

COGNITIVE SCIENCE

An Introduction to the Study of Mind

Jay Friedenber*g*
Manhattan College

Gordon Silverman
Manhattan College

 **SAGE Publications**
Thousand Oaks ■ London ■ New Delhi



Introduction: Exploring Inner Space

“The sciences have developed in an order the reverse of what might have been expected. What was most remote from ourselves was first brought under the domain of law, and then, gradually, what was nearer: first the heavens, next the earth, then animal and vegetable life, then the human body, and last of all (as yet very imperfectly) the human mind.”

—Bertrand Russell, 1935

A Brave New World

We are in the midst of a revolution. For centuries science has made great strides in our understanding of the external observable world. Physics revealed the motion of the planets, chemistry discovered the fundamental elements of matter, biology has told us how to understand and treat disease. But during much of this time, there were still many unanswered questions about something perhaps even more important to us. That something is the human mind.

What makes mind so difficult to study is that, unlike the phenomena described above, it is not something we can easily observe, measure, or manipulate. In addition, the mind is the most complex entity in the known universe.

To give you a sense of this complexity consider the following. The human brain is estimated to contain ten billion to one hundred billion individual nerve cells or neurons. Each of these neurons can have as many as ten thousand connections to other neurons. This vast web is the basis of mind, and gives rise to all of the equally amazing and difficult-to-understand mental phenomena such as perception, memory, and language.

The past several decades have seen the introduction of new technologies and methodologies for studying this intriguing organ. We have learned more about the mind in the past half-century than in all the time that came before that. This period of rapid discovery has coincided with an increase in the number of different disciplines—many of them entirely new—that study mind. Since then, a coordinated effort among the practitioners of these disciplines has come to pass. This interdisciplinary approach has since become known as cognitive science. Unlike the science that came before, which was focused on the world of external, observable phenomena, or “outer space,” this new endeavor turns its full attention now to the discovery of our fascinating mental world, or “inner space.”

What Is Cognitive Science?

Cognitive science can be roughly summed up as the scientific interdisciplinary study of the mind. Its primary methodology is the scientific method, although as we will see, many other methodologies also contribute. A hallmark of cognitive science is its interdisciplinary approach. It results from the efforts of researchers working in a wide array of fields. These include philosophy, psychology, linguistics, artificial intelligence, robotics, and neuroscience. Each field brings with it a unique set of tools and perspectives. One major goal of this book is to show that when it comes to studying something as complex as the mind, no single perspective is adequate. Instead, intercommunication and cooperation among the practitioners of these disciplines tell us much more.

The term *cognitive science* refers not so much to the sum of all these disciplines but to their intersection or converging work on specific problems. In this sense, cognitive science is not a unified field of study like each of the disciplines themselves, but a collaborative effort among researchers working in the various fields. The glue that holds cognitive science together is the topic of mind and, for the most part, the use of scientific methods. In the concluding chapter, we talk more about the issue of how unified cognitive science really is.

In order to really understand what cognitive science is all about we need to know what its theoretical perspective on the mind is. This perspective centers

on the idea of **computation**, which may alternatively be called information processing. Cognitive scientists view the mind as an information processor. Information processors must both represent and transform information. That is, a mind, according to this perspective, must incorporate some form of mental representation and processes that act on and manipulate that information. We will discuss these two ideas in greater detail later in this chapter.

Cognitive science is often credited with being influenced by the rise of the computer. Computers are of course information processors. Think for a minute about a personal computer. It performs a variety of information-processing tasks. Information gets into the computer via input devices, such as a keyboard or modem. That information can then be stored on the computer, for example, on a hard drive or other disk. The information can then be processed using software such as a text editor. The results of this processing may next serve as output, either to a monitor or printer. In like fashion, we may think of people performing similar tasks. Information is “input” into our minds through perception—what we see or hear. It is stored in our memories and processed in the form of thought. Our thoughts can then serve as the basis of “outputs,” such as language or physical behavior.

Of course this analogy between the human mind and computers is at a very high level of abstraction. The actual physical way in which data is stored on a computer bears little resemblance to human memory formation. But both systems are characterized by computation. In fact, it is not going too far to say that cognitive scientists view the mind as a machine or mechanism whose workings they are trying to understand.

Representation

As mentioned before, representation is fundamental to cognitive science. But what is a representation? Before listing the characteristics of a representation, it is helpful to describe briefly four categories of representation. A concept stands for a single entity or group of entities. Single words are good examples of concepts. The word “apple” denotes the concept of that particular type of fruit. Propositions are statements about the world and can be illustrated with sentences. The sentence “Mary has black hair” is a proposition that is itself made up of concepts. Rules are yet another form of representation that can specify the relationships between propositions. For example, the rule “If it is raining, I will bring my umbrella,” makes the second proposition contingent on the first. There are also analog representations. An analogy helps us to make comparisons between two similar situations. We will discuss all four of

these representations in greater detail in the In Depth section at the end of this chapter.

There are four crucial aspects of any representation (Hartshorne, Weiss & Burks, 1931–1958). First, a “representation bearer” such as a human or a computer must realize a representation. Second, a representation must have content—meaning it stands for one or more objects. The thing or things in the external world that a representation stands for are called **referents**. A representation must also be “grounded.” That is, there must be some way in which the representation and its referent come to be related. Fourth, a representation must be interpretable by some interpreter, either the representation bearer him or herself, or somebody else. These and other characteristics of representations are discussed next.

The fact that a representation stands for something else means it is **symbolic**. We are all familiar with symbols. We know for instance that the symbol “\$” is used to stand for money. The symbol itself is not the actual money, but instead is a surrogate that refers to its referent, which is actual money. In the case of mental representation, we say there is some symbolic entity “in the head” that stands for real money. Figure 1.1 shows some aspects of a mental representation of money. Mental representations can stand for many different types of things and are by no means limited to simple conceptual ideas such as “money.” Research suggests that there are more complex mental representations that can stand for rules, for example, knowing how to drive a car, and analogies, which may enable us to solve certain problems or notice similarities (Thagard, 2000). See the In Depth section for a more detailed discussion of these other forms of mental representation.

