2

What a language is and how it works

How essential is this chapter?

Surprisingly, it is NOT essential to read this chapter before going on to Chapter 3. Chapter 3, however, is absolutely essential. In fact, if you do not already have a solid background in linguistics you should probably just look briefly at Section 1 of this chapter, and then flick through the next three sections till you reach Section 5. That section gives a working example of all of the ideas introduced earlier. Alternatively, you could simply skipped the rest of this chapter and go straight to Chapter 3.

In other words - even though I have written this chapter as simply and directly as I can - most of it is quite demanding reading. This is because it tries to provide an answer to the two questions that are perhaps the most fundamental ones in linguistics, i.e. 'What is a language?' and 'How does language work?' In order to do this in twenty or so pages, I have had to introduce new and sometimes quite abstract concepts at a fairly rapid rate.

Yet this set of concepts will, in the long run, be absolutely fundamental to understanding the view of language that this book provides. This is especially true of Sections 2 to 5. So ***even if you decide to skip or skim this chapter at this point you will probably want to return to*** *it at one or more points, as you work your way through the rest of this book - and perhaps also to Figure 1 at the start of Chapter 1. This is because you will want to remind yourself of where, in the overall picture of language, the part that you are reading about at that moment fits in. And at some point - perhaps only by the end of the book - you will decide that you are ready for the brief summary of the fundamental concepts that this chapter provides. Nonetheless, it is here, at the start of the book, that it logically belongs.*

Finally, let me remind you that, as in Chapter 1, the FOOTNOTES (the notes at the bottom of a page, shown in the text by '1, 2, 3' etc) are for ALL readers, while the ENDNOTES (the notes at the end of the chapter, shown in the text by 'a, b, c' etc.) are for readers with more advanced experience of linguistics.

1 About the question 'What is a language?'

What is a language? Alternatively, we could ask: 'What is a language like?' And this, as we will discover in the course of this chapter, is a more useful starting point. So the purpose of the chapter is to answer the question 'What is a language?' by answering the second one - or at least to give a first outline of an answer. It will take both this book and its sister volume the *Functional Semantics Handbook* to give anything like a full answer - and even then it will be the answer to a related but ultimately rather different question, which is 'What do we need to know about a

language in order to describe a text-sentence that is spoken or written in that language?' Nonetheless it should be helpful to many readers to have, right from the start, a clear outline of the major components of the full answer.

But the question of what a language is cannot be separated from the question of how it works. We will therefore also find ourselves asking 'How does a language work?' - or, more specifically, 'How do we use a language to produce a text?' (We may also ask, of course, 'How do we use a language to understand a text?" - and the reason for emphasizing the question of how we produce a text is explained in Section 2.4. [a]

The answers to these questions that I will give here are based on my own overall picture of what a human language is like. Each person who tries to understand the fascinating phenomenon of human language has to decide what should be central in their own picture of language. The best help that I can give to you, as you work at the life-long process of trying to develop the best possible picture of language for yourself, is to show you the picture as I see it - and to do so as honestly and as straightforwardly as I can. You will then have to decide how much of it you feel should be incorporated into your own developing picture of language. I say 'developing', because however old or young you are and however much or little time you have put in on learning about language, our mental picture of what a language is should continue to change and grow throughout our life. Mine certainly does. Essentially, a linguist is an explorer.

2 Towards a good overall model of language

2.1 The need to have a good 'picture' of language

I have used the expression 'a picture of language'. The reason is that when we want to understand something - or to explain something to someone else - we often draw a picture of it. Sometimes it is sufficient simply to have a mental picture, but it often helps to draw a representation of it that is visible to the eye.[b] A good example would be a map drawn on the back of an envelope to show someone how to get to your home. Other examples are a flow chart that guides you through the task of filling out a government form, and even a neatly set out long division sum. All of these are, in their different ways, two-dimensional representations - i.e. 'pictures' - of things that are rather more abstract.

