Representing Information

"To be a smell is one thing, to be known as a smell, another." — John Dewey

Our brains can hold an enormous amount of information: Most people know how to read, write, and say thousands of words; we know how to get from our homes to dozens of different places; psychologists have shown that we can remember thousands of different pictures. A great deal of research into the brain focuses on how it stores and represents this wealth of data.

In the modern world all sorts of devices represent information. Some, like books and maps, have been around for thousands of years. At the other extreme the World Wide Web did not exist until 1993, although the Internet has been around for several years longer. However, even the oldest book is a newcomer compared to the brain. The human brain has been representing information for millions of years.

For thousands of years philosophers have been trying to figure out how the mind stores and represents information. About a hundred years ago psychologists began conducting experiments to answer this question. To do an experiment, you have to have some idea of what you will find. Scientists call such an idea a hypothesis. As psychologists experimented, they searched for suitable hypotheses to express how the mind

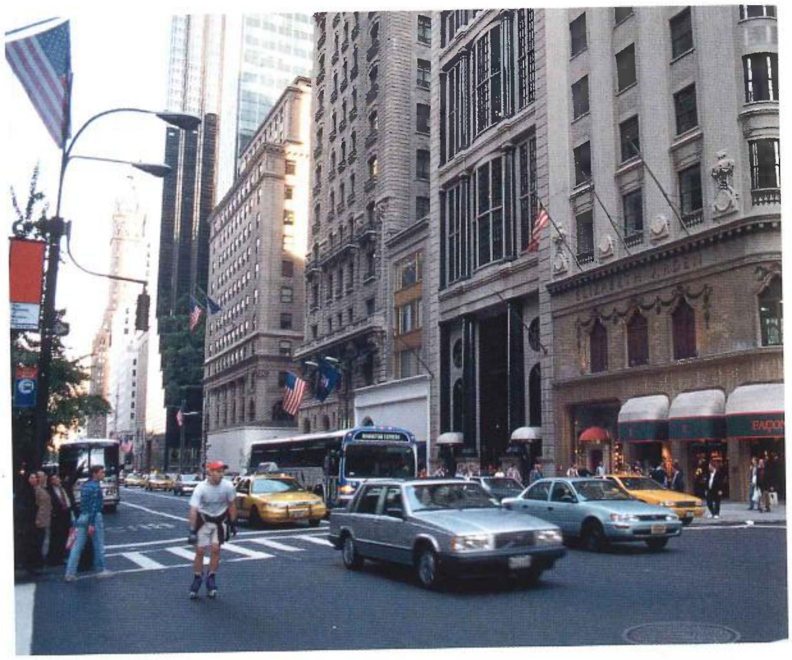

If this city street were mapped, not every detail would be shown—maps are external representations that omit unnecessary information.

KEY DATES

1883 Francis Galton investigates mental imagery.

1892 Gottlob Frege writes *On Concept and Object*.

1932 Frederick Bartlett publishes *Remembering*.

1958 Ludwig Wittgenstein publishes *Philosophical Investigations*.

1966 Ross Quillian submits his Ph.D thesis on semantic hierarchies.

1969 Brent Berlin and Paul Kay publish their findings on color terms in different languages.

1973 Researchers show that people can remember a minimum of several thousand images.

1973 Eleanor Rosch shows that categories have typical and atypical members.

1977 Roger Schank and Robert Abelson propose script theory.

1981 James McClelland publishes a network theory of concept representation.

1983 Stephen M. Kosslyn writes *Ghosts in the Mind's Machine*.

1985 James McClelland and David Rumelhart propose a network model of category learning.

1986 David Rumelhart and colleagues write *Parallel Distributed Processing*.

1993 Denis le Bihan and colleagues show that the parts of the brain used for imagining pictures are the same as those used for seeing them.

might represent information. Was the process of committing something to memory like drawing a picture inside your head? Were the stories people knew by heart stored in the mental equivalents of books? Did the brain represent our understanding of different words in the same way that a dictionary does?

People share information in many different ways. Books are written, pictures are painted, and maps are drawn. However, books, pictures, and maps are not the same as the things they stand for. A map of New York is not New York, for example. Maps, books, and pictures are representations. Representations are objects that give us useful information about the world. They also omit useless information. For example, a tourist map of New York does not mark the position of manhole covers. Tourists do not need to know where the city's manhole covers are. The inclusion of unnecessary information on a map makes it harder to read and less useful.

> *"I found that the great majority of the men of science . . . protested that mental imagery was unknown to them."*
> *—Francis Galton, 1880*

Psychologists describe maps, books, and pictures as external representations. They are distinguished from internal representations, which are the ways that the brain stores and displays potentially useful information.

PICTURES IN THE BRAIN

People have theorized about internal representations for centuries. The Greek philosopher Aristotle (384–322 B.C.) argued that memory was like storing pictures in the head. Philosophers have debated this point ever since, but scientists only joined the debate around 120 years ago. In 1883 the English scientist Francis Galton (1822–1911) investigated the imagery used by the brain by asking a number of eminent friends to imagine the way their breakfast table had looked that morning. Quite a few said they had no mental picture of their breakfast table. They could remember what they had eaten, but did not think they had a picture of the table in their heads.

KEY POINTS

- Maps, books, and pictures are external representations. They give us useful information about the world and deliberately leave out unimportant information.
- Internal representations are the way the brain stores potentially useful information.
- When we imagine a picture, we use some of the same parts of the brain as when we see the picture.
- Imagined pictures (or mental images) are similar in some ways to photographs.
- Our memory for pictures is affected by the way that we interpret them.
- Although we can remember thousands of pictures, our memory for detail fades fast.
- Mental maps can take months of experience to establish.
- Categories are groups of objects that may include nouns, verbs, or abstract concepts.
- The defining attribute view of concepts states that all concepts can be described by a list of attributes. Each attribute is necessary, and together they define the concept.
- Category membership is not all-or-none, and members differ in their typicality. Psychological research has revealed, for example, that people classify a robin as a typical bird, but not a penguin.
- When people think about a category, they tend to think of its typical members.
- The brain may store information about categories in a feature-associative network. An alternative theory suggests that information is stored in a network of specific examples.
- Human brains contain general information about the usual events that occur in particular situations. Roger Schank and Robert Abelson referred to this information as scripts.
- Stories are classed by people according to their broad themes, or schemata.
- Our expectations about what generally happens in a situation affect our memory for what actually happened.
- Our ability to remember stories is influenced by their meaningfulness to us.
- Connectionism is a way of thinking about the mind that takes into account biological knowledge of the brain.

EXPERIMENT

Mental Rotation

Imagine you are looking at two pictures of the same object, but from different angles. People are usually able to deduce that the object is the same in each picture, but how do they come to this conclusion? Many people feel as if they are turning one object in their mind's eye until it is the same way up as the other. They can then tell that the two objects are the same.

Do people really turn objects in their mind's eye to compare them? In 1971 psychologists Roger Shepard and Jacqueline Metzler conducted a series of experiments to find out. They produced a number of drawings of a pair of objects. In some drawings the objects were identical; some were drawn at the same angle, but others were drawn at angles varying between 20 and 180 degrees to each other. A second set of pictures also showed a pair of objects at various angles, but one was the mirror image of the other.

The researchers showed the drawings to a group of people and timed how long they took to decide whether the two objects were the same. When they looked at the resulting data, Shepard and Metzler noticed that for every extra degree the object had been rotated through, people took a little longer to decide whether they were the same. It seemed that people were able to turn the images in their minds at a rate of about 50 degrees a second.

In a later experiment the scientists added an arrow to the drawings that indicated which way to mentally turn the objects. Most of the time the arrow told the truth. If it pointed in a clockwise direction, then it was more efficient to go clockwise rather than counterclockwise. However, a small number of the arrows pointed in the wrong direction. This misled the subjects into mentally rotating the objects the wrong way. Again, the researchers found a close link between the angle the image was rotated (and therefore the effective distance) and the time taken to recognize the image.

Shepard and Metzler's work on mental rotation sparked many interesting research projects. In 1982 Juan Hollard and Valerie Delius performed a similar experiment with pigeons. In contrast to the human subjects of Shepard and Metzler, the pigeons did not seem to mentally rotate the images. The time the birds took to decide whether images showed the same objects was not affected by differences in the angles between them.

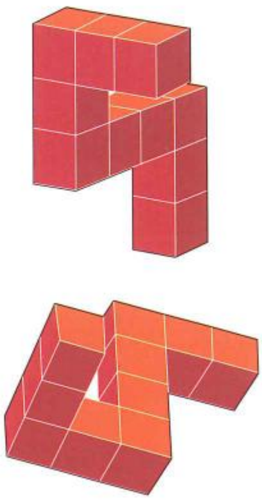

Shepard and Metzler showed images like these to subjects and asked whether or not they represented the same object at different angles. The researchers found a close link between the angle of rotation between the images and the time people took to decide whether or not the objects were identical.

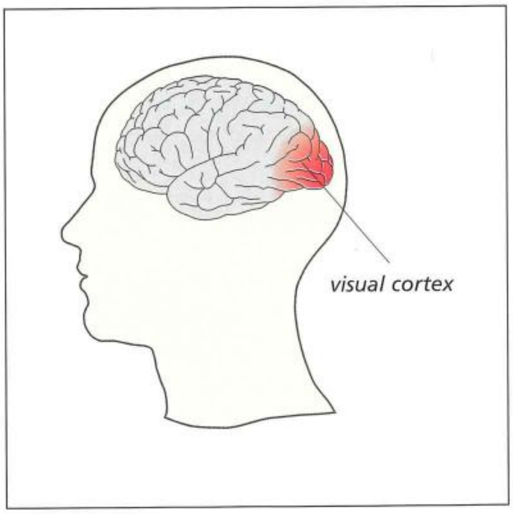

The visual cortex lies toward the rear of the brain. It is connected via the optic nerves to the retina, a structure at the back of the eyeball that translates light into electrical nerve impulses. The activity of the visual cortex is high both when looking at an image and when remembering it later on.

Psychologists now know that people can generate mental images. Techniques that map brain function, such as fMRI (*see* Vol. 2, pp. 20–39), show which parts of a person's brain are most active. When people look at pictures, a part of the brain called the primary visual cortex starts working hard. When you take the picture away, the primary visual cortex relaxes. When you ask people to imagine the picture they have just seen, the primary visual cortex starts working hard again. In fact, it works almost as hard as when the picture was present. This research shows that the same regions of the brain are highly active both when we see a picture and when we are are imagining it.

If imagining a picture we have just seen is just like seeing it, what about pictures we have never seen? People are good at forming mental images. Imagine a robin hopping across the ground. Now imagine that there is a cow standing behind it. The cow is bending its neck to look at the robin. Many people experience the same sequence of events as they imagine these pictures. First, they see a robin. The robin is large in their mental image, perhaps taking up half the picture. When they have to include the cow, they "zoom out" from the robin, or they make it smaller so there is enough space to fit the cow into the image.

> *"A word is like a key. When a word unlocks the correct stored memories, it is meaningful."*
> *—Stephen Kosslyn, 1999*

In 1975 the American psychologist Stephen Kosslyn asked people to imagine a particular animal with another one standing next to it. For example, he asked someone to imagine a rabbit sitting next to an elephant. He then asked a question about the rabbit, such as "Does the rabbit have a pointed nose?" Kosslyn then asked a different person to imagine a rabbit, but this time with a fly sitting next to it. He asked that person the same question. Kosslyn found that people took longer to answer questions about the rabbit if it was standing next to the elephant.

When the subjects in the experiment made mental images, they had to "zoom in" or "zoom out" to fit both the animals in. The mental image of the rabbit is larger when the animal is next to a fly than when it is next to an elephant. Kosslyn showed that the time taken to answer questions about mental images was closely related to the amount of "zoom" required to bring details into view.

If people are presented with large and small photographs of the same rabbit, they can spot whether the rabbit's nose is pointed more quickly when looking at the larger image. Kosslyn showed that the same was true for mental images. Just like photographs, the images we form in our minds have a limited size, and closer views are needed to determine small details.

It is tempting to describe mental images as photographs in the head. However, mental images do not represent what we have seen; rather, they represent our interpretation of what we have seen. In 1985 psychologists Deborah Chambers and Daniel Reisberg demonstrated this point in an elegant yet simple experiment. Show the image below to a friend very quickly before closing the book. Ask your friend what the picture showed, and whether there was anything else it could be. Next, ask your friend to draw the image as they remember it on a piece of paper, and then ask the questions again.

