# CHAPTER FOUR OUTLINE

## STUDENT LEARNING OUTCOMES

1. Compare and contrast decision support systems and geographic information systems.
2. Define expert systems and describe the types of problems to which they are applicable.
3. Define neural networks and fuzzy logic and the use of these AI tools.
4. Define genetic algorithms and list the concepts on which they are based and the types of problems they solve.
5. Describe the four types of agent-based technologies.

## PERSPECTIVES



## WEB SUPPORT

www.mhhe.com/haag

- Best in computer statistics and resources
- Consumer information
- Metadata
- Bureau of Labor and Statistics
- Demographics
- Exploring Google Earth
- Gold, silver, interest rates, and money
- Learning about investing
- Stock quotes

## SUPPORTING MODULES

**XLM/D Decision Analysis with Spreadsheet Software**
Extended Learning Module D provides hands-on instructions concerning how to use many of the powerful decision support features of Excel including Basic Filter, Custom Filter, conditional formatting, and pivot tables (in both two and three dimensions). Each of these takes only minutes to learn and requires just a few clicks.

**XLM/M Programming in Excel with VBA**

Extended Learning Module M covers the basics of learning how to write macros (short programs) in Excel using VBA, Visual Basic for Applications. It covers how to use the Visual Basic Editor (VBE), how to use the macro recorder, and how to write procedures, functions, if-then structures, and loops.

# CHAPTER FOUR

# Decision Support and Artificial Intelligence

## Brainpower for Your Business

### OPENING CASE STUDY: VISUALIZING INFORMATION IN MAP FORM CAN AID IN DECISION MAKING

What do shuttle debris, trees, and bus stops have in common? They can all be mapped and managed using a geographic information system (GIS). A GIS allows you to see information spatially. If you've ever explored our world using Google Earth, then you're familiar with a GIS. Being able to visualize the physical location of objects and their proximity to each other greatly aids the decision-making process.

On February 1, 2003, after 16 days of intensive research conducted while orbiting the earth, the space shuttle Columbia headed 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. The exact location of the pieces of debris was key, as researchers and scientists worked backward from that information to determine how and why the shuttle had exploded. Teams of people used GPS (global positioning system) gear to mark the exact location of each piece of debris as they came across it.

The city of Chattanooga, Tennessee, has created a GIS that maps the location of the 6,000 trees in and around its business district, along with 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 by species 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.

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, 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.[1,2,3,4]

### Questions

1. Do you use Web-based map services to get directions and find the location of buildings? If so, why?
2. In what ways could real estate agents take advantage of the features of a GIS?
3. How could GIS software benefit a bank wanting to determine the optimal placements for ATMs?

# Introduction

Decision making is crucial to business. Using all forms of intellectual assets—basic data to information to business intelligence to knowledge—decision makers in the business world ponder "make-or-break" decisions. Some are responding to changes in customer demographics (customer relationship management); others are considering the trade-offs of using low-cost suppliers as compared to higher-cost suppliers that deliver fewer defective products in faster time (supply chain management); still others are attempting to determine the best overall strategy for the organization. Does the organization position itself for overall cost leadership, differentiation, or focus? How does the organization combat an abundance of e-commerce businesses on the Web offering the same products and services (i.e., buyer power)? In this chapter, we'll investigate the tools that IT can provide to help you 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."[5]

In business, decision making has four distinct phases, as proposed by Herbert Simon (see Figure 4.1).[6] 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.

