

1

Introduction

1.1 Intelligent Agents

An *intelligent agent* is a computational or natural system that senses its environment and takes actions intelligently according to its goals. We focus on computational (versus natural) agents that act in the interests of their human principals. Such intelligent agents aid humans in making decisions. Intelligent agents can play several possible roles in the human decision process. They may play the roles of a consultant, an assistant, or a delegate. For simplicity, we will refer to intelligent agents as just *agents*.

When an agent acts as a consultant (Figure 1.1), it senses the environment but does not take actions directly. Instead, it tells the human principal what it thinks should be done. The final decision rests on the human principal. Many expert systems, such as medical expert systems (Teach and Shortliffe [75]), are used in this way. In one possible scenario, human doctors independently examine patients and arrive at their own opinions about the diseases in question. However, before the physicians finalize their diagnoses and treatments, the recommendations from expert systems are considered, possibly causing the doctors to revise their original opinions. Intelligent agents are used as consultants when the decision process can be conducted properly by humans with satisfactory results, the consequences of a bad decision are serious, and agent performance is comparable to that of humans but the agents have not been accorded high degrees of trust.

When an agent acts as an assistant (Figure 1.2), the raw data and observations are *directly* processed only by the agent. It either preprocesses the information and presents it to the human principal for further decision making or conducts the entire decision process and offers the recommendations to the human principal for approval and execution. Software systems commonly referred to as *decision support systems* (Druzdzel and Flynn [16]) are used as human assistants. For example, a corporate executive manager may use such a system to model past business data,




Figure 1.1: An agent as a consultant. The globe denotes the environment, and the computer denotes the agent.

analyze the consequences of alternative business actions, and arrive at a business decision. Due to the vast amount of business data and the complex interdependence among different aspects of business practice (such as material supply, production, personnel, marketing, sales, and investment), human cognitive capabilities limit




Figure 1.2: An agent as an assistant.

1.1 Intelligent Agents




Figure 1.3: An agent as a delegate.

the manager from effectively processing all available information and arriving at a rational decision. With the aid of decision support systems, the manager only needs to evaluate preprocessed information and final or semifinal decision recommendations, which are much smaller in volume and higher in abstraction level. Intelligent agents are used as assistants when the entire decision process is beyond human cognitive capabilities given the time and decision quality constraints, for intelligent agents are capable of producing quality decision recommendations in a short time.

When an agent acts as a delegate (Figure 1.3), it not only processes the observations and generates the decision but also directly executes the chosen action without requiring approval from the human principal. Autonomous robots, Internet infobots (who search Web sites and gather relevant information for human principals), intelligent tutors, and embedded intelligent systems (such as an autodriver in a smart car) are some examples of agents being used as delegates. For example, an intelligent tutor probes the mind of a student by posing queries, determines the student’s level of knowledge on the subject, detects misconceptions and weakness in skills, devises lessons and exercises to educate the student, corrects misconceptions, and strengthens the student’s skills. Throughout the decision process, the tutor receives no intervention from the student’s human teachers or parents (the principals of the agent). Intelligent agents are used as delegates for humans when the costs of humans’ performing the tasks are too high or the tasks are too dangerous to be performed by humans and the agents are capable of high-quality performance and have won human trust.

As the information available to human decision makers continues to grow beyond human cognitive capabilities, as the cost of computing power continues to drop, and as more intelligent agents with equihuman or superhuman performance are developed, we expect to see more agents deployed as human assistants and delegates that will play increasingly important roles in daily life.

1.2 Reasoning about the Environment

Whether an agent acts as a consultant, an assistant, or a delegate, its task contains a set of essential activities. We consider agents whose activities can be described in general as follows: The agent is associated with an environment or a problem domain of interest and carries a model or a representation or some prior knowledge about the domain. Its goal is to change (or maintain) the state of the domain in some desirable way (according to the interest of its human principal). To do so, it takes actions (or recommends actions to its human principal) from time to time. To take the proper action, it makes observations on the domain, guesses the state of the domain based on the observations and its prior knowledge about the domain, and determines the most appropriate action based on its belief and goal. We refer to the activity of guessing the state of the domain from prior knowledge and observations as *reasoning* or *inference*.

An agent’s activities may not fall under the preceding description. It may not separate reasoning about the state of the domain from choosing the actions but instead encodes its action directly as a function of observations. For example, an if-then rule in a rule-based expert system may have its *if* part specifying some observations and its *then* part specifying a desirable action when the observations are obtained. An if-then rule is made out of symbols (observations and actions are specified in terms of symbols). Hence, an agent constructed from a rule-based system uses a symbolic (Poole, Mackworth, and Goebel [55]) knowledge representation of the domain. An agent may not even use a symbolic representation. For example, an agent’s behavior may be based on an *artificial neural network* (Hassoun [23]) in which the observations are mapped to proper actions through network links and hidden units that do not have well-defined semantics.

