1 Introduction

1.1 Dimensions of Linguistic Information

Any natural language is a seriously complex system – in fact, learning your native language(s) could be considered your all-time most amazing intellectual achievement, even though you did not have to put much effort into it. The aim of any theory of grammar is to build an abstract model that can account for this complexity, that can predict what is and what isn't a grammatical sentence and that can, in some sense, explain why. What else is included in a model of grammar varies between theories; meaning may be included, and in some cases the pragmatic circumstances under which the particular linguistic element would be used. Some models are also committed to a particular explicit view of how children acquire language. The framework we will describe in this book – Lexical-Functional Grammar (LFG) – does not commit to any specific assumptions about language acquisition. However, when the ideas that led to the formal system that is LFG were first developed, this was in part a reaction against other approaches at the time that were perceived to be unrealistic as representations of how humans 'do' language. One of the founders of LFG, Joan Bresnan, wrote, 'If a given model of grammar cannot be successfully realized within a model of language use, it may be because it is psychologically unrealistic in significant respects and therefore inadequate in those respects as an empirical theory of the human faculty of language' (Bresnan, 1978, 2). In a similar vein, the first major work on LFG was entitled *The mental representation of grammatical relations* (Bresnan, 1982c).

LFG then aims to model the linguistic information that a native speaker has of their language; this information is complex and multifaceted, especially keeping in mind the variety that we find across the world's 6,500 or so languages. If you take a sentence like the one in (1), there are a lot of things we can say about it linguistically; there are many types of linguistic information associated with it.

(1) The cats devoured their breakfast.

Assuming that the sentence was spoken, there would be sounds, probably something like (2).

(2) / ðə kæts dɪvaʊəd ðeə brekfəst /

There may also have been emphasis on some part of the sentence to indicate what the speaker thought was the most important information. If there was emphasis on *their*, this would probably be because the speaker was contrasting *the cats' breakfast* with someone else's breakfast that the cats did not devour, so they might have added … *not the rats' breakfast*. In this case, the speaker would have used stress or intonation – together referred to as PROSODY – to indicate which information is new to the conversation and which is assumed to be known already. These distinctions with respect to the INFORMATION STATUS of constituents is yet another type of information about linguistic items. Many languages can use word order to indicate the information status of constituents, and when they do there is a tendency to organise sentences so that old information precedes new. However, English offers relatively little scope for changing the word order and so prosody is frequently used instead.

Different CATEGORIES of words can be identified; *cat* and *breakfast* are NOUNS and *devoured* is a VERB. The reason for saying that *cat* and *breakfast* belong to the same category is that they behave in similar ways formally. For instance, words like *cat* and *breakfast* can have a plural ending and can combine with *the*, whereas *devour* cannot. A word like *devour*, on the other hand, can occur in present or past tense – *devour* vs *devoured* – and can take a third person singular *-s* in present tense. The words *the* and *their* behave in similar ways and we will refer to such elements as DETERMINERS. Like many concepts we will use in our analysis, categories are idealisations; not all nouns have all the noun properties we can identify – they are not all equally "noun-y" – and the morphosyntactic properties of nouns vary across the languages of the world. The number of categories that can be recognised by formal criteria varies between languages; some languages may not, for instance, make a formal distinction between verbs and adjectives.

(1) is not just a flat string of words: some words belong more closely together, so *their* and *breakfast* form a unit, which in turn combines with *devour*. This yields the hierarchical structure that we refer to as CONSTITUENT STRUCTURE. The word-level categories form phrases, so that the NOUN *breakfast* forms a NOUN PHRASE with *their* and the VERB *devour* in turn forms a VERB PHRASE with that noun phrase. As we will soon see, the degree of structure that can be identified varies greatly across languages, with English being a particularly highly structured language. We will also have more to say about English clause structure specifically in Section 2.3.2.

Though *the cat* and *their breakfast* are both noun phrases, they take on different GRAMMATICAL FUNCTIONS – or GRAMMATICAL RELATIONS – within the sentence. The function of *the cat* in (1) is referred to as the SUBJECT and *their breakfast* is the OBJECT. Grammatical relations are central to LFG and we will introduce them more thoroughly in Section 2.2.