When we are working on paper we often use lines and boxes when we are picturing abstract things. We also use symbols of various kinds in such diagrams, such as arrows and boxes, and symbols borrowed from mathematics such as 'plus', 'minus' and 'equals' signs. And our diagrams often also include short written forms taken from language itself. Often we use a mixture of these.

Occasionally, if we have the resources, we go further and build a three-dimensional picture of the object, such as a model of the atom using straws and ping pong balls, or a model of the human heart made of lumps of coloured plastic. This is essentially what most children's toys are: three dimensional pictures of objects to play with, and so to use to learn about the world. But we are still essentially a '2D culture' - i.e. a culture dominated by two-dimensional paper-based

books, magazines, newspapers and billboards - and cinema, television, and computer screens. The result is that we typically settle for a two-dimensional picture - though a two-dimensional picture may of course be a representation of a three-dimensional object.[c]

This 'picture-making' is the most fundamental activity in human understanding - and so in human learning and in human explaining. 'Learning' means 'coming to understand', so it encompasses both the way in which a small child learns to make sense of his or her environment and the way in which a scientist makes progress at the frontiers of human understanding - and all the types of learning between the two - including learning a new method of analyzing texts (as in this handbook). And 'explaining' encompasses the transfer of ideas from one person to another at every stage of life and in every social context, from a university lecture to a map of how to reach your home drawn on the back of an envelope.

It is an interesting fact that, for the teacher or writer, 'learning' and 'explaining' merge into each other. This is because most teachers and writers continue to learn about their subject when they try to explain it to someone else. Indeed, it sometimes happens that the act of trying to provide a better picture of something for another person itself gives the speaker or writer a new insight into its nature. The most generally used word in science for the activity through which we try to understand - and which therefore also underlies both learning and explaining - is 'modelling'. So we can say that our task in this handbook and its companion volume is TO MODEL HUMAN LANGUAGE AND HOW IT IS USED.[1]

A model of language is like a model of anything else in that it is, by virtue of being a model, less complex than the phenomenon itself. If it wasn't simpler, it would not be a model of the thing that is being modelled, but a clone of it. So the essential characteristic of a model of something is that it is SELECTIVE. A model foregrounds THOSE ASPECTS OF THE PHENOMENON THAT THE MODELLER CONSIDERS TO BE MOST CENTRAL TO UNDERSTANDING THAT PHENOMENON. (I almost wrote 'feels' instead of 'considers', and it would not have been misleading to have done so - because ultimately such judgements rest on a gut feeling.)

2.2 Types of phenomenon and types of model

The many phenomena that humans wish to understand differ from each other in certain fundamental ways. So the first problem is to identify what type of phenomenon a human language is.

Sometimes the phenomenon that we are trying to understand is a PHYSICAL OBJECT, and sometimes it is an ABSTRACT OBJECT. Of the two types, it is easier to model physical objects. Normally we are quickly able to decide upon a possible way to start drawing or building a simplified representation of a physical object - in part because we can actually see it (through a telescope or a microscope, if

1. In everyday language the word *model* usually refers, rather narrowly, to a 'three-dimensional picture', as in *She made a matchstick model of their house*. But in science there is no such limitation. We will use a lot of two-dimensional diagrams to model language in this book, but from a more abstract viewpoint it could be argued that the model of language to be presented here has no less than EIGHT dimensions.

necessary). Modelling things of this type is sometimes simply a matter of selecting what you consider to be the salient properties of size, shape, colour, part-whole relationships, etc. But you can't do this with an abstract phenomenon, such as language. So the first reason why modelling language is difficult is that a language is something abstract.

Perhaps the most serious problem in finding a good model for an abstract phenomenon such as language is that the human mind finds abstract phenomena hard to deal with. Genetically, we are still stone age people, and our brains have developed to help us to solve the problems that beset stone age people - principally problems in the physical world, such as finding food. The result is that we tend to prefer models that have a physical presence; we like to be able to SEE our models. And this makes it quite a challenge for us to model abstract things well.