In this book, we consider agents using symbolic knowledge representations and reasoning explicitly about the state of the domain. Reasoning about the state of the domain can be a challenging task by itself, as we shall see. Separating the reasoning task from the selection of actions allows decomposing the decision process into two stages and working on each through divide-and-conquer techniques. To develop agents with quality performance, one approach is to ensure that agents’ behavior can be analyzed formally and rigorously. Separating inference about the state from choice of action also facilitates the analysis and explanation of agents’ behavior by analyzing and understanding each stage of the decision process individually.

1.3 Why Uncertain Reasoning?

The main focus of this book is on guessing (reasoning about) the state of the domain. We have used the word *guessing* for several reasons.

1. Some aspects of the domain are often unobservable and therefore must be estimated indirectly through observables. Consider a neurologist who needs to know which part of an epileptic brain is abnormal. Direct examination of the brain is not an option except as part of the surgical procedure after the diagnosis has been completed and surgery has been considered necessary. Instead, to estimate the state of the brain, the seizure behavior of the patient is observed, and the electroencephalogram (EEG) is recorded.

Complex equipment, such as automobiles, airplanes, automatic production lines, chemical plants, computer systems, and computer networks, is used widely in modern society. Each piece of equipment is made of many components, each of which is further composed of many devices and parts. Whether a device or component is functionally normal or faulty or on the verge of breaking down is usually not observable. For example, whether the bearing of a helicopter propeller is about to break down is not directly observable without disassembling the propeller, but such knowledge is crucial to prevent accidents. To estimate the state of the bearing, sensors are often used to collect the vibration patterns of the bearing.

In making financial and economic decisions as corporate executives and government officials do, it is often advantageous to take the upcoming economic climate into account. Is a recession or an economic boom on the horizon? Although this process is more concerned with predicting the future than evaluating current conditions, the trend can be viewed as one characteristic of the present economic state that is not directly observable.

In playing card games such as poker, knowing the opponent's hand allows a player to determine his or her best strategy with certainty. Because the opponent's hand is not observable, the best one can do is to guess it based on the cards revealed so far and the past behavior of the opponent.

2. The relations among domain events are often uncertain. In particular, the relationship between the observables and nonobservables is often uncertain. Therefore, the nonobservables cannot be inferred from the observables with certainty. The uncertain relation prevents a reasoner from arriving at a *unique* domain state from observation, which accounts for the guesswork. For instance, certain seizure behavior can be caused by lesions at different locations of the brain, which makes it difficult to determine exactly where surgery should take place.

Most equipment is intended to work deterministically. Given the input of a piece of equipment or, in general, the input history, a unique, desirable response or response sequence is expected. For example, when the brake pedal is pressed, a running car should slow down. After an enlarging scale has been entered into a photocopier, pressing the start button should cause the lens to reposition accordingly before photocopying is performed. However, equipment may fail due to failure of components and devices. Because failure is normally not designed, the faulty behavior of a piece of equipment is generally uncertain. The break down of different devices or parts of a piece of equipment

- may cause the same failure mode. On the other hand, the same faulty device may generate different failures owing to the mode of the equipment, the raw material or input being processed at the time, and other factors. A somewhat worn-out propeller bearing may last quite a while if the helicopter carries light loads and flies in good weather conditions, but the bearing may break down much sooner and cause an accident if the aircraft flies in severe weather with heavy loads.
- Earthquakes are often preceded by indicators such as sudden changes in water well levels, abnormal behavior of domestic animals, and instrumental indications. On February 4, 1975, a major earthquake (Holland [25]) struck a heavily populated area 400 miles northeast of Beijing, China. Ninety percent of the buildings in some areas were destroyed as was the entire town of Haicheng. On the basis of the indicators, a prediction was made, and nearly one million people who lived in the area were evacuated just hours before the earthquake. As a result, no one was killed when the earthquake struck. However, reliable predictions of major earthquakes are not yet a reality, and successful predictions such as this one have been rare. Consequently, disasters due to unpredicted major earthquakes do occur from time to time.
3. The observations themselves may be imprecise, ambiguous, vague, noisy, and unreliable. They introduce additional uncertainty to inferring the state of the domain. The EEG recorded from a patient cannot reflect the electromagnetic activities that are far from the scalp. On the other hand, the signals recorded from each electrode are the summation of the activities of many neurons (some normal and some lesioned) in the nearby brain plus artifacts (e.g., due to electrode movement or muscle activity).
- To monitor complex equipment, sensors are commonly used to extract information about the temperature, pressure, displacement, altitude, speed, vibration, and other physical quantities from the components or devices. The states of the components and the equipment are then inferred from the sensor outputs. However, sensors may not respond to the target quantities evenly under different conditions and hence can introduce errors. Sensors can pick up signals from nearby sources in addition to the target sources. Sensors can fail and consequently produce outputs correlated or uncorrelated to the source quantities. Sensor outputs can be contaminated by noise while being transmitted from the sensors to the processing units. For example, if the helicopter bearing vibration is monitored by a sensor and the sensor output is transmitted to a unit located in the cockpit, what is received by the unit may contain other nearby electromagnetic signals.
- When a witness testifies in court and states that he or she saw a suspect on site from a distance on the night of a murder, the reliability of this statement is to be judged by taking into account the illumination of the scene, the distance between the witness and the suspect at the time, the vision of the witness, and the relation between the witness and the suspect. Similarly, when asking for direction in an unfamiliar area, if you are told that your destination is “not too far,” a wide range of distance is still possible.
4. Even though many relevant events are observable, very often we do not have the resources to observe them all. Therefore, the state of the domain must be guessed based on incomplete information. In medicine, even though many laboratory tests may help improve the accuracy of a patient’s diagnosis, not all such tests will be performed on a