Most people think that the original picture shows either a duck or a rabbit. In fact, the picture is ambiguous. No one in the experiments could "see" both the duck and the rabbit in their mental image. However, almost all people could see the

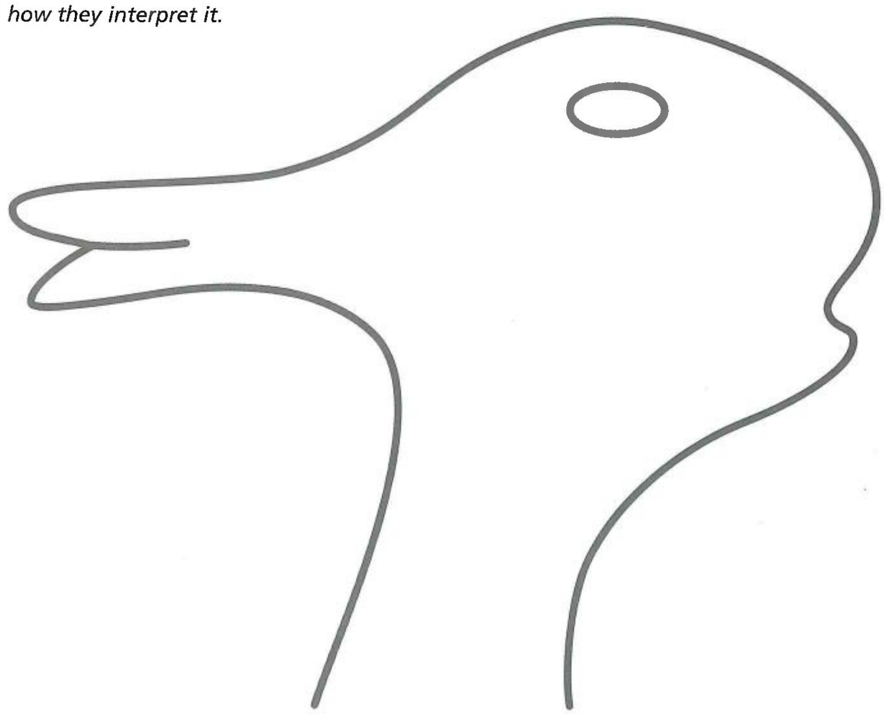

Is this a duck, or is it a rabbit? The duck–rabbit experiment works best if the subject has never seen the object before, so why not try showing it to friends to see how they interpret it.

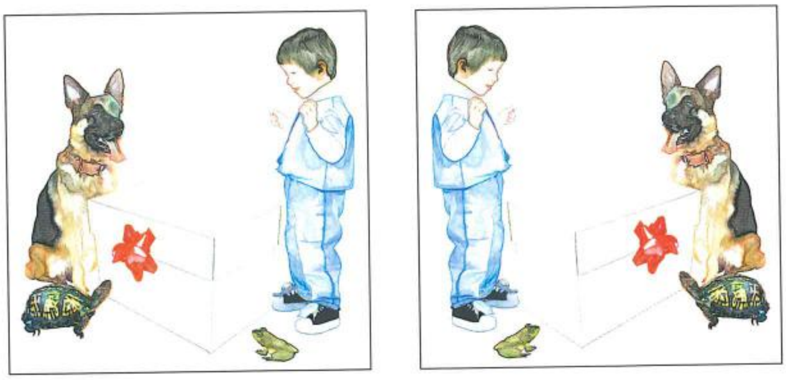

Gernsbacher showed one of a pair of images like these to subjects. She later showed them both images together and asked them which they had already seen. This research helped illustrate the short-term nature of mental pictures.

other animal once they had translated their mental image onto paper.

Mental images tend to have a fixed interpretation, while pictures and photographs in the outside world do not. The images that appear in our minds cannot be described simply as internal photographs. These pictures are internal representations, and meaning is an important part of that representation. Mental pictures are also short-lived.

An experiment by Morton Gernsbacher provided an excellent demonstration of this. Gernsbacher showed one of a pair of similar pictures, such as the ones on the left, to a subject. Ten seconds later she showed both pictures and asked which of the two was seen previously.

> *"Perception must, it seems, be a matter of seeing the present with stored objects from the past."*
> *—Richard Gregory, 1970*

After a gap of 10 seconds most got the question correct, but after 10 minutes the number of people getting the right answer did not differ statistically from guesswork alone (that is, 50 percent). The experiment worked best when the subjects did not know in advance what questions they were

EXPERIMENT

Droodles

Look at the pictures on the right. If you closed the book, do you think you could draw them accurately from memory? What if there were 28 pictures? Researcher Gordon Bower and his colleagues conducted an experiment using images like these in 1975. They showed people 28 different pictures and asked them to redraw as many as they could from memory.

People found this task tough. On average, they managed to reproduce about half of these nonsensical pictures. A second group of subjects, however, was shown the pictures with the addition of a series of short captions. For example, the caption for the picture on the left was *a midget playing a trombone in a telephone booth*. The caption for the picture on the right was *an early bird who caught a very strong worm*. When the pictures had captions that made them meaningful, the subjects found it significantly easier to remember and draw them.

Bower and his team showed that memories can be triggered by contextual cues. Bower described his pictures as droodles because without captions they were just doodles, but with captions they became drawings.

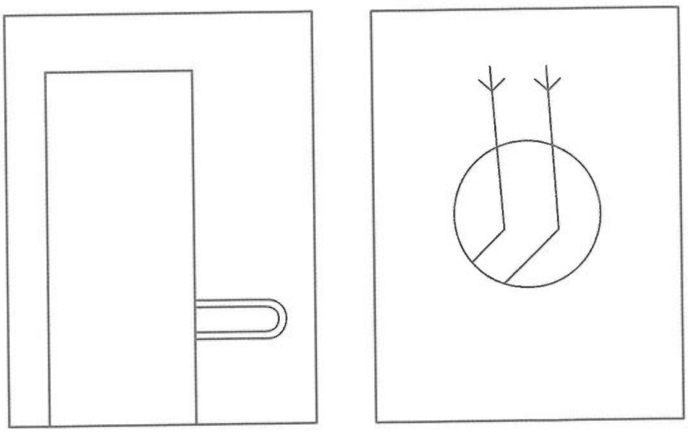

Two of the 28 droodles used by Bower and his team. People find these nonsensical images easier to remember and redraw when associated with contextual captions.

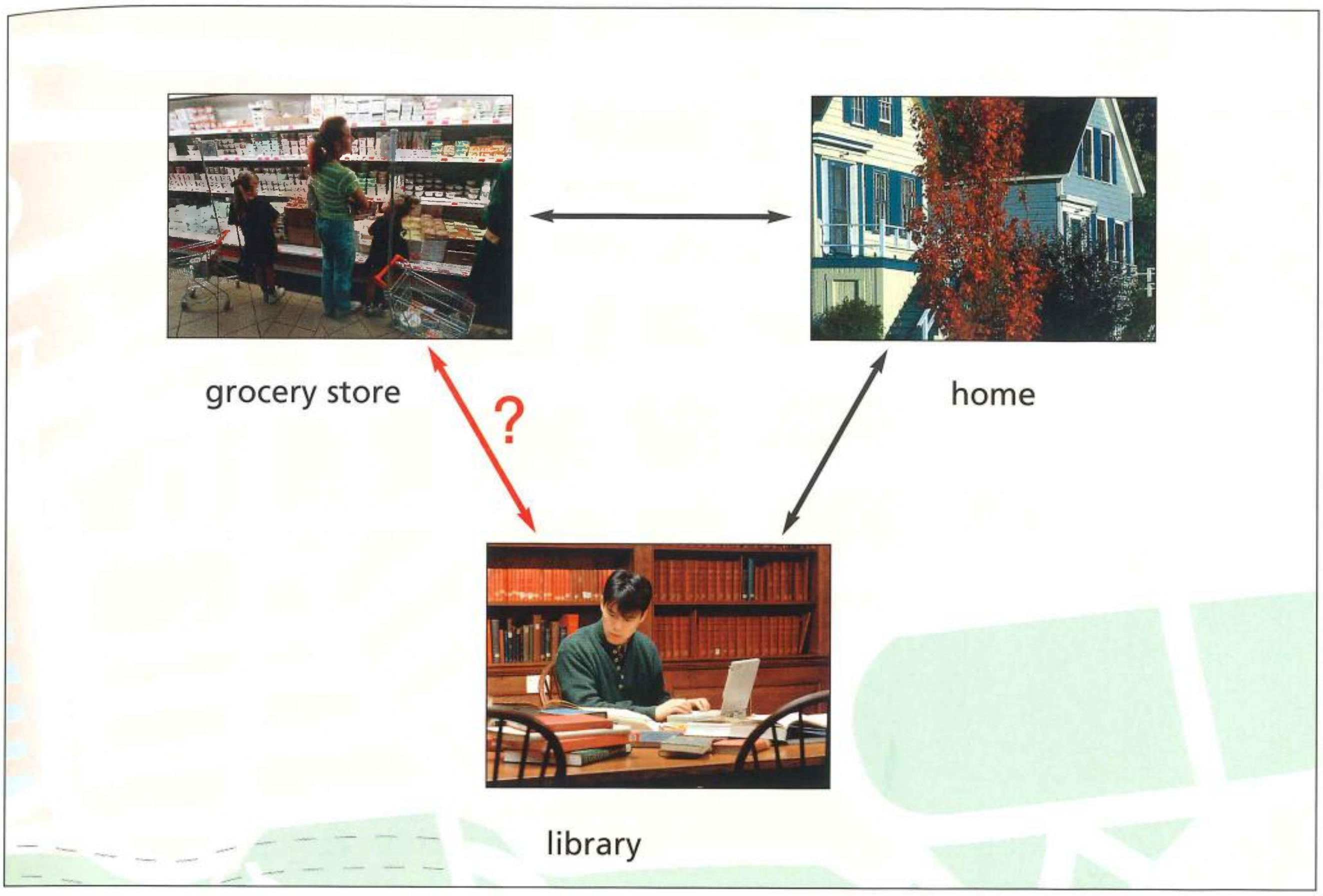

to be asked. It demonstrated that in the longer term people do not store photographlike images in their heads. They may do so initially but detail is soon lost. Some experiments suggest that mental images lose some of the information presented by a photograph within around two seconds.

> *"People remember nonsensical pictures much better if they can comprehend what they are about."*
> *—Gordon Bower, 1975*

People are good at saying whether they have seen a picture before, as long as the options they get to choose between represent different scenes or events. Later research showed that people who were shown 10,000 images could later correctly identify about 8,300 of them as pictures they had seen before.

Mental maps

Maps differ from photographs in many ways. The main difference is that they are less naturalistic. Maps tend to represent the minimum amount of information required by the user. As we have seen, mental images tend to lack fine detail in a similar way. Also, maps use false colors to aid interpretation. For example, wide and narrow roads are generally both gray in photographs, but they might be blue and green on a road map.

As we have seen, mental images are also as much about interpretation as photographic accuracy. So, if our brains represent photographs like maps, do they represent the information provided by external maps in a similar way? People are normally pretty good at remembering how to get from A to B. For example, you may know that to get from home to the subway station, you need to go down the hill, turn left at the corner, and the station is on the right. You may also know that

With instructions you know how to get from home to the grocery store and from the library back home, but could you point in the direction of the library from the grocery store? Without the help of a map the answer is "probably not."

BIOGRAPHY

Ludwig Wittgenstein

Ludwig Wittgenstein (1889–1951) was one of the most influential philosophers of the 20th century. He was born in Vienna, Austria, and originally trained as an engineer. In 1908 he moved to Manchester, England, where he was paid to experiment with kites. Later, he met the philosopher Gottlob Frege (1848–1925), who suggested he should study with the famous British philosopher Bertrand Russell (1872–1970). In 1922 Wittgenstein published *Tractatus Logico-Philosophicus*, a work that was to have a profound influence on many philosophers and psychologists. Feeling that his book had answered all the important philosophical questions, Wittgenstein gave up philosophy to become an elementary school teacher in Austria. In 1929 he returned to Cambridge University, England, to teach, and in 1939 he was awarded a chair in philosophy. As well as teaching, Wittgenstein wrote extensively. The summit of his achievements was *Philosophical Investigations*, which according to his wishes was published after his death.

