IN THIS CHAPTER

- Defining AI and its history
- Using AI for practical tasks
- Seeing through AI hype
- Connecting AI with computer technology

Chapter 1
Delving into What AI Means

Common apps, such as Google Assistant Alexa, and Siri, have all of us who are online every day, using artificial intelligence (AI) without even thinking about it. Productivity and creative apps such as ChatGPT, Synesthesia, and Gemini help us focus on the content rather than on how to get there. The media floods our entire social environment with so much information and disinformation that many people see AI as a kind of magic (which it most certainly isn't). So the best way to start this book is to define what AI is, what it isn't, and how it connects to computers today.

REMEMBER

Of course, the basis for what you expect from AI is a combination of how you define AI, the technology you have for implementing AI, and the goals you have for AI. Consequently, everyone sees AI differently. This book takes a middle-of-the-road approach by viewing AI from as many different perspectives as possible. We don't buy into the hype offered by proponents, nor do we indulge in the negativity espoused by detractors. Instead, we strive to give you the best possible view of AI as a technology. As a result, you may find that you have expectations somewhat different from those you encounter in this book, which is fine, but it's essential to consider what the technology can actually do for you — rather than expect something it can't.

Defining the Term AI

Before you can use a term in any meaningful and useful way, you must have a definition for it. After all, if nobody agrees on a meaning, the term has none; it's just a collection of characters. Defining the idiom (a term whose meaning isn't clear from the meanings of its constituent elements) is especially important with technical terms that have received more than a little press coverage at various times and in various ways.

REMEMBER

Saying that AI is an artificial intelligence doesn't tell you anything meaningful, which is why people have so many discussions and disagreements over this term. Yes, you can argue that what occurs is artificial, not having come from a natural source. However, the intelligence part is, at best, ambiguous. Even if you don't necessarily agree with the definition of AI as it appears in the sections that follow, this book uses AI according to that definition, and knowing it will help you follow the text more easily.

Discerning intelligence

People define intelligence in many different ways. However, you can say that intelligence involves certain mental activities composed of the following activities:

- **Learning:** Having the ability to obtain and process new information
- **Reasoning:** Being able to manipulate information in various ways
- **Understanding:** Considering the result of information manipulation
- **Grasping truths:** Determining the validity of the manipulated information
- **Seeing relationships:** Divining how validated data interacts with other data
- **Considering meanings:** Applying truths to particular situations in a manner consistent with their relationship
- **Separating fact from belief:** Determining whether the data is adequately supported by provable sources that can be demonstrated to be consistently valid

The list could easily grow quite long, but even this list is relatively prone to interpretation by anyone who accepts it as viable. As you can see from the list, however, intelligence often follows a process that a computer system can mimic as part of a simulation:

1. Set a goal (the information to process and the desired output) based on needs or wants.
2. Assess the value of any known information in support of the goal.
3. Gather additional information that could support the goal. The emphasis here is on information that *could* support the goal rather than on information you know *will* support the goal.
4. Manipulate the data such that it achieves a form consistent with existing information.
5. Define the relationships and truth values between existing and new information.
6. Determine whether the goal is achieved.
7. Modify the goal in light of the new data and its effect on the probability of success.
8. Repeat Steps 2 through 7 as needed until the goal is achieved (found true) or the possibilities for achieving it are exhausted (found false).

REMEMBER

Even though you can create algorithms and provide access to data in support of this process within a computer, a computer's capability to achieve intelligence is severely limited. For example, a computer is incapable of understanding anything because it relies on machine processes to manipulate data using pure math in a strictly mechanical fashion. Likewise, computers can't easily separate truth from mistruth (as described in Chapter 2). In fact, no computer can fully implement any of the mental activities described in the earlier list that describes intelligence.

As part of deciding what intelligence actually involves, categorizing intelligence is also helpful. Humans don't use just one type of intelligence; rather, they rely on multiple intelligences to perform tasks. Howard Gardner a Harvard psychologist has defined a number of these types of intelligence (for details, see the article "Multiple Intelligences" from Project Zero at Harvard University `https://pz.harvard.edu/resources/the-theory-of-multiple-intelligences`) and knowing them helps you relate them to the kinds of tasks a computer can simulate as intelligence. (See Table 1-1 for a modified version of these intelligences with additional description.)

