choose to engage in the activity involves the issue of intrinsic motivation, which was first defined and described in chapter 4. Lawler (1982) has suggested that microworlds, like those presented in LOGO, are successful because they produce "neat phenomena," or "phenomena that are inherently interesting to observe and interact with" (p. 141). However, constructivists offer little guidance on this issue to designers of microworlds. Turtle geometry, for example, may capture an innate human interest in the visual appeal of graphics.

Activities that are intrinsically motivating rely on student-centered incentives, rather than external lesson reinforcement. We will revisit the topic of intrinsic motivation in a later section of this chapter using a context that offers many similarities to microworlds — computer games. But first, we will consider an instructional format that offers the most similarities to microworlds — simulations.

Simulations and Their Relationship to Microworlds

When an instructional designer first hears a description of a microworld, the first reaction is usually to confuse it with a simulation. While characteristics of the two can heavily overlap, each can remain mutually exclusive. It all depends on design and, most important, how they are used in a learning environment. A microworld has two essential characteristics that distinguish it from a simulation. First, a microworld holds the simplest model of a system or domain that is still recognizable by an expert in that domain. Second, the parameters of a microworld are carefully designed to match the level, experience, and interest of the learner.

This second characteristic is the most important because it offers the user an entry point into the domain.

In contrast, a simulation is any attempt to mimic in some form a real or imaginary environment or system. Simulations have a long history in education. Box 8.2 describes the most recent “sibling” to simulations, most commonly called **virtual reality**, that uses the most sophisticated visualization techniques available (Rheingold, 1991). Conceptually, virtual reality systems “transport” the user from one reality to another so that what seems to be present really is not. An essential characteristic of any simulation is that there is a set of rules or model upon which the simulation is based (Willis, Hovey, & Hovey, 1987). Simulations serve two purposes. The first is to provide a means of studying a particular system, such as a scientific simulation. For example, a meteorologist may design a simulation of a tornado in order to better understand the conditions under which tornadoes form. An economist might construct a simulation of a free-market economy to understand the effects of government regulation. In both cases, the simulation would necessarily be based on some theory of the system. In other words, the simulation seeks to model theory. In this way, scientists can test and revise their theories of complex phenomena, because direct experience is either impossible, expensive, or dangerous.

The second purpose of simulations is educational — to teach someone about the system (Reigeluth & Schwartz, 1989). As with scientific simulations, educational simulations are used because there is some inherent reason not to have users experience the system directly. Typical reasons include cost, danger, and inaccessibility. Students learn about the system by observing the results of their actions or decisions through feedback generated by the simulation (Duchastel, 1990-1991). Computer simulations usually offer the advantage of providing the feedback to the student in real-time since the mathematical model of the system is programmed into the computer. Additionally, the computer can be programmed to speed up or slow down the process, a technique that is especially useful if the real system either occurs too fast (e.g., an internal combustion engine) or too slow (e.g., deforestation) for feedback to have any meaning.

Alessi and Trollip (1985) further distinguish simulations on the nature of their interactivity. Some types of simulations, similar to the scientific simulations described above, allow the user to choose or set the value of variables (such as the amount of gravity) and then watch the effects of their choices (such as how much time it takes an object to fall). Other simulations, such as the operation of a complex piece of machinery, give the user chances to learn the operating procedures without the risks and costs associated with its real use.

It is common for simulations to be visually based, although visuals are not an inherent characteristic of a simulation. Visuals may be used in order to provide greater similarity between the simulation and the actual system (e.g., a realistic visual of a cockpit and changing landscape for a flight simulation) (Alessi, 1988). It is hard to imagine an educational simulation without visuals, yet this is simply a design decision. The economics simulation described above might simply use a spreadsheet's row and column design to test a series of “what if” scenarios with the raw data over time. The visual design of a simulation's interface is probably best approached in terms of how the visuals provide

natural mapping between the users' execution and evaluation of intended actions while they are participating in the simulation (see chapter 7 for a discussion of this concept).

Box 8.2

Learning in a Virtual Reality

Simulations allow users to experience and participate in an environment that models some real (or imaginary) system. Simulations provide experiences in a *context* that is hoped to closely resemble the system that is being simulated. Simulations are often used because there is a reason why the system cannot be experience *for real*, such as time, cost, risk, complexity, or unavailability. However, every simulation places the user at a distance from the system being modeled. There is no mistaking the simulation for the real thing if only because users can always look away from the computer screen to remind themselves of the room in which they are sitting. But what about a computer simulation in which users cannot distinguish their real world from the simulated one? What if you looked up in your simulated world, say of the Space Shuttle, and saw the dingy dull ceiling of your office, but the Shuttle's overhead control panel? Proponents of an area of computer development, mostly commonly referred to as **virtual reality** (VR) hope to achieve the illusion that what appears to be present is not. The interface of a VR system is not the keyboard, not a joystick or a mouse, but your own body. VR blends many areas of interest and inquiry — computer science, computer visualization, cognitive and perceptual science, even sociology, philosophy, and ethics.

As first described in chapter one, one of the best characterizations of VR probably comes from the "holodec" on the television show *Star Trek: The Next Generation*. However, VR is not science fiction, though the crudeness of the graphics based on current limitations in computer processing power, leaves much to be desired as compared to its sci-fi counterpart. Unlike traditional computer simulations, VR attempts to remove and supplant all competing stimuli with computer-generated stimuli. In order to achieve this, current VR systems make the user wear a helmet containing two video displays, one for each eye, as well as a speaker for each ear. By sending separate visual images to each eye, each offset slightly from the other, the user experiences stereoscopic, or 3-dimensional, vision similar to that of a 3-D movie. The helmet also has sensors to relay information about the user's head movements to the computer. If the user looks right, left, up, or down the computer instructs the video displays to show the corresponding images from the VR world. The user also wears a DataGlove, a special glove with fiber optic sensors that can detect hand movements and relay the signals to the computer. Hence the user moves around the VR through special hand signals, such as pointing the index finger.

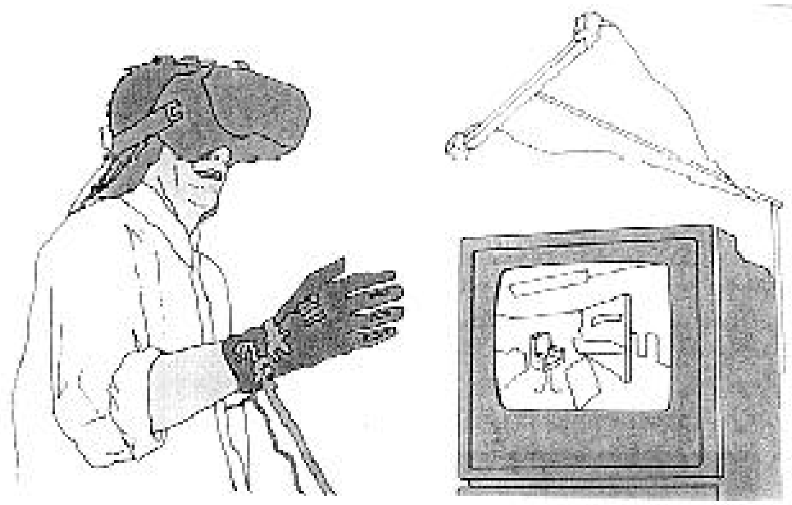

Reprinted with permission of *The Computing Teacher*. Copyright 1992, International Society for Technology in Education, Eugene, Oregon.

The physics of a VR world are determined by the computer program, so it is common for users to fly about or magically go through walls. A VR world might contain "reverse gravity" so that objects fly up to the ceiling when dropped. It is difficult to get a true understanding of VR just from reading about it — VR must be experienced. I had the opportunity to experience VR, courtesy of Meredith Bricken, a VR scientist at the Human Interface Technology Lab (HITL) at the University of Washington in Seattle. The VR world I visited was unique in that was designed and constructed by children ranging in age from about eight to 16 years old who were participating in a summer technology camp at Seattle's Pacific Science Center. Although the graphics from the video displays were slow (about 10 frames per second) and crude, and the hand movements were awkward, the fact that my visual system was completely dominated by the 3-D graphic display made it surprisingly easy to experience the "out of body" sensation. I felt like I had left the HITL and "entered" the children's VR world.

Whether or not virtual reality represent extensions of simulations or microworlds depends largely on how they are used. I still prefer the microworld label to distinguish environments, simulated or otherwise, on the basis of how they allow a user to enter and learn about the world at their level. An example of a VR simulation would be an architectural firm that uses VR to let clients visit and modify a building before it is constructed. However, an example of a *VR microworld* would be a medical school student learning about sinus cavities in the skull through actually visiting and exploring these recesses.

In another sense, virtual reality resembles neither simulations or microworlds. For example, one project being directed by Dr. Michael McGreevy of the NASA/Ames Research Center hopes to physically transport the necessary virtual reality equipment to

Mars, so that people can experience Mars without leaving earth. Strictly speaking, this could not be considered a simulation because the visual and tactile stimuli actually are *originating* from Mars, though the human experience is achieved on Earth (given, of course, the three-second delay to transmit the signal between Mars and Earth). As another example, a surgeon might practice a difficult procedure on a VR model of the patient before performing the *real* operation. Another prediction is that a surgeon with unique skills in San Francisco could *actually* perform the operation, VR style, on a *real* patient with only minutes to live in New York City. Space shuttle astronauts could practice the tricky procedures to share errant satellites before venturing out into space. Tired of putting up with traffic on your way to work in the morning? Then you might like working in a VR office. Although physically you and your office mates are at home, through VR you can "go to work."

What are the possibilities of VR for learning and training? If VR developers choose to follow microworld applications, then users must be able to both modify and construct their own VR worlds to match their learning needs (Bricken, 1991). Instead of watching a simulated object fall at varied rates by playing with a gravity setting, users would find *themselves* falling (although they would also experience some side effects, such as nausea, because the visual system is a more powerful trigger for physical reactions than most people realize — similar to the feeling one gets simply by watching a video of a roller coaster ride). In a VR microworld, the goal is to achieve the Zen-like sensation of *becoming* a molecule, a gear, or a neural synapse. In this way, VR microworlds create a sense of empathy between the user the system of interest. For better or worse, the first VR laboratories have only been in tinker mode with the most serious applications so far in the entertainment industry and the military (there is a Nintendo version of the DataGlove). Unfortunately, most of the general public's knowledge of VR has been of its dark side as shown in Stephen King's movie *Lawnmower Man*. Hopefully, as the cost and limitations of VR decrease, VR can be among the resources that education can casually call upon for environments in which users construct knowledge.

The degree of realism in a simulation, or the extent to which it resembles the actual experience, is referred to as its **fidelity**. The assumption that the best simulations are as realistic as possible is a false one. Simply increasing the fidelity of a simulation will *not* necessarily increase learning (Alessi, 1988). Instead, the relationship between learning and a simulation's fidelity is nonlinear and depends on the instructional level of the student. As shown in Figure 8.1, while it may be appropriate to provide experts with as realistic a simulation as possible, there appears to be optimal levels of fidelity for experienced and especially novice students. In other words, too much realism may cause more harm than good, especially for inexperienced students.

In trying to distinguish between microworlds and simulations, let's start with an example of a microworld that is not a simulation. Cuisenaire rods, a set of colored rods of varying lengths named after George Cuisenaire, the Belgian educator who developed them, act as a microworld for many mathematical ideas (Fuys & Tischler, 1979). Through their manipulation, many young children are introduced to a set of mathematical ideas that are fundamental to learning other, more sophisticated concepts. Despite their simplicity, even the ablest mathematician recognizes them as a mathematical tool. Cuisenaire rods offer mathematics at a level children can understand. On the other hand, Cuisenaire rods offer not a mathematical simulation, but permit *real* mathematics to take place.

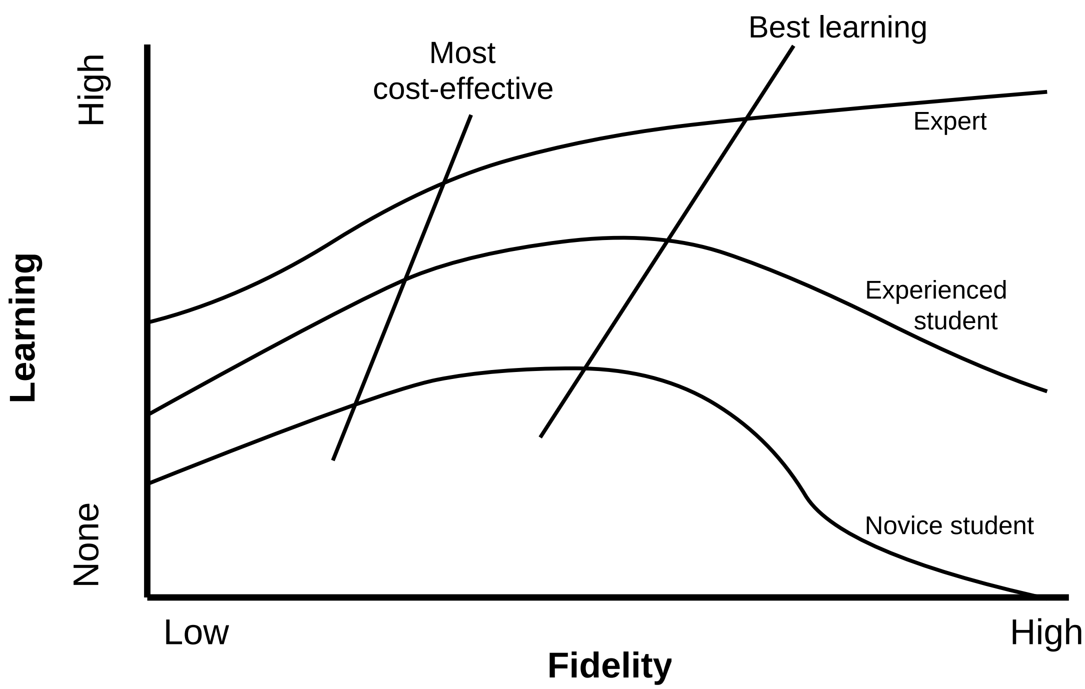

FIGURE 8.1
The relationship between learning and the fidelity of an instructional simulation on the basis of the experience level of students.

However, simulations can be designed that do not offer any significant difference from real-life experiences, such as sophisticated flight simulators used for training by the military and many major airlines. These simulations would not be considered microworlds for most people because they are designed to represent as many of the variables and factors of the real experience as possible. The simulation is not a microworld because the simulation does not match the user — the user must match the simulation. The feedback from this simulation would be largely meaningless and nonsensical to all but the most well-trained user.

Simulations start to become microworlds when they are designed to let a novice begin to understand the underlying model. A computer flight simulation can be designed to permit only limited control and manipulation with only one part of the aircraft, such as the rudder. In this sense, the simulation becomes a *rudder microworld*. Similarly, many microworlds can easily become simulations. Consider a mathematical microworld that involves estimating distances, such as by using the LOGO command FORWARD to move the turtle from one point to another on the screen with as few commands as possible. This microworld becomes a whale search simulation, simply by changing the turtle into an animated boat and the screen target into a whale. The mathematical microworld has not changed, only the context.

Games and Their Relationship to Microworlds and Simulations

Never underestimate the value of play. As adults, we tend to think of play as something that one has to give up when you grow up. However, play serves several cognitive functions in addition to being entertaining and reducing stress. For example, Piaget considered children's play as an assimilation strategy (Piaget, 1951). Through play, one practices a set of information over and over until the individual is completely comfortable and familiar with it. In one sense, play serves as a rehearsal strategy. The knowledge is played over and over in a variety of contexts generated by the individual.

On the other hand, Piaget considered imitation as an accommodation strategy. A child who imitates a parent going off to work by having a doll drive off to the “office” in a toy car with a brief case in the back seat, is reaching out to understand the “go to work” schema. A more detailed account of the value of play is outside the scope of this discussion, but suffice it to say that play *is* valuable for people of all ages. The instructional computer format of gaming closely parallels educational applications of play. Gaming also offers many similarities to microworlds and simulations, though gaming, too, can remain totally distinctive. The purpose of this section is to consider how to take advantage of gaming techniques in the design of microworlds and simulations.

The value of games is that they are fun. Of course, fun is an extremely abstract concept. One common characteristic of most games is competition, in the form of learner against learner, learner against computer, or learner against self (Hannafin & Peck, 1988). There are many negative aspects to competition, especially those involving learner versus learner. Students who constantly lose may become completely turned off to learning. Yet, there are ways to capture the positive aspects of competition by emphasizing a more enduring characteristic, namely challenge. As first described in chapter 4, Malone (1981) has proposed a framework of intrinsically motivating instruction based on the interplay of three characteristics: **challenge, curiosity**, and **fantasy**. In particular, Malone's model has been specifically applied to the design of computer games.

Challenge and curiosity are closely related, and both must be optimally maintained to be effective. Tasks that are too easy can be tedious and boring, and tasks that are too difficult are frustrating and intimidating. In either case, it is unlikely that a student would choose to engage in the activity for even short periods of time. Both challenge and curiosity often result when tasks are novel, moderately complex, or produce uncertain outcomes. An element of surprise results when the expected and actual results for an activity are different. In other words, such events trigger disequilibrium. As previously discussed, completion of challenging tasks can elicit feelings of confidence and competence (Weiner, 1979).

Malone (1981) provides several suggestions for optimizing challenge and curiosity in an educational game:

1. Design every game with a clear and simple goal.
2. Design games with uncertain outcomes.
3. Structure the game so that players can increase or decrease the difficulty to match their skill and interest.
4. Design the game with layers of complexity and a broad range of possible challenges.
5. Provide some clear measure of success for players, such as scorekeeping features, to let players know exactly how they are doing.
6. Clearly display feedback about a player's performance to make the feedback readily interpretable.
7. Provide players with some level of choice.

The element of fantasy is especially important. Fantasy is used to encourage students to imagine that they are completing the activity in a context in which they are really not present. Inducing fantasy relies on mental imagery of contexts that are very meaningful for a student. Fantasy is evident in the intense play of children, especially very young children. Malone (1981) describes two fundamental kinds of fantasies common in the design of computer games: intrinsic fantasies and extrinsic fantasies. As shown in Figure 8.2, the intrinsic or extrinsic nature of the game depends on the degree to which the skill and fantasy are related. Extrinsic fantasies simply overlay some general game context on an existing curriculum area. A common example is the popular “Hang Man” game in which incorrect answers lead to a man being hung, as shown in Figure 8.3. Extrinsic fantasies can re-use the same game design with any content area. There is no mistaking the game elements from the skill or educational value in a game that uses an extrinsic fantasy. In other words, students put up with the skill or educational value given an extrinsic fantasy.

On the other hand, intrinsic fantasies effectively combine or mix the game and skill being learned. The skill and fantasy depend on each other. This means that the skills to be learned are integrated into the fantasy, such as learning about how to use a compass to rescue a party of lost hikers. When effective, the fantasy becomes a meaningful context in which all subsequent instruction can be anchored, or situated (Cognition and Technology Group at Vanderbilt, 1990). Figure 8.4, for example, illustrates an intrinsic fantasy game where players learn about fractions. The game uses the fantasy context of a mine shaft where “some old miner left his ax down in the mine.” Players take turns trying to fetch the miner's ax. The mine shafts are meant to convey the concept of a number line. Each player has an elevator that ventures down the mine shaft to a distance equal to the fraction entered. The player who is closest to the ax is rewarded by having the ax loaded into his or her elevator, where it is transported back to the surface and dumped on his or her side. The player with the most axes at the end of the game wins.

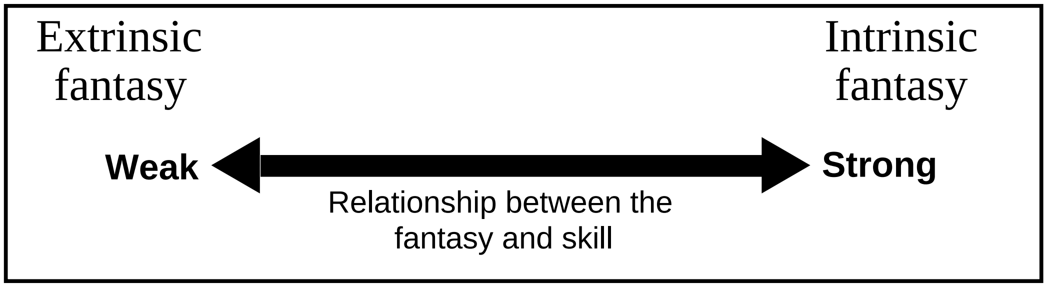

FIGURE 8.2
Games with an intrinsic fantasy have a strong bond between the game's educational value or skill and the game's fantasy whereas in an extrinsic fantasy this relationship is weak.

An interesting element of this game is that there are no wrong answers, just *better* answers. This helps the players explore the concept of fractions in a nonthreatening way. The game allows for any appropriate fraction and does not require lowest terms. Therefore, a player who enters a fraction like “100/200” will see how it is “related” to the fraction “1/2.” As players learn more about fractions, they learn how to tweak fractions, such as entering “499/800” or “501/800” to go just a little before or after “5/8.”

Figure 8.5 illustrates an example of an intrinsic fantasy for another version of the “mystery number” game first described in chapter 2. In this version, students try to locate the mystery number with a radar screen. The mystery number is at the center of the radar screen. As the student guesses, the distance away from the center is shown with a blip. Again, errors are a necessary and useful part of this game. Through careful successive guessing, students can pinpoint the identity of the mystery number. The game's fantasy can be extended to a context where the player is trying to help find a lost friend.

Designers are encouraged to provide an intrinsic fantasy in computer games whenever possible. Students who choose to participate in the game are also, therefore, choosing to participate in the instructional skill of the game. The trick, of course, is in finding appropriate intrinsic fantasies that have wide appeal. This recommendation is also extended to the design of microworlds and simulations. In general, the characteristics of intrinsic motivation — challenge, curiosity, and fantasy — are also relevant to the design of microworlds and simulations. Gaming contexts provide some of the easiest ways to apply these characteristics.

A final recommendation is made cautiously and with some hesitation. Much can be learned about the design of computer games by how children interact with video games. I dispute the widespread belief that fancy, high-resolution graphics and sound provide the strongest source of motivational appeal of video games. Instead, I contend that the best video games are inherently appealing and enduring because of their attention, accidental or otherwise, to challenge, curiosity, and fantasy. Instructional designers should spend a little more time in video arcades watching and talking to the clientele.