2.3 Static, functioning and behaving phenomena and their models

A second problem is that many of the most interesting things that we wish to understand are either BEHAVING phenomena or FUNCTIONING phenomena. A 'behaving phenomenon' is typically something that is alive, such as a human being, while a 'functioning phenomenon' is something that is used by a 'behaving phenomenon' (typically a human being) for some purpose, such as a camera or a car. In contrast with these two, we can call phenomena that don't 'behave' or 'function' STATIC phenomena.

But notice that, if we choose to, we can use a static phenomenon to represent a behaving or a functioning phenomenon - for example, a toy car whose wheels won't turn and which is so small that it cannot carry passengers).[d] However, a model of a functioning phenomenon that doesn't actually function is a poor model. Essentially, A LANGUAGE IS A FUNCTIONING PHENOMENON, IN THAT HUMAN BEINGS USE IT TO COMMUNICATE WITH EACH OTHER. It follows, then, that if we let ourselves be persuaded to adopt a static model of language we will find ourselves trying to understand language through a model that lacks its most vital characteristic. And yet, as we will see shortly, this is precisely what most current models of language offer us. at

So in the task of modelling language we need to consider both (1) the nature of the phenomenon of language itself and (2) the nature of the models of language that we might construct. And we need to do this in terms of the following two dimensions of variation:

(a) physical vs. abstract phenomena, and
(b) static vs. functioning vs. behaving phenomena.

What is a language like, then, in these terms? Firstly, then, a human language is clearly ABSTRACT**.** Secondly, it is equally clearly a FUNCTIONING PHENOMENON - i.e. it is a kind of 'abstract tool' that we use for accomplishing things. More precisely, its two main functions are (1) to turn a set of 'concepts' into a vocal or written output (in a broad sense of the potentially misleading term 'concept' which I will clarify very soon), and (2) to turn a vocal or written input into a set of such 'concepts'. When I express it like that, the case for treating language as an abstract

functioning phenomenon seems irrefutable - and it is certainly true that the general model of 'language as an instrument' is one that is used in very many introductory books about language.

Surprising as it may seem, many professional linguists would not be completely happy with this view of language as primarily a means of communication. Many linguists - especially those who have been influenced by what I will refer to as 'formal linguistics' - work with a rather different model, in that they see language as being essentially something that you KNOW, rather than something that you USE. Indeed, some linguists say quite openly that the reason why they are interested in language is that it is 'a window on the mind' (i.e. on the mind of an individual, thinking human) - and not because they want to understand how humans communicate with each other. So researchers who study language may have very different goals from each other.[2]

But there is a second reason why some linguists think of a language as something that we know rather than something that we use - and it is is that they have unconsciously accepted a long-established but seriously misleading model of language. This is the model of language in which it is seen as consisting essentially of a book, and so a physical object. Sometimes language is seen as essentially a grammar book (perhaps with additional notes about words and word classes), and sometimes as essentially a dictionary with occasional notes about grammar - and in the full version of this model there is both a grammar and a dictionary. Let us call this picture of language the 'two book' model.[3]

Even though most linguists would probably agree that a language is an abstract phenomenon, when theoretical linguists try to model the mental 'grammar' (i.e. the 'rules' of 'syntax' and 'morphology') they normally do so in a way that comes very close to being an actual 'grammar book', such as you might have on your bookshelf (though it is often dressed up in logical formalisms that you need to learn first in order to interpret it). It typically consists mainly of statements about part-whole relationships within the type of object known as a 'sentence'. But the view held by many generative-formalist linguists that (a) a language is essentially a set of sentences and (b) the task for linguistics is to write descriptions of how sentences are made up of words and morphemes is, as I hope to show very soon, a seriously impoverished view of language.[e]

2. In other words, language exists primarilty for communication, and when we use it to help us to think through some problem on our own, we are in effect communicating with ourselves. And when we write about difficult ideas, as I am doing now, the line between writing for myself and writing for you, the future reader, is not always clear - at least in the initial draft. (This is why the writer of material such as this should write and re-write the text - at least once and probably more often for themself, and at least once and probably more for the reader.)