TABLE 1-1 The Kinds of Human Intelligence and How AIs Simulate Them

Type	Simulation Potential	Human Tools	Description
Bodily kinesthetic	Moderate to High	Specialized equipment and real-life objects	Body movements, such as those used by a surgeon or a dancer, require precision and body awareness. Robots commonly use this kind of intelligence to perform repetitive tasks, often with higher precision than humans, but sometimes with less grace. It's essential to differentiate between human augmentation, such as a surgical device that provides a surgeon with enhanced physical ability, and true independent movement. The former is simply a demonstration of mathematical ability in that it depends on the surgeon for input.
Creative	None	Artistic output, new patterns of thought, inventions, new kinds of musical composition	Creativity is the act of developing a new pattern of thought that results in unique output in the form of art, music, or writing. A truly new kind of product is the result of creativity. An AI can simulate existing patterns of thought and even combine them to create what appears to be a unique presentation but is in reality just a mathematically based version of an existing pattern. In order to create, an AI would need to possess self-awareness, which would require intrapersonal intelligence.
Interpersonal	Low to Moderate	Telephone, audioconferencing, videoconferencing, writing, computer conferencing, email	Interacting with others occurs at several levels. The goal of this form of intelligence is to obtain, exchange, give, or manipulate information based on the experiences of others. Computers can answer basic questions because of keyword input, not because they understand the question. The intelligence occurs while obtaining information, locating suitable keywords, and then giving information based on those keywords. Cross-referencing terms in a lookup table and then acting on the instructions provided by the table demonstrates logical intelligence, not interpersonal intelligence.
Intrapersonal	None	Books, creative materials, diaries, privacy, time	Looking inward to understand one's own interests and then setting goals based on those interests is now a human-only kind of intelligence. As machines, computers have no desires, interests, wants, or creative abilities. An AI processes numeric input using a set of algorithms and provides an output; it isn't aware of anything it does, nor does it understand anything it does.

Type	Simulation Potential	Human Tools	Description
Linguistic (often divided into oral, aural, and written)	Low	Games, multimedia, books, voice recorders, spoken words	Working with words is an essential tool for communication because spoken and written information exchange is far faster than any other form. This form of intelligence includes understanding oral, aural, and written input, managing the input to develop an answer, and providing an understandable answer as output. Discerning just how capable computers are in this form of intelligence is difficult in light of AIs such as ChatGPT because it's all too easy to create tests where the AI produces nonsense answers.
Logical mathematical	High (potentially higher than humans)	Logic games, investigations, mysteries, brainteasers	Calculating results, performing comparisons, exploring patterns, and considering relationships are all areas in which computers now excel. When you see a computer defeat a human on a game show, this is the only form of intelligence you're seeing, out of eight kinds of intelligence. Yes, you might see small bits of other kinds of intelligence, but this is the focus. Basing an assessment of human-versus-computer intelligence on just one area isn't a good idea.
Naturalist	None	Identification, exploration, discovery, new tool creation	Humans rely on the ability to identify, classify, and manipulate their environment to interact with plants, animals, and other objects. This type of intelligence informs you that one piece of fruit is safe to eat though another is not. It also gives you a desire to learn how things work or to explore the universe and all that is in it.
Visual spatial	Moderate	Models, graphics, charts, photographs, drawings, 3D modeling, video, television, multimedia	Physical-environment intelligence is used by people like sailors and architects (among many others). To move around, humans need to understand their physical environment — that is, its dimensions and characteristics. Every robot or portable computer intelligence requires this capability, but the capability is often difficult to simulate (as with self-driving cars) or less than accurate (as with vacuums that rely as much on bumping as they do on moving intelligently).

Examining four ways to define AI

As described in the previous section, the first concept that's important to understand is that AI has little to do with human intelligence. Yes, some AI is modeled to simulate human intelligence, but that's what it is: a simulation. When thinking about AI, notice an interplay between goal seeking, data processing used to achieve

that goal, and data acquisition used to better understand the goal. AI relies on algorithms to achieve a result that may or may not have anything to do with human goals or methods of achieving those goals. With this in mind, you can categorize AI in four ways:

- Acting humanly
- Thinking humanly
- Thinking rationally
- Acting rationally

Acting humanly

When a computer acts like a human, it best reflects the *Turing test*, in which the computer succeeds when differentiation between the computer and a human isn't possible. (For details, see "The Turing test" at the Alan Turing Internet Scrapbook `www.turing.org.uk/scrapbook/test.html`). This category also reflects what most media would have you believe AI is all about. You see it employed for technologies such as natural language processing, knowledge representation, automated reasoning, and machine learning (all four of which must be present to pass the test). To pass the Turing test, an AI should have all four previous technologies and, possibly, integrate other solutions (such as expert systems).


