Language and Thought

12

Learning Objectives

By the end of this chapter you should appreciate that:

- language can be understood in terms of its grammar (syntax), its meaning (semantics) and its significance (pragmatic interpretation);

- a major issue is which language systems are modular and which interact with each other, and when;

- language consists not of separate sentences, but of connected discourse;

- normal reading depends upon eye movements, which can therefore be used to measure difficulties in reading texts;

- aphasia refers to language disorders that may affect the comprehension and production of language, and whether the language that is produced is grammatical or not;

- dyslexia is a disorder of reading and may be apparent during individual development or acquired later through brain injury;

- problem solving is characterized by the development of a mental representation of the problem, while logical reasoning is concerned with how people draw necessary conclusions from particular states of affairs;

- people often rely on simple heuristics to solve probability problems, sometimes leading to the wrong conclusions;

- on the other hand, if the information is in the right format, reasoning can be very effective.

INTRODUCTION

This chapter is about two central activities of human life, and as such we are confronted by a vast array of studies, almost all fascinating, and often important.

This chapter takes an *information-processing approach*, which essentially asks by what processes we can accomplish the tasks of thinking and using language. The information-processing approach has been the most successful of approaches to understanding cognition. Other aspects of language and thinking research, such as cross-cultural and cross-linguistic comparisons, provide rich data about the nature of what language is, but it is the information processing approach that has to be applied in order to understand how language actually works. Another important aspect of language is how processing develops in children: this is dealt with in chapter 9.

Nowhere is the astonishing capacity to perform acts of inference revealed more clearly than in the study of language, and nowhere are the limitations of inference-making rendered more obvious than in the study of thinking.

> **information-processing approach**
> understanding how something works by finding out the kinds of information involved and the steps through which it goes in order to accomplish a task

LANGUAGE

Language gives us the capacity to let others know things and do things that would otherwise be impossible. It enables us to share knowledge and ideas, and to extend our spheres of influence beyond the immediate.

Spoken language is the most basic form, especially dialogue, but most cultures have also developed written language systems. Written language not only allows the ready dissemination of information within our own culture, but also enables us to keep in touch with cultures that are remote in both time and place.

The psychology of language is concerned with the organization and processing of both written and spoken language. It is a complex field, at the interface of pure psychology, linguistics and communication studies. And as we examine how language is processed, it will soon become clear just how complex and mysterious the process is. For instance, a colleague of mine recently mentioned that he was feeling 'low' because he had just received some severe criticisms of a paper he had written. Why did I know immediately what he meant? Why did I not think he was simply nearer to the ground?

SYNTAX, SEMANTICS AND PRAGMATICS

It is conventional to divide up issues of language under the headings syntax, semantics and pragmatics.

Syntax is the set of rules or principles that govern word order, and which words can be combined with which. The rules and principles have been determined by scholars but, in a sense, they reflect the way the brain analyses language. An example of a syntax rule, in English, is that a sentence consists of a noun phrase plus a verb phrase. This can be written as:

syntax rules that govern the admissible orderings of letters within words, and words within the sentences of a language

$$S \rightarrow NP + VP$$

So with the sentence 'John loves Mary', 'John' is the noun phrase (NP) and 'loves Mary' is the verb phrase (VP).

Other descriptive rules specify what is an NP and a VP. The details are quite complex, but a descriptive grammar is one that allows only those strings of words that people accept as sentences.

Psycholinguistics has been especially concerned with how people *parse* sentences – that is, how they break them down into their correct grammatical structures. This has to be done because, otherwise, it would be impossible to interpret a sentence at all. Consider the following:

The horse raced past the barn fell.

Is this an acceptable English sentence? What does it mean? In fact, it is a classic illustration of the problem of parsing. People normally find this a hard sentence to understand, because the parsing

mechanism treats 'The horse' as an NP and 'raced' as the main verb, so it then expects more information consistent with the noun phrase. But the sentence actually contains what is called a reduced relative clause. Here it is in its unreduced version:

The horse that was raced past the barn fell.

By missing out the words 'that was', the relative clause is reduced. So, in fact, the structure of the sentence is:

NP: The horse (that was) raced past the barn
VP: fell.

The difficulty in understanding such sentences is ascribed to an initial misinterpretation, and is called a 'garden path' (see Frazier, 1987).

A large amount of time and effort has gone into studying the human parsing mechanism because it is central to language comprehension and production. By misparsing the sentence above, there is a resultant failure in comprehension at all levels. Another well known example is the sentence, 'The old man the boats.' Most people find this sentence ultimately intelligible but find there is a disturbance of understanding, because the string 'The old man' is parsed as an NP, and not as an NP + V ('The old' being a shortened version of 'The old people', and 'man' being a verb). Unless the sentence is properly parsed, it is unintelligible.

Semantics concerns aspects of meaning. For instance, while 'Green rain sinks frail grannies' has good syntax, it is meaningless. The meaning of a sentence is somehow assembled from the meanings of the individual words that make up the sentence. Meaning at the sentence level is vital for comprehension, just like syntax. Compare the following:

semantics the meaning of words and how they combine to give the meanings of sentences

Harry cooked dinner with his wife last night.
Harry cooked dinner with a wok last night.

In the first, 'his wife' is a co-agent, accompanying Harry, whereas in the second, 'a wok' is an instrument for cooking. To assign the wrong role (meaning) to 'his wife' would make Harry look like a cannibal!

And *pragmatics* concerns what we do with language. At the level of sentence meaning, 'Can you pass the salt?' is a simple question, and should be interpreted as a question about competence. But when a child is asked this at the table and replies 'Yes', everyone knows this is a game. This is because there is a distinction between semantics, or sentence meaning, and pragmatics, which is sometimes called speaker meaning, and concerns the meaning of an utterance, not just a sentence.

pragmatics the significance given to a sentence or utterance by relating its semantics to everyday knowledge of situations

The fact that sentence meaning is not sufficient to guide an interpretation led to a theory of speech acts (Searle, 1969), which

treated utterances as actions on the part of a speaker, with the actions requiring their own interpretation.

The introduction of pragmatics is essential to any account of language processing, and is especially obvious in cases where semantics (or literal meaning) appear to fall short.

There are two principal classes of phenomena that obviously require more than literal meaning. One is indirect speech acts, like the salt example above. The other is metaphor and related phenomena. For instance, if I say 'Adolf Hitler was a butcher', I do not mean it literally. Similarly, if I say 'John is really blue (or low) today', I do not mean that he is covered in blue dye, or has shrunk in height. I mean that he is depressed.

We appear to process many metaphors so readily that it is difficult to see what the problem is, but the processing problem is huge: not only does the processor have to parse sentences, but she has to determine their significance too. The psychology of language understanding is about just these issues.

Finally, interpretation proceeds by linking language to our knowledge about people and situations. Consider the following:

A. John was hungry. He went to a restaurant and ordered some nine-inch nails.
B. John was really hungry. At the restaurant, he ate some Crepe Suzette, and then ordered steak, followed by Moules.
C. Harry put the wallpaper on the wall. Then he sat his full coffee cup on that.
D. Harry put the wallpaper on the table. Then he put his coffee cup on that.

- In case A, a problem is detected because nine-inch nails are not edible. This information has to be retrieved in order to make use of it. It implies access to an almost encyclopedic knowledge of what one can eat.
- In case B, the procedure for determining the order in which things are eaten is accessed. In this case, one would not normally consume a sweet dish (Crepe Suzette) before a savoury dish (Moules). Schank and Abelson (1977) suggested that we have mental scripts for stereotyped sequences, which are accessed under the right conditions, and as a result we can spot anomalies when they occur. Without such stereotyped knowledge, we would not have any knowledge of social norms.
- In case C, wallpaper being on a wall puts it in a vertical position, so you cannot put your cup of coffee on it. Detecting the problem suggests that we set up a mental representation of what putting wallpaper on a wall entails.
- Case D is not a problem at all. But it is almost identical in linguistic terms to C; it is just that 'on the table' is taken to mean flat on the table, so sentence D is judged not to be problematic.

UNDERSTANDING LANGUAGE

Comprehension of language requires the processor to use knowledge of the language (syntax), meaning (semantics), and our knowledge of the world (scripts) and inferences about the intentions of speakers (pragmatics).

The central questions for the study of the processing system are:

How and when are these sources of information called upon?
How is the architecture of the system organized?
Is syntactic analysis carried out first, and then meaning and interpretations ascribed later? Or are they all used at any point they might be needed?

There are too many studies in this area to present a full overview here. Instead we present just two sample problems (word-sense retrieval and nonliteral meaning) to indicate how the issues may be addressed experimentally.