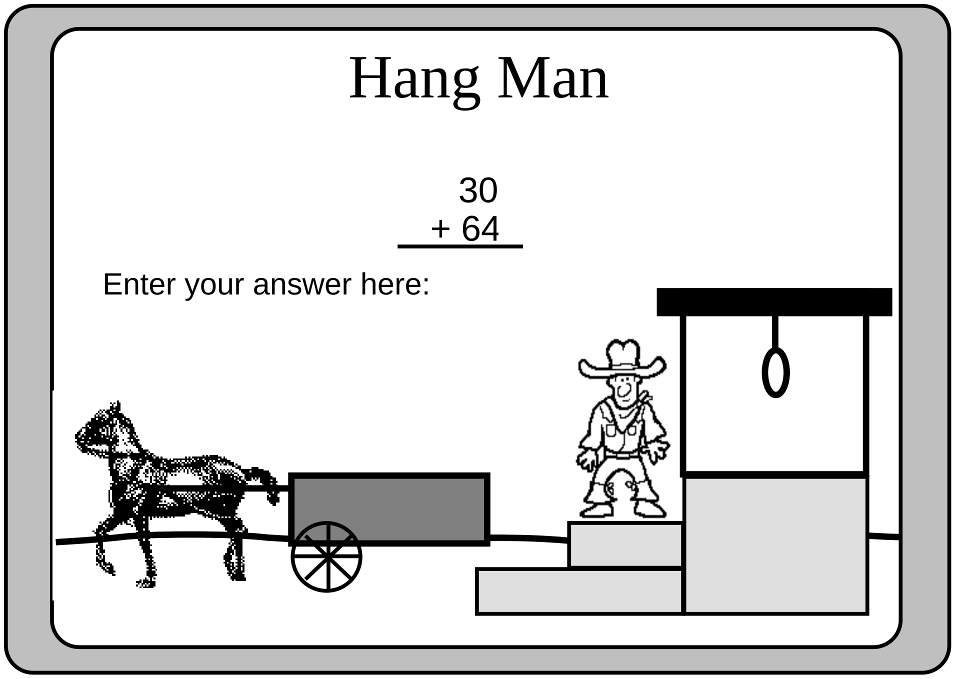

FIGURE 8.3
An example of an educational game using an extrinsic fantasy. The "hang man" game context is traditional and longstanding and has been used in all subject areas. This particular example uses the context for mathematics. If the student answers correctly, a piece of the "get-away" wagon is added. If the student answers incorrectly, the "desperado" moves one step closer to the "hangman's noose." In an extrinsic fantasy such as this, students who find the game's context enjoyable merely tolerate the educational value. This particular game has other problems, such as the ethical questions of promoting capital punishment and the idea of helping someone escape from their sentence.

In fact, many of the popular video games successfully combine the characteristics of microworlds in the gaming. Figures 8.6 and 8.7 show two simple examples of how this combination might be accomplished in designing highly interactive programs for young children or the learning-disabled. Figure 8.6 could be viewed as a “left hand versus right hand” microworld in the intrinsic game context of a treasure hunt. Figure 8.7 is a “1, 2, 3” microworld for learners to explore these simple, yet crucial principles from number theory.

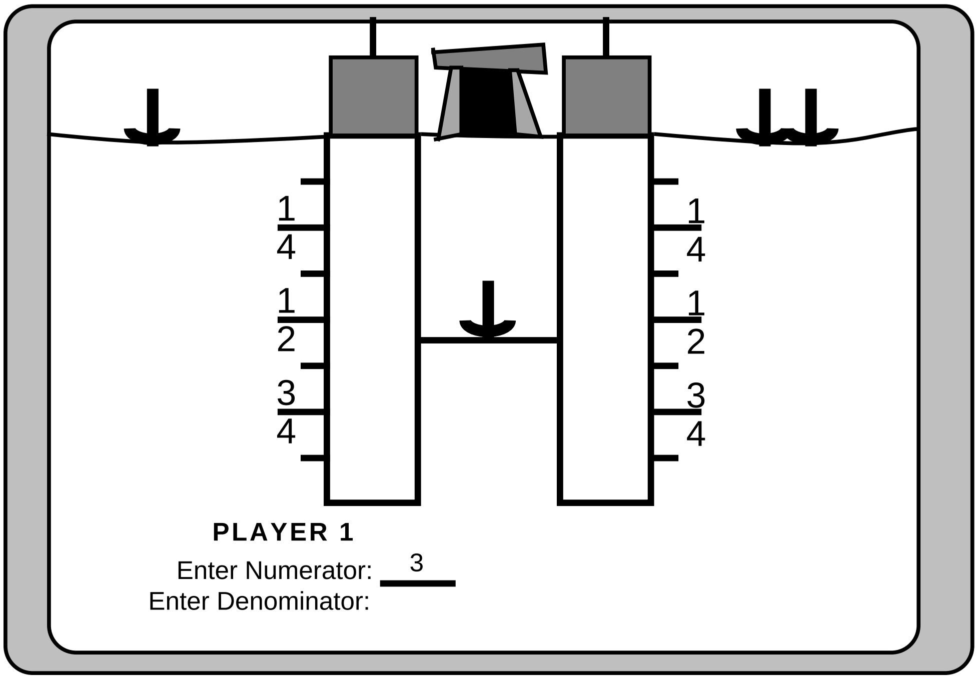

FIGURE 8.4
An example of an educational game using an intrinsic fantasy. The content of the game, fractions, is closely related to the fantasy of a mine shaft. Students play this game in pairs. Each tries to lower their "elevator" to the exact location of the miner's ax. Using animation, the elevator demonstrates the idea of fractions on a number line by dropping the elevator down from 0 (the surface) to the fraction chosen by each of the players. The ax gets loaded into the player's elevator that is closest to the "lost ax." Both elevators are "pulled" back up to the surface and the ax is "dumped" on to the winning player's side of the mine. The player with the most axes at the end of the game wins.

Space Shuttle Commander: Practical Constructivism

All of my trials, tribulations, and adventures in instructional technology have been a result of my own personal "disequilibrium" in my attempting to understand and apply many interpretations of the field. Interestingly, I began my career as a classroom teacher about the same time that microcomputers were invented. My training as a teacher was rooted in the Piagetian approach, yet I found it hard to translate theory into practice, given all of the typical constraints inherent in the public school classroom. But rather than give in to the frustration, I, too, adapted (i.e., assimilated and accommodated) by instituting a pattern of compromise between philosophical and practical circumstances. As a result, my view of instructional technology is an eclectic one.

As a case in point, the next section will describe my attempt at developing a software package, *Space Shuttle Commander* (SSC), that is meant to act as a prototype or model of how one might merge instructivist and constructivist goals and philosophies (Rieber, 1990c, 1992). The next section will provide an overview of SSC, followed by a set of design recommendations that helped guide its development. As you will see, computer graphics have been my main arsenal for realizing the blending of constructivist and instructivist goals in the design of interactive activities. Another way of looking at this approach is simply to blend the best ideas behind the design of microworlds, simulations, and games.

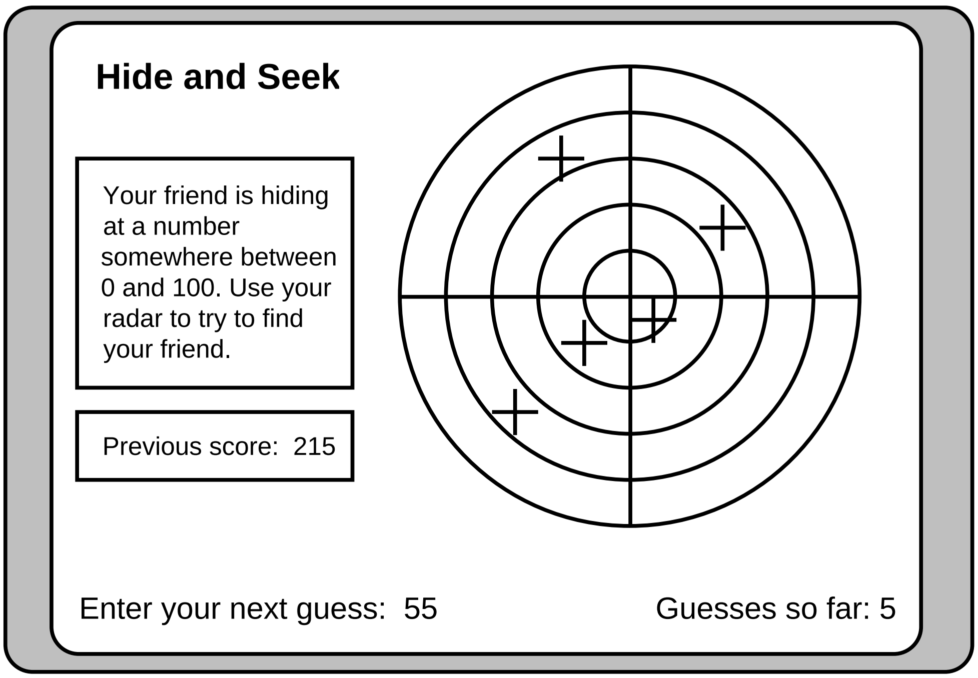

FIGURE 8.5
Another example of an educational game using an intrinsic fantasy. In this example, number theory and estimation skills are intertwined with the game's fantasy. In this example, the goal is to find a friend who is hiding at a number randomly chosen by the computer. Guesses show up as blips on the radar screen. The closer the blips are to the cross-hairs at the center of the radar, the better the guess. Students use the feedback from their guesses to triangulate, or pinpoint, the identity of the random number and, in so doing, find their friend.

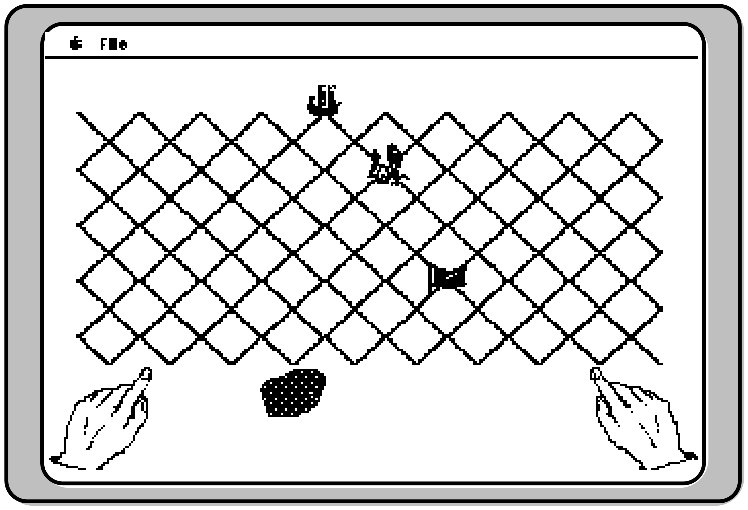

FIGURE 8.6
An example of combining the characteristics of microworlds and games to learn the simple directions of left and right. The student uses the mouse to click on either hand, followed by the computer saying “left” or “right” and the boat animated in the corresponding direction.

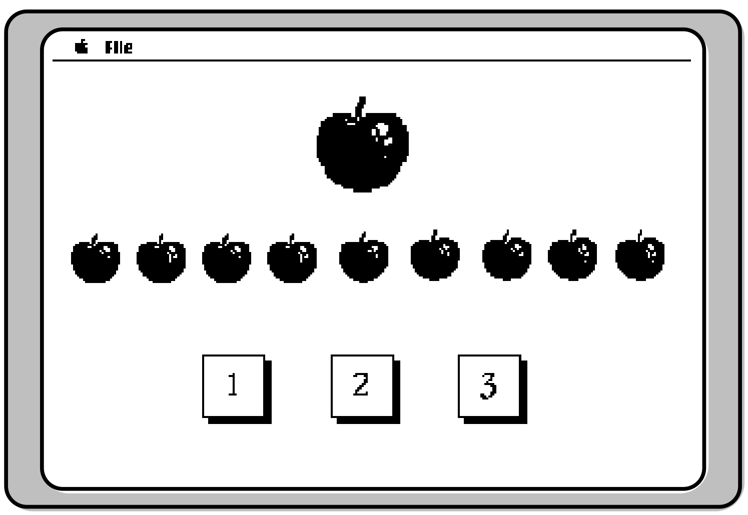

FIGURE 8.7
Another example of combining the characteristics of microworlds and games to learn the mathematical concepts of 1, 2, and 3. The student uses the mouse to click on one of the numbers and the computer, through animation, “eats” the corresponding number of apples. The goal is to eat all the apples.

An Overview of Space Shuttle Commander

SSC is designed for elementary and middle school students. The purpose of SSC is to help these students achieve a wide range of learning goals related to Newton's laws of motion. SSC was designed based on compromises between the extreme views of both instructivism and constructivism. For example, SSC accepts the constructivist position that learners should be given rich and powerful environments to build and transform mental structures. However, SSC also acknowledges the practicality of the current educational system, though this may be regarded as a necessary evil at the present.

SSC is a direct application of a physics microworld first designed by Andy diSessa (1982) involving a screen object called a *dynaturtle*. The dynaturtle closely resembles the more familiar LOGO turtle described in chapter 3, except that it has one more characteristic — velocity. By manipulating the dynaturtle, students can explore motion principles in a simulated frictionless, gravity-free environment. In SSC, the dynaturtle microworld becomes a simulation by placing it in the context of space travel. The dynaturtle becomes a “space shuttle” and students are encouraged to fantasize they are astronauts.

SSC tries to use both tutorials and simulation/gaming (called “flight lessons” and “missions,” respectively) in ways that maximize their strengths and minimize their weaknesses. For example, a tutorial is a good way to present large amounts of information in an organized way. However, tutorials are often dull and prone to promoting passive learning (cf. Merrill, Li, & Jones, 1990a; Jonassen, 1988b; Roblyer, 1988). Simulations and games are usually much more motivating and are well suited to discovery learning, though learning can be difficult to monitor and assess (Alessi & Trollip, 1985, 1991; Hannafin & Peck, 1988). Both the flight lessons and missions introduce students to the laws of motion in nonmathematical ways. The flight lessons present and explain the concepts in a structured, step-by-step way. The missions offer students a series of simulations, most with game-like features. Students pilot an animated shuttle, such as that represented in Figure 8.8.

Traditional instructional design usually promotes deductive learning strategies, such as presenting a concept to students followed by an assortment of examples and nonexamples and by practice (R. Gagné, 1985; Gagné, Briggs, & Wager, 1992). Constructivists, in contrast, promote interaction over explanation. Students are expected to discover, or induce, concepts and principles on their own based on experience and interpretation. Bruner (1986) referred to these inductive learning experiences as *learning by inventing* (p. 127).

The activities in SSC can be used for deductive or inductive approaches. Students can go through SSC in a deductive fashion simply by following the course structure, represented in Figure 8.9, starting with the first flight lesson. This approach takes full advantage of the learning hierarchy designed into SSC, where later skills build on those introduced earlier (Dunn, 1984; R. Gagné, 1985). Each flight lesson “teaches” the respective objectives according to conventional instructional design, and each “mission” acts as a suitable practice activity for each lesson.

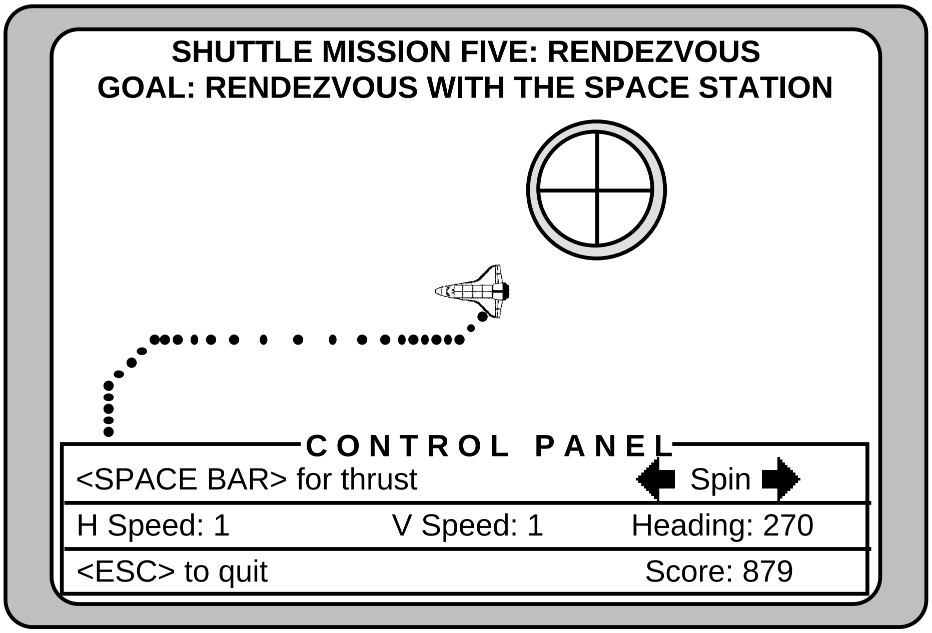

FIGURE 8.8
A representation of the computer screen during an episode of "Mission 5: Rendezvous." The animated shuttle is under student control. Arrow keys rotate the shuttle in 90 degree increments and the space bar gives the shuttle a "kick" or thrust in the direction it is pointing. The goal of this mission is to maneuver the shuttle to the space station.

On the other hand, each mission acts as a stand-alone microworld. Each mission simulates particular aspects of Newton's laws of motion. Early missions are very structured, with the number of learning variables minimized to help make fundamental ideas and concepts as explicit for students as possible. Later missions are very open-ended, but with the option of imposing or reducing structure and complexity. It is possible to have students begin to understand Newtonian mechanics by having them only explore the missions in SSC. Flight lessons would be consulted as additional resources, either as the result of curiosity or confusion.

It is expected that educators would use SSC depending on the philosophical orientation they hold. For example, an instructivist would probably focus on the flight lessons and only consider the missions as practice activities. On the other hand, a constructivist would probably only see value in the missions (and may even object to the presence of the flight lessons). A constructivist would let the learner determine sequence, whereas an instructivist would encourage or require the learner to closely follow SSC's course map. SSC affords a wide range of interpretations, by both teachers and students, on how it should be used.

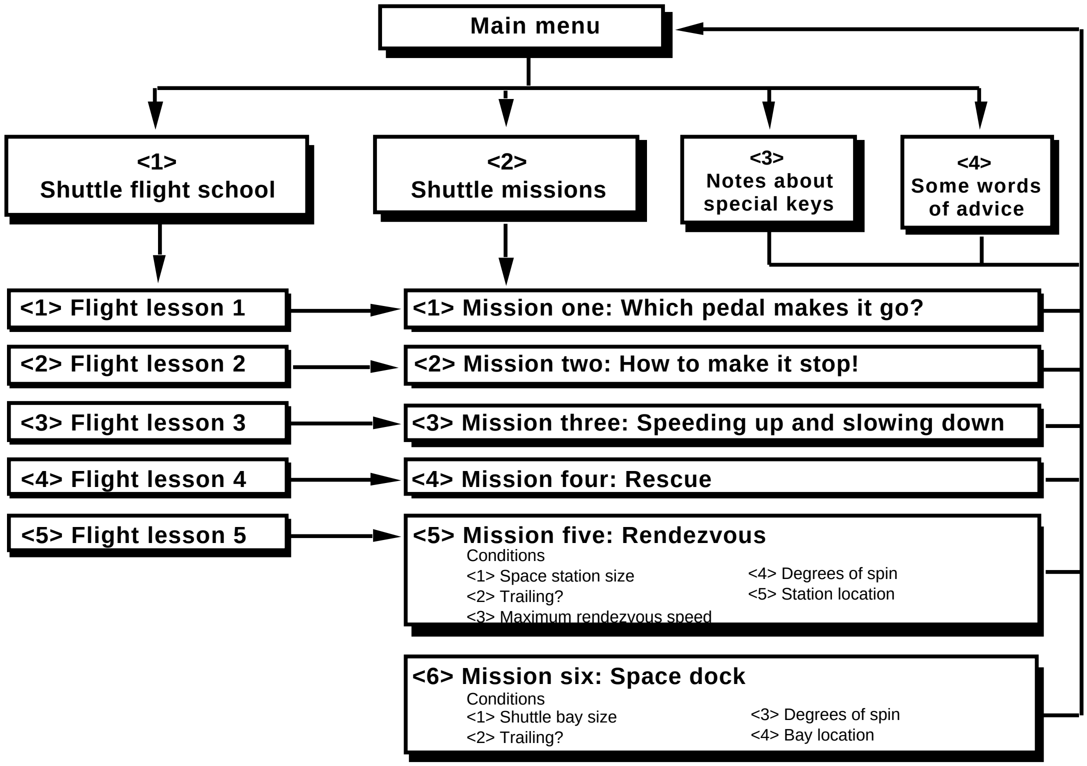

FIGURE 8.9
A course flowchart of *Space Shuttle Commande*r.

Another instructivist influence is on the overall design of the *series* of missions. The missions are hierarchically organized from simple to complex — early missions focus on the simplest ideas, and later missions combine and extend these ideas. For example, missions 1, 2, and 3 take the learner through a series of activities that introduce the simplest aspects of Newton's first and second laws. At first, structure is heavily imposed, but it is reduced as the learner establishes a foothold with these concepts. The first three missions further constrain the learner's experience to one dimension, as shown in Figure 8.10. However, mission 4, as illustrated in Figure 8.11, introduces the effects of two dimensions in a highly structured way so as to make the simplest relationships of two-dimensional motion as explicit as possible. This is accomplished by placing artificial restraints on the microworld. The student is told that the shuttle has been in a collision with a small asteroid, resulting in several malfunctions. For example, the student has no control over steering and must contend with the fact that the shuttle is pointing directly to the right (i.e., 90-degree heading). The shuttle is also coasting in space toward the 0-degree heading (i.e., from bottom to top), and the student only has enough fuel left for three maneuvers. The goal of the mission is for the student to “rescue” the shuttle by using the limited resources to fly the shuttle to the space station located, fortunately, nearby.

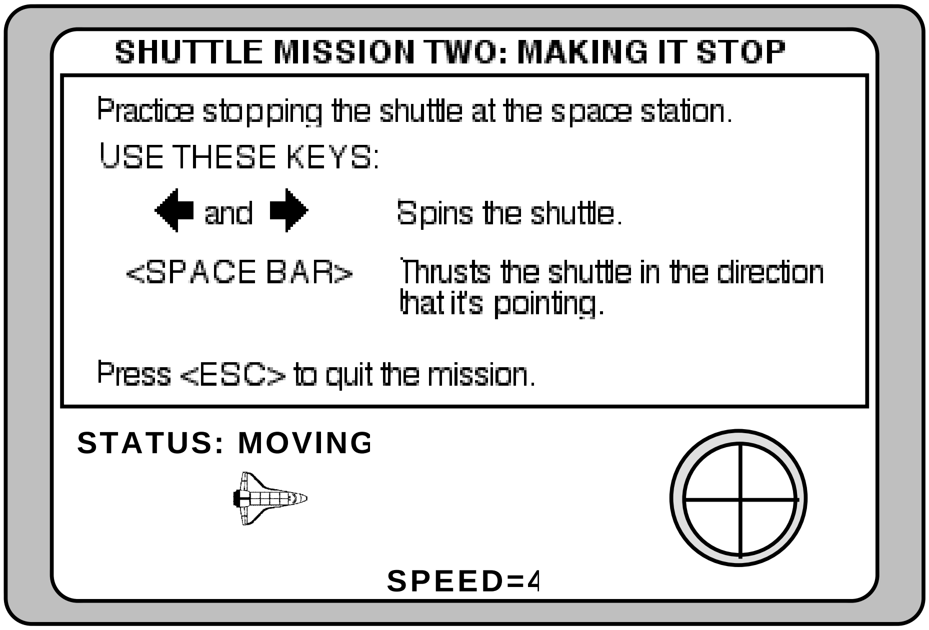

FIGURE 8.10
A representation of the computer during "Mission 2: Making It Stop." By programming the computer to constrain the "spin" to 180 degree increments, the shuttle is confined to one dimension. This makes it easier for the learner to explore and focus on the idea of how forces act on the shuttle's speed and direction.

The final two missions, “Rendezvous,” (see Figure 8.8) and “Space Dock” (see Figure 8.12), are highly detailed simulations. Students who unsuccessfully attempt these final two missions early on are encouraged by on-line coaching (and perhaps by the teacher) to go through earlier missions or flight lessons. Such use of coaching is a recognized instructivist strategy. Of course, constructivists would expect many students to discover much about Newton's laws of motion just by their experiences with these two missions.

