1 Grammar

1.1 Overview

In broad terms, this book is concerned with **syntax** in general, and with the syntax of English in particular. Syntax is one of the two key areas of what is traditionally called 'grammar' – the other being **morphology**. Morphology is the study of how words are formed out of smaller units (called **morphemes**) and so addresses questions such as 'What are the component morphemes of a word like *antidisestablishmentarianism*, and what is the nature of the operations by which they are combined together to form the overall word?' Syntax is the study of the way in which phrases and sentences are structured out of words, and so addresses questions like 'What is the structure of a sentence like *Where's the president going?* and what is the nature of the grammatical operations by which its component words are combined together to form the overall sentence structure?' This chapter begins (in §1.2) by looking at a range of approaches to the study of grammar, before going on (in §1.3) to look at how syntax was studied in **traditional grammar**: this also provides an opportunity to introduce some useful grammatical terminology. In the remainder of the chapter, we look at the approach to syntax adopted within the theory of **Universal Grammar/UG** developed by Chomsky over the past six decades. (Note that a convention used throughout the book is that key technical terms are in bold print when first introduced in a given chapter; such terms are generally given an entry in the Glossary at the end of the book if they are used in several different sections of the book, though not if they occur in only one section of the book and are glossed there.)

1.2 Approaches to grammar

A fundamental question that needs to be resolved at the outset concerns what kind of approach to adopt in studying grammar. There are two diametrically opposed answers to this question found in work on grammar. One sees the role of grammar as being essentially **prescriptive** (i.e. prescribing norms for grammatical correctness, linguistic purity and literary excellence); the other sees the role of grammar as being inherently **descriptive** (i.e describing the way people speak or write their native language). We can illustrate the differences between these two approaches in relation to the following dialogue between the

fictional Oxford detective Morse and his assistant Lewis, as they are looking at a dead body (where capital letters in the dialogue mark emphatic stress, and italics mark items of grammatical interest):

(1) MORSE: I think he was murdered, Lewis
LEWIS: *Who by*, sir?
MORSE: *By whom*, Lewis, *by whom*. Didn't they teach grammar at that COMPREHENSIVE school of yours?

Morse was educated at a *grammar* school – i.e. an elitist school which sought to give pupils a 'proper education' and taught them grammar, so that they could learn to speak and write 'properly' (i.e. in a prestigious form of Standard English). Lewis, by contrast, was educated at a *comprehensive* school – i.e. a more socially inclusive type of school which admitted pupils from a much broader social spectrum and didn't force-feed them with grammar. The linguistic skirmish between Lewis and Morse in (1) revolves round the grammar of an italicised phrase which comprises the preposition *by* and the pronoun *who(m)*. The differences between what the two men say relates to (i) the form of the pronoun (*who* or *whom*?), and (ii) the position of the pronoun (before or after the preposition *by*?). Lewis uses the pronoun form *who*, and positions the pronoun before the preposition when he asks 'Who by?' Morse corrects Lewis and instead uses the pronoun form *whom* and positions the pronoun after the preposition when he says 'By whom?' But why does Morse correct Lewis? The answer is that Morse was taught traditional prescriptive grammatical rules at his grammar school, including two which can be outlined informally as follows:

(2) (i) The form *who* is used as a subject of a finite verb, and *whom* as the object of a verb or preposition.
(ii) Never end a phrase, clause or sentence with a preposition.

When Lewis asks 'Who by?' he violates both rules. This is because the pronoun *who(m)* is the object of the preposition *by* and rule (2i) stipulates that *whom* must therefore be used, and rule (2ii) specifies that the preposition should not be positioned at the end of a phrase. The corrected form 'By whom?' produced by Morse obeys both rules, in that *whom* is used in conformity with rule (2i), and *by* is positioned in front of its object *whom* in conformity with rule (2ii).