Word sense retrieval

When reading or listening, it is important to retrieve word meaning, and that means retrieving the right sense of a word. This is an area where the role of background knowledge is important. For instance, in understanding 'John put his salary in the bank', it is necessary to select the appropriate sense of 'bank' – i.e. a place where financial transactions take place, not the side of a river. Context usually provides the solution to this problem, but the question is when during the sequence of processing? Is just one meaning of 'bank' selected at the outset, or are both meanings initially recruited, and then the right one selected later?

There are two main possibilities:

1. The *modular view* is that word meanings are stored in a way that is not context sensitive. When we encounter a string of letters that represents a word, we automatically look up and retrieve the meaning. If the string (such as 'bank') represents more than one word, then both meanings should be retrieved. The modular view is attractive because it keeps the mechanisms of looking up word meaning separate from context, and so is computationally simpler (see Fodor, 1983, for a discussion of this position).

   > **modular view** two processes are said to be modular when they occur independently of one another and do not interfere with one another

2. The *interactive view* suggests that word meaning information is connected to other processes of comprehension, so that which aspects of word meaning are active depends on context. This view is attractive because it implies a very adaptive organization of knowledge and word meaning, but at the cost of more computational complexity (e.g. see McClelland and Rumelhart, 1981; Morton, 1969).

   > **interactive view** two processes are interactive when the processing occurring in one of them depends on the processing occurring in the other

An important technique for finding the answer is priming (see Meyer & Schvaneveldt, 1971). When a word is read, it becomes easier to recognize words that are associated with it. So if you read the word 'nurse', you will then read the word 'doctor' more quickly than if you had just read an unrelated word, such as 'bread'. What will be primed after reading the word 'bank'? If

Research close-up 1

The Moses illusion and beyond

The research issue

One view of language processing is that the individual words of a sentence are read and their meaning discovered, and then these meanings are combined to produce the whole meaning of the sentence. So each word's meaning would be retrieved from memory as it is read.

But this is not necessarily the case. Answer the following question before reading on:

How many of each sort of animal did Moses put on the Ark?

Erickson and Mattson (1981) reported that many people simply respond 'Two' to this question, even though they 'know' that it was not Moses who put the animals on the Ark (it was, of course, Noah). This effect has been dubbed the 'Moses illusion'.

Barton and Sanford (1993) explored this effect further. They hypothesized that provided a word fits a context very well, then its meaning need not be 'fully' analysed, because it is easy to work out what it must mean from the utterance as a whole.

Design and procedure

Barton and Sanford had participants individually read and answer questions like this:

When an air crash occurs, where should the survivors be buried?

The question was asked as one of ten in a questionnaire about social customs that included questions like: 'After a death, who should officially be informed?' and 'At what age should people be permitted to hold a driver's licence?'

The data were the numbers of participants who spotted the fact that you simply do not bury survivors.

One group of participants got the air crash question and another group got a bicycle accident version of the same question:

When a bicycle accident occurs, where should the survivors be buried?

Barton and Sanford (1993) reasoned that people would have mental representations that included the concept of survivors for an air crash, but not for a bicycle accident, where deaths are less common. Consequently, the term 'survivors' would fit easily in the air crash context, and result in shallow analysis. In contrast, because the word does not fit so well in the bicycle accident context, it would be scutinized more closely, resulting in higher detections.

Results and implications

For the air crash scenario, detection rate was 33 per cent, while for the bicycle accident it was 80 per cent, confirming the hypothesis.

These results demonstrate that the extent to which the meaning of a word is processed is not all-or-none but is variable. It also depends on the fit of the word to the situation. If it fits well, subsequent analysis may be minimal; if it fits poorly, then the system analyses its meaning to a greater extent in order to achieve a fit. This enables anomalies to be detected with a higher probability if the word fits the situation relatively poorly.

More generally, psycholinguists are becoming increasingly interested in the extent to which the various processes under-lying sentence comprehension always occur as fully as previously thought, and how little work the system can get away with while sustaining comprehension.

Barton, S.B., & Sanford, A.J., 1993, 'A case-study of anomaly detection: Shallow semantic processing and cohesion estab-lishment', *Memory and Cognition*, 21, 477–87.

there is no biasing context, then target words relating to both senses should be primed, such as 'river' and 'money'.

Swinney (1979) presented participants with spoken passages like these:

   (a) Mary needed to buy some presents, so she went to the bank.
   (b) Mary found the river cold, so she swam to the bank.

Immediately after the presentation of the ambiguous word, he presented a single letter string on a screen. Participants had to decide whether the letter string was a word or not (a lexical decision). When the string was a word, it could either be related to the intended sense of the ambiguous word (e.g. 'money'), related to the other sense (e.g. 'mud'), or unrelated to either. The question was whether there would be a response time advantage

for the intended-sense associate alone, or whether there would also be an advantage for the other-sense associate of the word too.

It turned out that there was equal advantage (priming) for both senses. So context did not appear to affect initial sense selection. But if there was a delay of only 300 ms between hearing the ambiguous word and reading the letter string, the priming effect remained only with the intended (contextually cued) sense.

This work suggests that word meaning information is initially stored in a modular fashion, and its retrieval is uninfluenced by context. On the other hand, very shortly after a word has been processed, contextual cues inhibit the activation of word sense information that is inappropriate.

This one example represents a sample of work on the problem of modularity; research in this area remains very active (see Sanford, 1999, for a fuller review).

Nonliteral meaning

How do we understand sentences? One explanation is that we assign a literal meaning to them and then integrate this into the meaning of the discourse. But the literal meaning may not make any sense, especially if the sentence conveys an indirect speech act or a metaphor. For instance, if I say 'My job is a jail', I mean it restricts my freedom in a way that parallels being in jail. One prevalent view is that metaphors are first interpreted literally, then, if this fails, they are interpreted as nonliteral, or figurative (Searle, 1975, 1979).

As a series of processing operations, this may be formulated as follows (from Glucksberg & Keysar, 1990):

1. Derive a literal interpretation of the utterance.
2. Assess the interpretability of that interpretation against the context of that utterance.
3. If that literal meaning cannot be interpreted, then and only then derive an alternative nonliteral interpretation.

The sequence above suggests that in order to make an appropriate interpretation of a statement, we need to know whether it is meant to be literally true or not. But it also makes strong assumptions about the processes underlying comprehension that subsequent work has suggested may be incorrect.

The account has been examined for both indirect speech acts and metaphor comprehension. Gibbs (1979) showed that people take no longer to process indirect requests such as 'Must you open the window?' – meaning 'Don't open the window' – than to understand literal uses of the same expressions (in the present case, meaning 'Need you open the window?'). These data suggest that people do not need to obtain a literal meaning of an expression first in order to comprehend an indirect speech act. This runs against the traditional model (Glucksberg & Keysar, 1990).

Gibbs (1983) claimed, more strongly, that participants do not always derive a literal meaning at any point. To establish this would be another blow to the traditional model, since this model specifies that literal meanings are necessarily established. Gibbs had participants read stories that ended with critical sentences such as 'Can't you be friendly?' In different stories, the sentence was given a literal meaning ('Are you unable to be friendly?') or an indirect interpretation ('Please be friendly'). After reading a passage, participants had to decide whether a string of words was a grammatically correct sentence. Some of the strings were either the literal or the nonliteral interpretation of the critical sentence.

Gibbs predicted that the literal context would prime the literal interpretation, and the nonliteral context would prime the nonliteral interpretation. These results should be reflected in a priming effect on the subsequent sentence judgement task. In two experiments, the results confirmed these expectations. In particular, when the context biased the interpretation of the critical sentence towards a nonliteral interpretation, there was no priming of the literal interpretation.

These findings show that the applicability of the standard comprehension model (Glucksberg & Keysar, 1990) is at best limited, although it is worth noting that the comprehension of sentences in stories (such as have been used in most of the studies reported here) and actual interactions in dialogue are very different situations, so we must guard against simplistic conclusions. Nevertheless, work on indirect speech act comprehension reinforces the view that literal interpretation is not always necessary.

Similar findings have been obtained for metaphor comprehension. For example, Glucksberg, Gildea and Bookin (1982) asked participants to decide whether simple statements were literally true or false. For example, consider the statement 'Some desks are junkyards.' This is literally false, and so (according to the conventional model) the obvious metaphorical interpretation should not interfere with processing and the production of a 'no' response. Yet it does. A statement with an obvious figurative interpretation takes longer to reject as literally false than does a sentence with no obvious figurative meaning, such as 'Some desks are roads.' So, in the case of 'some desks are junkyards' it seems that the metaphorical meaning is computed automatically even though it is not needed, which indicates that testing for literal meaning cannot represent the previous, modular processing stage that the classic position would claim (see also Glucksberg and Keysar, 1990).