A clear and simple goal is made overt to each student in each mission. The goal provides a simple tool for students to evaluate their interaction during the mission. All of the continuous feedback received from the microworld — the motion of the shuttle, the trail left by the shuttle, and the verbal information from the control panel — can be used to compare the shuttle's current state against the desired state (i.e., the goal). For example, the goal from mission 5, as shown in Figure 8.8, is “fly the shuttle to the space station.” Besides its use as a game characteristic, successful completion of the goal also provides a means for teachers to evaluate performance across a group of students, similar to the use of performance objectives in most instructivist models.

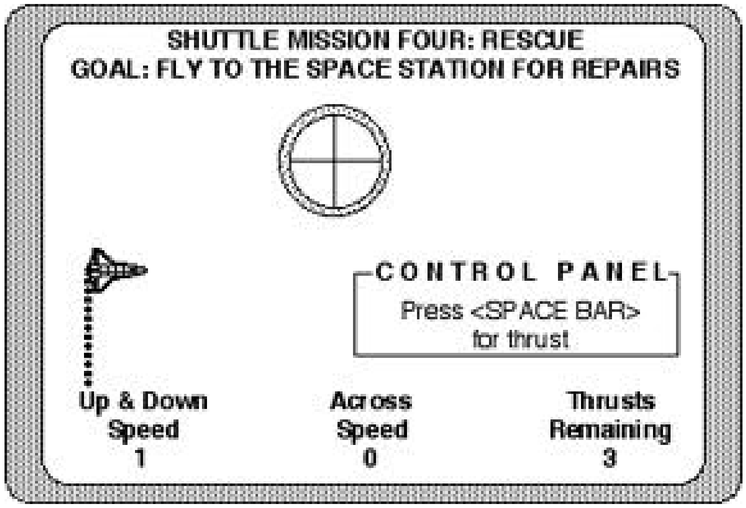

FIGURE 8.11
A representation of the computer screen during "Mission 4: Rescue." The goal of the mission is to maneuver a "disabled" shuttle to the space station. Students do not have control over the shuttle's rotation, and they only have enough fuel for three bursts of thrust. These constraints make the relationship of orthogonal forces (those occurring in 90 degree increments) more apparent.

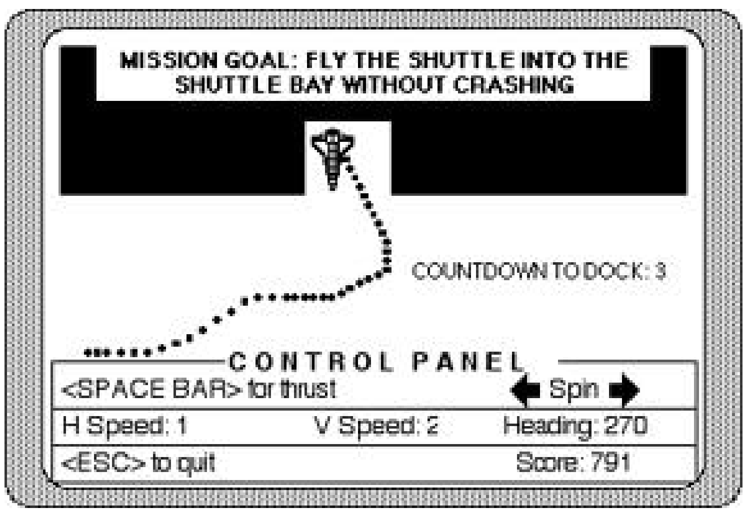

FIGURE 8.12
A representation of the computer screen during "Mission 5: Space Dock." The goal of this mission is to maneuver the shuttle to the shuttle "bay" of the space station without touching the walls of the bay or the station. This mission, even in its simplest form, is considerably more difficult than all of the other missions.

Graphics support many aspects of SSC. They induce the fantasy of piloting the space shuttle and also help explain the scientific concepts and principles. Perhaps most important, graphics are used as a critical source of continuous feedback to learners as they complete the missions. Compare, for example, Figure 8.8 with Figure 8.13, a hypothetical example where all of the graphics are replaced with pure verbal information. All of the raw information given to the learner is the same in both cases, yet the differences are stark and obvious. This intuitive graphical feedback is a natural mapping between the physics of SSC and the user.

Instructional Design Recommendations Rooted in Constructivism

Table 8.2 summarizes a series of considerations that guided the design and development of SSC. These guidelines have both instructivist and constructivist influences and, as such, are offered as a working compromise between these learning philosophies. However, more important, these guidelines also offer a means of understanding and incorporating constructivist goals in instruction. These guidelines are meant to complement the fourteen design recommendations discussed in chapter 7.

Table 8.2 Some considerations in the design of interactive learning environments based on characteristics of microworlds, simulations, and games
• Provide a meaningful learning context that supports intrinsically motivating and self-regulated learning. • Establish a pattern where the learner goes from the "known to the unknown." • Emphasize the usefulness of errors. • Provide a balance between deductive and inductive learning. • Anticipate and nurture incidental learning.

Reprinted from Rieber, L.P. (1992). Computer-based microworlds: A bridge between constructivism and direct instruction. Educational Technology Research and Development, 40(1), 93-106. Copyright 1992 by the Association for Educational Communications and Technology. Reprinted by permission of the publisher.

1. “*Provide a meaningful learning context that supports intrinsically motivating and self-regulated learning*“ (Rieber, 1992, p. 98). For learning to be a meaningful experience, it must directly connect to the learner's life in some way. The best examples of meaningful learning are anchored in contexts that are relevant and interesting for learners (Cognition and Technology Group at Vanderbilt, 1990, 1992). My experience has been that children and adults find it easy and fun to imagine they are astronauts. The context of space travel lends itself very well to both real-life drama and science fiction. In SSC, students are not watching someone else's experience in space, such as a video documentary about space travel. Instead, they can take their own imaginary trip aboard the shuttle through realistic fantasy. Students are not just along for the ride, they are in control.

The purpose of the missions is to induce and encourage the fantasy while helping the student to build awareness and competency in many science concepts and principles in a series of structured activities. Research shows initial evidence that the SSC missions hold intrinsically motivating appeal for elementary school children (Rieber, 1991b). Another example (and more widespread) of using a meaningful context for learning is the *Voyage of the Mimi*, a math and science curriculum set in the context of whale exploration (Bank Street College of Education, 1989; a sequel is based on Mayan archaeology).

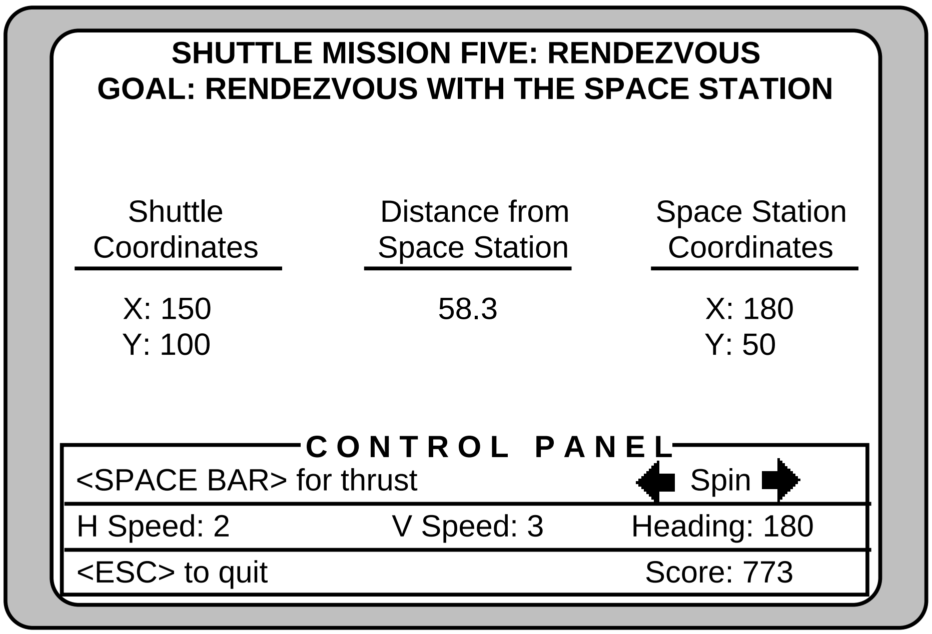

FIGURE 8.13

A hypothetical example of designing "Mission 5: Rendezvous" without the use of the graphical, animated feedback. This screen presents all of the same information presented in Figure 8.5, but in all verbal form. In order to interpret this feedback, the learner would have to mentally construct the corresponding visual elements.

Malone (1981) concludes that an activity needs to continually challenge students in order to maintain its intrinsic appeal. Several of the SSC missions, for example, provide students with the opportunity to vary the "mission conditions" (such as target size, target location, and shuttle rotation) in order to increase the difficulty of the activity, as shown in Figure 8.14. For example, it is much simpler to control the shuttle in two-dimensional space when the shuttle's rotation is kept to 90-degree increments. Control of the shuttle becomes much more difficult when control is changed to 45- or 30-degree increments. Scorekeeping features, another feature noted by Malone (1981) as a way to increase intrinsic appeal in computer games, also are provided.

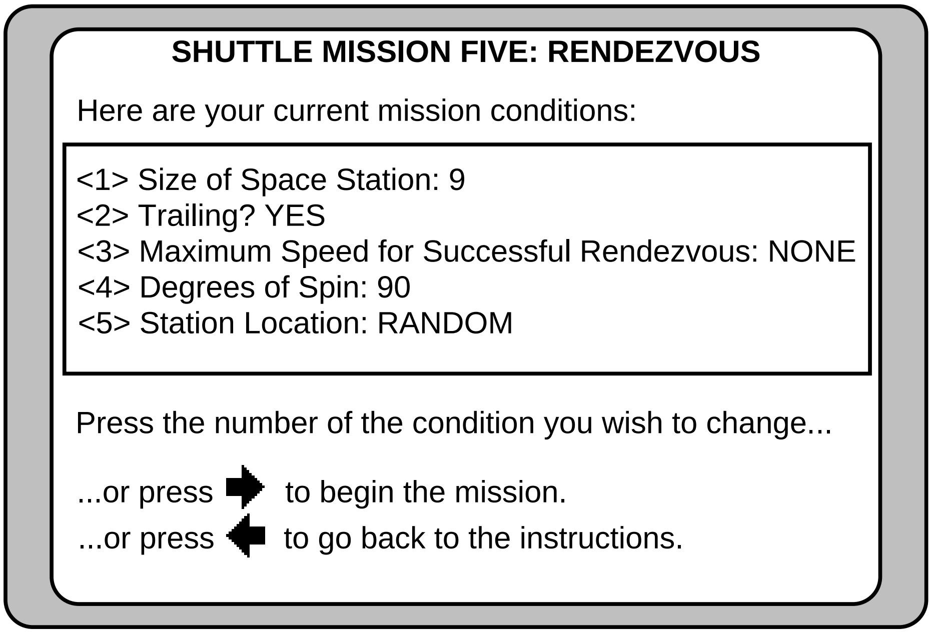

FIGURE 8.12
Sample screen showing the various mission "conditions" of Rendezvous that can be changed to increase the difficulty of the mission. A similar set of mission conditions is available for "Mission 5: Space Dock."

2. "*Establish a pattern where the learner goes from the "known to the unknown*" (Rieber, 1992, p. 100). Although research continually shows the importance of the learning context, meaningfulness can also be interpreted as the degree to which students can link new ideas to what they already know. In fact, the strength of the relationship between new information and prior knowledge may be among the most important determinants of learning (Ausubel, 1968). Many learning theorists have offered strategies for maximizing this relationship. Bruner (1966), for example, suggests a spiral approach where the simplest and most general ideas are introduced first to learners in highly interactive and concrete ways. These ideas are then reintroduced to students over and over at increasing levels of abstraction and detail. Similar models have been promoted where the most general ideas grasped early by learners are critical in helping them to comprehend much more detailed ideas introduced later (Ausubel, 1963; Reigeluth & Stein, 1983).

   SSC tries to provide students with a strong conceptual understanding of Newtonian principles in order to act as "anchoring posts" for later instruction. SSC can help bridge the learning of formal physics as it is usually taught in schools with the experiential learning of the dynaturtle. Paradoxically, traditional physics instruction usually tries to simplify learning by conveniently removing the influences of friction and gravity from the mathematical equations. Unfortunately, life without fraction and gravity is outside the experience of virtually every student. SSC can give students these experiences in a simulated context. The simulated space shuttle, like its cousin the dynaturtle, acts as a transitional object between the learner and Newtonian physics, thereby acting as an “object to think with” — a characteristic of microworlds discussed earlier.

3. “*Provide a balance between deductive and inductive learning*“ (Rieber, 1992, p. 101). Obviously, learning entails and requires both deductive and inductive approaches. An inexperienced homeowner may fix a leaky faucet by getting involved in the project outright or by consulting and carefully following a “how to” book. Extreme interpretations of either a deductive or inductive approach are obvious. Strict deductive approaches are prone to assigning a passive role to the learner. Instructional designs based on cookbook strategies usually lack both imagination and innovation. Lesson activities begin to resemble one another. Deductive approaches are much easier to apply for lower-level learning outcomes, such as fact learning, because there is little need for interpretation or “construction” on the part of the learner. On the other hand, strict inductive approaches are based on a “sink-or-swim” philosophy. Learners are at risk of becoming either frustrated or bored if they are unsuccessful or disinterested early on. Novices to a domain or content may also need structure or guidance that purely inductive experiences do not provide. Designers of hypertext, for example, report that novice learners are often prone to disorientation (Jonassen, 1986; Tripp & Roby, 1990). Inductive activities also require a playful attitude and a willingness to go exploring, conditions that older children and adults may resist (Seaman & Fellenz, 1989). The most successful learning environments carefully combine the strengths of direct instructional methods with some level of personal discovery and exploration (Sfondilias & Siegel, 1990). Though balancing deductive and inductive learning is not a simple task, it is an achievable and worthwhile goal and was inherent in the design of SSC.

4. “*Emphasize the usefulness of errors*“ (Rieber, 1992, p. 101). What are your memories of learning in school? For many people, it was a time of simply trying to find out the right answer while trying to avoid getting the wrong answer. Too much of a student's job is in just figuring out what the teacher already knows. Errors imply failure. This is unfortunate. Most of the learning in life that really counts is not only in discovering things that are totally new and original, but are also all at once challenging and complex. Rather than destructive, errors are essential for learning and are among the most instructive sources of information when an individual is engaged in problem solving (Fredericksen, 1984; Schimmel, 1988). Inductive learning theories promote the ability of learners to detect errors and then incorporate the information learned in subsequent trials. Error handling is a systematic process usually referred to as **debugging** in computer applications. Some of the most fundamental learning, such as concept formation, requires a student to isolate and control one or more variables while holding all other variables constant (see Mayer, 1983, for a review). In this way, a learner performs a series of “mini-experiments” to test a hypothesis. An example is someone installing a new light fixture who must first determine which breaker controls the electricity to that part of the house. A person's hypothesis testing may include the strategy of turning on lights or

appliances throughout the house to see which lose power when each breaker is tripped in isolation.

Unfortunately, many learning tasks contain so many individual variables that a novice would soon be inundated with information and become frustrated, such as learning how to use a new software package given only a cryptic reference manual. Microworlds offer a way to structure a learning experience so that only a limited number of variables are introduced at a given time in a context that is relevant and meaningful. This is sometimes referred to as the **project approach**. For example, it might help someone learning a word processor for the first time if a project with constrained boundaries are given, but with real and meaningful goals, such as writing a letter to a friend. When initial skills are mastered (such as entering, correcting, and saving text), additional variables can be introduced (such as how to underline, center, or print). In many physics problems, such as classical mechanics, problems are simplified when set in the context of one rather than two dimensions. The microworld of the SSC missions accomplishes this by turning on or off certain computer commands, depending on the mission. Some of the missions limit the rotation of the shuttle to 180 degrees, resulting in one-dimensional motion. Other missions begin by having the user rotate the shuttle in 90-degree increments, though the user can change this later if desired. This assures that the user first experiences how to maneuver the shuttle in two dimensions with the simplest case.

For errors to be useful, the goal of an activity must be clearly known to learners (Norman, 1988). If the goal is ambiguous, then all available feedback will be ambiguous as well. Again, the project approach is a useful strategy. In LOGO, for example, students often work toward completing a graphic they have designed beforehand, such as a house or a car. All of the graphical feedback given to them by the turtle is continually judged against their predetermined goals. This is referred to as goal monitoring, which, in the best of cases, is automatic and intrinsic (Schunk, 1990). The best advice is to provide the simplest and clearest goals when designing microworlds. Common examples of mission goals in SSC are “make the shuttle come to a stop” or “fly the shuttle to the space station.”

Errors are one important type of feedback that can take many forms, such as graphical, verbal, tactile, and aural. A variety of feedback types can also be used within a microworld so long as it all helps a learner to be successful in hypothesis forming and testing. For example, the user can choose to have the shuttle in SSC leave a trail as it moves. Verbal feedback, such as the information about the shuttle's speed, position, and heading, can also be presented along with the visual feedback. This kind of verbal feedback can be particularly helpful when students have difficulty seeing slight visual changes to the shuttle's speed or direction.

5. “*Anticipate and nurture incidental learning*“ (Rieber, 1992, p. 102). Constructivists recognize that learning rarely follows a fixed sequence that is the same for all learners. Indeed, some of the most worthwhile learning will not be anticipated. History is full of examples of scientific discoveries and technological innovations

that are the result of creative insight and unplanned tinkering. The trick is not only being in the right place at the right time, but also recognizing when one is confronted with a unique opportunity. In LOGO programming, for example, mistakes and other unintended or unexpected events often lead to interesting visual effects. Students often choose to revise (or abandon) the original programming project to pursue the unexpected results. Of course, it can be difficult to identify and document achievement in terms of unanticipated learning goals or competencies. There is also the risk that incidental learning will be irrelevant and trivial.

Instructivist approaches, on the other hand, try to take a group of learners through an instructional sequence designed to meet predetermined learning objectives. Learners are actually discouraged from exploring anything incidental to these objectives. Carefully designed microworlds help to balance the risks and incentives associated with both intentional and incidental learning. Incidental learning is expected and hoped to occur, but within design parameters. The teacher plays a critical role here. There is a need to channel incidental learning back toward the lesson objectives or to revise lesson objectives to accept unexpected learning outcomes. Of course, students who become sidetracked in unproductive ways should be redirected back to a relevant path. The teacher needs to act as the “safety valve” to make sure that a learner's actions are not counterproductive without thwarting or quelling potentially worthwhile incidental learning activities.

One of the studies discussed in chapter 6 (i.e., Rieber, 1991b) aptly shows both the potentials and risks of incidental learning. Recall that fourth-grade students successfully extracted incidental information about Newton's second law from an animated display, but also inappropriately applied this information to other contexts involving the law of gravity. Incidental learning must be carefully monitored and assessed so that it remains constructive and applicable to at least the broadest set of learning goals without contributing to misconceptions.

REVIEW

- At present, there are two divergent cognitive interpretations of instructional technology. One is associated with constructivism, and the other is associated with modifications to instructional systems development (ISD) (termed “instructivism”).
- Constructivists consider learning to be individual constructions of knowledge, which can be explained through the Piagetian process of equilibration — the enabling mechanisms of which are assimilation and accommodation.
- Learning is a natural consequence of an individual's interaction with and adaptation to an ever-changing environment.
- Externally provided or induced constructivistic learning environments are commonly called “microworlds.”
- Potentially powerful instructional designs can be effectively modeled on a blend of characteristics of microworlds, simulations, and games.
- The five design recommendations provided in this chapter are meant as a practical attempt to merge the advantages and strengths of several learning philosophies.

CHAPTER 9

Multimedia

OVERVIEW

This final chapter briefly considers the relationship among computers, graphics, and learning to the budding area of multimedia. Multimedia, and the closely related area of interactive video, is described in relation to design issues surrounding "hyper-" environments, such as hypertext and hypermedia. Some design considerations based on constructivism are also revisited.

OBJECTIVES

Comprehension

1. After reading this chapter, you should be able to:
2. Define multimedia, hypermedia, and interactive video.
3. Describe how arguments and issues related to constructivism, instructional design, and learner control may extend to multimedia learning environments.
4. Describe five levels of interactivity associated with interactive video.
5. Summarize some of the research associated with multimedia, interactive video, and hypermedia.

Application

After reading this chapter, you should be able to:

1. Apply the theory, research, and design principles considered in previous chapters to multimedia.

It should be clear by now that, if used appropriately, graphics can be an important part of learning and instruction. We have critically analyzed functions served by graphics in learning and instruction, especially when designed, developed, and delivered by computer. As computer technologies continue to flourish in education and as computers increase in graphical ability, there will be a strong need for designers to remain in control of how these technologies will be applied to learning environments. History has shown that the temptation to have machine technologies drive instructional design is difficult to resist. There will always be the tendency by some to believe that they have the right to abandon and ignore established knowledge bases as they apply new technologies. There will also be others who feel a need to maintain the status quo. This book has presented arguments suggesting that current knowledge bases of instructional design derived by theory, research, and experience are crucial elements to consider. Yet, there should be no mistaking the need to extend and elaborate these knowledge bases to take into account new philosophies and techniques. Additionally, none of these knowledge bases negates the need for designers to be inventive, creative, and willing to take risks. Increasing instructional wisdom in light of emerging technologies has been slow and gradual. One foot should remain in what is known and understood (e.g., available theory and research), while the other foot carefully explores uncharted areas of design.

The purpose of this last chapter is to briefly discuss several computer applications that should be particularly relevant for those interested in instructional visualization in the years to come. The future of visually based computer technologies is at once exciting, inspiring, and intimidating. Fortunately, the case can be made that the principles covered in this book can offer the best offense and defense for tapping the potential of future technologies while avoiding being swept away by the flood of options and considerations. One of the areas we will consider is the broadly defined area of multimedia. First, however, we will revisit and put relative closure on a topic from the last chapter — constructivism — which has much to do with many issues related to multimedia, hypermedia, and interactive video.

CONSTRUCTIVISM REVISITED

As described and discussed in the previous chapter, constructivism asserts that learning is a continuous and never-ending process of building and reshaping mental structures. Knowledge cannot be imposed on an individual; rather, knowledge is itself constructed by each person. The inherent antagonism between direct instruction and constructivism was not meant to be resolved in the previous chapter. Instead, a working compromise was offered, which may allow instructional designers to tap the strengths of direct instructional methods and constructivism in designing learning environments. These learning environments share attributes of gaming, simulations, and microworlds.

The previous chapter asked designers to consider how to merge features of gaming, simulations, and microworlds as *they* construct learning environments *for* students. Of course, these learning environments should, at the very least, be flexible enough for learners to be able to appropriately alter the environment to match their abilities and interests. *Space Shuttle Commander* was offered as one example of how one might design computer software according to this confluence of instructivist and constructivist views. Radical constructivists, on the other hand, seek ways to have *learners* construct their own learning environments. Rather than consider if SSC is suitable for a group of learners, a radical constructivist would probably prefer to focus on how SSC's designer learned physics as a result of building it. Similarly, LOGO was not meant for designers to develop structured lessons, but rather for students to use as a kind of mathematical “erector set.”

The same debate can be initiated for any authoring tool, graphical or otherwise. On one hand, graphical software packages are viewed as tools or resources for instructional designers to use as they develop instruction for learners. On the other hand, these same packages should be considered as graphical tools and resources for *students* to construct their own learning materials and experiences. These issues are not foreign to instructional designers and developers. Cognitive orientations to instructional design frequently call for **generative learning strategies** (Wittrock, 1974, 1978), where learners are asked to deliberately take action to create meaning from material. Rather than viewing students as passive agents who “receive” instruction, generative learning assumes and requires learners to be active participants in their own learning. The generative learning hypothesis creates a learning “partnership” between the instruction and the learner. Learners are given much authority and responsibility for their learning, but are guided by and through instruction.