Human mental representations, especially linguistic ones, are said to be **semantic**, which is to say they have meaning. Exactly what constitutes meaning and how a representation can come to be meaningful are topics of debate. According to one view, a representation’s meaning is derived from the relationship between the representation and what it is about. The term that describes this relation is **intentionality**. Intentionality means “directed upon an object.” Mental states and events are intentional. They refer to some actual thing or things in the world. If you think about your brother, then the thought of your brother is directed toward him, not toward your sister, a cloud, or some other object.

Intentionality is considered to have at least two properties. The first is **isomorphism**, or similarity of structure between a representation and its referent. This similarity means one can map different aspects of a representation onto its referent. Analog visual images, discussed further below, are good examples of this property. This is because they are believed to preserve the spatial






Figure 1.1 Different aspects of the symbolic representation of money

characteristics of the referent. A visual image of a cruise ship, for instance, would have greater horizontal than vertical extent because these boats are much longer than they are tall. The researcher Stephen Kosslyn has shown that it takes longer to “scan” a visual image across a dimension where distances between points in the object are greater and relatively less time across a dimension where such distances are shorter. The section on visual imagery contains more on the methods and results of this experiment and others that demonstrate the isomorphic characteristics of images.

A second characteristic of intentionality has to do with the relationship between inputs and outputs to the world. An intentional representation must be triggered by its referent or things related to it. Consequently, activation of a representation (i.e., thinking about it) should cause behaviors or actions that are somehow related to the referent. For example, if your friend Sally told you about a cruise she took around the Caribbean last December, an image of a cruise ship would probably pop to mind. This might then cause you to ask her if the food on board was good. Sally’s mention of the cruise was the stimulus input that activated the internal representation of the ship in your mind. Once



A digital representation of time



An analog representation of time

Figure 1.2 Digital and analog clocks represent time in fundamentally different ways

it was activated, it caused the behavior of asking about the food. This relation between inputs and outputs is known as an **appropriate causal relation**.

Digital Representations

In a **digital representation**, sometimes also known as a symbolic representation, information is coded in a discrete way with set values. A digital clock for example, represents time discretely (see Figure 1.2). It displays a separate number for each hour, minute, or year. There are distinct advantages to digital representations. They specify values exactly. The symbols used in digital representations, such as numbers, can be operated on by a more general set of processes than analog structures. In mathematics, a wide range of operators such as addition, division, or squaring can be applied to digital number representations. The results of these operations are new numbers that can themselves be transformed by additional operations.

Language can serve as an example of a digital mental representation, and in fact verbal concepts seem to be the system of human symbolic representation that is most commonly used. The basic elements of written language are letters. These are discrete symbols that are combined according to a set of rules. The combinations, or words, have meaning and are themselves combined into other higher-order units, sentences, which also have semantic content. The rules by which these word elements are combined and transformed in language are called **syntax**. Syntax constitutes the set of permissible operations on the word elements. It is the elements themselves that are the mental representations. In

the chapter on linguistics, we talk more about linguistic representations and syntax.

Analog Representations

Analog representations, in contrast, represent information in a continuous way. Information in an analog system can theoretically take on any value not limited by resolution. Resolution refers to the amount of detail contained in an analog representation. Representations with high resolution have correspondingly more information. An analog clock represents time through the movement of its various hands. The positions of these hands on the dial indicate the time (see Figure 1.2). In addition to being able to represent large numbers of values, analog representations have the advantage of providing simple, direct solutions to some problems. They do however have a greater computational margin of error and because of the smaller number of operations that can be performed on them, are more limited for use in problem-solving.

Visual images are the best example of mental analog representations. Researchers in cognitive psychology have conducted numerous experiments that strongly suggest we represent visual information in an analog fashion. Stop reading for a moment and close your eyes. Imagine a picture of a palm tree on a sunny beach. Can you see the pattern on the bark? What about the coconuts? Images capture many of the same properties as their referents, such as distances between corresponding sets of points. The types of transformations that can be performed on images are also the kinds of changes that physical objects in the external observable world undergo. These include rotations, translations, and reflections. In the section on visual imagery in the cognitive chapter, we elaborate on the nature of visual images and discuss experiments that reveal the kinds of operations that can be performed on them.

The Dual-Coding Hypothesis

The use of both digital/symbolic and image representations collectively has been referred to as the **dual-code hypothesis** (Paivio, 1971). Alan Paivio believes that many ideas can be represented in either of these two forms interchangeably. This is especially true for a specific concrete concept, such as “elephant,” for which we can form a visual image or a verbal representation. However, there are some concepts for which a symbolic code seems more appropriate. Take the idea of “justice.” This is abstract, and although we could attach an image to it, such as that of a court building, there is no unambiguous and unique identifying image.

Evidence in support of dual-code theory comes from studies in which better recall is demonstrated for words representing concrete concepts, as compared to words representing abstract concepts (Paivio, 1971). According to Paivio, the reason for this is that two codes are better than one. Let's assume a subject in a memory experiment is presented with the word "elephant" and forms two codes to remember it. If he or she has forgotten one code later on at recall, he or she should still be able to access and retrieve the other. In this case, the image of the elephant may come to mind even if its symbolic word representation has faded.

Propositional Representations

Propositions are a third major category of representation, in addition to symbolic and imaginal codes (Pylyshyn, 1973). According to the **propositional hypothesis**, mental representations take the form of abstract sentence-like structures. Propositions are good at capturing the relationships between concepts. For example, the sentence "Mary looked at John" specifies a type of relationship between Mary and John, and that relationship can then be translated into either a verbal symbolic code, as in the actual form of a sentence, or an image code.