Philosophical Investigations made a major contribution to our understanding of mental representation. Before its publication psychologists thought that all concepts could be represented by a set of defining attributes. A bachelor was "adult," "single," and "male." In the same way there were defining attributes for birds, chairs, and democracy. In *Philosophical Investigations* Wittgenstein challenged this idea. He wrote: "Consider for example the proceedings we call 'games.' I mean board games, card games, Olympic games, and so on. What is common to them all? Don't say: 'There must be something in common, or they would not be called games' but look and see whether there is anything common to all. For if you look at them you will not see something that is common to all, but similarities, relationships, and a whole series of them at that . . . I can think of no better expression to characterize these similarities than 'family resemblances.'"

Wittgenstein was saying that many concepts do not have defining attributes. Even if you can think of a defining attribute, then there is often the problem of "overinclusion." Games tend to involve an element of competition (even if it is with yourself), but there are all manner of things in life that involve a degree of competition that people would not normally refer to as a game.

Wittgenstein's way of thinking about concepts opened up a whole new area of psychological research. The work of a number of later researchers, such as Eleanor Rosch, can be seen as developments of the ideas of Ludwig Wittgenstein.

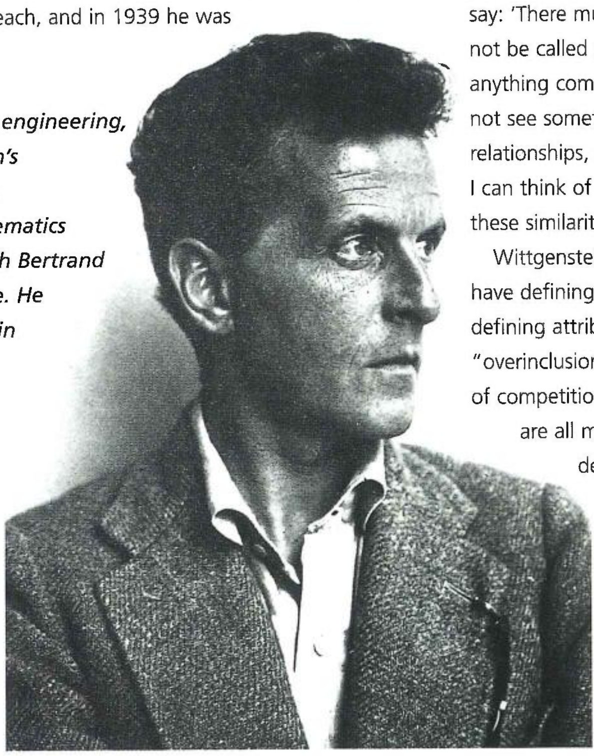

Although trained in engineering, Ludwig Wittgenstein's fascination with the philosophy of mathematics led him to study with Bertrand Russell in Cambridge. He served as an officer in the Austrian army during World War I (1914–1918) and wrote the bulk of his doctoral dissertation in an Italian prisoner-of-war camp. This work was later published as Tractatus Logico-Philosophicus.

to get from the swimming pool to your house, you go over a bridge, up the hill, follow the street around, and turn right at the corner. People remember and use this sort of information every day.

It is tempting to think that the brain holds these sets of memories as a series of mental maps. Maps, however, contain a lot more information than a set of directions and landmarks. If you were at a friend's house, you would be able to follow their directions and make your way to the grocery store. If you were at the library, you would be able to get to your friend's house. However, if you were at the library, would you be able to point in the direction of the grocery store? The answer is "probably not" unless you had a map. In most cases only people with considerable experience of a town or those who have studied a map beforehand have a good memory for such information.

In 1982 Perry Thorndyke and Barbara Hayes-Roth demonstrated the inaccuracy of most people's mental maps. They interviewed secretaries working in a particularly large and complex office building. They found that secretaries who had recently arrived could accurately describe how to get from A to B. For example, they had no difficulty giving directions from the coffee room to the computer center.

However, the new secretaries were often unable to indicate the straight-line direction of the coffee room from the computer center. Generally speaking, only secretaries who had worked in the building for several years could do this.

> *"Concept is a vague concept."*
> *—Ludwig Wittgenstein, 1958*

Even people who have had years of experience with an external map often make errors unless they have the map in front of them. If you live in the United States or Canada, ask yourself whether Montreal is farther north than Seattle. If you live in Europe, ask yourself whether London is farther north than Berlin. The answer to both questions is no, but most people would answer yes. Canada is north of the United States, but the Canadian border reaches farther south in the east of the United States than in the west of the country. Most of the United Kingdom lies north of Germany, but southern England is on the same latitude as the north of Germany.

People often make these sorts of mistakes, suggesting that the brain does not represent locations in a truly maplike way. People seem to reason about where cities are, for example, from the location of the broader regions that contain them. This often leads to mistakes.

A dictionary definition of an elephant would probably include the terms "large," "gray," "mammal," and "tusks." They are all necessary defining attributes. Put them together, and they become sufficient to define the animal.

DICTIONARIES IN THE BRAIN

Dictionaries store information about the properties of objects. They also store information about actions (verbs) and abstract concepts like democracy. People also store much of this information in the brain. Does the brain represent this information in the same way that dictionaries do? Psychologists usually focus on objects such as cats, shoes, or hammers. They may also look at ill-defined categories—"patients suffering from a mental disorder," for example.

Writers of dictionary entries aim to present a list of defining attributes or features. For example, the *Cambridge*

English Dictionary defines an elephant as "a very large gray mammal which has a long nose (trunk) with which it can pick things up." Gottlob Frege (1848–1925) was the first to suggest that all concepts could be described with a set of defining attributes. The "defining attributes" theory is best explained by example. Think of the word *bachelor*. The defining attributes of this concept are "male," "single," and "adult." Each attribute is "necessary." If any is missing, the person cannot be a bachelor. Together the three attributes are "sufficient." If you know someone is an adult single male, you can be sure that he is a bachelor—no further information is required. The idea that all visible objects and all concepts could be represented by defining attributes came to dominate philosophical and psychological thinking for a time but was strongly opposed by Ludwig Wittgenstein (*see* box p. 70).

Psychologists describe groups of objects that share certain defining characteristics as "categories." The objects that make up categories are called "members." The views of Frege led to the conclusion that all objects must either be classified as members or nonmembers of a category.

For example, all objects either are or are not members of the category "furniture." Membership in a category is all-or-nothing; there are no shades of gray. However, the decisions people make when allocating objects to categories do not seem to follow this rule. Psychologists Michael McClosky and Sam Glucksberg asked people whether certain objects belonged to the category "furniture." Everyone agreed that chairs were furniture and that cucumbers were not. When they came to bookends, however, some people thought that they were classed as furniture, while others did not. In addition, people were inconsistent about their definitions. The researchers asked people about objects like bookends on a number of occasions. Some people said that bookends were furniture the first time they were asked but not the second, or on the second time but not the first.

If peoples' mental dictionaries contained lists of defining attributes, the results of the experiment should have provided complete agreement about whether bookends were furniture or not. We would expect decisions about common categories to remain constant from one day to the next.

The research of Eleanor Rosch revealed further problems with the defining attribute view. If the mental dictionary is simply a list of defining attributes, there should be no such thing as a good or a bad example of something, such as a bird. All objects should either be birds or not birds. Rosch asked people to rate how typical they thought various members of categories were. People generally agreed on typical

A robin and a penguin—which is the more "typical" bird? The work of Eleanor Rosch suggests that most people think of a robin as an example of a typical bird, but do not think the same of a penguin. This is despite the fact that penguins have feathers, and their females lay eggs—zoologically, they are as much birds as robins are.

Are Some Numbers Odder than Others?

EXPERIMENT

When people are asked to think about categories, they tend to think about typical members. If you are asked to think of a sport, you are more likely to think of football than of weight lifting. This depends on how quickly we categorize things—typical category members are categorized more quickly and therefore more easily.

Are all categories like this? What about odd numbers? Odd numbers are those that leave a remainder when divided by two. The number 106 is even because when it is divided by two, the answer is a whole number, 53. The number 23 is odd because when it is divided by two, the answer, 11.5, is not a whole number. All numbers are either odd or even. There is no doubt and no debate.

Anyone who can divide by two can tell whether a number is odd or even. Odd numbers seem to be a perfect example of a category that is defined by a clear rule. Does this mean there is no such thing as a typical odd number? Psychologist Sharon Armstrong and colleagues decided to investigate. They gave some people a list of odd numbers such as 501, 3, and 57. They then asked them to rate each number on the list for how "good" an odd number it was. People tended to agree that 3 and 7 were good examples of odd numbers. They also agreed that 501 and 447 were bad examples of odd numbers. The same sort of result was found for even numbers. The numbers 4, 8, and 10 were "good" even numbers, while the numbers 34 and 106 were "bad" even numbers. Psychologists disagree on why this should happen. One theory is that people do not really use the "any number that does not leave a whole number when divided by two" rule to decide whether a number is odd. For a start, most people know that 1, 3, 5, 7, and 9 are odd numbers without having to figure it out. They have learned this information by rote in math lessons. As a result, 1, 3, 5, 7, and 9 come to mind when we think about odd numbers. This leads us to rate them as "good" or typical odd numbers.

When we see a number like 501, we do not know immediately whether it is odd, but we can figure it out swiftly. If the final digit is odd, then the number is odd. Although we can see that 501 is odd, it is not typical because it is encountered relatively infrequently.

and atypical members. For example, people agreed that robins were typical birds but that penguins were not. If people's mental dictionaries were as Frege had suggested, there should be no such thing as a typical bird. The question should be meaningless, forcing people to guess. When people guess, they tend not to agree. The fact that people did agree suggested there was something more to their concepts than a series of defining attributes.

Rosch wanted to show that typicality was central to the way people thought about categories. She did so by showing sentences like this to college students:

- A robin is a bird.
- A chicken is a bird.

The students had to decide as quickly as possible whether each sentence was true or false. They made their decisions more quickly when the object was a typical example of its category. For example, they took less time to agree that "a robin is a bird" than to agree that "a chicken is a bird." Obviously, both of these questions are easy to answer, but people do take measurably longer to answer the second question, although the time difference is measured in fractions of a second.

> *"Most, if not all, categories do not have clear-cut boundaries."*
> *—Eleanor Rosch, 1978*

Rosch suggested that when people are asked to think about a category, they do not think of a list of defining attributes. Instead, they think of typical members of that category. If you are asked to think about "birds," you tend to think of some typical birds. Perhaps a robin comes to mind. If you are asked whether a robin

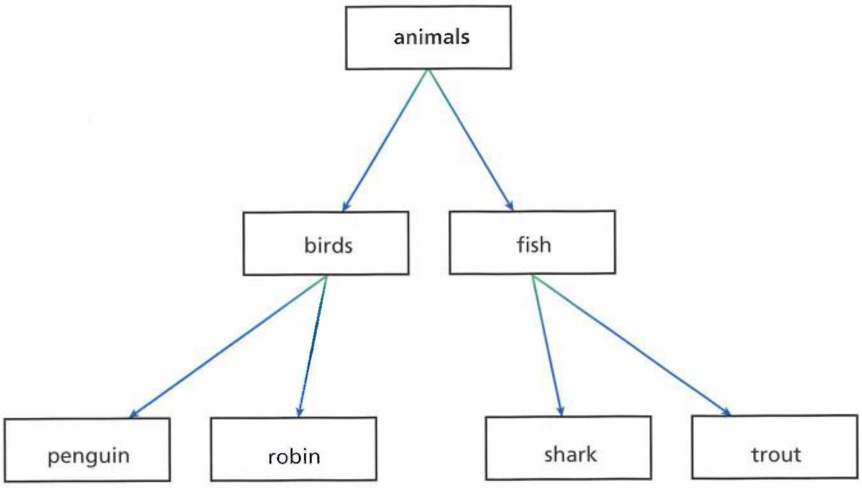

An example of a hierarchy. "Animals" sits at the top of the tree. "Birds" and "fish" are categories of animals, and they can be subdivided further into more specific examples, such as "trout" or "robin."

is a bird, the answer is easy because the term "bird" brings robins to mind. If you are asked whether a dolphin is a mammal, this takes longer because the word *mammal* is more likely to make you think of other more typical mammals.

Even when categories are easily defined by attributes, people still seem to be influenced by typicality. We have seen that "bachelor" can be defined by the attributes "single," "adult," and "male." Yet people tend to agree that some bachelors are more typical than others. For example, Tarzan is not a typical bachelor because he had no chance to marry when he lived in the jungle. Even concepts like numbers differ in typicality (*see* box p. 73).