Our sample of work on the comprehension of metaphors shows how simple response time studies can be used to evaluate the sequence of language processing events. The conclusions suggest that the straightforward classical view that literal interpretation takes place first, and then nonliteral interpretation takes place later if needed, is wrong.

DISCOURSE, AND A RETURN TO UNDERSTANDING

Language consists of more than disconnected utterances. When sentences are put together to make a sensible message, the result is *discourse*. A substantial part of the psychology of language deals with discourse processing, especially when it concerns text. Many theories of discourse processing have been developed, for example by Gernsbacher (1990), Kintsch (1988), and Sanford and Garrod (1981).

discourse a set of sentences that bear a sensible relationship to one another and so form a message

The primary feature of discourse is that it is coherent – in other words, the individual sentences fit together in a meaningful way and do not contain any contradictions. Sometimes sentences are connected by explicit devices, called cohesion markers. Consider the following:

John fell off the cliff because he was walking too near the edge.

There are two cohesion markers here:

1. the connective 'because' indicates that the sentence 'John fell off the cliff' portrays the result of a cause – i.e. 'he was walking too near the edge'; and
2. the pronoun 'he' signals that some individual who is singular and male has been mentioned. The only thing that fits the bill is 'John', so we take it that it was 'John' who 'was walking near the edge'.

But the establishment of coherence does not always rely on cues such as these. For instance:

John was hit by a train. He had been walking down the track.

This is coherent because 'walking down the track' was the condition that enabled 'John' to be 'hit by a train'. But there is no explicit connector ('because'): the connection is inferred. Coherence establishment may sometimes make use of cues in the text, but always relies on some degree of inference.

As a final example, consider the following single sentence:

John lent Harry some money because he was hard up.

What is the referent of 'he'? Obviously it is 'Harry', not 'John'. Why? Because money is lent to people who are 'hard up', and this inference is automatically drawn and used to solve the reference problem.

These few examples show the complexity of the computational operations that underlie even the most mundane language processing at the discourse level, and they represent just a small sample of the issues.

Inferences vs. scenarios

Experimental work shows that it takes time to make inferences. Haviland and Clark (1974) asked people to read short texts made up of two sentences, and then measured the reading times for the second sentences. Compare the following pairs:

Inference version: Herb took the picnic supplies from the car.
                       The beer was warm.
Explicit control: Herb took the beer from the car.
                       The beer was warm.

The reading time for the second sentence was longer in the inference version, because participants had to infer that 'The beer' is part of the 'picnic supplies'. The text demands that an inference be made, which demands cognitive resources.

But sometimes knowledge may be automatically recovered and included in the mental representation of the sentence. For instance, given 'Harry drove to London', there may be a default representation of the fact that a car was used. Subsequent mention of a car would not be a problem, because its default is already in the representation resulting from the sentence. This is what Garrod and Sanford (1982; 1983) found to be the case. In a fuller theory, Sanford and Garrod (1981; 1998) argued that we automatically relate what is being said to background knowledge, and that background knowledge is organized in long-term memory about specific situations. They called these structures 'scenarios', and argued that the basic, most fundamental operation of understanding is to recognize the situation in which the message is set. So, because we are retrieving further situation information from memory, sentences can lead to representations that go beyond their content.

As one final example of a study that seems to support this view, Garnham (1979) required participants to try to remember sentences they had seen previously: e.g. 'The housewife cooked the chips.' He found that participants remembered this sentence better if they saw the cue 'fried' than if they saw the cue 'cooked', even though 'cooked' is actually part of the original sentence. According to the scenario theory, this is because cooking chips has been implicitly represented as a situation in which frying is taking place. Of course, another possibility is that the word 'fried' simply provided more information, in terms of a cue for remembering.

The ultimate questions

For discourse studies, the ultimate questions are just *which* inferences are made (i.e. what knowledge is recruited) and *when* during language processing. Some theorists believe that sometimes there might not be much inferential activity taking place during natural discourse (McKoon & Ratcliff, 1992), and that inferences and elaborations will only take place when the relevant background knowledge is highly available in memory. Sanford and Garrod (1998) take the view that it is the task of the writer or speaker to say things in such a way that a scenario can easily be found, because this is essential for good message-level interpretation.

Whatever they think about component processes, there would be few scientists who would disagree that understanding is based on bringing language input into co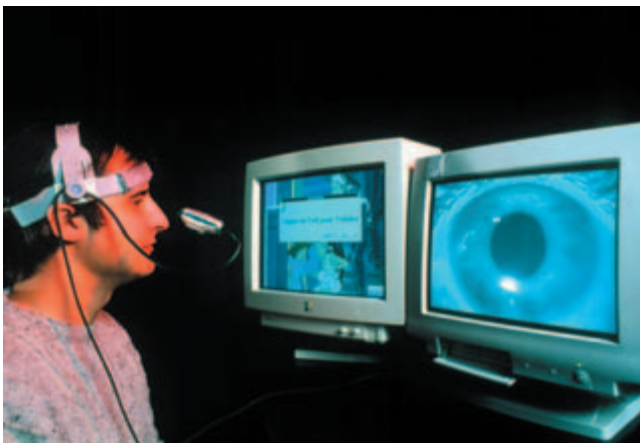ntact with world knowledge – the basic question being how this is done. Noam Chomsky has been at the forefront of international thought over the past several decades regarding the individual development and intergenerational heritability of language. The classic Chomskian sentence 'Curious green ideas sleep furiously' is not intelligible at the message level, simply because it is hard to relate to anything we know about. But 'The housewife cooked the chips' is intelligible because we can easily relate it to many things we know about.

TRACKING THE READING PROCESS

Reading is a complex process, which can be broken down into a variety of activities:

Figure 12.2

A sample showing eye-tracking of a single sentence.
Source: Martin (1998).

- fixating words with our eyes;
- processing words in accordance with syntax and semantics;
- representing meaning; and
- understanding the significance of what is read.

Until now, we have focused on the last three activities – how we come to understand language. Now we will take a look at the first point in the process.
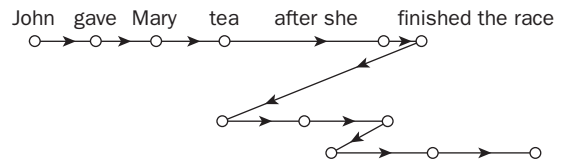
Some of the oldest studies of the reading process were concerned with the pattern of eye movements that occurs when we read text. Even today, many of our insights come from studies of eye-tracking.

Using modern eye-tracking equipment (see figure 12.1), it is possible to establish where the most sensitive part of the eye (the fovea) is fixating within a piece of text. Although we have the impression of a smooth process when we read, in fact the eye moves in jumps, called saccades, and then fixates, or remains stationary, upon successive pieces of text (see chapter 7). Figure 12.2 shows an eye-tracking record for a person reading a piece of text.

The dots are fixation points, and the lines are saccades. When the line moves back towards an earlier part of the sentence, this is a regression. Word information is only encoded when the eye is stationary, and then only about 15 letters can be encoded within a single fixation.

From the perspective of understanding, it is interesting to note that small words are not always fixated. So a word such as 'he' may only be fixated 30 per cent of the time. Content words, on the other hand, are nearly always fixated.

At one time it was thought that where the eyes fixated was simply a mechanical process, but now it is clear that eye movements are under the control of some of the complex processes underlying language understanding (Rayner & Polletsek, 1989). For instance, when someone has difficulty comprehending a piece of text, regressive eye movements take place – in other words, their eyes move back to earlier parts of the text. These movements are quite common, even in reading straightforward text, as a means of checking earlier information to aid interpretation.

LANGUAGE DISORDERS

Aphasia

Loss of language function is called aphasia (see chapter 1) – strictly dysphasia when there is partial language loss, but the term 'aphasia' is commonly used for all types of language loss.
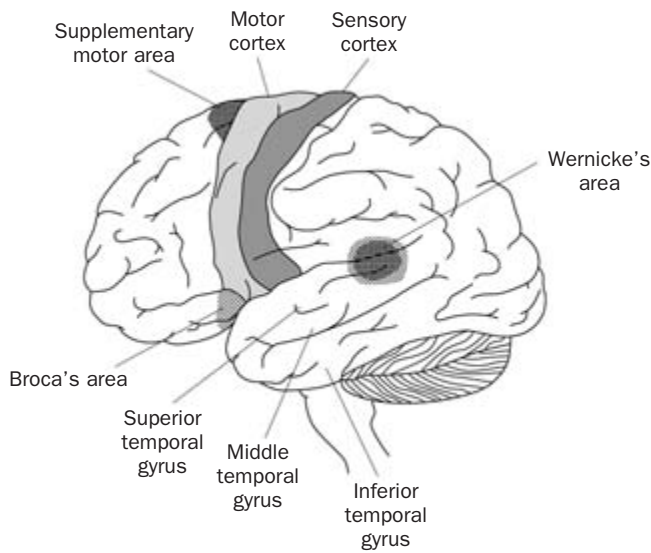
Aphasia is diagnosed when there are language difficulties that occur in the absence of sensory impairments or thought disturbances – in other words, the symptoms are specific to language.