Simple examples of generative learning activities include underlining meaningful parts of text, note-taking strategies, paraphrasing, and outlining textual passages. Other activities can be more elaborate, such as student- generated questions and the creation of mnemonic learning aids and concept maps. Any activity can promote generative learning in which learners are required to “consciously and intentionally . . . relate new information to their existing knowledge rather than responding to material without using personal, contextual knowledge” (Jonassen, 1988b, p. 154). Instead of presenting students with ready-made representational, analogical, or arbitrary graphics, a generative approach would ask students to create their own graphics, with guidance, for the same purpose: to clarify relationships, and to facilitate understanding, establish meaning, and promote motivation.

Experience shows that there are times when learners need (and perhaps want) some imposed structure and times when learners should be given more freedom and responsibility to direct and design their own learning paths. Research on learner control in CBI is inconsistent (Milheim & Martin, 1991). One pool of research generally indicates that total learner control of computer-based instruction is usually not advisable unless paired with some sort of coaching or advisement strategy (see review by Steinberg, 1989). Yet, other research indicates that learner control is an important characteristic of successful instruction (Kinzie, Sullivan, & Berdel, 1988). It seems that a full understanding of learner control must simultaneously take into account performance and motivation variables. Related theories of motivation and attribution suggest that learners should be provided with some level of control over the selection, sequence, and pacing of content in order to reinforce the belief that they personally control their own success (Milheim & Martin, 1991). Obviously, the issue of learner control must be based on a combination of perspectives — some cognitive and some motivational.

Figure 9.1 illustrates an example of an instructional computer activity that attempts to balance these issues. The activity's main cognitive objective is to help a learning-disabled student develop a working sight-word vocabulary. A secondary cognitive objective is to help the student with shape recognition. Both cognitive objectives are set in a highly motivating graphical context. When any of the words on the right is selected, the computer pronounces the word and displays a simple graphic of the word comprised of a combination of the three basic shapes. The student can then fill in the graphic by moving shapes onto the graphic's outline. The student is free to choose any of the words on the right at any time or to just doodle. Words can be added to or deleted from the list. This activity provides a variety of sensory inputs, including tactile, to help the student associate the written and spoken word. The activity, by nature of the available words, is guided, yet the student is free to explore and make individual choices. Also, the student can choose at any time to “clean up” the shapes and start over. The computer animates all the shapes back to their starting positions, which has proven to be a very motivating feature.

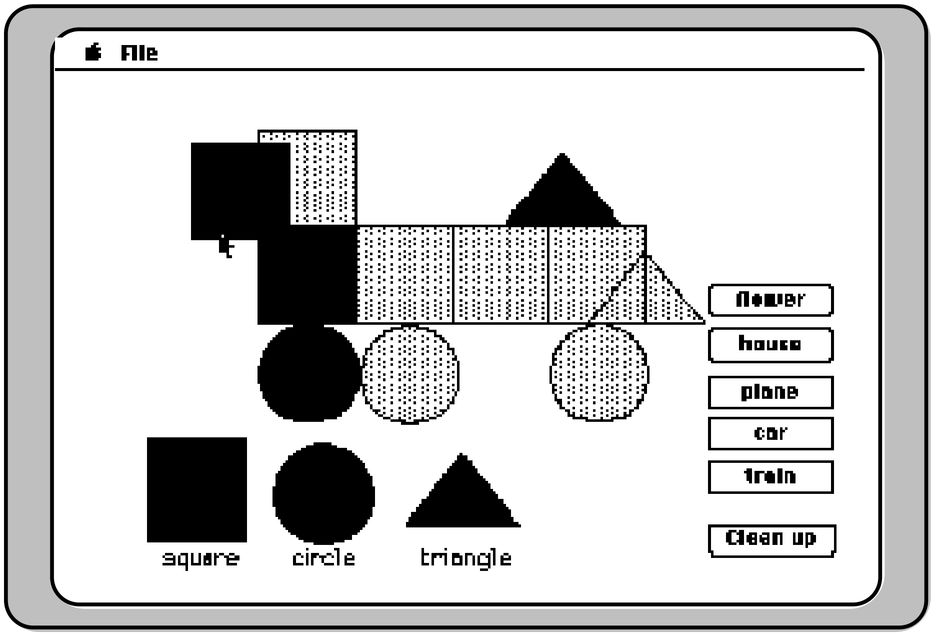

FIGURE 9.1
An example of a computer activity that balances the guidance of instruction with a student's purposeful construction of ideas and concepts.

The decision of when to use a more instructivistic or constructivistic orientation will depend largely on the interplay of the learner's experience or background in a particular domain and the learner's ability to self-regulate his or her learning in the domain. Novices will be especially prone to disorientation and confusion if left without guidance. Conversely, as students become more experienced and confident in a domain, they may become more resistant to imposed control on what they learn, how they learn, and when they learn. These same arguments are currently being played out in the case of multimedia (Locatis, Letourneau, & Banvard, 1989).

MULTIMEDIA

Multimedia is one of the latest buzzwords in educational technology. As such, its meaning has been stretched to fit almost any situation in which a variety of educational media are used. In its most general sense, it refers to any instructional delivery system that includes two or more media components, such as print-based, computer-based, and video-based. A traditional instructional setting combining lecture with a slide/tape presentation could be considered multimedia, for example. In its more common usage, it refers to integrated instructional systems that deliver a wide range of visual and verbal stimuli, usually through or in tandem with computer-based technologies. Although the computer is not necessarily a prerequisite component of multimedia, it is usually the focus of the instructional system. The most common multimedia systems are highly interactive computer-managed video/audio systems. Figure 9.2 illustrates a high-end work station to *produce* multimedia, and Figure 9.3 shows a system design for the *delivery* of multimedia. Ambron and Hooper (1990) define interactive multimedia as “a collection of computer-centered technologies that give a user the capability to access and manipulate text, sounds, and images” (p. xi). Although this book is not about multimedia, per se, it can be argued that the concepts and principles that have been discussed are directly relevant to multimedia. The principles of instructional visualization on which we have focused must be seriously considered in educational applications of multimedia.

Multimedia is not a new concept. However, enthusiasm for multimedia has grown as manufacturers rapidly expand computer hardware to use, integrate, and standardize video and audio formats in their systems. But the greatest enthusiasm for multimedia is probably due to the ease with which text, graphics, sounds, and video can be incorporated and accessed in instructional systems. As discussed briefly in chapter 3, the range of video and audio capabilities of desktop computers is evolving at a tremendous rate. For better or for worse, there seems to be a commitment among hardware manufacturers to conquer the sizable memory and processing hurdles inherent with making video and audio an integrated part of the desktop computer. The current proliferation of compact disc (CD) technology is a case in point.

Much of the enthusiasm for multimedia is centered on hardware. Proponents of multimedia in education usually refer to the old argument that merely increasing the external modes of delivery will result in increases in learning. Other proponents use surface-level arguments that students learn best when given a great variety of stimuli and instructional strategies. While there has been some initial work done to shift multimedia research from media to psychology (see Nix & Spiro, 1990, for example), rarely do the most popular media arguments extend beyond novelty effects. Although there is cause for enthusiasm, given the increasing number of options available to the instructional developer by computer-based multimedia advances, many of these arguments are in danger of falling into the “technocentric design” traps discussed in chapter 1. Enthusiasts also risk the unfortunate mindset that the past 50 years of experience (both successes and failures) with educational media do not apply to multimedia. There is also the curious dilemma of the hardware evolving faster than instructional designers, developers, and researchers are able to test and apply the resulting applications. Unfortunately, this pattern of forgetting the old while not being able to keep up with the new has been often repeated. There is also the continual danger (and paradox) that software designers will not be able to have direct influence on hardware advances.

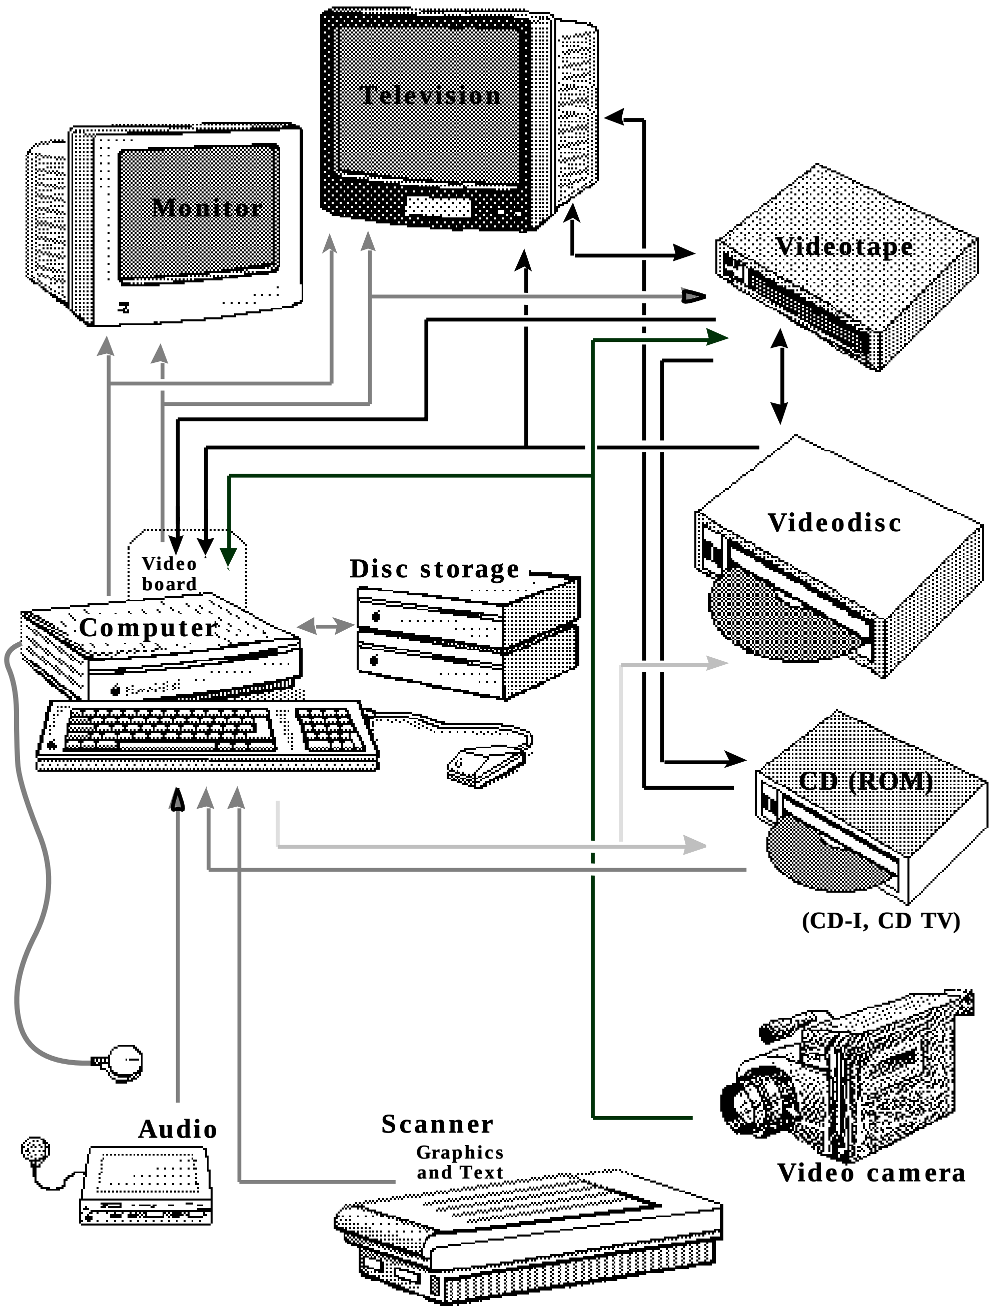

FIGURE 9.2
A sample computer configuration to *produce* instructional multimedia material. Copyright 1992 by R.D. Zellner and reprinted with permission.

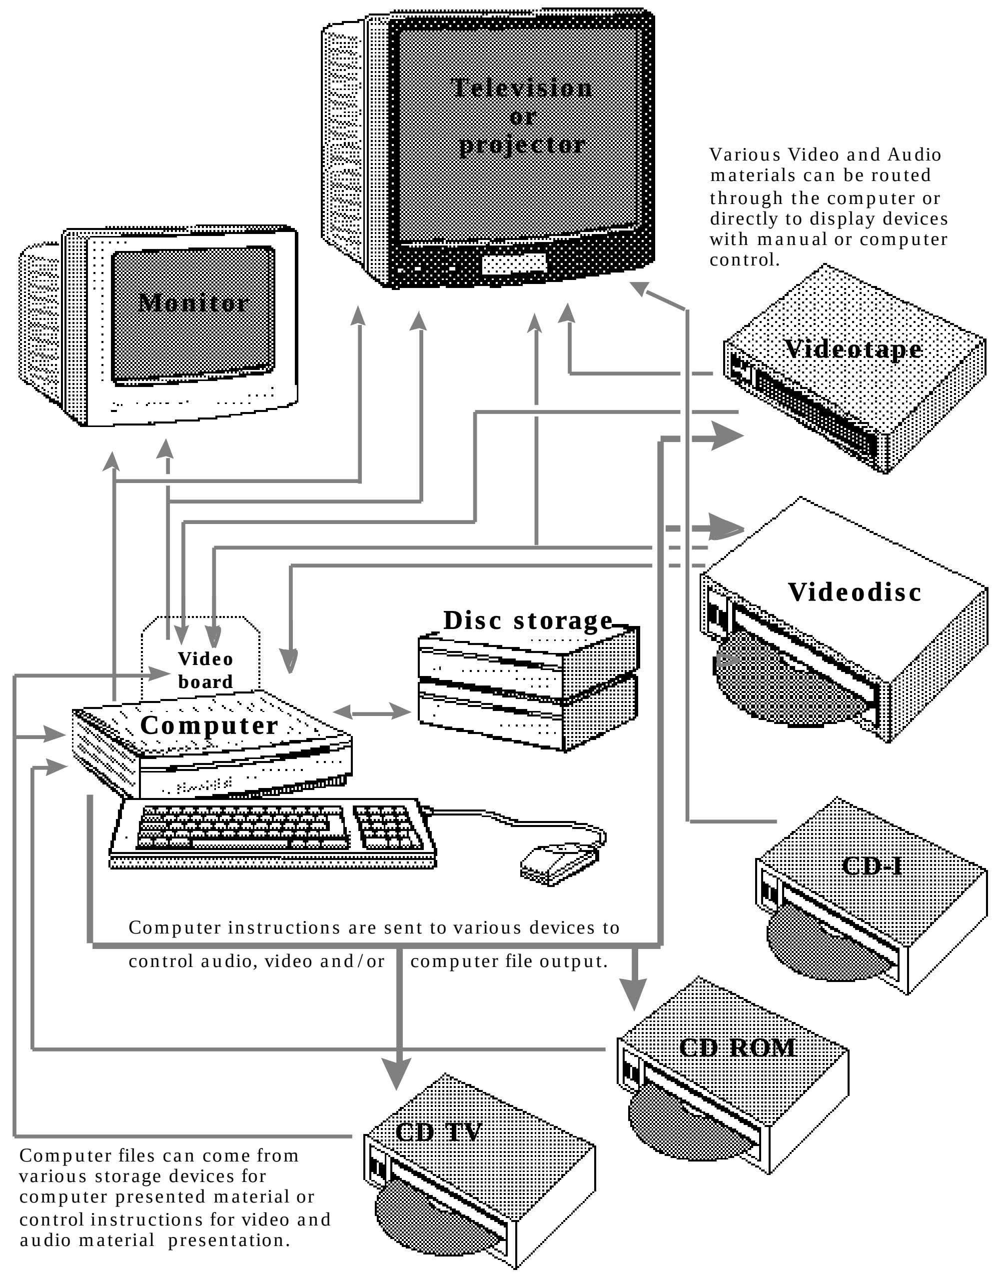

FIGURE 9.3
A sample computer configuration needed to *deliver* instructional multimedia material. Copyright 1992 by R.D. Zellner and reprinted with permission.

Multimedia and Hypermedia

Interestingly, most current discussions of multimedia are linked with the development of **hypertext** tools, such as Apple's *HyperCard* and IBM's *Linkway* and *ToolBook*. The term “hyper” translates simply as “link” and has been extended to include hypermedia environments, or systems that link various media, such as computer and video (Locatis, Letourneau, & Banvard, 1989). The origin of hypertext is usually traced to an article by Vannevar Bush (1945), then president of the Massachusetts Institute of Technology, entitled *As We May Think* (the term “hypertext” was actually coined by Ted Nelson in the 1960s). Bush described a hypothetical device, called the Memex, which would allow people to explore ideas in nonlinear ways. The technology has only recently been able to catch up with Bush's original ideas.

The philosophy behind hypertext and hypermedia environments is that informational and instructional systems can be built to allow users to explore knowledge bases in ways that may mirror how people actually think. In contrast to CBI, founded on some external instructional design model, hypertext is meant to allow users to create their own knowledge representations in a particular domain. Hypermedia proponents argue that human thinking is not linear, so, therefore, users should be able to explore informational systems by selecting and sequencing their own paths in a domain.

There is relative support for hypertext from the field of cognitive science. Hypertext environments seem to closely conform to the idea of propositional networks discussed in chapter 4, which suggest that cognition involves an ever-transforming network of nodes and links. Nodes represent one of many different kinds of informational units, and links represent how the nodes are related or associated. As the basic unit of information, nodes can be represented verbally or visually. Therefore, a hypermedia environment would allow for the knowledge base to be represented through a variety of stimuli, including text, graphics, video, and sound. In addition, users would be able to add or reconfigure the hypermedia environment, thus promoting an interactive and dynamic system. However, the cognitive power of hypermedia is derived much more from the links than the nodes. Higher-order applications of human memory, such as problem-solving, are believed to be a function of the strength (in terms of meaning) of the association between informational units.

Hypermedia and CBI go in seemingly opposite design directions. In CBI, authors and designers make decisions on how information will be related, and that representation is subsequently imposed on learners. Proponents of hypermedia argue that since there may be as many representations of a knowledge base as there are learners, one interpretation given by the author is more or less arbitrary. Therefore, they suggest it is better to allow users to make their own associations. However, cognitive interpretations of instructional design show many similarities with principles of hypermedia (Jonassen, 1991b).

The serious research on hypertext has only begun. Considering the high expectations, most current reports have been discouraging for proponents (an often-repeated pattern for new educational media innovations). It appears that without guidance, novices have a difficult time knowing how to explore a hypertext environment and often become disoriented (Tripp

& Roby, 1990; Jonassen, 1988c). Novices have limited cues or strategies for how to allocate their limited attentional resources. Also, as users allocate cognitive processing to certain tasks, there is the risk that their performance on other potentially rewarding tasks will deteriorate. Based on related research on human cognition, it is likely that hypertext environments would be more facilitative as a user becomes more familiar with a content area or domain. Hypertext environments are probably not good instructional systems for introducing novices to an area, but may be good environments in which users can subsequently organize and integrate information. It may be that users simply are not accustomed to the nature (and responsibility) of having total navigational control within a knowledge base. Simply providing an environment that allows users to customize knowledge for themselves does not necessarily mean that they will be able to do so.

Applied properly, hypermedia environments have much potential in education. However, it is clear that significant learning will not occur simply by haphazardly planting these environments in educational settings. There is surely a need for *both* structured and unstructured learning environments in training and education. The potentials of hypermedia closely parallel the arguments calling for a balance between deductive and inductive learning strategies that were presented in the previous chapter in the context of simulations, games, and microworlds.

Multimedia and Interactive Video

Interactive video is the most longstanding application area that most closely resembles current interpretations of multimedia. Interactive video is best thought of as the marriage of computer and video technologies. It is typical for users to emphasize either the video or computer component of interactive video, such as considering it as computer-managed video or as CBI with a video component. However, many point to the increased visualization capabilities as the chief advantage of interactive video systems (Locatis, Charuhas, & Banvard, 1990), especially in regard to increasing a learner's control over the video material (Hannafin & Peck, 1988). In most systems, a computer is directly cabled, or interfaced, with a video play-back unit, most frequently a laser disc player. One side of a standard 12-inch videodisc typically contains 54,000 individual frames that can be referenced directly, allowing for 30 minutes of linear video footage (figuring 30 frames per second). However, interactive video can also include computer-controlled videotape players. Despite the utility of interactive videotape, laser disc systems are far more popular for several reasons. Laser discs allow for the random access of any frame with a typical delay of no more than about two seconds, whereas tape systems take relatively large amounts of time while the system either fast-forwards or rewinds to the proper location. Laser discs also provide unlimited playbacks, including freeze frames, with no deterioration of quality since there is no physical contact with the medium — the digital information on the disc is accessed by the reflection of a laser beam. In contrast, the read/write heads of a videotape system are in continuous contact with the magnetically charged tape, resulting in a relatively quick drop in signal quality after a series of plays. Although the physical differences between tape and disc may be offset so that no instructional differences may be experienced (e.g., Hannafin & Phillips, 1987), videodisc applications are far more common.

The growth and improvement of disc technology, as evidenced by compact discs, seem to indicate that disc technology will ride the crest of applications in the immediate future.

A taxonomy of the levels of interactive video has been offered (Daynes & Butler, 1984; Gayeski & Williams, 1985; Parsloe, 1983). Although this taxonomy is largely hardware-based, it speaks to the interactive opportunities made possible through different hardware configurations. The first two levels of the taxonomy, levels 0 and 1, includes only video technology — no computer technology is involved. At level 0 there is no *overt* interactivity between the video materials and students, such as in linear video presentations or broadcast television. There is no opportunity to interrupt the video presentation once it has begun. Level 1 interactive video includes manual interruptions of a stand-alone videodisc or videotape, usually stop/start, by either the teacher or student. Level 1 also includes manual branching and searching of segments by the teacher or student through the manual controls of the video unit.

Level 2 interactive video is the first level at which computer technology begins to play a role. Through an onboard microprocessor, a videodisc player is able to run a program encoded onto the videodisc itself. Such programs would allow for simple conditional branching based on a student's input to the videodisc's keypad. However, level 3 is the level at which the video player is interfaced to a *separate* computer. This is the level which is usually considered for mainstream interactive video or multimedia applications. Most level 3 systems have one monitor for the computer and one for the video, although some more expensive systems allow one monitor to be switched between the video and computer signal. The latest technology allows a computer monitor to include video “windows,” in which a small portion of the monitor displays video material. It is hoped that such video windows will be as easily accessed and manipulated in a graphical user interface (GUI) as text and graphic windows are now.

Level 4 usually defines the last level of interactive video (although some taxonomies include several more levels). One might refer to this as the “dare to dream” level. Typically, level 4 systems include a creative assortment of hardware, such as multiple video units, sound synthesizers, voice recognition, touch screens, etc. In light of new technologies, such as virtual reality (see chapter 8), interactive video takes on a completely new definition. One might learn about the relationship between enzymes and proteins not by simply interacting with a computer-based multimedia station, but by reaching out, grabbing, and manipulating a particular molecule or even perhaps by “becoming” the molecule.

As you can see, these levels are heavily hardware-oriented. While such a taxonomy makes it easy for newcomers to understand the system configurations, it also unfortunately promotes technocentric design, as discussed in chapter 1. While advances in hardware technologies provide wonderful opportunities for instruction and learning, it is only through the software or *idea technologies*, such as instructional design, that the potentials can be realized. For this reason, most of the serious research and developmental work in interactive video has carefully concerned the research and theory most related to the parent technologies of instructional design, CBI, and video instruction. Although there are those who consider interactive video and multimedia as completely unique technologies, a more realistic view is that interactive video and multimedia will be best appropriated for learning when based on careful analysis of learner and instructional design attributes.