Propositions are believed to lie in a deep format that is neither visual nor verbal. This format can best be described as a logical relationship among constituent elements and is denoted by a **predicate calculus**. A predicate calculus is a general system of logic that accurately expresses a large variety of assertions and modes of reasoning. The proposition "Mary looked at John" can be represented by a predicate calculus such as:

$$[\text{Relationship between elements}] ([\text{Subject element}], [\text{Object element}])$$

where "Mary" is the subject element, "John" is the object element, and "looking" is the relationship between elements. What is nice about a predicate calculus is that it captures the essential logical structure of a complex idea independent of its actual elements. Any number of subjects, objects, and relationships can be inserted into the abstract format of a proposition. A proposition is thus believed to capture the basic meaning of a complex idea. This basic meaning, when translated back into a symbolic or visual code, can then be expressed in a variety of ways. For example, the sentences "Mary looked at John" and "John was looked at by Mary" are two alternate verbal codes for the same proposition. Likewise, one could form several different visual images to convey the one proposition.

Although a predicate calculus is a nice way of expressing a proposition, it doesn't mean that the proposition actually assumes this format in our brains. In fact, it is not clear exactly how propositions are mentally instantiated or realized. They do, however, serve as very useful hypothetical constructs because they are concise and can specify virtually all of the possible relationships between concepts.

To sum up this section, mental representations are powerful. They allow for the creation of an inner world that we can think about. The byproducts of these thoughts allow us to understand and interact successfully with the environment. Rather than knocking about in the world and making mistakes or taking risks, we can use representations to plan and carry out appropriate actions. Furthermore, the formal implementation of representations in a set of symbols, such as we envisage in mental pictures or language, allows us to communicate our thoughts to others. This in turn gives rise to more complex and adaptive forms of social cooperation.

Computation

As mentioned earlier, representations are only the first key component of the cognitive science view of mental processes. Representations by themselves are of little use unless something can be done with them. Having the concept of money doesn't do much for us unless we know how to calculate a tip or can give back the correct amount of change to someone. In the cognitive science view, the mind performs computations on representations. It is therefore important to understand how and why these mental mechanisms operate.

What sorts of mental operations does the mind perform? If we wanted to get detailed about it, the list would be endless. Take the example of mathematical ability. If there were a separate mental operation for each step in a mathematical process, we could say the mind adds, subtracts, divides, and so on. Likewise, with language we could say there are separate mental operations for making a noun plural, putting a verb into past tense, and so on. It is better, then, to think of mental operations as falling into broad categories. These categories can be defined by the type of operation that is performed or by the type of information acted upon. An incomplete list of these operations would include sensation, perception, attention, memory, language, mathematical reasoning, logical reasoning, decision making, and problem-solving. Many of these categories may incorporate virtually identical or similar sub-processes, for example, scanning, matching, sorting, and retrieving. Figure 1.3 shows the kinds of mental processes that may be involved in solving a simple addition problem.

| | | |
|--|--|---|
| $\begin{array}{r} 36 \\ + 47 \\ \hline 83 \end{array}$ | <p>Computational Steps</p> <ol style="list-style-type: none"> 1. $6 + 7 = 13$ 2. 3 3. 1 4. $3 + 4 = 7$ 5. $7 + 1 = 8$ 6. 8 7. 38 | <p>Add right column Store three Carry one Add left column Add one Store eight Record result</p> |
|--|--|---|

Figure 1.3 Some of the computational steps involved in solving an addition problem

The Tri-Level Hypothesis

Any given information process can be described at several different levels. According to the tri-level hypothesis, mental or artificial information-processing events can be evaluated on at least three different levels (Marr, 1982). The highest or most abstract level of analysis is the **computational level**. At this level, one is concerned with two tasks. The first is a clear specification of what the problem is. Taking the problem as it may have originally been posed, in a vague manner perhaps, and breaking it down into its main constituents or parts can bring about this clarity. It means describing the problem in a precise way such that the problem can be investigated using formal methods. It is like asking the questions: What exactly is this problem? What does this problem entail? The second task one encounters at the computational level concerns the purpose or reason for the process. The second task consists of asking: Why is this process here in the first place? Inherent in this analysis is the idea of adaptiveness—the idea that human mental processes are learned or have evolved to enable the human organism to solve a problem it faces. This is the primary explanatory perspective used in the evolutionary approach. We describe a number of cognitive processes and the putative reasons for their evolution in the chapter devoted to that approach.

Stepping down one level of abstraction, we can next inquire about the actual way in which an information process is carried out. To do this we need an **algorithm**, a formal procedure or system that acts on informational representations. It is important to note that algorithms can be carried out regardless of a representation's meaning; algorithms act on the form, not the meaning, of the symbols they transform. One way to think of algorithms is that they are “actions” used to manipulate and change representations. Algorithms are formal, meaning they are well-defined. We know exactly what occurs at each step of an

algorithm and how a particular step changes the information being acted on. A mathematical formula is a good example of an algorithm. A formula specifies how the data is to be transformed, what the steps are, and what the order of steps is. This type of description is put together at the **algorithmic level**, sometimes also called the programming level. It is equivalent to asking the question: What information-processing steps are being used to solve the problem? If we were to draw an analogy with computers, the algorithmic level is like software, because software contains instructions for the processing of data.

The most specific and concrete type of description is formulated at the **implementational level**. Here we ask: What is the information processor made of? What types of physical or material changes underlie changes in the processing of the information? This level is sometimes referred to as the hardware level, since, in computer parlance, the hardware is the physical “stuff” the computer is made of. This would include its various parts—a monitor, hard-drive, keyboard, and mouse. At a smaller scale, computer hardware consists of circuits and even the flow of electrons through the circuits. The hardware in human or animal cognition is the brain and, at a smaller scale, the neurons and activities of those neurons.

At this point, one might wonder: Why do we even need an algorithmic or formal level of analysis? Why not just map the physical processes at the implementational level onto a computational description of the problem, or alternatively, onto the behaviors or actions of the organism or device? This seems simpler, and we need not resort to the idea of information and representation. The reason is that the algorithmic level tells us how a particular system performs a computation. Not all computational systems solve a problem in the same way. Computers and humans can both perform addition, but do so in drastically different fashions. This is true at the implementational level obviously, but understanding the difference formally tells us much about alternative problem-solving approaches. It also gives us insights into how these systems might compute solutions to other novel problems that we might not understand.