The traumatic event of a stroke often results in an inability to use language to some degree, and is a sadly common occurrence. Strokes (cerebrovascular accidents) affecting those parts of the brain that support language processing account for 85 per cent of aphasia cases.

The main areas of the brain implicated in aphasia are shown in figure 12.3.

The left hemisphere has long been known to be associated with language function. Damage the left hemisphere, and language dysfunction is likely to result. In particular, two areas of the brain have long been associated with specific aphasic symptoms: Broca's area, and Wernicke's area (see chapter 3).

Figure 12.1

Eye-tracking apparatus enables researchers to establish how the eye moves and where it fixates within a piece of text.

Figure 12.3

The main areas of the brain implicated in aphasias.

Broca's aphasia (or production aphasia) Broca's area (see figure 12.3) is found to be damaged in patients with Broca's aphasia. These patients have difficulty in the production of language, some being unable speak at all, others only with difficulty. When language is produced, it lacks fluency and is slow. Speech may consist of just one or two words, with no grammar and often an absence of verbs necessary for the production of well-formed sentences. Broca's aphasics can understand language, though. This is demonstrated by their capacity to follow instructions or to verify whether scenes match sentences.

Wernicke's aphasia (or sensory aphasia) Patients with Wernicke's aphasia have a problem in comprehending the speech of others, and although they can produce a fluent flow of speech, it is usually garbled, containing many pseudo-words (so-called jargon). Because they cannot understand the speech of others, they also may not be aware that they are not making sense. They suffer from word retrieval deficits and cannot properly parse sentences. These effects result from lesions to Wernicke's area (see figure 12.3).

Other types of aphasia include the debilitating global aphasia, in which heard speech cannot be comprehended or even repeated, there is no capacity to produce speech, and even objects cannot be named. Another category is conduction aphasia, in which patients have an apparently normal capacity to understand and produce speech. But they have difficulty repeating word-strings and 'nonsense' words. This condition has been attributed to damage to fibres connecting Broca's and Wernicke's areas.

Psychologists who study the changes that occur in aphasia will explore specifics, such as whether the patient has difficulty finding the right words in normal speech, repeating words and sentences, using grammar so that they can understand sentences, or producing grammatical outputs themselves. For further information on treatments of aphasia, see Zurif and Swinney (1994).

Dyslexia

Dyslexia means impaired reading. There are two broad categories: acquired dyslexia and developmental dyslexia.

> **dyslexia** impaired reading due to trauma or developmental factors

1 Acquired dyslexia Brain damage in people who could previously read well can lead to acquired dyslexia. There are four main classes of this disorder:

1. People with visual form dyslexia might not be able to recognize all the individual letters. So they might read 'mat' as 'cat'.
2. Those with phonological dyslexia have difficulty reading pronounceable pseudo words, like 'pleke', but they are good at reading real words. This shows that their problem is caused by damage to the mechanism that connects how a word looks (its orthography) to how it sounds (its phonology). By contrast, when they read well known real words, these patients can use direct routes between the whole word pattern and its sound – these direct routes are established when we learn to read.
3. Surface dyslexia is the opposite way round to phonological dyslexia. People with this disorder are unable to use this direct route to recognize words on the basis of their overall appearance, but they can read words by using orthographic knowledge. This means that they make errors pronouncing words that are irregular in the mapping between the letters and the sound, like 'pint' or 'yacht'.
4. Deep dyslexia forms a very interesting category. On being asked to repeat concrete nouns, such as 'uncle', the patient may say 'aunt' instead: i.e. they substitute a semantically related item. These patients cannot read abstract words and pronounceable pseudo words. Deep dyslexia is associated with widespread left hemisphere damage, and tends to co-occur with aphasia.

2 Developmental dyslexia This refers to a developmental difficulty with reading, despite adequate intelligence. Attempts to match the reading difficulties to the categories of acquired dyslexia have led the division of syndromes into two main types: those associated with difficulties in 'sounding out' (as in acquired phonological dyslexia) and those related to difficulties in recognizing word forms (as in surface dyslexia). But one prevalent problem for most developmental dyslexics is poor phonological awareness: so they perform badly on tests of rhyme awareness, segmenting words into individual sounds (spelling out) and producing rhymes.

The detailed study of dyslexia entails the application of well developed psycholinguistic techniques. For a review of one hundred years of work in this area, see Miles and Miles (1999).

Everyday Psychology

The tip-of-the-tongue (TOT) phenomenon

'I left it on the, you know, what's its name?' . . . We have all had the experience of being unable to find a word that we want, when we want it. It is there, somewhere, we can feel it, but we cannot find it or speak it. We say it is 'on the tip of my tongue'.

The tip-of-the-tongue (TOT) phenomenon is defined as the temporary inability to retrieve a word that is well known to the speaker. How have psychologists investigated TOTs? When are they more likely to occur? Why do they occur? What does it tell us about everyday speech production and memory?

Researchers have used two main approaches to study TOTs, both rely on introspection. Researchers have either elicited TOTs experimentally in the laboratory, or asked speakers to record spontaneously occurring TOTs in a diary. In laboratory studies (e.g. Brown & McNeill, 1966) speakers are presented with a definition of a rare word (e.g. 'a place where bees are raised for honey' [word = *apiary*]), then asked to provide the name. Those speakers who are unable to say the word, but who report that they know the word and feel that it is on the verge of occurring to them are in a TOT state. They are then asked to provide information about the unavailable word (e.g. other words that come to mind, the initial letter etc). In diary studies (e.g. Burke et al., 1991), similar questions are asked, as well as which strategies the speaker uses to resolve TOT states.

The incidence of TOTs is related to two main classes of factors (see Brown, 1991), speaker-related factors and word-related factors. Speaker-related factors include age and brain damage. TOTs are more common for older than younger speakers (Burke et al., 1991), and much higher in patients with *anomia* (language-specific disturbance arising after brain damage; Vigliocco et al., 1999). Word-related factors include word type (TOTs were highest for proper names; Burke et al., 1991; see also chapter 11), how frequently a word is used and how recently it has been used.

TOT can be interpreted as a failure of memory retrieval, rather than a problem of either memory encoding or storage. Speakers in a TOT state are able to report some information about the word at better than chance levels: its meaning, its grammatical features, its number of syllables, and its beginning sound or letter. This evidence has been interpreted in models of speech production as supporting a dual-stage word retrieval process: TOT states do not occur during the first step, retrieving memory and syntax, but during the second, retrieval of form (i.e. the sound pattern of the word).

Thus this mundane speech error has provided psychologists with fascinating material about everyday language and cognition. TOTs have also been used as a tool for studying 'meta-cognition' (i.e. what people 'know that they know'; see Koriat, 1993), and for improving our understanding of psychological deficits associated with brain damage in patients with anomia or Parkinson's disease (Matison et al., 1982; Vigliocco et al., 1999).

Brown, A., 1991, 'A review of the tip-of-the-tongue experience', *Psychological Bulletin*, 109, 204–33.

THE POWER OF METAPHOR

There is an interesting theory that the natural metaphors we use to talk about things influence our descriptions and the way we think. Over the past 20 years or so, Lakoff and his colleagues (Lakoff, 1987; Lakoff & Johnson, 1980) have presented a remarkable set of observations about the role that metaphorical systems play in both our thinking and our language. In general, the Lakovian claim is that the conceptual system relies on metaphor because this is equivalent to setting up mental models, and that these then constrain the way we think and communicate.

Metaphors are much more prevalent than you might think (e.g. Cacciari & Glucksberg, 1994). Far from being restricted to specialist literary uses, they permeate our language in such a way that they surely must reflect something about the way our conceptual structures support understanding in general. Lakoff suggests that there are certain fundamental ways in which we think about things. This kind of analogical thinking finds its way into our language in striking ways.

For example, Lakoff (1987) considers the conceptions we have about anger. There are many expressions relating to anger, which, if taken literally, make no sense at all:

John was so angry, he hit the ceiling [roof]
Mary blew her stack when she heard the news.
When John broke the news, Mary exploded.
There was steam coming out of his ears.

