Cambridge University Press
978-1-107-05162-1 — Cognitive Science
2nd Edition
Excerpt
More Information

PART I

HISTORICAL LANDMARKS

INTRODUCTION

Here is a short, but accurate, definition of cognitive science: Cognitive science is the science of the mind. Much of this book is devoted to explaining what this means. As with any area of science, cognitive scientists have a set of problems that they are trying to solve and a set of phenomena that they are trying to model and explain. These problems and phenomena are part of what makes cognitive science a distinctive discipline. Equally important, cognitive scientists share a number of basic assumptions about how to go about tackling those problems. They share a very general conception of what the mind is and how it works. The most fundamental driving assumption of cognitive science is that minds are information processors. As we will see, this basic idea can be developed in many different ways, since there are many different ways of thinking about what information is and how it might be processed by the mind.

The chapters in this first section of the book introduce the picture of the mind as an information processor by sketching out some of the key moments in the history of cognitive science. Each chapter is organized around a selection of influential books and articles that illustrate some of the important concepts, tools, and models that we will be looking at in more detail later on in the book. We will see how the basic idea that the mind is an information processor emerged and look at some of the very different ways in which it has been developed.

We begin in Chapter 1 by surveying some of the basic ideas and currents of thought that we can, in retrospect, see as feeding into what subsequently emerged as cognitive science. These ideas and currents of thought emerged during the 1930s, 1940s, and 1950s in very different and seemingly unrelated areas. The examples we will look at range from experiments on problem-solving in rats to fundamental breakthroughs in mathematical logic, and from studies of the grammatical structure of language to information-processing models of how input from the senses is processed by the mind.

The early flourishing of cognitive science in the 1960s and 1970s was marked by a series of powerful and influential studies of particular aspects of mental functioning. In Chapter 2 we survey three examples, each of which has been taken by many to be a paradigm of cognitive science in action. These include the studies of mental imagery carried out by Roger Shepherd and various collaborators; Terry Winograd's computer program SHRDLU; and David Marr's tri-level model of the early visual system.

The latter decades of the twentieth century saw challenges to some of the basic assumptions of the "founding fathers" of cognitive science. This was cognitive science's "turn to the brain." A crucial factor here was the development of new techniques for studying the brain. These include the possibility of studying the responses of individual neurons, as well as of mapping changing patterns of activation in different brain areas. In Chapter 3 we look at two pioneering sets of experiments. The first is Ungerleider and Mishkin's initial development of the hypothesis that there are two different pathways along which visual information travels through the brain. The second is the elegant use of positron emission tomography (PET) technology by Steve Petersen and collaborators to map how information about individual words is processed in the human brain. Another important factor was the emergence of a new type of model for thinking about cognition, variously known as connectionism or parallel distributed processing. This is also introduced in Chapter 3.

CHAPTER ONE

The prehistory of cognitive science

Overview

In the late 1970s cognitive science became an established part of the intellectual landscape. At that time an academic field crystallized around a basic set of problems, techniques, and theoretical assumptions. These problems, techniques, and theoretical assumptions came from many different disciplines and areas. Many of them had been around for a fairly long time. What was new was the idea of putting them together as a way of studying the mind.

Cognitive science is at heart an interdisciplinary endeavor. In interdisciplinary research great innovations come about simply because people see how to combine things that are already out there but have never been put together before. One of the best ways to understand cognitive science is to try to think your way back until you can see how things might have looked to its early pioneers. They were exploring a landscape in which certain regions were well mapped and well understood, but where there were no standard ways of getting from one region to another. An important part of what they did was to show how these different regions could be connected in order to create an interdisciplinary science of the mind.

In this chapter we go back to the 1930s, 1940s, and 1950s – to explore the *prehistory* of cognitive science. We will be looking at some of the basic ideas and currents of thought that, in retrospect, we can see as feeding into what came to be known as cognitive science. As we shall see in more detail later on in this book, *the guiding idea of cognitive science is that mental operations involve processing information*, and hence that we can study how the mind works by studying how information is processed. This basic idea of the mind as an information processor has a number of very specific roots, in areas that seem on the face of it to have little in common. The prehistory of cognitive science involves parallel, and largely independent, developments in psychology, linguistics, and mathematical logic. We will be looking at four of these developments:

- The reaction against behaviorism in psychology (section 1.1)
- The idea of algorithmic computation in mathematical logic (section 1.2)
- The emergence of linguistics as the formal analysis of language (section 1.3)
- The emergence of information-processing models in psychology (section 1.4)

In concentrating on these four developments we will be passing over other important influences, such as neuroscience and neuropsychology. This is because until quite recently the direct study of the brain had a relatively minor role to play in cognitive science. Almost all cognitive scientists are convinced that in some fundamental sense the mind just is the brain, so that everything that happens in the mind is happening in the brain. Few, if any, cognitive scientists are *dualists*, who think that the mind and the brain are two separate and distinct things. But for a long time in the history of cognitive science it was widely held that we are better off studying the mind by abstracting away from the details of what is going on in the brain. This changed only with the emergence in the 1970s and 1980s of new technologies for studying neural activity and of new ways of modeling cognitive abilities. Until then many cognitive scientists believed that the mind could be studied without studying the brain.