There is also information about the internal structure of words to be captured. A complete description of (1) will have something to say about how the words are built up and the role the different parts play in the sentence. There is MORPHOLOGICAL information which tells us that the words *cats* and *devoured*

both consist of two parts, the main meaning part – *cat* and *devour*, respectively – and an inflectional morpheme indicating plural (*-s*) and past tense (*-ed*).

We also know things about the meaning of words – about the SEMANTICS – and how the meaning of individual words is combined to form the meaning of phrases and sentences. The meaning of *cat* is intuitively clear in that we agree pretty much on which entities in the world can be referred to appropriately by using *cat*. We also have ways of describing the semantics of words like *the*; we use it roughly speaking when we think the entity referred to by the noun following it is known to the hearer. The semantics of a verb such as *devour* is a bit more complex in that knowing about its meaning also means knowing that it needs to combine with other elements; you cannot get any devouring going unless you have someone to do the devouring and something (or someone) to be devoured. We will refer to this as the verb requiring ARGUMENTS. The arguments have different roles with respect to the semantics of the verb; for instance, in (3a) the roles of *the dog* and *the postman* are different – *the dog* carries out an action whereas *the postman* undergoes the action. The role of each of the noun phrases is also different between the sentences in (3a) and (3b) – *the dog* no longer carries out an action in (3b) but experiences an emotion, and *the postman* may not even be aware of the feelings, whereas he would be aware of the biting in (3a). These roles are referred to as THEMATIC ROLES or θ-ROLES, and when we record information about the number of arguments a predicate takes, we also need to include the θ-roles of those arguments.

(3) a. The dog bit the postman.
b. The dog loves the postman.

As we have seen, there are many types of information that a full description of a language would need to refer to. **How** different models do this varies a fair bit. Some models assume that certain types of linguistic information should not be included in the model of grammar itself, but should be dealt with outside the actual grammar. LFG includes all these dimensions of linguistic information in the model. The different dimensions capture the information in different ways, so whereas a hierarchical tree is a natural way of representing constituent structure, this is not the best way of representing information about functions. The dimensions are then linked by mapping rules. The fact that dimensions of information are represented separately, but are linked, is a fundamental characteristic of LFG, and for this reason it is referred to as a PARALLEL CORRESPONDENCE ARCHITECTURE. LFG is formally explicit, which means that analyses can be tested computationally.

1.2 The Architecture of LFG

We have established that we need information about different aspects of linguistic elements and that this information may take different shapes. The LFG architecture contains the following dimensions:

- argument structure and thematic roles: A-STRUCTURE (for argument structure)
- syntactic categories and constituency: C-STRUCTURE (for constituent or category structure)
- grammatical functions and other functional features: F-STRUCTURE (for functional structure)
- morphological properties: M-STRUCTURE
- prosody: P-STRUCTURE
- semantics: S-STRUCTURE (usually modelled using 'glue semantics' as the theory of the syntax–semantics interface)
- information structure (or discourse structure): I-STRUCTURE (or D-STRUCTURE)

These different dimensions of information need to be linked – or in LFG terms MAPPED. The mapping will make explicit how a particular c-structure is linked to an f-structure, or how a p-structure is linked to an i-structure etc. If this linking was not part of a language user's knowledge, we would not be able to connect the correct semantic interpretation to a particular string of sounds. To take a simple example, a native speaker (or anyone who has learned English to even a basic level) knows that the sentence in (3a) involves *the dog* and not the *the postman* doing the biting. In LFG terms, this means that the c-structure of (3a) must be mapped to an f-structure in which *the dog* is the subject and *the postman* the object.

The dimensions that will be central in this book are a-structure, f-structure and c-structure, and we will briefly outline them here. Much more will be said about them in Chapters 2, 3 and 8. These are the most well-established components of LFG. For some of the other dimensions there are alternative ways of representing them within LFG, but we will not go into the details here; instead, we will provide useful references at the end of the chapter.

C-structure is represented as category-labelled trees. This means that the trees capture both category and constituent structure. We'll consider English first. Take the sentence in (4).

(4) The dog did bite the boy.

As referred to in Section 1.1, there is a hierarchical structure to a sentence like this in English. The two instances of *the* form constituents with *dog* and *boy*, respectively. Similarly *bite* forms a constituent with *the boy* and *did* with *bite the boy*. We can apply constituency tests to show that these strings function as units structurally. The resulting structure for (4) can be represented as in (5).