Lakoff claims that mental models of anger result from simple observations, like an increase in internal pressure (blood pressure, heart pounding), becoming hot and sweaty, etc. These observations can be understood in terms of familiar everyday experiences with the material world, such as heating things up in containers. So Lakoff suggests that one way in which we conceptualize anger is in terms of heat being applied to a container that may contain a liquid (e.g. 'she was boiling/seething').

Once the model of heat being applied to a container is being used, it is generative – that is, it leads to outcomes, like steam. To keep the steam in, a lid is normally used. So we get expressions

like 'Contain your anger' and 'Put a lid on it'. A lid on a container generates other possibilities, too. For instance, increased pressure leads to the lid coming off – 'He flipped his lid' – and ultimate explosion – 'John exploded with rage', 'Mary blew her top'.

In his case study of anger, Lakoff suggests many more metaphors may be produced on this well known, simple basis. And we can understand statements like 'I thought he was going to erupt' because of these well worn conceptual connections. If you overheard this statement in a conversation, you would likely infer that it was about anger.

Lakoff's basic argument is that we have very simple but significant and repeated experiences of certain types. For instance, we all go on journeys, we all use pots and containers, and we all engage in some sort of competition. We are also familiar with conflict and war, albeit to different degrees. These common experiences are used as models for the interpretation of a wide range of phenomena. So, in the anger case, containers play a central role. In the same way, the idea of a journey can form the basis of understanding relationships – 'This relationship is at a dead end/ isn't going anywhere' – or arguments – 'At least we are getting near the conclusion.'

These attractive ideas merit very careful consideration, not just because they have the potential to explain the wide variety of metaphorical features of language, but because of the influence they have on thought and communication.

It was a very deliberate act of dehumanization when the Nazi propagandists portrayed Jews as rats in films, justifying the treatment of people in an inhuman way. More recently, in Rwanda, propaganda by one group, Hutu, described the other group, Tutsi, as 'cockroaches'. In a similar vein, many things that require action are thought of in terms of war. For instance, Lakoff and Johnson (1980) cite Jimmy Carter, one-time president of the USA, as reacting to the energy crisis by declaring 'the moral equivalent of war'. They point out that this opens up a set of analogues of war concepts. So there will be an 'enemy', a 'target' will be set, 'sacrifices' will be called for, and so on.

THOUGHT

The study of thinking concerns how we come to understand the world, what our intuitions are, and how we reason, solve problems and make judgements and decisions about things.

The cognitive approach to thinking attempts to discover the processing activities that characterize these mental acts. As with language, a great deal of progress has been made as a result of adopting a procedural approach. But the most striking thing to emerge from the study of thinking is that, as a species, although we can solve some amazingly difficult problems, we can also fail on others that seem quite simple.

Two main strands have coexisted in the study of thinking for many years: problem solving and reasoning. Problem solving has revolved around the study of puzzles and how people solve them, while reasoning has been more concerned with what conclusions people draw, logical or otherwise, on the basis of knowledge and evidence.

More recently, studies in both areas have stressed the nature of the representation that results from trying to understand what a problem is about. This has led to the suggestion that people form mental models of problems, which represent, as far as possible, the crucial aspects of the problems. In this way, mental model theory links thinking to language comprehension (Johnson-Laird, 1983), placing great emphasis on how problems are both understood and represented.

PROBLEM SOLVING

If I ask you, 'What is 6 + 6?', unless you are a young schoolchild, you will be able to retrieve the answer 12 straight from memory. On the other hand, if I ask you, 'What is 37 + 568?', you have to do some problem solving.

Being numerate means that you know how to solve this problem: it calls for a standard application of arithmetic procedures, and these procedures can be drawn from memory. This kind of problem-solving is called simply routine problem solving. In contrast, creative problem solving cannot be done according to a formula because there are no standard procedures in memory.

As we experience the same problem type over and over again, what was at first creative may become routine, of course.

Search space

Consider this anagram problem:

> What two-word phrase can be made out of these letters: SPROGOLIBVELM?

What strategies would you employ to solve it? The simplest is blind search, in which you just move the letters around blindly until a phrase appears. The possibilities here are enormous, so blind search is clearly not a very smart way to proceed. But how do we constrain the search?

There are some sequences of letters in English that are legal and commonplace (like 'pro'), some that are rare (like 'goli'), and some that are downright impossible (like 'blvm'). So a smarter strategy is to try constructing fragments from common grammatically legal combinations, then trying sequences that are more and more rare. Fragments will serve to cue word possibilities that you know, which will help speed up the search. With practice, people who like anagrams in crosswords develop a number of ways to constrain the *search space*.

> **search space** a space of possible reasoning, within which we search for a path linking the problem with the solution

All problems can be construed in terms of search spaces, though this is more obvious with some problems than with others. In their classic book *Human Problem Solving*, Newell and Simon (1972) illustrated the problems of search space more thoroughly than anyone had before. One problem they studied in some detail is the following (cover the solution and try the problem first):

For each letter, substitute one digit, such that the whole thing fits the laws of (base 10) arithmetic; in the example below D = 5:

```
  DONALD
+ GERALD
  ROBERT
```

Solution:

```
  526485
+ 197485
  723970
```

You will notice that your perception of what is involved in the problem increases as you work on it. For instance, to begin with, you may not have noticed the problem of carrying. That is, you will need to add 1 to a column left of the one you are working on if the sum exceeds 9.

speak-aloud protocols a description of our own processes of thinking during a problem solving task

Newell and Simon (1972) collected *speak-aloud protocols* – they required people to say aloud what they were doing while they were attempting problems like this. This helped them to analyse in detail the steps people go through in problem solving. There were two main findings:

1. People set up initial representations of problems, which influence the search space.
2. They employ general purpose techniques, called heuristics, which help constrain the search space.

So, with the problem above, Newell and Simon found several possible representations. For instance, some people saw it as being one based on word meaning. Suppose the puzzle was:

```
  BILL
+ WAS
= KING
```

A person might reason that BILL = William the conqueror → 1066, therefore B = 1, I = 0, L = 6. This kind of reasoning turns out to be inappropriate for our particular problem given above. Other examples might be described as typographic – E looks a bit like 3, etc. – and cryptographic – using some sort of systematic code, like A = 1, B = 2, etc. Neither applies to our particular example, but the important point here is that our initial conception of the problem can alter the way in which we attempt to solve it.

Understanding how people develop a problem space – the representation of a problem in the head of an individual – is a major aspect of work on problem solving. (The more general idea of a Mental Model is discussed later in this chapter.) For instance, when we learn how to problem solve, we must first recognize when seemingly different problems have a common logical structure.

Looking for a common structure

A classic study of how underlying common structure might be spotted was carried out by Gick and Holyoak (1980; 1983). They examined how experience with a puzzle called the 'military problem' (Holyoak, 1984) affected performance on a second problem, called the 'radiation problem' (Duncker, 1945):

The military problem
A general wishes to capture a fortress located at the center of a country. There are many roads radiating out from the fortress. All have been mined, so that although small groups of men can pass over the roads safely, any large force will detonate the mines. A full-scale direct attack is therefore impossible. What should the general do? (Holyoak, 1984, p. 205)

The radiation problem
Imagine that you are a doctor treating a patient with a malignant stomach tumor. You cannot operate because of the severity of the cancer, but you must destroy the cancer. You could use high-intensity X-rays. However, the intensity needed is such that the beam would destroy the healthy tissue that the rays must pass through. A lower intensity beam would not harm the healthy tissue, but would also not destroy the cancer. How can you use X-rays to destroy the tumor without destroying the healthy tissue? (adapted from Duncker, 1945)

The solution to the two problems is very similar. In the case of the radiation problem, the solution is to direct weak X-rays from a number of different points outside of the body, and to set the sources up so that the beams converge at the site of the tumor. That way, no single beam is strong enough to cause damage to healthy tissue, but the combined effect on the tumour is enough to destroy it. The military problem has a solution based on the same principle: small groups of soldiers are sent along different roads at the same time, converging as one big army at the fortress.

Gick and Holyoak had participants do the military problem first. One group of participants simply read the problem in the belief that they were just to recall the wording. Under those circumstances, only 30 per cent derived the correct solution to the radiation problem. However, if the participants were given two similar problems before the radiation problem, then there was more transfer. In general, though, the more superficially similar problems are, the better the transfer (Holyoak, 1990). So spotting the similarity of problems is far from automatic.