1.1 The reaction against behaviorism in psychology

Behaviorism was (and in some quarters still is) an influential movement in psychology. It takes many different forms, but they all share the basic assumption that psychologists should confine themselves to studying observable phenomena and measurable behavior. They should avoid speculating about unobservable mental states, and should instead rely on non-psychological mechanisms linking particular stimuli with particular responses. These mechanisms are the product of conditioning. For examples of conditioning, think of Pavlov's dogs being conditioned to salivate at the sound of the bell, or the rewards/punishments that animal trainers use to encourage/discourage certain types of behavior.

According to behaviorists, psychology is really the science of behavior. This way of thinking about psychology leaves little room for cognitive science as the scientific study of cognition and the mind. Cognitive science could not even get started until behaviorism ceased to be the dominant approach within psychology. Psychology's move from behaviorism was a lengthy and drawn-out process (and some would say that it has not yet been completed). We can appreciate some of the ideas that proved important for the later development of cognitive science by looking at three landmark papers. Each was an important statement of the idea that various types of behavior could not be explained in terms of stimulus-response mechanisms. Instead, psychologists need to think about organisms as storing and processing information about their environment, rather than as responding mechanically to reinforcers and stimuli. This idea of organisms as information processors is the single most fundamental idea of cognitive science.

Learning without reinforcement: Tolman and Honzik, "'Insight' in rats" (1930)

Edward Tolman (1886–1959) was a behaviorist psychologist studying problem-solving and learning in rats (among other things). As with most psychologists of the time, he started off with two standard behaviorist assumptions about learning. The first assumption is that all learning is the result of *conditioning*. The second assumption is that conditioning depends upon processes of *association* and *reinforcement*.

We can understand these two assumptions by thinking about a rat in what is known as a Skinner box, after the celebrated behaviorist B. F. Skinner. A typical Skinner box is illustrated in Figure 1.1. The rat receives a reward for behaving in a particular way (pressing a lever, for example, or pushing a button). Each time the rat performs the relevant behavior it receives the reward. The reward *reinforces* the behavior. This means that the association between the behavior and the reward is strengthened and the rat's performing the behavior again becomes more likely. The rat becomes *conditioned* to perform the behavior.

The basic idea of behaviorism is that all learning is either reinforcement learning of this general type, or the even simpler form of associative learning often called classical conditioning.

In classical conditioning what is strengthened is the association between a *conditioned stimulus* (such as the typically neutral sound of a bell ringing) and an *unconditioned stimulus* (such as the presentation of food). The unconditioned stimulus is *not* neutral for the organism and typically provokes a behavioral response, such as salivation. What happens during classical conditioning is that the strengthening of the association between conditioned stimulus and unconditioned stimulus eventually

