CHAPTER FOUR OUTLINE

STUDENT LEARNING OUTCOMES

1. Define decision support system, list its components, and identify the type of applications it's suited to.
2. Define geographic information systems and state how they differ from other decision support tools.
3. Define artificial intelligence and list the different types that are used in business.
4. Define expert systems and describe the types of problems to which they are applicable.
5. Define neural networks and fuzzy logic and the uses of these AI tools.
6. Define genetic algorithms and list the concepts on which they are based and the types of problems they solve.
7. Define intelligent agents, list the four types, and identify the types of problems they solve.
8. Define agent-based modeling and swarm intelligence.

INDUSTRY PERSPECTIVES

FEATURES

WEB SUPPORT

www.mhhe.com/haag

- Learning about investing
- Researching the company behind the stock
- Retrieving stock quotes
- Computer-aided decision support

CHAPTER FOUR

Decision Support and Artificial Intelligence

Brainpower for Your Business

OPENING CASE STUDY: DECISION SUPPORT SYSTEM—THE RESIDENT SECOND OPINION

Doctors in the near future will not dream of going into an examination room without automated decision support. So says Dr. C. Martin Harris, CIO at the Cleveland Clinic, who believes that IT will soon become as central to patient care as the stethoscope. Giving this trend a nudge in the right direction, the Cleveland Clinic and the Mayo Clinic have each begun work on data warehouses (see Chapter 3) that will help health care professionals make analyses crucial to providing better patient care.

At the Cleveland Clinic, the first part of this ambitious project is to improve the treatment and survival rate among patients with abdominal aortic aneurysms that are usually spotted only when they are about to erupt, at which time more than half of patients die. An aneurysm is a weak spot in a blood vessel that causes big problems when it ruptures. Predicting which patients are at risk for this medical condition will help doctors to provide treatment to the patient early enough to avert disaster.

The Cleveland Clinic system ties together all the databases in the hospital, including patient histories, genetic makeup records, and countless lab results, among others. This will allow doctors to quickly compare a sick patient with thousands of others to find variables and correlations that would be highly impractical to find manually.

At the Mayo Clinic, using a similar system that has information on 4 million patients, doctors have already made some very useful discoveries. For instance, a new drug was being tested for lung cancer but was producing mixed results until a researcher discovered that a subset of the lung-cancer patients—those with a specific gene makeup—responded very well. This is the type of medical business intelligence that is vital in treating serious illnesses.

The type of computer-aided decision support systems that the Mayo Clinic and the Cleveland Clinic are using incorporate data-mining techniques and neural networks (an artificial intelligence system that identifies patterns and categorizes individual entities into categories), among other tools. With so much information available, treatment can be tailored to each individual's case, taking into account as much information about the patient as possible including the patient's demographic information.

The type of IT system we are discussing here is one type of computer-aided decision-making tool. Computer-aided decision support includes any IT system that allows you to analyze information to reach a decision or that provides the decision or answer for you. These tools for decision making are the focus of this chapter.[1]

Introduction

The objective of decision support systems is to help you analyze information to find business intelligence, which, as you learned in Chapters 2 and 3, is the extraction of the true meaning of information so that you can take creative and powerful steps to gain a competitive advantage. In the opening case study, you saw how decision support systems are starting to be used in the medical profession. Clinics and doctors use the "business intelligence" in their databases to make faster and better diagnoses.

In this chapter we'll investigate the tools that IT can provide to help you to transform business information into business intelligence and make good decisions. According to *Management Review,* the big winners in tomorrow's business race will be those organizations that are "big of brain and small of mass."[2]

For many years, computers have been crunching numbers faster and more accurately than people can. A computer can unerringly calculate a payroll for 1,000 people in the time it takes a pencil to fall from your desk to the floor. Because of IT, knowledge workers have been freed from much of the drudgery of manually handling day-to-day transactions. And now, IT power is augmenting brainpower and thought processes in ways previously seen only in science fiction. In some cases, IT power is actually *replacing* human brainpower to a limited degree.

Businesses, like individuals, use brainpower to make decisions, some big, some small, some relatively simple, and some very complex. As an effective knowledge worker, you'll have to make decisions on issues such as whether to expand the workforce, extend business hours, use different raw materials, or start a new product line. IT can help you in most, if not all, of these decisions. The extended brainpower that IT offers you as a decision maker comes in the form of decision support systems and artificial intelligence.

Whether to use a decision support system or some form of artificial intelligence depends on the type of decision you have to make and how you plan to go about making it. So let's first look at different types of decisions and the process you go through to make a decision. Then we'll discuss decision support systems and artificial intelligence—IT brainpower (see Figure 4.1). To learn even more about decision support systems and artificial intelligence, visit the Web site that supports this text at www.mhhe.com/haag.

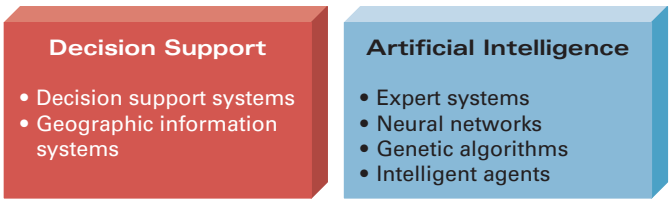

Figure 4.1

The Two Categories of Computer-Aided Decision Support

Decisions, Decisions, Decisions

You make many decisions every day from simple to complex. Contrast a decision about which mozzarella cheese to buy based on cost with a decision about which job offer to take. Choosing the right job is definitely a more complex decision because it has multiple decision criteria, not all of which are quantifiable (unlike the price of cheese). Therefore, it's much more difficult to select among job offers.

Decision making is crucial to business. Organizations devote vast resources of time and money to decision-making processes. In this section, we'll consider the phases of decision making and different decision types to help you better understand how IT can benefit the decision-making process.

HOW YOU MAKE A DECISION

In business, decision making has four distinct phases (see Figure 4.2).[3] These four phases are:

1. ***Intelligence*** (find what to fix): Find or recognize a problem, need, or opportunity (also called the diagnostic phase of decision making). The intelligence phase involves detecting and interpreting signs that indicate a situation which needs your attention. These "signs" come in many forms: consistent customer requests for new-product features, the threat of new competition, declining sales, rising costs, an offer from a company to handle your distribution needs, and so on.
2. ***Design*** (find fixes): Consider possible ways of solving the problem, filling the need, or taking advantage of the opportunity. In this phase, you develop all the possible solutions you can.
3. ***Choice*** (pick a fix): Examine and weigh the merits of each solution, estimate the consequences of each, and choose the best one (which may be to do nothing at all). The "best" solution may depend on such factors as cost, ease of implementation, staffing requirements, and timing. This is the prescriptive phase of decision making—it's the stage at which a course of action is prescribed.
4. ***Implementation*** (apply the fix): Carry out the chosen solution, monitor the results, and make adjustments as necessary. Simply implementing a solution is seldom enough. Your chosen solution will always need fine-tuning, especially for complex problems or changing environments.

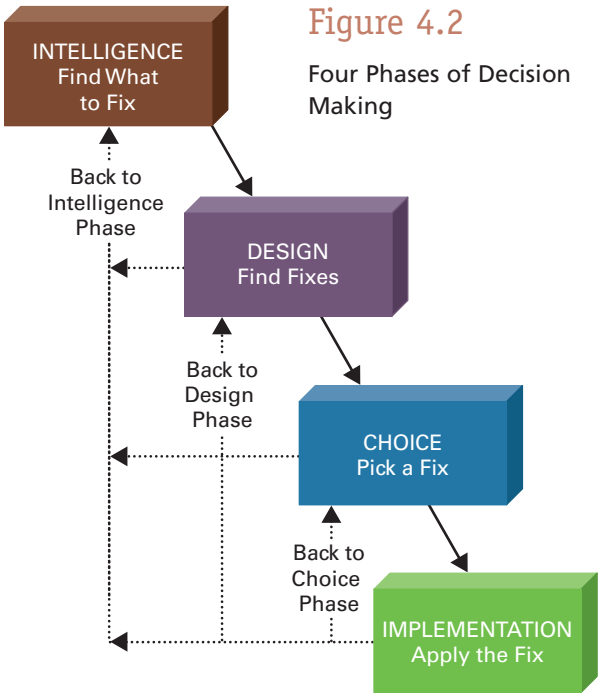

Figure 4.2
Four Phases of Decision Making

This four-phase process is not necessarily linear: You'll often find it useful or necessary to cycle back to an earlier phase. When choosing an alternative in the choice phase, for example, you might become aware of another possible solution. Then you would go back to the design phase, include the newly found solution, return to the choice phase, and compare the new solution to the others you generated.

TYPES OF DECISIONS YOU FACE

It's pretty clear that deciding which cheese to buy when you want the cheapest is a decision with a simple comparison that leads to a correct answer. Thus, it is an example of a structured decision, whereas choosing the right job is an example of a decision with nonstructured and structured elements. That is, some parts are quantifiable and some are not.

A ***structured decision*** involves processing a certain kind of information in a specified way so that you will always get the right answer. No "feel" or intuition is necessary. These are the kinds of decisions that you can program—if you use a certain set of inputs and process them in a precise way, you'll arrive at the correct result. Calculating gross pay for hourly workers is an example. You can easily automate these types of structured decisions with IT.

A ***nonstructured decision*** is one for which there may be several "right" answers, and there is no precise way to get a right answer. No rules or criteria exist that guarantee you a good solution. Deciding whether to introduce a new product line, employ a new marketing campaign, or change the corporate image are all examples of decisions with nonstructured elements.

In reality, most decisions fall somewhere between structured and nonstructured. The job choice decision is an example (see Figure 4.3). In choosing the right job, the salary part of the decision is structured, whereas the other criteria involve nonstructured aspects (for example, your perception of which job has the best advancement opportunity). Stock market investment analysis is another example of "somewhere in between" because you can calculate financial ratios and use past performance indicators. However, you still have to consider nonstructured aspects of the companies, such as projected prime interest rate, unemployment rates, and competition.

Another way to view decisions is by the frequency with which the decision has to be made. The decision as to which job to take is the sort of decision you don't make on a regular basis; this is a nonrecurring, or ad hoc, decision. On the other hand, determining pay for hourly employees is a routine decision that businesses face periodically. Therefore, determining gross pay for hourly employees is a recurring decision.

A ***recurring decision*** is one that happens repeatedly, and often periodically, whether weekly, monthly, quarterly, or yearly. You'll usually use the same set of rules each time. When you calculate pay for hourly employees, the calculation is always the same regardless of the employee or time period. A ***nonrecurring,*** or ***ad hoc, decision*** is one that you make infrequently (perhaps only once), and you may even have different criteria for determining the best solution each time. A company merger is an example. These don't happen often, although they are becoming more frequent. And if the managers of a company need to make the merger decision more than once, they will most likely have to evaluate a different set of criteria each time. The criteria depend on the needs of the companies considering the merger, the comparability of their products and services, their debt structure, and so on.

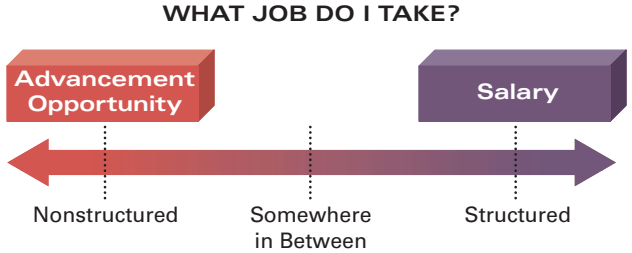

Figure 4.3

Viewing Structured versus Nonstructured Decision Making as a Continuum

INDUSTRY PERSPECTIVE

USING A DSS TO KEEP THE COST OF GAS DOWN

Think you spend a lot on gas for your car every year? J. B. Hunt Transportation Inc. spends a lot more. J. B. Hunt moves freight around the country on its 10,000 trucks and 48,000 trailers. The company spent $250 million in 2004 on fuel. That figure was up by 40 percent over the previous year. Diesel fuel is the company's second-largest expense (drivers' wages is the largest), and the freight hauler wanted to find a way to reduce that. Part of the answer lay, as it often does, in IT.

In 2000, J. B. Hunt installed a decision support system that provides drivers with help in deciding which gas station to stop at for refueling. Using satellite communications, the system beams diesel-fuel prices from all over the country straight into the cabs of the trucks. The software accesses a database with local taxes for each area of the country and then calculates for the drivers how much refueling will actually cost.

J. B. Hunt doesn't require drivers to use this system, but provides incentives for those who do. The company estimates that the system saves about $1 million annually.[4]

Decision Support Systems

In Chapter 3, you saw how data mining can help you make business decisions by giving you the ability to slice and dice your way through massive amounts of information. Actually, a data warehouse with data-mining tools is a form of decision support. The term *decision support system,* used broadly, means any computerized system that helps you make decisions. There's also a more restrictive definition, however. It's rather like the term *medicine.* Medicine can mean the whole health care industry or it can mean cough syrup, depending on the context.