TECHNICAL STUFF

The original Turing test didn't include any physical contact. Harnad's Total Turing Test does include physical contact, in the form of perceptual ability interrogation, which means that the computer must also employ both computer vision and robotics to succeed. Here's a quick overview of other Turing test alternatives:

- **Reverse Turing test:** A human tries to prove to a computer that the human is not a computer (for example, the Completely Automated Public Turing Test to Tell Computers and Humans Apart, or CAPTCHA).
- **Minimum intelligent signal test:** Only true/false and yes/no questions are given.
- **Marcus test:** A computer program simulates watching a television show, and the program is tested with meaningful questions about the show's content.
- **Lovelace test 2.0:** A test detects AI by examining its ability to create art.
- **Winograd schema challenge:** This test asks multiple-choice questions in a specific format.

IS THE TURING TEST OUTDATED?

Current discussions about the Turing test have researchers Philip Johnson-Laird, a retired psychology professor from Princeton University, and Marco Ragni, a researcher at the Germany-based Chemnitz University of Technology, asking whether the test is outdated. For example, If AI is making the Turing test obsolete, what might be better? This issue poses several problems with the Turing test and offers a potential solution in the form of a psychological-like evaluation. These tests would use the following three-step process to better test AIs, such as Google's LaMDA and OpenAI's ChatGPT:

- Use tests to check the AI's underlying inferences.
- Verify that the AI understands its own way of reasoning.
- Examine the underlying source code, when possible.

Modern techniques include the idea of achieving the goal rather than mimicking humans completely. For example, the Wright brothers didn't succeed in creating an airplane by precisely copying the flight of birds; rather, the birds provided ideas that led to studying aerodynamics, which eventually led to human flight. The goal is to fly. Both birds and humans achieve this goal, but they use different approaches.

Thinking humanly

A computer that thinks like a human performs tasks that require intelligence (as contrasted with rote procedures) from a human to succeed, such as driving a car. To determine whether a program thinks like a human, you must have some method of determining how humans think, which the cognitive modeling approach defines. This model relies on these three techniques:

- **Introspection:** Detecting and documenting the techniques used to achieve goals by monitoring one's own thought processes.
- **Psychological testing:** Observing a person's behavior and adding it to a database of similar behaviors from other persons given a similar set of circumstances, goals, resources, and environmental conditions (among other factors).
- **Brain imaging:** Monitoring brain activity directly through various mechanical means, such as computerized axial tomography (CAT), positron emission tomography (PET), magnetic resonance imaging (MRI), and magnetoencephalography (MEG).

After creating a model, you can write a program that simulates the model. Given the amount of variability among human thought processes and the difficulty of accurately representing these thought processes as part of a program, the results are experimental at best. This category of thinking humanly is often used in psychology and other fields in which modeling the human thought process to create realistic simulations is essential.

Thinking rationally

Studying how humans think using an established standard enables the creation of guidelines that describe typical human behaviors. A person is considered rational when following these behaviors within certain levels of deviation. A computer that thinks rationally relies on the recorded behaviors to create a guide to how to interact with an environment based on the data at hand.

The goal of this approach is to solve problems logically, when possible. In many cases, this approach would enable the creation of a baseline technique for solving a problem, which would then be modified to actually solve the problem. In other words, the solving of a problem in principle is often different from solving it in practice, but you still need a starting point.

Acting rationally

Studying how humans act in given situations under specific constraints enables you to determine which techniques are both efficient and effective. A computer that acts rationally relies on the recorded actions to interact with an environment based on conditions, environmental factors, and existing data.

As with rational thought, rational acts depend on a solution in principle, which may not prove useful in practice. However, rational acts do provide a baseline on which a computer can begin negotiating the successful completion of a goal.

HUMAN-VERSUS-RATIONAL PROCESSES

Human processes differ from rational processes in their outcome. A process is *rational* if it always does the right thing based on the current information, given an ideal performance measure. In short, rational processes go by the book and assume that the book is correct. Human processes involve instinct, intuition, and other variables that don't necessarily reflect the book and may not even consider the existing data. As an example, the rational way to drive a car is to always follow the law. However, traffic isn't rational. If you follow the law precisely, you end up stuck somewhere because other drivers aren't following the law precisely. To be successful, a self-driving car must therefore act humanly rather than rationally.

