Introduction

The major perspective we adopt in this book regards a language as a *cognitive* system which is part of any normal human being's mental or psychological structure. An alternative to which we shall also give some attention emphasises the *social* nature of language, for instance studying the relationships between social structure and different dialects or varieties of a language.

The cognitive view has been greatly influenced over the past five decades by the ideas of the American linguist and political commentator Noam Chomsky. The central proposal which guides Chomsky's approach to the study of language is that when we assert that Tom is a speaker of English, we are ascribing to Tom a certain mental structure. This structure is somehow represented in Tom's brain, so we are also implicitly saying that Tom's brain is in a certain state. If Clare is also a speaker of English, it is reasonable to suppose that Clare's linguistic cognitive system is *similar* to Tom's. By contrast, Jacques, a speaker of French, has a cognitive system which is *different* in important respects from those of Tom and Clare, and different again to that of Guo, a speaker of Chinese. This proposal raises four fundamental research questions:

(1)   What is the nature of the cognitive system which we identify with knowing a language?

(2)   How do we acquire such a system?

(3)   How is this system used in our production and comprehension of speech?

(4)   How is this system represented in the brain?

Pursuit of these questions defines four areas of enquiry: linguistics itself, developmental linguistics, psycholinguistics and neurolinguistics.

At the outset, it is important to be clear that an answer to question (1) is *logically* prior to answers to questions (2), (3) and (4); unless we have a view on the nature of the relevant cognitive system, it makes no sense to enquire into its acquisition, its use in production and comprehension and its representation in the brain.

Question (1), with its reference to a *cognitive* system, looks as if it ought to fall in the domain of the cognitive psychologist. However, the Chomskian approach maintains that we can formulate and evaluate proposals about the nature of the human mind by *doing linguistics*, and much of this book is intended to establish the plausibility of this view. In order to do linguistics, we usually rely on native speakers of a language who act as informants and provide us with data; and it is

1

with respect to such data that we test our hypotheses about native speakers' linguistic cognitive systems. Often, linguists, as native speakers of some language or other, rely on themselves as informants. Linguists (as opposed to psycholinguists, see below) do not conduct controlled experiments on large numbers of subjects under laboratory conditions. This is a major *methodological* difference between linguists and cognitive psychologists in their study of the human mind, and some critics might see it as making linguistics unscientific or subjective. However, it is important to point out that the data with which linguists work (supplied by themselves or by other native speakers) usually have such clear properties as to render controlled experimentation pointless. For instance, consider the examples in (5):

(5) a. The dog chased the cat
    b. *Cat the dog chased the

A native speaker of English will tell us that (5a) is a possible sentence of English but (5b) is not (the * is conventionally used to indicate this latter judgement). Of course, we could design experiments with large numbers of native speakers to establish the reliability of these claims, but there is no reason to believe that such experiments would be anything other than a colossal waste of time. Native speakers have vast amounts of data readily available to them, and it would be perverse for linguists not to take advantage of this. Notice that above we said that the data supplied by native speakers *usually* have very clear properties. When this is not the case (and an example will arise in our discussion of psycholinguistics below), we proceed with more caution, trying to understand the source of difficulty.

The logical priority of question (1) should not lead to the conclusion that we must have a *complete* answer to this question before considering our other questions. Although question (2) requires some view on the cognitive linguistic system, there is no reason why acquisition studies of small children should not themselves lead to modifications in this view. In such a case, pursuit of question (2) will be contributing towards answering question (1), and similar possibilities exist for (3) and (4). In practice, many linguists, developmental linguists, psycholinguists and neurolinguists are familiar with each other's work, and there is a constant interchange of ideas between those working on our four questions.

Our questions foster different approaches to linguistic issues, and in this introduction we shall first take a preliminary look at these. Having done this, we shall turn to the social perspective mentioned at the outset and offer some initial remarks on how this is pursued.

