CHAPTER

1 Introduction

PREVIEW

KEY TERMS

mental lexicon
process
psycholinguistics
representation

This chapter provides an overview both of the field of psycholinguistics and of the book itself. By the end of the chapter you should have a better understanding of, amongst other things, the appeal, significance and subject matter of psycholinguistics.

1.1 Introduction

Although there is little agreement on how many words an adult native speaker of English might know, let us take a conservative estimate of 20,000 (based on Nation, 2006, but see Section 8.2), and let us also assume that *splundle* is not one of them. Imagine such a native speaker searching through their 20,000-word **mental lexicon**, as we call the dictionary in our heads. Let's assume that they do so at a seemingly impressive rate of 100 words per second (which incidentally is 20 times faster than a good reading pace of 300 words per minute). At this rate, if they searched exhaustively through their mental lexicon, it would still take over 3 minutes to confirm that *splundle* is not there. Yet if you ask someone if they know the word *splundle*, they will be able to tell you more or less instantly that they do not. Clearly something is wrong with the assumptions we have just made about how rapidly we can look up words in our mental lexicon, or about the way in which we search through it, or both.

Ask someone to tell you what they had for breakfast, and then ask them to explain the workings of their coffee machine, and you are likely to find that their speech is much more hesitant in the second task than in the first, with more errors and restarts as well as pauses and *uhms* and *ahs*. The differences here reveal something about the nature of planning involved in different speaking tasks. The locations of the hesitations might also tell you which kinds of words and/or sentence structures the speaker finds more difficult to find or to put together.

Have a conversation with someone while they are carrying out a difficult task like driving a car, and you will find that both their language production and their language comprehension is less fluent than at other times (Becic *et al.*, 2010).

If you listen out for speech errors, or slips of the tongue, you are much more likely to hear an error where the beginning sounds of two words have been swapped over, as in 'tip of the slongue', than one where the beginning of one word and the end of another have been exchanged, as in 'slit of the pongue'.

When you see a sign like that in Figure 1.1, or one that says 'Please go slowly round the bend', you might well chuckle, but your experience of misreading the sign shows that there are certain preferred patterns of analysing the structure and meaning of sentences in English.

Consider what is wrong with the following interaction between speakers A and B, in which <SILENCE> indicates noticeably long periods of silence, and CAPITALS indicates the part of the word that receives word stress:

A. Could you tell me the time of the next train to Palmerston North and how much a single fare costs?
B. <SILENCE>
A. Hello?
B. <SILENCE> Invisible tables sleeeeep violently

Figure 1.1
Thank goodness our young ones are not breaking the speed limit.

A. When does the next train go to Palmerston North and what is a single ticket?
B. <SILENCE> PalMERston North trains from platform 3 leave. A single ticket licences a one-way journey.

There are of course a number of odd features here – there is meaningless and irrelevant content, there are long silences with odd patterns of pausing, there are infelicities of vocabulary and structure ('A single ticket licences a one-way journey', 'Palmerston North trains from platform 3 leave'), there is strange stress placement ('PalMERston') and another odd pronunciation ('sleeeeep'). Looking at 'odd' speech like this (the example is fabricated but the principle remains valid) sheds some light on what we need to do in order to speak. That is, speakers usually aim to produce utterances that have appropriate meaningful content, that use appropriate lexical items and grammatical structures, and that have appropriate pronunciation, intonation and phrasing. Speakers aim to do this fluently and in real time, and rarely have opportunity to rehearse. They also have to relate what they say to the context, including to previous speech in a conversation.

Speakers are generally very good at doing this. For example, one experiment was explicitly designed to elicit subject-verb agreement errors like (1.1), where the subject is plural (*efforts*) but is followed by a singular verb (*is*). The reason for the error is that the most recent noun before the verb is singular (*language*). In the experiment, fewer than 5% of the stimuli designed to produce agreement errors like this actually did so (Bock & Miller, 1991).

(1.1) Efforts to make English the official language is gaining strength throughout the US.

In spontaneous speech, it has been found that sound errors (e.g. saying 'par cark' for 'car park') occur only about 1.5 times per 10,000 words, and word errors (where the wrong word is chosen or words change positions in a sentence) occur only about 2.5 times per 10,000 words (Deese, 1984). Clearly we are generally pretty good at what we do when we produce and understand language.