(5)

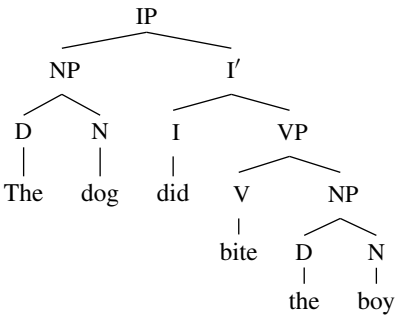

There are some things about this tree that may appear odd at this stage: the fact that *did* heads the phrase *did bite the boy*, that this involves a category I and that there is a category I′, for instance. We will explain these choices in some detail in Section 2.3, where we will also take another look at the NPs. At this point all we want to do is illustrate the fact that the hierarchical structure is represented as a tree, and that the constituents in the tree are assigned a category label. The tree in (5) consists exclusively of ENDOCENTRIC constituents, that is, all constituents have a head, the IP has a head I, the NP a head N etc. This is quite a striking fact about English; it is highly CONFIGURATIONAL. This means that the hierarchical structure plays an important role in expressing linguistic information. We will show how this is captured in LFG in Section 3.2. It is also the case in (5) that each node has at most two daughters – the tree in (5) is BINARY BRANCHING. This is quite common in English, but it is not exclusively so; some nodes may have only one daughter, or three or more, and in other languages, non-binary branching is the norm.

Not all languages rely as heavily on syntactic structure as English does. Consider the data from Latin in (6), for instance. (Since we will be using a range of languages in this book, it is important that the reader gets used to using the glossing to understand examples: see the List of Abbreviations at the beginning of the book.)

(6) a. Caesar suas copias in
Caesar.NOM REFL.ACC.F.PL troop.ACC.F.PL in
proximum collem subducit
nearest.ACC.M.SG hill.ACC.M.SG withdraw.PRS.3SG
'Caesar withdraws his (own) forces to the nearest hill.' [*De bello Gallico* 1.22]

b. copias suas Caesar in
troop.ACC.F.PL REFL.ACC.F.PL Caesar.NOM in
proximum collem subducit
nearest.ACC.M.SG hill.ACC.M.SG withdraw.PRS.3SG
[*De bello Gallico* 1.24]

c. copias suas in proximum
troop.ACC.F.PL REFL.ACC.F.PL in nearest.ACC.M.SG
collem subducit Caesar
hill.ACC.M.SG withdraw.PRS.3SG Caesar.NOM
[constructed]

d. non respuit condicionem Caesar
NEG reject.PRS.3SG proposal.ACC.F.SG Caesar.NOM
'Caesar did not reject the proposal.' [*De bello Gallico* 1.42]

There are some similarities with English here. For instance, notice that *in proximum collem* is not separated in any of the first three sentences and occurs in that order, just like English *to the nearest hill* would, so Latin appears to have a prepositional phrase here much like English does. On the other hand, though *suas* 'his own' and *copias* 'troops' occur together, both orders *suas copias* and *copias suas* are possible, and we can conclude that Latin has noun phrases, just like English, but that the order within them is somewhat freer. These conclusions are based on the assumption that the data in (6) is representative, obviously; we would need a lot more data to confirm this initial assessment. The big difference between English and Latin lies in the order of the major constituents. If we consider the verb, the subject 'Caesar' and the object 'his own troops' in (6a)–(6c) or 'the proposal' in (6d), then the order is [Subject Object V] in (6a), [Object Subject V] in (6b), [Object V Subject] in (6c) and [V Object Subject] in (6d). Though (6c) is constructed, there is good evidence that this order existed under the right pragmatic circumstances. Since the verb and the object do not have a fixed position with respect to each other-indeed they do not even have to be adjacent-there is no reason to include a VP in our c-structure for Latin. A language that lacks evidence for a VP is often described as NON-CONFIGURATIONAL. There are also no arguments in Latin for assuming that there is an I and an IP, but we will get back to the issue of what motivates the use of IP for a clause in Section 2.3.2. This means that trees like those in (7) to (10) are more appropriate for the Latin sentences in (6) (to keep the trees simple, we have ignored the adverbial). The key differences compared to the tree for English in (5) are that there is no VP, that the sentences are of the EXOCENTRIC – that is non-headed – category S and that different orders are possible. We end up with a flat tree at clausal level. Unlike in English, it is not the position that identifies the subject and the object, but case markers on the noun phrases – NOMinative for subject and ACCusative for object – and the agreement on the verb – here third person singular. Both these means identify *Caesar* as the subject in this example. We will see how this works in more detail in Chapter 4.