LOGICAL AND CONDITIONAL REASONING

Logical reasoning

A special form of problem solving is *logical reasoning*. In these kinds of task, people are required to draw conclusions that necessarily follow

logical reasoning reasoning about issues whose conclusions necessarily follow from what is given

from a given, but not to draw conclusions about what might possibly follow. For example, in this syllogistic reasoning task, two premises enable conclusions to be drawn:

If all men have blood, and John is a man, then, necessarily, John has blood.

But syllogisms are not always this easy, and some can lead to false conclusions. For example:

> If some As are Bs, and some Bs are Cs, what can be said about the relation of As to Cs?

A common error is to say: Some As are Cs. But while this may be case, it is not necessarily true. Those Bs that are Cs might be the ones that are not As. Johnson-Laird (1983) suggested that when people get this wrong, it is not because they are not 'logical'; it is because they have an inadequate representation of the problem – what he calls a mental model (see figure 12.4). Johnson-Laird was able to show that forming such models is harder with some premises than others, and that the harder it is (i.e. the more complex the mental models), the more likely it is that we will make an error.

Conditional reasoning

Another much studied type of logical reasoning is *conditional reasoning*, which deals with 'if–then' statements. For instance:

> If a green light comes on, then the ball has rolled left.

Suppose the ball has rolled left. What can we conclude? A common error is to conclude that the green light must have come on (Rips & Marcus, 1977), but this is not a necessary conclusion. The ball could have rolled left for any number of other reasons.

(a)   (b)
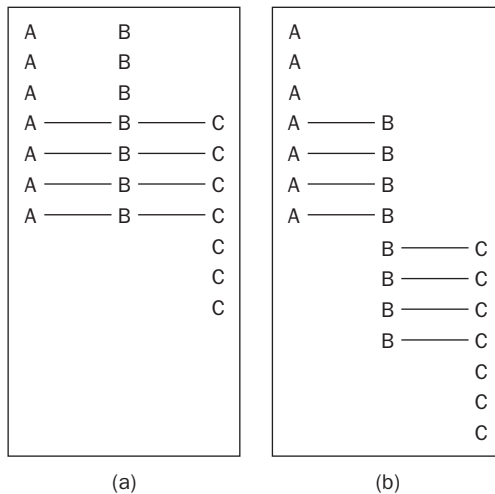
Figure 12.4

A representation of two mental models for the syllogism.
Some As are Bs
Some Bs are Cs
What follows?
Both (a) and (b) are possible representations. The difference is that in (a), the Bs that are connected to Cs are also connected to As. In (b), the Bs that are connected to Cs are not connected to As. This means that although a possible conclusion is that 'Some As are Cs', this is not necessarily true (as in case (b)). Johnson-Laird's idea is that we set up such mental models when reasoning about syllogisms, and that unless we set up all possibilities, we may erroneously consider a possible conclusion to be logically necessary.

This error is called 'confirming the antecedent'. Does the fact that the error occurs mean that people are not logical? This is the wrong way of thinking about the issue. Like the logical error, what it means is that some people have the wrong representation of the problem, and this leads to false conclusions. For instance, the abstract form of the problem, 'If A then B. B, so . . . ?', suggests that there is only A and B to consider, in which case it is reasonable to suppose that if B, then A. But, in general, there can always be some other cause for B – it simply is not stated. So it is easy to confirm the antecedent. For instance, if you commit murder, you go to jail. But if you go to jail . . . this does not mean you committed murder!

Detecting cheats

A very important way of testing if–then statements is known as the Wason Selection – or four-card problem (Wason, 1966). In this task, the participant is given a rule, and four cards are laid out that have information written on both sides. For example:

> Rule: If a card has a vowel on one side, then it has an even number on the other side.
> Card 1: A
> Card 2: D
> Card 3: 4
> Card 4: 7

The task is to verify (i.e. test) whether the rule holds by turning over the two cards that will enable this to be determined. Which cards would you turn over to verify the rule? Try it before you continue reading.

The most frequent response is to check A and 4. Turning A will provide information that is consistent with the rule if there is an even number on the other side of the card, and will falsify the rule if there is an uneven number, so that is fine. But turning 4 will achieve nothing, because the rule does not say, 'If a card has an even number on one side, it will have a vowel on the other.' Turning this card is very much like confirming the antecedent. In fact, the crucial second card to turn is the card with the 7, because if this has a vowel on it, then the rule is false.

This problem is hard to think about. But real-life versions can be much easier. For instance:

> If a student is drinking beer, then they are over 18.
> Card 1: Over 18
> Card 2: Drinking beer
> Card 3: Drinking Coke
> Card 4: Under 18

How would you test the rule? Most people would now think the crucial card to turn was card 4, 'Under 18', because if that had 'Drinking beer' on the other side, there is a clear violation of the rule. This is because testing for under-age drinking is an example of detecting cheating, which is something we appear to be good at (Cosmides, 1989; Gigerenzer & Hug, 1992). The argument is that we have social rules to live by, and that we are naturally attuned to be able to test whether these rules are being broken.

Clearly the representation of the problem is crucial to how reasoning takes place. When a concrete form of the problem is used, we can bring in specific procedures that we have access to for detecting cheats, which is something that is socially important. With an abstract version of the task, this is not possible.

HEURISTIC REASONING – OR TAKING A SHORT CUT

Thinking, understanding and decision-making take place in the real world, where there are usually time pressures and rarely a full range of information available to support a complete appraisal of the problem at hand.

For instance, suppose you are buying a new washing machine. A good basis for the decision might include comparative data on reliability, ease of servicing, servicing and repair costs, ease of use, even noise levels during operation. The list could go on and on. Although sometimes data of this sort might be available, and sometimes it might be published in magazines, it is more likely that you will have to cut corners. In other words, you might not be able to obtain a machine that fulfils all of your desirable features, but you will instead settle for the closest that is available.

heuristic reasoning solving a problem by using a method that is likely to give the right answer, although there is no guarantee

Kahneman, Slovic and Tversky (1982) popularized the term *heuristic reasoning* for thinking and decision making that involves these types of short cuts. They also suggested that these short cuts are so common that they should be considered part of the machinery of thought itself.

Availability

availability relies on the possibility that a solution (from heuristic reasoning) that readily or quickly comes to mind may be the correct one

Perhaps the simplest kind of heuristic reasoning is *availability*. The availability heuristic is a method of estimating the likelihood of something based

---

> ## *Pioneer*
>
> **Daniel Kahneman** (1934–) has conducted highly influential work over the last several decades into human reasoning, specifically regarding the role of heuristics (i.e. reasoning short cuts, using strategies that generally work but are not guaranteed to work). To a large extent, heuristic reasoning overlaps considerably with the everyday idea of intuition. Kahneman and colleagues have suggested that these heuristic short cuts are so common that they should be considered part of the machinery of thought itself. For example, the availability heuristic is a method of estimating the likelihood of something based on how easily it comes to mind. The representativeness heuristic is based on the principle that we can estimate the likelihood of something by seeing how well it fits a prototype of which it may be an exemplar. For his body of work investigating human judgement and decision-making under conditions of uncertainty, Kahneman was awarded the Nobel Prize in 2002.

on how easily it comes to mind. For instance, we might assess the divorce rate by thinking of its prevalence amongst people we know personally. Or when buying a car, we might estimate reliability from comments made by acquaintances and colleagues. Because there will generally be a correspondence between what comes to mind easily and the likelihood of the underlying event, this heuristic can be useful.

Kahneman et al. (1982) point to two mechanisms that come under the availability rubric: ease of recalling relevant instances and ease of constructing representations. For instance, someone's estimate of how many flower names they know will directly depend on how many they can think of in a short time – say, two minutes (Tversky & Kahneman, 1973). In this case, there is generally a good correspondence between initial rate of retrieval and the total number known. But this is not always the case. For instance, it is easier to recall the names of famous people than ordinary people. So if participants hear lists of names containing equal numbers of famous and non-famous names, they will typically believe that there are more famous people on the list than ordinary ones (Tversky & Kahneman, 1973). Here, the heuristic leads to a biased outcome.

Another example of bias occurs through the construction of representations. Consider the following problem:

> A group of ten people want to form a committee with only two people in it. How many possible committees are there?

Now try this:

> A group of ten people want to form a committee with eight people in it. How many possible committees are there?

Most people produce a higher figure for the first question than for the second, even though they are actually equivalent questions

(because 8 + 2 = 10, so for every committee of 2 that is formed there is an equivalent committee of 8 formed from among the same group of 10 people). Tversky and Kahneman argue that this is because it is easier to imagine several committees of two than several committees of eight. (This seems reasonable if we suppose that it is easier to form and manipulate a mental model with two rather than eight tokens in it.)