The more general question raised by the dialogue in (1) is the following. When studying grammar, should we adopt a descriptive approach and *describe* what ordinary people like Lewis actually say, or should we adopt a prescriptive approach and *prescribe* what people like Morse think they ought to say? There are several reasons for rejecting the prescriptive approach. For one thing, it is elitist and socially divisive, in that a privileged elite attempts to lay down grammatical norms and impose them on everyone else in society. Secondly, the grammatical norms which prescriptivists seek to impose are often derived from structures found in 'dead' languages like Latin, which is somehow regarded as a model of grammatical precision and linguistic purity: and so, because Latin made

a distinction between subject and object forms of pronouns, English must do so as well; and because Latin (generally) positioned prepositions before their objects, English must do so as well. Such an approach fails to recognise typological diversity in languages – i.e. that there are many different types of structure found in the world's 8,000 or so known languages. Thirdly, the prescriptive approach fails to recognise sociolinguistic variation – i.e. that different types of structure are found in different styles and varieties of English (e.g. *Who by?* is used in colloquial English, and *By whom?* in formal styles of English). And fourthly, the prescriptive approach also fails to recognise linguistic change – i.e. that languages are constantly evolving, and that structures used centuries ago may no longer be in use today (e.g. *whom* is an archaic form which has largely dropped out of use and is no longer part of the grammar of teenagers today). For reasons such as these, the approach taken to grammar in work over the past seventy years or so has been descriptive.

What this means is that in attempting to devise a grammar of (e.g.) English, we aim to describe the range of grammatical structures found in present-day English. But how do we determine what is or isn't grammatical in present-day English? One approach is to study **usage** (i.e. the range of structures used by people when they speak or write). Linguists who adopt this kind of approach rely on data from a **corpus** (e.g. a computerised database such as the British National Corpus) containing authentic examples of spoken or written English. Such corpora offer the advantage that they contain millions of sentences, and the sentences have usually been codified/tagged by a team of researchers, so simplifying the task of searching for examples of a particular construction. Some linguists treat the Web as a form of corpus, and use a search engine to find examples from the Internet of the kinds of structures they are interested in.

The usefulness of corpora can be illustrated as follows. One of the ways I collect data on spoken English is by listening to live, unscripted radio and TV broadcasts, and noting down unusual structures (often using them as exercise material in my syntax books). An interesting sentence which I heard one day (reported in Radford 2004: 429; 8a) is the following:

(3) What is thought has happened to him? (Interviewer, BBC Radio 5 Live)

When I first heard the sentence in (3), I wasn't sure what to make of it. One possibility that occurred to me is that it might be an accidental speech error (perhaps induced by the pressure of live broadcasting), representing a **blend** of the two different structures in (4), formed by combining the italicised part of (4a) with the italicised part of (4b):

(4) (a) *What is thought* to have happened to him?
(b) What is it thought *has happened to him*?

However, an alternative possibility is that the kind of structure in (3) is not a speech error but rather a productive structure – albeit not one described in standard grammars of English. To check on whether (3) is a productive structure

or not, I searched for similar structures on the Internet, and found hundreds of them. For example, I came across 116 examples of sentences containing the string (i.e. sequence of words) *is/are thought may*, like those below:

(5) (a) The toxicology issue is thought may have arisen because of a pre-existing health issue in the animals
(b) Police are investigating the cause of the blazes, which is thought may be arson
(c) Curiously, about one-third of adults 60 and over are thought may have antibodies that may help protect against the virus
(d) The user enters one or more search words which are thought may exist in the definition of the word sought

The fact that I was able to locate thousands of similar examples of the structure in (3) makes it more likely that (3) is a (hitherto unreported) grammatical structure in English, and not a one-off 'slip of the tongue'.

What the discussion here illustrates is that usage data (from corpora or from the Web) provide a very useful source of information about the productivity of a given type of structure (i.e. how often it is used). However, there are also downsides to the usage-based approach. For one thing, a corpus may contain relatively few examples of low-frequency structures. Secondly, it is generally not possible to ask the speakers who produced them questions about the sentences in the corpus (e.g. 'How would you negate this sentence?'). Thirdly, a corpus may contain examples of production errors (slips of the tongue, or pen, or keyboard) which would probably be judged as unacceptable even by the people who produced them. And (in the case of internet examples), it is sometimes unclear whether someone producing a given sentence (who may use an identity-concealing pseudonym like *CutiePie* or *MasterBlaster* as their name) is a native speaker of English or not, and if so what variety/dialect of English they speak.