(7)

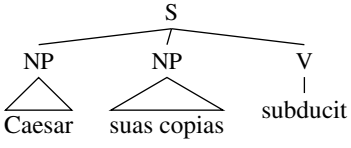

(8)

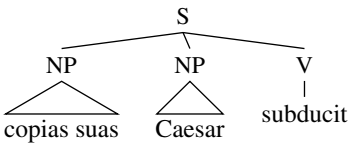

(9)

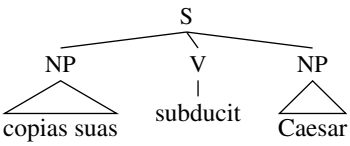

(10)

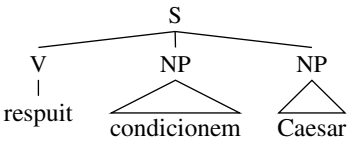

We have shown here that the extent to which a language relies on structure to identify functions varies and that this variation is represented as different c-structures in LFG. We will see more c-structure variation in Section 2.3.4.

If we turn to arguments and θ-roles, we find a different picture. Take *bite* in (4): if you know what this means, you know that it involves two participants. As we said in Section 1.1, one argument carries out an action and the other one suffers from that action. This is all part of the meaning of *bite*, and indeed of any other word meaning the same, whether it is Latin *mordeo* or Dutch *bijten* or indeed a verb meaning the same thing in any language. It is common to generalise over θ-roles, so rather than refer to them as "biter" and "bitten one", we use AGENT and PATIENT for any verb where there is some participant acting on another participant. We will return to θ-roles briefly in Section 2.1 and more extensively in Chapter 8. The a-structure for the word meaning 'bite' in any of these languages is then as in (11).

(11) < Agent, Patient >

Similarly with f-structure; whatever means a language uses to identify what is the subject and what is the object, the Agent will be identified as the subject and the Patient as the object in the translation of (4). This means that the f-structure for any of the sentences in (4) or (12) will be as in (13), in spite of the obvious difference even between these three relatively closely related languages (ignoring details such as definiteness and number, as well as the presence of *did* in (4), it will be clear in Section 2.3.2 why we included it in this example). We hope that the ideas behind (13) are clear to the reader, though PRED might require some explanation; it can be said to capture the semantic form of the predicate and the respective roles of the SUBJ and OBJ within it. SUBJ and OBJ are features with the respective values [PRED 'DOG'] and [PRED 'BOY']. We will have more to say about PRED and f-structure in Section 2.2.

(12) a. De hond heeft de jongen gebeten. (Dutch)
the dog have.PRS.3SG the boy bite.PPTCP

b. Canis puerum momordit. (Latin)
dog.NOM.M.SG boy.ACC.M.SG bite.PRF.3SG

(13)
$$\begin{bmatrix} \text{PRED} & \text{`BITE <SUBJ,OBJ>'} \\ \text{SUBJ} & \begin{bmatrix} \text{PRED} & \text{`DOG'} \end{bmatrix} \\ \text{OBJ} & \begin{bmatrix} \text{PRED} & \text{`BOY'} \end{bmatrix} \end{bmatrix}$$

1.3 Mapping between Dimensions

In the chapters that follow, we will have a lot more to say about the mapping between the dimensions of information. However, at this point, we would like to give you a rough idea of how it will work, at least in a language like English that relies on structure for the identification of grammatical relations and contrast it with a language like Latin that relies less on structure. Take the c-structure for *The dog did bite the boy*, which was provided in (5). Because *the dog* is in the structural position it is in the tree; we know it is the subject. Hence the mapping function must provide a link between that position and the value of SUBJ in the f-structure. Once a speaker has identified that *the dog* is the subject, she also knows that it is the Agent of the verb in this example; this is how a speaker of English can tell who is biting whom. This means there must also be a mapping (or correspondence) between the SUBJ in the f-structure and the a-structure Agent of *bite*. Similarly, the position of *the boy* in the tree must be mapped to the value of OBJ in the f-structure, which corresponds to the Patient in the a-structure. Considering the Latin sentence in (12b), we know from the discussion around the examples in (6) that the word order could have been different in (12b), so that it is not the position of *canis* that tells the hearer that it is the subject. Instead it is the NOMinative case marker that provides the crucial clue and hence the mapping to SUBJ must link to the case-marked noun. English and Latin thus share the same associations, or mapping, between a-structure and f-structure, as shown in (14), but have quite different c-structures, as shown in (15).

(14)
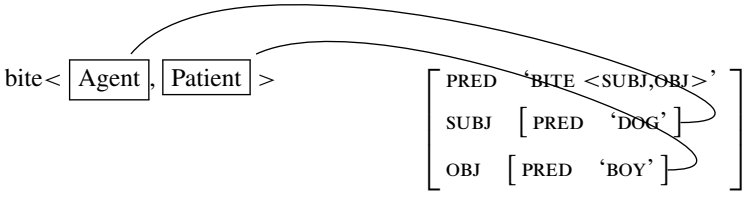

(15)

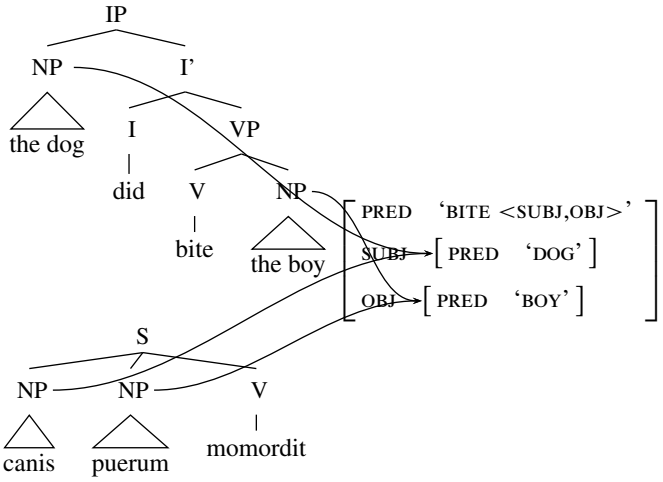

1.4 Why Different Dimensions?

We have shown here that the linking between categories, functions and arguments varies between languages so that there is cross-linguistic evidence for the need to distinguish the different dimensions of linguistic information; there are arguments in favour of a parallel correspondence architecture. In fact, data internal to English also supports this distinction.

Let's consider category first and compare the COMPLEMENTS of some verbs. We use 'complements' as a cover term for all elements required by the verb apart from the subject. The examples in (16) show that the verb *believe* can take a clausal complement or a noun phrase complement, either a pronoun or a full noun phrase. We have labelled the clause here as '*that*-clause'; this is just informal, we will come back to details about the category of clauses in Section 2.3.2 and Chapter 6. Do note that *the fact that there are fairies at the bottom of the garden* is a noun phrase built up around the noun *fact*, which takes as its complement the clause *that there are fairies at the bottom of the garden*. The verb *smile*, on the other hand, cannot take any complement, as shown by the ungrammaticality of the examples in (17).

(16) a. Oscar believes [that there are fairies at the bottom of the garden]$_{\textit{that}-\text{clause}}$
b. Oscar believes [it]$_{\text{NP}}$
c. Oscar believes [the fact that there are fairies at the bottom of the garden]$_{\text{NP}}$

(17) a. *Oscar smiles [that there are fairies at the bottom of the garden] $_{\text{that}-\textit{clause}}$
b. *Oscar smiles [it]$_{\text{NP}}$
c. *Oscar smiles [the fact that there are fairies at the bottom of the garden]$_{\text{NP}}$

This correlation is not remarkable; the distribution of complement types appears to be predictable from argument structure. If you *believe*, you have to

believe some sort of proposition and it may be represented by a clause or a noun phrase. It only takes one entity to *smile*, so no complement is required.