This partitioning of the analysis of information-processing events into three levels has been criticized as being fundamentally simplistic, since each level can in turn be further subdivided into levels (Churchland, Koch & Sejnowski, 1990). Figure 1.4 depicts one possible organization of the many structural levels of analysis in the nervous system. Starting at the top, we might consider the brain as one organizational unit; brain regions as corresponding to another organizational unit one step down in spatial scale; and then neural networks, individual neurons, and so on. Similarly, we could divide algorithmic steps into different sub-steps, and problems into sub-problems. To compound all this, it is not entirely clear how to map one level of analysis onto another.

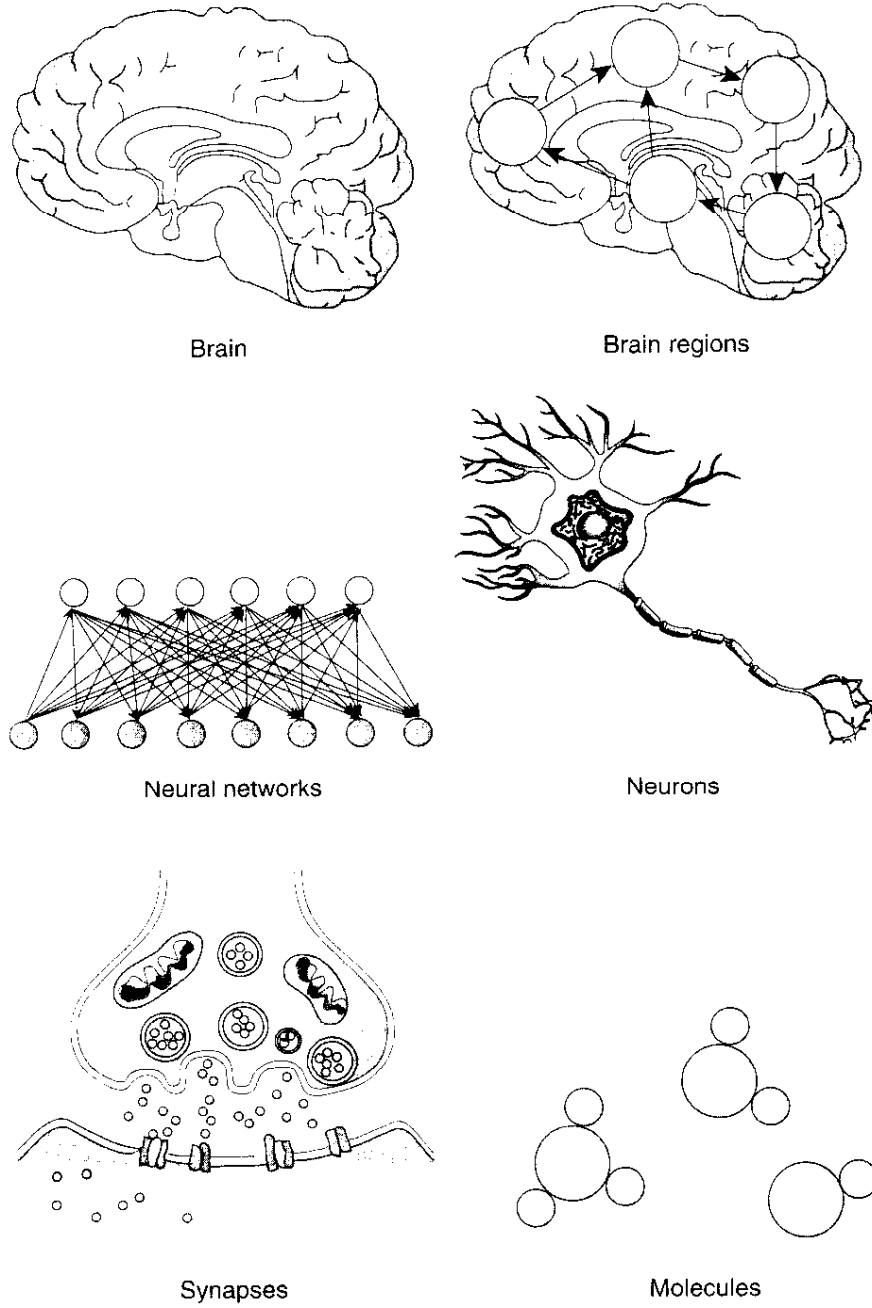


Figure 1.4 Structural levels of analysis in the nervous system

We may be able to clearly specify how an algorithm executes, but be at a loss to say exactly where or how this is achieved with respect to the nervous system.

The Classical and Connectionist Views of Computation

Before finishing our discussion of computation, it is important to differentiate between two distinct conceptions of what it is. So far, we have been talking about computation as being based on the formal systems notion. In this view a computer is a **formal symbol manipulator**. Let's break this definition down into its component parts. A system is formal if it is syntactic or rule-governed. The rules of language and mathematics are formal systems because they specify which types of allowable changes can be made to symbols. Formal systems also operate on representations independent of the content of those representations. In other words, a process can be applied to a symbol regardless of its meaning or semantic content. A symbol, as we have already indicated, is a form of representation and can assume a wide variety of forms. Manipulation here implies that computation is an active, embodied process that takes place over time. That is, manipulations are actions, they occur physically in some type of computing device, and they take some time to occur, that is, they don't happen instantaneously.

But this is not the only conception of what computation is. The network approach to computation differs from the classical formal systems approach in cognitive science in several ways. In the classical view, knowledge is represented locally, in the form of symbols. In the connectionist view knowledge is represented as a pattern of activation or weights that is distributed throughout a network. Processing style is also different in each approach. The classical view has processing occurring in discrete stages, whereas in connectionism, processing occurs in parallel through the simultaneous activation of nodes. Some cognitive scientists downplay these differences, arguing that information processing occurs in both systems and that the tri-level hypothesis can be applied equally to both (Dawson, 1998). We compare and contrast the classical and connectionist views at the beginning of the network approach chapter.