1.2 What is psycholinguistics?

The rather disparate observations in the preceding section illustrate just a few of the areas of interest in psycholinguistics. Psycholinguistics can be defined as the study of the mental **representations** and **processes** involved in language use, including the production, comprehension and storage of spoken and written language. A number of issues arise from this definition. Some are to do with *representations*, such as:

- How are words stored in the mental lexicon, i.e. the dictionary in our heads? Is the mental lexicon like a dictionary, or more like a thesaurus? For instance, is *cat* listed near the similar sounding word *catch* or near the meaning-related word *dog*? Or neither? Or both?
- Do we have phoneme-sized chunks of language in our heads? That is, as part of recognising the word *cat* do we also recognise the component sounds /k/, /æ/ and /t/?
- Do literate people have letter-sized chunks filling equivalent roles for the processing of written language?
- How is the meaning of a sentence represented in our memory?
- Is *government* a single word or *govern* + *ment*?
- Is the plural form *cats* represented in the lexicon, or just the singular *cat*?

Other questions concern the *processes* that might operate on those representations:

- How do we recognise words so effortlessly?
- Do we analyse the speech signal phoneme-by-phoneme or do we identify complete syllables or even larger units?
- Do we recognise *government* as a complete form or do we have to construct it from *govern* + *ment*?
- If *cats* is not represented in the lexicon, does that mean that we use a rule to get the plural form of *cat*, and how does this work for irregular plurals like *children*?
- When we speak, how do we convert an idea into an utterance?
- As listeners, how do we get from hearing an utterance to developing our own representation of the idea(s) being expressed by that utterance?
- What stages do we have to go through during the construction of utterances? For example, do we first generate a sentence structure and only then populate it with words from our mental lexicon, or do we first choose words and then build a structure around those words?
- Do the processes involved in language production and comprehension influence one another, and if so in what ways?

The outline sketch of language use in Figure 1.2 gives an overview of areas of interest in psycholinguistics as well as providing the basic structure of this book, which has chapters on language production (roughly following the progression shown in the second column) and on comprehension

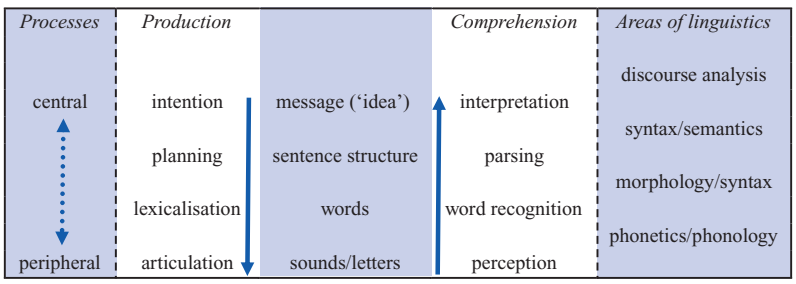

Figure 1.2
An outline sketch of language use.

(fourth column). Psycholinguistics clearly has links to other areas of linguistic study, and some of these are shown in the final column.

From the language producer's (speaker's, writer's) perspective, the production of a message takes us from an underlying intention, through stages of planning sentence structures and selecting words, to the articulation of that intention as a sequence of sounds or letters, as shown by the arrow.

From the comprehender's (listener's, reader's) viewpoint, the goal is to perceive or recognise elements such as letters and sounds in the input, to recognise words and to work out the connections between these words in sentence structures, in order to arrive at a message-level interpretation. The arrow in the 'comprehension' column shows such a '**bottom-up**' flow of information from the input to an interpretation. This is a simplification, though, as there is evidence for '**top-down**' information flow too, e.g. when a listener starts to gain an understanding of the sentence they are hearing this can influence the efficiency with which they recognise subsequent words in the sentence. Most psycholinguists today support the idea of **interactive** processing in both production and comprehension, with information flowing in both directions (bottom-up and top-down) as well as between elements at the same level (so recognising one word has an effect on the likelihood of recognising similar words).

Bottom-up: processing based on information flow from lower levels of processing to higher levels, e.g. from the sensory input (the speech signal) to the lexicon (words).