patient due to the cost involved and the potential side effects. Hence, the diagnosis must be made based on the routine physical examination and limited laboratory tests.

On battlefields, gathering more intelligence reports can result in more accurate knowledge about the enemy's location, movement, and intention and hence can lead to a more informed strategy. Because of the time and risk involved in gathering these reports, the enemy's state must very often be guessed using only limited reports.

5. Even though event relations are certain in some domains, very often it is impractical to analyze all of them explicitly. Consequently, the state of the domain is estimated from computationally more efficient but less reliable means. In board games, the configuration is certain, and the consequences of all legal moves of the player and the opponent are also certain. However, for many board games, due to the huge number of combinations of legal moves it is not feasible to analyze each of them (to an endgame) before making the current move. As a result, less reliable but more efficient heuristic functions are used to evaluate each board configuration and potential move.

A mechanical workshop manufacturing parts on contract needs to schedule which machine processes which part at which time slot. As the number of machines, the number of different parts to be manufactured, and the operations to be performed on each part increase, finding the optimal schedule for manufacturing all parts in the shortest time becomes impracticable.

In the light of these factors and others, the reasoner's task is not one of deterministic deduction but rather uncertain reasoning. That is, the reasoner must infer the state of the domain based on incomplete and uncertain knowledge about the domain and incomplete and uncertain observations. Many competing theories exist on how to perform such uncertain reasoning. This book focuses on methodologies founded on Bayesian probability theory. In other words, it focuses on *probabilistic reasoning*. A body of literature exists that compares the relative merits of alternatives.

1.4 Multiagent Systems

We have considered an agent that makes observations on a domain, performs probabilistic inference based on its knowledge about the relations among domain events, and estimates the state of the domain. However, a single agent is limited by its knowledge, its perspective, and its computational resources. As the domain becomes larger and more complex, open, and distributed, a set of cooperating agents is needed to address the reasoning task effectively.

Imagine a smart house (Boman, Davidsson, and Younes [4]), which is likely to be available in the near future. Compared with existing houses, the appliances and other components in a smart house are operated more energy efficiently and provide better comfort to the occupants. Curtains are closed automatically in the evening when it becomes dark, and lights are turned on. In the winter, heating is

reduced when people are out working and is restored to the previous setting before the occupants return. Air conditioning is handled similarly on hot summer days. The temperature in each room is adjusted individually according to the preference of the occupant(s) and the activity conducted. Perhaps the room for fitness is kept cooler than the family room. The refrigerator sends an order to the local grocery supplier through the Internet whenever the supply of milk and eggs is low. The oven starts cooking in the late afternoon before family members are home so that the dinner will be ready at the right time. The sprinkler system turns on after several dry summer days but will save the water when it rains from time to time.

The proper operation of these components depends on believing that certain events have happened or are about to happen such as “people are out working,” “people will return shortly,” “the grocery supply is low,” and “it’s been dry for quite a while.” Knowledge and reasoning generate such beliefs. Consider estimating the occurrence of “people are out working.” A simple timing based on a rigid schedule is not sufficient because one family member may be sick at home on a certain day and not follow the regular schedule. Motion detectors are not foolproof either. A dog may be wandering around the house when no one else is home. The dog may cause the motion detector to believe that a family member is home; hence, heating will not be reduced as expected. A patient may have caught the flu and be in bed without much motion during the day. Because no motion is detected and no one is believed at home, heating is reduced. This may worsen the patient’s condition. Prior knowledge about each family member’s normal work schedule, the expected activities during sickness, the existence of a pet and its behavior, as well as outputs from different sensors can all contribute, through reasoning, to the belief “people are out working.” The available prior knowledge is generally uncertain. For example, how much remaining grocery supply constitutes a “low” supply depends on the eating habits of the family, the day of the week (people may eat differently on weekdays than weekends), and other factors.