The Interdisciplinary Perspective

There is an old fable about five blind men who stumble upon an elephant (see Figure 1.5). Not knowing what it is, they start to feel the animal. One man feels only the elephant's tusk and thinks he is feeling a giant carrot. A second man,




Figure 1.5 If you were the blind man, would you know it is an elephant?

feeling the ears, believes the object is a big fan. The third feels the trunk and proclaims it is a pestle, while a fourth touching only the leg believes it is a mortar. The fifth man, touching the tail, has yet another opinion: he believes it to be a rope. Obviously, all five men are wrong in their conclusions because each has only examined one aspect of the elephant. If the five men had gotten together and shared their findings, they may have easily pieced together what kind of creature it was. This story serves as a nice metaphor for cognitive science. We can think of the elephant as the mind and the blind men as researchers in different disciplines in cognitive science. Each individual discipline may make great strides in understanding its particular subject matter, but, if it cannot compare its results to those of other related disciplines, may miss out on understanding the real nature of what it is that is being investigated.

The key, then, to figuring out something as mysterious and complex as mind is communication and cooperation among disciplines. This is what's meant when one talks about cognitive science—not the sum of each of the disciplines or approaches, but their union. Recent years have seen an increase in this cooperation. A number of major universities have established interdisciplinary cognitive science centers, where researchers in such diverse areas as philosophy, neuroscience, and cognitive psychology are encouraged to work together on common problems. Each area can then contribute its unique strength to the phenomenon under study. The philosophers can pose broad questions and hypotheses, the neuroscientists can measure physiological performance and brain activity, while the cognitive psychologists can design and carry out experiments. The

consequent exchange of results and ideas then leads to fruitful synergies between these disciplines, accelerating progress with respect to finding solutions to the problem and yielding insights into other research questions.

We have alluded to some of the different approaches in cognitive science. Because this book is about explaining each approach and its major theoretical contributions, it is worth describing each now in terms of its perspective, history, and methodology. In the following sections we will also provide a brief preview of the issues addressed by each approach.

The Philosophical Approach

Philosophy is the oldest of all the disciplines in cognitive science. It traces its roots back to the ancient Greeks. Philosophers have been active throughout much of recorded human history, attempting to formulate and to answer basic questions about the universe. This approach is free to study virtually any sort of important question on virtually any subject, ranging from the nature of existence to the acquisition of knowledge, to politics, ethics, and beauty. Philosophers of mind narrow their focus to specific problems concerning the nature and the characteristics of mind. They might ask questions like: What is mind? How do we come to know things? How is mental knowledge organized?

The primary method of philosophical inquiry is reasoning, both deductive and inductive. **Deductive reasoning** involves the application of the rules of logic to statements about the world. Given an initial set of statements assumed to be true, philosophers can derive other statements that logically must be correct. For example, if the statement “College students study three hours every night” is true and the statement “Mary is a college student” is true, we can then conclude that “Mary will study three hours every night.” Philosophers also engage in **inductive reasoning**. They make observations about specific instances in the world, notice commonalities among them, and draw conclusions. An example of inductive reasoning would be: “Whiskers the cat has four legs,” “Scruffy the cat has four legs,” therefore “All cats have four legs.” However, philosophers do not use a systematic form of induction known as the scientific method. That is employed within the other cognitive science disciplines.

In Chapter 2, we summarize several of the fundamental issues facing philosophers of mind. With respect to the mind-body problem, philosophers wrangle over what exactly a mind is. Is the mind something physical like a rock or a chair, or is it nonphysical? Can minds exist only in brains or can they emerge from the operation of other complex entities such as computers? In

the free will–determinism debate we explore whether our actions can ever be completely known and/or predicted beforehand. The knowledge acquisition problem deals with how we come to know things. Is knowledge a product of one's genetic endowment or does it arise through one's interaction with the environment? How much does each of these factors contribute to any given mental ability? We also look into one of the most fascinating and enigmatic mysteries of mind, that of consciousness. What is consciousness? Are we really conscious at all?

The Psychological Approach

Compared to philosophy, psychology is a relatively young discipline. It can be considered to be old though, particularly when it is compared to some of the more recent newcomers to the cognitive science scene, for example, artificial intelligence and robotics. Psychology arose in the late 19th century and was the first discipline in which the scientific method was applied exclusively to the study of mental phenomena. Early psychologists established experimental laboratories that would enable them to catalog mental ideas and to investigate various mental capacities, such as vision and memory. Psychologists apply the scientific method to both mind and behavior. That is, they attempt to understand not just internal mental phenomena, such as thoughts, but also the external behaviors that these internal phenomena can give rise to.

The scientific method is a way of getting hold of valid knowledge about the world. One starts with a hypothesis or idea about how the world works and then designs an experiment to see if the hypothesis has validity. In an experiment, one essentially makes observations under a set of controlled conditions. The resulting data then either support or fail to support the hypothesis. This procedure, employed within psychology and cognitive science in general, is described more fully at the start of Chapter 3.

The field of psychology is broad and encompasses many subdisciplines, each one having its unique theoretical orientations. Each discipline has a different take on what mind is. The earliest psychologists, that is, the voluntarists and structuralists, viewed the mind as a kind of test tube in which chemical reactions between mental elements took place. In contrast, functionalism viewed mind not according to its constituent parts, but according to what its operations were—what it could do. The Gestaltists again went back to a vision of mind as composed of parts, but emphasized that it was the combination and interaction of the parts, which give rise to new wholes, that was important. Psychoanalytic psychology conceives of mind as a collection of

competing entities, while behaviorism sees it as a device that maps stimuli onto behaviors.

The Cognitive Approach

Starting in the 1960s a new form of psychology arrived on the scene. Known as cognitive psychology, it came into being in part as a backlash against the behaviorist movement and its profound emphasis on behavior. Cognitive psychologists placed renewed emphasis on the study of internal mental operations. They adopted the computer as a metaphor for mind, and described mental functioning in terms of representation and computation. They believed that the mind, like a computer, could be understood in terms of information processing.

The cognitive approach was also better able to explain phenomena such as language acquisition, for which behaviorists did not have good accounts. At around the same time, new technologies that allowed better measurement of mental activity were being developed. This promoted a movement away from the behaviorist's emphasis on external observable behaviors toward the cognitive scientist's emphasis on internal functions, as these could, for the first time, be observed with reasonable precision.