3. While we may consult the grammar book and the dictionary as an aid to communicating, they are not, in themselves, a means of communication, as language itself is. However, they do function as texts when we read them as messages from their writers to us, their readers; see the distinction between 'potential' and 'instance' that is made in Section 3.2. In these terms, 'instances' at the level of 'form' are 'objects'). In practice, however, grammar books are only consulted very rarely (except by linguists and learners of a language), and even dictionaries are used relatively rarely in everyday life. So in practical terms the grammar and the dictionary are essentially static physical objects.

The second 'book' in the 'two book' model is an abstract, mental 'dictionary'. It is usually referred to in linguistics as a 'lexicon' - but this technical term simply means 'dictionary', so that adopting it adds nothing except a misleading implication that a 'lexicon' is in some way superior as a model to a 'dictionary'. The lexicon contains lists of spellings, pronunciations and characterizations of the meanings of words, together with how they fit into the grammar, and it is typically arranged in alphabetical order. In other words, the 'dictionary' model of language is an even more familiar physical object than the grammar - - the dictionary on your bookshelf.

The physical 'two book' model of language is still accepted by very many linguists, apparently quite uncritically. And this is despite the fact that one of the indisputable advances in linguistics of the last few decades has been the clear establishment of the fact that grammar (in the sense of 'syntax') and words (more technically, 'lexis') are inextricably interwoven. But what makes the 'two book' model of language even more misleading is that it typically leaves out altogether two other types of 'form' in language which are also interrelated, i.e. intonation and punctuation. There are, of course, honourable exceptions, such as Quirk et al (1985.)

Sadly, it seems from the books that are currently being published in linguistics that relatively few linguists are actively exploring alternatives to the 'two book' model. Too often a language is presented as if it were an abstract static object, rather than a functioning phenomenon - and too often it is seen as consisting essentially of something very like a grammar book or a dictionary or, in the full 'two book' model, both.

The model that I will set out in what follows is very different. If there is a single VISIBLE image that should replace the 'book' image, it is that of a network - a FUNCTIONING network. But it must be located at the appropriate place in the abstract functioning model that I will describe now.

2.4 Modelling language as a computer program

In the last two decades of the twentieth century a totally new means of modelling abstract functioning phenomena became widely available. This was the computer - or, more accurately, the COMPUTER PROGRAM. A group of researchers at Cardiff University have been working for the last fifteen years to develop a model of language which is (1) CONCEIVED OF AS BEING LIKE A COMPUTER PROGRAM - and which is also (2) actually IMPLEMENTED AS A COMPUTER PROGRAM. This handbook draws directly on this work - as well as on the research involving very extensive **text analysis** that was mentioned in Chapter 1.

The field of modelling the human mind in computers is called **artificial intelligence**, and the sub-field that is concerned with the role played by natural human languages is called **natural language processing** - or, from the viewpoint of linguistics, **computational linguistics**. Within this field of study there are two complementary sub-fields. The first is **natural language understanding** (NLU), and this field has been in existence for over half a century. The second has developed only in the last twenty-five years or so, and it is **natural language generation** (NLG). The task of NLG is to model, in the form of a computer program, the way in which we turn a person's 'communicative purpose' and its

associated 'propositional content' into a spoken or written text - and to do this in a way that is appropriate to the various types of situation in which we communicate with our fellow human beings. This is an extremely hard task, and we who are working in this field still have a long way to go before we can report anything approaching complete success. (The same, incidentally, is true of the field of NLU, even though it began much earlier and has many more researchers working in it.)