A very different approach to studying grammar has been adopted by Noam Chomsky and his followers in work over the past sixty years. For Chomsky, the goal of studying the grammar of a language is to determine what it is that native speakers *know* about the grammar of their native language which enables them to speak and understand the language: hence, in studying language, we are studying a specific kind of cognition (i.e. human knowledge). In a fairly obvious sense, any native speaker of a language can be said to *know* the grammar of his or her native language. For example, any native speaker of English can tell you that the negative counterpart of *I like syntax* is *I don't like syntax*, and not, e.g., **I no like syntax.* (Note that a prefixed star/asterisk in front of a phrase or sentence indicates that it is ungrammatical.) In other words, native speakers know how to form phrases and sentences in their native language. Likewise, any native speaker of English can tell you that a sentence like *She loves me more than you* is ambiguous and has two **interpretations** which can be paraphrased as 'She loves me more than she loves you' and 'She loves me more than you love me': in other words, native speakers also know how to **interpret** (i.e. assign meaning

to) expressions in their language. However, it is important to emphasise that this grammatical knowledge of how to form and interpret expressions in your native language is **tacit** (i.e. subconscious) rather than **explicit** (i.e. conscious): so, it's no good asking a native speaker of English a question such as 'How do you form negative sentences in English?' since human beings have no conscious awareness of the processes involved in speaking and understanding their native language. To introduce a technical term devised by Chomsky, we can say that native speakers have grammatical **competence** in their native language: by this, we mean that they have tacit knowledge of the grammar of their language – i.e. of how to form and interpret words, phrases and sentences in the language.

In work in the 1960s, Chomsky drew a distinction between competence (the native speaker's tacit knowledge of his or her language) and **performance** (what people actually say or understand by what someone else says on a given occasion). Competence is 'the speaker-hearer's knowledge of his language', while performance is 'the actual use of language in concrete situations' (Chomsky 1965: 4). Very often, performance is an imperfect reflection of competence: we all make occasional slips of the tongue, or occasionally misinterpret something which someone else says to us. However, this doesn't mean that we don't know our native language or that we don't have *competence* in it. Misproductions and misinterpretations are **performance errors**, attributable to a variety of performance factors like tiredness, boredom, drunkenness, drugs, external distractions and so forth. A grammar of a language tells you what you need to know in order to have native-like competence in the language (i.e. to be able to speak the language like a fluent native speaker): hence, it is clear that grammar is concerned with competence rather than performance. This is not to deny the interest of performance as a field of study, but merely to assert that performance is more properly studied within the different – though related – discipline of psycholinguistics, which studies the psychological processes underlying speech production and comprehension. (It should, however, be acknowledged that performance errors can provide us with clues about the nature of competence, and we will see some examples of this in later chapters.)

When we study grammatical competence, we're studying a cognitive system internalised within the brain/mind of native speakers which is the product of a 'cognitive organ' which is 'shared among humans and in crucial respects unique to them' (Chomsky 2007: 1). In the terminology adopted by Chomsky (1986a: 19–56), our ultimate goal in studying competence is to characterise the nature of this 'internal language' or **I-language** (to use a term employed by Chomsky) which makes native speakers proficient in their native language.

Although native speakers only have tacit knowledge of the grammar of their language, they do have **intuitions** about grammaticality (i.e. 'gut feelings' about whether a particular sentence is or isn't grammatical in their native language) – e.g. as noted above, any native speaker of English would readily accept *I don't like syntax* as a grammatical sentence of English, but not **I no like syntax*. Consequently, an approach widely used by linguists over the past sixty years

has been to devise grammars on the basis of native-speaker intuitions about grammaticality. Where linguists are describing aspects of their own native language, they often rely primarily on their own intuitions/introspective grammaticality judgments.