Narrowly defined, a ***decision support system (DSS)*** is a highly flexible and interactive IT system that is designed to support decision making when the problem is not structured. A DSS is an alliance between you, the decision maker, and specialized support provided by IT (see Figure 4.4). IT brings speed, vast amounts of information, and

What You Bring	Advantages of a DSS	What IT Brings
Experience	Increased productivity	Speed
Intuition	Increased understanding	Information
Judgment	Increased speed	Processing capabilities
Knowledge	Increased flexibility	
	Reduced problem complexity	
	Reduced cost	

Figure 4.4

The Alliance between You and a Decision Support System

sophisticated processing capabilities to help you create information useful in making a decision. You bring know-how in the form of your experience, intuition, judgment, and knowledge of the relevant factors. IT provides great power, but you—as the decision maker—must know what kinds of questions to ask of the information and how to process the information to get those questions answered. In fact, the primary objective of a DSS is to improve your effectiveness as a decision maker by providing you with assistance that will complement your insights. This union of your know-how and IT power helps you generate business intelligence so that you can quickly respond to changes in the marketplace and manage resources in the most effective and efficient ways possible. Following are some examples of the varied applications of DSSs:

- A national insurance company uses a DSS to analyze its risk exposure when insuring drivers with histories of driving under the influence. The DSS revealed that married male homeowners in their forties with one DUI conviction were rarely repeat offenders. By lowering its rates to this group the company increased it market share without increasing its risk exposure.[5]
- Burlington Northern and Santa Fe (BNSF) railroad regularly tests the rails its trains ride on to prevent accidents. Worn out or defective rails result in hundreds of derailments every year, so it's important to address the problem. Using a decision support system to schedule rail testing, BNSF decreased its rail-caused derailments by 33 percent in 2000, while the other three large railroad companies had a 16 percent rise in such accidents.[6]
- Customer relationship management (CRM), as you saw in Chapter 2, is an important part of any successful company's strategy. Decision support is an important part of CRM. On Wall Street, retail brokerage companies analyze customers' behaviors and goals with decision support, which highlights opportunities and alerts brokers to beginning problems.[7]

COMPONENTS OF A DECISION SUPPORT SYSTEM

DSSs vary greatly in application and complexity, but they all share specific features. A typical DSS has three components (see below and Figure 4.5): model management, data management, and user interface management.

Before we look at these three components individually, let's get a quick overview of how they work together. When you begin your analysis, you tell the DSS, using the user interface management component, which model (in the model management component) to use on what information (in the data management component). The model requests the information from the data management component, analyzes that information, and sends the result to the user interface management component, which in turn passes the results back to you (see Figure 4.5). Here's an example of a decision support system at Lands' End clothing business.

- *Model management:* The DSS at Lands' End has to have models to analyze information. The models create new information that decision makers need to plan product lines and inventory levels. For example, Lands' End uses a statistical model called regression analysis to determine trends in customer buying patterns and forecasting models to predict sales levels.
- *Data management:* The DSS's data management component stores Lands' End's customer and product information. In addition to this organizational information, the company also needs external information, such as demographic information and industry and style trend information.

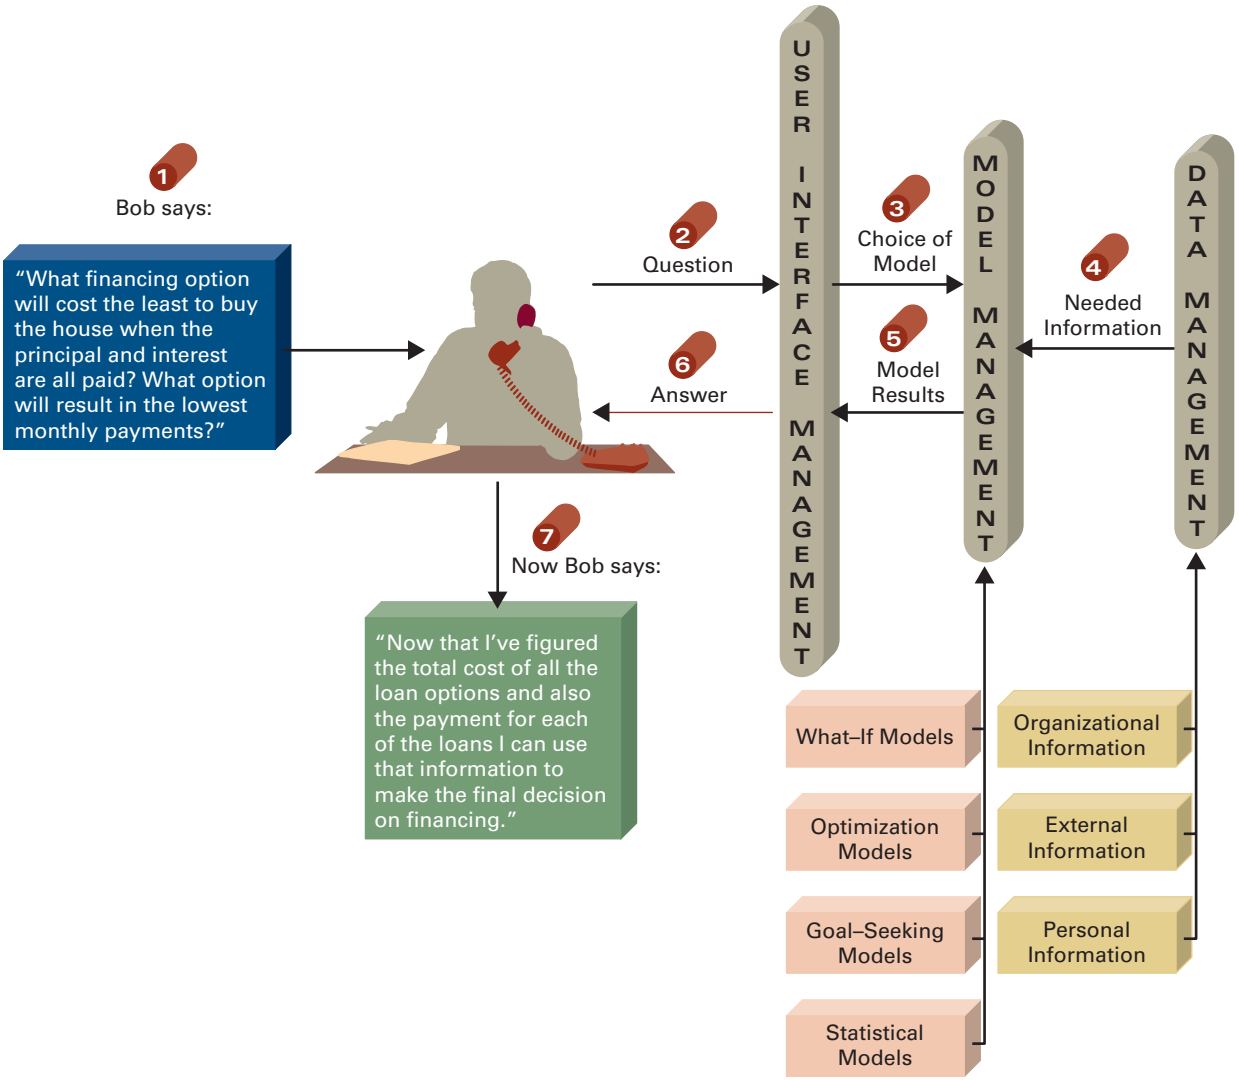

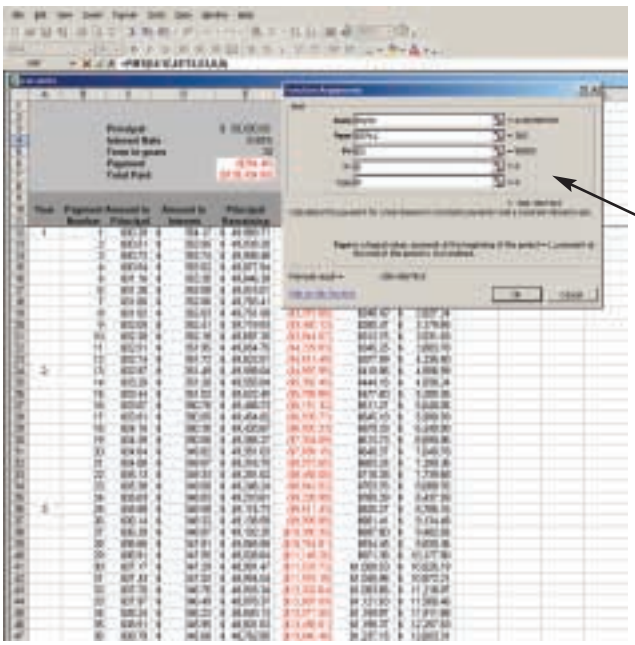

The PMT function calculates the monthly payment on the mortgage. The Rate (interest rate) is usually quoted as an annual rate so you have to divide by 12 to get the monthly rate. Similarly, the NPER (number of periods) is in years so you have to multiply by 12 for monthly payments.

Figure 4.5

Components of a Decision Support System

- *User interface management:* A user interface enables Lands' End decision makers to access information and specify the models they want to use to create the information they need.

Now we'll examine the three DSS components in more general terms.

MODEL MANAGEMENT COMPONENT The ***model management*** component consists of both the DSS models and the DSS model management system. A model is a representation of some event, fact, or situation. Businesses use models to represent variables and their relationships. For example, you would use a statistical model called analysis of variance to determine whether newspaper, television, and billboard advertising are equally effective in increasing sales. DSSs help in various decision-making situations by using models that allow you to analyze information in many different ways. The models you use in a DSS depend on the decision you're making and, consequently, the kind of analysis you require. For example, you would use what-if analysis to see what effect the change of one or more variables will have on other variables, or optimization to find the most profitable solution given operating restrictions and limited resources. You can use spreadsheet software such as Excel to create a simple DSS for what-if analysis. Look again at Figure 4.5 for an example of a spreadsheet DSS you might build to compare how much you'd pay for a house at different interest rates and payback periods.

The model management system stores and maintains the DSS's models. Its function of managing models is similar to that of a database management system. The model management component can't select the best model for you to use for a particular problem—that requires your expertise—but it can help you create and manipulate models quickly and easily.

DATA MANAGEMENT COMPONENT The ***data management*** component performs the function of storing and maintaining the information that you want your DSS to use. The data management component, therefore, consists of both the DSS information and the DSS database management system. The information you use in your DSS comes from one or more of three sources:

1. *Organizational information:* You may want to use virtually any information available in the organization for your DSS. You can design your DSS to access this information directly from your company's databases and data warehouses.
2. *External information:* Some decisions require input from external sources of information. Various branches of the federal government, Dow Jones, and the Internet, to mention just a few, can provide additional information for use with a DSS.
3. *Personal information:* You can incorporate your own insights and experience—your personal information—into your DSS.

USER INTERFACE MANAGEMENT COMPONENT The ***user interface management*** component allows you to communicate with the DSS. It consists of the user interface and the user interface management system. This is the component that allows you to combine your know-how with the storage and processing capabilities of the computer. The user interface is the part of the system you see; through it you enter information, commands, and models. If you have a DSS with a poorly designed user interface—if it's too rigid or too cumbersome to use—you simply won't use it no matter what its capabilities. The best user interface uses your terminology and methods and is flexible, consistent, simple, and adaptable.

GLOBAL PERSPECTIVE

GOT MILK? THEY DO IN BRITISH COLUMBIA

The British Columbia Milk Marketing Board has the responsibility of collecting raw milk from farmers and transporting it to the processing plants. The Marketing Board transfers payments from the processors to the farmers—to the tune of $17 million every two weeks. Milk is picked up from 350 farms per day by a fleet of 70 trucks, so that all 700 farms in the system see collection every other day. The milk goes to 26 processors across the vast province.

The paper-based system of record keeping was collapsing under its own weight. Upon pumping a farmer's milk into his truck, a driver recorded the number of liters collected on a four-part paper form and gave the farmer one copy. To fill a two-trailer truck, a driver collected milk from five farms and kept track, on paper, of how much milk came from each one. When the trailers were full of milk, the driver headed for the processing plant where a receipt for the quantity of milk was filled out. The truck then proceeded to the next farm and repeated the process.

Each day's work generated 350 producer slips from the farms and about 70 truck slips used to keep track of payments to the truck companies. To enter and process all this information, even using a computer, took many hours each day, especially with the problems of slips that truck drivers lost or forgot to turn over and mistakes made entering data onto forms that had to be investigated. The cost of the paper system was about $100,000 per year. The basic accounting of incoming milk and outgoing payments took so much time that there wasn't much left for the in-depth analysis necessary for the type of decision making that would have led to more effective use of resources.

The new decision support system handles data entry in real time. Each driver uses a handheld wireless scanning unit to send information to the central system. Each farm and processing plant has a unique bar code that the driver just has to scan. All that's left is to enter the volume of milk and add any comments like the temperature of the milk and so on. A wireless printer in the truck's cab can produce a printed copy of the transaction for the farmer or processing plant. The data is then sent over the cellular data network to the Marketing Board's Internet access point, and from there to the main computer for processing.

Quite apart from collecting data faster and more accurately, enabling more efficient decision making, the IT system allows the British Columbia Milk Marketing Board to quickly identify any farm that was the source of a substandard shipment so that corrective action can be taken immediately. And it's not just management that benefits from the decision support aspect of the new system. IT also supplies drivers with useful information, such as how much space is left in each trailer and how a load should be divided between trailers. The new decision support system not only brings costs down; it increases profits too.[8]

Geographic Information Systems

In 1992, Hurricane Andrew attacked the east coast of the United States leaving devastation throughout several states in its wake. One of the places hardest hit was Miami, Florida, where the hurricane came on land smashing businesses and private buildings and causing billions of dollars in damage. Reporters at the *Miami Herald* believed that not all the damage was due to Andrew. They hypothesized that at least some of the harm was a result of shoddy construction of homes built after 1980.

Four months after Hurricane Andrew, the paper ran a series of reports and used geographic information system maps to make its point. A ***geographic information system (GIS)*** is a decision support system designed specifically to analyze spatial information. Spatial information is any information that can be shown in map form, such as roads, the distribution of the bald eagle population, sewer systems, or the path of a hurricane.