Hierarchies

We have seen that a dictionary defines an elephant as "a very large gray mammal." In dictionary definitions words like *mammal* are quite common. Dictionary writers try to define objects as part of a "hierarchy." If you look at the diagram above, you will see that at the top of the hierarchy is the term "animals." Both birds and fish are types of animals, so they sit below "animals" on the hierarchy and are connected to it by downward arrows. Robins and penguins are both types of birds, so they are connected to "bird." In the same way "trout" and "shark" are both "fish" and are connected to "fish." Dictionary writers use hierarchies because they help shorten definitions.

If the dictionary states that "a robin is a bird," the reader knows that a robin has feathers and wings, and that the female lays eggs. The dictionary does not need to include "a female robin lays eggs" in the definition because the phrase "a robin is a bird" already tells the reader that. Might the brain use the same trick to reduce the amount of information it has to store?

Allan Collins and Ross Quillian argued that the answer was "yes." They presented a series of sentences like this to students:

- Canaries can sing.
- Canaries have feathers.

The students were quick to agree that canaries could sing. They took longer to agree that canaries had feathers. If the brain were organized like a dictionary, that is exactly what you would expect to find. If you imagine you know nothing about birds, you will need a dictionary to check whether "canaries can sing." If you look up "canary," the dictionary will tell you that "canaries can sing." That is because not all birds can sing, so singing has to be part of the definition. However, the dictionary does not mention feathers.

Not all species of birds are able to sing melodious tunes, but among those that can are canaries. A dictionary definition will need to include the phrase "canaries can sing." However, the phrases "canaries have feathers" and "female canaries lay eggs" are redundant as long as the dictionary definition mentions that "canaries are birds."

CASE STUDY

Color in Different Cultures

How many different colors are there? One answer is "about seven million." That is the number of different colors that our eyes can detect. Modern computer screens can display at least 16 million different colors. Compare these huge numbers to the number of different words you know. Very few people know more than about 80,000 words in their native tongue. Most of us can see about 100 times more colors than we have words in our language. Even if every word in your language described a color, each word would have to represent about 100 different colors.

How many different color words are there? In some languages, like English, there are quite a lot. However, there are many we do not use very often. Magenta (a dark purple-red) is one example. There are also some that only apply to certain types of object. For example, "blond" is really only used to describe hair (and some types of beers). In a 1969 paper psychologists Brent Berlin and Paul Kay suggested that there were just 11 basic color terms. They are: black, white, red, green, yellow, blue, brown, purple, pink, orange, and gray.

These 11 basic color terms occur in most languages and mean basically the same thing in all of them. Berlin and Kay visited many different countries with many different languages. They showed people in these countries more than 300 little colored squares. They then asked them to pick out the best red, the best green, and so on. They asked this for all 11 basic color terms in 20 different languages. People from all over the world agreed on the best representatives of the basic colors.

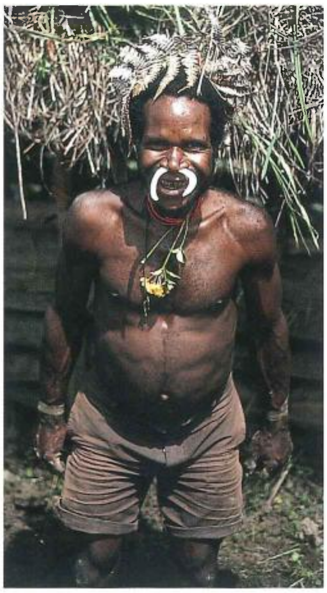

The Dani people live in the highlands of Irian Jaya, New Guinea, where they grow sweet potatoes and raise pigs. The Dani language is unusual because it does not contain any color terms—there are words only for "light" and "dark."

But not all languages have color terms. The Dani people of Irian Jaya, New Guinea, have only two color-related words in their language. They have one word for dark things and another for light things. Eleanor Rosch used Berlin and Kay's colored squares technique to teach some of the Dani the names for some basic colors. The colors that people around the world had agreed were the best examples of basic colors were the ones that the Dani found easiest to learn.

Berlin and Kay argued that basic color terms seem to represent something common to all people. The brain is geared to select certain frequencies of light as typical examples of colors. This is mirrored in people around the world regardless of ethnic group or culture.

It tells you that a canary is a bird. If you look up "bird," the dictionary tells you that it has feathers. You have your answer, but only after looking in two different places, which has taken a measurably longer period of time.

Collins and Quillian believed that the human brain organizes information into dictionarylike hierarchies. Many psychologists liked this idea, and for a time it was popular. However, it soon became apparent that Collins and Quillian were wrong. Another group of psychologists—Edward Smith, Edward Shoben, and Lance Rips—gave students a slightly different series of sentences. Two of the sentences that the researchers used were:

- A chicken is a bird.
- A chicken is an animal.

If the brain was like a dictionary, it should take longer to check the second sentence than the first. To check that a chicken is a bird, you just need to look up the definition of "chicken." To check that a chicken is an animal, you also need to look up the definition of "bird." The researchers showed that the opposite

David E. Rumelhart

BIOGRAPHY

David Rumelhart remains an influential figure in many fields of psychological research. In the 1970s he worked on the interpretation of stories, publishing a major work in 1975 on the "grammar" of storytelling. He also wrote *Explorations in Cognition* in collaboration with Don Norman in 1975. This work had a great influence on many future developments in cognitive psychology.

In 1981 Rumelhart worked with James McClelland on a theory about how we read words. In 1985 they wrote about how the brain might learn and store information about categories. Both works are modern classics of cognitive psychology, and their influence is still felt today.

In 1986 Rumelhart and McClelland joined about a dozen other scientists to write a massive, two-volume book called *Parallel Distributed Processing: Explorations in the Microstructure of Cognition*. This book is often credited with taking mainstream psychology in a completely new direction. This new field became known as connectionism (see box p. 81). In the 1980s and 1990s connectionism took cognitive psychology by storm. Many commentators consider it to be one of the most important developments in the history of psychology.

As well as being part of a team responsible for bringing connectionism into mainstream psychology, Rumelhart was an author of one of its best-known articles, "Learning Internal Representations by Error Propagation." This paper contained a mathematical equation called the back-prop algorithm without which many of today's connectionist theories would not exist.

In 1998 David Rumelhart developed a progressive neurodegenerative illness called Pick's disease. In 2000 an Internet multimillionaire called Robert Glushko set up a prize in his honor. Glushko had been one of Rumelhart's graduate students. Although he did not stay in research, he never forgot David Rumelhart. The David E. Rumelhart prize is a cash prize of $100,000 that is awarded annually to the team or person making the most important contribution to the study of cognition.

was true. People took longer to agree that a chicken was a bird than they took to agree that a chicken was an animal. Why should this happen?

> *"The essence of memory organization is classification."*
> *—Bill Estes, 1994*

Remember how Eleanor Rosch showed that some category members were more typical than others? According to her research, a robin is a typical bird, but a chicken is not. When asked to think about birds, people do not often think of chickens. As a result, checking the sentence "a chicken is a bird" takes longer.

Now look again at the second sentence, "A chicken is an animal." When asked to think about animals, chickens sometimes spring to mind. It therefore takes less time to check and agree with the sentence "A chicken is an animal." The same sort of argument can be applied to the original results of Collins and Quillian. When you think about canaries, singing is probably one of the first things to come to mind. Having feathers is also part of being a canary, but that is probably not the first thing you think of. People are quicker to confirm that a canary sings than to confirm that it has feathers because singing is a more "typical" characteristic of canaries than are feathers.

Mental dictionaries

We are not sure how the brain stores information. One popular idea is that the brain's dictionary is pretty disorganized. Our mental dictionaries do not contain a long, neat list of definitions. Instead, our knowledge is held within a mass of connections between small chunks of information. Psychologists call these chunks features. Some of the features of being a dog might be "furry," "four-legged," and "having a wagging tail." We learn things like this about dogs when we are young. Our brains store this

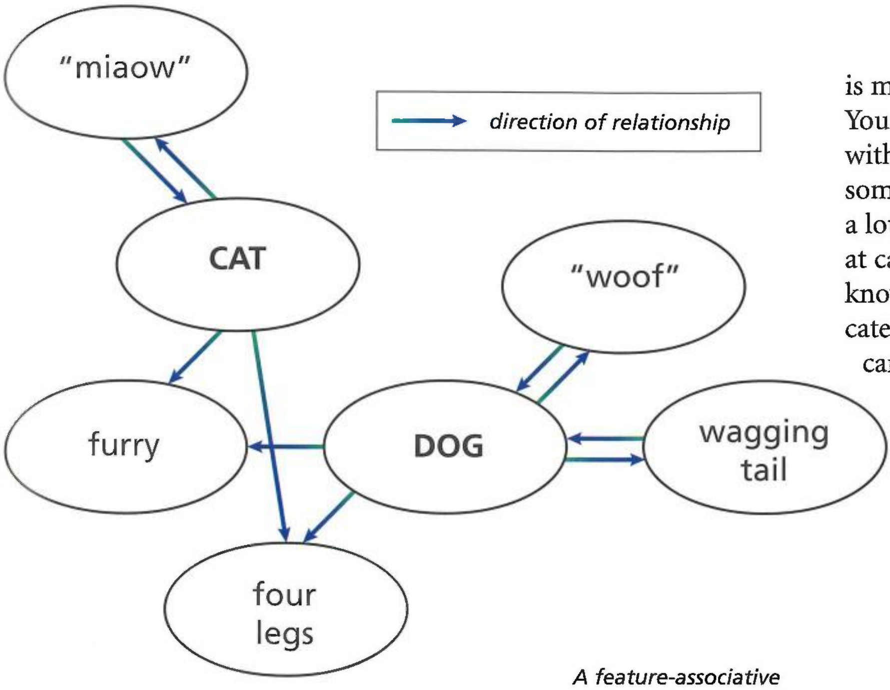

A feature-associative network (above). "Dog" and "cat" are both linked to features common to each, such as "four legs," and to unique characteristics, such as "miaow."

information by forming links between features like "wagging tail" and labels like "dog." The diagram above shows part of our mental dictionary. This sort of association is described by psychologists as a feature-associative network.

How do you "read" this sort of mental dictionary? The easiest explanation is to think of the circles in the diagram as lights. If you want to know whether a dog is furry, you light up "dog." There is a link from dog to furry, so "furry" also lights up. You have your answer—dogs are furry.

Another popular idea about our mental dictionary is that it is filled with examples. According to this theory, your dictionary entry for "dog" is just a collection of particular dogs you have met. It might include a description of your pet dog, Lassie, your neighbor's dog, and the guard dog you once saw at a factory. Your dictionary entry for "cat" is similar. It might contain a description of your grandmother's cat, a friend's cat, and a cat you have seen on a TV commercial.