The availability heuristic has been used to explain many, many phenomena. In risk perception, for example, people tend to over-estimate car accidents, tornadoes and homicide as causes of death, and underestimate death from complications due to diabetes, stroke and smallpox vaccination. Furthermore, studies show a good correlation between the prevalence of these events in news reports (availability) and estimated likelihood as a cause of personal death (Slovic, Fischhoff & Lichtenstein, 1979).

Social psychology research has established that individuals tend to think that they initiated arguments with significant others more than 50 per cent of the time, and that they did more than 50 per cent of the work in domestic situations. This applies to both partners! It is argued that this is because we each have ready access to information about our own contributions in these situations, so we are more likely to register and remember these than our partner's contributions (because of the higher availability of the former) (Ross, 1981; Ross & Sicoly, 1979).

Representativeness

This heuristic is based on the principle that we can estimate the likelihood of something by seeing how well it fits a prototype of which it may be an exemplar. For instance, if you are trying to decide whether a person is a Christian, the more properties they have that fit your model of how Christians behave, and the fewer they have that do not fit, the more confident you would be that the person is a Christian.

representativeness used in heuristic reasoning to decide whether something is likely because it corresponds to an idea of what is typical in that situation

Like availability, *representativeness* is a double-edged weapon – it can lead to fallacious reasoning. Many of the examples Kahneman and Tversky (1972) give are about reasoning with distributions, such as the 'Exact Birth Order Problem':

> All families of six children in a city were surveyed. In 72 families, the exact order of boys and girls was GBGBBG. What is your estimate of the number of families found in which the exact order was BGBBBB?

The majority of participants thought that the first sequence was much more likely. In fact, the two orders are almost equally likely because, on any occasion, either a boy or a girl could be born with approximately equal probability. Both of these orders fulfil this requirement. From an intuitive viewpoint, the first seems much more likely because there is an equal number of girls and boys. But the equal number gives the impression of being more likely seemingly because it is judged to be more representative.

The impact of representativeness on exact order judgements can be seen even more clearly with the following:

> Which is more likely to occur: GGGBBB or GBBGBG?

Most people think it is the latter, because it is more 'random-looking' than the former. Yet on a random draw basis, both examples are equally likely.

To make this clearer, draw out all the possible sequences that could occur using three boys and three girls. Although the sequences are all equally likely, there are more 'mixed up' ones like the second one above, and only one other (BBBGGG) that looks more extreme (and therefore less representative). Yet these possibilities are all equally likely.

Another example shows how representativeness can apparently obscure the use of what is termed base-rate information. Consider the following scenario:

> 100 people, comprising 70 lawyers and 30 engineers, apply for a job. One of the applicants, Dick, is a 30-year-old man, married with no children. A man of high ability and motivation, he is likely to be quite successful in his field. He is well liked by his colleagues.

Is Dick more likely to be an engineer, a lawyer or equally likely to be either? Kahneman and Tversky (1972) found that the predominant answer given was 'equally likely' because the information does not discriminate between the two. Yet the prior odds are 70:30 in favour of Dick being a lawyer, so this should be the answer in the case where there is insufficient extra evidence in the description. In such cases, it is as if the representativeness of the description dominates the thinking of participants – a typical illustration of what is widely known as the 'fallacy of ignoring the base-rate'.

INTUITION

Heuristics provide a means of reasoning, but they are short cuts, using strategies that generally work but are not guaranteed to work. At the same time, they can induce quite high levels of confidence in us regarding our decisions, even when we are wrong.

To a large extent, heuristic reasoning overlaps considerably with the everyday idea of intuition. Intuitive thought is automatic, often fast and not derived from detailed analysis. It involves a strong feeling of conviction but – like heuristic reasoning – tends to be hard to justify.

Problem-to-model mapping

The mappings from the description of a problem to an automatic conception of that problem can be very strong, and constitute the basis of some very strong feelings of the intuitive 'correctness' of our understanding. Try the following problem (schematically illustrated in figure 12.5):

> Suppose there are three cups in front of you and the experimenter puts a coin under one of the cups. You don't know which one it is
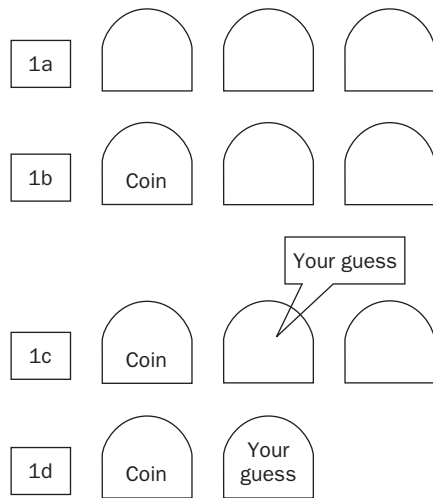
Figure 12.5

The 'Three Cups Problem', showing the sequence of events. A description of the task is given in the text.

under. Next you try to choose the cup you think the coin might be under. Rather than tell you whether you are right or wrong, the experimenter removes one of the cups, but not the one you pointed at, and not the one the coin was under (which may be different). The question is, would you have a greater chance of getting the coin if you stuck to your original choice, or shifted?

Participants usually believe that they have a 1:3 chance of being correct when they start, and then that they have a 1:2 chance of being right once there are just the two cups left. They usually declare that there is no point in changing because after the cup has been removed thay have a 50/50 chance of being correct (and if they changed their choice at this stage they would still only have a 50/50 chance of being correct). This behaviour fits a simple mental model: with N choices, the chance of being correct is 1:N. The situation is mapped onto this simple model, and the result is coherent and compelling.

Despite this, the answer is that you should shift. In the first place, the chance of being correct was 1:3, and the chance of being incorrect was 2:3. But the important point is that the experimenter does not remove the cup at random, and – the key point – she never moves the cup that contains the coin. So the chance of being wrong by sticking to the original decision is still 2:3 (as per the original decision), even though there are only two cups now left. But since there is only one other cup now remaining, the chance of that being the wrong choice is in fact 1:3 (because there is only one other of the original three cups under which the coin could now be located), so it makes sense to change. In fact, the odds in favour of changing are 2:1.

This is a very difficult puzzle to think about (e.g. see Granberg & Brown, 1995). The usual mental model people set up does not have the capacity to deal with the correct solution, and yet it is very compelling.

There is an intuitive way of making the point about shifting, though. Suppose there are 100 cups (each numbered), and one has a coin under it. The chance of your being incorrect in your choice is 99:100. You choose a cup – say, number 15. Now the experimenter takes away all of the cups except the one you chose and one other (say number 78), but you know she never takes the one with the coin under it. Do you now think that there are even odds on your having selected the correct one, or would you prefer to shift? Most people think it appropriate to shift under those circumstances.

The 'Three Cups Problem' is a good illustration of a strong mapping between a state of affairs (two cups are left) and a pre-existing mental model (if there are two cups, one with a coin under it, then the odds on choosing the correct one are 50:50). The intuitive belief that goes with these problem-to-model mappings is very strong. Try it on your friends.

The hindsight bias

Just as discourse makes sense if it portrays a series of connected events that match some plausible possible world, so facts about things make sense if they fit a coherent scenario. Also, once we know the facts, it is often easy to find a way of linking them.

Nowhere is this clearer than with the *hindsight bias* (see chapter 1), in which people believe that they had a prior insight ('I knew it all along') and that an event was therefore not surprising. Hindsight judgements are made 'after the fact'.

> **hindsight bias** falsely overestimating the probability with which we would have predicted an outcome after we know it has already occurred

In a typical hindsight experiment (Fischhoff, 1977; Slovic & Fischhoff, 1977), participants first answer binary-choice general knowledge questions, such as:

Was Aladdin (a) Chinese?
            (b) Persian?

Subsequently, they are presented with the questions again, this time with the correct alternative marked, and are asked to say whether they got each one right on the previous occasion.

In general, participants tend to falsely remember getting more right than they actually did, as though the correct answer interferes with their memories. Even if participants are paid for remembering correctly, the effect still occurs, so strong are the intuitions the paradigm generates.

A major consequence of the hindsight bias is that things appear to be more obvious than they should. Before new experiments are carried out, it is never clear what the outcome will be – otherwise they would not be original experiments. Yet in one interesting study, the same information was presented to two groups of participants concerning an experiment with rats. One group was told that one result occurred, while the other group was told that another occurred. Although the two sets of results were quite different, both groups of participants rated the outcome as obvious (Slovic & Fischhoff, 1977).

ARE HUMANS POOR AT REASONING?

This brief overview of the literature on thinking might lead us to wonder whether we are capable of being rational and logical, or whether we fall short of that 'ideal'.