Views on the effectiveness of interactive video range from highly enthusiastic (Debloois, 1982) to cautious (Hannafin, Garhart, Rieber, & Phillips, 1985). Most of the evidence for interactive video has come from developmental projects that have been largely atheoretical (Cronin & Cronin, 1992). The most credible research has investigated instructional design issues *with* interactive video, rather than studying the medium itself (see for example, Hannafin, 1985; Hannafin, 1992; and Hannafin & Hughes, 1986).

Arguments for interactive video and multimedia, apart from the interactive components of CBI, are best understood as times when video provides the best source of instructional delivery. Some rationales for video are rooted in the cognitive domain, such as the use of high-fidelity video images to demonstrate what a particular chemical reaction will look like without exposing students to highly volatile chemicals (Smith & Jones, 1991) or medical education where real-life situations can be better represented with video than text and graphics (Nashel & Martin, 1991).

One of the most compelling justifications for video may be its dramatic and immediate ability to elicit an emotional response from an individual. Such a reaction can provide a strong motivational incentive to choose and persist in a task. For example, compare the differences between hearing or reading an account of a bridge collapse and actually watching the video footage of the bridge oscillating wildly before disintegrating and crashing into the water below. (See Footnote 1)

Other examples combine cognitive and motivational elements by using video to provide a meaningful context for learning, called “anchored instruction” (Cognition and Technology Group at Vanderbilt, 1990). An example would be showing students a video segment where Indiana Jones carefully replaces a golden idol with a sandbag to prevent the booby traps from being triggered as a context for understanding the relationship between volume and weight (Bransford, Sherwood, & Hasselbring, 1986; also see descriptions of the Jasper Woodbury Problem Solving Series by the Cognition and Technology Group at Vanderbilt, 1992). Other applications point to social and language applications, such as using interactive video in bilingual training (Reed, 1991). Still others have a more constructivistic flavor, where children build their own interactive video materials to learn about science and social studies (Gerrish, 1991).

Despite the allure of interactive video and other multimedia environments, there is every reason to believe that the instructional design of these systems should be based on a careful analysis of the many interrelated and interdependent elements discussed throughout this book. These include psychological foundations of the individual, especially those related to visual learning, and the instructional design of interactive learning systems.

A FINAL WORD

This book has been but a beginning in the effort to tie together the theory, research, and practice of instructional visualization in the computer age. We have considered information from many different areas and points of view. This book was written to summarize, organize, and synthesize a wide spectrum of ideas related to computers, graphics, and learning. It is hoped that you now feel empowered, not overwhelmed. Much is written and known about these three topics when they are considered individually or in pairs, but little when taken collectively. Designers who have searched the literature for guidance have probably found themselves either with the feeling that no integrated literature is available or swamped with the idea that everything they read seems to apply. A fundamental principle of learning is that when people have too much or too little to do or consider, they seek either to "turn off" the task or look for another. In either case, the result is invariably the same — they stop trying the task at hand. It is the goal of this book to provide a compromise in both cases. For some, this book may have opened up an understanding of how many areas are relevant that were not previously considered. For others, this book may have organized the flood of available information, thereby increasing its potential to be understood and used. Ultimately, it is up to you to decide how to best apply what we know about computers, graphics, and learning in the design of instructional systems and learning environments.

Despite this book's frequently cautious tone, there is much cause for enthusiasm. Desktop computer technology gives designers and developers access to impressive graphical power. This does not diminish the role of graphic designers, artists, programmers, and technicians. Instead, it places more power into the hands of people who more directly influence the construction of learning environments. Although this book was written for instructional designers and developers, it is hoped that this book has also provided a window for other interested professionals to glimpse at the task of putting computers and graphics to use in instructional settings. We all come to instructional design with our own strengths, interests, and experiences. The challenge is to take advantage of all these diverse abilities to achieve the common goal of enhancing an individual's learning in a particular domain.

This is a good time to repeat that this book is not meant to be a substitute for a thorough introduction to the many ideas and areas that have been included here. Probably the most important of these are learning theory, instructional design, computer-based instruction, and computer graphics. Each of these areas can be considered as distinct and sophisticated fields of inquiry requiring years of study to adequately understand. However, this book has tried to bring these areas together without demanding that readers be experts in any one. An obvious next step is for you to further explore these areas.

Finally, it is recognized that many areas have not been adequately explored in this book. State-of-the-art computer visualization is a particular case in point where highly realistic (though perhaps imaginary) three-dimensional graphics are modeled, rendered, and animated on high-end computer graphics work stations (see Figure 9.4 for an intriguing example of 3-D graphics and Figure 9.5 on the next page for the solution). The development of virtual reality is another example. These have been deliberate omissions for two main reasons. First, most of these areas have not, as yet, been sufficiently applied to educational settings. For example, computer visualization has been applied most frequently in fields such as architecture, medicine, art, and commercial television. Second, few of these areas have migrated to desktop computer applications. As educational practice catches up to the potentials of computer graphics technology, this will surely change.

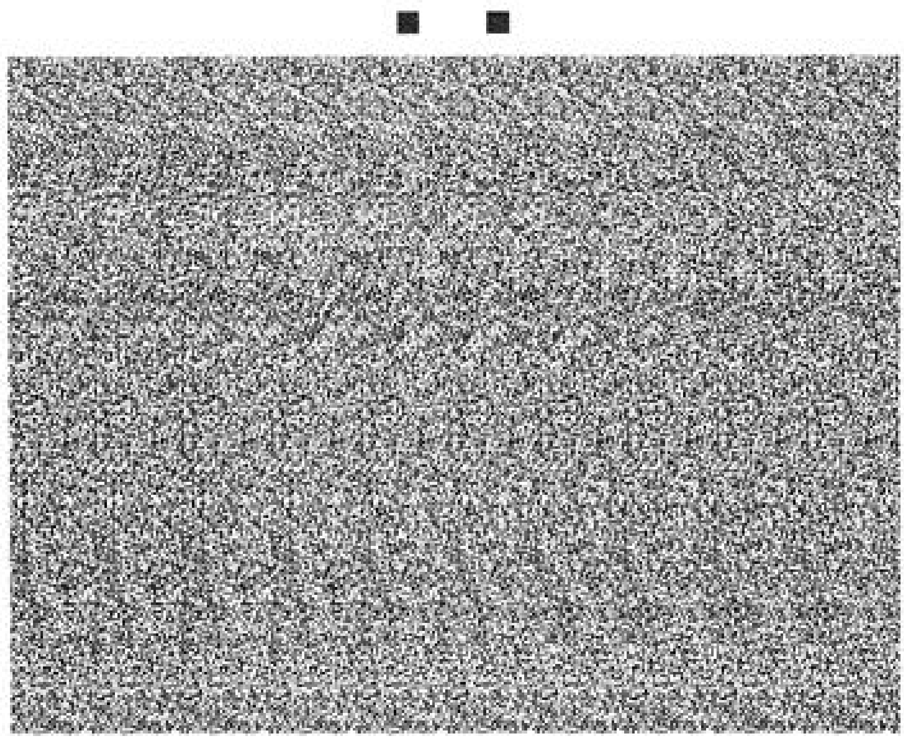

FIGURE 9.4
This figure contains a 3-D message. To see it, you need to stare through the figure as though it were a window. As you relax your eyes, the two dots on top will blend into three dots (you may need to adjust your viewing distance). Be patient because it will take some time and practice for the image to appear. Beware, not all people report seeing the image. (See Figure 9.5 at the end of the chapter for the solution.)

It is fitting that this book end with an image of what it represents. We choose to compare an instructional designer applying knowledge about computers, graphics, and learning to an explorer who wants to verify which reports of a faraway land are accurate and which are pretend, imagined, or lies. The motivation to go will be part decreed, part economic, part curiosity, and part self-satisfaction. Theory and research related to learning and graphics are like basic training about many fundamental ideas and skills, such as knowledge about survival, first aid, navigation, map reading, and how to use a compass.

Instructional design is like plotting a course and setting out on the journey. The computer and other instructional materials are like the explorer's gear, which has been deliberately chosen for the trip. Important decisions must be made, and the consequences of good and bad choices made before and during the trip must be recognized and dealt with as they are encountered. Along the way, oceans are crossed, rivers are forded, and mountains are climbed; however, each step is meant to be taken in the charted direction. At times, the explorer travels already established routes and other times must blaze a new trail. There comes the time when the explorer must determine if the intended destination or another land has been reached. It is even possible that an entirely new world will be discovered. The journey is a success only when knowledge gained from it is shared, understood, and used by others. However, unlike some actual historical examples of this metaphor, our explorer celebrates the journey, as well as the destination, without exploiting or destroying that which is encountered along the way.

REVIEW

- In contrast to instructional designers using computer graphic tools to develop instructional systems, constructivists would probably choose to focus on how *students* use these tools to represent and reconstruct knowledge.
- Although multimedia can refer to any instructional system that uses two or more media to deliver a wide range of verbal and visual stimuli to students, multimedia is usually described as a highly interactive *computer-managed* video/audio system.
- Hypermedia describes multimedia systems that are designed based on hypertext principles where “hyper-” translates as “link.”
- Interactive video is the most studied form of computer-managed multimedia systems.
- Several taxonomies describing levels of interactivity within an interactive video system, usually focusing on hardware, have been offered. Level 3 systems are those in which a separate computer is interfaced, or cabled, to a video unit (usually videodisc).

NOTES

1. Such a disaster, the collapse of the Tacoma Narrows Bridge in Washington in the 1940s, was actually captured on film. The Nebraska Videodisc Production/Design Group designed an early interactive videodisc project that described the collapse and the physical science principles explaining why it occurred.

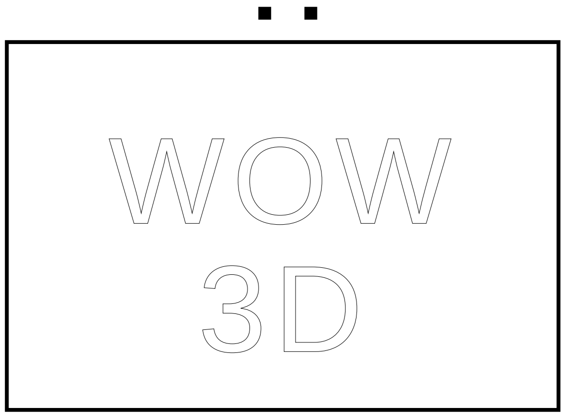

FIGURE 9.5
This is the solution to Figure 9.4.

BIBLIOGRAPHY

Abelson, H., & diSessa, A. (1981). *Turtle geometry: The computer as a medium for exploring mathematics.* Cambridge, MA: MIT Press.

Acker, S. R., & Klein, E. L. (1986). Visualizing spatial tasks: A comparison of computer graphic and full-band video displays. *Educational Communications and Technology Journal*, *34*(1), 21-30.

Ackerman, B. (1985). The effects of specific and categorical orienting on children's incidental and intentional memory for pictures and words. *Journal of Experimental Child Psychology*, *39*(2), 300-325.

Alesandrini, K. (1984). Pictures and adult learning. *Instructional Science, 13,* 63-77.

Alesandrini, K. (1987). Computer graphics in learning and instruction. In H. Houghton & D. Willows (Eds.), *The psychology of illustration, volume 2: Instructional issues* (pp. 159-188). New York: Springer-Verlag.

Alessi, S. M. (1988). Fidelity in the design of instructional simulations. *Journal of Computer-Based Instruction, 15*(2), 40-47.

Alessi, S. M., & Trollip, S. R. (1985). *Computer-based instruction: Methods and development.* Englewood Cliffs, NJ: Prentice-Hall.

Alessi, S. M., & Trollip, S. R. (1991). *Computer-based instruction: Methods and development* (2nd ed.). Englewood Cliffs, NJ: Prentice-Hall.

Ambron, S., & Hooper, K. (Eds.). (1990). *Learning with interactive multimedia: Developing and using multimedia tools in education.* Redmond, WA: Microsoft Press.

Anderson, J. (1978). Arguments concerning representations for mental imagery. *Psychological Review, 85*(4), 249-277.

Anderson, J. R. (1980). *Cognitive psychology and its implications.* San Francisco, CA: W.H. Freeman.

Anderson, R., & Biddle, W. (1975). On asking people questions about what they are reading. In G. Bower (Ed.), *Psychology of learning and motivation (Vol. 9)* (pp. 89-132). New York: Academic Press.

Anderson, R., Kulhavy, R., & Andre, T. (1971). Feedback procedures in programmed instruction. *Journal of Educational Psychology, 62,* 148-156.

Anderson, R., Kulhavy, R., & Andre, T. (1972). Conditions under which feedback facilitates learning from programmed lessons. *Journal of Educational Psychology, 63,* 186-188.

Andrews, D., & Goodson, L. (1980). A comparative analysis of models of instructional design. *Journal of Instructional Development, 3*(4), 2-16.

Anglin, G. (1986). Prose-relevant pictures and older learners' recall of written prose. *Educational Communications and Technology Journal, 34*(3), 131-136.

Anglin, G. J. (1987). Effect of pictures on recall of written prose: How durable are picture effects? *Educational Communications and Technology Journal, 35* (1), 25-30.

Artwick, B. (1985). *Microcomputer displays, graphics, and animation.* Englewood Cliffs, NJ: Prentice-Hall.

Atkinson, M., & Burton, J. (1991). Measuring the effectiveness of a microcomputer simulation. *Journal of Computer-Based Instruction, 18* (2), 63-65.

Ausubel, D. (1963). *The psychology of meaningful verbal learning.* New York: Grune & Stratton.

Ausubel, D. (1968). *Educational psychology: A cognitive view.* New York: Holt, Rinehart & Winston.

Ausubel, D. (1978). In defense of advance organizers: A reply to the critics. *Review of Educational Research, 48,* 251-257.

Baek, Y., & Layne, B. (1988). Color, graphics, and animation in a computer-assisted learning tutorial lesson. *Journal of Computer-Based Instruction, 15* (4), 131-135.

Bajer, C. L., Jr., Baydala, A., & Zeitouni, N. (1989). Optimal displacement in apparent motion. *Vision Research, 29*(7), 849-859.

Banathy, B. (1987). Instructional systems design. In R. Gagné (Ed.), *Instructional Technology: Foundations* (pp. 85-112). Hillsdale, NJ: Lawrence Erlbaum Associates.

Bangert-Drowns, R., Kulik, J., & Kulik, C.-L. C. (1985). Effectiveness of computer-based education in secondary schools, *Journal of Computer-Based Instruction, 12* (3), 59-68.

Bangert-Drowns, R. L., Kulik, C. C., Kulik, J. A., & Morgan, M. (1991). The instructional effect of feedback in test-like events. *Review of Educational Research, 61,* 213-238.

Bank Street College of Education, (1989). *The Voyage of the Mimi* [Computer Software]. Pleasantville, NY: Sunburst Communications.

Barnes, B., & Clawson, E. (1975). Do advance organizers facilitate learning? Recommendations for further research based on an analysis of 32 studies. *Review of Educational Research, 45,* 637-659.
Bell, M. (1981). Information-processing theory: The missing component in instructional design. *Journal of Instructional Psychology, 8* (4), 165-169.
Ben-Chaim, D., Lappan, G., & Houang, R. T. (1988). The effect of instruction on spatial visualization skills of middle school boys and girls. *American Educational Research Journal, 25,* 51-71.
Berlyne, D. (1965). *Structure and direction in thinking.* New York: John Wiley.
Bertin, J. (1983). *Semiology of graphics.* Madison, WI: The University of Wisconsin Press.
Blake, T. (1977). Motion in instructional media: Some subject-display mode interactions. *Perceptual and Motor Skills, 44,* 975-985.
Bloom, B. (Ed.). (1956). *Taxonomy of educational objectives.* Handbook I: Cognitive Domain. New York: Longman.
Bonner, J. (1988). Implications of cognitive theory for instructional design: Revisited. *Educational Communications and Technology Journal, 36* (1), 3-14.
Borg, W. R., & Gall, M. D. (1989). *Educational research: An introduction* (5th ed.). White Plains, NY: Longman.
Bork, A. (1981). *Learning with computers.* Billerica, MA: Digital Press.
Borkowski, J. G., Carr, M., Rellinger, E., & Pressley, M. (1990). Self-regulated cognition: Interdependence of metacognition, attributions, and self-esteem. In B. F. Jones & L. Idol (Eds.), *Dimensions of thinking and cognitive instruction* (pp. 53-92). Hillsdale, NJ: Lawrence Erlbaum Associates.
Boulton, J. C., & Hess, R. F. (1989). Luminance contrast and motion detection. *Vision Research, 30* (1), 175-179.
Bower, G. (1970). Analysis of a mnemonic device. *American Scientist, 58,* 496-510.
Bower, G. (1972). A selective review of organizational factors in memory. In E. Tulving & W. Donaldson (Eds.), *Organization and memory* (pp. 93-137). New York: Academic Press.
Boyce, M., Rieber, L. P., & Phillips, T. (1990, February). *From theory to practice: Examples of CAI lessons utilizing cognitive learning strategies.* Paper presented at the meeting of the Association of Educational Communications and Technology, Anaheim, CA.
Braddick, O. (1974). A short range process in apparent motion. *Vision Research, 14,* 519-527.
Brainerd, C. (1981). Working memory and the developmental analysis of probability judgement. *Psychological Review, 88,* 463-502.
Bransford, J. (1979). Human cognition: *Learning, understanding, and remembering.* Belmont, CA: Wadsworth Publishing.
Bransford, J., Sherwood, R., & Hasselbring, T. (1986). *Computers, videodiscs, and the teaching of thinking.* (Technical Report No. 86.1.1). Nashville, TN: Learning Technology Center.
Branson, R., & Grow, G. (1987). Instructional systems development. In R. Gagné (Ed.), *Instructional Technology: Foundations* (pp. 397-428). Hillsdale, NJ: Lawrence Erlbaum Associates.
Braun, C. (1969). Interest-loading and modality effects on textual response acquisition. *Reading Research Quarterly, 4,* 428-444.
Bricken, M., (1991). *Virtual Reality Learning Environments: Potentials and Challenges* (Tech. Pub. No. HITL-P-91-5). Seattle: University of Washington, Human Interface Technology Laboratory.
Briggs, L., & Wager, W. (1981). *Handbook of procedures for the design of instruction.* Englewood Cliffs, NJ: Educational Technology Publications.
Broadbent, D. E. (1971). *Decision and stress.* New York: Academic Press.
Brown, J. (1983). Learning-by-doing revisited for electronic learning environments. In M. A. White (Ed.), *The future of electronic learning* (pp. 13-32). Hillsdale, NJ: Lawrence Erlbaum Associates.
Brown, J., & Cunningham, S. (1990). Visualization in higher education. *Academic Computing, 4* (6), 24-25, 41-45.
Brown, J. S., Collins, A., & Duguid, P. (1989). Situated cognition and the culture of learning. *Educational Researcher, 18* (1), 32-42.
Bruner, J. (1966). *Toward a theory of instruction.* New York: Norton.
Bruner, J. (1986). *Actual minds, possible worlds.* Cambridge, MA: Harvard University Press.
Bruner, J. (1990). *Acts of meaning.* Cambridge, MA: Harvard University Press.
Bunderson, C., & Inouye, D. (1987). The evolution of computer-aided educational delivery systems. In R. Gagné (Ed.), *Instructional technology: Foundations* (pp. 283-318). Hillsdale, NJ: Lawrence Erlbaum Associates.

Burke, J. (1985). *The day the universe changed.* Boston: Little, Brown, and Company.
Burke, R. (1982). *CAI sourcebook.* Englewood Cliffs, NJ: Prentice-Hall.
Bush, V. (1945). As we may think. *Atlantic Monthly, 176* (1), 101-108.
Butler, R. J., Markulis, P. M., & Strang, D. R. (1988). Where are we? An analysis of methods and focus of the research on simulation gaming, *Simulation & Gaming, 19* (1), 3-26.
Butts, R. E., & Brown, J. R. (Eds.). (1989). *Constructivism and science.* Norwell, MA: Kluwer Academic Publishers.
Caldwell, R. (1980). Guidelines for developing basic skills instructional materials for use with microcomputer technology. *Educational Technology, 20* (10), 7-12.
Calvert, S., Huston, A., Watkins, B., & Wright, J. (1982). The relation between selective attention to television forms and children's comprehension of content. *Child Development, 53,* 601-610.
Canelos, J. (1987). Research justification and implementation of the program of systematic evaluation (PSE). In F. Dwyer (Ed.), *Enhancing visualized instruction: Recommendations for practitioners* (pp. 1-3). State College, PA: Learning Services.
Canelos, J., Taylor, W., & Dwyer, F. (1985). The effects of recall cue and cognitive trace compatibility when learning from visualized instructions: Application of encoding specificity. *International Journal of Instructional Media, 12,* 167-178.
Caraballo, A. (1985). *An experimental study to investigate the effects of computer animation on the understanding and retention of selected levels of learning outcomes.* Unpublished doctoral dissertation, The Pennsylvania State University.
Caraballo, J. (1985). *The effect of various visual display modes in computer-based instruction and language background upon achievement of selected educational objectives.* Unpublished doctoral dissertation, The Pennsylvania State University.
Carney, R. N., Levin, J. R., & Morrison, C. R. (1988). Mnemonic learning of artists and their paintings. *American Educational Research Journal, 25,* 107-125.
Carpenter, C. R. (1953). *A theoretical orientation for instruction film research. AV Communication Review, 1,* 38-52.
Carpenter, P., & Just, M. (1975). Sentence comprehension: A psycholinguistic processing model of verification. *Psychological Review, 82,* 45-73.
Carroll, J. M. (1990). *The Nurnberg funnel: Designing minimalist instruction for practical computer skill.* Cambridge, MA: MIT Press.
Carterette, E. C., & Friedman, M. P. (Eds.). (1975). *Handbook of Perception Volume 5: Seeing,* 387-424, London: Academic Press.
Cavanagh, P., Arguin, M., & von Grunau, M. (1989). Interattribute apparent motion. *Vision Research, 29* (9), 1197-1204.
Case, R., & Bereiter, C. (1984). From behaviourism to cognitive behaviourism to cognitive development: Steps in the evolution of instructional design. *Instructional Science, 13,* 141-158.
Chambers, J., & Sprecher, J. (1980). Computer-assisted instruction: Current trends and critical issues. In D. Walker & R. Hess (Eds.), *Instructional software: Principles and perspectives for design and use* (pp. 6-20). Belmont, CA: Wadsworth.
Champagne, A., Klopfer, L., & Anderson, J. (1980). Factors influencing the learning of classical mechanics. *American Journal of Physics, 48,* 1074-1079.
Champagne, A., Klopfer, L., & Gunstone, R. (1982). Cognitive research and the design of science instruction. *Educational Psychologist, 17* (1), 31-53.
Chase, W., & Clark, H. (1972). Mental operations in the comparison of sentences and pictures. In L.W. Gregg (Ed.), *Cognition in learning and memory* (pp. 205-232). New York: Wiley.
Chase, W., & Ericsson, K. (1981). Skilled Memory. In J. Anderson (Ed.), *Cognitive skills and their acquisition* (pp. 141-189). Hillsdale, NJ: Lawrence Erlbaum Associates.
Chi, M., & Koeske, R. (1983). Network representation of a child's dinosaur knowledge. *Developmental Psychology, 19,* 29-39.
Chi, M., Feltovich, P., & Glaser, R. (1981). Categorization and representation of physics problems by experts and novices. *Cognitive Science, 5,* 121-152.
Chu, G., & Schramm, W. (1979). *Learning from television: What the research says* (3rd printing). Washington, DC: National Association of Educational Broadcasters.
Clark, J. M., & Paivio, A. (1991). Dual coding theory and education. *Educational Psychology Review, 3,* 149-210.