Imagine you are walking down the street and see a four-legged animal coming toward you. Is it a cat or a dog? Quickly you compare the animal in front of you to the cats in your dictionary and the dogs in your dictionary. Overall it is more like the cats than the dogs. You decide that it is a cat. The problem with this idea is that each time you see something, you have to cross-refer with a lot of examples. You cannot just look at cat examples because you do not yet know it is a cat. You need to look at all categories, comparing the object to dogs, cars, cucumbers, refrigerators, and so on. And yet we can decide whether an object is a cat in a fraction of a second. If the brain had to make all these comparisons one at a time, this decision would take

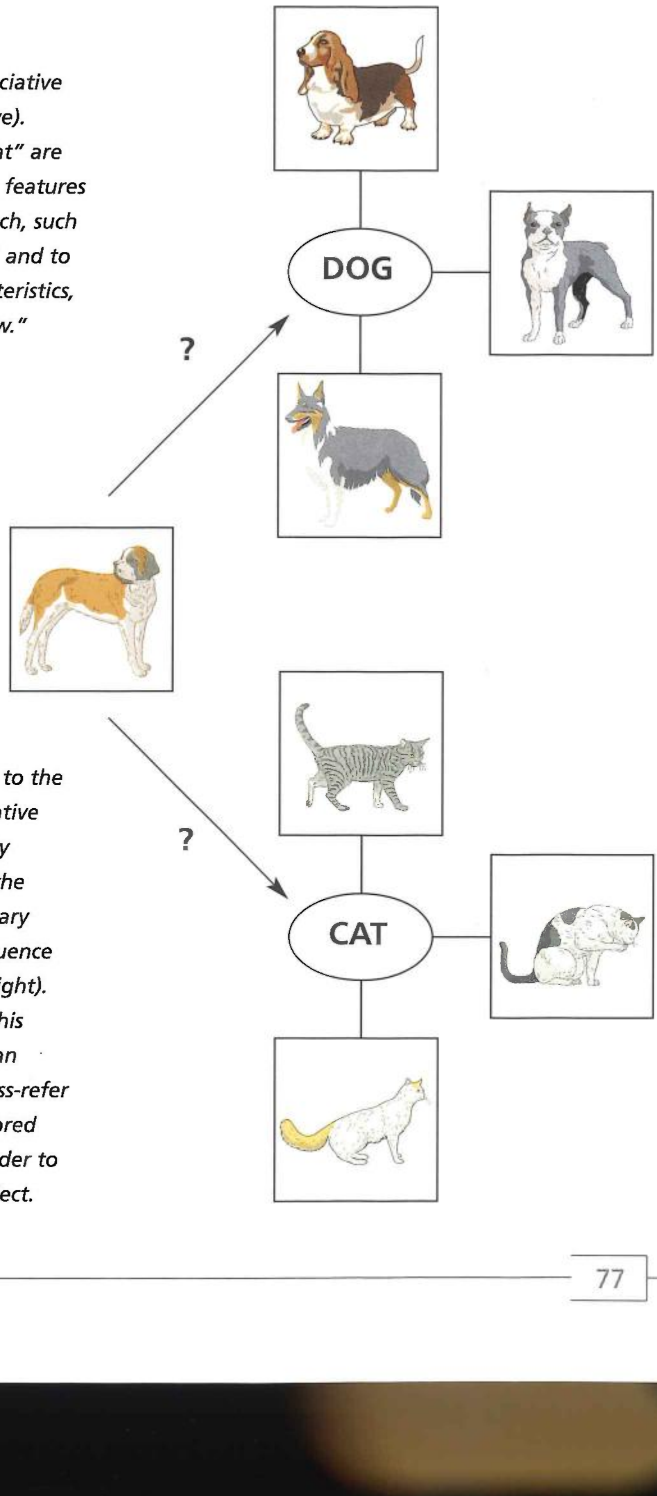

An alternative to the feature-associative network theory suggests that the mental dictionary contains a sequence of examples (right). According to this idea, the human brain must cross-refer with all the stored examples in order to identify an object.

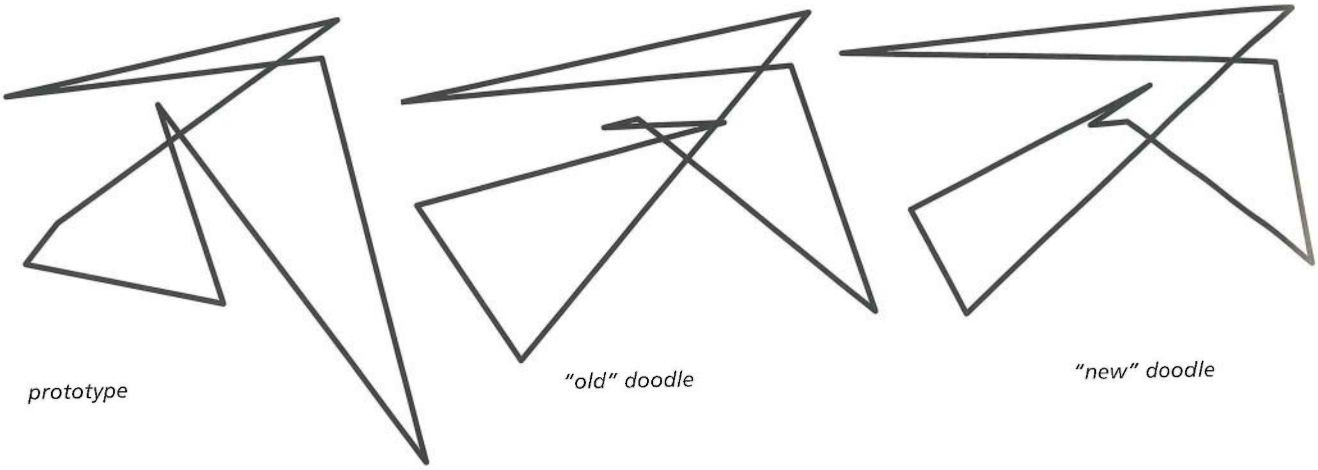

Homa and his colleagues designed different categories of doodles, each based on a different prototype. Subjects learned to associate "old" doodles with the correct prototype categories. Later, the subjects were presented with previously unseen "new" doodles. The subjects were good at categorizing the new doodles, but not as good as they were at categorizing the old ones. That is because the old doodles had been incorporated into the mental dictionaries of the subjects, while the new ones had not.

much longer. We know that the brain is very good at doing a lot of things at the same time. If the comparisons can be made contiguously, it is possible that our mental dictionaries are indeed simply a collection of examples. Research in this field is currently focused on deciding which of these ideas is correct.

Writing a mental dictionary

Dictionaries do not just come into existence; they have to be written. The same is true for our mental dictionaries. People are not born with a complete set of information about the objects around them, so they have to learn as they go along. We have already looked at how the information that is present in the mental dictionary is organized, but how does this information get there in the first place?

One way to study this is to teach adults new categories. To make sure the categories are new for everyone, psychologists often use made-up ones. Made-up categories allow us to answer questions that are difficult to answer with real categories. A 1981 experiment by Donald Homa, Sharon Sterling, and Lawrence Trepel provides a good example of this. The researchers invented some categories of doodles. Making each doodle category involved two steps. First, they created a prototype doodle. The prototype was the most typical member of a category.

Second, they created other members of the category by moving the points of the doodle around a little. They were also part of the prototype doodle's category, but they were less typical than the prototype. The psychologists created three different doodle categories in this way. Homa and his colleagues took some of the doodles they had made and put them to one side. They showed the rest of the doodles to college students and taught them which category each came from. These doodles were termed "old" doodles.

> *"To categorize is to . . . group the objects and events and people around us into classes."*
> *—Jerome S. Bruner, 1956*

When the students had mastered this, the psychologists took the doodles they had put to one side and asked the students to decide which category each of these "new" doodles had come from. The college students were fairly good at doing this, but not as good as they were with the old doodles.

The students found the old doodles easier to deal with because they had information about them in their mental dictionaries. They had not seen the new doodles before, and so they had not been entered into their dictionaries. Many psychologists consider results like these to be good evidence that mental dictionaries are collections of specific examples. On the other hand, some psychologists argue that these results can be explained by feature-associative networks. The true answer remains unknown, but research in this area continues to progress rapidly.

SCRIPTS AND THEMES

A dictionary will tell you what eggs and flour are, but it will not tell you how to bake a cake. To find that out, you need to look in a cookbook. Cookbooks are just one example of the wide range of instruction manuals that people rely on. Home maintenance books and car repair manuals are two other common examples. Instruction manuals tell us, step by step, what we have to do to complete a task. When we are familiar with a task, we do not need to use an instruction manual—we can rely on our memory to tell us what to do. Few people, for example, need an instruction manual to tell them how to get dressed each morning.

Does the brain store information about everyday events in the same way that an instruction manual would? Roger Schank and Robert Abelson suggested that people use mental scripts for occasions such as going to a restaurant. Scripts are a list of the typical events that occur in a specific situation. For example, the script for a trip to a restaurant might be:

> *Enter restaurant. Go to table. Sit down. Get menu. Look at menu. Choose food. Give order. Wait and talk. Waiter delivers food. Eat and talk. Receive check. Pay check. Leave.*

Obviously, not all restaurants are like this. In some restaurants, for example, you are asked to pay before you eat. A script does

Schank and Abelson suggested that people use mental scripts when dealing with everyday situations, such as going to a restaurant. Scripts allow people to have a good idea of what to expect from the experience.

not tell you for sure what will happen, but it does tell you what is likely to happen most of the time.

Scripts also help us communicate with other people more efficiently. If you ask someone what they did last night, and they reply, "I went to a restaurant," your restaurant script will provide you with an idea of the series of events the person experienced. Schank and Abelson thought that scripts were for specific events. For example, if you have visited a doctor, then you might have a "visit doctor" script and through that know roughly what to expect. If you have never visited a dentist, then you will not have a "visit dentist" script and will not know what to expect. The "visit doctor" script cannot help because you are not visiting a doctor.

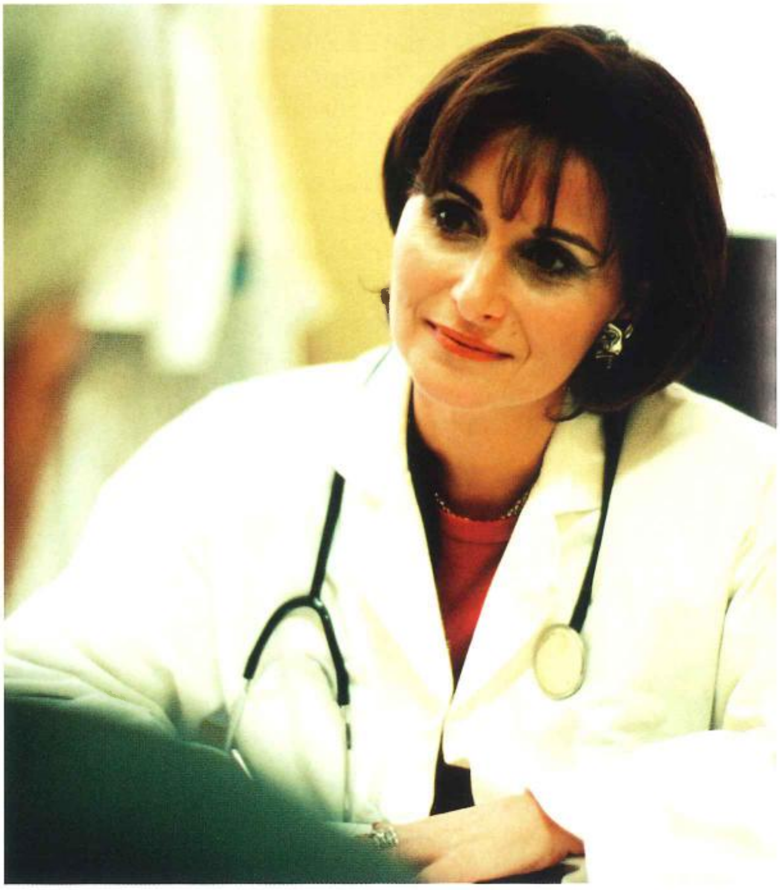

However, it seems likely that our expectations about events are a little broader than this. When we visit any healthcare professional, there are a number of steps that we can expect. They include making an appointment, describing the problem, and receiving treatment. In some situations these steps also include writing a check. If we have visited a doctor, we can predict some of the things that will happen on a trip to a dentist's office, even if we have never visited one ourselves. People certainly seem to have a great deal of shared knowledge about some events, such as going to a restaurant.

Someone visiting a dentist's office for the first time (below left) cannot rely on a script relating to a trip to the doctor (above). The specific details of a visit vary from one profession to the next. However, visits to healthcare professionals involve a number of steps that are usually common to all, such as making an appointment and describing symptoms. These steps can be incorporated into a general "healthcare" script that helps the patient anticipate the general sequence of events at such visits.

EXPERIMENT

Remembering Stories

People have the ability to listen to stories, remember them, and then tell them to others. Long ago this was the only way we ever experienced stories. They would be told by one generation and learned by the next. These days we rely much more on books than on our memories. Most of us, though, know some stories. They may be tales we were read as a child or the plots of Hollywood blockbusters. Although it is not necessarily true that everyone has a novel within them, most of us have at least one story to tell.

Is remembering a story like writing a book in our heads? If it were, it would not matter whether the story made sense or not; we could still write it in our mental story book and read it back later. Read the story below to a friend, then ask your friend to recall the story without referring back to the text.

"If the balloons popped, the sound wouldn't be able to carry far, since everything would be too far away from the correct floor. A closed window would also prevent the sound from carrying, since most buildings tend to be well insulated. Since the whole operation depends on a steady flow of electricity, a break in the middle of the wire would also cause problems. Of course, the fellow could shout, but the human voice is not loud enough to carry that far. An additional problem is that a string could break on the instrument. There would be no accompaniment to the message. It is clear that the best solution would involve less distance. Then there would be fewer potential problems. With face-to-face contact the least number of things could go wrong."