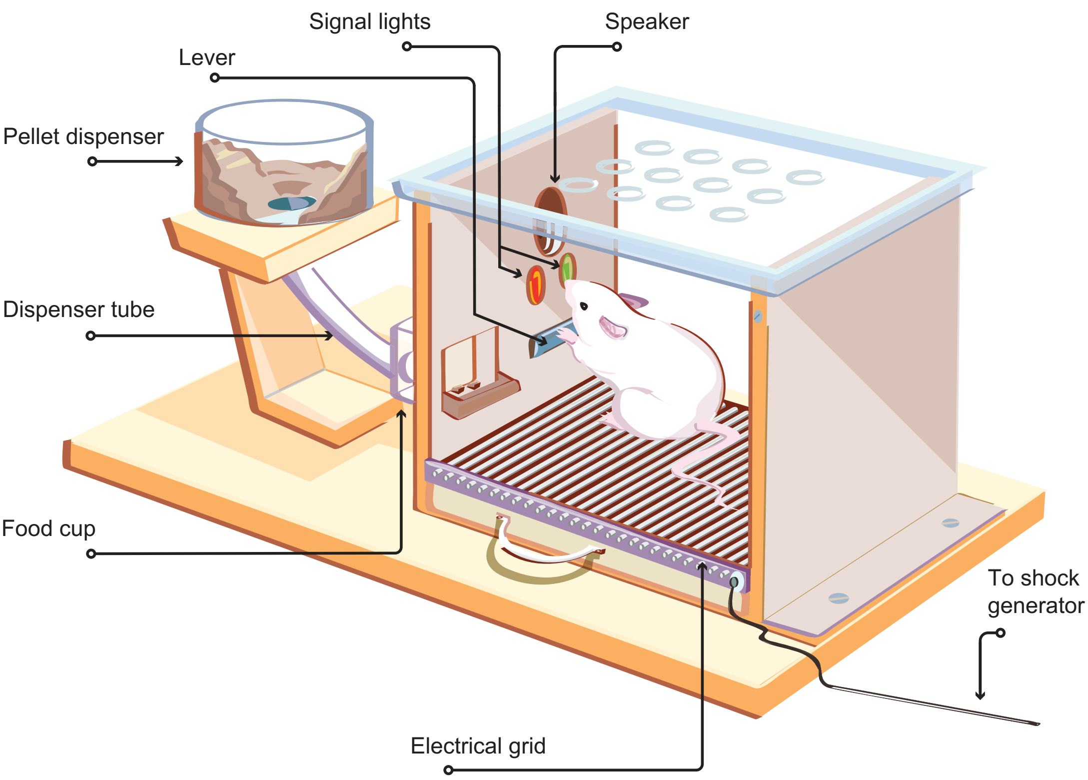

Figure 1.1 A rat in a Skinner box. The rat has a response lever controlling the delivery of food, as well as devices allowing different types of stimuli to be produced. (Adapted from Spivey 2007)

leads the organism to produce the unconditioned response to the conditioned stimulus alone, without the presence of the unconditioned stimulus. The most famous example of classical conditioning is Pavlov's dogs, who were conditioned to salivate to the sound of a bell by the simple technique of using the bell to signal the arrival of food.

So, it is a basic principle of behaviorism that all learning, whether by rats or by human beings, takes place through processes of reinforcement and conditioning. What the studies reported by Tolman and Honzik in 1930 seemed to show, however, is that this is not true even for rats.

Tolman and Honzik were interested in how rats learnt to navigate mazes. They ran three groups of rats through a maze of the type illustrated in Figure 1.2. The first group received a reward each time they successfully ran the maze. The second group never received a reward. The third group was unrewarded for the first ten days and then began to be rewarded. As behaviorism would predict, the rewarded rats quickly learnt to run the maze, while both groups of unrewarded rats simply wandered