Linguistics

To begin to answer question (1), Chomsky identifies knowing a language with having a mentally represented **grammar**. This grammar constitutes the native speaker's **competence** in that language, and on this view, the key to

understanding what it means to know a language is to understand the nature of such a grammar. Competence is contrasted with **performance**, the perception and production of speech, the study of which falls under psycholinguistics (see below). Since this is a fundamental distinction that underlies a great deal of what we shall be discussing, it is worth trying to get a clear grasp of it as early as possible. Consider the situation of a native speaker of English who suffers a blow to the head and, as a consequence, loses the ability to speak, write, read and understand English. In fortunate cases, such a loss of ability can be short-lived, and the ability to use English in the familiar ways reappears quite rapidly. What cognitive functions are impaired during the time when there is no use of language? Obviously, the ability to use language, i.e. to perform in various ways, is not available through this period, but what about knowledge of English, i.e. linguistic competence? If we suppose that this is lost, then we would expect to see a long period corresponding to the initial acquisition of language as it is regained, rather than the rapid re-emergence which sometimes occurs. It makes more sense to suppose that knowledge of language remains intact throughout such an episode; the problem is one of accessing this knowledge and putting it to use in speaking, etc. As soon as this problem is overcome, full knowledge of English is available, and the various abilities are rapidly reinstated.

What does a grammar consist of? The traditional view is that a grammar tells us how to combine words to form phrases and sentences. For example, by combining a word like *to* with a word like *Paris* we form the phrase *to Paris*, which can be used as a reply to the question asked by speaker A in the dialogue below:

(6)   SPEAKER A: Where have you been?
             SPEAKER B: *To Paris*.

By combining the phrase *to Paris* with the word *flown* we form the larger phrase *flown to Paris*, which can serve as a reply to the question asked by speaker A in (7):

(7)   SPEAKER A: What's he done?
             SPEAKER B: *Flown to Paris*.

And by combining the phrase *flown to Paris* with words like *has* and *he*, we can form the sentence in (8):

(8)   He has flown to Paris

On this view, a grammar of a language specifies how to combine words to form phrases and sentences, and it seems entirely appropriate to suggest that native speakers of English and of other languages have access to cognitive systems which somehow specify these possibilities for combination (***exercise 1***). A very important aspect of this way of looking at things is that it enables us to make sense of how a cognitive system (necessarily *finite*, since it is represented in a brain) can somehow characterise an *infinite* set of objects (the phrases and sentences in a natural language). That natural languages are infinite in this sense is easy to see by considering examples such as those in (9):

(9) a.   Smith believes that the earth is flat
    b.   Brown believes that Smith believes that the earth is flat
    c.   Smith believes that Brown believes that Smith believes that the earth is flat
    d.   Brown believes that Smith believes that Brown believes that Smith believes that the earth is flat

A native speaker of English will recognise that such a sequence of sentences could be indefinitely extended, and the same point can be made in connection with a variety of other constructions in English and other languages (*exercise 2*). But the infinite nature of the set of English sentences, exemplified by those in (9), does not entail that the *principles of combination* used in constructing these sentences are also infinite; and it is these principles which form part of a grammar.

The view we have introduced above implies that a grammar contains two components: (i) a **lexicon** (or dictionary), which lists all the words found in the language, and (ii) a **syntactic component**, which specifies how to combine words together to form phrases and sentences. Each **lexical entry** (i.e. each item listed in the lexicon) will tell us about the linguistic properties of a word. For example, the entry for the word *man* will specify its **phonological** (= sound) properties (namely that it is pronounced /man/ – for the significance of the slashes, see section 5), its **grammatical** properties (e.g. that it can function as a noun and that when it does, it has the irregular plural form *men*) and its **semantic** (i.e. meaning) properties (namely that it denotes an adult male human being). The linguistic properties of words, including the nature of lexical entries, form the subject matter of part II of this book, while syntax (i.e. the study of how words are combined together to form phrases and sentences) provides the focus for part III. A grammar can be said to **generate** (i.e. specify how to form) a set of phrases and sentences, and using this terminology, we can view the task of the linguist as that of developing a theory of **generative grammar** (i.e. a theory about how phrases and sentences are formed).