Figure 4.6

San Diego Shown in ArcExplorer 2, GIS Software from ESRI[9]

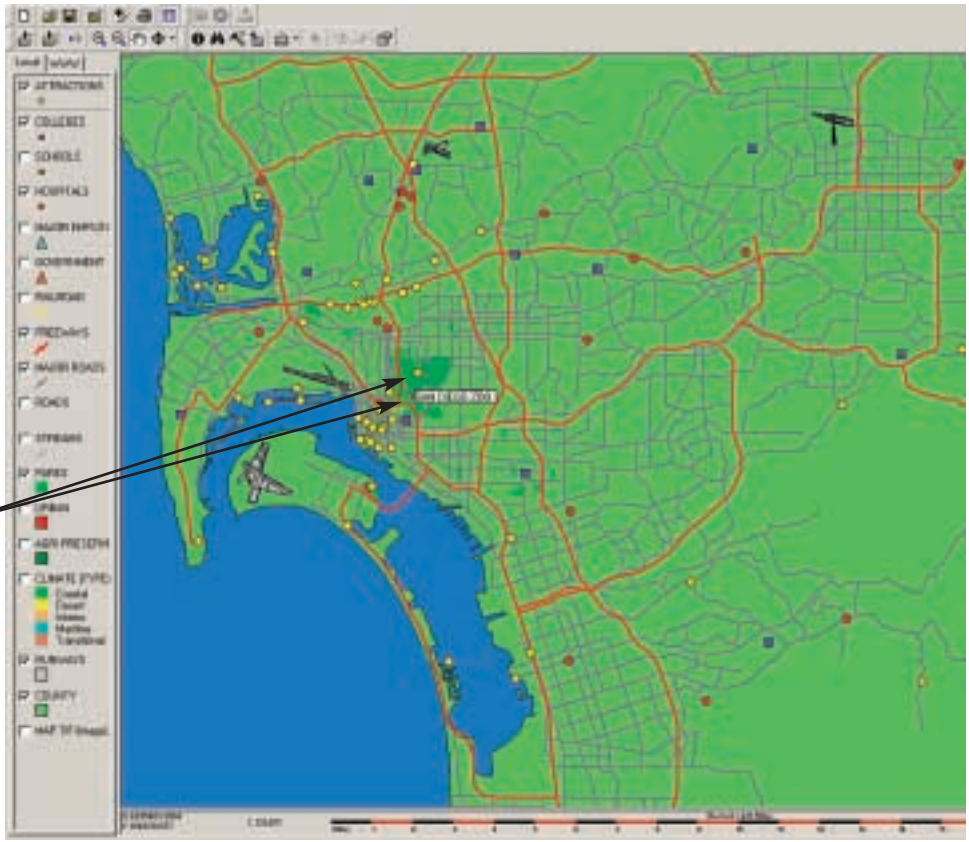

The ArcExplorer software allows you to choose which features you want to view. The Active Theme, which in this case is ATTRACTIONS, allows you to see the name of attraction by placing the cursor on the symbol.

The *Miami Herald* plotted the arrival point of the hurricane where winds were strongest and charted the progress of Andrew inland, where it lost some of its initial punch. Then the reporters plotted the damaged houses with one dot representing 10 homes. These dots were color-coded to show level of damage—blue showed 10 repairable homes and orange dots represented 10 destroyed homes. When the two maps were laid over each other, it was clear that the wind strength did not match up with damage as it should have. That is, the reporters showed that Hurricane Andrew alone was not responsible for all the devastation. The net result was that building codes in Dade County were tightened so that contractors had to use more nails and install stronger windows, doors, and shutters on homes. The *Miami Herald* received the Pulitzer Prize for its investigative work.

Businesses use GIS software too to analyze information, generate business intelligence, and make decisions. GIS is a powerful combination of database and graphics technology. There is virtually no limit to the sort of information you can plot with a GIS, including the placement of roads, the course of rivers, income levels, health conditions, areas of high or low crime, and so on. Of course, you can do this with paper maps too. The strength of an electronic GIS is in the ability you get to layer information with a mouse click.

Figure 4.6 is a map of San Diego shown in ESRI's ArcExplorer 2 GIS software. On the left-hand side you can see the layers (or type of information) that can be shown on the map. In our figure we have shown the *Attractions, Colleges, Hospitals, Freeways, Major Roads, Parks,* and *Runways.* Each type of information has its own symbol and is called a *theme.* When a theme is active, as *Attractions* is in our example, and you move the cursor over the little yellow circle, you can see the name of the attraction—in our case, the San Diego Zoo.

This is an example of the feature of GIS software that represents information thematically (i.e., in map or theme form). With themes, you can show the layers in combination as we did in our example and as the *Miami Herald* reporters did for Hurricane Andrew. You can represent either statistical information, such as the average salary of homeowners, or point information, such as the location of a bank's customers.

When businesses use GIS software to generate maps showing information of interest to them, we call it *business geography*. You can find a wealth of information to incorporate into your GIS from various sources. For example, the U.S. Census Bureau has a vast database of demographic information and the Bureau of Labor Statistics has employment information. Both of these would be statistical information. In the private sector, many research companies would be happy to sell you consumer habit information.

Studies show that how information is presented significantly impacts the effectiveness and efficiency of the decision-making process.[10] Here are other examples of GISs in use (see also this chapter's Closing Case Study One for more examples):

- Clean Harbor is the company that hauled away potentially anthrax-tainted debris from the offices of the NCB during the anthrax scare of 2001. Since hazardous waste removal is such a dangerous job, the company had to keep track of the shipment every minute. Clean Harbor used software that incorporated a GIS map and GPS information to monitor the trucks along their journey.[11]
- Branson, Missouri, is a small town in the Ozark Mountains that boasts more than two dozen theaters offering 75 performances per day to the 5 to 7 million visitors the town hosts every year. The town's motto is "the show must go on." But the show can't go on without electricity. The White River Valley Electric Cooperative (WRVEC) that supplies power to the area has a Web-based GIS where employees can access all the information they need, such as the precise location of poles, meters, transformers, capacitors, and underground facilities.[12]

Artificial Intelligence

DSSs and GISs are IT systems that augment business brainpower. IT can further expand business brainpower by means of artificial intelligence—the techniques and software that enable computers to mimic human behavior in various ways. Financial analysts use a variety of artificial intelligence systems to manage assets, invest in the stock market, and perform other financial operations.[13] Hospitals use artificial intelligence in many capacities, from scheduling staff, to assigning beds to patients, to diagnosing and treating illness. Many government agencies use artificial intelligence, including the IRS and the armed forces. Credit card companies use artificial intelligence to detect credit card fraud, and insurance companies use artificial intelligence to spot fraudulent claims.[14] Artificial intelligence lends itself to tasks as diverse as airline ticket pricing, food preparation, oil exploration, and child protection. It is widely used in the insurance, meteorology, engineering, and aerospace industries and by the military. It was artificial intelligence that guided cruise missiles during the Gulf War in 1991.[15]

Artificial intelligence (AI) is the science of making machines imitate human thinking and behavior. For example, an expert system is an artificial intelligence system that makes computers capable of reasoning through a problem to reach a conclusion. We use the process of reasoning to find out, from what we already know, something that we don't know.

Today computers can see, hear, smell, and, important for business, think (in a manner of speaking). Robots are a well-known form of AI. A ***robot*** is a mechanical device equipped with simulated human senses and the capability of taking action on its own (in contrast to a mechanical device such as an automobile, which requires direction from the

driver for its every action). Robots are in use in many industries. For example, Piedmont Hospital's Pharmacy Dosage Dispenser is a robotic prescription-filling system. Using bar code technology, this pharmaceutical robot receives medication orders online, retrieves prepackaged doses of drugs, and sends them to hospital patients.[16] One of the most exciting new areas of research in robotics is the development of microrobots that can be introduced into human veins and arteries to perform surgery.

A recent U.S. Commerce Department survey reported that 70 percent of the top 500 companies use artificial intelligence as part of decision support, and the sale of artificial intelligence software is rapidly approaching the $1 billion mark. The AI systems that businesses use most can be classified into the following major categories:

- *Expert systems,* which reason through problems and offer advice in the form of a conclusion or recommendation.
- *Neural networks,* which can be "trained" to recognize patterns, and *fuzzy logic,* which is a way of representing ambiguous or subjective information in computerized analysis.
- *Genetic algorithms,* which can generate increasingly better solutions to problems by generating many, many solutions, choosing the best ones, and using those to generate even better solutions.
- *Intelligent agents,* which are adaptive systems that work independently, carrying out specific, repetitive, or predictable tasks.

Expert Systems

Suppose you own a real estate business and generate more than 40 percent of your revenue appraising commercial real estate. Suppose only one person in your firm is capable of performing these appraisals. What if that person quits? How do you replace her or his expertise? How fast can you find someone else? How much business will you lose if it takes you a month to find a suitable replacement?

In business, people are valuable because they perform important business tasks. Many business tasks require expertise, and people often carry this expertise in their heads—often that's the only place it can be found in the organization. AI can provide you with an expert system that can capture expertise, thus making it available to those who are not experts so that they can use it, either to solve a problem or to learn how to solve a problem.

An ***expert system,*** also called a ***knowledge-based system,*** is an artificial intelligence system that applies reasoning capabilities to reach a conclusion. Expert systems are excellent for diagnostic and prescriptive problems. Diagnostic problems are those requiring an answer to the question, "What's wrong?" and correspond to the intelligence phase of decision making. Prescriptive problems are those that require an answer to the question, "What to do?" and correspond to the choice phase of decision making.

An expert system is usually built for a specific application area called a *domain.* You can find expert systems in the following domains, among others:

- *Accounting*—for auditing, tax planning, management consulting, and training.
- *Medicine*—to prescribe antibiotics where many considerations must be taken into account (such as the patient's medical history, the source of the infection, and the price of available drugs).

INDUSTRY PERSPECTIVE

WHAT MEDICAL PROCEDURE WOULD SUIT YOU?

In the opening case study for this chapter you saw how the Mayo Clinic and the Cleveland Clinic are developing IT systems to help doctors and other health care workers decide on the best treatment for their patients, but what about the patients themselves? Health care professionals consistently advise us to be informed partners with our providers and to take responsibility for our own health care.

To help people do that, HealthShare (www.HealthShare.com) provides decision support in the form of an expert system with information about surgery, risks, specialist credentials, posttreatment options, and recovery times. This information is gleaned from Medicare data on 177 diagnoses and procedures in 14 medical and surgical areas, including adult asthma, diabetes, heart valve replacement, and strokes. Links to this site can be found on the sites of health insurance company Web sites like CIGNA HealthCare, for example, where a member who needs a hip replacement can review options such as the locations of hospitals where the surgery can be performed.

The CIGNA subscriber can also find out the volume of hip replacements done at each hospital and even mortality rates. This figure is adjusted for the severity of illness so that hospitals that treat the most difficult cases don't appear more careless just because they're the last stop in a long battle with a deadly illness.

The system also suggests topics or questions that potential patients should discuss with their doctors prior to deciding on a course of treatment. HealthShare emphasizes that its mission is to provide information and analysis to help people think, and not to provide answers to what are highly personal, complex decisions.[17]

- *Process control*—to control offset lithographic printing, for example.
- *Human resource management*—to help personnel managers determine whether they are in compliance with an array of federal employment laws.
- *Financial management*—to identify delinquency-prone accounts in the loan departments of banks.
- *Production*—to guide the manufacture of all sorts of products, such as aircraft parts.
- *Forestry management*—to help with harvesting timber on forest lands.

A DSS sometimes incorporates expert systems, but an expert system is fundamentally different from a DSS. To use a DSS, you must have considerable knowledge or expertise about the situation with which you're dealing. As you saw earlier in this chapter, a DSS *assists* you in making decisions. That means that you must know how to reason through the problem. You must know which questions to ask, how to get the answers, and how to proceed to the next step. When you use an expert system, however, the know-how is in the system—you need only provide the expert system with the facts and symptoms of the problem for which you need an answer. The know-how, or expertise, that actually solves the problem came from someone else—an expert in the field. What does it mean to have expertise? When someone has expertise in a given subject, that person not only knows a lot of facts about the topic but also can apply that knowledge to analyze and make judgments about related topics. It's this human expertise that an expert system captures.

Rule	Symptom or Fact	Yes	No	Explanation
1	Is the light green?	Go through the intersection.	Go to Rule 2.	Should be safe if light is green. If not, need more information.
2	Is the light red?	Go to Rule 4.	Go to Rule 3.	Should stop, may not be able to.
3	Is the light likely to change to red before you get through the intersection?	Go to Rule 4.	Go through the intersection.	Will only reach this point if light is yellow, then you'll have two choices.
4	Can you stop before entering the intersection?	Stop.	Go to Rule 5.	Should stop, but there may be a problem if you can't.
5	Is traffic approaching from either side?	Prepare to crash.	Go through the intersection.	Unless the intersection is clear of traffic, you're likely to crash.

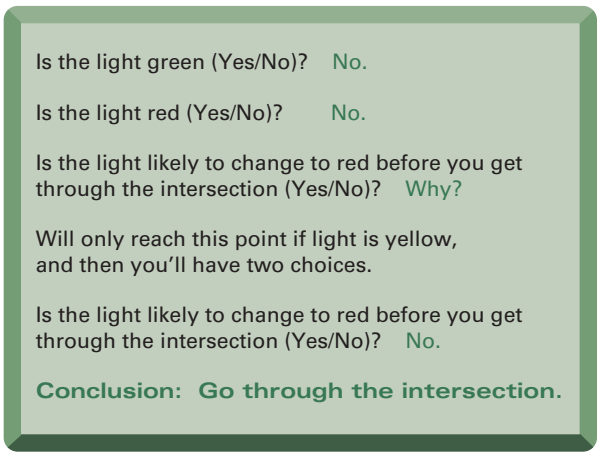

Figure 4.7

Traffic Light Expert System Rules

Let's look at a very simple expert system that would tell a driver what to do when approaching a traffic light. Dealing with traffic lights is an example of the type of problem to which an expert system is well-suited. It is a recurring problem, and to solve it you follow a well-defined set of steps. You've probably gone through the following mental question-and-answer session hundreds of times without even realizing it (see Figure 4.7).