Figure 4.1

Four Phases of Decision Making

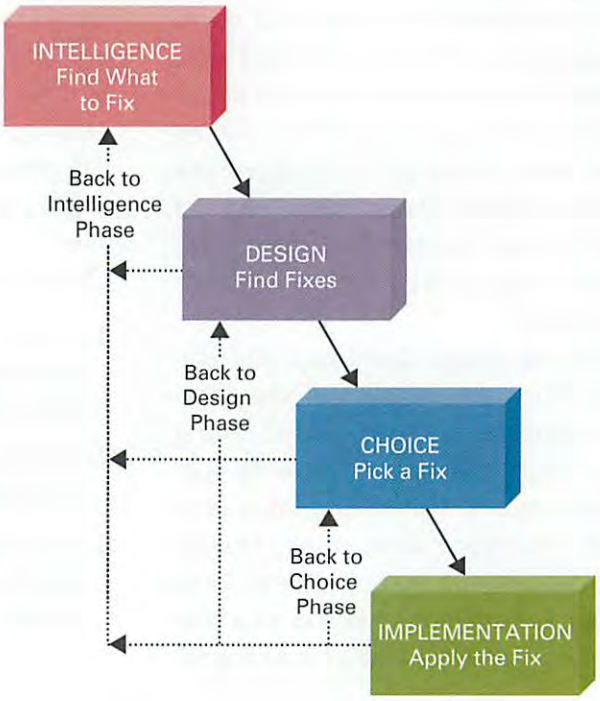

# CHAPTER FOUR

## Decision Support and Artificial Intelligence

### Brainpower for Your Business

### OPENING CASE STUDY: VISUALIZING INFORMATION IN MAP FORM CAN AID IN DECISION MAKING

What do shuttle debris, trees, and bus stops have in common? They can all be mapped and managed using a geographic information system (GIS). A GIS allows you to see information spatially. If you've ever explored our world using Google Earth, then you're familiar with a GIS. Being able to visualize the physical location of objects and their proximity to each other greatly aids the decision-making process.

On February 1, 2003, after 16 days of intensive research conducted while orbiting the earth, the space shuttle Columbia headed 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. The exact location of the pieces of debris was key, as researchers and scientists worked backward from that information to determine how and why the shuttle had exploded. Teams of people used GPS (global positioning system) gear to mark the exact location of each piece of debris as they came across it.

The city of Chattanooga, Tennessee, has created a GIS that maps the location of the 6,000 trees in and around its business district, along with 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 by species 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.

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, 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.[1,2,3,4]

#### Questions

1. Do you use Web-based map services to get directions and find the location of buildings? If so, why?
2. In what ways could real estate agents take advantage of the features of a GIS?
3. How could GIS software benefit a bank wanting to determine the optimal placements for ATMs?

# Introduction

Decision making is crucial to business. Using all forms of intellectual assets—basic data to information to business intelligence to knowledge—decision makers in the business world ponder "make-or-break" decisions. Some are responding to changes in customer demographics (customer relationship management); others are considering the trade-offs of using low-cost suppliers as compared to higher-cost suppliers that deliver fewer defective products in faster time (supply chain management); still others are attempting to determine the best overall strategy for the organization. Does the organization position itself for overall cost leadership, differentiation, or focus? How does the organization combat an abundance of e-commerce businesses on the Web offering the same products and services (i.e., buyer power)? In this chapter, we'll investigate the tools that IT can provide to help you 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."[5]

In business, decision making has four distinct phases, as proposed by Herbert Simon (see Figure 4.1).[6] 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.

Figure 4.1

Four Phases of Decision Making

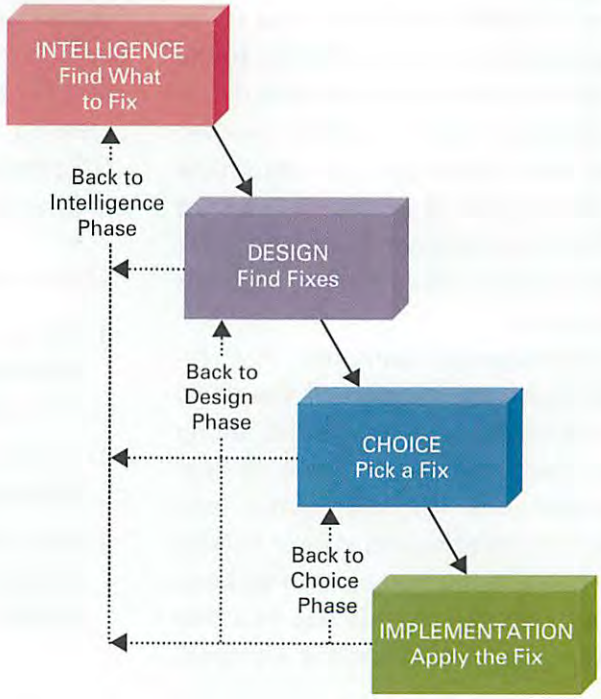

3. ***Choice*** (pick a fix): Examine and weigh the merits of each solution, estimate the consequences of each, and choose the best one (which may be to do nothing at all). The "best" solution may depend on such factors as cost, ease of implementation, staffing requirements, and timing. This is the prescriptive 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.

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.

A second model of decision making, also proposed by Simon, is *satisficing*, which differs from the four-phase process. ***Satisficing*** is making a choice that meets your needs and is satisfactory without necessarily being the best possible choice available. The term is a combination of the words "satisfied" and "sufficient."

Organizations in the private and public sectors are "satisficing" all the time in setting goals such as "fair price" or "reasonable profit." There's a fundamental difference between setting a goal of "high growth" and one of "maximum growth." "Maximum growth" is an optimizing strategy while "high growth" is a satisficing strategy. Usually a term like "high growth" is precisely defined. It may be 3 percent or 30 percent, but the idea is that when you reach that level, you can declare success.

In both business and your personal life, you'll face decisions that are some combination of four main types of decisions (see Figure 4.2): The first type is a ***structured decision***, which 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 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.


Figure 4.2

Categorizing Decisions by Type

On the other hand, 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 unstructured, for example, choosing a job. Structured elements of choosing a job include consideration of such things as salary and signing bonus. Unstructured elements of such a decision include things like the potential for advancement.

Another type of decision regards frequency with which the decision is made. A ***recurring decision*** is one that happens repeatedly, and often periodically, whether weekly, monthly, quarterly, or yearly. Deciding how much inventory to carry and deciding at what price to sell the inventory are recurring decisions. 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. Deciding where to build a distribution center or company mergers are examples of nonrecurring or ad hoc decisions (although, the general trend in business today is for companies to consider mergers on a more consistent basis).

## Decision Support Systems

LEARNING OUTCOME 1

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. 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 4.3). IT brings speed, vast amounts of information, and 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

Figure 4.3

The Alliance between You and a Decision Support System

| 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 | |

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.[7]
- Burlington Northern and Santa Fe (BNSF) railroad regularly tests the rails its trains ride on to prevent accidents. Wornout 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.[8]
- 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.[9]

## 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 4.4 on the next page): 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. 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.
- *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