The proper operation of these components also depends on the knowledge of, and belief in, the functionality and expected behavior of appliances. When to turn off the sprinkler depends on determining whether the lawn has been watered sufficiently. This in turn depends on the knowledge of the sprinkler system’s capacity and the size of the lawn and the belief about the lawn’s degree of dryness. How early to start cooking before the family is back from work depends on the belief about what is to be cooked and knowledge of how long it takes the oven to do the cooking. Clearly, diverse knowledge about household components is needed. With new appliances installed or upgraded, an open-ended set of knowledge needs to be managed and maintained.

Often, activities in different sections of the house need to be coordinated. After the dishwasher in the kitchen is loaded and ready to wash, it may be better to

delay dish washing if it is believed that a family member will start taking a shower upstairs in the washroom. If both activities are going forward, it is likely that the hot water will run out before the person finishes the shower. Such coordination can be achieved if relevant events occurring in different areas of the house are assessed properly.

Building and maintaining a single intelligent agent to manage such complex, distributed, and open-ended activities would be very costly if not impossible. An alternative to the single-agent paradigm is the paradigm of *multiagent systems*. A multiagent system consists of a set of agents acting in a problem domain. Each agent carries only a partial knowledge representation about the domain and can observe the domain from a partial perspective. Although an agent in a multiagent system can reason and act autonomously as in the single-agent paradigm, to overcome its limit in domain knowledge, perspective, and computational resource, it can benefit from other agents' knowledge, perspectives, and computational resources through communication and coordination. This multiagent paradigm is promising for overcoming the limitation of the single-agent paradigm, as discussed in the following paragraphs.

1. In large and complex domains, diverse knowledge is required. In the smart house domain, the knowledge of different household appliances and components and of human activities and behaviors is needed to operate intelligently. A powerful tool to handle such complexity and diversity is *modularity*. Under the multiagent paradigm, for each appliance or component we can construct an agent capable of operating the unit. Because such an agent requires only limited knowledge, this approach simplifies development. The interdependence between units is handled by coordination among agents.

As a different example, consider equipment monitoring and diagnosis (M&D). The total complexity of a piece of complex modern equipment (e.g., an airplane or a chemical plant) is usually beyond the comprehension of a single person or even a single team. One reflection is that it is increasingly common for the manufacturer of a particular piece of equipment to purchase half or even more of the components from other vendors (Parunak [47]). Let R be a manufacturer who needs a component c for its product and S be the supplier of c . Then R must have the knowledge of how c should function in terms of its input–output relation so that R can integrate c with other components – purchased or manufactured. However, more detailed knowledge about the internals of c , which is necessary to monitor and diagnose c , may not be available from S . Even if it is available, R may not want to bother with it in order to manage its core business more efficiently. In such a case, to build an M&D system, R may instead use an M&D agent for c developed by S and let it cooperate with agents responsible for other components.

2. In large and complex domains, sensors are often distributed. In a smart house, sensors can collect data on temperature, humidity, object movement, lighting, water usage, and other events in each room and near the house. Components in a complex system are often physically distributed (e.g., heaters and compressors in a chemical plant). Sensors to collect




Figure 1.4: Major plants in a fertilizer factory.

observations are therefore also distributed. Figure 1.4 shows a top-level decomposition of a fertilizer factory consisting of four plants. Each plant can be further decomposed into several major components. Figure 1.5 shows the decomposition of the water treatment plant. A large number of sensors are normally placed throughout the factory to collect observations about components.

Traditionally, the observations from distributed sensors are transmitted to a central place. There, they are processed by a single agent, the necessary actions are decided, and the control signals are transmitted to the locations where the actions take place. Transmission of the observations and action control signals, however, is limited by the available bandwidth, the time delay, and potential interruption due to failure of the communication channels. The multiagent paradigm suggests deployment of multiple agents, each near a component or a small group of nearby components, so that the sensor outputs can be processed on site and actions be taken more promptly (e.g., opening a local valve to release pressure believed to be beyond the safety level).

- 3. Many complex problem domains are open-ended. Each smart house will have a different set of appliances and components in accord with the family’s need and budget. As time goes by, new appliances may be installed and existing appliances may be replaced or upgraded. Similar situations happen with equipment. New functional components may




Figure 1.5: Major equipment in a water treatment plant.