Inherent in the cognitive approach is the idea of modularity. Modules are functionally independent mental units that receive inputs from other modules, perform a specific processing task, and pass the results of their computation onto yet additional modules. The influence of the modular approach can be seen in the use of process models or flow diagrams. These depict a given mental activity via the use of boxes and arrows, where boxes depict modules and arrows the flow of information among them. The techniques used in this approach are the experimental method and computational modeling. Computational modeling involves carrying out a formal (typically software-based) implementation of a proposed cognitive process. Researchers can run the modeling process so as to simulate how the process might operate in a human mind. They can then alter various parameters of the model or change its structure in an effort to achieve results as close as possible to those obtained in human experiments. This use of modeling and comparison with experimental data is a unique characteristic of cognitive psychology and is also used in the artificial intelligence and network approaches.

Cognitive psychologists have studied a wide variety of mental processes. These include pattern recognition, attention, memory, imagery, and problem-solving. Theoretical accounts and processing models for each of these are given in Chapters 4 and 5. Language is within the purview of cognitive psychology,

but because the approach to language is also multidisciplinary, we describe it separately in Chapter 9.

The Neuroscience Approach

Brain anatomy and physiology have been studied for quite some time. Recent times however have seen tremendous advances in our understanding of the brain, especially in terms of how neuronal processes can account for cognitive phenomena. The general study of the brain and endocrine system is called neuroscience. The attempt to explain cognitive processes in terms of underlying brain mechanisms is known as cognitive neuroscience.

Neuroscience, first and foremost, provides a description of mental events at the implementational level. It attempts to describe the biological “hardware” upon which mental “software” supposedly runs. However, as discussed above, there are many levels of scale when it comes to describing the brain, and it is not always clear which level provides the best explanation for any given cognitive process. Neuroscientists, however, investigate at each of these levels. They study the cell biology of individual neurons and of neuron-to-neuron synaptic transmission, the patterns of activity in local cell populations, and the interrelations of larger brain areas.

A reason for many of the recent developments in neuroscience is, again, the development of new technologies. Neuroscientists employ a wide variety of machines to measure the performance of the brain at work. These include positron emission tomography (PET) scanners, computerized axial tomography (CAT) scanners, and magnetic resonance imaging (MRI) machines. Studies that use these devices have participants perform a cognitive task; the brain activity that is concurrent with the performance of the task is recorded. For example, a participant may be asked to form a visual image of a word that appears on a computer screen. The researchers can then determine which parts of the brain became active during imagery and in what order. Neuroscientists use other techniques as well. They study brain-damaged patients and the effects of lesions in laboratory animals, and use single- and multiple-cell recording techniques.

The Network Approach

The network approach is at least partially derived from neuroscience. In this perspective, mind is seen as a collection of individual computing units. These

units are connected to one another and mutually influence one other's activity via the connections. Although each of the units is believed to perform a relatively simple computation, for example, a neuron's either firing or not firing, the connectivity of the units can give rise to representational and computational complexity.

Chapter 7, which outlines the network approach, has two parts. The first involves the construction of artificial neural networks. Most artificial neural networks are computer software simulations that have been designed to mimic the way actual brain networks operate, or the functioning of neural cell populations. Artificial neural networks that can perform arithmetic, learn concepts, and read out loud now exist. A wide variety of network architectures have developed over the last thirty years.

The second part of the network chapter is more theoretical and focuses on knowledge representation—on how meaningful information may be mentally coded and processed. In semantic networks, nodes standing for concepts are connected to one another in such a way that activation of one node causes activation of other related nodes. Semantic networks have been constructed to explain how conceptual information in memory is organized and recalled. They are often used to predict and explain data obtained from experiments with human participants in cognitive psychology.

The Evolutionary Approach

The theory of natural selection proposed by Charles Darwin in 1859 revolutionized our way of thinking about biology. Natural selection holds that adaptive features enable the animals that possess them to survive and pass these features on to future generations. The environment in this view is seen as selecting from among a variety of traits those that serve a functional purpose.

The evolutionary approach can be considered in a quite general way and used to explain phenomena outside of biology. The field of evolutionary psychology applies selection theory to account for human mental processes. It attempts to elucidate the selection forces that acted on our ancestors and how those forces gave rise to the cognitive structures we now possess. Evolutionary psychologists also adopt a modular approach to mind. In this case, the modules correspond to “favored” cognitive capacities that were used by ancestors successful at solving certain problems. Evolutionary theories have been proposed to account for experimental results across a wide range of capacities, from categorization to memory, to logical and probabilistic reasoning, language, and cognitive differences between the sexes.

A variant on this theme is evolutionary computing, in which the rules of evolution are applied to create successful computer algorithms. An offshoot of this form of computing is artificial life. These are software simulations that mimic biological ecosystems. There is also neural Darwinism, which uses evolution to explain the formation of neural circuits. See Chapter 8 for more on these.

The Linguistic Approach

Linguistics is an area that focuses exclusively on the domain of language. It is concerned with all questions concerning language ability, such as: What is language? How do we acquire language? What parts of the brain underlie language use? As we have seen, language is a topic studied within other disciplines, for example, cognitive psychology and neuroscience. Because so many different researchers in different disciplines have taken on the problem of language, we consider it here as a separate discipline, united more by topic than by perspective or methodology.

Part of the difficulty in studying language is the fact that language itself is so complex. Much research has been devoted to understanding its nature. This work looks at the properties all languages share, the elements of language, and how those elements are used during communication. Other foci of linguistic investigation center on primate language use, language acquisition, deficits in language acquisition caused by early sensory deprivation or brain damage, the relationship between language and thought, and the development of speech recognition systems.

Linguistics, perhaps more than any other perspective discussed here, adopts a very eclectic methodological approach. Language researchers employ experiments and computer models, study brain-damaged patients, track how language ability changes during development, and compare diverse languages.

The Artificial Intelligence Approach

Researchers have been building devices that attempt to mimic human and animal function for many centuries. But it is only in the past few decades that computer scientists have seriously attempted to build devices that mimic complex thought processes. This area is now known as artificial intelligence (AI). Researchers in AI are concerned with getting computers to perform tasks that have heretofore required human intelligence. As such they construct programs to do the sorts of things that require complex reasoning on our part. AI programs have been developed that can diagnose medical disorders, use language, and play chess.