Careful reflection shows that a grammar must contain more than just a lexicon and a syntax. One reason for this is based on the observation that many words change their **phonetic form** (i.e. the way they are pronounced) in connected speech, such sound changes being determined by the nature of neighbouring sounds within a word, phrase or sentence. These changes are effected by native speakers in a perfectly natural and unreflective way, suggesting that whatever principles determine them must be part of the relevant system of mental representation (i.e. grammar). We can illustrate what we mean here by considering examples of changes which result from the operation of regular **phonological processes**. One such process is **elision**, whereby a sound in a particular position can be dropped and hence not pronounced. For instance, the 'f' in the word *of* (which is pronounced /v/) can be elided in colloquial speech before a word beginning with a consonant (but not before a word beginning with a vowel): hence we say 'pint o' milk' (sometimes written *pinta milk*) eliding /v/ before the /m/ of the word *milk*, but 'pint of ale' (not 'pint o' ale') where the /v/ can't be elided because the word *ale* begins with a vowel. A second regular phonological

process is **assimilation**, a process by which one sound takes on some or all the characteristics of a neighbouring sound. For example, in colloquial speech styles, the final 'd' of a word like *bad* is assimilated to the initial sound of an immediately following word beginning with a consonant: hence, *bad boy* is pronounced as if it were written *bab boy* and *bad girl* as if it were written *bag girl* (*exercise 3*).

The fact that there are regular phonological processes such as those briefly described above suggests that in addition to a lexicon and a syntactic component, a grammar must also contain a **phonological component**: since this determines the phonetic form (= PF) of words in connected speech, it is also referred to as the **PF component**. **Phonology**, the study of sound systems and processes affecting the way words are pronounced, forms the subject matter of part I of this book.

So far, then, we have proposed that a grammar of a language contains three components, but it is easy to see that a fourth component must be added, as native speakers not only have the ability to *form* sentences, but also the ability to *interpret* (i.e. assign meaning to) them. Accordingly, a grammar of a language should also answer the question 'How are the meanings of sentences determined?' A commonsense answer would be that the meaning of a sentence is derived by combining the meanings of the words which it contains. However, there's clearly more involved than this, as we see from the fact that sentence (10) below is ambiguous (i.e. has more than one interpretation):

(10)   She loves me more than you

Specifically, (10) has the two interpretations paraphrased in (11a, b):

(11) a.   She loves me more than you love me
      b.   She loves me more than she loves you

The ambiguity in (10) is not due to the meanings of the individual words in the sentence. In this respect, it contrasts with (12):

(12)   He has lost the match

In (12), the word *match* is itself ambiguous, referring either to a sporting encounter or a small piece of wood tipped with easily ignitable material, and this observation is sufficient to account for the fact that (12) also has two interpretations. But (10) contains no such ambiguous word, and to understand the ambiguity here, we need to have some way of representing the logical (i.e. meaning) relations between the words in the sentence. The ambiguity of (10) resides in the relationship between the words *you* and *loves*; to get the interpretation in (11a), *you* must be seen as the **logical subject** of *loves* (representing the person giving love), whereas for (11b), it must function as the **logical object** of *loves* (representing the person receiving love). On the basis of such observations, we can say that a grammar must also contain a component which determines the **logical form** (= LF) of sentences in the language. For obvious reasons, this component is referred to as the **LF component**, and this is a topic which is discussed in section 23 of this book (*exercise 4*).

Our discussion has led us to the conclusion that a grammar of a language comprises (at least) four components: a lexicon, a syntactic component, a PF component and an LF component. A major task for the linguist is to discover the nature of such grammars.