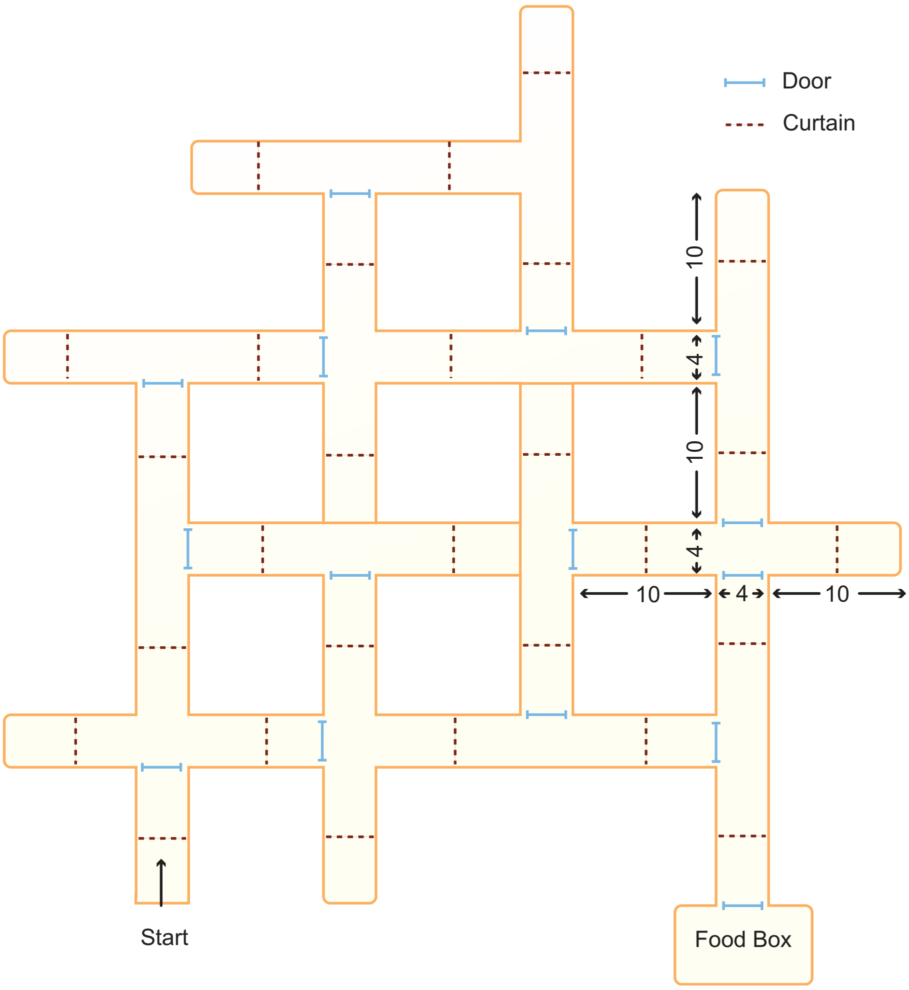

Figure 1.2 A 14-unit T-Alley maze (measurements in inches). Note the blocked passages and dead ends. (Adapted from Elliott 1928)

around aimlessly. The striking fact, however, was that when the third group of rats started to receive rewards they learnt to run the maze far more quickly than the first group had.

Tolman and Honzik argued that the rats must have been learning about the layout of the maze during the period when they were not being rewarded. This type of *latent learning* seemed to show that reinforcement was not necessary for learning, and that the rats must have been picking up and storing information about the layout of the maze when they were wandering around it, even though there was no reward and hence no reinforcement. They were later able to use this information to navigate the maze.

Exercise 1.1 Explain in your own words why latent learning seems to be incompatible with the two basic assumptions of behaviorism.

Suppose, then, that organisms are capable of latent learning – that they can store information for later use without any process of reinforcement. One important follow-up question is: What sort of information is being stored? In particular, are the rats storing information about the spatial layout of the maze? Or are they simply "remembering" the sequences of movements (responses) that they made while wandering around the maze? And so, when the rats in the latent-learning experiments start running the maze successfully, are they simply repeating their earlier sequences of movements, or are they using their "knowledge" of how the different parts of the maze fit together?

Tolman and his students and collaborators designed many experiments during the 1930s and 1940s to try to decide between *place learning* and *response learning* accounts of how rats learn to run a maze. Some of these experiments were reported in a famous article in 1946.

Cognitive maps in rats? Tolman, Ritchie, and Kalish, "Studies in spatial learning" (1946)

One experiment used a cross-maze with four end-points (North, South, East, West), like that illustrated in Figure 1.3. Rats were started at North and South on alternate trials. One group of rats was rewarded by food that was located at the same end-point, say East. The relevant feature of the map for this group was that the same turning response would not invariably return them to the reward. To get from North to East the rat needed to make a left-hand turn, whereas a right-hand turn was required to get from South to East. For the second group the location of the food reward was shifted between East and West so that, whether they started at North or South, the same turning response was required to obtain the reward. A rat in the second group starting from North would find the reward at East, while the same rat starting from South would find the reward at West. Whether it started at North or South a left turn would always take it to the reward.

This simple experiment shows very clearly the distinction between place learning and response learning. Consider the first group of rats (those for which the food was always in the same place, although their starting-points differed). In order to learn to run the maze and obtain the reward they had to represent the reward as being at a particular place and control their movements accordingly. If they merely repeated the same response they would only succeed in reaching the food reward on half of the trials. For the second group, though, repeating the same turning response would invariably bring them to the reward, irrespective of the starting-point.

Tolman found that the first group of rats learnt to run the maze much more quickly than the second group. From this he drew conclusions about the nature of