AI secondarily gives us insights into the function of human mental operations. Designing a computer program that can visually recognize an object often proves useful in understanding how we may perform the same task ourselves. An even more exciting outcome of AI research is that someday we may be able to create an artificial person who will possess all or many of the features that we consider uniquely human, such as consciousness, the ability to make decisions, and so on.

It is the development of computer algorithms and their testing, their comparison with empirical data or performance standards, and their subsequent modification that constitute the methodology of the AI perspective. Not all computer programs are alike, however. Researchers have employed a wide range of approaches. An early attempt at getting computers to reason involved the application of logical rules to propositional statements. Later on, expert systems, scripts, and fuzzy logic procedures, among others, were used. Chapters 10 and 11 give detailed descriptions of these techniques.

The Robotics Approach

Finally, we consider robotics. Robotics may be considered a familial relation to AI and has appeared on the scene as a formal discipline just as recently. Whereas AI workers build devices that “think,” robotics researchers build machines that must also “act.” Investigators in this field build autonomous or semi-autonomous mechanical devices that have been designed to perform a physical task in a real world environment. Examples of things that robots can do presently include navigating around a cluttered room, welding or manipulating parts on an assembly line, and defusing bombs.

The robotics approach has much to contribute to cognitive science and to theories of mind. Robots, like people and animals, must demonstrate successful goal-oriented behaviors under complex, changing, and uncertain environmental conditions. Robotics therefore helps us to think about the kinds of minds that underlie and produce such behaviors.

In Chapter 12 we outline different paradigms in robotics. Some of these approaches differ radically from one another. The hierarchical paradigm offers a “top down” perspective, according to which a robot is programmed with knowledge about the world. The robot then uses this model or internal representation to guide its actions. The reactive paradigm, on the other hand, is “bottom up.” Robots that use this architecture respond in a simple way to environmental stimuli: they react reflexively to a stimulus input and there is little in the way of intervening knowledge.

In Depth: Categories of Mental Representation

We have said that there are three broad classes of mental representation—digital, analog, and propositional—each having its own characteristics, and we gave examples of each. However, the history of research in cognition suggests that there are also numerous forms of mental representation. Paul Thagard, in *Mind: Introduction to Cognitive Science* (2000), proposes four. These are concepts, propositions, rules, and analogies. Although some of these have already been alluded to and are described elsewhere in the book, they are central to many ideas in cognitive science. It is therefore useful to sketch out some of their major characteristics here.

A **concept** is perhaps the most basic form of mental representation. A concept is an idea that represents things we have grouped together. The concept “chair” does not refer to a specific chair, such as the one you are sitting in now, but is more general than that. It refers to all possible chairs no matter what their colors, sizes, and shapes. Concepts need not refer to concrete items. They can stand for abstract ideas, for example, “justice” or “love.” Concepts can be related to one another in complex ways. They can be related in a hierarchical fashion, where a concept at one level of organization stands for all members of the class just below it. “Golden retrievers” belongs to the category of “dogs,” which in turn belongs to the category of “animals.” We discuss a hierarchical model of concept representation in the network approach chapter. The question of whether concepts are innate or learned is discussed in the philosophical approach chapter. The artificial intelligence chapter outlines the use of structures called frames as a means of representing conceptual knowledge.

A **proposition** is a statement or assertion typically posed in the form of a simple sentence. An essential feature of a proposition is that it can be proved true or false. For instance, the statement “The moon is made out of cheese” is grammatically correct and may represent a belief that some people hold, but it is a false statement. We can apply the rules of formal logic to propositions to determine the validity of those propositions. One logical inference is called a syllogism. A syllogism consists of three propositions. The first two are premises and the last is a conclusion. Take the following syllogism:

All men like football.

John is a man.

John likes football.

Obviously, the conclusion can be wrong if either of the two premises is wrong. If it is not true that all men like football, then it might not be true that John likes football, even if he is a man. If John is not a man, then he may or may not like football, assuming all men like it. Syllogistic reasoning of this sort is the same as deductive reasoning, mentioned earlier.

You may have noticed that propositions are representations that incorporate concepts. The proposition "All men like football" incorporates the concepts "men" and "football." Propositions are more sophisticated representations than concepts because they express relationships, sometimes very complex ones, between concepts. The rules of logic are best thought of as computational processes that can be applied to propositions in order to determine their validity. However, logical relations between propositions may themselves be considered a separate type of representation. The evolutionary approach chapter provides an interesting account of why logical reasoning, which is difficult for many people, is easier under certain circumstances.

Logic is not the only system for performing operations on propositions. Rules do this as well. A **production rule** is a conditional statement of the form: "If x , then y ," where x and y are propositions. The "if" part of the rule is called the condition. The "then" part is called the action. If the proposition that is contained in the condition (x) is true, then the action that is specified by the second proposition (y) should be carried out, according to the rule. The following rules help us drive our cars:

If the light is red, then step on the brakes.

If the light is green, then step on the accelerator.

Notice that, in the first rule, the two propositions are "the light is red" and "step on the brakes." We can also form more complex rules by linking propositions with "and" and "or" statements:

If the light is red or the light is yellow, then step on the brakes.

If the light is green and nobody is in the crosswalk, then step on the accelerator.

The "or" that links the two propositions in the first part of the rule specifies that if either proposition is true, the action should be carried out. If an "and" links these two propositions, the rule specifies that both must be true before the action can occur.

Rules bring up the question of what knowledge really is. We usually think of knowledge as factual. Indeed, a proposition such as “Candy is sweet,” if validated, does provide factual information. The proposition is then an example of declarative knowledge. **Declarative knowledge** is used to represent facts. It tells us what is and is demonstrated by verbal communication. **Procedural knowledge**, on the other hand, represents skill. It tells us how to do something and is demonstrated by action. If we say that World War II was fought during the period 1939–1945, we have demonstrated a fact learned in history class. If we ski down a snowy mountain slope in the winter, we have demonstrated that we possess a specific skill. It is therefore very important that information-processing systems have some way of representing actions if they are to help an organism or machine to perform those actions. Rules are one way of representing procedural knowledge. We discuss two cognitive rule-based systems, the Atomic Components of Thought (ACT) and SOAR models, in the cognitive approach chapters.