Trying to remember this story without looking back is pretty difficult. Researchers John Bransford and Marcia Johnson found that people generally only remembered about three or four things from this story. The story does not make much sense, and so it is hard to remember. Now show your friend the illustration on the left, and try reading the story again. That time it should make a lot more sense. Bransford and Johnson found that people who saw the illustration first remembered about eight different things about the story. This is about twice as many as people who had not seen the picture. This shows that our memory for stories depends heavily on our ability to understand them.

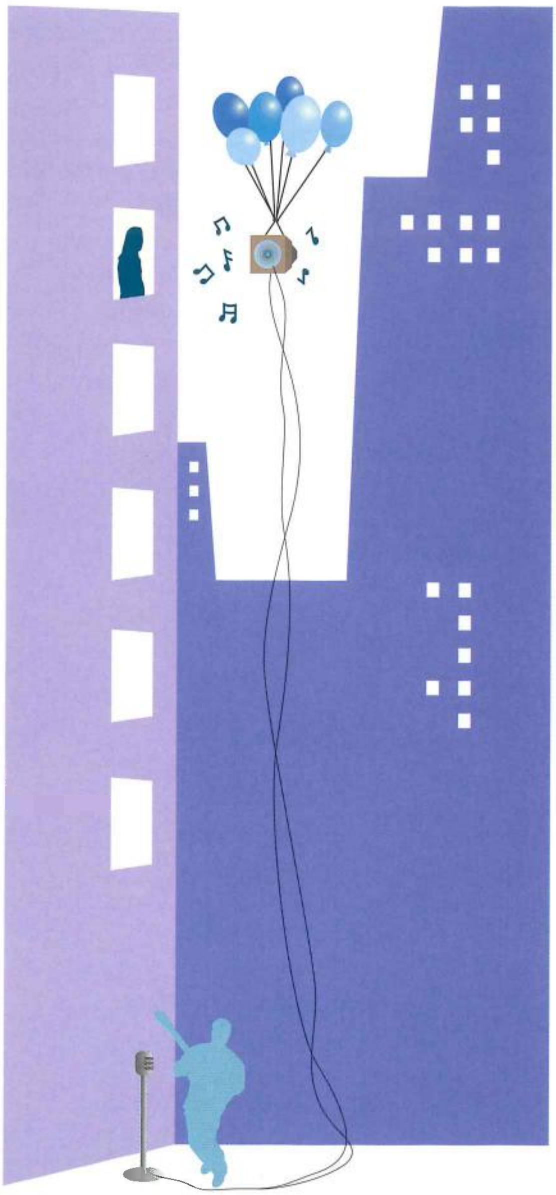

This diagram adds context to the odd little story above. When read in conjunction with the picture, more details of the tale can be remembered.

Psychologists Gordon Bower, John Black, and Terrance Turner asked people to list about 20 things that usually happened when they went to a restaurant. Almost three-quarters of the people included five key events. They were: looking at the menu, ordering, eating, paying the check, and leaving. Almost

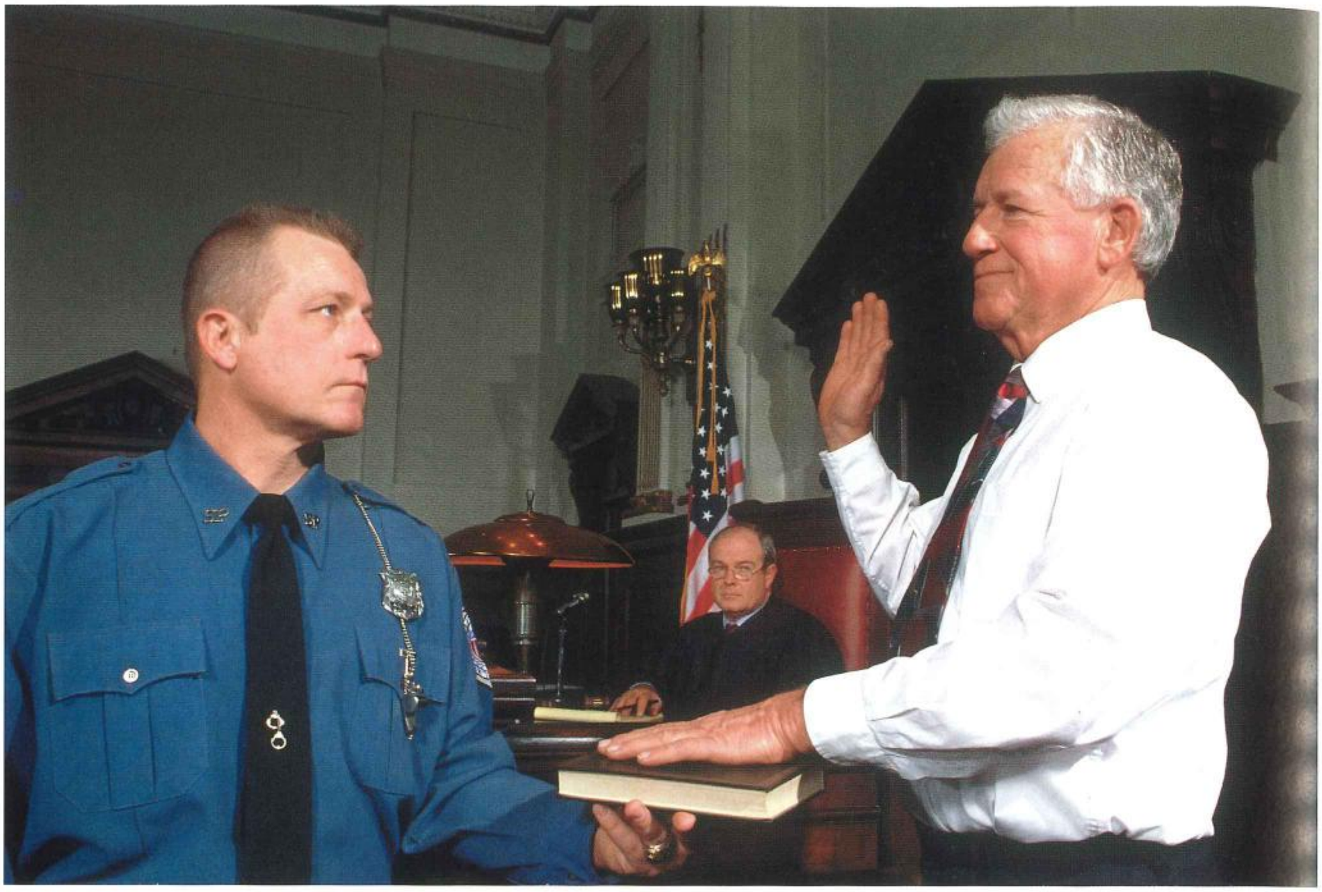

half of the people asked included seven further events. They were: ordering drinks, discussing the menu, talking, eating a salad or soup, ordering dessert, eating dessert, and leaving a tip.

Peoples' memories of specific events are affected by their mental scripts. Bower's team gave people some stories to read. They were based on scripts like "going to a restaurant," but the psychologists had jumbled the order of some of the events. So, for example, one story might involve someone going to a restaurant, paying the bill, and then sitting down and ordering some food. Next, they eat their food and then look at the menu. Finally, they leave.

When people were asked to remember the stories, they often described what normally happens in a restaurant rather than what actually happened in the story. So the story would typically be remembered as someone going to a restaurant, sitting down, looking at the menu, ordering their food, paying the check, and then leaving. Mental scripts help us anticipate what happens in certain situations; they may also color our recall of what happens in reality.

Can testimony be fully trusted? Holst and Pezdek showed that scripts can affect memories of certain events. They gave subjects facts about a hypothetical robbery. A week later they asked the subjects for their version of events. Rather than detailing the facts as presented, the subjects' stories often fitted around a generic "robbery" script. This finding has wide-ranging legal implications, providing a psychological framework for false identification and recall on the witness stand.

> *"We read for the gist, and very quickly forget the details"*
> *—Trevor Harley, 1995*

Valerie Holst and Kathy Pezdek showed that people who witness crimes can experience the same problem. They showed that when people try to remember what really happened in a crime they witnessed, they sometimes refer to a mental script and recall what typically happens instead. In another experiment Gordon Bower and colleagues gave people several different stories to read. Later, they gave the same people a different set of stories. Some stories were exactly the same as before; others were new. The researchers asked the volunteer readers to decide which stories were new. The volunteers were generally good at this, but a certain type of new story caused problems.

If a story was new but described an event similar to an old story, the subjects sometimes thought they had read it before. They became confused by stories that had the same script and were also nonplussed by stories that had different but related scripts. For example, one of the original stories involved a trip to the dentist. Later, the subjects read a new story about a visit to a doctor. The volunteers would often believe that they had read the story before. They had not,

Stills from the movies Romeo and Juliet *(1936), with Leslie Howard and Moira Shearer (below), and* West Side Story *(1961), starring Natalie Wood and Richard Beymer (right). Although set in different eras and places, they share many themes.*

but they had read a story with a similar theme. This suggests people remember stories in terms of general themes. These organizing themes are less tied to a specific situation than a script and can be very generalized. For example, most people would say that the 20th-century *West Side Story* is similar to William Shakespeare's *Romeo and Juliet*, despite the fact that they are set in different countries and in different

BIOGRAPHY

FREDERICK BARTLETT

Sir Frederick Bartlett (1886–1969) was one of the founding fathers of modern psychology. The Laboratory of Experimental Psychology at Cambridge University, England, had been open for about a year when Bartlett joined as an assistant in 1914. By 1931 Bartlett was a professor and head of the laboratory.

By the close of World War II (1939–1945) Bartlett's former students were running the majority of university psychology departments across the United Kingdom. Bartlett fought hard to get psychology recognized as a scientific subject. His success had a profound influence on 20th-century psychology.

Historians of science believe that no living psychologist will be able to have so great an influence on psychology as Bartlett. In the early 20th century psychology was still a new subject. Bartlett was at his most productive during a time of transition; psychology moved from being a relatively obscure discipline to a major scientific field.

Frederick Bartlett is probably best remembered for his ideas about schemata. A schema is a piece of general information stored in the brain that helps us understand the world around us. For example, the restaurant script (*see* p. 79) is a schema. Bartlett argued that people from different cultures can have different story schemata. A story schema provides information about the sort of general events that we expect to find in stories and the order in which we expect them to happen.

In perhaps his most famous experiment Bartlett translated a Native American folktale called "War of the Ghosts" into English. He then asked some of his students to memorize it. As shown by this excerpt, Native American folktales are very different from the stories that were familiar to 1930s Cambridge students:

> *"He told it all and then became quiet. When the sun rose, he fell down. Something black came out of his mouth. His face became contorted . . . . He was dead."*

When the students later had to retell the stories, they altered them to sound much more like English tales. Bartlett believed that remembering was not like writing things in a book and reading them out again. He saw memory as a process of reconstruction, and our expectations about how things normally are affect our memories of what happened. In 1932 he published his theories on memory in a book called *Remembering: A Study of Experimental and Social Psychology*.

Bartlett's ideas about schemata were largely ignored in the United States at the time but have subsequently been included in many textbooks. In 1952 he retired from university research, but he continued to theorize and write, and in 1958 Bartlett published another important book, *Thinking*.

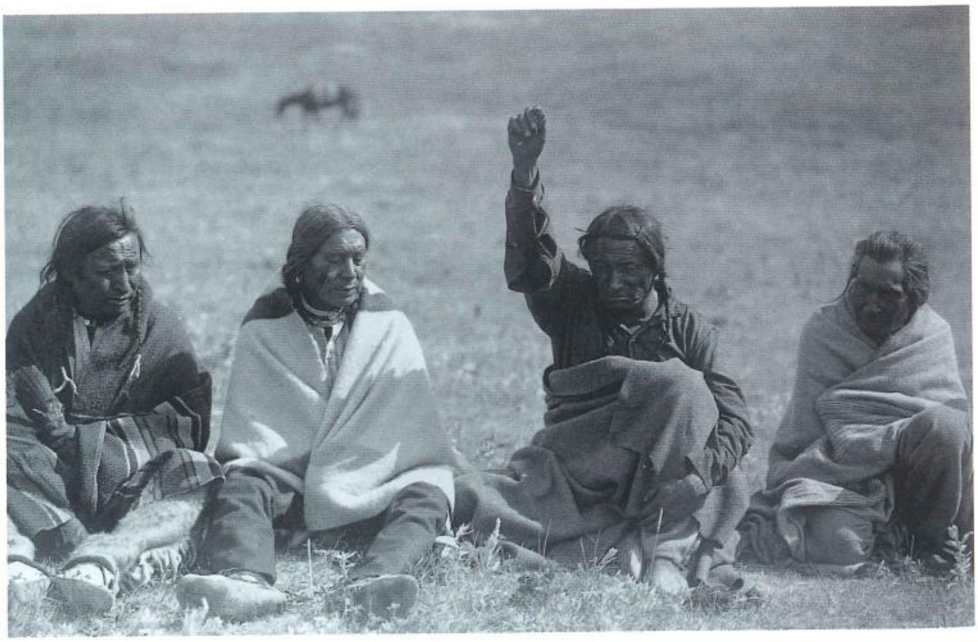

Bartlett translated a Native American folktale and asked students to learn and recite it. The students altered details of the story to resemble their own cultural context more closely than the original version.