Top-down: processing guided by information flow from higher levels (e.g. sentence interpretation) to lower levels (e.g. words).

Interactive: processing that involves a combination of bottom-up and top-down flow of information, as well as links between elements at the same level.

It is reasonable to claim that the main focal areas of psycholinguistics have tended to be *sentences* and *words*. So production studies have focused on the generation of sentence structure and on syntactic planning, as well as on word finding and word building. Similarly, much of the study of comprehension has dealt with word recognition and sentence parsing (working out the syntactic structure of sentences). The study of these processes has in turn involved consideration of the representation of words and of grammar.

In effect there are many more 'fields' than this, dealing with many subprocesses, such as the perception of letters and of speech sounds, the processing of different kinds of word endings, and so on. In addition, these fields are not always clearly distinct, neither in terms of research nor in terms of the ways in which the processor operates. For instance does our choice of words influence our choice of sentence structure, does sentence-level interpretation influence phonetic perception, does the spoken form of a word affect our silent reading of it, and so on?

Although the discussion above has tried to be inclusive of both spoken and written language, the content of this book shows a bias towards the production and comprehension of spoken language. To an extent this is a personal bias, since the author's own research interests are in this area, but it also reflects the primacy of the spoken language in the development both of the individual and of the species, as well as a general bias in the research literature.

1.3 Who does psycholinguistics?

The term *psycholinguistics* appears to have first been used to refer to the psychology of language in the 1920s. However, the birth of psycholinguistics as a discipline is often linked to a seminar at Cornell University in 1951. The interdisciplinary nature of the field is reflected in the fact that this seminar was reported in both a psychological journal (Osgood & Sebeok, 1954a) and a linguistics one (Osgood & Sebeok, 1954b). Psycholinguistics continues to be taught in both of these disciplinary areas, and is influenced by research and teaching practices and methodologies in both. Courses and research in psycholinguistics are also found in multidisciplinary cognitive science units throughout the world, and the issues studied in this area are also of interest to computer scientists (particularly those interested in Artificial Intelligence and Human–Computer Interaction), philosophers (especially in areas concerning the development and representation of meaning) and others.

Of course, our reasons for studying the psychology of language may differ, according to our starting point, and this will be reflected in the nature of courses and of research agendas. Thus a characterisation of the psychological perspective might be that studying language use can contribute to our knowledge of the workings of the mind. Neuropsychologists might especially be interested in locating the language faculties within the physical structures of the brain. A more linguistic perspective may be that studying language use can inform our theories of language structure, it can provide the performance data to support theories of competence, it can provide psychological validity for linguistic constructs, and so on.

1.4 How do psycholinguists do psycholinguistics?

Although some of our knowledge in this field comes from introspection and the observation of daily behaviour, most of the major insights have come through the generation and testing of hypotheses through experiment. More recently, there has been a marked increase in high-tech observation, measuring brain activity while participants are engaged in language-related tasks. Because of this combination of types of evidence, psycholinguistics tends to blend the theoretical and descriptive insights of linguistics with the experimental methodology and rigour of psychology.

Observation

Sophisticated experimental procedures and equipment have only become available relatively recently, and so it is no surprise that early discoveries in psycholinguistics were based on more observational approaches. Early approaches also focused in particular on speech production, since the spoken output is most easily observed. For instance, we can note and analyse the occurrence of pauses and hesitations in speech, and assess these as indicators of the planning and sequencing of the spoken output. We can observe where speakers make errors, and relate these errors to hypotheses about the speech planning and production processes. We can also look at the correction of errors as an indication that speakers are monitoring their own output. Rather less often, we can note when speakers are in a 'tip-of-the-tongue' state and find out from them what they can remember of an elusive word, and see what this might tell us about the processes involved in finding words.

Some of these observational techniques can be combined with experimental intervention. We will see in later chapters, for example, that researchers have run experiments designed to elicit speech errors or to induce the tip-of-the-tongue state. A range of additional tasks has been used to study both production and comprehension.