However, there is an additional concern for the linguist. Suppose grammars are produced for a variety of languages by specifying the components introduced above. Naturally, we would expect these grammars to exhibit certain differences (a grammar of English will be different to a grammar of Japanese), but we might also discover that they have some properties in common. If these properties appear in grammars for a wide range of languages, standard scientific practice leads us to hypothesise that they are common to the grammars of *all* natural languages, and this means that an additional goal for the linguist is the development of a theory of **Universal Grammar (UG)**. A great deal of contemporary linguistic theory can be viewed as testing hypotheses about UG on an ever-wider class of languages.

As described above, UG is viewed as emerging from the linguist's study of individual grammars, but there is a different way to introduce this concept which affords it a much more important and fundamental position in the work of linguists. To appreciate this, we need to turn to the second of our questions, namely, 'How do we acquire a grammar?'

Developmental linguistics

Readers familiar with small children will know that they generally produce their first recognisable word (e.g. *Dada* or *Mama*) round about their first birthday; from then until the age of about one year, six months, children's speech consists largely of single words spoken in isolation (e.g. a child wanting an apple will typically say 'Apple'). At this point, children start to form elementary phrases and sentences, so that a child wanting an apple at this stage might say 'Want apple'. From then on, we see a rapid growth in children's grammatical develop-ment, so that by the age of two years, six months, most children are able to produce adult-like sentences such as 'Can I have an apple?'

From this rough characterisation of development, a number of tasks emerge for the developmental linguist. Firstly, it is necessary to *describe* the child's devel-opment in terms of a sequence of grammars. After all, we know that children become adults, and we are supposing that, as adults, they are native speakers who have access to a mentally represented grammar. The natural assumption is that they move towards this grammar through a sequence of 'incomplete' or 'imma-ture' grammars. Secondly, it is important to try to *explain* how it is that after a period of a year and a half in which there is no obvious sign of children being able to form sentences, between one-and-a-half and two-and-a-half years of age there is a 'spurt' as children start to form more and more complex sentences, and a phenomenal growth in children's grammatical development. This uniformity and (once the 'spurt' has started) rapidity in the pattern of children's linguistic

development are central facts which a theory of language acquisition must seek to explain. But how?

Chomsky maintains that the most plausible explanation for the uniformity and rapidity of first language acquisition is to posit that the course of acquisition is determined by a biologically endowed innate **language faculty** (or *language acquisition program*, to borrow a computer software metaphor) within the human brain. This provides children with a genetically transmitted set of procedures for developing a grammar which enables them to produce and understand sentences in the language they are acquiring on the basis of their *linguistic experience* (i.e. on the basis of the speech input they receive). The way in which Chomsky visualises the acquisition process can be represented schematically as in (13) below (where L is the language being acquired):

(13)

experience → language → grammar
of L   faculty   of L
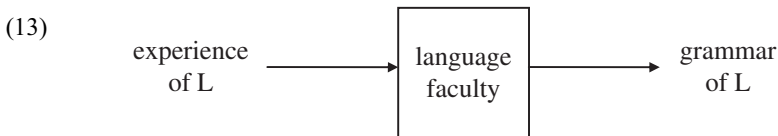
Children acquiring a language will observe people around them using the language, and the set of expressions in the language which the child hears (and the contexts in which they are used) in the course of acquiring the language constitute the child's linguistic experience of the language. This experience serves as input to the child's language faculty, which provides the child with a set of procedures for analysing the experience in such a way as to devise a grammar of the language being acquired. Chomsky's hypothesis that the course of language acquisition is determined by an innate language faculty is known popularly as the **innateness hypothesis**.

Invocation of an innate language faculty becoming available to the child only at some genetically determined point may constitute a plausible approach to the questions of uniformity and rapidity, but there is an additional observation which suggests that some version of the innateness hypothesis *must be correct*. This is that the knowledge of a language represented by an adult grammar appears *to go beyond anything supplied by the child's linguistic experience*. A simple demonstration of this is provided by the fact that adult native speakers are not only capable of combining words and phrases in acceptable ways but also of recognising unacceptable combinations (see 5b above and exercise 1). The interesting question this raises is: where does this ability come from? An obvious answer to this question is: that the child's linguistic experience provides information on unacceptable combinations of words and phrases. But this is incorrect. Why do we assert this with such confidence?