Reviewing AI categories

The categories used to define AI offer a way to consider various uses or ways to apply AI. Some of the systems used to classify AI by type are arbitrary and indistinct. For example, some groups view AI as either strong (generalized intelligence that can adapt to a variety of situations) or weak (specific intelligence designed to perform a particular task well).

The problem with strong AI is that it doesn't perform any task well, whereas weak AI is too specific to perform tasks independently. Even so, just two type classifications won't do the job, even in a general sense. The four classification types promoted by Arend Hintze form a better basis for understanding AI:

- **Reactive machines:** The machines you see defeating humans at chess or playing on game shows are examples of reactive machines. A reactive machine has no memory or experience on which to base a decision. Instead, it relies on pure computational power and smart algorithms to re-create every decision every time. This is an example of a weak AI used for a specific purpose.
- **Limited memory:** A self-driving (SD) car or an autonomous robot can't afford the time to make every decision from scratch. These machines rely on a small amount of memory to provide experiential knowledge of various situations. When the machine sees the same situation, it can rely on experience to reduce reaction time and provide more resources for making new decisions that haven't yet been made. This is an example of the current level of strong AI.
- **Theory of mind:** A machine that can assess both its required goals and the potential goals of other entities in the same environment has a kind of understanding that is feasible to some extent today, but not in any commercial form. However, for SD cars to become truly autonomous, this level of AI must be fully developed. An SD car would need to not only know that it must move from one point to another but also intuit the potentially conflicting goals of drivers around it and react accordingly. (Robot soccer, at `www.cs.cmu.edu/~robosoccer/main` and `www.robocup.org`, is another example of this kind of understanding, but at a simple level.)
- **Self-awareness:** This is the sort of AI you see in movies. However, it requires technologies that aren't even remotely possible now because such a machine would have a sense of both self and consciousness. In addition, rather than merely intuit the goals of others based on environment and other entity reactions, this type of machine would be able to infer the intent of others based on experiential knowledge.

For more on these classification types, check out "Understanding the four types of AI, from reactive robots to self-aware beings" at `theconversation.com/understanding-the-four-types-of-ai-from-reactive-robots-to-self-aware-beings-67616`. It's several years old but still pertinent.

Understanding the History of AI

Earlier sections of this chapter help you understand intelligence from the human perspective and see how modern computers are woefully inadequate for simulating such intelligence, much less actually becoming intelligent themselves. However, the desire to create intelligent machines (or, in ancient times, idols) is as old as humans. The desire not to be alone in the universe, to have something with which to communicate without the inconsistencies of other humans, is a strong one. Of course, a single book can't contemplate all of human history, so Figure 1-1 provides a brief, pertinent overview of the history of modern AI attempts.

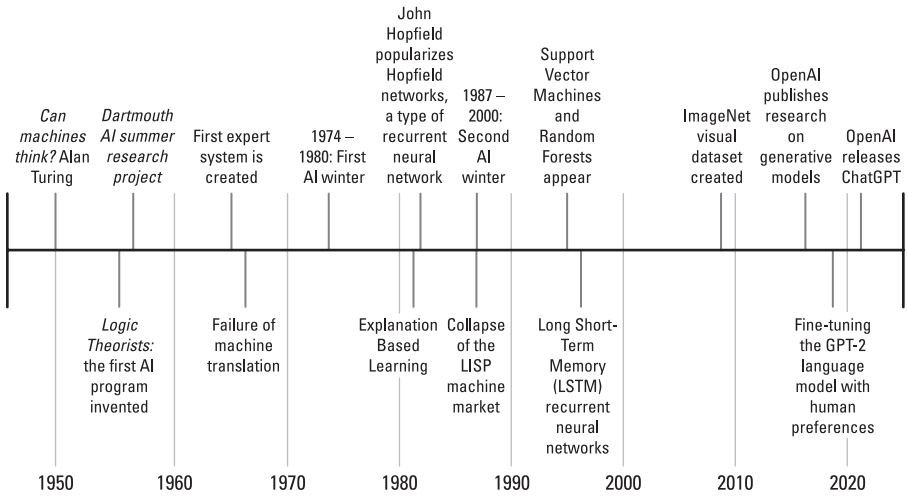

FIGURE 1-1: An overview of the history of AI.