Clark, R. (1982). Antagonism between achievement and enjoyment in ATI studies. *Educational Psychologist, 17,* 92-101.
Clark, R. (1983). Reconsidering research on learning from media. *Review of Educational Research, 17* (2), 92-101.
Clark, R. (1984a). Research on student thought processes during computer-based instruction. *Journal of Instructional Development, 7* (3), 2-5.
Clark, R. (1984b). “Where's the beef?”: A reply to Heinich. *Educational Communications and Technology Journal, 32* (4), 229-232.
Clark, R. (1985). Confounding in educational computing research. *Journal of Educational Computing Research, 1* (2), 137-148.
Clark, R. (1989). Current progress and future directions for research in instructional technology. *Educational Technology Research and Development, 37* (1), 57-66.
Clark, R., & Snow, R. (1975). Alternative designs for instructional technology research. *AV Communications Review*, *23* (4), 373-394.
Clark, R., & Voogel, A. (1985). Transfer of training principles for instructional design. *Educational Communications and Technology Journal, 33,* 113-123.
Clayson, J. (1982). Computer games teach problem solving. *Impact of Science on Society, 32,* 435-444.
Cleary, R. (1990). Contrast dependence of short-range apparent motion. *Vision Research, 30* (3), 463-478.
Cleary, R., & Braddick, O. J. (1990). Masking of low frequency information in short-range apparent motion. *Vision Research, 30* (2), 317-327.
Clements, D. H. (1985). Research on LOGO in education: Is the turtle slow but steady, or not even in the race? *Computers in the Schools, 2*(2/3), 55-71.
Clements, D. H. (1988). Social and cognitive interactions in educational computer environments. *American Educational Research Journal, 25*, 87-106.
Clements, D. H. (1991). Enhancement of creativity in computer environments. *American Educational Research Journal, 28*, 173-187.
Cognition and Technology Group at Vanderbilt (1990). Anchored instruction and its relationship to situated cognition. *Educational Researcher, 19*(6), 10.
Cognition and Technology Group at Vanderbilt (1992). The Jasper experiment: An exploration of issues in learning and instructional design. *Educational Technology Research and Development, 40*(1), 65-80.
Cohen, R. (1987). Implementing LOGO in the grade two classroom: Acquisition of basic programming concepts. *Journal of Computer-Based Instruction, 14*(4), 124-132.
Cohen, R. (1988). Formative evaluation of pre-LOGO programming environments: A collaborative effort of researchers, teachers, and children. *Journal of Computer-Based Instruction, 15*(4), 112-122.
Collins, A., Adams, M., & Pew, R. (1978). Effectiveness of an interactive map display in tutoring geography. *Journal of Educational Psychology, 70* (1), 1-7.
Collins, A., Brown, J. S., & Newman, S. E. (1987). Cognitive apprenticeship: Teaching the craft of reading, writing, and mathematics. In L. Resnick (Ed.), *Cognition and instruction: Issues and agendas.* Hillsdale, NJ: Lawrence Erlbaum Associates.
Collins, A., & Quinlan, M. (1969). Retrieval time from semantic memory. *Journal of Verbal Learning and Verbal Behavior, 8*, 240-247.
Condry, J. (1977). Enemies of exploration: Self-initiated versus other-initiated learning. *Journal of Personality and Social Psychology, 35*, 459-477.
Conrac Corporation (1985). *Raster graphics handbook.* New York: Van Nostrand Reinhold.
Cooper, L., & Shepherd, A. (1973). Chronometric studies of the rotation of mental images. In W. Chase (Ed.), *Visual information processing* (pp. 95-176). New York: Academic Press.
Craik, F., & Lockhart, R. (1972). Levels of processing: A framework for memory research. *Journal of Verbal Learning and Verbal Behavior, 11*, 761-784.
Cron, M. (1983, December). Trouble in Logoland. *Teaching, Learning, Computing,* pp. 59-65.
Cronbach, L. (1957). The two disciplines of scientific psychology. *American Psychologist, 12,* 671-684.
Cronin, M. W., & Cronin, K. A. (1992). A critical analysis of the theoretical foundations of interactive video instruction. *Journal of Computer-Based Instruction, 19*(2), 37-41.
Cupchik, G. C., & Berlyne, D. E. (1979). The perception of collative properties in visual stimuli. *Scandinavian Journal of Psychology, 20*, 93-104.
Curtis, R. V., & Reigeluth, C. M. (1984). The use of analogies in written text. *Instructional Science, 13,* 99-117.

Dale, E. (1946). *Audio-visual methods in teaching*. New York: Dryden Press.
Damarin, S. (1990). *Unthinking educational technology*. Proceedings of the 1990 Annual Convention of the Association for Educational Communication and Technology, Anaheim, CA, 179-191.
Dawson, G. O. (Principal Investigator). (1991). *ScienceVision*. Boston: Houghton Mifflin.
Day, J. C., & Belleza, F. S. (1983). The relationship between visual imagery mediators and recall. *Memory & Cognition, 11*(3), 251-257.
Daynes, R., & Butler, B. (1984). *The videodisc book: A guide and directory*. New York: John Wiley and Sons.
Debloois, M. (Ed.). (1982). *Videodisc/microcomputer courseware design*. Englewood Cliffs, NJ: Educational Technology Publications.
Deci, E. L. (1975). *Intrinsic motivation*. New York: Plenum Press.
Deci, E. L. (1985). *Intrinsic motivation and self-determination in human behavior*. New York: Plenum Press.
Dede, C. (1987). Empowering environments, hypermedia and microworlds. *The Computing Teacher, 15*(3), 20-24, 61.
Dick, W., & Carey, L. (1985). *The systematic design of instruction* (2nd ed.). Glenview, IL: Scott, Foresman, and Co.
diSessa, A. (1982). Unlearning Aristotelian physics: A study of knowledge-based learning. *Cognitive Science, 6*(1), 37-75.
diSessa, A., & Abelson, H. (1986). Boxer: A reconstructible computational medium. *Communications of the ACM, 29*(9), 859-868.
diSessa, A., & White, B. (1982). Learning physics from a dynaturtle. *Byte, 7*(8), 324.
Di Vesta, F. (1987). The cognitive movement and education. In J. Glover and R. Royce (Eds.), *Historical foundations of education psychology* (pp. 203-233). New York: Plenum Press.
Di Vesta, F., & Peverly, T. (1984). The effects of encoding variability, processing activity and rule example sequences on the transfer of conceptual rules. *Journal of Educational Psychology, 76*, 108-119.
Di Vesta, F., & Rieber, L. P. (1987). Characteristics of cognitive engineering: The next generation of instructional systems. *Educational Communications and Technology Journal, 35*(4), 213-230.
Dodd, D. H., & White, R. M. (1980). *Cognition, mental structures and processes.* Boston: Allyn and Bacon.
Duchastel, P. (1990-1991). Instructional strategies for simulation-based learning. *Journal of Educational Technology Systems, 19*(3), 265-275.
Duffield, J. A. (1991). Designing computer software for problem-solving instruction. *Educational Technology Research and Development, 39*(1), 50-62.
Dunn, T. (1984). Learning hierarchies and cognitive psychology: An important link for instructional psychology. *Educational Psychologist, 19*(2), 75-93.
Dwyer, F. M. (1969). Motion as an instructional cue. *AV Communication Review*, 3, 20-21, 41-43.
Dwyer, F. M. (1972). *A guide for improving visualized instruction*. State College, PA: Learning Services.
Dwyer, F. M. (1977). *Effect of method in cueing televised instruction*. Unpublished research report, The Pennsylvania State University.
Dwyer, F. M. (1978). *Strategies for improving visual learning*. State College, PA: Learning Services.
Dwyer, F. M. (Ed.). (1987). *Enhancing visualized instruction: Recommendations for practitioners*. State College, PA: Learning Services.
Dwyer, F. M., & Lamberski, R. J. (1982-83). A review of the research on the effects of the use of color in the teaching-learning process. *International Journal of Instructional Media, 10*, 303-328.
Edwards, J., Norton, S., Taylor, S., Weiss, M., & Dusseldorp, R. (1975). How effective is CAI? A review of the research. *Educational Leadership, 33*, 147-153.
Eylon, B., & Linn, M. (1988). Learning and instruction: An examination of four research perspectives in science education. *Review of Educational Research, 58*(3), 251-301.
Falio, T., & DeBloois, M. (1988). Designing a visual factors-based screen display interface: The new role of the graphic technology. *Educational Technology*, 28, 12-21.
Feaver, W. (1977). *When we were young: Two centuries of children's book illustration*. London: Thames & Hudson.
Flavell, J. (1985). *Cognitive development* (2nd Ed.). Englewood Cliffs, NJ: Prentice-Hall.
Fleming, M. (1987). Displays and communications. In R. Gagné (Ed.), *Instructional technology: Foundations* (pp. 233-260). Hillsdale, NJ: Lawrence Erlbaum Associates.
Fletcher, J. D., Hawley, D. E., & Piele, P. K. (1990). Costs, effects, and utility of microcomputer assisted instruction in the classroom. *American Educational Research Journal*, 27, 783-806.

Floyd, S., & Floyd, B. (1982). *Handbook of interactive video.* White Plains, NY: Knowledge Industry Publications.
Forman, G., & Pufall, P. (Eds.). (1988a). *Constructivism in the computer age*. Hillsdale, NJ: Lawrence Erlbaum Associates.
Forman, G., & Pufall, P. (1988b). Constructivism in the computer age: A reconstructive epilogue. In G. Forman & P. Pufall (Eds.), *Constructivism in the computer age* (pp. 235-250). Hillsdale, NJ: Lawrence Erlbaum Associates.
Fosnot, C. T. (1989). *Enquiring teachers, enquiring learners: A constructivist approach for teaching*. New York: Teacher's College Press.
Frase, L. (1973). Integration of written text. *Journal of Educational Psychology*, 65, 252-261.
Fredericksen, N. (1984). Implications of cognitive theory for instruction in problem-solving. *Review of Educational Research, 54*, 363-407.
Fry, E. (1972). *Reading instruction for classroom and clinic*. New York: McGraw-Hill.
Fuson, K. C., & Willis, G. B. (1989). Second graders's use of schematic drawings in solving addition and subtraction word problems. *Journal of Educational Psychology, 81*(4), 514-520.
Fuys, D. J., & Tischler, R. W. (1979). *Teaching mathematics in the elementary school*. Boston: Little, Brown, and Company.
Gagné, E. (1985). *The cognitive psychology of school learning*. Boston: Little, Brown, and Company.
Gagné, R. M. (1985). *The conditions of learning* (4th ed.). New York: Holt, Rinehart, & Winston.
Gagné, R. M. (Ed.). (1987). *Instructional Technology: Foundations*. Hillsdale, NJ: Lawrence Erlbaum Associates.
Gagné, R. M. (1989). *Studies of learning: 50 years of research.* Tallahassee: Florida State University, Learning Systems Institute.
Gagné, R. M., Briggs, L., & Wager, W. (1992). *Principles of instructional design* (4th ed.). New York: Holt, Rinehart, & Winston.
Gagné, R. M., & Dick, W. (1983). Instructional psychology. *Annual Review of Psychology*, 34, 261-295.
Gagné, R. M., & Glaser, R. (1987). Foundations in learning research. In R. Gagné (Ed.), *Instructional Technology: Foundations* (pp. 49-83). Hillsdale, NJ: Lawrence Erlbaum Associates.
Gagné, R. M., & Merrill, M. (1988, January). *Presidential session: A conversation on instructional design*. Presented at the annual convention of the Association for Educational Communications and Technology, Dallas, TX.
Gagné, R. M., & Merrill, M. D. (1990). Integrative goals for instructional design. *Educational Technology Research and Development, 38*(1), 23-30.
Gagné, R. M., Wager, W., & Rojas, A. (1981). Planning and authoring computer-assisted instruction lessons. *Educational Technology,* 21(9), 17-26.
Garhart, C., & Hannafin, M. (1986). The accuracy of cognitive monitoring during computer-based instruction. *Journal of Computer-Based Instruction*, 13, 88-93.
Gay, G. (1987). Interaction of learner control and prior understanding in computer-assisted video instruction. *Journal of Educational Psychology*, 78, 225-227.
Gayeski, D., & Williams, D. (1985). *Interactive media*. Englewood Cliffs, NJ: Prentice-Hall.
Gentner, D., & Stevens, A. (Eds.). (1983). *Mental models*. Hillsdale, NJ: Lawrence Erlbaum Associates.
Gerrish, J. (1991). New generation of “wiz kids” develop multimedia curriculum. *Technological Horizons in Education, 19*(3), 93-95.
Ghatala, E., & Levin, J. (1981). Children's incidental memory for pictures: Item processing versus list organizations. *Journal of Experimental Child Psychology, 31*(2), 231-244.
Gleick, J. (1987). *Chaos: Making a new science*. New York: Penguin Books.
Glynn, S., & Di Vesta, F. (1977). Outline and hierarchical organization as aids for study and retrieval. *Journal of Educational Psychology*, 69, 89-95.
Goldberg, F. (1974). Effects of imagery on learning incidental material in the classroom. *Journal of Educational Psychology, 66*(2), 233-237.
Goldsmith, E. (1987). The analysis of illustration in theory and practice. In H. Houghton & D. Willows (Eds.), *The psychology of illustration, volume 2: Instructional issues* (pp. 53-85). New York: Springer-Verlag.
Goldstein, A., Chance, J., Hoisington, M., & Buescher, K. (1982). Recognition memory for pictures: Dynamic vs. static stimuli. *Bulletin of the Psychonomic Society, 20*(1), 37-40.
Good, T. L., & Brophy, J. E. (1990). *Educational psychology: A realistic approach* (4th ed.). White Plains, NY: Longman.

Goodman, N. (1984). *Of mind and other matters.* Cambridge, MA: Harvard University Press.
Grabinger, R. S. (1984). CRT text design: Psychological attributes underlying the evaluation of models of CRT text displays. *Journal of Visual and Verbal Languaging, 4*(1), 17-39.
Green, M., & Odom, J. V. (1986). Correspondence matching in apparent motion: Evidence for three dimensional spatial representation. *Science*, 233, 1427-1429.
Green, M., & von Grunau, M. (1983). *Real and Apparent motion: One mechanism or two?* Proceedings of the ACM SIGGRAPH/SIGART Interdisciplinary Workshop on Motion: Representation and Perception (pp. 99-104). Toronto, Canada, New York: North-Holland.
Greene, D., & Lepper, M. (1974). How to turn play into work. *Psychology Today*, 8, 49-54.
Gropper, G. (1983). A behavioral approach to instructional prescription. In C.M. Reigeluth (Ed.), *Instructional-design theories and models: An overview of their current status* (pp. 106-161). Hillsdale, NJ: Lawrence Erlbaum Associates.
Gustafson, K. (1981). *Survey of instructional development models*. ERIC Clearinghouse on Information Resources, Syracuse University. (ERIC Document Reproduction Service No. ED 211 097)
Gustafson, K., & Powell, G. (1991). *Survey of instructional models with an annotated ERIC bibliography* (2nd ed.). (ERIC Document Reproduction Service No. ED 335 027)
Guttman, J., Levin, J. & Pressley, M. (1977). Pictures, partial pictures, and young children's oral prose learning. *Journal of Educational Psychology*, 69, 473-480.
Hale, M., Okey, J., Shaw, E., & Burns, J. (1985). Using computer animation in science testing. *Computers in the Schools, 2*(1), 83-90.
Halpern, D. F., Hansen, C., & Riefer, D. (1990). Analogies as an aid to understanding and memory. *Journal of Educational Psychology, 82*(2), 298-305.
Hamaker, C. (1986). The effects of adjunct questions on prose learning. *Review of Educational Research*, 56, 212-242.
Hannafin, M. (1983). The effects of instructional stimulus loading on the recall of abstract and concrete prose. *Educational Communications and Technology Journal*, 31, 103-109.
Hannafin, M. (1984). Guidelines for determining instructional locus of control in the design of computer-assisted instruction. *Journal of Instructional Development, 7*(3), 6-10.
Hannafin, M. (1985). Empirical issues in the study of computer-assisted interactive video. *Educational Communication and Technology Journal*, 33, 235-247.
Hannafin, M. (1992). Emerging technologies, ISD, and learning environments: Critical perspectives. *Educational Technology Research and Development, 40*(1), 49-63.
Hannafin, M., Garhart, C., Rieber, L. P., & Phillips, T. (1985). Keeping interactive video in perspective: Tentative guidelines and cautions in the design of interactive video. In E. Miller (Ed.), *Educational media and technology yearbook* (pp. 13-25). Littleton, CO: Libraries Unlimited.
Hannafin, M., & Hooper, S. (1989). An integrated framework for CBI screen design and layout. *Computers in Human Behavior*, 5, 155-165.
Hannafin, M., & Hughes, C. (1986). A framework for incorporating orienting activities in computer-based interactive video. *Instructional Science*, 15, 239-255.
Hannafin, M., & Peck, K. (1988). *The design, development, and evaluation of instructional software*. New York: MacMillan.
Hannafin, M., & Phillips, T. (1987). Perspectives in the design of interactive video: Beyond tape versus disc. *Journal of Research and Development in Education. 21*(1), 44-60.
Hannafin, M., Phillips, T., Rieber, L. P., & Garhart, C. (1987). The effects of orienting activities and cognitive processing time on factual and inferential learning. *Educational Communications and Technology Journal, 35*(2), 75-84.
Hannafin, M., Phillips, T., & Tripp, S. (1986). The effects of orienting, processing, and practicing activities on learning from interactive video. *Journal of Computer-Based Instruction*, 13, 134-139.
Hannafin, M., & Rieber, L. P. (1986, November). *The role of cognitive processing during computer-based instruction*. Paper presented at the annual meeting of the Association for the Development of Computer-Based Instructional Systems, Washington, DC.
Hannafin, M., & Rieber, L. P. (1989a). Psychological foundations of instructional design for emerging computer-based instructional technologies: Part I. *Educational Technology Research and Development, 37*(2), 91-101.

Hannafin, M., & Rieber, L. P. (1989b). Psychological foundations of instructional design for emerging computer-based instructional technologies: Part II. *Educational Technology Research and Development, 37*(2), 102-114.
Hartley, J. (1987). Designing electronic text: The role of print-based research. *Educational Communications and Technology Journal, 35*(1), 3-17.
Hartley, J. (1992). A postscript to Wainer's "Understanding graphs and tables". *Educational Researcher, 21*(5), 25-26.
Hartley, J., & Lovell, K. (1984). The psychological principles underlying the design of computer-based instructional systems. In D. Walker & R. Hess (Eds.), *Instructional software: Principles and perspectives for design and use* (pp. 38-56). Belmont, CA: Wadsworth.
Hawking, S. W. (1988). *A brief history of time: From the big bang to black holes.* New York: Bantam Books.
Hayes, D., & Tierney, R. (1982). Developing readers' knowledge through analogy. *Reading Research Quarterly*, 17, 256-280.
Heines, J. M. (1984). *Screen design strategies for computer-assisted instruction*. Bedford, MA: Digital Press.
Heinich, R. (1984). The proper study of instructional technology. *Educational Communications and Technology Journal, 32*(2), 67-87.
Hicken, S. (1991). Learning in tasks that incorporate natural versus unnatural mappings: An example using aircraft instrument comprehension. *Educational Technology Research and Development, 39*(1), 63-71.
Holley, C., & Dansereau, D. (Eds.). (1984). *Spatial learning strategies*. New York: Academic Press.
Hooper, S., & Hannafin, M. J. (1986). Variables affecting the legibility of computer generated text. *Journal of Instructional Development*, 9, 22-29.
Intraub, H. (1979). The role of implicit naming in pictorial encoding. *Journal of Experimental Psychology: Human Learning and Memory*, 5, 78-87.
Jacobs, R. L., & Baum, M. (1987). Simulations and games in training and development: Status and concerns about their use. *Simulations & Games, 18*(3), 385-394.
Jenkins, J. J. (1978). Four points to remember: A tetrahedral model of memory experiments. In L. S. Cermak & F. I. M. Craik (Eds.), *Levels of processing and human memory* (pp. 429-446). Hillsdale, NJ: Lawrence Erlbaum Associates.
Jonassen, D. (1982). *The technology of text*. Englewood Cliffs, NJ: Educational Technology Publications.
Jonassen, D. (Ed.). (1984). *The technology of text: Principles for structuring, designing, and displaying text* (Vol. 2). Englewood Cliffs, NJ: Educational Technology Publications.
Jonassen, D. (1986). Hypertext principles for text and courseware design. *Educational Psychologist, 21*(4), 269-292.
Jonassen, D. (Ed.). (1988a). *Instructional design for microcomputer courseware*. Hillsdale, NJ: Lawrence Erlbaum Associates.
Jonassen, D. (1988b). Integrating learning strategies into courseware to facilitate deeper processing. In D. Jonassen (Ed.), *Instructional designs for microcomputer courseware* (pp. 151-181). Hillsdale, NJ: Lawrence Erlbaum Associates.
Jonassen, D. (1988c, November). Designing structured hypertext and structuring access to hypertext. *Educational Technology*, 13-16.
Jonassen, D. (1991a). Objectivism versus constructivism: Do we need a new philosophical paradigm? *Educational Technology Research and Development, 39*(3), 5-14.
Jonassen, D. (1991b). Hypertext as instructional design. *Educational Technology Research and Development, 39*(1), 83-92.
Joyce, B., & Weil, M. (1980). *Models of teaching* (2nd ed.). Englewood Cliffs, NJ: Prentice-Hall.
Julesz, B. (1971). *Foundations of cyclopean perception*. Chicago: University of Chicago Press.
Just, M., & Carpenter, P. (1987). *The psychology of reading and language comprehension*. Newton, MA: Allyn and Bacon.
Kaufman, L. (1974). *Sight and mind: An introduction to visual perception*. London: Oxford University Press.
Kazdin, A. E. (1980). *Behavior modification in applied settings.* Homewood, IL: Dorsey.
Keller, J. M. (1983). Motivational design of instruction. In C. M. Reigeluth (Ed.), *Instructional-design theories and models: An overview of their current states* (pp. 383-434). Hillsdale, NJ: Lawrence Erlbaum Associates.
Keller, J. M., & Suzuki, K. (1988). Use of the ARCS motivation model in courseware design. In D. Jonassen (Ed.), *Instructional designs for microcomputer courseware* (pp. 401-434). Hillsdale, NJ: Lawrence Erlbaum Associates.

King, W. (1975). *A comparison of three combinations of text and graphics for concept learning*. (Report No. NPRDC-TR-76-16). San Diego, CA: Navy Personnel Research and Development Center. (ERIC Document Reproduction Service No. ED 112 936)

Kinzie, M. B. (1990). Requirements and benefits of effective interactive instruction: Learner control, self-regulation, and continuing motivation. *Educational Technology Research and Development, 38*(1), 5-21.

Kinzie, M. B., & Berdel, R. L. (1990). Design and use of hypermedia systems. *Educational Technology Research and Development*, *38*(3), 61-68.

Kinzie, M. B., & Sullivan, H. J. (1989). Continuing motivation, learning control, and CAI. *Educational Technology Research and Development*, *37*(2), 5-14.