Obviously, when people speak, they do make mistakes (although research has shown that language addressed to children is *almost* completely free of such mistakes). However, when this happens, there is no clear signal to the child indicating that an adult utterance contains a mistake, that is, as far as the child is

concerned, an utterance containing a mistake is just another piece of linguistic experience to be treated on a par with error-free utterances. Furthermore, it has been shown that adults' 'corrections' of children's own speech do not take systematic account of whether children are producing syntactically acceptable or unacceptable combinations of words and phrases; parents do 'correct' their children, but when they do this, it is to ensure that children speak *truthfully*; grammatical correctness is not their target. Overall, there is compelling evidence that children do *not* receive systematic exposure to information about unacceptable sequences, and it follows that in this respect the child's linguistic experience is *not sufficient* to justify the adult grammar. From this **poverty of the stimulus** argument it follows that something must supplement linguistic experience and the innate language faculty fulfils this role (***exercise 5***).
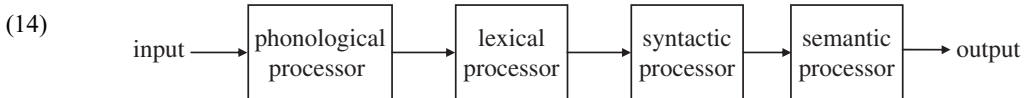
Now, it is important to underline the fact that children have the ability to acquire *any* natural language, given appropriate experience of the language: for example, a British child born of monolingual English-speaking parents and brought up by monolingual Japanese-speaking parents in a Japanese-speaking community will acquire Japanese as a native language. From this it follows that the contents of the language faculty *must not be specific to any one human language*: if the language faculty accounts for the uniformity and rapidity of the acquisition of English, it must also account for the uniformity and rapidity of the acquisition of Japanese, Russian, Swahili, etc.; and if the language faculty makes up for the insufficiency of a child's experience of English in acquiring a grammar of English, it must also make up for the insufficiency of a child's experience of Japanese in acquiring a grammar of Japanese, for the insufficiency of a child's experience of Russian in acquiring a grammar of Russian, for the insufficiency of a child's experience of Swahili in acquiring a grammar of Swahili, etc. This entails, then, that the language faculty must incorporate a set of **UG principles** (i.e. principles of Universal Grammar) which enable the child to form and interpret sentences in any natural language. Thus, we see an important convergence of the interests of the linguist and the developmental linguist, with the former seeking to formulate UG principles on the basis of the detailed study of the grammars of adult languages and the latter aiming to uncover such principles by examining children's grammars and the conditions under which they emerge.

In the previous paragraph, we have preceded 'language' with the modifier 'human', and genetic transmission suggests that a similar modifier is appropriate for 'language faculty'. The language faculty is *species-specific* and the ability to develop a grammar of a language is *unique to human beings*. This ability distinguishes us from even our nearest primate cousins, the great apes such as chimpanzees and gorillas, and in studying it we are therefore focusing attention on one of the defining characteristics of what it means to be a human being. There have been numerous attempts to teach language to other species, and success in this area would seriously challenge the assertion we have just made. Indeed, it has proved possible to teach chimpanzees a number of signs similar to those employed in the Sign Languages used as native languages by the deaf, and it has been

reported that pigmy chimpanzees can understand some words of spoken English, and even follow a number of simple commands. Such research arouses strong emotions, and, of course, we are not in a position to assert that it will *never* produce dramatic results. At the moment, however, we can maintain that all attempts, however intensive, to teach grammatical knowledge to apes have been spectacular failures when the apes' accomplishments are set alongside those of a normal three-year-old child. As things stand, the evidence is firmly in favour of the species-specificity of the language faculty.