Experiment

Experimental methods have become highly sophisticated over the past half century, in particular with the ready availability of software that allows a high level of control over the presentation of stimuli and collection of data, with precise timing. Advantages of experimental over observational approaches include the reduction of observer bias and the increased control over what participants are required to do. This last point means that researchers are more easily able to isolate aspects of the production or comprehension processes that they are interested in. The disadvantage of experimental approaches, however, is the relative lack of ecological validity – it is usually only in experimental settings that participants are asked to make an explicit judgement about whether or not a stimulus they have heard is a word of their language, or makes sense in their language.

Typically, experimental techniques involve some kind of response time measure, where the time taken by participants to carry out some language-related task is recorded, often together with their accuracy in that task. Examples of some of these tasks can be found on the website for this book. Some response time tasks involve the notion of competing attention. That is, it is assumed that the time taken by the participant to complete the task depends on how much spare capacity they have, which will in turn depend on the nature of the linguistic processing going on at the same time. So for instance in the **click location** task participants listen to an utterance that has had a non-linguistic click or beep sound added to it. Their task is to listen to and understand the utterance and at the same time to mark on a transcript of the utterance where they think the click occurred. Results from such tasks show 'migration' of the click to

Table 1.1 Approximate timing of components involved in picture naming (based on Indefrey & Levelt, 2004).	
Conceptual preparation	175 msec
Word retrieval	75 msec
Phonological code retrieval	80 msec
Syllabification	25 msec per phoneme
Phonetic encoding	145 msec

locations other than where it actually occurred. This migration is at least in part a consequence of the fact that participants' attention was devoted to linguistic processing. Similar effects are found in **dichotic switch monitoring**. In this task participants have to indicate when the speech they are listening to moves from one headphone channel to the other. They do not always do this accurately because of the linguistic processing load at the point of the switch. (See Chapters 10 and 11 for examples of such tasks and how they contribute to our understanding of sentence comprehension.)

Other response time tasks use reaction times as a measure of the processing speed for the item being attended to. In its simplest form, perhaps, the response time task requires participants to press a button as quickly as possible when they see or hear a word. There may be a choice of buttons, one for a positive response when the word is an actual word of the language and the other for a negative response, when the word is not an actual word. Chapters 8 and 9 discuss how the speed and accuracy of participants' responses depend on factors such as familiarity – well-known words result in faster and more accurate responses than less well-known words.

In another button-pressing task, a participant reads a sentence on-screen, with presentation of each successive word controlled by the participant's pressing of the button. Chapters 10 and 11 give illustrations of how difficulty in interpreting the sentence at any point is reflected in the time taken to press the button to see the next word.

Carrying out a response task involves many components, some of which are assumed to be relatively invariant across different stimuli (e.g. the time it takes to send a motor command from the brain to the hand to execute a button press). One study (Indefrey & Levelt, 2004) broke down the supposedly simple task of saying a word that names a pictured object into the components listed in Table 1.1, starting with a mapping from the picture onto a mental concept and finishing with the initiation of the spoken word. Details will be given in the following chapters, but note for now that the most obvious influence on the time taken to initiate the response is the length of the word, since the stage labelled syllabification in the table has a duration that depends on the number of phonemes, or speech sounds, in the word. But in addition, it is important to remember that the other stages will have variable durations, depending for example on the accessibility of the word, which will vary with factors such as how often or how recently that word has been used. These are the factors that

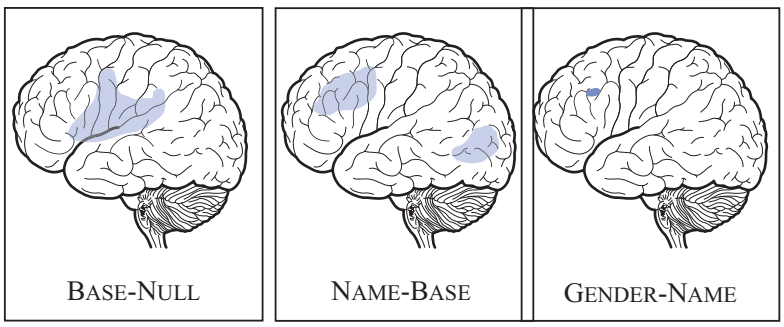

Figure 1.3
Schematic representation of differences in brain activity during different tasks (based on fMRI data published by Heim *et al.*, 2002; see text for explanation).

are often explored in experiments, and which require other aspects of the experiment to be held constant.