When you approach a green traffic light, you proceed on through. If the light is red, you need to stop. If you're unable to stop, and if traffic is approaching from either side, you'll surely be in trouble. Similarly, if the light is yellow, you may be able to make it through the intersection before the light turns red. If not, you will again be faced with the problem of approaching traffic.

WHAT EXPERT SYSTEMS CAN AND CAN'T DO

An expert system uses IT to capture and apply human expertise. For problems with clear rules and procedures, expert systems work very well and can provide your company with great advantages. An expert system can

- Handle massive amounts of information
- Reduce errors
- Aggregate information from various sources
- Improve customer service

ON YOUR OWN

TRAFFIC LIGHTS REVISITED

Create a table similar to Figure 4.7 to extend the traffic light expert system. Include the following situations in the table:

1. There is a wreck in the middle of the intersection.
2. You are turning left at the intersection.
3. You are turning right at the intersection.
4. A pedestrian is crossing in front of you.
5. A dog has wandered into the intersection.
6. A ball belonging to children playing near the intersection has rolled into the street.
7. The car in front of you has stalled.

- Provide consistency in decision making
- Provide new information
- Decrease personnel time spent on tasks
- Reduce cost

You can run into trouble, however, in building and using an expert system. Difficulties can include the following:

1. Transferring domain expertise to the expert system is sometimes difficult because domain experts cannot always explain how they know what they know. Often experts are not aware of their complete reasoning processes. Experience has given them a feel for the problem, and they just "know."
2. Even if the domain expert can explain the whole reasoning process, automating that process may be impossible. The process may be too complex, requiring an excessive number of rules, or it may be too vague or imprecise. In using an expert system, keep in mind that it can solve only the problems for which it was designed. It cannot deal with inconsistency or a newly encountered problem situation. An expert system can't learn from previous experience and can't apply previously acquired expertise to new problems the way humans can.
3. An expert system has no common sense or judgment. One of the early expert systems built into an F-16 fighter plane allowed the pilot to retract the landing gear while the plane was still on the ground and to jettison bombs while the plane was flying upside down.

Neural Networks and Fuzzy Logic

Suppose you see a breed of dog you've never encountered before. Would you know it's a dog? For that matter, would you know it's an animal? Probably so. You know, because you've learned by example. You've seen lots of living things, have learned to classify them, and so can recognize a dog when you see one. A neural network simulates this human ability to classify things without taking prescribed steps leading to the solution. A ***neural network*** (often called an ***artificial neural network*** or ***ANN***) is an artificial intelligence system that is capable of finding and differentiating patterns. Your brain has learned to consider many factors in combination to recognize and differentiate objects. This is also the case with a neural network. A neural network can learn by example and can adapt to new concepts and knowledge. Neural networks are widely used for visual

TEAM WORK

HOW WOULD YOU CLASSIFY PEOPLE?

Some people have suggested that neural networks could be applied to people to indicate how likely they are to develop disease or even become criminals. The idea is to input a child's personal characteristics, demographics, and genealogy into a neural network, and the neural network will classify that youngster as being at risk for a disease or for aberrant behavior.

Choose either susceptibility to disease or to criminal behavior, discuss it with your group, and make the following lists, explaining why you chose each one.

1. What personal characteristics would be useful?
2. What demographic factors would strongly influence a person's future?
3. What, if any, inherited characteristics can predict a child's future?

Would such classification on a large scale be legal? Would it be ethical? Would it be effective? Why or why not (to all three questions)?

pattern and speech recognition systems. If you've used a PDA that deciphered your handwriting, it was probably a neural network that analyzed the characters you wrote.[18]

Neural networks are useful in a variety of situations. For example, bomb detection systems in U.S. airports use neural networks that sense trace elements in the air that may indicate the presence of explosives. The Chicago Police Department uses neural networks to identify corruption within its ranks.[19] In medicine, neural networks check 50 million electrocardiograms per year, check for drug interactions, and detect anomalies in tissue samples that may signify the onset of cancer and other diseases. Neural networks can detect heart attacks and even differentiate between the subtly different symptoms of heart attacks in men and women.[20,21,22] In business, neural networks are very popular for securities trading, fraud detection, real estate appraisal, evaluating loan applications, and target marketing, to mention a few. Neural networks are used to control machinery, adjust temperature settings, and identify malfunctioning machinery.

Neural networks are most useful for identification, classification, and prediction when a vast amount of information is available. By examining hundreds, or even thousands of examples, a neural network detects important relationships and patterns in the information. For example, if you provide a neural network with the details of numerous credit card transactions and tell it which ones are fraudulent, eventually it will learn to identify suspicious transaction patterns.

Here are some examples of the uses of neural networks:

- Many banks and financial institutions use neural networks. Citibank uses neural networks to find opportunities in financial markets.[23] By carefully examining historical stock market data with neural network software, Citibank financial managers learn of interesting coincidences or small anomalies (called market inefficiencies). For example, it could be that whenever IBM stock goes up, so does Unisys stock. Or it might be that a U.S. Treasury note is selling for 1 cent less in Japan than it is in the United States. These snippets of information can make a big difference to Citibank's bottom line in a very competitive financial market.
- In Westminster, California, a community of 87,000 people, police use neural network software to fight crime. With crime reports as input, the system detects and maps local crime patterns. Police say that with this system they can better predict crime trends, improve patrol assignments, and develop better crime-prevention programs.[24]

INDUSTRY PERSPECTIVE

NEURAL NETWORKS TO SAVE ENERGY

After the energy crisis that began in California in 1999, many companies, including the Bank of America, started to look for energy management systems. The bank installed a neural network–based system to control energy consumption in its 78 California office buildings. The neural network allows managers to float heating and air conditioning controls so that temperatures remain comfortable for the office occupants without being hotter or colder than necessary.

The Bank hoped to save 5 percent on energy cost but surpassed that figure and paid for the system within a year. Not only is the organization able to save energy, it now has a better understanding of its energy consumption.

The California Department of General Services also uses neural network software to save energy. The software was installed in 65 state-owned offices in 2001 and the agency found that it could save up to 20 percent annually on its energy costs.[25,26]

- Fingerhut, the mail order company based in Minnesota, has 6 million people on its customer list. To determine which customers were and were not likely to order from its catalog, Fingerhut recently switched to neural network software. The company finds that the new software is more effective and expects to generate millions of dollars by fine-tuning its mailing lists.[27]
- Fraud detection is one of the areas in which neural networks are used the most. Visa, MasterCard, and many other credit card companies use a neural network to spot peculiarities in individual accounts. MasterCard estimates neural networks save them $50 million annually.[28]
- Many insurance companies (Cigna, AIG, Travelers, Liberty Mutual, Hartford) along with state compensation funds and other carriers use neural network software to identify fraud. The system searches for patterns in billing charges, laboratory tests, and frequency of office visits. A claim for which the diagnosis was a sprained ankle and which included an electrocardiogram would be flagged for the account manager.[29]
- FleetBoston Financial Corporation uses a neural network to watch transactions with customers. The neural network can detect patterns that may indicate a customer's growing dissatisfaction with the company. The neural network looks for signs like decreases in the number of transactions or in the account balance of one of Fleet's high-value customers.[30]

All of the above situations have pattern recognition in common. They all require identification and/or classification, which may then be used to predict a finding or outcome. Neural networks are often called predictive systems since they can see patterns in huge volumes of information.

INSIDE A NEURAL NETWORK

Neural networks are so called because they attempt to mimic the structure and functioning of the human brain. Conceptually, neural networks consist of three layers of virtual nerve cells, or neurons. There's an input layer and an output layer and between them is a hidden layer, although there may be more than one hidden layer. The input and output layers are connected to the middle layer(s) by connections called weights of various

Figure 4.8

The Layers of a Neural Network

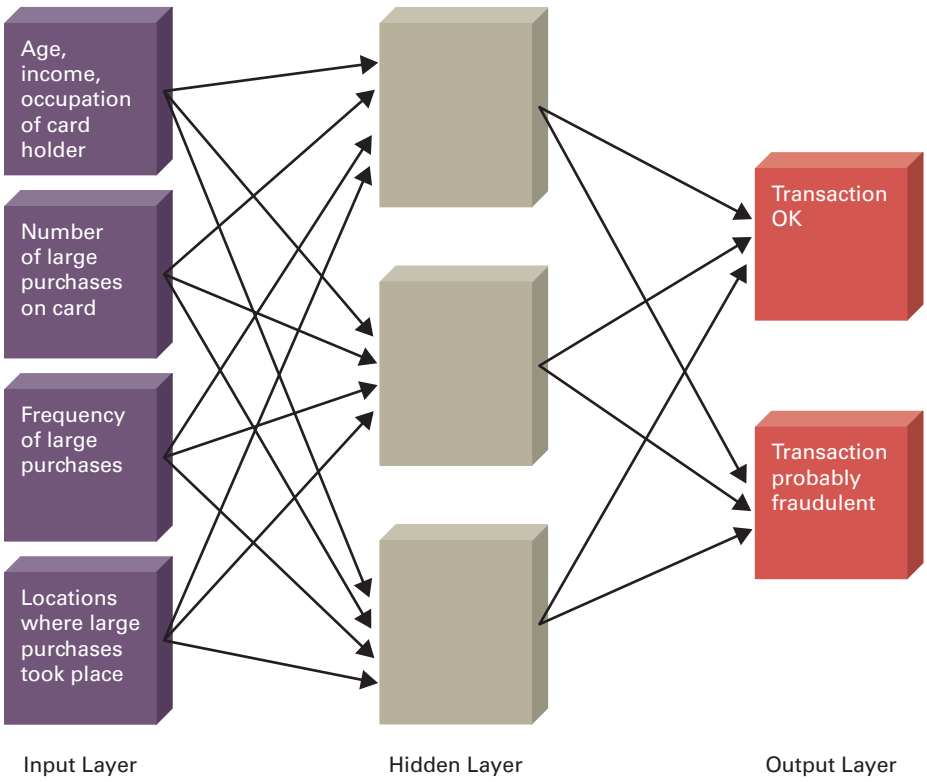

strengths (see Figure 4.8). If you were to train a neural network to recognize a "good" stock portfolio, you would input many, many examples of good and bad portfolios, telling the neural network which was which. As the neural network is learning to differentiate between good and bad, the weights change. The flow of information to the output layer also changes. After you have fed the system enough examples, the weights stabilize, and the neural network then consistently classifies portfolios correctly.

So, you may be asking, how is a neural network different from an expert system, since both can take input and produce an answer as to which group the input belongs? An expert system, as we saw, can also classify; that is, it asks questions and, based on the answers, can diagnose or prescribe. The difference is that an expert system does not adjust by itself and is rigid in its application of the rules. For example, if a credit card fraud detection expert system had a rule that said to flag a purchase over a certain amount on certain types of accounts, the expert system would flag a transaction that was even one penny over. A neural network, on the other hand, would learn the spending behavior of cardholders and would be better able to evaluate whether deviations were large enough to be queried or not. A neural network can even adjust to situations not explicitly used in training. For example, if when the neural network was learning, mortgage rates were between 6 percent and 10 percent, the system could interpolate if the rate were to drop to 5 percent.

Neural networks have many advantages. For example, neural networks can

- Learn and adjust to new circumstances on their own.
- Lend themselves to massive parallel processing.
- Function without complete or well-structured information.
- Cope with huge volumes of information with many dependent variables.
- Analyze nonlinear relationships in information (they've been called fancy regression analysis systems).

The biggest problem with neural networks to date has been the fact that the hidden layers are "hidden." That is, you can't see how the neural network is learning and how the neurons are interacting. Newer neural networks no longer hide the middle layers. With these systems you can manually adjust the weights or connections giving you more flexibility and control.

FUZZY LOGIC

Fuzzy logic is a way of reaching conclusions based on ambiguous or vague information. Humans tend to make decisions based on approximate information, since not every type of information can be separated into mutually exclusive categories. For example, you might consider 20 degrees Fahrenheit to be very cold, 40 degrees to be cold, 60 degrees to be warm, 90 degrees to be hot, and 105 to be very hot, but you can't really specify exact degree values that would separate these categories of warmth (or lack of it). Compounding the problem is the fact that somebody else might experience it differently; subjectively, a temperature some people consider warm others would judge to be hot—or cool. Thus, you need a way to translate an approximate or vague judgment into something that a computer, requiring a precise assignment of numbers to all events, can handle. This is the type of situation in which fuzzy logic is very effective.

Fuzzy logic is a mathematical method of handling imprecise or subjective information. The basic approach is to assign values between 0 and 1 to vague or ambiguous information. The higher the value, the closer it is to 1. For example, you might assign the value of 0.8 to the value "hot." Then you would construct rules and processes, called *algorithms,* to describe the interdependence among variables. A fuzzy logic algorithm is a set of steps that relate variables representing inexact information or personal perceptions.

Fuzzy logic and neural networks are often combined to express complicated and subjective concepts in a form that makes it possible to simplify the problem and to apply rules. The rules are executed with a level of certainty. This is similar to, but not the same as, confidence levels in statistics. In statistics, probability is used to estimate the likelihood of an outcome, whereas fuzzy logic describes the data point itself while incorporating subjective perception.