Kinzie, M. B., Sullivan, H. J., & Berdel, R. L. (1988). Learner control and achievement in science computer-assisted instruction. *Journal of Educational Psychology, 80*(3), 299-303.

Klauer, K. (1984). Intentional and incidental learning with instructional texts: A meta-analysis for 1970-1980. *American Educational Research Journal, 21*(2), 323-339.

Klausmeier, H. (1990). Conceptualizing. In B. Jones & L. Idol (Eds.), *Dimensions of thinking and cognitive instruction* (pp. 93-138). Hillsdale, NJ: Lawrence Erlbaum Associates.

Klein, D. (1987, November). *Conditions affecting the effectiveness of animated and non-animated displays in computer-based instruction.* Paper presented at the annual meeting of the Association for the Development of Computer-Based Instructional Systems, Oakland, CA.

Knirk, F., & Gustafson, K. (1986). *Instructional technology: A systematic approach to education.* New York: Holt, Rinehart, & Winston.

Kobayashi, S. (1986). Theoretical issues concerning superiority of pictures over words and sentences in memory. *Perceptual and Motor Skills*, 63, 783-792.

Koenke, K. (1968, February). *The roles of pictures and readability in comprehension of the main idea of a paragraph.* Paper presented at the annual meeting of the American Educational Research Association, Chicago.

Kolers, P. A. (1963). Some differences between real and apparent visual movement. *Vision Research*, 3, 191-206.

Kolers, P. A. (1964). The illusion of movement. *Scientific American, 211*(4), 98-106.

Kolody, G. (1977). How students learn velocity and acceleration. *Journal of College Science Teaching*, 6, 224-227.

Korte, A. (1915). Kinematoskopische Untersuchungen. *Z. Psychol*, 72, 193-206.

Kosel, M., & Fish, M. (1983). *The factory* [computer program]. Pleasantville, NY: Sunburst Communications.

Kosslyn, S., Holyoak, K., & Huffman, C. (1976). A processing approach to the dual coding hypothesis. *Journal of Experimental Psychology*, 2, 223-233.

Kozma, R. B. (1991). Learning with media. *Review of Educational Research*, 61, 179-211.

Krayrallah, M., & Van Den Meiraker, M. (1987). LOGO programming and the acquisition of cognitive skills. *Journal of Computer-Based Instruction, 14*(4), 133-137.

Kulhavy, R. (1977). Feedback in written instruction. *Review of Educational Research*, 47, 211-232.

Kulik, C.-L., & Kulik, J. (1985). Effectiveness of computer-based education in colleges. *AEDS Journal, 19* (2/3), 81-108.

Kulik, C.-L., Kulik, J., & Shwalb, B. (1986). Effectiveness of computer-based adult education: A meta-analysis. *Journal of Educational Computing Research*, 2, 235-252.

Kulik, J., Kulik, C.-L., & Bangert-Drowns, R. (1985). Effectiveness of computer-based education in elementary schools. *Computers and Human Behavior*, 1, 59-74.

Lagerspetz, J. M. J., & Engblom, P. (1979). Immediate reactions to TV-violence by Finnish pre-school children of different personality types. *Scandinavian Journal of Psychology*, 20, 43-53.

Lamb, M. (1982). Interactive graphical modeling game for teaching musical concepts. *Journal of Computer-Based Instruction, 9*(2), 59-63.

Lane, D. (1980). Incidental learning and the development of selective attention. *Psychological Review, 87*(3), 316-319.

Larkin, J., McDermott, J., Simon, D., & Simon, H. (1980). Expert and novice performance in solving physics problems. *Science*, 208, 1335-1342.

Lawler, R. (1982). Designing computer-based microworlds. *Byte, 7*(8), 138-160.

Lee, H. (1980). The effects of review questions and review passages on transfer skills. *Journal of Educational Research. 73*(6), 330-335.

Lehrer, R., & Littlefield, J. (1991). Misconceptions and errors in LOGO: The role of instruction. *Journal of Educational Psychology, 83*(1), 124-133.
Leinhardt, G., Zaslavsky, O., & Stein, M. K. (1990). Functions, graphs, and graphing: Tasks, learning, and teaching. *Review of Educational Research*, 60, 1-64.
Lepper, M. (1985). Microcomputers in education: Motivational and social issues. *American Psychologist, 40*(1), 1-18.
Lepper, M., Greene, D., & Nisbett, R. (1973). Undermining children's intrinsic interest with extrinsic rewards: A test of the overjustification hypothesis. *Journal of Personality and Social Psychology*, 28, 129-137.
Lesgold, A. (1983). A rationale for computer-based reading instruction. In A. C. Wilkinson (Ed.), *Classroom computers and cognitive science*, (pp. 167-181) New York: Academic Press.
Lesgold, A., Levin, J., Shimron, J., & Guttman, J. (1975). Pictures and young children's learning from oral prose. *Journal of Educational Psychology*, 67, 636-642.
Lesgold, A., McCormick, C., & Golinkoff, R. (1975). Imagery training and children's prose learning. *Journal of Educational Psychology*, 67, 663-667.
Levie, W. (1987). Research on pictures: A guide to the literature. In D. Willows & H. Houghton (Eds.), *The psychology of illustration, volume 1: Instructional issues* (pp. 1-50). New York: Springer-Verlag.
Levie, W., & Lentz, R. (1982). Effects of text illustrations: A review of research. *Educational Communications and Technology Journal*, 30, 195-232.
Levin, H. M., Glass, G. V., & Mesiter, G. R. (1987). Cost-effectiveness of computer-assisted instruction. *Evaluation Review*, 11, 50-71. Cited in Fletcher, Hawley, & Piele (1990).
Levin, J., Anglin, G., & Carney, R. (1987). On empirically validating functions of pictures in prose. In D. Willows & H. Houghton (Eds.), *The psychology of illustration, volume 1: Instructional issues* (pp. 51-85). New York: Springer-Verlag.
Levin, J., & Divine-Hawkins, P. (1974). Visual imagery as a prose-learning process. *Journal of Reading Behavior*, 6, 23-30.
Levin, J., & Lesgold, A. (1978). On pictures in prose. *Educational Communications and Technology Journal, 26*(3), 233-243.
Levin, M., & Levin, J. (1990). Scientific mnemonomies: Methods for maximizing more than memory. *American Educational Research Journal*, 27, 301-321.
Lincoln, Y. S., & Guba, E. G. (1985). *Naturalistic research*. Newbury Park, CA: Sage Publications.
Linden, M., & Wittrock, M. (1981). The teaching of reading comprehension according to the model of generative learning. *Reading Research Quarterly*, 17, 44-57.
Lindseth, M. (1969). *The use of pictures to answer comprehension questions in a selected developmental reading series*. Unpublished master's thesis, University of Minnesota.
Locatis, C., Charuhas, J., & Banvard, R. (1990). Hypervideo. *Educational Technology Research and Development, 38*(2), 41-49.
Locatis, C., Letourneau, G., & Banvard, R. (1989). Hypermedia and instruction. *Educational Technology Research and Development, 37*(4), 65-77.
Lockard, J., Abrams, P., & Many, W. (1990). *Microcomputers for educators* (2nd ed.). Glenview, IL: Scott, Foresman/Little, Brown Higher Education.
Lorch, E., Anderson, D., & Levin, S. (1979). The relationship of visual attention to children's comprehension of television. *Child Development*, 50, 722-727.
Lukas, G., & Lukas, J. (1986). *LOGO: Principles, programming, and projects*. Monterey, CA: Brooks/Cole Publishing.
Mack, A., Klein, L., Hill, J., & Palumbo, D. (1989). Apparent motion: Evidence of the influence of shape, slant, and size on the correspondence process. *Perception and Psychophysics, 46*(2), 201-206.
Madigan, S. (1983). Picture memory. In J. C. Yuille (Ed.), *Imagery, memory, and cognition: Essays in honor of Allan Paivio* (pp. 65-89). Hillsdale, NJ: Lawrence Erlbaum Associates.
Maehr, M. (1976). Continuing motivation: An analysis of a seldom considered educational outcome. *Review of Educational Research*, 46, 443-462.
Malone, T. (1981). Toward a theory of intrinsically motivating instruction. *Cognitive Science*, 4, 333-369.
Mayer, R. (1979a). Can advance organizers influence meaningful learning? *Review of Educational Research, 49*(2), 371-383.
Mayer, R. (1979b). Twenty years of research on advance organizers: Assimilation theory is still the best predictor of results. *Instructional Science*, 8, 133-167.
Mayer, R. (1983). *Thinking, problem solving, cognition*. New York: W. H. Freeman & Company.

Mayer, R. (1984). Aids to text comprehension. *Educational Psychologist, 19*(1), 30-42.
Mayer, R. E. (1989a). Models for understanding. *Review of Educational Research*, *59*, 43-64.
Mayer, R. E. (1989b). Systematic thinking fostered by illustrations in scientific text. *Journal of Educational Psychology*, *81*, 240-246.
Mayer, R. E., & Anderson, R. B. (1991). Animations need narrations: An experimental test of a dual-coding hypothesis. *Journal of Educational Psychology, 83*(4), 484-490.
Mayer, R. E., & Gallini, J. K. (1990). When is an illustration worth ten thousand words? *Journal of Educational Psychology*, *82*, 715-726.
McDermott, V., & Combs, E. (1991, September). Breaking the cycle: Multimedia motivates high I.Q. underachievers. *Supplement to Technological Horizons in Education*, 24-26.
McNeil, B. J., & Nelson, K. R. (1991). Meta-analysis of interactive video instruction: A 10-year review of achievement effects. *Journal of Computer-Based Instruction, 18*(1), 1-6.
Merickel, M. L. (1990). The creative technologies project: Will training in 2D/3D graphics enhance kids' cognitive skills? *Technological Horizons in Education*, *18*(5), 55-58.
Merrill, M. (1983). Component display theory. In C. Reigeluth (Ed.), *Instructional-design theories and models: An overview of their status* (pp. 279-334). Hillsdale, NJ: Lawrence Erlbaum Associates.
Merrill, M. D., Li, Z., & Jones, M. K. (1990a). Limitations of first generation instructional design. *Educational Technology, 30*(1), 7-11.
Merrill, M. D., Li, Z., & Jones, M. K. (1990b). Second generation instructional design (ID2). *Educational Technology, 30*(2), 7-14.
Merrill, P., & Bunderson, C. (1981). Preliminary guidelines for employing graphics in instruction. *Journal of Instructional Development*, *4*(4), 2-9.
Meyer, B. J. (1975). Identification of the structure of prose and its implications for the study of reading and memory. *Journal of Reading Behavior, 7*, 7-47.
Mielke, K. W. (1990). Research and development at the Children's Television Workshop. *Educational Technology Research and Development*, *38*(4), 7-16.
Milheim, W., & Martin, B. (1991). Theoretical bases for the use of learner control: Three different perspectives. *Journal of Computer-Based Instruction, 18*(3), 99-105.
Miller, G. A. (1956). The magical number seven, plus or minus two: some limits on our capacity for processing information. *Psychological Review*, *63*, 81-97.
Montague, W. E. (1988). Promoting cognitive processing and learning by designing the learning environment. In D. Jonassen (Ed.), *Instructional designs for microcomputer courseware* (pp. 125-149). Hillsdale, NJ: Lawrence Erlbaum Associates.
Moore, M., Nawrocki, L., & Simutis, Z. (1979). *The instructional effectiveness of three levels of graphics displays for computer-assisted instruction.* (Report No. ARI-TP-359). Arlington, VA: Army Research Institute for the Behavioral and Social Sciences. (ERIC Document Service No. ED 178 057)
Morris, C. W. (1946). *Signs, languages and behavior.* New York: Prentice-Hall.
Morrison, G. R., Ross, S. M., & O'Dell, J. K. (1988). Text density level as a design variable in instructional displays. *Educational Communication and Technology Journal*, *36*, 103-115.
Morrison, G. R., Ross, S. M., Schultz, C. W., & O'Dell, J. K. (1989). Learner preferences for varying screen densities using realistic stimulus materials with single and multiple designs. *Educational Technology Research and Development*, *37*(3), 53-60.
Nashel, D. J., & Martin, J. J. (1991, September). Images in rheumatology (IR): A multimedia program for medical education. *Supplement to Technological Horizons in Education*, 14-16.
National Task Force on Educational Technology. (1986). Transforming American education: Reducing the risk to the nation, a report to the Secretary of Education, United States Department of Education. *T.H.E. Journal*, *14*(1), 58-67.
Nelson, D. (1979). Remembering pictures and words: Appearance, significance, and name. In L. S. Cermak & F. Craik (Eds.), *Levels of processing in human memory* (pp. 45-76). Hillsdale, NJ: Lawrence Erlbaum Associates.
Neuman, D. (1989). Naturalistic inquiry and computer-based instruction: Rationale, procedures, and potential. *Educational Technology Research and Development, 37*(3), 39-51.
Neuman, S. B. (1992). Is learning from media distinctive? Examining children's inferencing strategies. *American Educational Research Journal, 29*(1), 119-140.
Newby, T. J., & Stepich, D. A. (1987). Learning abstract concepts: The use of analogies as a mediational strategy. *Journal of Instructional Development, 10*(1), 20-26.

Newman, D. (1990). Opportunities for research on the organizational impact of school computers. *Educational Research, 19*(3), 8-13.
Nicholls, J. G., & Thorkildsen, T. A. (1989). Intellectual conventions versus matters of substance: Elementary school students as curriculum theorists. *American Educational Research Journal, 26*, 533-544.
Nickerson, R. (1965). Short-term memory for complex meaningful visual configurations: A demonstration of capacity. *Canadian Journal of Psychology, 19*, 155-160.
Nitsch, K. (1977). *Structuring decontextualized forms of knowledge*. Unpublished doctoral dissertation, Vanderbilt University.
Nix, D., & Spiro, R. (Eds.) (1990). *Cognition, education, and multimedia*. Hillsdale, NJ: Lawrence Erlbaum Associates.
Norman, D. (1978). Notes toward a theory of complex learning. In A. Lesgold, J. Pellegrino, S. Fokkema, & R. Glaser (Eds.), *Cognitive psychology and instruction* (pp. 39-48). New York: Plenum Press.
Norman, D. (1982). *Learning and memory*. San Francisco: W.H. Freeman & Company.
Norman, D. (1983). Some observations on mental models. In D. Gentner, & A. Stevens (Eds.), *Mental models* (pp. 7-14). Hillsdale, NJ: Lawrence Erlbaum Associates.
Norman, D. (1988). *The psychology of everyday things*. New York: Basic Books.
Norman, D. A. (1992). *Turn signals are the facial expressions of automobiles*. Reading, Mass: Addison-Wesley.
Norman, D. A., & Bobrow, D. G. (1975). On data-limited and resource-limited processes. *Cognitive Psychology, 7*, 44-64.
Novak, J., & Gowin, D. (1984). *Learning how to learn*. New York: Cambridge Books.
Nugent, G. (1987). Innovations in telecommunications. In R. Gagné (Ed.), *Instructional technology: Foundations* (pp. 261-282). Hillsdale, NJ: Lawrence Erlbaum Associates.
Orbach, E. (1979). Simulation games and motivation for learning. *Simulation & Games, 10*(1), 3-39.
Paivio, A. (1979). *Imagery and verbal processes*. Hillsdale, NJ: Lawrence Erlbaum Associates.
Paivio, A. (1983). Empirical case for dual coding. In J. Yuille (Ed.), *Imagery, memory, and cognition: Essays in honor of Allan Paivio* (pp. 307-332). Hillsdale, NJ: Lawrence Erlbaum Associates.
Paivio, A. (1986). *Mental representations: A dual coding approach*. New York: Oxford University Press.
Paivio, A. (1990). *Mental representations: A dual coding approach* (2nd ed.). New York: Oxford University Press.
Paivio, A. (1991). Dual coding theory: Retrospect and current status. *Canadian Journal of Psychology, 45*, 255-287.
Paivio, A., & Csapo, K. (1973). Picture superiority in free recall: Imagery or dual coding? *Cognitive Psychology, 5*, 176-206.
Pantle, A. J., & Picciano, J. V. (1976). A multistable movement display: Evidence for two separate motion systems in human vision. *Science, 193*, 500-502.
Papert, S. (1980). *Mindstorms: Children, computers, and powerful ideas*. New York: Basic Books.
Papert, S. (1981). Computer-based microworlds as incubators for powerful ideas. In R. Taylor (Ed.), *The computer in the school: Tutor, tool, tutee* (pp. 203-210). New York: Teacher's College Press.
Papert, S. (1987). Computer criticism vs. technocentric thinking. *Educational Researcher, 16*(1), 22-30.
Papert, S. (1988). The conservation of Piaget: The computer as grist to the constructivist mill. In G. Forman & P. Pufall (Eds.), *Constructivism in the computer age* (pp. 3-13). Hillsdale, NJ: Lawrence Erlbaum Associates.
Parsloe, E. (1983). *Interactive video*. Wilmslow, Cheshire (UK): Sigma Technical Press.
Pauline, R., & Hannafin, M. (1987, June). Interactive slide-sound instruction: Incorporate the power of the computer with high fidelity visual and aural images. *Educational Technology, 27*(6), 27-31.
Peng, C., & Levin, J. (1979). Pictures and children's story recall: Some questions of durability. *Educational Communication and Technology Journal, 27*, 39-44.
Perkins, D. N., & Simmons, R. (1988). Patterns of misunderstanding: An integrative model for science, math, and programming. *Review of Educational Research, 58*, 303-326.
Petersik, J. T. (1989). The two process distinction of apparent motion. *Psychological Bulletin, 106*(1), 107-123.
Petkovich, M., & Tennyson, R. (1984). Clark's "learning from media:" A critique. *Educational Communications and Technology Journal, 32*(4), 233-241.
Pettersson, R. (1989). *Visuals for information: Research and practice*. Englewood Cliffs, NJ: Educational Technology Publications.

Piaget, J. (1951). *Play, dreams, and imitation in childhood*. New York: Norton.
Piaget, J. (1952). *The origins of intelligence in children*. New York: International University Press.
Piaget, J. (1970). *Genetic epistemology*. New York: Columbia University Press.
Pinker, S. (1984). Visual cognition: An introduction. *Cognition, 18*, 1-63.
Pokorny, C. K., & Gerald, C. F. (1989). *Computer graphics: The principles behind the art and science*. Irvine, CA: Franklin, Beedle & Associates.
Polya, G. (1957). *How to solve it* (2nd ed.). New York: Doubleday.
Poole, L. (1988). *HyperTalk*. Redmond, WA: Microsoft Press.
Pressley, M. (1976). Mental imagery helps eight-year-olds remember what they read. *Journal of Educational Psychology, 68*, 355-359.
Pressley, M. (1977). Imagery and children's learning: Putting the picture in developmental perspective. *Review of Educational Research, 47*(4), 585-622.
Pressley, M., & Levin, J. R. (1978). Developmental constraints associated with children's use of the keyword method of foreign language vocabulary learning. *Journal of Experimental Child Psychology, 26*, 359-373.
Proffitt, D., & Kaiser, M. (1986). The use of computer graphics animation in motion perception research. *Behavior Research Methods, Instruments, & Computers, 18*(6), 487-492.
Pylyshyn, Z. (1973). What the mind's eye tells the mind's brain: A critique of mental imagery. *Psychological Bulletin, 80*, 1-24.
Pylyshyn, Z. (1981). The imagery debate: Analogue media versus tacit knowledge. *Psychological Review, 88*, 16-45.
Ragan, T. J. (1989). Educational technology in the schools: The technoromantic movement. *Educational Technology, 29*(11), 32-34.
Ramachandran, V. S., & Anstis, S. M. (1986). The perception of apparent motion. *Scientific American, 254*, 102-109.
Rankin, R.O. (1989). The development of an illustration design model. *Educational Technology Research and Development, 37*(2), 25-46.
Rathus, S. (1987). *Psychology* (3rd ed.). New York: Holt, Rinehart, and Winston.
Reder, L. (1982). Elaborations: when do they help and when do they hurt? *Text, 2*, 211-244.
Reed, M. (1991). Videodiscs help American Indians learning English and study heritage. *Technological Horizons in Education, 19*(3), 96-97.
Reed, S. (1985). Effect of computer graphics on improving estimates to algebra word problems. *Journal of Educational Psychology, 77*(3), 285-298.
Reigeluth, C. (Ed.). (1983a). *Instructional-design theories and models: An overview of their current status*. Hillsdale, NJ: Lawrence Erlbaum Associates.
Reigeluth, C. (1983b). Instructional design: What is it and why is it? In C. Reigeluth (Ed.), *Instructional-design theories and models: An overview of their current status* (pp. 3-36). Hillsdale, NJ: Lawrence Erlbaum Associates.
Reigeluth, C., & Curtis, R. (1987). Learning situations and instructional models. In R. Gagné (Ed.), *Instructional technology: Foundations* (pp. 175-206). Hillsdale, NJ: Lawrence Erlbaum Associates.
Reigeluth, C., & Schwartz, E. (1989). An instructional theory for the design of computer-based simulation. *Journal of Computer-Based Instruction, 16*(1), 1-10.
Reigeluth, C., & Stein, F. (1983). The elaboration theory of instruction. In C. Reigeluth (Ed.), *Instructional-design theories and models: An overview of their status* (pp. 335-382). Hillsdale, NJ: Lawrence Erlbaum Associates.
Reiser, R. (1987). Instructional technology: A history. In R. Gagné (Ed.), *Instructional Technology: Foundations* (pp. 11-48). Hillsdale, NJ: Lawrence Erlbaum Associates.
Reitman, J., & Reuter, H. (1980). Organization revealed by recall orders and confirmed by pauses. *Cognitive Psychology, 12*, 554-581.
Reynolds, G. (1968). *A primer of operant conditioning*. Glenview, IL: Scott, Foresman, and Company.
Rheingold, H. (1991). *Virtual reality*. New York: Summit Books.
Rickards, J. P. (1978). Instructional psychology: From a behavioral to a cognitive orientation. *Improving Human Performance Quarterly, 7*(4), 256-266.
Riding, R., & Tite, H. (1985). The use of computer graphics to facilitate story telling in young children. *Educational Studies, 11*(3), 203-210.
Rieber, L. P. (1987). LOGO and its promise: A research report. *Educational Technology, 27*(2), 12-16.