Caution is needed with this question. Survival depends on being good at doing things that confront us in the real world. Rather than think of rationality as an absolute, Herbert Simon (1991) introduced the idea of *satisficing* – that is, performing optimally with the limited data and time available to us. This is known as *bounded rationality* – it is about as close to the idea of being rational as we are likely to get, and is the best we could expect from any system with finite resources.

It has also been argued that many of the tasks used in laboratories are artificial, and that they lack *ecological validity*. In other words, they are not typical of the kinds of problem humans have to solve. (For a discussion of this important idea, see Cosmides & Tooby, 1997; Gigerenzer & Hoffrage, 1995.) Gigerenzer and Hoffrage show, for example, that when information is presented in terms of frequencies (like 95 out of 100) rather than probabilities (like 0.95), people do better at a range of reasoning tasks, and ignore base-rates to a lesser degree. They argue that this is because we are naturally adapted to frequency information because we tend to collect instances one at a time. These authors are working on a program of investigation into *evolutionary cognition*, which attempts to establish whether we are good at certain ways of thinking because we have evolved that way to adapt to our evolutionary environment (see also Piatelli-Palmarini, 1992).

satisficing making a judgement based on bounded rationality

bounded rationality being rational by making a rational judgements, but based on only part of the evidence

ecological validity the extent to which a task is typical of tasks that people have to solve in everyday life

evolutionary cognition cognitive processes that are established by evolution

Pioneer

Herbert Simon (1916–2001) was a true cognitive scientist, crossing disciplinary boundaries in his efforts to understand human problem solving and decision making. He was awarded the Nobel Prize in economics for his work on administrative behaviour, but is best known in psychology for his work on the representation of problems, and problem-solving heuristics (with the eminent cognitive scientist Alan Newell). In the early 1950s, Simon and Newell conceived the idea that the best way to study problem-solving was to simulate it with computer programs. Computer simulation of human cognition subsequently became Simon's central research interest, which he pursued until his death in 2001.

Research close-up 2

Is there a 'natural' way of thinking about probabilities?

The research issue

Extracting regularities from the environment – learning what goes with what, and how frequently or probably things will co-occur – is an important aspect of our survival skills. Yet there is a mass of evidence that people are rather poor at solving the kinds of problem that use this information.

Are humans incapable of using probability information? Gigerenzer (2000; Gigerenzer & Hoffrage, 1995) suggest that it all depends on how the information is presented.

Design and procedure

Gigerenzer presented two versions of a puzzle to two groups of 24 physicians. One version used a frequency format:

A doctor screens women for breast cancer using a mammogram. He knows that:
Ten out of every 1000 women have breast cancer.
Of these ten women with breast cancer, eight will have a positive mammogram.
Of the remaining 990 women without breast cancer, 99 will still have a positive mammogram.

Imagine a sample of women (aged 40–50, no symptoms) who have positive mammograms in your breast cancer screening. How many of these women do actually have breast cancer?

_____ out of_____

Try this before reading on.

The other version of the same problem, presented to a second group of 24 physicians, used a probability format (most commonly used in these types of studies):

The probability that a woman has breast cancer is 1 per cent if she is in the 40–50 age range.
If a woman has breast cancer, the probability is 80 per cent that she will have a positive mammogram.
If a woman does not have breast cancer, the probability is 10 per cent that she will still have a positive mammogram.
Imagine a woman (aged 40–50, no symptoms) who has a positive mammogram in your screening. What is the probability that she actually has breast cancer?

_____%

Try this before reading on.

Results and implications

The two formats contain the same information, presented in different ways. Gigerenzer found that using the first format, 46 per cent of physicians gave the correct answer (around 7.7 per cent). But using the second format, only two of the physicians (8 per cent) gave the correct answer, the median estimate being 70 per cent!

The results suggest that using a frequency format enables people to set up representations for the problem that enable an easier solution than the (mathematically equivalent) probability format.

Discovering which ways are natural and easy for people to represent problems, and which ones are difficult, is important if psychologists are to understand the evolution of human understanding and thinking skills. Using the right formats for providing information about uncertainty and probability is also of great practical significance for many walks of life. Gigerenzer's insights are a valuable step along this road.

Gigerenzer, G., & Hoffrage, U., 1995, 'How to improve Bayesian reasoning without instruction: Frequency formats', *Psychological Review*, 102, 684–704.

FINAL THOUGHTS

It is possible to look at language and thinking as processes, and the attempt to determine the steps involved in the execution of mental activity is crucial to a proper understanding of that activity. In the case of language, the challenge is to specify how squiggles on a page, or complex sound stimuli, become interpreted as messages about the world. We have attempted to show how some of those stages may be isolated, and we have illustrated some of the techniques used to measure them. How language understanding works is one of the great scientific mysteries of the present age.

Thinking can seem even more intractable than language, being more difficult to break down into component processes. Yet this was exactly what Newell and Simon set out to do in their work on problem solving. Studies of intuition, something normally considered to be quite different from logical problem solving, have also benefited from carefully considering just what processes a thinker must go through in order to have an intuition.

The psychology of language and thinking is of great intrinsic interest, but there are practical reasons for adopting an information-processing approach. Designing intelligent programs that can understand language, or solve problems, has benefited from psychological modelling, and psychology has benefited from the computer scientists' need to be clear about what is happening at every stage of processing. These topics are at the core of interdisciplinary efforts to understand the nature of mental activity.

Summary

- Language can be understood in terms of its grammar (syntax), its meaning (semantics) and its significance (pragmatic interpretation).
- A major research issue is which language systems are modular and which interact with each other, and when.
- Language consists not of separate sentences, but of connected discourse. In this chapter, we have considered how meaning and significance are established with multi-sentence written texts.
- Normal reading depends upon eye movements; eye movements can be used to measure difficulties in reading texts.
- Aphasia refers to language disorders that may affect the comprehension and production of language, and whether the language that is produced is grammatical or not. Aphasia results from brain injury. Dyslexia is a disorder of reading, and may be detected during individual development, or may be acquired later through brain injury.
- Problem solving is characterized by the development of a mental representation of the problem: the problem-space.
- Logical reasoning is concerned with how people draw necessary conclusions from particular states of affairs. Central to reasoning is the concept of a mental model of the problem. Socially important activities like detecting rule-breaking are also important.
- Lakoff has argued that all thinking is to some extent metaphorical. We considered some examples.
- Several studies show that humans can be poor at solving probability problems. They can often rely on simple heuristics, which sometimes lead to the wrong conclusions.
- Availability and representativeness are two important heuristics that people often use when reasoning.
- Intuition can be seen as part of the mechanism of mapping problems onto existing cognitive representations. If the fit is good, the intuition seems strong.
- Humans are not necessarily poor at reasoning. If they have the information in the right format, reasoning can be very effective.

REVISION QUESTIONS

1. What are the distinctions between syntax, semantics and pragmatics? Can you think of examples from your daily life where one of these components of language may have failed you in the past?
2. Why is modularity an important concept for language?
3. How would you define discourse? Is 'communication' different from 'language'?
4. 'Speed reading' courses often emphasize eye movements during reading. What do you think is the basis for these techniques? Do you think they are effective?
5. What is aphasia? Do you think that we all suffer from some form of aphasia from time to time?
6. Does acquired dyslexia have the same functional characteristics as developmental dyslexia?
7. What types of problem solving do you think humans are good at, and why? What about computers?
8. Sometimes people believe that the results of psychology experiments are 'obvious', and that the research was not worthwhile. What pattern of reasoning would lead to such a judgement, even in cases when, strictly speaking, the results of the experiment could not have been predicted?
9. Can you think of an example from the everyday life of you or your friends where you have used a) the availability heuristic or b) the representativeness heuristic during problem solving?

FURTHER READING

Altmann, G.T.M. (1997). *The Ascent of Babel*. Oxford: Oxford University Press.
A wide-ranging discussion of the psychology of language.

Davidson, J.E., & Sternberg, R.J. (eds) (2003). *The Psychology of Problem Solving*. London: Cambridge University Press.
Current issues in problem solving, alongside historical background.

Garrod, S.C., & Pickering, M.J. (eds) (1999). *Language Processing*. Hove: Psychology Press.
A collection of readings describing state-of-the-art research in a number of key areas.

Gigerenzer, G. (2002). *Reckoning with Risk*. London: Penguin.
An exciting new view of thought processes associated with handling risk.

Johnson-Laird, P.N. (1983). *Mental Models*. Cambridge: Cambridge University Press.
A classic book expounding a key theory of understanding and reasoning.

Rayner, K., & Polletsek, A. (1989). *The Psychology of Reading*. Englewood Cliffs, NJ: Prentice Hall.
Another classic text, describing eye movement and more general research in reading.

Contributing author:
Tony Sanford