In the business world, fuzzy logic has been applied to financial analysis, the pharmaceutical industry, the wood processing and metal cutting industries, the manufacture of antilock brakes, and washing machines that determine by themselves how much water to use or how long to wash. (They wash until the water is "clean.") In accounting and finance, for example, fuzzy logic allows you to analyze information with subjective financial values (say, on an important intangible resource like goodwill) that are very important considerations in economic analyses.

Fuzzy logic is used by Google to find answers to your search terms, which makes sense, since your perception of a topic often influences how you phrase your query, hence determining the relevance of the Web pages that Google delivers.

Genetic Algorithms

Have you ever wondered how chefs around the world create recipes for great-tasting foods? For example, how did the Chinese discover that cashew nuts and chicken taste good when combined? How did Mexican chefs arrive at combining tomatoes, onions, cilantro, and other spices to create pica de gallo? All those great recipes came about through *evolutionary processes.* Someone decided to put together a few ingredients and taste the result. Undoubtedly, many of those combinations resulted in unpalatable

concoctions that were quickly discarded. Others were tasty enough to warrant further experimentation of combinations.

Today significant research in AI is devoted to creating software capable of following a similar trial-and-error process, leading to the evolution of a good result. Such a software system is called a genetic algorithm. A ***genetic algorithm*** is an artificial intelligence system that mimics the evolutionary, survival-of-the-fittest process to generate increasingly better solutions to a problem. In other words, a genetic algorithm is an optimizing system: It finds the combination of inputs that give the best outputs.

Here's an example. Suppose you were trying to decide what to put into your stock portfolio. You have countless stocks to choose from but a limited amount of money to invest. You might decide that you'd like to start with 20 stocks and you want a portfolio growth rate of 7.5 percent.

Probably you'd start by examining historic information on the stocks. You would take some number of stocks and combine them, 20 at a time, to see what happens with each grouping. If you wanted to choose from a pool of 30 stocks, you would have to examine 30,045,015 different combinations. For a 40-stock pool, the number of combinations rises to 137,846,500,000. It would be an impossibly time-consuming, not to mention numbingly tedious, task to look at this many combinations and evaluate your overall return for each one. This is just the sort of repetitive number-crunching task at which computers excel, however.

So, instead of a pencil, paper, and calculator, you might use a genetic algorithm. You could input the appropriate information on the stocks, including the number of years the company has been in business, the performance of the stock over the last five years, price to earnings ratios, and other information.

You would also have to tell the genetic algorithm your exact "success" criteria. For example, you might use a growth rate in the company over the last year of at least 10 percent, a presence in the marketplace going back at least three years, a connection to the computer industry, and so forth. The genetic algorithm would simply combine and recombine stocks eliminating any combinations that don't fit your criteria and continuing to the next iteration with the acceptable combinations—those that give an aggregate growth rate of at least 7.5 percent while aiming for as high a growth rate as possible.

Genetic algorithms use three concepts of evolution:

1. ***Selection***—or survival of the fittest. The key to selection is to give preference to better outcomes.
2. ***Crossover***—or combining portions of good outcomes in the hope of creating an even better outcome.
3. ***Mutation***—or randomly trying combinations and evaluating the success (or failure) of the outcome.

Genetic algorithms are best suited to decision-making environments in which thousands, or perhaps millions, of solutions are possible. Genetic algorithms can find and evaluate solutions intelligently and can get through many more possibilities more thoroughly and faster than a human can. As you might imagine, businesses face decision-making environments for all sorts of problems like engineering design, computer graphics, strategies for game playing, anything, in fact, that requires optimization techniques. Here are some other examples.

- Genetic algorithms are used by business executives to help them decide which combination of projects a firm should invest in, taking complicated tax considerations into account.[31]
- They're used by investment companies to help in trading choices and decisions.[32]

TEAM WORK

BE A GENETIC ALGORITHM AND PUT NAILS IN BOXES

This project involves packaging nails so that you make the most profit possible (this is a profit maximizing problem). Say you have six types of nails and can make as many as you need of each. These are 4-inch, 3.5-inch, 3-inch, 2.5-inch, 2-inch, and 1.5-inch nails. The cost of making each type of nail depends on how big a nail it is. Those cost and selling prices are listed in the table below along with the weights. The nails will be sold in boxes of up to 30 nails. There must be no more than 10, but no less than 5, of each of three types of nails in each box. The nails in each box should weigh no more than 20 ounces. You're looking for the combination with the highest profit using a trial-and-error method.

A spreadsheet would be helpful for completing this project. You'll most likely find that you identify some promising paths to follow right away and will concentrate on those to reach the best one.

Nail	Weight	Cost	Selling price
4 inch	1 oz	4 cents	8 cents
3.5 inch	0.85 oz	3.5 cents	6 cents
3 inch	0.7 oz	3 cents	5 cents
2.5 inch	0.5 oz	2.5 cents	4 cents
2 inch	0.25 oz	2 cents	3 cents
1.5 inch	0.1 oz	1.5 cents	2 cents

- In any garment that you buy, the fabric alone accounts for between 35 percent and 40 percent of the selling price. So, when cutting out the fabric to make the garment, it's important that there be as little waste as possible. Genetic algorithms are used to solve this problem of laying out the pieces of the garment and cutting fabric in a way that leaves as little waste as possible.[33]
- US West uses a genetic algorithm to determine the optimal configuration of fiber-optic cable in a network that may include as many as 100,000 connection points. By using selection, crossover, and mutation, the genetic algorithm can generate and evaluate millions of cable configurations and select the one that uses the least amount of cable. At US West, this process used to take an experienced design engineer almost two months. US West's genetic algorithm can solve the problem in two days and saves the company $1 million to $10 million each time it's used.[34]

Genetic algorithms are good for these types of problems because they use selection, crossover, and mutation as methods of exploring countless solutions and the respective worth of each.

You have to tell the genetic algorithm what constitutes a "good" solution. That could be low cost, high return, among other factors, since many potential solutions are useless or absurd. If you created a genetic algorithm to make bread, for example, it might try to boil flour to create moistness. That obviously won't work, so the genetic algorithm would simply throw away that solution and try something else. Other solutions would eventually be good, and some of them would even be wonderful. According

to David Goldbert, a genetic algorithm pioneer at the University of Illinois at Urbana–Champaign, evolution is the oldest and most powerful algorithm there is, and "three billion years of evolution can't be wrong!"[35]

Intelligent Agents

Do you have a favorite restaurant? Is there someone there who knows you and remembers that you like Italian dressing, but not croutons, on your salad; and ice cream and a slice of cheddar cheese with your apple pie? Does this person familiar with your tastes put a glass of diet cola on your favorite table when you come in the door? If so, he or she has the qualities that artificial intelligence scientists are working on incorporating into intelligent agents. An ***intelligent agent*** is software that assists you, or acts on your behalf, in performing repetitive computer-related tasks. Future intelligent agents will most likely be autonomous, acting independently, and will learn and adapt to changing circumstances.

You may not realize it, but you're probably already familiar with a primitive type of intelligent agent—the shifty-eyed paper clip that pops up in some versions of Word. For example, if your document looks as if it is going to be a business letter—that is, you type in a date, name, and address—the animated paper clip will offer helpful suggestions on how to proceed.

You can find hundreds of intelligent agents, or bots, for a wide variety of tasks. The BotSpot and SmartBot Web sites at www.botspot.com and www.smartbots.com are good places to get an idea of the many different types of agents that are available.

Essentially there are four types of intelligent agents:

- Information agents (including buyer agents or shopping bots)
- Monitoring-and-surveillance agents
- Data-mining agents
- User or personal agents

INFORMATION AGENTS

Information agents are intelligent agents that search for information of some kind and bring it back. The best known information agents are buyer agents. A ***buyer agent,*** also known as a ***shopping bot,*** is an intelligent agent on a Web site that helps you, the customer, find products and services that you need. They work very efficiently for commodity products such as CDs, books, electronic components, and other one-size-fits-all products. Amazon.com uses intelligent technology to show you a list of books or other products that you might like. The Web site classifies you into a category of people with similar tastes and, based on that category, presents you with a list of products that Amazon hopes you will find appealing enough to buy.

Shopping bots make money by selling advertising space, from special promotions in cooperation with merchants, or by charging click-through fees, which are payments to the site that provided the link to the merchant site. Some shopping bots give preference to certain sites for a financial consideration. The people who run shopping bot sites have two, sometimes competing, objectives. They want to present as many listings as possible to the consumer in the most useful way, but they also want to make money doing it.

MySimon.com is the most successful shopping bot to date with more than a million visitors a month according to Nielsen/NetRatings. MySimon searches for millions of products on thousands of Web sites.[36]

Both the Google and Ask Jeeves Web sites use information agents to find information—and not just when you request it. The URL server at Google sends thousands of

INDUSTRY PERSPECTIVE

INTELLIGENT AGENTS RUN EXPERIMENTS ABOARD SPACECRAFT

We all know that NASA sends various types of craft into space for a variety of purposes. We've heard of the Hubble telescope that has been sending back terrific pictures for years. A lesser known but similarly spectacular observation craft—although in this case its mission is to observe the earth—is the Earth Observing-1 (or simply EO1) satellite.

During 2004, NASA uploaded intelligent agent software to EO1 to run experiments and even the spacecraft itself. One of EO1's tasks was to avoid wasting fuel without disrupting the onboard experiments or otherwise compromising the mission. This was a job that used to be handled by ground control, but now the people at mission headquarters are free to concentrate on tasks other than routine maintenance while the agents do the job in space.

EO1 was launched in 2001 on a one-year mission to observe the earth from space. It was part of NASA's New Millennium Program mission and was originally intended as a pilot project to test new space technologies. Now, four years later it's still going strong collecting and sending back to us valuable information about our earth with about 20 times more detail than any previous Earth-observing satellites. The information that EO1 sends back encompasses all manner of happenings on earth, like the spread of forest and bush fires, the impact of cattle grazing in South America, the state of the rain forest, and the spread of harmful plant species.

The beauty of incorporating intelligent agents into the spacecraft is that they can learn and adapt to changing and unexpected conditions. NASA's software engineers have designed these agents to achieve certain goals rather than to react to prespecified situations, making them able to handle complex interactions. Since the software can "learn" and function autonomously, it can react to unexpected situations, allowing scientists to conduct more complex and interesting research on board.[37]

Googlebots out to surf the Web sites of all the sites in Google's index. They copy individual pages, at the rate of more than 100 per second, to Google's repository, where the Google software indexes them. This means that when you perform a Google search, the search engine builds a list of all the pages that have the keywords you specify and presents them to you in PageRank order. Google's PageRanking process sorts the pages in order of links pointing to each page. That is, the more links on the Web that point to a Web site, the higher that Web site will be in the list.[38]

Government sites have information agents you can use to get the information you need. FERRET (Federal Electronic Research and Review Extraction Tool) was developed jointly by the Census Bureau and the Bureau of Labor Statistics. With FERRET you can find information on employment, health care, education, race and ethnicity, health insurance, housing, income and poverty, aging, and marriage and family. Other types of information agents include intelligent agents that scan Web pages and highlight relevant text for you, and still others can assemble customized news reports. There are several versions of these. A CNN Custom News bot will gather news from CNN on the topics you want to read about—and only those.

MONITORING-AND-SURVEILLANCE AGENTS

Monitoring-and-surveillance agents (also called ***predictive agents***) are intelligent agents that constantly observe and report on some entity of interest, a network, or manufacturing equipment, for example. NASA's Jet Propulsion Laboratory has an agent that monitors inventory, planning, and the ordering of scheduled equipment to keep costs down.[39] Other monitoring-and-surveillance agents work on the manufacturing shop

floor, finding equipment problems and locating other machinery that can perform the same job.

Monitoring-and-surveillance agents are often used to monitor complex computer networks. Allstate Insurance has a network with 2,000 computers. The company uses a network monitoring agent from Computer Associates International called Neugent that watches its huge networks 24 hours a day. Every five seconds, the agent measures 1,200 data points and can predict a system crash 45 minutes before it happens. Neugent combines intelligent agent technology with neural network technology to look for patterns of activity or problems. The neural network part can learn what conditions predict a downturn in network efficiency or a slowing in network traffic. Neugent also watches for electronic attacks and can detect them early so that they can be stopped.

Another type of monitoring-and-surveillance agent is one that works on computer networks keeping track of the configuration of each computer connected to the network. It tracks and updates the central configuration database when anything on any computer changes, like the number or type of disk drives. An important task in managing networks is prioritizing traffic and shaping bandwidth. That means sending enough network capacity or bandwidth to the most important tasks versus those that are secondary. At a university, for example, processing end-of-semester grades might take precedence.

Some other types of monitoring-and-surveillance agents include

- Agents that watch your competition and bring back price changes and special offer information.
- Agents that monitor Internet sites, discussion groups, mailing lists, and so on, for stock manipulation, insider training, and rumors that might affect stock prices.
- Agents that monitor sites for updated information on the topic of your choice.
- Agents that watch particular products and bring back price or terms changes.
- Agents that monitor auction sites for products or prices that you want.

DATA-MINING AGENTS

A ***data-mining agent*** operates in a data warehouse discovering information. A data warehouse brings together information from lots of different sources. Data mining is the process of looking through the data warehouse to find information that you can use to take action—like ways to increase sales or to keep customers who are considering defecting. Data mining is so called because you have to sift through a lot of information for the gold nuggets that will affect the bottom line (or top line). This sort of nugget spotting is similar to what the FBI and CIA do when they bring together little bits of information from diverse sources and use the overall pattern to spot trouble brewing.