Which of the two is the most fundamental - generation or understanding? It seems clear to me that it is the process of generation, or production. What is the evidence for this claim? Firstly, notice that is possible to produce a text without anyone's receiving it - but you can't receive a text unless someone has produced one. For example, you can shout a message to someone who doesn't hear it, or you can write a letter to someone and then decide not to send it.[4] Secondly, the starting point and the end point of an act of communication is the representation of the communicative purpose and the content of the message in the minds of the 'Performer' and the 'Addressee'. (See Chapter 3 for the reasons for this choice of terms.) And the communicative purpose and the propositional content are clearly those that are put into the text by the PERFORMER - rather than by the Addressee (until it is his or her turn to be the Performer). It therefore follows that an outsider who is seeking to understand any given text will get more from an understanding of the process of how the Performer produced it than from how the Addressee might go about understanding it. In other words, programs of the type that researchers in the computer generation of language develop provide the best currently available model for understanding the nature of language and text. This, essentially, is why my colleagues and I now do a large part of our linguistics in the framework of developing a very large computer model of NLG.

The exploration of this new medium for modelling language has brought many insights. One of the most significant of these is the confirmation that the view of language found in Michael Halliday's **systemic functional linguistics** (SFL) is, as I had long thought, the most helpful of all existing theories of language. However, Halliday's explicitly 'social' view of language needs to be complemented by an equal emphasis on the fact that it is through our individual human minds that we communicate with, typically, other individual human minds - i.e. we need a 'cognitive-interactive' approach to language.[f] In my view the key concept in understanding human language is that it is not simply an abstract object that is stored in the human mind, but a PROGRAM FOR SOLVING PROBLEMS - a program that guides us, in the very greatest detail, through that part of our social behaviour that is enacted though language. At its heart lie a set of an algorithms (or 'programs') with the basic structure of a 'decision tree' (but with certain vital additions). From our present viewpoint, the most important of these decision trees is the one which offers the user of a language the set of CHOICES BETWEEN MEANINGS that are expressed in words, grammatical structures and intonation, etc - and in SFL this is called a

4. Of course, it may happen that you come across a letter that is not addressed to you, and in such cases you can probably get some sort of sense out of it. But the fact remains that the letter will normally have been written by a particular writer with a particular message for a particular reader at a particular time, etc., and any adequate explanation of the letter will require you to relate the text to these factors - if only by guesswork.

'system network'. And, as is common in modelling decision-making, we know about the general probabilities of making one choice rather than another - and we can incorporate these **probabilities** in the system network.

That last paragraph introduced several important new concepts, and we will explain them more fully very soon - in Section 3.[g]

2.5 A new problem: how to describe, in a book, something that is like a computer program

I have said that the most helpful model of language is a computer program. If this is so - an I believe that it is - I have an immediate problem because of the fact that the medium though which I am now communicating with you is a book. In particular, I must avoid falling into the trap of presenting you with the 'two books' model of language.[h]

Clearly, we must make the best of the modelling tools that are available to us. And, whatever the advantages of computers for various tasks, books such as this one still have their uses. (For a start, you don't have to plug them into an electricity supply at regular intervals to use them.) So how can we best represent a computer program in a book?

We can make a useful start by recognizing that computer programs are detailed implementations of higher level programs, and that these are often represented in **flow charts**. And the diagrams that follow in the next section can be regarded as roughly equivalent to a high-level flow chart.

However, most of the diagrams in the rest of this book are very different. What they represent is not language itself but THE RESULTS OF THE USE OF LANGUAGE. They show the internal structure of short texts - usually texts of the length of a clause, but sometimes even shorter texts. These, as we will see, are representations of the **input** to or **output** from one of the boxes in which the work of generating or understanding a text is done, as shown in the overall diagrams of language to which we will come shortly. In other words, these are not programs, but objects - though objects of a pretty abstract nature.