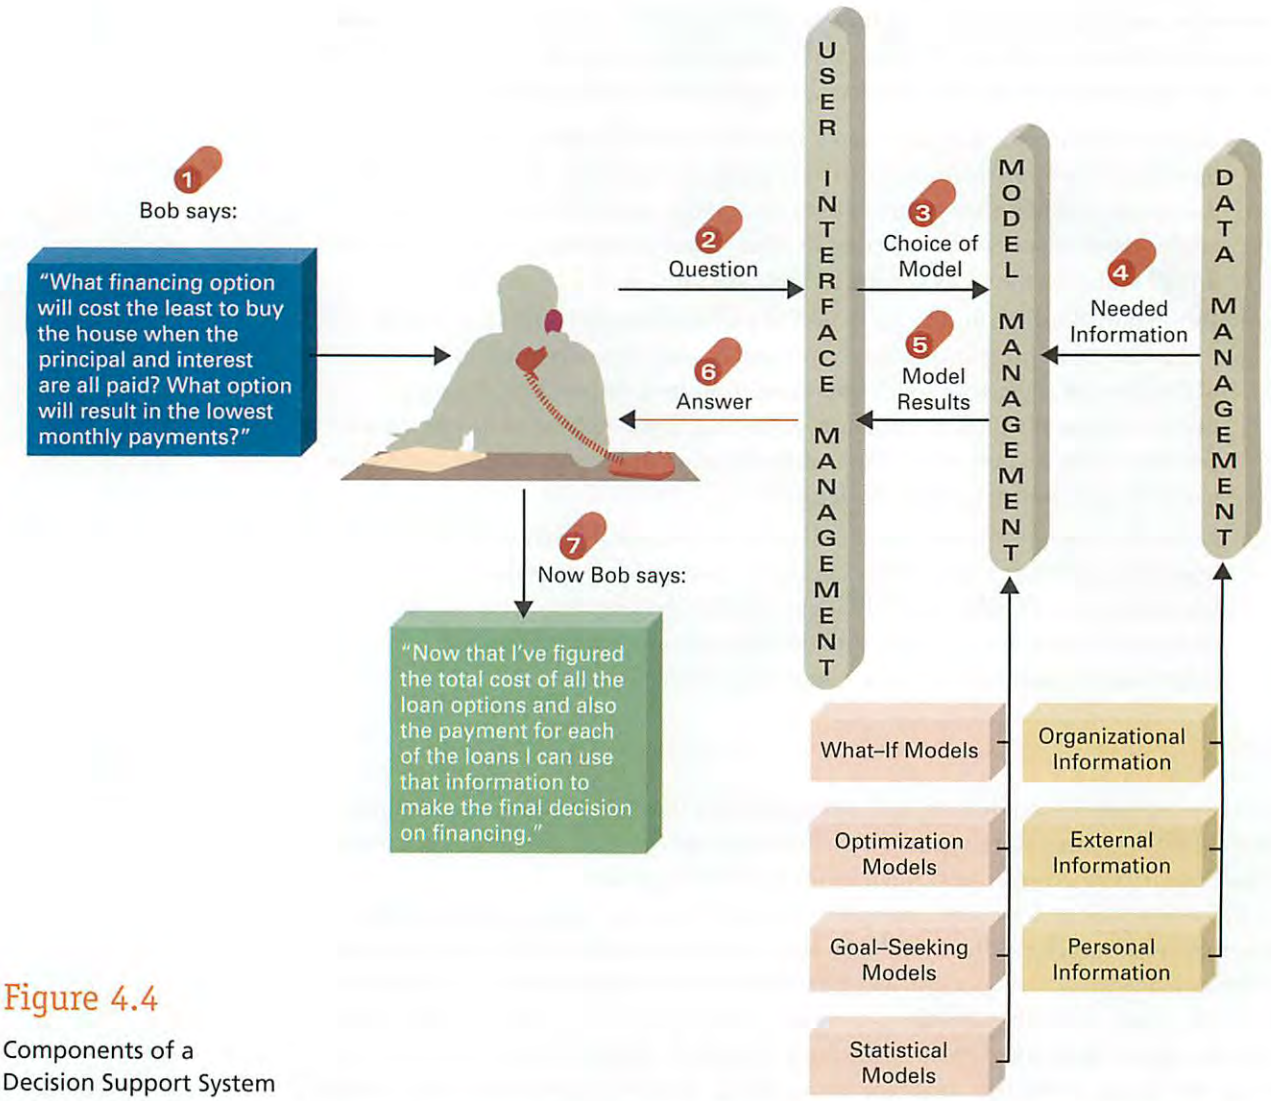


Figure 4.4

Components of a Decision Support System

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.

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.

## 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.[10]

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.

## 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.

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 4.5 is a map of the Denver area shown using the real estate search function in Zillow (www.zillow.com). On the left-hand side you can see the layers (or types of information) that can be shown on the map—homes for sale, homes recently sold, and so on. The map itself is color-coded to show the estimated dollars per square foot for residential areas. Color shadings of red through maroon are in the $500 to $700 per square foot range and variations of blue are in the $0 to $200 per square foot range.

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 the *Miami Herald* reporters did for Hurricane Andrew. You can represent either statistical information, for instance, 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

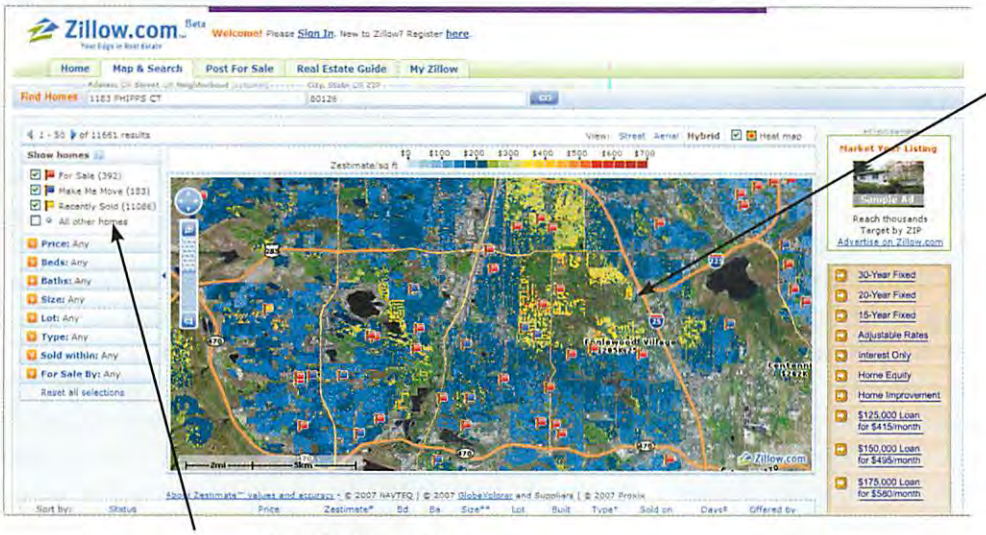


Figure 4.5

Square Footage Costs in Denver according to Zillow

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. Here are other examples of GISs in use:

- The city of Chicago has 40 departments and a fleet of 7,300 vehicles, ranging from snowplows to sanitation trucks to construction vehicles. The city implemented a system that improved the accountability of the departments and the efficiency of the use of its diverse vehicles. Combining GIS technology with GPS and a wireless network, the city has a system that allows supervisors to use the city's intranet to view the location of any vehicle. The system keeps information on where each truck, van, or earth mover goes, its speed, and the path it took to get to its current location. This allows the city to reassign vehicles for greater effectiveness and to tell where each one is or was at any particular time.[11]
- During the bovine spongiform encephalopathy (BSE, or more commonly, mad cow disease) outbreak in 2003 it became clear that a national tracking system of cattle, from birth on, was needed to help keep the beef supply—and the people who consume it—safe. During 2003, only 87 cattle and 15 calves were at risk of being susceptible to BSE, but over 400 were slaughtered because officials did not have complete enough records on the movement of these animals. To correct that problem, the Kansas Animal Health Department (KAHD) has developed a tracking system to log every producer, sales barn, feed yard, and processing facility through which each animal passes. This is a gargantuan task since every year millions of cattle come through the state on their way from one part of the country to the other. KAHD's tracking system combines RFID (radio frequency identification, see Chapter 9) tags on the ears of cattle, GPS, GIS, and a wireless network to track each animal's movements and transmit the information to a national database.[12]

DSSs and GISs are IT systems that augment business brainpower. IT can further expand business brainpower by means of ***artificial intelligence (AI)***—the science of making machines imitate human thinking and behavior. Financial analysts use a variety of artificial intelligence systems to manage assets, invest in the stock market, and perform other financial operations. Hospitals use artificial intelligence in many capacities, from

## INDUSTRY PERSPECTIVE

### LOCATION MASHUP = GIS + THE INFORMATION YOU WANT

Broadly defined, a ***mashup*** is a combination of content from more than one source. Due in large part to the technologies of Web 2.0 (see *Extended Learning Module B*), people produce and post mashups of all kinds—music mashups (a single song using parts of numerous other songs), video mashups (just like a music mashup but using video), and even location mashups. A ***location mashup*** is a geographic information system (GIS) that displays a particular geographic area and then overlays content according to the user's desires. For example, at www.gawker.com/stalker, you can view a mashup showing the latest travels and locations of celebrities like Madonna.

Location mashups are actually quite easy to create using many new and emerging tools. Pipes (by Yahoo!), GeoCommons (by FortiusOne), and GeoIQ (also by FortiusOne) are among the many tools you can use to create your own location mashup. Many of them come with location data that you can instantly add to a GIS map to create a location mashup. GeoIQ, for example, includes a wealth of data such as the location of spammers by street address and incidents of the West Nile virus.

Location mashups are not limited to the personal space. Many businesses are using them too. British Petroleum, UPS, Best Buy, Century 21, DaimlerChrysler, FedEx's Kinko's, Ford, General Motors, H&R Block, Starbucks, and Target are among the many businesses taking advantage of the capabilities of location mashups. Some are using them for internal efficiency and decision-making efforts (e.g., BP uses them to provide updated weather patterns to employees in the Gulf of Mexico), while others are using them in the form of a customer-facing technology (e.g., Starbucks' location mashup helps customers find their next caffeine fix).[13]

scheduling staff, to assigning beds to patients, to diagnosing and treating illness. Many government agencies, including the IRS and the armed forces, use artificial intelligence. Credit card companies use artificial intelligence to detect credit card fraud, and insurance companies use artificial intelligence to spot fraudulent claims. 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. The AI systems that businesses use most can be classified into the following major categories:

1. Expert systems
2. Neural networks (and fuzzy logic)
3. Genetic algorithms
4. Intelligent agents (or agent-based technologies)

## Expert Systems

LEARNING OUTCOME 2

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).
- *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.

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 4.6 on the next page).

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

| 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. |

Is the light green (Yes/No)? No.

Is the light red (Yes/No)? No.

Is the light likely to change to red before you get through the intersection (Yes/No)? Why?

Will only reach this point if light is yellow, and then you'll have two choices.

Is the light likely to change to red before you get through the intersection (Yes/No)? No.