However, although extensively used, this approach of relying on introspective judgments about the grammaticality of sentences has been criticised by some as being unscientific (hence yielding potentially unreliable results), particularly in relation to judgments about **marginal** sentences – i.e. sentences of uncertain grammaticality. This is for a number of reasons. Firstly, different individuals may disagree in their judgments of particular sentences (and may have different tolerance thresholds): this means that relying on the intuitions of one person alone may give misleading results. Secondly, the same individual may sometimes give conflicting judgments about the same sentence on different occasions. Thirdly, it can sometimes be very difficult to judge the grammaticality of a sentence in isolation (without an appropriate context). Fourthly (as we will see below), grammaticality is often a matter of degree rather than an absolute property (e.g. a given sentence may be more acceptable than some sentences but less acceptable than others). Fifthly, native speakers who are non-linguists very often have no clear idea what it means for a sentence to be 'grammatical' or not (since **grammaticality** is a technical term which non-linguists have little conception of). Rather, all that non-experts can do is say how acceptable they find a sentence, and this may depend on a range of factors which have little to do with grammaticality, including how frequent a given structure is, whether it contains taboo language or concepts and so on. And sixthly, linguists who rely on their own grammaticality judgments tend to give more extreme judgments than non-linguists and are vulnerable to the accusation that (however unwittingly) they may tailor their grammaticality judgments to fit their analysis (e.g. they may judge a given sentence to be grammatical because their analysis predicts that it should be): indeed, Chomsky (1957: 14) even suggested that (as a matter of principle) we should 'let the grammar itself decide' about the status of marginal sentences.

Because of the potential unreliability of informal intuitions, some linguists prefer to elicit native-speaker judgments experimentally, particularly when dealing with 'marginal' structures whose grammaticality status is not clear-cut. Such experimental studies sometimes produce judgments which are at variance with the intuitions of linguists. By way of illustration, consider the claim made by Chomsky (2008) that an (italicised) constituent can be **extracted** out of a passive subject like that bracketed in (6a) below (i.e. moved from the gap position marked by - - - to the italicised position), but not out of the corresponding active subject in (6b):

(6) (a) *Of which car* was [the driver - - -] awarded a prize?
(b) **Of which car* did [the driver - - -] cause a scandal?

In traditional terms, a sentence like that in (6a) is said to be in the **passive voice**, while that in (6b) is in the **active voice**. An assumption made by Chomsky is that

the italicised phrase in (6a,b) originates in the gap (- - -) position within the bracketed subject and is subsequently extracted out of the bracketed subject and moved to the front of the overall sentence. Chomsky develops an elaborate account of why extraction is possible out of a passive subject in sentences like (6a) but not out of an active subject in sentences like (6b). However, the robustness of Chomsky's syntactic analysis ultimately depends on the robustness of his judgment that (6a) is grammatical and (6b) is ungrammatical. But is this judgment shared by others?

To test this, Jurka (2010) ran an online experiment in which thirty-seven native speakers of English were asked to judge the acceptability of a number of sentences, including those in (7) below, where (7a) involves **extraction** of an italicised phrase out of a bracketed passive subject, and (7b) out of a bracketed active subject:

(7) (a) John wondered *which man* [a book about - - -] was released last year
(b) John wondered *which man* [a book about - - -] caused a scandal last year

Participants were asked to rate the tested sentences on a 7-point scale, and to give a score of 6 or 7 to sentences they found perfectly acceptable, 3–5 to sentences they found not totally unacceptable but also not completely perfect, and 1 or 2 to sentences they found completely unacceptable. Jurka found that extraction from a passive subject in sentences like (7a) yielded a mean score of 2.68, and that this was not significantly different statistically from the score of 2.55 for extraction out of the corresponding active subject in sentences like (7b). The moral of the story would seem to be that the grammaticality judgments of individual linguists need to be treated with caution in cases which are not clear-cut.

At the same time, however, it should be acknowledged that there are a number of drawbacks to experimental studies. For one thing, they require considerable time and money to set up: it can take months to design an experiment, collect the data, and process the results; and a design flaw (or problematic results) may require the whole experiment to be re-designed and re-run subsequently. Moreover, it is in the nature of experiments that (in order to meet stringent methodological requirements on experimental design) they can only be used to collect data relating to a specific (and narrow) set of phenomena. Furthermore, experiments can sometimes produce results which are skewed by the design of the experiment. In addition, how acceptable (or otherwise) people perceive a sentence to be may depend on a whole range of extraneous factors other than its grammaticality: these extraneous factors include, for example, how interesting it is, how long it is, how plausible it is, how frequent the relevant type of structure is, how easy it is to imagine a context where it could be used, whether or not the sentence expresses ideas which offend cultural or religious sensibilities or contains taboo words, etc. Furthermore, the results which experiments yield can be far from straightforward to interpret: for example, they sometimes produce results which represent acceptability in terms of many different shades of grey, rather than as a black-and-white issue. Moreover, in order to achieve statistical

significance in results, it may be necessary to discard outliers (i.e. atypical results).