Another specific type of mental representation is the **analogy**, although, as is pointed out below, the analogy can also be classified as a form of reasoning. Thinking analogically involves applying one’s familiarity with an old situation to a new situation. Suppose you had never ridden on a train before, but had taken buses numerous times. You could use your understanding of bus riding to figure out how to take a ride on a train. Applying knowledge that you already possess and that is relevant to both scenarios would enable you to accomplish this. Based on prior experience, you would already know that you have to first determine the schedule, perhaps decide between express and local service, purchase a ticket, wait in line, board, stow your luggage, find a seat, and so on.

Analogies are a useful form of representation because they allow us to generalize our learning. Not every situation in life is entirely new. We can apply what we have already learned to similar situations without having to figure everything out all over again. Several models of analogical reasoning have been proposed (Forbus, Gentner & Law, 1995; Holyoak & Thagard, 1995). We outline some features of analogical reasoning in the Minds On section of this chapter. You can turn to this now if you want to try to solve an analogical reasoning problem. The application of the analogical approach in artificial intelligence is called case-based reasoning and is later described in the artificial intelligence chapter.

Minds On Exercise: Analogical Reasoning

To give you a sense of what it is like to reason analogically, we present here a classic problem, first posed by Duncker (1945). It is called “the tumor problem.” Take a

moment to read it. After reading it, see if you can come up with the solution. If stumped, read past the next paragraph for the solution as well as an account of how the problem can be solved.

Suppose you are a doctor faced with a patient who has a malignant tumor in his stomach. To operate on the patient is impossible, but unless the tumor is destroyed, the patient will die. A kind of ray, at a sufficiently high intensity, can destroy the tumor. Unfortunately, at this intensity the healthy tissue that the rays pass through on the way to the tumor will also be destroyed. At lower intensities the rays are harmless to healthy tissue, but will not affect the tumor. How can the rays be used to destroy the tumor without injuring the healthy tissue?

And here is another story. This one is called "the general and fortress problem." Please read it. Does it help you to come up with a solution to the tumor problem? How?

A small country was ruled from a strong fortress by a dictator. The fortress was situated in the middle of the country, surrounded by farms and villages. Many roads led to the fortress through the countryside. A rebel general vowed to capture the fortress. The general knew that an attack by his entire army would capture the fortress. He gathered his army at the head of one of the roads, ready to launch a full-scale direct attack. However, the general then learned that the dictator had planted mines on each of the roads. The mines were set so that small bodies of men could pass over them safely, since the dictator needed to move his own troops and workers to and from the fortress. However, any large force would detonate the mines. Not only would this blow up the road, but it would also destroy many neighboring villages. It seemed impossible to capture the fortress. However, the general devised a simple plan. He divided his army into small groups and dispatched each group to the head of a different road. When all was ready, he gave the signal and each group marched down a different road. Each group continued down its road to the fortress, so that the entire army arrived together at the fortress at the same time. In this way, the general captured the fortress and overthrew the dictator.

You may have noticed a number of similarities between these two stories. The tumor is similar to the fortress. The rays that were to be used to destroy the tumor are like the soldiers sent to capture the fortress. The healthy tissue in the first story can be likened to the villages in the second. Noticing these similarities, you may have then applied a solution similar that used by the rebel general to the problem of eradicating the tumor. Like the solution of dividing up the army and sending its

soldiers down separate roads to converge on the fortress, the solution to the tumor problem involves dividing up the high-intensity ray into multiple low-intensity rays and then targeting them on the tumor at different angles. In this fashion, the rays converge on the tumor in a show of strength and destroy it without damaging any of the surrounding healthy tissue. Gick and Holyoak (1980) found that only 10% of participants in their study could solve the tumor problem correctly in the absence of their being provided with the general and fortress story. A full 75% of participants solved it correctly when they were provided with the story.

Models of analogical reasoning usually posit a novel analog that stands for a new situation or problem to be solved. That is the tumor problem in this example. They also employ an existing analog that has been derived from another learned situation. That is the general and fortress problem. The analogy is the systematic relationship between these two analogs and includes the similarities pointed out above. Most models of analogical reasoning reveal that there are four stages to the process of analogical reasoning. First is comprehension of the target problem. Second is remembering a similar source problem for which a solution is already known. Next, the source and target problems are compared. This is equivalent to mapping out similarities in their corresponding structures. As a last step, the source problem is adapted to produce a solution to the target problem.

Food for Thought: Discussion Questions

1. Many metaphors have been proposed for thinking about the mind. These range from water pumps to telephone systems. Can a corporate office building serve as a metaphor for mind? Why or why not?
2. Might concrete concepts such as "snake" be represented differently from abstract concepts such as "democracy"? Which kind of concept lends itself more easily to an analog representation? Why?
3. Describe how a handheld pocket calculator performs division at computational, algorithmic, and implementational levels of analysis.
4. Images, concepts, propositions, rules, and analogies are all forms of mental representation. Can you think of other examples?
5. Think of an instance in everyday life in which you used analogical reasoning. Describe in as much detail as possible the target and source problems, and the similarities between them.



CHAPTER REVIEW AND EXTENSIONS

Log on to the student study site at <http://www.sagepub.com/csstudy> for electronic flashcards, review quizzes, and a list of Web resources to aid you in further exploring the field of cognitive science.

Suggested Readings

- Franklin, S. (1995). *Artificial minds*. Cambridge, MA: MIT Press.
- Paivio, A. (1971). *Imagery and verbal processes*. New York: Holt, Rinehart & Winston.
- Sobel, C. P. (2001). *The cognitive sciences: An interdisciplinary approach*. Mountain View, CA: Mayfield.
- Thagard, P. (2000). *Mind: Introduction to cognitive science*. Cambridge, MA: MIT Press.