As you learned in Chapter 3, database queries answer questions like "How much did we spend on transportation in March of this year?" Multidimensional analysis is the next step in complexity and answers questions like "How much did we spend on transportation in the southeast during March of the last five years?" Data mining goes deeper and may suggest questions you may not even have thought to ask like the retail manager we mentioned in Chapter 1 who thought "What else do young men buy on Friday afternoons when they come in to buy diapers?" (Remember from Chapter 1 the answer to that one, suggested by data-mining tools, was beer.)[40]

One of the most common types of data mining is classification, which finds patterns in information and categorizes items into those classes. You may remember that this is just what neural networks do best. So, not surprisingly, neural networks are part of many data-mining tools. And data-mining agents are another integral part, since these intelligent agents search for information in a data warehouse.

ON YOUR OWN

GO BARGAIN HUNTING ONLINE

Try out shopping bots for yourself. Choose three items to search for: one music item, one item of clothing, and one household item. Search for them with each of the following sites.

- Bottom Dollar at www.bottomdollar.com
- MySimon at www.mysimon.com
- R U Sure at www.rusure.com
- Yahoo! Shopping at shopping.yahoo.com
- Prescan at www.prescan.com

Answer these questions . . .

- How many hits did you get at each site for each item?
- Are tax, postage, and handling charges included in the quoted price?
- Can you sort in order of price?
- Does the shopping site specialize in a particular kind of item?

A data-mining agent may detect a major shift in a trend or a key indicator. It can also detect the presence of new information and alert you. Volkswagen uses an intelligent agent system that acts as an early-warning system about market conditions. If conditions become such that the assumptions underlying the company's strategy are no longer true, the intelligent agent alerts managers.[41] For example, the intelligent agent might see a problem in some part of the country that is about to or will shortly cause payments to slow down. Having that information early lets managers formulate a plan to protect themselves.

USER AGENTS

User agents (sometimes called ***personal agents***) are intelligent agents that take action on your behalf. In this category belong those intelligent agents that already perform, or will shortly perform, the following tasks:

- Check your e-mail, sort it according to priority (your priority), and alert you when good stuff comes through—like college acceptance letters.
- Play computer games as your opponent or patrol game areas for you.
- Fill out forms on the Web automatically for you. They even store your information for future reference.
- "Discuss" topics with you from your deepest fears to your favorite sports.

One expanding application of intelligent agent technology is in automating business functions. For example, Mission Hockey, a company that manufacturers and distributes in-line and ice hockey skates and other gear, uses software from Sweden called Movex that has a user-agent component. Movex will search the Internet or a company intranet or extranet to negotiate and make deals with suppliers and distributors. In this case, the intelligent agent is incorporated into an enterprise resource planning system. Enterprise resource planning (or ERP) is a very important concept in today's business world. The term refers to a method of getting and keeping an overview on every part of the business (a bird's-eye view, so to speak), so that production, development, selling, and servicing of goods and services will all be coordinated to contribute to the company's goals and objectives. You'll learn more about the concept of ERP in Chapter 7.

Multi-Agent Systems and Agent-Based Modeling

What do cargo transport systems, book distribution centers, the video game market, a flu epidemic, and an ant colony have in common? They are all complex adaptive systems and thus share some common characteristics. By observing parts of the ecosystem, like ant or bee colonies, artificial intelligence scientists can use hardware and software models that incorporate insect characteristics and behavior to (1) learn how people-based systems behave; (2) predict how they will behave under a given set of circumstances; and (3) improve human systems to make them more efficient and effective. This concept of learning from ecosystems and adapting their characteristics to human and organizational situations is called ***biomimicry.***

Complex organizations, those of humans and ants alike, are made up of separate groups of individuals. The groups have different roles, follow their own specific sets of rules, and react to changing conditions; what's more, not every individual in a group acts in precisely the same way.

Let's examine a simplified version of the retail video game market for instance. One group of people designs and writes the code for the games (and within that group are smaller subsets of people), a second group manufactures the CDs or DVDs and distributes them to the retail outlets (the third group), then the last group (consumers) choose to buy, or not to buy, the games to play on their gaming devices. What each group does influences the others.

The traditional approach to studying a system such as this was to examine each group of people and model that group according to its habits and behavior. However, this type of research has serious limitations because the interaction of one group with one or more of the others must be ignored, or studied separately. The difficulty of studying the system as a whole is that it's very complex, making it extremely difficult to keep track of all the variables and their interactions. The traditional model becomes unmanageable very quickly.

In the last few years, AI research has made much progress in modeling complex organizations as a whole with the help of multi-agent systems. In a ***multi-agent system*** groups of intelligent agents have the ability to work independently and to interact with each other. The simulation of a human organization using a multi-agent system is called agent-based modeling. ***Agent-based modeling*** is a way of simulating human organizations using multiple intelligent agents, each of which follows a set of simple rules and can adapt to changing conditions.

Agent-based modeling systems are being used to model stock market fluctuations, predict the escape routes that people seek in a burning building, estimate the effects of interest rates on consumers with different types of debt, and anticipate how changes in conditions will affect the supply chain, to name just a few. See Figure 4.9 for examples of companies that have used agent-based modeling to their advantage.

ANT COLONIES AND SWARM INTELLIGENCE

The ant ecosystem is one of the most widely used types of simulations in business problems. If you've ever tried to remove ants from your home, you know how determined and effective ant colonies are. Individual ants are autonomous, acting and reacting independently. (If you drop a crumb into the middle of a group of ants, they'll all scatter in different directions.) However, ants are unusual insects in that they are social. (Less than 2 percent of insects are social, with termites being the only other entirely social species,

- Southwest Airlines—to optimize cargo routing.
- Procter & Gamble—to overhaul its handling of what the company calls its "supply network" of 5 billion consumers in 140 countries.
- Air Liquide America—to reduce production and distribution costs of liquefied industrial gases.
- Merck & Co.—to find more efficient ways of distributing anti-AIDS drugs in Africa.
- Ford Motor Co.—to build a model of consumer preferences and find the best balance between production costs and customers' demands.
- Edison Chouest Offshore LLC—to find the best way to deploy its service and supply vessels in the Gulf of Mexico.

Figure 4.9

Companies That Use Agent-Based Modeling[42]

although some types of bees and wasps are, too.) The term "social" implies that all the members of a colony work together to establish and maintain a global system that's efficient and stable. So, even though the ants are autonomous, each ant contributes to the system as a whole. Ants have been on Earth for 40 million years, compared to the relatively short human occupation of 100 thousand years, and their extraordinary evolutionary success is the result of ants' collective behavior, known as swarm intelligence.

Swarm (collective) intelligence is the collective behavior of groups of simple agents that are capable of devising solutions to problems as they arise, eventually leading to coherent global patterns.[43] That is to say, complex collective behavior can result from the individuals in the system consistently following a small number of simple rules. Swarm intelligence allows the creation and maintenance of systems that have the following four characteristics:

1. *Flexibility,* so that the system can respond to changes, both large and small, in the environment around it. In an ant colony, for example, if you move the food, the ants will find it again very quickly.
2. *Robustness,* so that even if some individual members of the system don't succeed, the work gets done. For example, if you remove some of the ants, others will step in and continue the work.
3. *Decentralization,* in that each individual has a relatively simple job to do and performs that job without supervision. In the ant colony there are forager ants, soldier ants who protect the nest, queens who produce the new generations, ants who take care of and feed the cocoons, and so on.
4. *Self-organization,* in that the methods of problem solving are not prescribed from a central authority, but rather are developed as problem-solving strategies by the individuals in the group who are responsible for the completion of the work. For example, if an ant finds a food morsel that's too large for one ant to carry, others come to help and they run around changing positions until they have the morsel balanced well enough that they can carry it off. See the Web site at www.scottcamazine.com/personal/research/index.htm for other examples of self-organization in nature.[44]

So, how are the workings of ant colonies related to information technology in modern business? Swarm intelligence gives us a way to examine collective systems where groups of individuals have certain goals, solve problems, and make decisions without centralized control or a common plan. Think of our video game market. No one directs the whole system dictating that buyers buy specific games or that manufacturers sell a certain number of games or sell only in a certain region. The four groups of participants

act in ways that they believe will further their own goals, and the result of their interaction is a system with its own features that finds its own level of balance between supply and demand.

A comparison of the activities of forager ants and those of the cargo-handling arm of Southwest Airlines affords a striking example of the similarities between ecosystems and human organizations, which we will consider shortly. There are some uncanny parallels that surprised Southwest's management. First, though, let's ponder the ants.

Forager ants have the sole responsibility of providing food to the colony. They don't form committees and discuss strategies or look to a central authority for direction; they just find food and bring it back to the nest, and in doing so follow a simple procedure (see Figure 4.10).

Say two ants leave the same point to search for food. Figure 4.10, panel 1, shows that Ant A finds the food first because of the shorter path. Having found a food source, Ant A returns to the nest (see Figure 4.10, panel 2), leaving behind a trail of pheromones (a biological breadcrumb trail) so that Ant A and the other ants will know what path to take. The first ant that returns "lays the trail" first so that's the one that other ants take. Then the other ants strengthen the pheromone trail on their return journey by leaving their own pheromone tracks (see Figure 4.10, panel 3).

Meanwhile, Ant B arrives back at the nest after the shorter path has already been established. The other ants that are already on the move don't change their route. Additionally, the pheromone trail on the unused path evaporates after a certain length of time so that it's effectively deleted from the system as a desirable route to food. The approach is straightforward but effective, and can be expressed as the following rules:

- Rule 1: Follow the trail if one exists, otherwise create one.
- Rule 2: Find food.
- Rule 3: Return to the nest, making a pheromone trail.

If changes occur (say, for example, that the food source is removed), the ants cease returning to the place where the food used to be, and the trail disappears. Then the process begins again, and proceeds relentlessly, with forager ants finding a new food source and creating pheromone corridors that lead the way.

The problem that the ants have just solved is one of the oldest problems that humans (as well as ants) have faced. It's known as "the shortest path problem" or the "traveling salesman problem." Anyone who schedules drop-off and pick-up routes for delivery trucks, or schedules jobs on the factory floor, or even colors maps, making sure that no two adjacent components have the same color, has had to find a solution to the same type of problem.

Taking their cue from nature, AI researchers built sets of small robots and incorporated software that allowed the robots to follow rules and interact with each other in the same basic ways as the ants. They also dispensed with the physical forms altogether, creating virtual ants in the form of small autonomous blocks of code that we call intelligent agents. And each code block could follow certain rules, interact, and adapt. These virtual ants were then arranged into multi-agent systems that were further refined into agent-based models. Enter Southwest Airlines as a case in point.

Even though cargo is a small part of Southwest's business, it was causing management headaches and bottlenecks at busy airports. Southwest consulted with swarm intelligence experts, who used a virtual model of foraging ants to simulate the cargo-handling process. And that was how Southwest managers discovered, to their surprise, that there were actually better ways to handle cargo than to put it on the first plane flying in the right direction. Surprisingly, the computer's swarm intelligence model showed that it might actually be better to leave cargo on a plane heading in the wrong direction. For example, cargo headed from Chicago to Boston would be better left on a plane going from