REMEMBER

Figure 1-1 shows you some highlights, nothing like a complete history of AI. One thing you should notice is that the early years were met with a lot of disappointment from overhyping what the technology would do. Yes, people can do amazing things with AI today, but that's because the people creating the underlying technology just kept trying, no matter how often they failed.

Considering AI Uses

You can find AI used in a great many applications today. The only problem is that the technology works so well that you don't know it even exists. In fact, you might be surprised to find that many home devices already make use of AI. For example, some smart thermostats automatically create schedules for you based on how you manually control the temperature. Likewise, voice input that is used to control certain devices learns how you speak so that it can better interact with you. AI definitely appears in your car and most especially in the workplace. In fact, the uses for AI number in the millions — all safely out of sight even when they're quite dramatic in nature. Here are just a few of the ways in which you might see AI used:

- **Fraud detection:** You receive a call from your credit card company asking whether you made a particular purchase. The credit card company isn't being nosy; it's simply alerting you to the fact that someone else might be making a purchase using your card. The AI embedded within the credit card company's code detected an unfamiliar spending pattern and alerted someone to it.
- **Resource scheduling:** Many organizations need to schedule the use of resources efficiently. For example, a hospital may have to determine which room to assign a patient to based on the patient's needs, the availability of skilled experts, and the length of time the doctor expects the patient to be in the hospital.
- **Complex analysis:** Humans often need help with complex analysis because there are literally too many factors to consider. For example, the same set of symptoms might indicate more than one illness. A doctor or another expert might need help making a timely diagnosis to save a patient's life.
- **Automation:** Any form of automation can benefit from the addition of AI to handle unexpected changes or events. A problem with some types of automation is that an unexpected event, such as an object appearing in the wrong place, can cause the automation to stop. Adding AI to the automation can allow the automation to handle unexpected events and continue as if nothing happened.
- **Customer service:** The customer service line you call may not even have a human behind it. The automation is good enough to follow scripts and use various resources to handle the vast majority of your questions. After hearing good voice inflection (provided by AI as well), you may not even be able to tell that you're talking with a computer.
- **Safety systems:** Many of the safety systems now found in machines of various sorts rely on AI to take over operation of the vehicle in a time of crisis. For example, many automatic braking systems (ABSs) rely on AI to stop the

car based on all the inputs a vehicle can provide, such as the direction of a skid. Computerized ABS is, at 40 years, relatively old from a technology perspective.

- **Machine efficiency:** AI can help control a machine to obtain maximum efficiency. The AI controls the use of resources so that the system avoids overshooting speed or other goals. Every ounce of power is used precisely as needed to provide the desired services.
- **Content generation:** When people consider content generation, they often think about ChatGPT because it's in the public eye. However, content generation can exist deep within an application to provide specific functionality. For example, given a photo of the user, how will a new outfit look?

Avoiding AI Hype and Overestimation

You've no doubt seen and heard lots of hype about AI and its potential impact. If you've seen movies such as *Her* and *Ex Machina*, you might be led to believe that AI is further along than it is. The problem is that AI is actually in its infancy, and any sort of application such as those shown in the movies is the creative output of an overactive imagination. The following sections help you understand how hype and overestimation are skewing the goals you can achieve using AI today.

Defining the five tribes and the master algorithm

You may have heard of a concept called the singularity, which is responsible for the potential claims presented in the movies and other media. The *singularity (*when computer intelligence surpasses human intelligence) (when computer intelligence surpasses human intelligence*)* is essentially a master algorithm that encompasses all five "tribes" of learning used within machine learning. To achieve what these sources are telling you, the machine must be able to learn as a human would — as specified by the eight kinds of intelligence discussed in the section "Discerning intelligence," early in this chapter. Here are the five tribes of learning:

- **Symbologists:** The origin of this tribe is in logic and philosophy. It relies on inverse deduction to solve problems.
- **Connectionists:** This tribe's origin is in neuroscience, and the group relies on backpropagation to solve problems.

- **Evolutionaries:** The Evolutionaries' tribe originates in evolutionary biology, relying on genetic programming to solve problems.
- **Bayesians:** This tribe's origin is in statistics and relies on probabilistic inference to solve problems.
- **Analogizers:** The origin of this tribe is in psychology. The group relies on kernel machines to solve problems.