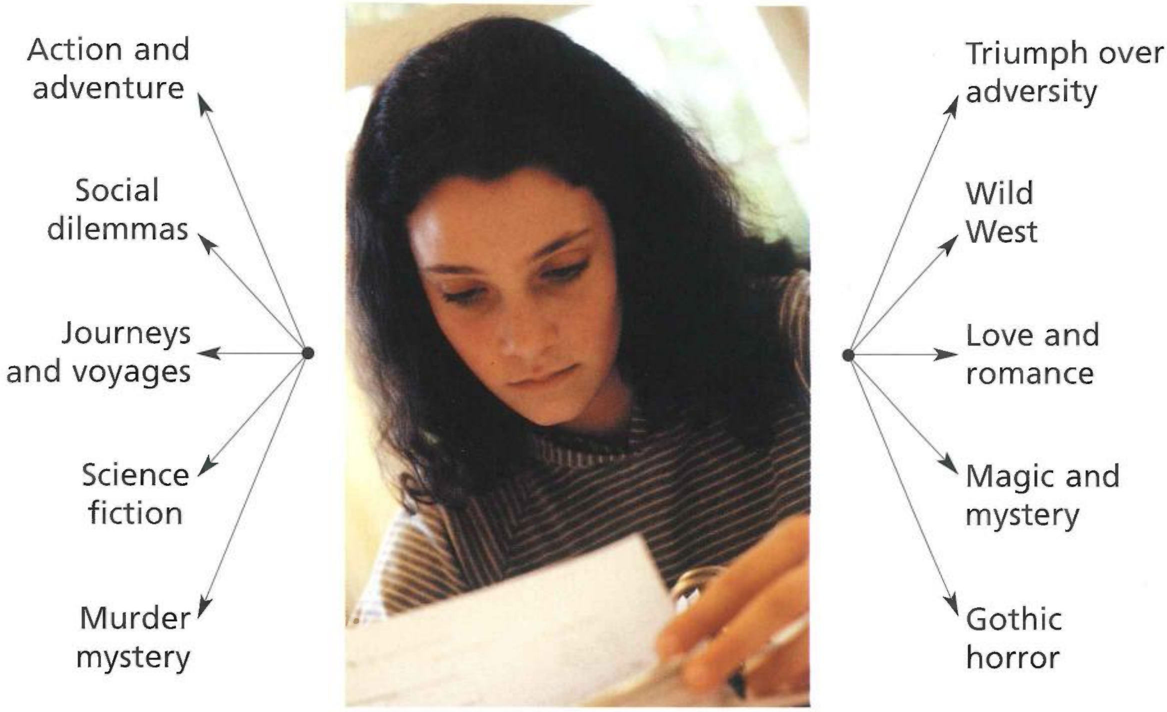

centuries. *West Side Story* is a musical, while *Romeo and Juliet,* which was written around 1595, is a play. (In fact, *West Side Story* was based on *Romeo and Juliet.*)

Roger Schank has argued that the stories share the common theme of "mutual goal pursuit against outside opposition." Romeo and Juliet love each other and so want to be together. Their togetherness is the mutual goal that they pursue. Their parents oppose the relationship, and so Romeo and Juliet pursue this goal against outside opposition. The theme of *West Side Story* is exactly the same.

Psychologist Colleen Seifert and colleagues showed people a series of stories that differed in many details, but all had the same general theme. Once the subjects had read the stories, the researchers asked them to write out similar stories. Most of the subjects wrote stories that contained different details but had the same general theme. Seifert's team also gave people a set of stories to sort into different piles. The subjects were allowed to do this in any way they wanted, but most sorted them into common themes.

Seifert and her team conducted an experiment that showed that people categorize what they read into themes. People use such organizing themes to categorize much of the information that enters the mind.

> ***"In order to observe one must learn how to compare."***
> ***—Bertolt Brecht, 1949***

INFORMATION AND THE BRAIN

For most of the 20th century psychologists relied on metaphors to explain how the brain stores information. In literature metaphors are words that describe people and objects by likening them to things they are not. The phrase "the city is a jungle" is a metaphor. Cities are, of course, not jungles. A "born-again" Christian has not literally been born for a second time, and so on.

Psychologists have used objects constructed by people as metaphors for the minds that made them—the mind has been equated with a photograph album, a dictionary, and the script of a play. In the end, though, the mind is none of these things. It is simply the mind.

Many psychologists are now coming to the conclusion that the mind should be examined for what it is. One step in this direction is what some people have described as the "connectionist revolution." Connectionism is not so much a specific theory as a way of thinking about psychology (*see* box p. 76). Connectionists believe that theories about the mind should take into account how the brain really works. The brain does not contain dictionaries, maps, pictures, or instruction manuals. It contains nerve cells, called neurons, that communicate through electrical signals—the nerve impulse. We know quite a lot about how neurons interact and a little about how they store information. For example, we know that compared to modern computers, neurons work very, very slowly.

We also know that neurons work in a "massively parallel" way. When you look at a picture, some neurons detect horizontal lines, some detect vertical lines, and others look for diagonal lines. They all do this at the same time, as well as carrying out a multitude of other functions. Connectionist theory

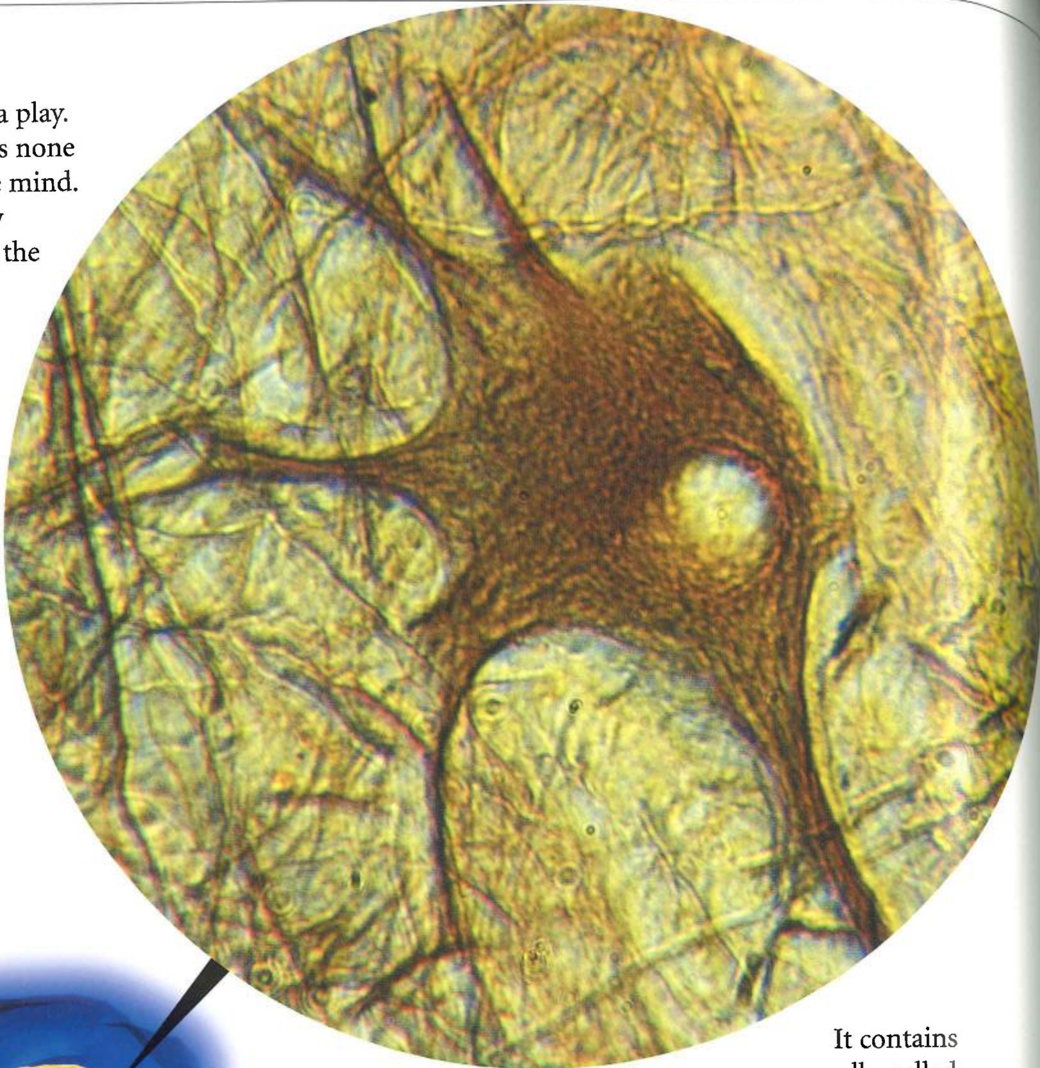

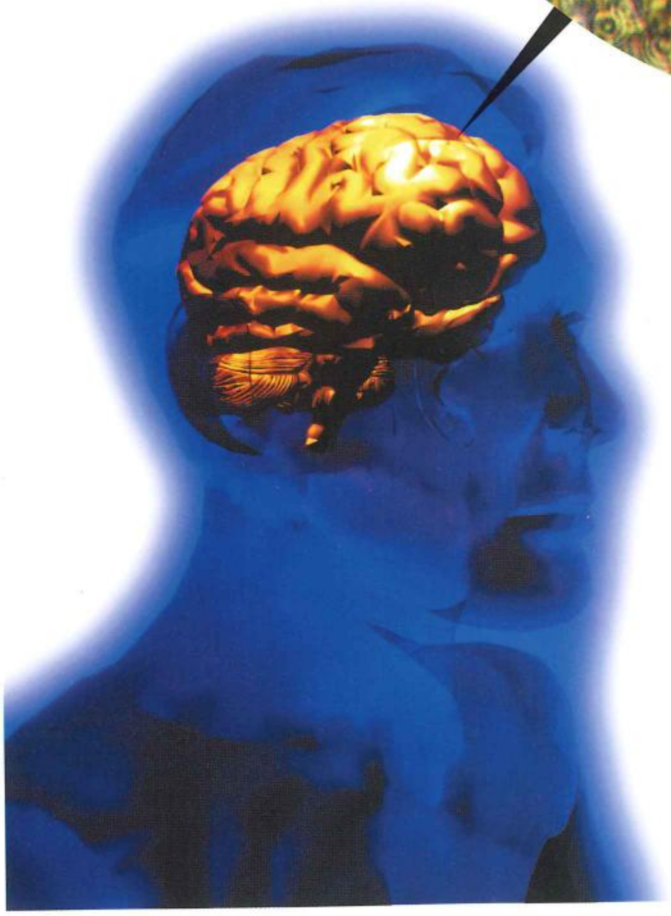

The brain (left) contains a network of nerve cells, or neurons (above), of staggering complexity. Scientists are gradually uncovering the electrical and chemical interactions that allow these cells to communicate, and psychologists are using the results to revise many theories of the mind.

FOCUS ON

Cognitive Science

Despite recent technological advances, a small child can perform feats of mental agility that remain far beyond the capabilities of the most powerful computers.

Throughout the 20th century psychologists conducted experiments to find out how the mind represents information. Philosophers, neuroscientists, and computer programmers have also explored the same issues. Some psychologists attempted to distance themselves from philosophy, but philosophers like Ludwi Wittgenstein (*see* p. 70) showed that philosophers have an importantpart to play in the search for an understanding of the mind.