So the main visual models used in this book will be old-fashioned two-dimensional pictures, and they will consist of lines and other symbols that will be explained, with labels drawn from a well-defined **metalanguage** (i.e. a language for talking about language). However, just as a two-dimensional picture can represent a three-dimensional physical object, we will find that these diagrams are able to represent several 'dimensions of meaning' at the same time. This is especially true of the *Functional Semantics Handbook*, in which I introduce the **multi-strand analysis** approach to analyzing the meanings of texts. But the present volume will also illustrate, at many points, the multifunctional nature of language.

I pointed out in Chapter 1 that a particular feature of this book is the very frequent use of diagrams. Normally these figures will be supported by explanations in ordinary language - but these explanations will assume, at every point, that THE TEXTS THAT WE ARE EXAMINING ARE THE PRODUCT OF A PROCESS OF TEXT GENERATION IN A HUMAN MIND, FOR THE PURPOSE OF COMMUNICATING WITH ANOTHER PERSON. And the best model available to us for language itself, at the start of the twenty-first century,

is that of a fairly complex computer program.

Let's now begin on the process of building a model of language.

3 Towards a systemic functional model of lexicogrammar

3.1 Saussure: form and meaning

Saussure is widely recognized as the 'father' of modern linguistics, and, in various degrees, his ideas have influenced most of the theories of language developed in the 20th century. His influence on the version of Systemic Functional Grammar presented here will quickly become clear in what follows.

Saussure's central concept is the 'linguistic sign'. For Saussure, any 'sign' consists of a 'signifier' and a 'signified' - i.e. a **form** and a **meaning**. Indeed, it was he who showed us that forms and meanings are mutually defining.[i]

It follows that, if our goal is to model language, we need to model a very large SET of linguistic signs. In other words, just as a single SIGN has both a form and a meaning, so too a SIGN SYSTEM such as a natural human language has the two levels of form and meaning. Figure 1 can therefore be seen as a very simple model of A LANGUAGE AS A WHOLE. But it is deficient in at least one key respect..

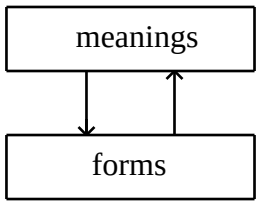

Figure 1: a simple model of any sign system

Because a sign system is not a static object but something that is designed to be used, Figure 1 includes two arrows. These indicate - rather inadequately, as we will see in a moment - that the function of a sign system is to turn meanings into forms, and forms into meanings. For example, a simple traffic control system has just two **forms** - a red disk and a green disk - and two **meanings** - which we can express crudely as 'stop' and 'go'. A human language is of course very much more complex than this, in both its forms and its meanings. But it is often useful to recall this simple, two-level model when you are working on a problem within some sub-component of the overall system. The concept of a language as a 'sign system' summarizes neatly the vital point that we cannot expect to understand the **forms** of language without considering the **meanings** of language - and vice versa.

3.2 Potential and instance (and Saussure's *langue* and *parole*)

In adding the concept of a 'sign system' to the concept of a 'sign', we have already introduced a second vital distinction made by Saussure (though he made it in terms of *langue* and *parole*). This is the distinction between a **language** taken as

a whole - i.e. as a **resource** for communicating meanings to our fellow human beings - and an **output** arising from the use of that resource - i.e. a **text**. A text is defined here as AN INSTANCE OF LANGUAGE IN USE, and it may of course be spoken or written. It therefore follows that, just as a language has both a level of **form** and a level of **meaning**, so too does a text - as is illustrated in Figure 2.

In **systemic functional linguistics** (SFL) we also express this distinction between a language and a text as the distinction between a **potential** and an **instance** of that potential. These two concepts can be applied to both the level of meaning and the level of form. It therefore follows that our model of language will have at the level of meaning a component that specifies what Halliday has aptly termed the **meaning potential** of the language. And it is here that the **system network** of choices between **semantic features** is located - i.e. the system networks model the **potential** at the level of **meaning**. This, then, is the core of a systemic functional grammar.