Language and the brain

It is becoming more and more usual to find studies which include monitoring of the patterns of blood flow and/or electrical activity associated with neural stimulation in the brain while a participant is carrying out some language production or comprehension task. These measures have a clear advantage over behavioural methods such as response button pressing, since they provide a more direct measure of psycholinguistic processing. Imaging techniques that measure blood flow give an indication of the areas of the brain involved in the tasks being carried out (and of brain areas that have been damaged after strokes or accidents). However, they typically have relatively poor temporal resolution, and so are not ideally suited to tracking the time course of the very fast processes involved in language processing. On the other hand, techniques that measure electrical activity provide better temporal resolution but poor spatial resolution. They are therefore better suited to the measurement of the time course of processing than to the precise determination of which brain areas are responsible for various tasks.

As an example of what we can learn from **neurophysiological** techniques, Figure 1.3 gives a very approximate depiction of differences in brain activity levels measured in a range of tasks carried out by German participants (Heim, Opitz & Friederici, 2002). Brain activity was measured using functional Magnetic Resonance Imaging (**fMRI**), which tracks changes in blood flow (see sidebar). An important aspect of German that is being examined here is the fact that nouns have grammatical gender which is marked through the forms of articles (e.g. the word for *the*). For instance the word for *cat* has feminine gender, so *the cat* is 'die Katze', while the word for *dog* has masculine gender and *the dog* is 'der Hund'.

The shaded shapes in each panel show differences in brain activity when two tasks are compared. The first panel shows areas with an increase in activity in a simple baseline speaking task (BASE), compared to when the participants are at rest (NULL). The baseline task is simply to say 'ja ja' ('yes yes') when they see a simple visual stimulus. The second panel shows areas that are active in a picture naming task (NAME), compared with the

fMRI – functional Magnetic Resonance Imaging

Neurons that are active consume more oxygen from their blood supply than those that are inactive. Oxygenated and deoxygenated blood have different magnetic properties, and though these differences are small they are sufficient to show up in a scanner that generates high magnetic fields. The participant lies in a scanner and changes in neural activity in different brain areas are measured while they perform some language task.

baseline task. In a naming task the participants are not just speaking but are also retrieving a word from memory. For instance, if they see a picture of a cat, then they have to retrieve and produce the word that names that picture, i.e. 'Katze' (*cat*). Clearly additional areas of the brain are involved in this task when compared with the baseline task. Finally, the third panel shows a small additional area (when compared with the naming task) that is active when the participant produces the appropriate definite article. As this is marked for the grammatical gender of the noun in German, the correct form has to be selected, and this seems to involve this small area of the brain (hence the label GENDER in the figure).

fMRI, as used to obtain the data above, is one of the imaging techniques that measures blood flow in the brain. Other types include **PET scans** (positron emission tomography), which measure emissions from radioactive substances injected into the bloodstream, and **CAT scans** (computed axial tomography), which are based on a series of X-rays taken from different angles.

Imaging techniques measuring electrical activity in the brain include **MEG** (magnetoencephalography) and **EEG** (electroencephalography). The former measures the magnetic fields created by naturally occurring electrical activity in the brain, while the latter measures electrical activity on the scalp. An advantage of the MEG technique is that the magnetic fields it measures are not distorted by surrounding tissues, which can affect the electrical activity measured by EEG. These techniques measure event-related potentials (**ERP**), the electrical currents passing through the fluid that surrounds neurons as they respond to an event, or event-related magnetic fields (**ERF**) which measure currents over larger neuronal structures.

Despite the technological advances that now allow us to look in great detail at neurophysiological aspects of language production and comprehension, there are areas of psycholinguistics in which informal observation and relatively simple experimental techniques still provide appropriate means for learning more about the production, comprehension and representation of language. A clear advantage of such techniques for many researchers, especially those housed in linguistics departments in humanities faculties, is their low cost compared with the hugely expensive imaging techniques mentioned above. Where possible, such approaches have been used in the exercise material included in each chapter of this book.

Summary

This chapter has given a short overview of some of the subject matter of psycholinguistics. It has given some examples of the types of phenomena that psycholinguistics explores, as well as examples of how psycholinguists carry out that exploration.

Some of the areas of psycholinguistics that have been highlighted as of relevance to this book are: