

1 Grammar

1.0 Overview

In broad terms, this book is concerned with aspects of grammar. Grammar is traditionally subdivided into two different but interrelated areas of study – **morphology** and **syntax**. 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 parts/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 §1.1) by looking at a range of approaches to the study of grammar, before going on (in §1.2 and §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 developed by Chomsky over the past six decades.

1.1 Approaches to Grammar

A fundamental question that needs to be resolved at the outset concerns what kind of approach to adopt in studying grammar. One traditional view sees the role of grammar as being essentially **prescriptive** (i.e. prescribing norms for grammatical correctness, linguistic purity and literary excellence). However, a more modern view 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 TV dialogue between the fictional Oxford detective Morse and his assistant Lewis, as they are looking at a dead body (where 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 type of school that 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 that 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 around the grammar of an italicised phrase which comprises the **preposition** *by* and the **pronoun** *who(m)*. (See the Glossary at the end of the book for any terms that are unfamiliar to you.) The differences between what the two men say relate 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 the 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 discussion here 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. *By whom?* is used in formal styles of English, and *Who by?* or *By who* in non-formal styles). 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). And fifthly, prescriptive rules are very often oversimplistic, in the sense that they paint a vastly oversimplified picture of what is in fact a more complex linguistic reality (as our discussion of *by who/m* and *who by* illustrates). For reasons such as these, the approach taken to grammar in work over the past sixty years or so has been descriptive.

What this means is that in attempting to devise a grammar of, for example, 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). Contemporary 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 kind of structures they are interested in.

However, although usage data (from corpora or the Web) provide a useful source of information about what people say or write, there are some downsides associated with a 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 (i.e. someone who has acquired and used English as a first language in an English-speaking environment from birth or early childhood, and who speaks the language fluently), and if so what variety/dialect of English they speak.

A very different (essentially mentalist) 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, for example, **I no like syntax* (Note that an 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** that 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 (determine the meaning of) 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 that 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**, 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). For example, 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 judgements.

However, although extensively used, this approach of relying on introspective judgements about the grammaticality of sentences has been criticised by some as being unscientific (hence yielding potentially unreliable results), particularly in relation to judgements about marginal sentences, i.e. sentences of dubious grammaticality, such as the following:

- (3) a. He ought to apologise, *oughtn't he?*
 b. He ought to apologise, *shouldn't he?*

This type of structure is generally referred to as a tag question (with the italicised part of the sentence following the comma being the tag). Normally in tag questions like this, the auxiliary in the tag is a (contracted negative) copy of that used in the **main clause** (i.e. the part of the sentence preceding the comma). However, use of the *oughtn't* tag in (3a) results in a relatively degraded sentence for speakers like me, simply because the contracted negative form *oughtn't* is obsolete in my variety. So, if you asked me to rate a sentence like (3a) on a 5-point scale (where 5 denotes 'completely acceptable' and 1 denotes 'completely unacceptable'), I'd give it a low score of 2. Instead, I'd prefer to use *should* in the tag (which can freely have the contracted negative form *shouldn't*), as in (3b): but since this results in a structure with mismatching auxiliaries (*ought* in the main clause, *shouldn't* in the tag), it feels ungainly, so I'd probably rate it as 3 (and thus treat it as marginal, i.e. a sentence which is neither clearly grammatical nor clearly ungrammatical).

Marginal sentences pose a problem for the use of introspective judgements about sentence acceptability for a number of reasons. Firstly, different individuals may disagree in their judgements 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 judgements 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, grammaticality is sometimes 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 may 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 judgements tend to give different judgements from non-linguists, and are vulnerable to the accusation that (however unwittingly) they may tailor their grammaticality judgements to fit their analysis (e.g. they may judge a given sentence to be grammatical because their analysis predicts that it should be).

Because of the potential unreliability of informal intuitions, some linguists prefer to elicit native speaker judgements experimentally, particularly when dealing with marginal structures whose grammaticality status is not clear-cut. 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 subsequently re-run. 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).

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, for example, 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 judgements given by individual native speakers; and (iii) experimental studies eliciting acceptability judgements from groups of speakers. There are heated debates in the research literature about what is the ‘best’ way of collecting data. For the most part, the judgements presented in this book will be based on my own intuitions about grammaticality (as a native speaker and experienced linguist), and on data I have collected from broadcast media or the internet: I will highlight cases that I am aware of where my intuitions differ markedly from those of other native speakers.

Although (on the basis of considerations such as those above) it is tempting to dismiss prescriptive grammar as anachronistic, pedantic or even worthless, it should be pointed out that (a mild form of) prescriptive grammar can play a useful role in society. After all, prescriptive grammar sets out standards for ‘good English’, and none of us would want to find that a job application we make is turned down because it is written in what the potential employer regards as ‘bad English’. Hence, being aware of whether a given type of structure is (or isn’t) generally accepted in ‘standard’ varieties of English is a key life skill.

A further point to note is that although the distinction between descriptive and prescriptive approaches to grammar might seem to be clear-cut, even descriptive grammars can sometimes be implicitly prescriptive. This is because descriptive grammars generally try and characterise so-called standard languages, and this can involve making value judgements about whether a given type of structure is found in standard or non-standard varieties; this in turn can be regarded as tantamount to prescribing what can and can’t be said in the standard language.

A final point to note is that while this book investigates aspects of ‘the grammar of English’, the use of the definite article *the* is potentially misleading, in that it carries the implicit assumption that all native speakers of English somehow have ‘the same grammar’, and thus that there is a single grammar of English which all native speakers share. However, it will become abundantly clear from some of the data presented in the book (especially authentic data sourced from radio and TV programmes or from the internet) that speakers of different varieties of English have grammars which differ (and are incompatible) in certain respects: for example, some speakers accept sentences like ‘He jumped *out* the window’, while others do not and instead say ‘He jumped *out of* the window’; conversely, some speakers accept ‘He jumped *off of* the table’, while others do not and instead say ‘He jumped *off* the table.’ Indeed, such variation often leads to heated debates on internet forums about whether or not a given type

of structure is grammatical in English. The more general point to underline is that there are a wide range of different varieties and registers of English, and while these share a great deal in common, they sometimes differ in specific areas.

1.2 Words, Categories and Features

Contemporary syntactic theory makes use of a wide range of concepts and constructs rooted in centuries of earlier grammatical tradition, as well as introducing new techniques, terminology and perspectives of its own. For this reason, in this section and the next I'm going to look at key ideas from traditional grammar (as reflected e.g. in reference grammars, or pedagogical grammars for second language learners), and note how some of these ideas have evolved in contemporary work.

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:

(4) 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 **noun** and *protested* a **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 for example 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 (4) in which a subject comprising a noun is followed by a predicate comprising a verb. In this section, we'll take a look at grammatical categories, while in the next we'll look at grammatical functions.

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. their **distribution** – that is, the positions they can occupy within phrases and sentences): a set of words which belong to the same category thus have a number of semantic, morphological and syntactic properties in common. An important point to note at the outset is that the properties of a given category of items may differ between languages: for example, in English we can differentiate **nouns** from **adjectives** in that we can add the plural ending *-s* to regular nouns (as in *fools*) but not to adjectives (cf. **foolishes*). However, this is not a criterion which we could apply to a language like Italian, where both adjectives and nouns can inflect for number and gender: for example, in a phrase like *belle_{beautiful} macchine_{cars}* ‘beautiful cars’, both the noun and the adjective carry the feminine plural ending *-e*. For this reason, I shall only make use of the properties of words in English to determine the set of word categories found in English: you should bear in mind that other languages may differ from English in certain ways.

I’ll begin by looking at what are sometimes called the **major categories** of English, i.e. those categories which have dozens, hundreds or even thousands of members: in more recent work, these are termed **lexical categories** because their members are **content words/contentives** which have idiosyncratic descriptive semantic content which varies from one **lexical item**/word to another. One such major/lexical category is that of noun, abbreviated to **N** (capital letters being used by convention for abbreviated category labels). 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/lexical 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, verbs have the property that they typically have up to five potentially distinct forms: for example, alongside the **base** (i.e. uninflected) **form** *show* we find the **perfect participle+passive participle** form *shown*, the **past tense** form *showed*, the (third person singular) **present tense** form *shows*, and the **progressive participle+gerund** form *showing*. (See the relevant entries in the Glossary if you are not familiar with some of the grammatical terms used here or elsewhere in the book.) However, for many verbs the distinction between two or more different forms is **syncretised/neutralised** (e.g. *cut* serves as a base form, perfect/passive participle, and a past tense form). This is shown by the table of verb forms below (where PERF/PASS denotes a form which can be used as a perfect or passive participle, PAST denotes a past tense form, 3.SG.PRES denotes a third person singular present tense form, and PROG/GER denotes a form which can be used as a progressive participle or gerund):

(5) **Table of verb forms**

BASE	PERF/PASS	PAST	3.SG PRES	PROG/GER
show	shown	showed	shows	showing
go	gone	went	goes	going
speak	spoken	spoke	speaks	speaking
see	seen	saw	sees	seeing
come		came	comes	coming
wait	waited		waits	waiting
meet	met		meets	meeting
cut			cuts	cutting

The largest class of verbs in English are regular verbs which have the morphological characteristics of *wait*, and so have past, perfect and passive forms ending in the suffix *-(e)d*. However, the picture becomes more complicated for the irregular verb BE, which has eight distinct forms (the base form *be*, the perfect form *been*, the progressive/gerund form *being*, the past tense forms *was/were*, and the present tense forms *am/are/is*). The most regular verb suffix in English is *-ing*, which can be attached to the **base form** of almost any verb (though a handful of defective verbs like *beware* are exceptions).

A third major/lexical category found in English is that of **adjective** (= A). Adjectives are traditionally said to have the semantic property of denoting states or attributes (e.g. *happy*, *tired*, *conscientious*, *red*, *cruel*, *ill*, *old* etc.). Many (but not all) adjectives have the morphological property that they have negative *un-* counterparts (e.g. *unhappy*), **comparative** forms ending in *-er* and **superlative** forms ending in *-est* (e.g. *happy/happier/happiest*), and have noun derivatives in *-ness* (e.g. *happiness*). Those adjectives denoting a **gradable** property which can exist in varying degrees have the property that they can be modified by