However, things are not as straightforward as this. Most verbs are particular about what category their complement is: *enjoy* in (18) can take an NP, but not a *that*-clause; *hope*, on the other hand, can take a *that*-clause but not an NP, as (19) shows. Still, in semantic terms, the complement of both verbs describes a situation that you *enjoy* or *hope* for.

(18) a. *Oscar enjoys [that there are fairies at the bottom of the garden]$_{that-\text{clause}}$
b. Oscar enjoys [the fact that there are fairies at the bottom of the gardens]$_{\text{NP}}$
c. Oscar enjoys [it]$_{\text{NP}}$

(19) a. Oscar hopes [that there are fairies at the bottom of the garden]$_{that-\text{clause}}$
b. *Oscar hopes [the fact that there are fairies at the bottom of the garden]$_{\text{NP}}$
c. *Oscar hopes [it]$_{\text{NP}}$

More generally, an element like a verb can determine not just how many complements there can be, but also the category of these complements. We cannot derive information about category from our knowledge of how many arguments there are; we need information separately about arguments and categories.

Turning now to the relation between category and function, so far all subjects we have seen have been noun phrases. If a particular function is always filled by a particular category, then we might think we only need to have one of the two types of information specified. However, subjects can be of categories other than noun phrase as in (20).

(20) a. [The young man]$_{\text{NP}}$ surprised his mother.
b. [On the beach]$_{\text{PP}}$ is a great place to be.
c. [That Oscar had finally found a girlfriend]$_{that-\text{clause}}$ surprised his friends.

The subjects in (20b) and (20c) may have some unusual properties in comparison to noun phrase subjects, but in many respects they do behave like subjects.

It is maybe more obvious, but also worth pointing out, that noun phrases can bear functions other than subject, as in (21). We use traditional labels for the functions here, but we will return to LFG terminology in Section 2.2.

(21) a. Oscar saw [the film]$_{\text{Object}}$
b. Oscar gave [his sister]$_{\text{Object}_{ind}}$ [a present]$_{\text{Object}_{dir}}$
c. Oscar left [the day before yesterday]$_{\text{Adverbial}}$

This is an indication that we need to specify category and function separately.

How about grammatical functions and θ-roles, do we need both? Often it is the subject which is the Agent, but in (18b) and (18c) the subject of *enjoys* can be described as an EXPERIENCER. In fact, subjects can have any number of other roles as the examples in (22) show.

(22) a. [Oscar]$_{Recipient}$ received the present from his sister.
b. [The hammer]$_{Instrument}$ drove the nail into the wood.
c. [On the beach]$_{Location}$ is a great place to be.

There are certain syntactic processes which reorganise sentences so that the same function can be linked to different θ-roles in different versions. One such process is passivisation. The subject of a passive sentence has the θ-role that is associated with the object in a corresponding active sentence, so that a Patient noun phrase can also be a subject, as in (23), or indeed any other θ-role that can occur as an object.

(23) [Sarah]$_{Patient}$ was tickled by Oscar.

In some cases, the same verb, in the same form, can also have subjects with different θ-roles; one such example in English is *break* as in (24).

(24) a. [The boy]$_{Agent}$ broke the window with a stone.
b. [The stone]$_{Instrument}$ broke the window.
c. [The window]$_{Patient}$ broke.

There is then evidence that the relation between categories, functions and θ-roles is not a one-to-one relation and that as a consequence we need to keep information about categories, functions and θ-roles separate. The different types of information are connected by mapping rules which do not necessarily predict a one-to-one relation between the dimensions. This realisation is crucial to the architecture of LFG.

Reading

If you are interested in the earliest formulations of LFG, then you might want to look at Bresnan (1978) or the articles in Bresnan (1982c), but keep in mind that the approach has developed since then, so some aspects of the formal notation may be outdated and some analyses may have been updated in more recent publications.

In Section 1.2, we introduced seven dimensions of information. In this book, we shall have a lot to say about three of these: a-structure, f-structure and c-structure. In Chapter 9, we say a little more about p-structure, m-structure, s-structure and i-structure, but, most importantly, there we provide further references to the literature for you to follow up if you are particularly interested in one of these dimensions.

A summary of issues that arise around non-configurationality can be found in Nordlinger (1998); we will return to how to analyse such languages in Chapter 4. If you are interested in the Latin data, you can find an LFG analysis of aspects of Latin in Jøhndal (2012).