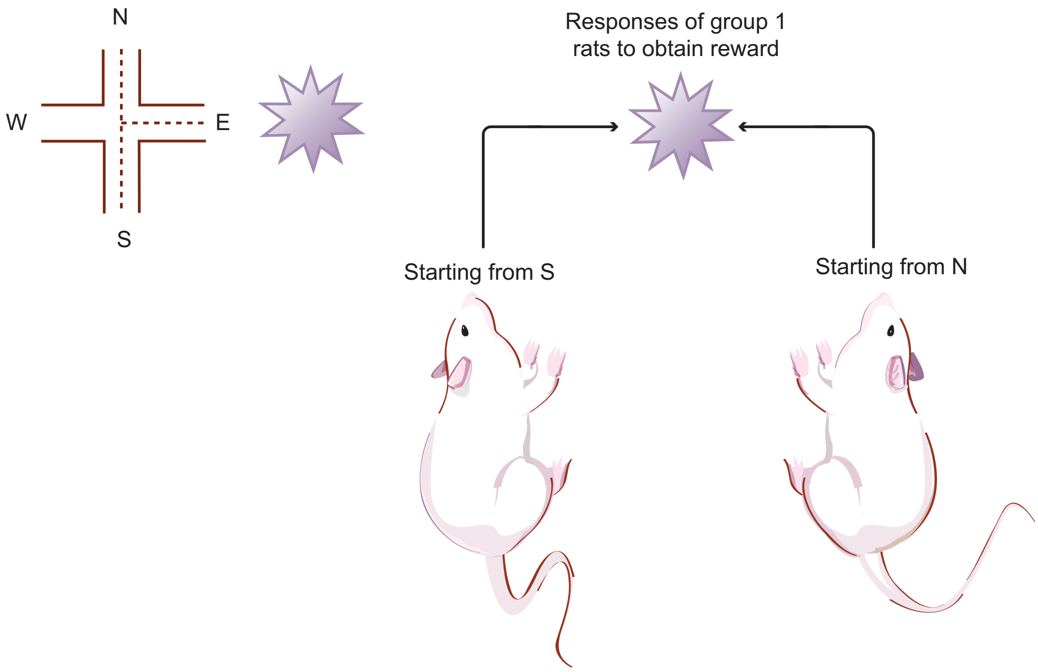

Figure 1.3 A cross-maze, as used in Tolman, Ritchie, and Kalish (1946). The left-hand part of the figure illustrates the maze, with a star indicating the location of the food reward. The right-hand side illustrates how the group 1 rats had to make different sequences of movements in order to reach the reward, depending on where they started.

animal learning in general – namely, that it was easier for animals to code spatial information in terms of places rather than in terms of particular sequences of movements.

Exercise 1.2 Explain in your own words why the experimental results seem to show that rats engage in place learning rather than response learning.

Tolman took his place-learning experiments as evidence that animals form high-level representations of how their environment is laid out – what he called *cognitive maps*. Tolman's cognitive maps were one of the first proposals for explaining behavior in terms of *representations* (stored information about the environment). Representations are one of the fundamental explanatory tools of cognitive science. Cognitive scientists regularly explain particular cognitive achievements (such as the navigational achievements of rats in mazes) by modeling how the organism is using representations of the environment. Throughout this book we will be looking at different ways of thinking about how representations code information about the environment, and about how those representations are manipulated and transformed as the organism negotiates and engages with its environment.

Plans and complex behaviors: Lashley, "The problem of serial order in behavior" (1951)

At the same time as Tolman was casting doubt on standard behaviorist models of spatial navigation, the psychologist and physiologist Karl Lashley was thinking more generally about the problem of explaining complex behavior.

Much of human and animal behavior has a very complex structure. It involves highly organized sequences of movements. Stimulus–response behaviorists have limited resources for thinking about these complex behaviors. They have to view them as linked sequences of responses – as a sort of chain with each link determined by the link immediately preceding it. This is the basic idea behind response-learning models of how rats run mazes. The standard behaviorist view is that rats learn to chain together a series of movements that leads to the reward. Tolman showed that this is not the right way to think about what happens when rats learn to run mazes. Lashley made the far more general point that this seems to be completely the wrong way to think about many complex behaviors.

Think of the complicated set of movements involved in uttering a sentence of English, for example. Or playing a game of tennis. In neither of these cases is what happens at a particular moment solely determined by what has just happened – or prompted by what is going on in the environment and influencing the organism. What happens at any given point in the sequence is often a function of what will happen later in the sequence, as well as of the overall goal of the behavior. According to Lashley, we should think about many of these complex behaviors as products of prior planning and organization. The behaviors are organized hierarchically (rather than linearly). An overall plan (say, walking over to the table to pick up the glass) is implemented by simpler plans (the walking plan and the reaching plan), each of which can be broken down into simpler plans, and so on. Very little (if any) of this planning takes place at the conscious level.

Exercise 1.3 Give your own example of a hierarchically organized behavior.

Lashley's essay contains the seeds of two ideas that have proved very important for cognitive science. The first is the idea that much of what we do is under the control of planning and information-processing mechanisms that operate below the threshold of awareness. This is the *hypothesis of subconscious information processing*. Even though we are often conscious of our high-level plans and goals (of what goes on at the top of the hierarchy), we tend not to be aware of the information processing that translates those plans and goals into actions. So, for example, you might consciously form an intention to pick up a glass of water. But carrying out the intention requires calculating very precisely the trajectory that your arm must take, as well as ensuring that your hand is open to the right degree to take hold of the glass. These calculations are carried out by information-processing systems operating far below the threshold of conscious awareness.