**Conclusion: Go through the intersection.**

Figure 4.6

Traffic Light Expert System Rules

- 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

LEARNING OUTCOME 3

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 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.[14]

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.[15] 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.[16,17,18] 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.[19] 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.[20]

- 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.[21]
- 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.[22]
- 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.[23]
- 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.[24]

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.

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

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;

## INDUSTRY PERSPECTIVE

### HOW PROFITABLE WOULD A NEURAL NETWORK SAY YOU ARE?

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.

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.[25,26]

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

LEARNING OUTCOME 4

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.[27]
- They're used by investment companies to help in trading choices and decisions.[28]
- 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.[29]
- 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.[30]

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!"[31]

## Intelligent Agents

LEARNING OUTCOME 5

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.

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 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.[32]

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

## 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.[33]

down.[34] 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?"[35]

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.

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.[36] 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. We discussed ERPs in Chapter 2 and will explore them further 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.***

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 4.7 for examples of companies that have used agent-based modeling to their advantage.

Figure 4.7

Companies That Use Agent-Based Modeling[37]

- 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.

## 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, 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 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.[38] 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.[39]

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.

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 they follow a simple procedure.

Say two ants (A and B) leave the same point to search for food. Ant A finds food first because ant B has to traverse around several rocks before finding food (i.e., Ant A found a shorter route to the food). Having found a food source, Ant A returns to the nest by the same route, leaving behind a trail of pheromones (a biological breadcrumb trail) so that it will know what path to take next time and so will the other ants. 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 along the path Ant A found.

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. The pheromone trail on the unused path (that left by ant B) 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 in 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 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.[40]

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, 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. 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. **Compare and contrast decision support systems and geographic information systems.** 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 ***geographic information system (GIS)*** is a decision support system designed specifically to analyze spatial information. So, they both are designed to support decision-making efforts. While traditional DSSs mainly use text and numeric data, GISs represent many types of information in spatial or map form.
2. **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.
3. **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.
4. **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.
5. **Describe the four types of agent-based technologies.** 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
   - Two emerging variants within agent-based technologies are (1) ***multi-agent systems,*** groups of intelligent agents that have the ability to work independently and to interact with each other, and (2) ***agent-based modeling,*** 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. Both are based on ***biomimicry,*** learning from ecosystems and adapting their characteristics to human and organizational situations.

## CLOSING CASE STUDY ONE

# CRYSTAL BALL, CLAIRVOYANT, FORTUNE TELLING. . . CAN PREDICTIVE ANALYTICS DELIVER THE FUTURE?

In the Tom Cruise movie *Minority Report,* police were able to accurately predict a crime, its location, and the criminal in advance of the event in time to send police to prevent the crime from occurring. Science fiction at its best, huh? Actually, that's somewhat of a reality now through predictive analytics.

***Predictive analytics*** uses a variety of decision tools and techniques—such as neural networks, data mining, decision trees, and Bayesian networks—to analyze current and historical data and make predictions about the likelihood of the occurrence of future events. Along the lines of *Minority Report,* police in Richmond, Virginia, are using predictive analytics to determine the likelihood (probability) that a particular type of crime will occur in a specific neighborhood at a specific time.

Using the system, the mobile task force of 30 officers is deployed to the areas with the greatest likelihood of crimes occurring. According to Richmond Police Chief, Rodney Moore, "Based on the predictive models, we deploy them [the mobile task force] almost every three or four hours." Sixteen fugitives have been arrested directly as a result of the system's prediction of the next time and location of a crime. Moreover, in the first week of May in 2006, no homicides occurred, compared to three in the same week of the previous year.

The predictive analytics system uses large databases that contain information on past calls to police, arrests, crime logs, current weather data, and local festivals and sporting and other events. From an IT point of view, the system is a combination of software—SPSS's Clementine predictive analysis software and reporting and visualization tools from Information Builder—and decision support and predictive models developed by RTI International.

The Richmond police afford just one of many examples of the use predictive analytics. Some others include the following:

- *Blue Cross Blue Shield of Tennessee*—uses a neural network predictive model to predict which health care resources will be needed by which postoperative patients months and even years into the future. According to Soyal Momin, manager of research and development at Blue Cross Blue Shield, "If we're seeing a pattern that predicts heart failure, kidney failure, or diabetes, we want to know that as soon as possible."
- *FedEx*—uses a predictive analytics system that is delivering real and true results 65 to 90 percent of the time. The system predicts how customers will respond to new services and price changes. It also predicts which customers will no longer use FedEx as a result of a price increase and how much additional revenue the company will generate from proposed drop-box locations.
- *University of Utah*—uses a predictive analytics system to generate alumni donations. The system determines which of its 300,000 alumni are most likely to respond to an annual donation appeal. This is particularly appealing to most higher-education institutions as they have limited resources to devote to the all-important task of fund raising. Donations increased 73 percent in 2005 for the University of Utah's David Eccles School of Business as a result of the system.