As a case in point, consider the following data reported in Radford and Iwasaki (2015), in a study of the syntax of structures like *Who by?* in sentences like (1). This type of structure is traditionally considered to be grammatical in colloquial English only if the object of the preposition is interrogative (e.g. if it is a question word like *who?*). In order to check whether it is indeed the case that this kind of structure is only acceptable when the object of the preposition is interrogative, we asked Philip Hofmeister to run an online experiment to test the acceptability of the sentences below, each containing a bracketed structure in which the italicised object of a bold-printed preposition is positioned in front of the preposition:

(8) (a) I wonder where she bought that awful tie, and [*who* **for**]
(b) Whenever we argue and [*whatever* **about]**, we always make up afterwards
(c) I'm amazed at how much he bought on eBay and [*how little* **for]**!
(d) I'm going away, but [*not long* **for**]
(e) The fewer presents we send and [*the fewer people* **to]**, the happier Scrooge will be
(f) So hard has he trained and [*so long* **for]** that he is sure to win the race

The italicised object of the preposition is interrogative in (8a), unconditional (in the terminology of Rawlins 2008) in (8b), exclamative in (8c), negative in (8d), comparative in (8e) and consecutive in (8f). Sixty-four native speakers of American English were asked to rate the acceptability of a set of sentences including those in (8) using a 7-point numerical scale on which 7 denotes 'extremely natural' and 1 denotes 'extremely unnatural'. The mean scores for each sentence were as follows: 5.13 for the interrogative (8a); 4.95 for the unconditional (8b); 4.50 for the exclamative (8c); 4.27 for the negative (8d); 3.92 for the comparative (8e) and 2.69 for the consecutive (8f). The results are problematic for the traditional black-and-white view that 'inverted' structures of the form OBJECT+PREPOSITION are grammatical if the object is interrogative, and ungrammatical if it is not. They suggest that we have to recognise different degrees of grammaticality.

The general conclusions to be drawn from our discussion in this section are the following. Contemporary work in grammar is descriptive in orientation rather than prescriptive, so that (e.g.) a grammar of contemporary English seeks to describe the structures found in present-day (spoken and written) English. A grammar is said to be **descriptively adequate** (or achieve **descriptive adequacy**) if it provides a comprehensive description of the full range of structures found in a given language. Three main sources of data are used to devise grammars: (i) usage-based data derived from corpora or the Web; (ii) introspective grammaticality judgments given by individual native speakers; and (iii) experimental studies eliciting acceptability judgments from groups of speakers. There are heated (but ultimately inconclusive) debates in the research literature about what is the 'best' way of collecting data. For the most part, the judgments presented in this book will be

based on my own intuitions about grammaticality (as a native speaker and experienced linguist): I will highlight cases that I am aware of where my intuitions differ markedly from those of other native speakers.

Although the distinction between descriptive and prescriptive approaches to grammar might seem to be clear-cut, it should be noted that even descriptive grammars can sometimes be implicitly prescriptive. This is because descriptive grammars generally try and characterise so-called standard languages. This can involve making judgments about whether a given type of structure is found in standard or non-standard varieties; and this in turn can be regarded as tantamount to *prescribing* what can and can't be said in the standard language.

1.3 Traditional grammar

Contemporary syntactic theory makes use of a wide range of concepts and constructs which are rooted in centuries of earlier grammatical tradition, as well as introducing new techniques, terminology and perspectives of its own. For this reason, a useful starting point for any book about syntax is to look at key ideas from traditional grammar (as reflected, for example, in reference grammars, and in pedagogical grammars for second language learners).