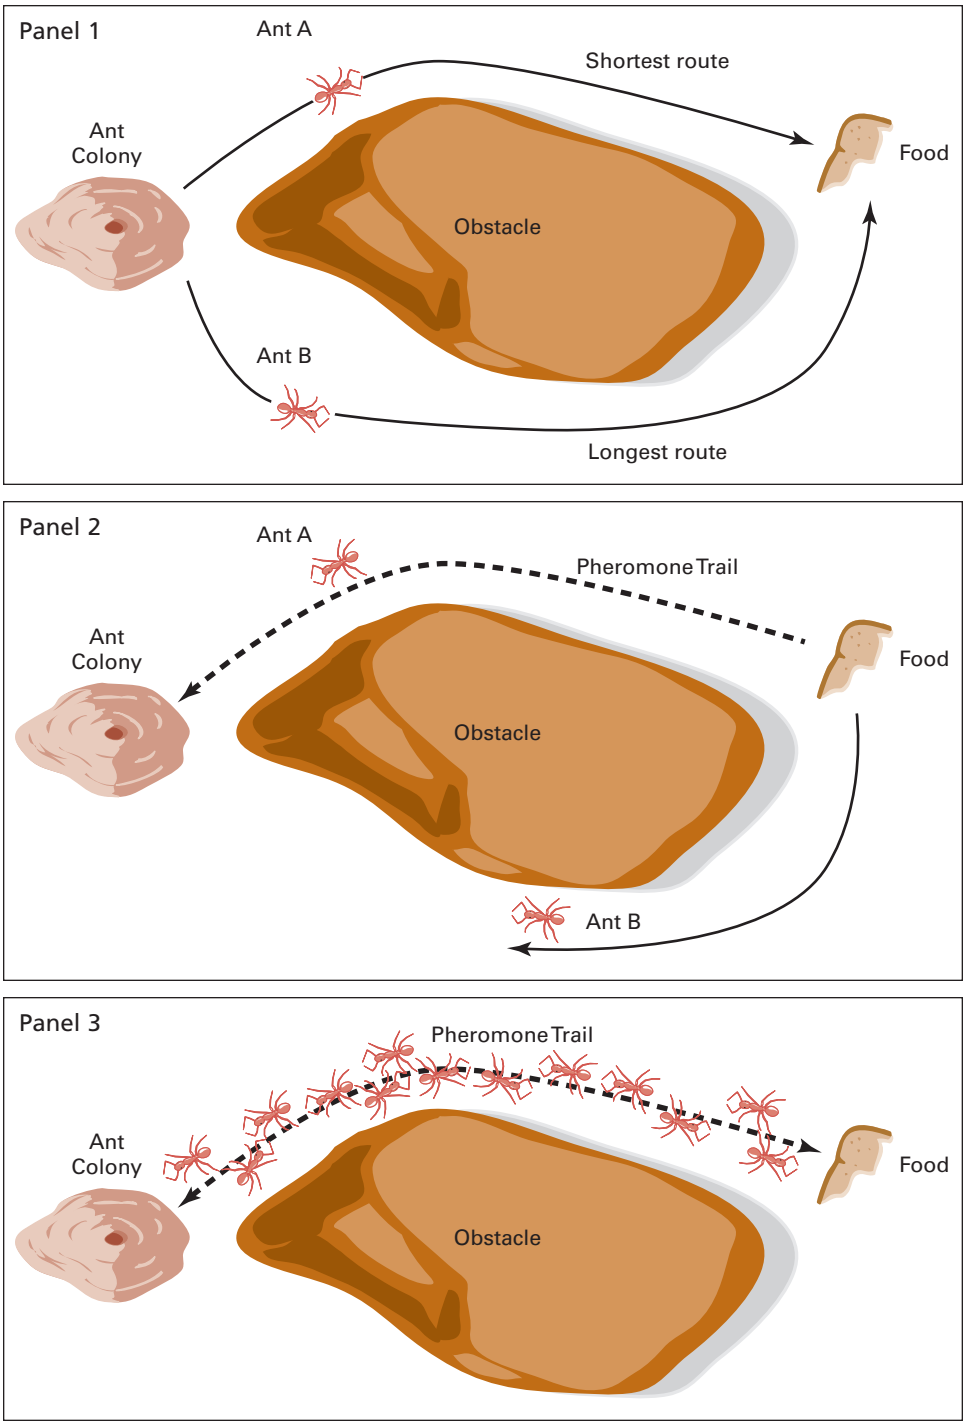

Figure 4.10
Ants Finding Food

Chicago to Atlanta and then reloaded onto a flight to Boston, requiring less unloading and reloading. Following the ant model, Southwest decreased its cargo transfer rates by 80 percent, reduced the workload of cargo employees by 20 percent, and also found that there was spare cargo space on flights that were previously full, enabling the company to accept more business. The overall gain to Southwest was in excess of $10 million per year.[45] For more examples of business lessons taken from nature, see the Industry Perspective entitled "Automated Truck-Painting Booths Schedule Themselves."

The future will see many more uses of intelligent agents. It's a pretty safe bet that these applications will include swarm intelligence and agent-based modeling. Already,

INDUSTRY PERSPECTIVE

AUTOMATED TRUCK-PAINTING BOOTHS SCHEDULE THEMSELVES

Honeybees have a very flexible work allocation system. For example, when food is scarce, bees normally assigned other tasks like nursing the young will switch roles to help to find food. The same principle can be applied to painting new trucks. In the factory, as a truck comes off the assembly line, one of the paint booths applies paint. Booths can change the color they're using, but it's a costly process to make the changeover, mainly because of the extraordinary expense of cleaning the lines, spray nozzles, etc. So, using the honeybee principle, each booth is programmed to paint the truck a standard color unless a pressing need arises. Rather than following some centralized schedule generated by the central scheduling system, each booth determines its own workload. The result is very effective. If a large batch of trucks all need to be painted red, the booths that usually paint in blue and black change to red to accommodate the trucks that are waiting. Also, when a booth breaks down, the others take over the load.

Another work allocation system comes from seed-harvester ants whose job it is to carry food back to their nest. They do this in relay-race fashion with each ant handing off the food to the next one. There the analogy ends, however, as there are no predetermined hand-off points; each ant carries the food until it meets the next ant, then hands over the food, and turns back to meet the new load from the ant before it and on up the line. Only the location of the food source and the nest are fixed. This process is known as the "bucket brigade," and is used in fast-food restaurants and book distribution centers.

Say, for example, when a book order came in, the first person put in all the geography books, the next one put in the IT books, and so on. This is called the zone approach and it pretty much guarantees that there will always be some workers waiting around for others to finish their tasks. The solution, taken from the ants, is to reorganize the workflow so that an employee fills his/her own order—stacks a certain number of books—until the person next in line takes over, and then the first employee goes back up the line to help someone else. Using this model, workers became 30 percent more productive than they had been when using the zone approach. Many distribution companies, including McGraw-Hill and Blockbuster Music, have used variations on the bucket brigade approach with great success.[46]

swarm intelligence is being implemented widely for scheduling, resource allocation, and routing. Other applications in the early stages include networks that have self-organizing components and robots that assemble themselves. There must be many, many more that have not yet been dreamt of.[47] Some people believe that intelligent agents will replace many of the other types of simulations in the future since swarm intelligence supports individuality, flexibility, and entities that can adapt quickly and effectively in a fast-changing business environment.

Summary: Student Learning Outcomes Revisited

1. **Define decision support system, list its components, and identify the type of applications it's suited to.** A ***decision support system (DSS)*** is a highly flexible and interactive IT system that is designed to support decision making when the problem is not structured. A DSS has three components: model management, data management, and user interface management. It's primarily an analysis tool to support your decision making, but you make the final decision.

2. **Define geographic information systems and state how they differ from other decision support tools.** A ***geographic information system (GIS)*** is a decision support system designed specifically to work with spatial information. It's used for the analysis of information in map form. Information is stored in layers which can be overlaid as appropriate. It's the layering and presentation that separates a GIS from other decision support tools.
3. **Define artificial intelligence and list the different types that are used in business.** ***Artificial intelligence*** is the science of making machines imitate human thinking and behavior. The types used in business include expert systems, neural networks, genetic algorithms, and intelligent agents.
4. **Define expert systems and describe the types of problems to which they are applicable.** An ***expert system*** (or ***knowledge-based system***) is an artificial intelligence system that applies reasoning capabilities to reach a conclusion. A rule-based expert system asks the user questions and, based on the answers, asks other questions until it has enough information to make a decision or a recommendation. Expert systems are good for diagnostic (what's wrong) and prescriptive problems (what to do). For example, you could use an expert system to diagnose illness or to figure out why a machine is malfunctioning. And you could use an expert system to determine what to do about the problem.
5. **Define neural networks and fuzzy logic and the uses of these AI tools.** A ***neural network*** (also called an ***artificial neural network*** or ***ANN***) is an artificial intelligence system that is capable of finding and differentiating patterns. Neural networks are good for finding commonalities in situations that have many variables. ***Fuzzy logic*** is a mathematical method of handling imprecise or subjective information. It is used to represent relative terms such as "hot" and "cold" so that a computer can use them in processing.
6. **Define genetic algorithms and list the concepts on which they are based and the types of problems they solve.** A ***genetic algorithm*** is an artificial intelligence system that mimics the evolutionary, survival-of-the-fittest process to generate increasingly better solutions to a problem. Genetic algorithms use the principles of ***selection, crossover,*** and ***mutation*** from evolution theory. These systems are best suited to problems where hundreds or thousands of solutions are possible and you need an optimum solution.
7. **Define intelligent agents, list the four types, and identify the types of problems they solve.** An ***intelligent agent*** is software that assists you, or acts on your behalf, in performing repetitive computer-related tasks. The four types are
   - ***Information agents***—the most common are ***buyer agents*** (or ***shopping bots***) that search the Web for products and services
   - ***Monitoring-and-surveillance agents*** (or ***predictive agents***) track conditions, perhaps on a network, and signal changes or troublesome conditions
   - ***Data-mining agents*** search data warehouses to discover information
   - ***User agents*** (or ***personal agents***) take action for you, particularly in repetitive tasks like sorting e-mail
8. **Define agent-based modeling and swarm intelligence.** ***Agent-based modeling*** is a way of simulating human organizations and behaviors using multiple intelligent agents, each of which follows a set of simple rules and can adapt to changing conditions. ***Swarm (collective) intelligence*** is the collective behavior of groups of simple agents that are capable of devising solutions to problems as they arise, eventually leading to coherent global patterns.

CLOSING CASE STUDY ONE

SHUTTLE DEBRIS, TREES, AND SNIPER ATTACKS

What do shuttle debris, trees, and possible locations of sniper attacks have in common? They can all be mapped and managed using a geographic information system (GIS). A GIS, as you learned in this chapter, allows you to see information spatially. Being able to visualize the physical location of objects and their proximity to each other greatly aids the decision-making process. Following are some examples that will give you an idea of how diverse the applications of GIS are.

THE ILL-FATED SPACESHIP COLUMBIA

On January 31, 2003, after 16 days of intensive research conducted while orbiting the earth, the seven-member crew of the space shuttle *Columbia* shut down all experiments, packed up their stuff, and ran tests of the systems they would need to land the next day. On February 1, the shuttle headed for home, but something went terribly wrong and the shuttle exploded over East Texas and arrived on earth in small pieces. To figure out what had happened, it was necessary to gather the pieces and try to reconstruct the sequence of events that led to the disaster. So, within hours of the terrible tragedy, while many of us were still too stunned to think, researchers and students at Stephen F. Austin State University in Nacogdoches, Texas, collected their GPS (global positioning system) gear and headed out looking for the debris. The field crews wanted to pinpoint exactly where all the debris fell—they didn't move or even touch anything they found. Within days there were 60 to 70 teams with up to 200 people collecting the information.

The GPS system they used is accurate to 100 feet, but with further processing the accuracy can be increased to within 3 feet. Using ArchInfo from ESRI, the teams were able to plot all the data points on a digital map that showed topological features, as well as roads and other man-made structures. In all, the volunteers collected 70 megabytes of information that they turned over to the federal agencies that had set up a central command post in Lufkin, Texas.

TREES IN CHATTANOOGA

In Chattanooga, people take the natural beauty and environmental benefits of their trees seriously. The city has created a GIS that maps the location of the 6,000 trees located in and around the business district. Along with this basic information, the GIS also has details on each tree that include its species, tree-pit dimensions, irrigation status, and trunk diameter. This information is very helpful when working out a maintenance plan to keep the trees healthy. For example, the size of a tree determines how many pruning hours it will need. Having an accurate map of where trees are allows the city to plan for future foliage. The city doesn't want to have more than 10 percent of any one species in such a small area since any insect infestation or disease that takes hold on one tree could spread more easily and potentially eliminate the whole tree population.

With the GIS system, which cost about $7,000 in labor to develop, the city can make more accurate estimates of the cost of materials and labor that are necessary to maintain the trees. This keeps the city from wasting resources or getting nasty surprises at the end of the fiscal reporting period.

RICHMOND MAPS BUS STOPS AND TRACKS BUSES

When GRTC Transit, the public transportation agency serving the City of Richmond and Chesterfield County, in Virginia, created a GIS of its bus routes and stops, it was not part of the plan to use it to protect its ridership from snipers. The idea was to improve its planning process designed to serve its rapidly expanding customer base.

GRTC's first step was to map its 2,500 bus stops. The agency wanted to know the details about each one, such as curb length, type and condition of the signs, benches, shelters, ramps, and trashcans at each stop. Within months of collecting this information, GRTC was using the GIS to analyze its bus routes and to consider requests from customers for changes and additions to the routes and also for amenities at the stops. With the ease of viewing lots of information in context that the GIS provides, the agency has implemented many improvements to its service.

In November 2001, the northern and central parts of Virginia were beset by sniper attacks, and no one knew where they were coming from, why they were occur-

ring, or how to stop them. GRTC Transit was able to use its GIS to quickly provide the police with information on bus stops where passengers' lives could have been at risk from sniper attacks because of on/off highway ramps close by.[48,49,50,51]

Questions

1. You saw examples in this case of a city that used a GIS to keep track of trees and one that mapped its bus stops to better serve its residents. Cities provide many other benefits, for example, firefighting and ambulance services. What sort of city features would a city incorporate into a GIS so that emergency dispatchers could more efficiently send out ambulances and fire trucks?
2. How would NASA use the information that students and other volunteers collected on the widely dispersed bits of the space shuttle *Columbia*? Why is it so important for NASA and for the Federal Aviation Authority (FAA) to map the location of, collect, and reassemble, as far as possible, this debris?
3. How might a business such as a bank use a GIS to determine where to put a new branch? What sort of information would the bank want to have about the area surrounding the proposed site?
4. Thematic GIS maps show a statistical value for certain locations and link those locations to an underlying geographic feature, such as the distribution of population within a county, or the areas in a nature reserve where birds and animals congregate. A point map shows the location of specific data items, like where your customers live or where fire hydrants are located. For each of the three examples in this case, specify whether each one uses thematic or point maps, or some combination.

CLOSING CASE STUDY TWO

USING NEURAL NETWORKS TO CATEGORIZE PEOPLE

Would your banker give you an A, B, or C? What about your supermarket? You know you're being graded in your classes, but did you know that you're also being graded by businesses?

Special treatment for certain customers is not new. Airline customers who fly first class have always received preferential treatment, even when flights were cancelled or delayed. You won't find them napping on a stone floor with their backpacks as pillows. This makes business sense to the airlines, since these are the customers who are most profitable.

Although companies have always offered preferential treatment to their more profitable customers, the speed and capacity of computers today are making the segmenting of customers possible to a degree unheard of just a few years ago. Part of the reason for this is neural networks. Using neural network software, businesses now have the ability to look for patterns in their customer information and classify customers according to how they affect the company's bottom line and thus to gauge whether it's worth the trouble of making them happy.

BANKS

The First Union Bank uses software that categorizes people into red, green, and yellow classes depending on the customer's history and value to the bank. Customers who are green might get better credit card rates than customers who are red and are judged to add less to the bank's bottom line.

Say you called the bank that issued you your credit card and said that you didn't want to pay the annual fee anymore. The bank could look at your credit card activity and decide whether it's more profitable to the bank to waive your fee rather than risk your not using the credit card anymore.

CREDIT CARD COMPANIES

Visa has saved millions of dollars using neural network software to spot fraud and to determine which of their customers might default or go bankrupt. Neural networks are good at spotting patterns, and if your profile looks like that of people who have defaulted, you'll be tossed into that category.