The future of predictive analytics is very bright. Sales of predictive analytics software are expected to reach $3 billion by 2008. Moreover, businesses are beginning to build predictive analytics into mainstream, operational applications—such as CRM, SCM, and inventory management—which will further increase their use. According to Scott Burk, senior statistician and technical lead for marketing analytics at Overstock.com, "Predictive analytics is going to become more operational. We're definitely doing things a lot smarter than we were six months ago." Overstock.com uses its predictive analytics system to predict demand levels for products at various price points.[41]

### Questions

1. Many predictive analytic models are based on neural network technologies. What is the role of neural networks in predictive analytics? How

can neural networks help predict the likelihood of future events. In answering these questions, specifically reference Blue Cross Blue Shield of Tennessee.

2. What if the Richmond police began to add demographic data to its predictive analytics system to further attempt to determine the type of person (by demographic) who would in all likelihood commit a crime. Is predicting the type of person who would commit a crime by demographic data (ethnicity, gender, income level, and so on) good or bad?
3. In the movie *Gattaca,* predictive analytics were used to determine the most successful career for a person. Based on DNA information, the system determined whether or not an individual was able to advance through an educational track to become something like an engineer or if the person should complete only a lower level of education and become a janitor. The government then acted on the system's recommendations and placed people in various career tracks. Is this a good or bad use of technology? How is this different from the variety of personal tests you can take that inform you of your aptitude for different careers?
4. What role can geographic information systems (GISs) play in the use of predictive analytics? As you answer this question, specifically reference FedEx's use of predictive analytics to (1) determine which customers will not respond positively to a price increase and (2) project additional revenues from proposed drop-box locations.
5. The Department of Defense (DoD) and the Pacific Northwest National Laboratory are combining predictive analytics with visualization technologies to predict the probability that a terrorist attack will occur. For example, suspected terrorists caught on security cameras who loiter too long in a given place might signal their intent to carry out a terrorist attack. How can this type of predictive analytics be used in an airport? At what other buildings and structures might this be used?

## CLOSING CASE STUDY TWO

# CLOSING THE GREAT HEALTH CARE DIVIDE WITH PATTERN RECOGNITION AND DATA-MINING TECHNOLOGIES

In today's world, people, cultures, and nations are distinguished by which side of the "*great divides*" they are on. There are the financial and economic divides, which differentiate between the haves and have-nots with respect to wealth, income, and prosperity. There are the educational divides that distinguish among people on the basis of their access to education and according to whether a college degree is achieved. There is a "great digital divide," on either side of which are those who have or do not have access to technology.

Worldwide—and certainly in the United States—is a health care divide between those who can and cannot afford good health care. With health care costs spiraling out of control, the health care divide is widening and the population with access to affordable health care is shrinking. Therefore, many health care providers and technology providers such as IBM, the Mayo Clinic, and the Cleveland Clinic are collaborating on IT-enabled strategies to reduce the cost of health care while providing better health care than ever before.

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 1)? 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.

# CHAPTER PROJECTS

## Group Projects

- Assessing the Value of Customer Relationship Management: Trevor Toy Auto Mechanics (p. 464)
- Analyzing the Value of Information: Affordable Homes Real Estate (p. 465)
- Executive Information System Reporting Political Campaign Finance (p. 466)
- Building a Decision Support System: Creating an Investment Portfolio (p. 470)
- Creating a Decision Support System: Buy versus Lease (p. 476)
- Building a Decision Support System: Break-Even Analysis (p. 482)
- Building a Scheduling Decision Support System: Airline Crew Scheduling (p. 484)

## e-Commerce Projects

- Best in Computer Statistics and Resources (p. 488)
- Consumer Information (p. 488)
- Metadata (p. 489)
- Bureau of Labor and Statistics (p. 490)
- Demographics (p. 490)
- Exploring Google Earth (p. 492)
- Gold, Silver, Interest Rates, and Money (p. 494)
- Learning about Investing (p. 496)
- Stock Quotes (p. 497)