The 20th century saw great advances in our understanding of the brain. We know much more about its composition and its connections, and neuroscientists continue to make important discoveries every year. The more we know about the brain, the better we will understand the mind. In the second half of the 20th century digital computers developed from clumsy devices that filled several rooms to very powerful tools that fit on a desk, a lap, or in a pocket. As computers became more powerful, people thought of more and more things they could do with them. However, psychologists and programmers alike soon realized that a two-year-old or a rat could do things that were beyond the most powerful of computers. If we better understood how people and other animals did things, we could build truly remarkable computers that would make today's machines obsolete. With the arrival of connectionist ideas theories of the mind began to use more equations and fewer words, involving mathematicians and physicists as well as psychologists. Over the last 20 years it has become obvious that psychologists, philosophers, neuroscientists, computer scientists, and mathematicians were all working on the same sorts of questions; but because they were from different disciplines, they seldom communicated. To solve this problem, universities began to set up a new subject called cognitive science. Cognitive science departments include people with different backgrounds and skills who are united by their desire to answer the sort of questions raised by this chapter. Cognitive science is now a popular degree course at a number of universities around the world.

incorporates biological features of the brain. The theories often include mathematical principles to describe how neurons communicate with and learn from each other. Some of them include the concept of drawing a relationship between the working of the brain and a computer.

There are connectionist theories associated with much of the research discussed in this chapter.

CONNECTIONS

- Attention and Information Processing: pp. 24–43
- Storing Information: pp. 88–113
- Neuropsychology: Volume 1, pp. 90–95
- Cognitive Psychology: Volume 1, pp. 104–117
- Biology of the Brain: Volume 2, pp. 20–39
- The Mind: Volume 2, pp. 40–61
- Perception: Volume 2, pp. 62–87

THINKING and KNOWING

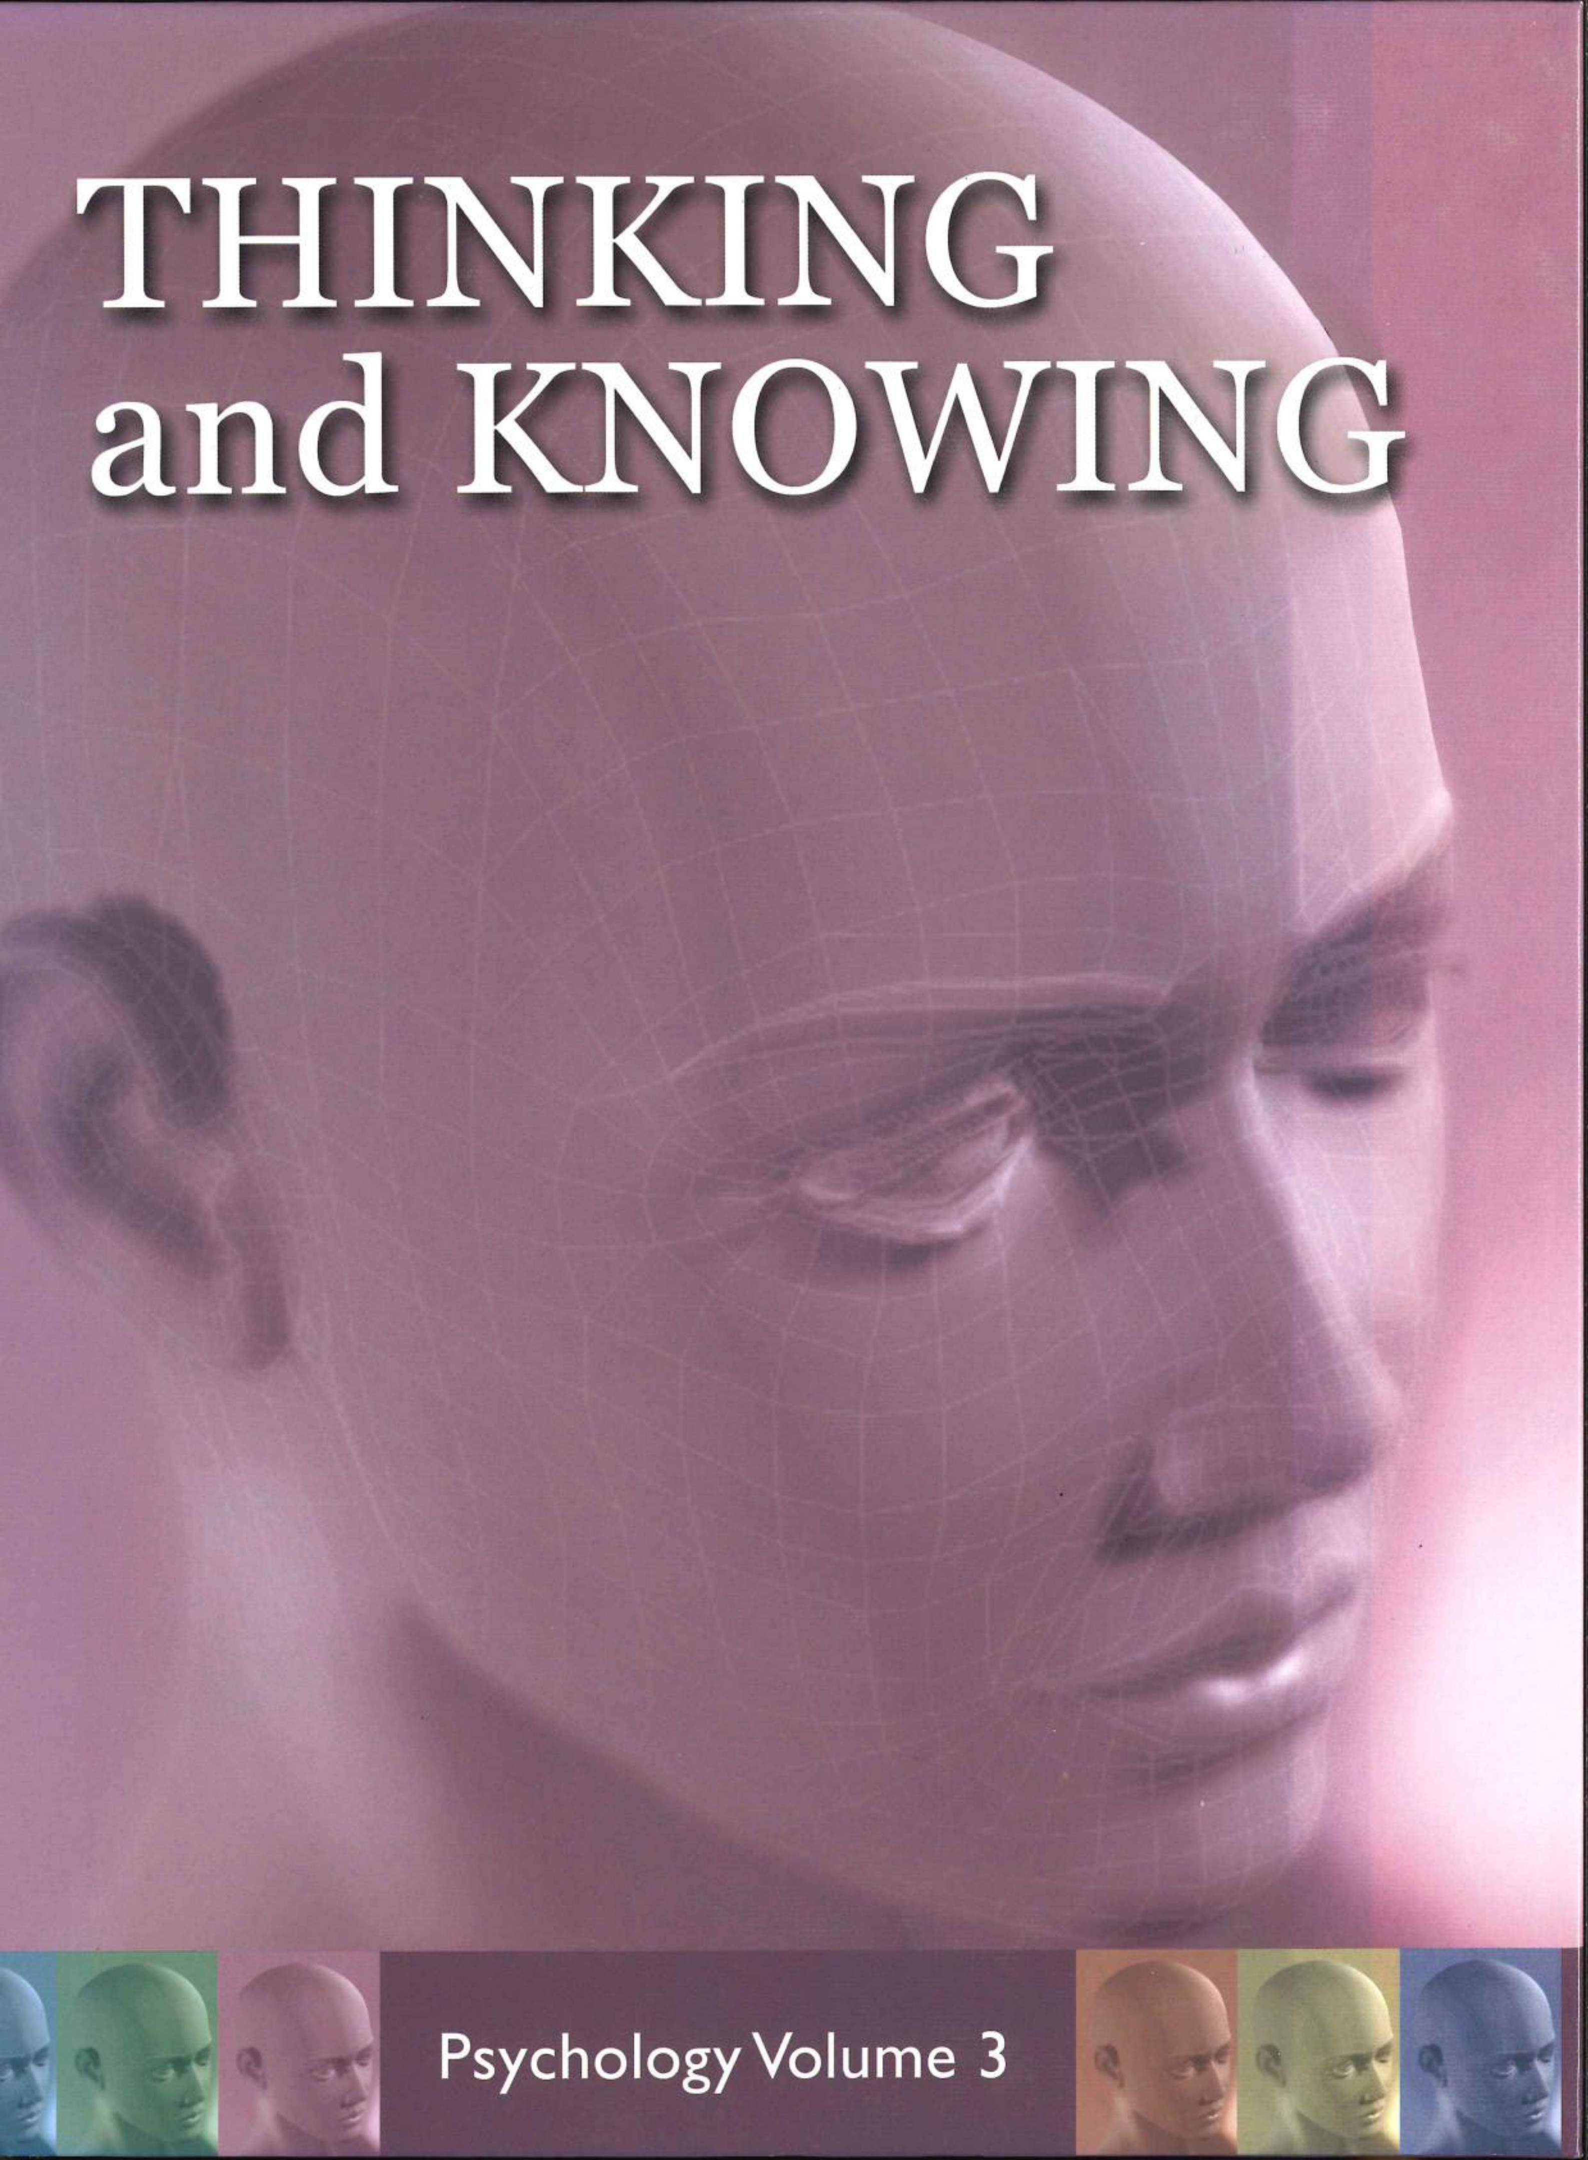

Psychology Volume 3

GROLIER
EDUCATIONAL
Sherman Turnpike, Danbury,
Connecticut 06816

SET ISBN 0-7172-5662-6
VOL ISBN 0-7172-5658-8