Within traditional grammar, the syntax of a language is described in terms of a taxonomy (i.e. classificatory list) of the range of different types of syntactic structure found in the language. The central assumption underpinning syntactic analysis in traditional grammar is that phrases and sentences are built up of a series of **constituents** (i.e. syntactic units), each of which belongs to a specific **grammatical category** and serves a specific **grammatical function**. Given this assumption, the task of the linguist in analysing the syntactic structure of any given type of sentence is to identify each of the constituents in the sentence, and (for each constituent) to say what category it belongs to and what function it serves. For example, in relation to the syntax of a simple sentence like:

(9) Students protested

it would traditionally be said that the sentence consists of two constituents (the word *students* and the word *protested*), that each of these constituents belongs to a specific grammatical category (*students* being a plural **noun** and *protested* a past tense **verb**) and that each serves a specific grammatical function (*students* being the **subject** of the sentence, and *protested* being the **predicate**). The overall sentence *Students protested* has the status of a **clause** which is **finite** in nature (by virtue of denoting an event taking place at a specific time) and has the semantic function of expressing a **proposition** which is **declarative** in type (in that it is used to make a statement rather than, e.g., ask a question or issue an order). Accordingly, a traditional grammar of English would tell us that the simplest type of finite declarative clause found in English is a sentence like (9) in which a

nominal subject is followed by a verbal predicate. Let's briefly look at some of the terminology used here.

In traditional grammar, words are assigned to grammatical **categories** (called **parts of speech**) on the basis of their **semantic** properties (i.e. meaning), **morphological** properties (i.e. the range of different forms they have) and **syntactic** properties (i.e. word-order properties relating to the positions they can occupy within sentences): a set of words which belong to the same category thus have a number of semantic, morphological and syntactic properties in common. To illustrate this, let's begin by looking at what are sometimes called the **major** categories of English – i.e. those categories which have dozens, hundred or thousands of members. Let's start by looking at the category of **noun** (conventionally abbreviated to N).

Nouns are traditionally said to have the semantic property that they denote entities: so, *bottle* is a noun (since it denotes a type of object used to contain liquids), *water* is a noun (since it denotes a type of liquid) and *John* is a noun (since it denotes a specific person). There are a number of distinct subtypes of noun: for example a noun like *chair* is a **count noun** in that it can be counted (cf. *one chair, two chairs* . . .), whereas a noun like *furniture* is a **non-count** (or **mass**) **noun** in that it denotes an uncountable mass (hence the ungrammaticality of **one furniture, *two furnitures*: recall that a prefixed star/asterisk is used to indicate that an expression is ungrammatical). Likewise, a distinction is traditionally drawn between a **common noun** like *boy* (which can be modified by a **determiner** like *the*, as in '*The boy* is lying') and a **proper noun** like *Andrew* (which can't be used in the same way in English, as we see from the ungrammaticality of *'*The Andrew* is lying'). Typical count nouns exhibit the morphological property of having two different forms: a **singular** form (like *horse* in *one horse*) used to denote a single entity, and a **plural** form (like *horses* in *two horses*) used to denote more than one entity. Common nouns have the syntactic property that only (an appropriate kind of) common noun can be used to end a sentence such as *They have no* . . . In place of the dots here we could insert a singular count noun like *car*, or a plural count noun like *friends*, or a mass noun like *money*, but not other types of word (e.g. not *see* or *slowly* or *up*, as these are not nouns).

A second major category is that of **verb** (= V). Verbs are traditionally said to have the semantic property that they denote actions or events: so *eat, sing, pull, resign* and *die* are all verbs. From a syntactic point of view, verbs have the property that only an appropriate kind of verb (in its uninflected **infinitive** form) can be used to complete a sentence such as *They/It can* . . . So words like *stay, leave, hide, die, starve* and *cry* are all verbs and hence can be used in place of the dots here (but words like *apple, under, pink* and *if* aren't). From a morphological point of view, regular verbs like *cry* in English have the property that they have four distinct forms: e.g. alongside the **bare** (i.e. uninflected) **form** *cry* we find the **present tense** form *cries*, the **past tense+perfect participle+passive participle** form *cried* and the **progressive participle+gerund** form *crying*. (See