Rieber, L. P. (1989). The effects of computer animated elaboration strategies and practice on factual and application learning in an elementary science lesson. *Journal of Educational Computing Research, 5*(4), 431-444.
Rieber, L. P. (1990a). Animation in computer-based instruction. *Educational Technology Research and Development, 38*(1), 77-86.
Rieber, L. P. (1990b). Using animation in science instruction with young children. *Journal of Educational Psychology, 82*(1), 135-140.
Rieber, L. P. (1990c). *Space shuttle commander* [computer software]. Washington, DC: NASA Educational Technology Branch.
Rieber, L. P. (1991a). Effects of visual grouping strategies of computer animated presentations on selective attention in science. *Educational Technology Research and Development, 39*(4), 5-15.
Rieber, L. P. (1991b). Computer animation, incidental learning, and continuing motivation. *Journal of Educational Psychology, 83*(3), 318-328.
Rieber, L. P. (1992). Computer-based microworlds: A bridge between constructivism and direct instruction. *Educational Technology Research and Development, 40*(1), 93-106.
Rieber, L. P. (1992b). Review of Donald Norman's "The psychology of everyday things." *Educational Technology Research and Development, 40*(4), 107-109.
Rieber, L. P. (in press). A pragmatic view of instructional technology. In K. Tobin (ed.), *Constructivist perspectives on science education*. Washington, DC: AAAS Press.
Rieber, L. P., Boyce, M., & Alkindi, M. (1991). The effects of computer-based interactive visuals as orienting and practicing activities on retrieval tasks in science. *The International Journal of Instructional Media, 16*(1).
Rieber, L. P., Boyce, M., & Assad, C. (1990). The effects of computer animation on adult learning and retrieval tasks. *Journal of Computer-Based Instruction, 17*(2), 46-52.
Rieber, L. P., & Hannafin, M. (1988). The effects of textual and animated orienting activities and practice on learning from computer-based instruction. *Computers in the Schools, 5*(1/2), 77-89.
Rieber, L. P., & Kini, A. (1991). Theoretical foundations of instructional applications of computer-generated animated visuals. *Journal of Computer-Based Instruction, 18*(3), 83-88.
Rieber, L. P., & Parmley, W. (1992, February). *Effects of animated computer simulations on inductive learning with adults*. Paper presented at the meeting of the Association of Educational Communications and Technology, Washington, DC.
Rieber, L. P., & Welliver, P. (1989). Infusing educational technology into mainstream educational computing. *International Journal of Instructional Media, 16*(1), 21-32.
Rigney, J., & Lutz [Alesandrini], K. (1975). *The effects of interactive graphic analogies on recall of concepts in chemistry* (Technical Report No. 75). Washington, DC: Office of Naval Research. (ERIC Document Service No. ED 109 639)
Roblyer, M. D. (1988). Fundamental problems and principles of designing effective courseware. In D. Jonassen (Ed.), *Instructional designs for microcomputer courseware* (pp. 7-33). Hillsdale, NJ: Lawrence Erlbaum Associates.
Romiszowski, A. J. (1981). *Designing instructional systems*. London: Kogan Page.
Ross, S. (1984). Matching the lesson to the student: Alternative adaptive designs for individualized learning systems. *Journal of Computer-Based Instruction, 11*(2), 42-48.
Rothkopf, E. (1966). Learning from written instructive material: An exploration of the control of inspection behavior by test-like events. *American Educational Research Journal, 3*, 241-249.
Rotter, J. (1954). *Social learning and clinical psychology*. Englewood Cliffs, NJ: Prentice-Hall.
Royer, J. (1979). Theories of the transfer of learning. *Educational Psychologist, 14*, 53-69.
Ruben, B. D., & Lederman, J. M. (1982). Instructional simulation gaming: Validity, reliability, and utility. *Simulation & Games, 13*, 233-244.
Rumelhart, D. (1980). Schemata: The building blocks of cognition. In R. Spiro, B. Bruce, & W. Brewer (Eds.), *Theoretical issues in reading comprehension* (pp. 33-59). Hillsdale, NJ: Lawrence Erlbaum Associates.
Rumelhart, D., & Norman, D. (1981). Analogical processes in learning. In J. Anderson (Ed.), *Cognitive skills and their acquisition*. Hillsdale, NJ: Lawrence Erlbaum Associates.
Rumelhart, D., & Ortony, A. (1977). The representation of knowledge in memory. In R. C. Anderson, R. J. Spiro, & W. E. Montague (Eds.), *Schooling and the acquisition of knowledge*. Hillsdale, NJ: Lawrence Erlbaum Associates.

Ryan, M. (1975). The impact of television news film on perceived media credibility. *Journal of Applied Communications Research, 3*, 69-75.

Sadoski, M., Paivio, A., & Goetz, E. (1991). A critique of schema theory in reading and a dual coding alternative. *Reading Research Quarterly, 26*(4), 463-484.

Saettler, P. (1990). *The evolution of American educational technology*. Denver: Libraries Unlimited.

Sagaria, S., & Di Vesta, F. (1978). Learner expectations induced by adjunct questions and the retrieval of intentional and incidental information. *Journal of Educational Psychology, 70*(3), 280-288.

Sales, G. C., Tsai, B., & MacLeod, R. (1991). The evolution of K-12 instructional software: An analysis of leading microcomputer-based programs from 1981-1988. *Journal of Computer-Based Instruction, 18*(2), 41-47.

Salisbury, D. (1988). Effective drill and practice strategies. In D. Jonassen (Ed.), *Instructional designs for microcomputer courseware* (pp. 103-124). Hillsdale, NJ: Lawrence Erlbaum Associates.

Salisbury, D., Richards, B., & Klein, J. (1985). Designing practice: A review of prescriptions and recommendations from instructional design theories. *Journal of Instructional Design, 8*(4), 9-19.

Salomon, G. (1979). *Interaction of media, cognition, and learning*. San Francisco: Jossey Bass.

Salomon, G. (1983). The differential investment of mental effort in learning from different sources. *Educational Psychologist, 18*(1), 42-51.

Salomon, G. (1991). Transcending the qualitative-quantitative debate: The analytic and systemic approaches to educational research. *Educational Researcher, 20*(6), 10-18.

Salomon, G., & Clark, R. (1977). Reexamining the methodology of research on media and technology in education. *Review of Educational Research, 47*(1), 99-120.

Salomon, G., & Gardner, H. (1986). The computer as educator: Lessons from television research. *Educational Researcher, 15*(1), 13-19.

Salomon, G., Perkins, D., & Globerson, T. (1991). Partners in cognition: Extending human intelligence with intelligent technologies. *Educational Researcher, 20*(3), 2-9.

Samuels, S. (1967). Attentional process in reading: The effect of pictures on the acquisition of reading responses. *Journal of Educational Psychology, 58*, 337-342.

Samuels, S. (1970). Effects of pictures on learning to read, comprehension, and attitudes. *Review of Educational Research, 40*, 397-407.

Schiffman, H. R. (1976). *Sensation and perception: An integrated approach*. London: John Wiley and Sons.

Schimmel, B. (1988). Providing meaningful feedback in courseware. In D. Jonassen (Ed.), *Instructional designs for microcomputer courseware* (pp. 183-195). Hillsdale, NJ: Lawrence Erlbaum Associates.

Schouten, J. F. (1967). Subjective stroboscopy and model of visual motion detectors. In W. Wathen-Dunn (Ed.), *Models for the perception of speech and visual form* (pp. 44-55). Cambridge, MA: MIT Press/ Wiley.

Schunk, D. H. (1990). Goal setting and self-efficacy during self-regulated learning. *Educational Psychology, 25*(1), 71-86.

Seaman, D. F., & Fellenz, R. A. (1989). *Effective strategies for teaching adults*. Columbus, OH: Merrill.

Sekular, R., & Blake, R. (1985). *Perception*. New York: Alfred A. Knopf.

Sekuler, R. (1975). Seeing. In E. C. Carterette & M. P. Friedman (Eds.), *Handbook of Perception: Volume 5* (pp. 387-424). London: Academic Press.

Sfondilias, J. S., & Siegel, M. A. (1990). Combining discovery and direct instruction strategies in computer-based teaching of mathematical problem solving. *Journal of Computer-Based Instruction, 17*(4), 130-134.

Shepard, R. (1967). Recognition memory for words, sentences, and pictures. *Journal of Verbal Learning and Verbal Behavior, 6*, 156-163.

Sherwood, C. (1991). Adventure games in the classroom: A far cry from A says Apple . . . *Computers Education, 17*(4), 309-315.

Shimron, J. (1975). *Imagery and comprehension of prose by elementary school children*. Dissertation Abstracts International, 36, 795-A. (University Microfilms No. 75-18, 254)

Shirey, L. L., & Reynolds, R. E. (1988). Effect of interest on attention and learning. *Journal of Educational Psychology, 80*(2), 159-166.

Shuell, T. (1986). Cognitive conceptions of learning. *Review of Educational Research, 56*(4), 411-436.

Siegler, R. (1989). Strategy diversity and cognitive assessment. *Educational Researcher, 18*(9), 15-20.

Simon, H. A. (1981). Studying human intelligence by creating artificial intelligence. *American Scientist, 69*(3), 300-309.

Simonsen, M., Clark, R., Kulik, J., Tennyson, R., & Winn, W. (1987, February). "*. . . Mere vehicles . . .": A symposium on the status of research in instructional technology*. Proceedings of the 1987 Annual Convention of the Association for Educational Communications and Technology, Atlanta, GA, 1-43.
Skinner, B. (1968). *Technology of teaching*. New York: Meredith Publishing.
Slavin, R. (1988). *Educational psychology: Theory into practice*. Englewood Cliffs, NJ: Prentice-Hall.
Slythe, R. (1970). *The art of illustration 1750-1900*. London: The Library Association.
Smith, P. (1988, January). *Cognitive psychology: Implications for selection of techniques and measures in instructional research*. Proceedings of the 1988 Annual Convention of the Association for Educational Communications and Technology, 10, 706-719.
Smith, S. G., & Jones, L. L. (1991, September). The acid test: Five years of multimedia chemistry. *Supplement to Technological Horizons in Education*, 21-23.
Snow, R. (1977). Individual differences and instructional design. *Journal of Instructional Development, 1*(1), 23-26.
Spangenberg, R. (1973). The motion variable in procedural learning. *AV Communication Review, 21*(4), 419-436.
Special issue on Research and Development in Instructional Interactive Video, 1986, *Journal of Computer-Based Instruction, 13*(4).
Special issue/section on qualitative methodology. *American Educational Research Journal, 24*(2).
Sperling, G. (1960). The information available in brief visual presentation. *Psychological Monographs, 74*, 498.
Standing, L., Conezio, J., & Haber, R. (1970). Perception and memory for pictures: Single trial learning of 2500 visual stimuli. *Psychonomic Science, 19*, 73-74.
Steinbeck, J. (1939). *The grapes of wrath*. New York: The Viking Press.
Steinberg, E. R. (1977). Review of student control in computer-assisted instruction. *Journal of Computer-Based Instruction, 3*, 84-90.
Steinberg, E. R. (1989). Cognition and learner control: A literature review, 1977-1988. *Journal of Computer-Based Instruction, 16*, 117-121.
Steinberg, E. R. (1991). *Computer-assisted instruction: A synthesis of theory, practice, and technology*. Hillsdale, NJ: Lawrence Erlbaum Associates.
Streibel, M. (1986). A critical analysis of the use of computers in education. *Educational Communications and Technology Journal, 34*(3), 137-161.
Strommen, E. F., & Revelle, G. L. (1990). Research in interactive technologies at the Children's Television Workshop. *Educational Technology Research and Development, 38*(4), 65-80.
Surber, J., & Leeder, J. (1988). The effect of graphic feedback on student motivation. *Journal of Computer-Based Instruction, 15*(1), 14-17.
Sutton, R. E. (1991). Equity and computers in the schools: A decade of research. *Review of Educational Research, 61*, 475-503.
Szabo, M., Dwyer, F., & De Melo, H. (1981). Visual testing: Visual literacy's second dimension. *Educational Communications and Technology Journal, 29*, 177-187.
Tatsuoka, M. M. (1986). Graph theory and its applications in educational research: A review and integration. *Review of Educational Research, 56*, 291-329.
Taylor, W., Canelos, J., Belland, J., Mayton, G., & Dwyer, F. (1989, February). *The effects of MCBI program embedded cognitive strategies for performance on verbal and visual testing*. Proceedings of the 1989 Annual Convention of the Association for Educational Communications and Technology, Dallas, TX, 575-591.
Tennyson, R. D. (1980). Instructional control strategies and content structure as design variables in concept acquisition using computer-based instruction. *Journal of Educational Psychology, 72*(4), 525-532.
Tennyson, R. D. (1984). Application of artificial intelligence methods to computer-based instructional design: The Minnesota Adaptive Instructional System. *Journal of Instructional Development, 7*(3), 17-22.
Tennyson, R. D., & Buttrey, T. (1980). Advisement and management strategies as design variables in computer-assisted instruction. *Educational Communications and Technology Journal, 28*, 169-176.
Tennyson, R. D., & Cocchiarella, M. J. (1986). An empirically based instructional design theory for teaching concepts. *Review of Educational Research, 56*(1), 40-71.
Tennyson, R. D., & Park, O. (1980). The teaching of concepts: A review of instructional design research literature. *Review of Educational Research, 50*(1), 55-70.

Tennyson, R. D., & Park, O. (1987). Artificial intelligence and computer-based learning. In R. Gagné (Ed.), *Instructional technology: Foundations* (pp. 319-342). Hillsdale, NJ: Lawrence Erlbaum Associates.
Tennyson, R. D., & Rothen, W. (1979). Management of computer-based instruction. *Journal of Computer-Based Instruction, 5*, 126-134.
Tetenbaum, T., & Mulkeen, T. (1984, November). LOGO and the teaching of problem solving: A call for a moratorium. *Educational Technology, 24*(11), 16-19.
Thatcher, D. C. (1986). Promoting learning through games and simulations. *Simulation & Gaming, 16*(4), 144-154.
Thorndyke, P. (1977). Cognitive structures in comprehension and memory of narrative discourse. *Cognitive Psychology, 9*, 77-110.
Tobias, S. (1987). Learner characteristics. In R. Gagné (Ed.), *Instructional technology: Foundations* (pp. 207-232). Hillsdale, NJ: Lawrence Erlbaum Associates.
Tobin, K., & Dawson, G. (1992). Constraints to curriculum reform: Teachers and the myths of schooling. *Educational Technology Research and Development, 40*(1), 81-92.
Tripp, S., & Bichelmeyer, B. (1990). Rapid prototyping: An alternative instructional design strategy. *Educational Technology Research and Development, 38*(1), 31-44.
Tripp, S., & Roby, W. (1990). Orientation and disorientation in a hypertext lexicon. *Journal of Computer-Based Instruction, 17*, 120-124.
Trollip, S. R., & Lippert, R. C. (1987). Constructing knowledge bases: A promising instructional tool. *Journal of Computer-Based Instruction, 14*(2), 44-48.
Trowbridge, D., & McDermott, L. (1981). Investigation of student understanding of the concept of acceleration in one dimension. *American Journal of Physics, 49*, 242-253.
Tufte, E. R. (1983). *The visual display of quantitative information.* Cheshire, CT: Graphics Press.
Tufte, E. R. (1990). *Envisioning information*. Cheshire, CT: Graphics Press.
Tulving, E. (1979). Relation between encoding specificity and levels of processing. In L. S. Cermak & F. I. M. Craik (Eds.), *Levels of Processing in Human Memory* (pp. 405-428). Hillsdale, NJ: Lawrence Erlbaum Associates.
Ullman, S. (1979). *The interpretation of visual motion*. Cambridge, MA: MIT Press.
Underwood, B. (1963). Stimulus selection in verbal learning. In C. Cofer & B. Musgrave (Eds.), *Verbal Behavior and Learning, Problems and Processes* (pp. 33-48). New York: McGraw-Hill.
Vgotsky, L. S. (1978). *Mind in society: The development of higher mental processes*. Cambridge, MA: Harvard University Press.
Vgotsky, L. S. (1986). *Thought and language*. Cambridge, MA: MIT Press.
Vuyk, R. (1981). *Overview and critique of Piaget's genetic epistemology*, 1965-1980. New York: Academic Press.
Wager, W., & Wager, S. (1985). Presenting questions, processing responses, and providing feedback in CAI. *Journal of Instructional Development, 8*(4), 2-8.
Wainer, H. (1992). Understanding graphs and tables. *Educational Researcher, 21*(1), 14-23.
Wang, A. (1983). Individual differences in learning speed. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 9*, 300-311.
Wanner, H. E. (1968). *On remembering, forgetting, and understanding sentences: a study of the deep structure hypothesis*. Unpublished doctoral dissertation, Harvard University.
Watzlawick, P. (Ed.). (1984). *The invented reality*. New York: W.W. Norton and Company.
Weiner, B. (1979). A theory of motivation for some classroom experiences. *Journal of Educational Psychology, 71*(1), 3-25.
Weiner, B. (1985). An attribution theory of achievement motivation and emotion. *Psychological Review, 92*, 548-573.
Weiner, B. (1990). History of motivational research in education. *Journal of Educational Psychology, 82*(4), 616-622.
Weintraub, S. (1960). *The effect of pictures on the comprehension of a second grade basal reader*. Unpublished doctoral dissertation, University of Illinois.
White, B. (1984). Designing computer games to help physics students understand Newton's laws of motion. *Cognition and Instruction, 1*(1), 69-108.
White, B., & Horwitz, P. (1987). *ThinkerTools: Enabling children to understand physical laws* (Report No. 6470). Cambridge, MA: BBN Laboratories.

Wildman, T., & Burton, J. (1981). Integrating learning theory with instructional technology. *Journal of Instructional Development, 4*(3), 5-14.
William, R. (1981). *Taxonomy of Visual Processes*. Hillsdale, NJ: Lawrence Erlbaum Associates.
Willis, G. B., & Fuson, K. C. (1988). Teaching children to use schematic drawings to solve addition and subtraction word problems. *Journal of Educational Psychology, 80*(2), 192-201.
Willis, J., Hovey, L., & Hovey, K. (1987). *Computer simulations: A source book to learning in an electronic environment*. New York: Garland Publishing.
Willows, D. (1978). A picture is not always worth a thousand words: Pictures as distractors in reading. *Journal of Educational Psychology, 70*(2), 255-262.
Winer, B. J. (1971). *Statistical principles in experimental design* (2nd ed.). New York: McGraw Hill.
Winkles, J. (1986). Achievement, understanding, and transfer in a learning hierarchy. *American Educational Research Journal, 23*(2), 275-288.
Winn, W. D. (1987). Charts, graphs, and diagrams in educational materials. In D. Willows & H. Houghton (Eds.), *The psychology of illustration, volume 1: Instructional issues* (pp. 152-198). New York: Springer-Verlag.
Winn, W. D., Li, T., & Schill, D. (1991). Diagrams as aids to problem solving: Their role in facilitating search and computation. *Educational Technology Research and Development, 39*(1), 17-29.
Winn, W. D., & Sutherland, S. W. (1989). Factors influencing the recall of elements in maps and diagrams and the strategies used to encode them. *Journal of Educational Psychology, 81*, 33-39.
Withrow, M. (1978). Computer animation and language instruction. *American Annals of the Deaf, 123*(6), 723-725.
Withrow, M. (1979). Illustrating language through computer generated animation. *American Annals of the Deaf, 124*(5), 549-552.
Wittrock, M. (1974). Learning as a generative process. *Educational Psychologist, 11,* 87-95.
Wittrock, M. C. (1978). The cognitive movement in instruction. *Educational Psychologist, 15*, 15-29.
Wood, E., Pressley, M., Turnure, J. E., & Walton, R. (1987). Enriching children's recall of picture-dictionary definitions with interrogation and elaborated pictures. *Educational Communications and Technology Journal, 35*(1), 43-52.
Wood, L. E., & Stewart, P. W. (1987). Improvement of practical reasoning skills with a computer game. *Journal of Computer-Based Instruction, 14*(2), 49-53.
Woodward, J., Carnine, D., & Gersten, R. (1988). Teaching problem solving through computer simulations. *American Educational Research Journal, 25,* 72-86.
Woolfolk, A. (1987). *Educational psychology* (3rd ed). Englewood Cliffs, NJ: Prentice-Hall.
Wright, J., & Huston, A. (1983). A matter of form: Potentials of television for young viewers. *American Psychologist, 38*, 835-843.
Yerkes, R. M., & Dodson, J. D. (1908). The relation of strength of stimulus to rapidity of habit- formation. *Journal of Comparative Neurological Psychology, 18*, 459-482.
Zavotka, S. (1987). Three-dimensional computer animated graphics: A tool for spatial skill instruction. *Educational Communication and Technology Journal, 35*(3), 133-144.
Zimmerman, B. J. (1989). A social cognitive view of self-regulated academic learning. *Journal of Educational Psychology, 81*, 329-339.
Zimmerman, B. J. (1990). Self-regulated learning and academic achievement: An overview. *Educational Psychologist, 25*(1), 3-17.
Zook, K. B., & Di Vesta, F. J. (1991). Instructional analogies and conceptual misrepresentations. *Journal of Educational Psychology, 83*(2), 246-252.

CREDITS

Chapter 1
Figure 1.2 Copyright 1992 by Authorware, Inc. Authorware Professional and Movie Editor are registered trademarks of Authorware, Inc.

Chapter 2
Figure 2.3 This advertisement provided courtesy of the U.S. Council for Energy Awareness, Washington, D.C. **Figure 2.8** Edward R. Rufte, *The visual display of quantitative information* (Cheshire, Connecticut: Graphics Press, 1983). **Figure 2.10** Copyright 1993 by Becky Rieber. Reproduced with permission. **Figure 2.11** Copyright 1992 by Texas A&M University and reprinted with permission. **Figure 2.18** Copyright 1993 by the American Heart Association. Reproduced with permission. **Figure 2.19** Copyright 1987 by Electronic Arts and reproduced with permission.

Chapter 3
Figure 3.10 Copyright 1987-1989 Claris Corporation. All Rights Reserved. MacDraw is a registered trademark of Claris Corporation. **Figure 3.14** Copyright 1992 by Computer Associates International, Inc. CA-Crickett Graph III is a registered trademark of Computer Associates International, Inc. **Figure 3.18** Copyright 1992 by Authorware, Inc. Authorware Professional and Movie Editor are registered trademarks of Authorware, Inc. **Figure 3.19** Copyright 1992 by Macromedia, Inc. Macromedia ® Director is a product of Macromedia, Inc. **Figure 3.20** Copyright 1992 by Authorware, Inc. Authorware Professional and Movie Editor are registered trademarks of Authorware, Inc. **Figure 3.21** Copyright 1992 by Authorware, Inc. Authorware Professional and Movie Editor are registered trademarks of Authorware, Inc.

Chapter 4
Figure 4.12 Larkin, J., McDermott, J., Simon, D., & Simon, H., 1980, Expert and novice performance in solving physics problems, *Science*, 208, 1335-1342. Copyright 1980 by the American Association for the Advancement of Science.

Chapter 6
Figure 6.3 Reprinted from Rieber, L.P. (1991). Effects of visual grouping strategies of computer-animated presentations on selective attention in science. *Educational Technology Research and Development*, *39*(4), 5-15. Copyright 1991 by the Association for Educational Communications and Technology. Reprinted by permission of the publisher. **Figure 6.4** Reprinted from Rieber, L.P. (1990). Using computer animated graphics in science instruction with children. *Journal of Educational Psychology*, *82*(1), 135-140. Copyright 1990 by the American Psychological Association. Reprinted by permission of the publisher. **Figure 6.6** Reprinted from Rieber, L.P. (1991). Animation, incidental learning, and continuing motivation. *Journal of Educational Psychology*, *83*(3), 318-328. Copyright 1991 by the American Psychological Association. Reprinted by permission of the publisher.

Chapter 7
Figure 7.4 Adapted from Tripp, S.D., & Bichelmeyer, B. (1990). Rapid prototyping: An alternative instructional design strategy. *Educational Technology Research and Development*, *38*(1), 31-44. Copyright 1991 by the Association for Educational Communications and Technology.

Chapter 8
Figure 8.1 Taken from Alessi, S. (1988). Fidelity in the design of instructional simulations. *Journal of Computer-Based Instruction*, *15*(2), 40-47. Copyright 1988 by the Association for the Development of Computer-Based Instructional Systems and reproduced with permission. **Figure 8.14** Adapted from Rieber, L.P. (1992). Computer-based microworlds: A bridge between constructivism and direct instruction. *Educational Technology Research and Development*, *40*(1), 93-106. Copyright 1992 by the Association for Educational Communications and Technology. Adapted by permission of the publisher.

Chapter 9

Figure 9.2 Copyright 1992 by R.D. Zellner and reprinted with permission. **Figure 9.3** Copyright 1992 by R.D. Zellner and reprinted with permission. **Figure 9.4** This 3-D illusion provided courtesy of John Williamson.