SUPPLIERS

Neural network classifying software can be applied to finding the best suppliers, too. Weyerhaueser Corporation's door factory executives have software to rank suppliers and distributors based on price, speed of delivery, and innovation. Using this information, Weyerhaueser doubled its sales and increased its return on net assets from 2 percent to 24 percent.

SUPERMARKETS

Catalina Supermarkets keeps track of which customers buy which products, how frequently, and what price they pay. Using neural network software, the supermarket chain can identify high-value customers and work at retaining them with offers of services such as free home delivery.

MOVIES

Even the movie business is getting in on the act. Twentieth Century Fox slices and dices its information in its databases to determine the most popular movies, actors, and plots in certain theaters, cities, and areas of the country. The aim is to show movies in those areas that will add the most to the bottom line. The result may be that people in certain areas will not get the chance to see certain movies.

There was a time when certain neighborhoods or geographic regions were redlined. That meant that banks and other businesses wouldn't deal with anyone who lived there. Some people think that this sort of market segmentation is a new form of redlining. Do you? Following are some questions for you to answer regarding this practice.[52,53]

Questions

1. A neural network learns to recognize patterns based on past information. Is this fair or reliable when applied to people? How accurate is it for a business to predict the future behavior of customers on the basis of historic information? Don't people change? Have you ever changed your behavior in the course of your life?
2. Customers are not likely ever to see the information that companies are using to pigeonhole them. Even the company executives may not know what criteria the neural network uses. How important are the assumptions underlying the software (i.e., the facts that the neural network is given about customers)? Even the IT specialists who design neural networks can't always vouch for their accuracy or specify exactly how the neural network reaches its conclusions. Is this safe for businesses? What are the possible business consequences of using neural networks without assurances of their reliability?
3. Businesses can use segmenting to suggest products and services to you, or if you request it, prevent your getting junk mail you don't want. Is that good? Would receiving wanted information or avoiding junk mail be worth the price of being categorized?
4. Say you run a business that supplies medical equipment (not prescription drugs)—wheelchairs, hospital beds, heating packs. You're trying to determine which customers you should give preferential treatment to. What assumptions or variables would you use (for example, age, income, and so on) to segment your customer population?
5. Do you think that this segmentation practice is fair? First, consider the business's stockholders, then consider the customers. Does it matter whether it's fair or not? Why or why not? Should there be laws against it, or laws controlling it, or none at all? Explain and justify your answer.
6. Does this differentiating practice make business sense? If you owned stock in a company, how would you feel about this practice? Do you think you should get better treatment if you're a better customer? Do you think people who are not such good customers should get the same deal that you get? Would it make any difference whether the company collected the information and did the neural network analysis itself, or bought the information or the whole package from a third party?
7. Is this practice of classifying the same as "redlining," or is it okay because it looks at behavior to classify people rather than assuming characteristics based on membership in a particular group?

Key Terms and Concepts

Agent-based modeling, 204
Artificial intelligence (AI), 189
Biomimicry, 204
Buyer agent (shopping bot), 200
Choice, 181
Crossover, 198
Data management, 186
Data-mining agent, 202
Decision support system (DSS), 183
Design, 181
Expert system (knowledge-based system), 190
Fuzzy logic, 197
Genetic algorithm, 198
Geographic information system (GIS), 187
Implementation, 181
Information agent, 200
Intelligence, 181
Intelligent agent, 200
Model management, 186
Monitoring-and-surveillance agent (predictive agent), 201
Multi-agent system, 204
Mutation, 198
Neural network (artificial neural network, ANN), 193
Nonrecurring (ad hoc) decision, 182
Nonstructured decision, 182
Recurring decision, 182
Robot, 189
Selection, 198
Structured decision, 182
Swarm (collective) intelligence, 205
User agent (personal agent), 203
User interface management, 186

Short-Answer Questions

1. What are the four types of decisions discussed in this chapter? Give an example of each.
2. What are the four steps in making a decision?
3. What is a DSS? Describe its components.
4. What is a geographic information system used for?
5. How is information represented in a geographic information system?
6. What is artificial intelligence? Name the artificial intelligence systems used widely in business.
7. What are the advantages of an expert system?
8. What sort of problems is an expert system used for?
9. How does a neural network work?
10. What three concepts of evolution are used by the genetic algorithm?
11. What are intelligent agents? What tasks can they perform?
12. What is a multi-agent system?
13. What do monitoring-and-surveillance agents do?

Assignments and Exercises

1. **MAKE A GIS** Make a GIS-type map using transparencies. Draw a map of your campus on one plastic transparency sheet. You can use software or felt-tip pens to do the actual drawing of the map. Next, use a second sheet as an overlay and mark on it what classes you have taken in what buildings. Take a third sheet and enter the type of classroom you had the course in (i.e., auditorium, lab, small, medium, large room). Make a fourth layer with special facilities, like a computer lab or a biology lab, and so on. What problems did you encounter while designing your GIS? What other information would you like to see in a real GIS of this type? Would this handmade GIS be helpful for new students? What layers would you keep for general use? What layers would you keep for sentimental value when your college days are over?
2. **CHOOSE A FINANCING OPTION** Using a spreadsheet (like Excel, for example) evaluate your options for a $12,000 car. Compare the payments (use the =pmt function in Excel), the total amount of interest, and the total you'll pay for the car under the following four options:

a. 3 years at 0 percent interest
b. 2 years at 1.99 percent annual percent rate (APR)
c. 4 years at 5 percent APR
d. 6 years at 6 percent APR

What other considerations would you take into account if you were going to buy a new car? Are there considerations other than the interest rate and the other parts that can be calculated? What are they? How is a car different from other purchases, such as CDs or TV sets or computers?

3. **WHICH SOFTWARE WOULD YOU USE?** Which type or types of computer-aided decision support software would you use for each of the situations in the table below? Note why you think each of your choices is appropriate. The decision support alternatives are
   - Decision support system
   - Geographic information system
   - Expert system
   - Neural network
   - Genetic algorithm
   - Intelligent agent

Problem	Type of Decision Support
You and another marketing executive on a different continent want to develop a new pricing structure for products	
You want to predict when customers are about to take their business elsewhere	
You want to fill out a short tax form	
You want to determine the fastest route for package delivery to 23 different addresses in a city	
You want to decide where to spend advertising dollars (TV, radio, newspaper, direct mail, e-mail)	
You want to keep track of competitors' prices for comparable goods and services	

4. **WHAT SHOULD THE MUSIC STORE OWNER DO?** A music store owner wants to have enough of the hottest CDs in stock so that people who come in to buy a particular CD won't be disappointed—and the store won't lose the profit. CDs that are not sold within a certain length of time go onto the sale table where they may have to be sold at cost, if they sell at all.

   The owner wants to design a decision support system to predict how many copies she should purchase and what information she will need. List some of the considerations that would go into such a system. Here are a couple to start you off: (1) the population of the target market; (2) sales for particular types of music in similar markets.

Discussion Questions

1. Some experts claim that if a business gets 52 percent of its decisions right, it will be successful. Would using a decision support system guarantee better results? Why or why not? What does the quality of any decision depend on? Do you think it matters what type of decisions are included in this 52 percent? For example, would getting the right type of paper clips be as influential a decision as deciding where to locate the business? Can you think of a situation where the type of paper clip matters a great deal?

2. Consider the topic of data warehouses in Chapter 3. In the future, AI systems will be increasingly applied to data warehouse processing. Which AI systems do you think might be helpful? For which tasks, or situations, might they best be applied? Do you think that AI systems will someday play a greater role in the design of databases and data warehouses? Why or why not?
3. Consider the differences and similarities among the four AI techniques discussed in this chapter. Name some problems that might be amenable to more than one type of AI system. Say you sell baseballs from your Web site. What types of AI systems could you use to generate information that would be useful to you in deciding what direction to take your company in the future? If you were pretty successful at selling baseballs, would you expect to have the amount of information on customers that, say, Wal-Mart has? Why or why not?
4. AI systems are relatively new approaches to solving business problems. What are the difficulties with new IT approaches in general? For each of the systems we discussed, identify some advantages and disadvantages of AI systems over traditional business processes. Say you were selling specialty teas and had both brick and click stores. Would you use the same type of AI systems for each part of your business? In what way would you use them or why would you not? Is there a place for decision support and artificial intelligence techniques in small specialty businesses? In what way would decision support add value? Can you think of how a DSS or an AI system would be value reducing (in terms of the value chain concept we discussed in Chapter 2)? What do you see as the major differences between running a mammoth concern and a small specialty business?
5. Neural networks recognize and categorize patterns. If someone were to have a neural network that could scan information on all aspects of your life, where would that neural network potentially be able to find information about you? Consider confidential (doctor's office) as well as publicly available (department of motor vehicles) information.
6. What type of AI systems could your school use to help with registration? Intelligent agents find vast amounts of information very quickly. Neural networks can classify patterns instantaneously. What sorts of information might your school administration be able to generate using these (or other AI systems) with all of its student data?
7. For which activities that are part of college life could you use agent-based modeling to simulate what happens? Describe three such scenarios.

Electronic COMMERCE

Finding Investment Opportunities on the Internet

When you buy stock in a company, you're betting on its success. Sometimes that bet is a good one, and sometimes it's not. Finding a company that's a good bet involves lots of research. To further complicate matters, some people prefer investing in large, safe companies, whereas others prefer the higher return of a small, more risky firm. So how do you make sense of all the options? Well, now you have access to financial information that professional investors use to evaluate stocks. The Internet brings together information-hungry investors with companies that have been anxiously looking to reach out to investors online. More than 900 companies now offer investment information on the World Wide Web, and the number is increasing rapidly. Remember, though, you must proceed with caution. Do your best to verify the source of any information.

You'll find many links on the Web site that supports this textbook (www.mhhe.com/haag, and select "Electronic Commerce Projects").

LEARNING ABOUT INVESTING

Investing can be as simple as finding a company that performs well financially and buying some of their stock. Or, if you want to spread your investment over a number of stocks and you don't want to select each stock personally, you can invest in a mutual fund. Of course, there are thousands of mutual funds with all types of investment objectives. So, any way you go you must pick your investment wisely. To help you get up to speed quickly, you'll find many helpful Web sites on the Internet.

For starters, you might explore the National Association of Securities Dealers (NASD) at www.nasdr.com. Check out their Investor Resources with its Education and Tools. You might also want to retrieve more general information from the online versions of traditional print media such as *The Wall Street Journal* or *Money* magazine.

Find three investment reference sites and explore what information is available. Then answer the following questions.

A. Is the site designed for first-time investors or those that are more experienced?

B. Can you search for a specific topic?

C. Are specific stocks or mutual funds reviewed or evaluated?

D. Does the site provide direct links to brokerage or stock quoting sites?

E. Is a forum for submitting questions available? If so, are frequently asked questions (FAQs) posted?

F. Who sponsors the site? Does it seem as if the sponsor is using the site to advertise its own products or services?

G. Can you download reference documents to read later?

RESEARCHING THE COMPANY BEHIND THE STOCK

One excellent way to pick a stock investment is to research the company behind that stock. Focusing on items such as sales revenues and profits to pick a stock is called *fundamental research.* So you might choose to invest in Hughes stock because you've discovered their sales revenues have been climbing steadily for the last three years. Or you might initially consider buying some Disney stock but change your mind when you find that EuroDisney revenues have been below expectations.

Now that you're ready to research a stock investment, connect to four different company sites. You can find your own or go to the Web site that supports this text where you will find a list of many other company sites. As you connect to the four sites, look up each company's financials and answer the questions that follow. You'll probably want to include at least two companies with which you are familiar and two that are new to you. In addition to reviewing company financials, look around each company site and see to what degree the site is investor oriented.

A. Do all the company sites offer financial information?
B. Is the information targeted at investors? How can you tell?
C. Can you download financial information to your computer and use it in a spreadsheet?
D. Can you download the company's annual report? Is it a full-color version?
E. Does the site provide direct links to e-mail addresses for requesting additional information?
F. Do the companies provide comparisons to others in their industry?
G. Does the site provide stock quotes as well as financials?
H. Can you search the site for financial-related information such as press releases?
I. Was there a charge for retrieving the financial information?

RETRIEVING STOCK QUOTES

Once you find the right stock to buy, you'll then be asking yourself, How much will this stock cost me? Stocks and mutual funds are both offered by the share and so you can easily buy as much or as little of the stock or mutual fund as you like. Still, some individual shares are priced in the hundreds or thousands of dollars, and that alone might make the purchase undesirable to you.

In addition to pricing individual shares to assess the affordability of an investment, you'll probably want to see how the price has varied over time. Even though most financial advisors will tell you that historical price variations provide no indication of future performance, most everyone uses price history to get a feel for whether the investment is trading at all-time highs or lows. So finding a chart of a stock price online might be helpful when deciding to make your purchase.

And even after you've made your purchase, you'll probably want to follow how your investment is doing. The thrill of realizing a "paper profit" is enough to keep many investors checking their investments daily. Of course, realizing a "paper loss" can be equally disappointing. And even if daily tracking isn't for you, you'll certainly want to check on your investments regularly, and doing so online can be quick and painless.

Pick three stock quoting services, examine what it takes to retrieve a stock or mutual fund quote, and answer the following questions.

A. Are the quotes provided free of charge or for a fee?
B. Does the site require a ticker symbol (the abbreviation used by experienced investors) or can you type in a company name?
C. Are the quotes in real time or are they delayed (15 to 20 minutes old)?
D. Does the site require registration?
E. Are historical prices available?
F. Are price charts available? Can you customize the chart display?
G. Can you create and save a personal portfolio of stocks?