REMEMBER

The ultimate goal of machine learning is to combine the technologies and strategies embraced by the five tribes to create a single algorithm (the *master algorithm*) that can learn anything. Of course, achieving that goal is a long way off. Even so, scientists such as Pedro Domingos at the University of Washington are working toward that goal.

To make things even less clear, the five tribes may not be able to provide enough information to actually solve the problem of human intelligence, so creating master algorithms for all five tribes may still not yield the singularity. At this point, you should be amazed at just how little people know about how they think or why they think in a certain manner.

REMEMBER

Any rumors you hear about AI taking over the world or becoming superior to people are just plain false.

Considering sources of hype

Many sources of AI hype are out there. Quite a bit of the hype comes from the media and is presented by persons who have no idea of what AI is all about, except perhaps from a sci-fi novel they read a few years back. So it's not just movies or television that cause problems with AI hype — it's all sorts of other media sources as well. You can often find news reports presenting AI as being able to do something it can't possibly do because the reporter doesn't understand the technology. Oddly enough, many news articles are now written entirely by AI like ChatGPT, so what you end up with is a recycling of the incorrect information.

Some products should be tested much more before being placed on the market. The article "2020 in Review: 10 AI Failures" at `SyncedReview.com` (`syncedreview.com/2021/01/01/2020-in-review-10-ai-failures/`) discusses ten products, hyped by their developers, that fell flat on their faces. Some of these failures are huge and reflect badly on the ability of AI to perform tasks as a whole. However, something to consider with a few of these failures is that people may have interfered with the device using the AI. Obviously, testing procedures need to start considering the possibility of people purposely tampering with the AI as a potential source of errors. Until that happens, the AI will fail to perform as expected

because people will continue to fiddle with the software in an attempt to cause it to fail in a humorous manner.

WARNING

Another cause of problems stems from asking the wrong person about AI — not every scientist, no matter how smart, knows enough about AI to provide a competent opinion about the technology and the direction it will take in the future. Asking a biologist about the future of AI in general is akin to asking your dentist to perform brain surgery — it simply isn't a good idea. Yet many stories appear with people like these as the information source.

TIP

To discover the future direction of AI, ask a computer scientist or data scientist with a strong background in AI research.

Managing user overestimation

Because of hype (and sometimes laziness or fatigue), users continually overestimate the ability of AI to perform tasks. For example, a Tesla owner was recently found sleeping in his car while the car zoomed along the highway at 90 mph (see "Tesla owner in Canada charged with 'sleeping' while driving over 90 mph"). However, even with the user significantly overestimating the ability of the technology to drive a car, it does apparently work well enough (at least, for this driver) to avoid a complete failure.

WARNING

Be aware that there are also cases where auto drive failed and killed people such. (See the article at `www.washingtonpost.com/technology/interactive/2023/tesla-autopilot-crash-analysis`.)

However, you need not be speeding down a highway at 90 mph to encounter user overestimation. Robot vacuums can also fail to meet expectations, usually because users believe they can just plug in the device and then never think about vacuuming again. After all, movies portray the devices working precisely in this manner, but unfortunately, they still need human intervention. Our point is that most robots eventually need human intervention because they simply lack the knowledge to go it alone.

Connecting AI to the Underlying Computer

To see AI at work, you need to have some sort of computing system, an application that contains the required software, and a knowledge base. The computing system can be anything with a chip inside; in fact, a smartphone does just as well as a desktop computer for certain applications. Of course, if you're Amazon and you

want to provide advice on a particular person's next buying decision, the smartphone won't do — you need a *big* computing system for that application. The size of the computing system is directly proportional to the amount of work you expect the AI to perform.

The application can also vary in size, complexity, and even location. For example, if you're a business owner and you want to analyze client data to determine how best to make a sales pitch, you might rely on a server-based application to perform the task. On the other hand, if you're a customer and you want to find products on Amazon to complement your current purchase items, the application doesn't even reside on your computer; you access it via a web-based application located on Amazon's servers.

The knowledge base (a database that holds information about the facts, assumptions, and rules that the AI can use), varies in location and size as well.) The more complex the data, the more insight you can obtain from it, but the more you need to manipulate the data as well. You get no free lunch when it comes to knowledge management. The interplay between location and time is also important: A network connection affords you access to a large knowledge base online but costs you in time because of the latency of network connections. However, localized databases, though fast, tend to lack details in many cases.