Psycholinguistics

As noted above, the psycholinguist addresses the question of how the mentally represented grammar (linguistic competence) is employed in the production and comprehension of speech (linguistic performance). The most direct way to approach this relationship is to adopt the hypothesis that a generative grammar can simply be regarded as itself providing an account of how we understand and produce sentences in real time. From the point of view of language comprehension, this gives rise to the following (highly simplified) model, where the input is a stretch of spoken or written language such as a particular sentence:

(14)

input ⟶ [phonological processor] ⟶ [lexical processor] ⟶ [syntactic processor] ⟶ [semantic processor] ⟶ output

In terms of this rather crude model, the first step in language comprehension is to use the *phonological processor* to identify the sounds (or written symbols) occurring in the input. Then, the *lexical processor* identifies the component words. The next step is for the *syntactic processor* (also called the *parser*, and incorporating the syntactic component of the grammar) to provide a syntactic representation of the sentence (i.e. a representation of how the sentence is structured out of phrases and the phrases out of words). The last step is for the semantic processor to compute a meaning representation for the sentence, on the basis of the syntactic and lexical information supplied by earlier stages in the process. The relevant meaning representation serves as the output of the model: once this has been computed, we have understood the sentence.

An important characteristic of (14), as of all models of psycholinguistic processing, is that its various stages are to be viewed as taking place in real time, and a consequence of this is that psycholinguists can utilise their experimental techniques to try to measure the duration of specific parts of the process and link these measurements to levels of complexity as defined by the grammar itself. In fact, it is fairly easy to see that the idea that the grammar can, without any additional

considerations, serve as a model of sentence comprehension is implausible. A sentence such as (15) is known as a **garden-path sentence**:

(15)   The soldiers marched across the parade ground are a disgrace

A common reaction to (15) from native speakers of English is that it is *not* an acceptable sentence. However, this reaction can often be modified by asking native speakers to consider the sentences in (16) (recall our observation that not all linguistic data have immediately obvious properties):

(16) a.   The soldiers who were driven across the parade ground are a disgrace
     b.   The soldiers driven across the parade ground are a disgrace
     c.   The soldiers who were marched across the parade ground are a disgrace

Sentence (16a) should be regarded as entirely straightforward, and we can view (16b) as 'derived' from it by deleting the sequence of words *who were*. Now, if we delete *who were* from sentence (16c), which should also be recognised as an acceptable English sentence, we 'derive' (15), and at this point many readers are likely to change their reaction to (15): it *is* an acceptable English sentence, so long as it is interpreted with the phrase *the soldiers* as the logical object of *marched* (see p. 5 above). When we read (15) for the first time, we immediately interpret *the soldiers* as the logical subject of *marched* – the soldiers are marching rather than being marched; as a consequence, the sequence *the soldiers marched across the parade ground* is interpreted as a complete sentence and the sentence processor doesn't know what to do with *are a disgrace*. The sentence processor has been 'garden-pathed', i.e. sent down the wrong analysis route (***exercise 6***).

What is important about garden-path sentences is that they show that sentence comprehension *must* involve something in addition to the grammar. As far as the grammar is concerned, (15) is an acceptable structure with only one interpretation. However, it appears that this structure and interpretation are not readily available in sentence processing, suggesting that the parser must rely (to its detriment in this case) on something beyond the principles which determine acceptable combinations of words and phrases.

There are other aspects of (14) which are controversial and have given rise to large numbers of experimental psycholinguistic studies. For instance, there is no place in (14) for *non-linguistic general knowledge about the world*; according to (14), interpretations are computed entirely on the basis of linguistic properties of expressions without taking any account of their plausibility, and an alternative would allow encyclopaedic general knowledge to 'penetrate' sentence perception and guide it to more likely interpretations. A further assumption in (14) is that the different sub-components are *serially ordered* (in that the first stage is phonological processing which does its job before handing on to lexical processing, etc.) An alternative would allow syntactic and semantic factors to influence phonological and lexical processing, for semantic factors to influence syntactic processing, etc. These issues, along with several others, will be discussed in sections 14 and 26.