Since there is a **potential** at this level, we should logically expect that there will also be **instances** at the level of meaning - and indeed there are. These take the form of the set of semantic features that are chosen ON ANY ONE PASS through the system network - this set of features being called a **selection expression.** In Figure 2, the two top boxes show (1) that the potential is a system network of semantic features and (2) that an instance of that potential consists of the set of semantic features selected on a traversal of the network. The component of the language is shown in a box with right-angled corners, and the output from that component is shown in a smaller box with cut-off corners. In this handbook, where our topic is syntax, I will only occasionally refer to the system networks, but the main treatment of these comes in the *Functional Semantics Handbook*. However, the little **lexicogrammar** given in the next section provides an introductory picture of what a system network is and how it operates. (Perhaps I may remind you again that, in SFL, the term 'grammar' is simply a short form for 'lexicogrammar.')

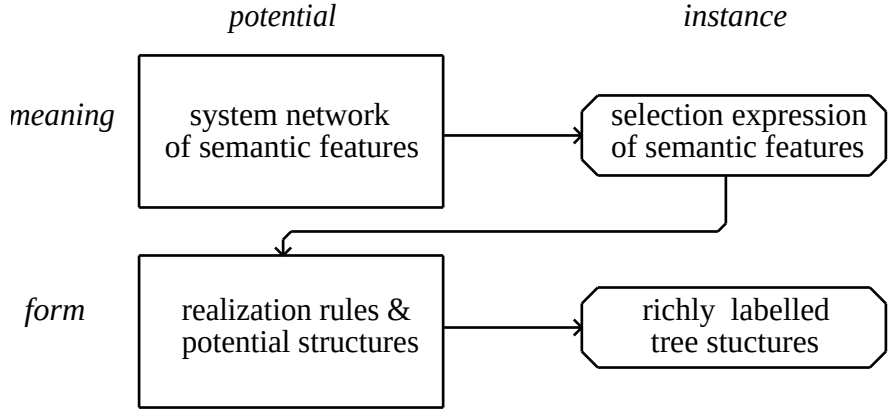

Figure 2: the core components of a simplified systemic functional grammar

We turn now to the level of **form**. 'Form' is a general term that includes **items**, **syntax** and (depending on whether the medium is speech or writing) **intonation** or **punctuation**. Figure 2 shows that here too - as we would logically

expect - there is both a potential and an instance. The **form potential** of a language consists of the **realization rules** and **potential structures -** concepts that are better explained through examples, as in done in the next section. Here we should simply note that the two levels of form and meaning are connected through the fact that the **output** from the level of **meaning** is the **input** to the level of **form** - more precisely, to the **form potential**. The role of the realization rules is to convert the selection expression of semantic features into a layer of the tree diagram representation of the sentence that is being built up. Thus the little grammar in Section 5 has the potential to generate eighteen different nominal groups. These output diagrams are labelled sufficiently richly for the analyst - or a computer parser - to be able to work out the various functions that each element serves, and it is on these elements - and specifically the elements of the clause - that this paper focusses. These, then, are the **instances** at the level of **form**.

The most typical instance at the level of form is a 'sentence' - a sentence frequently being a single **clause** - such as *I've been discussing that new student with Peter*. But we can also treat a **group** of words such as *that new student* as an instance, or even a single **word** such as *that, new* or *student.*

Figure 2 therefore brings together, in a single diagram, Saussure's two key pairs of concepts: **meaning** and **form,** and **potentia**l and **instance**. This simple model is in fact also the outline of a **systemic functional grammar**, and it summarises the model of language within whose framework the entire contents of this handbook are to be understood. However, since the book focusses on the analysis of text, the parts of Figure 2 that we will be most concerned with are the instances. This handbook covers instances at the level of **form** and the *Functional Semantics Handbook* covers instances at the level of **meaning**. But even in the present volume we will frequently be reminded of the interdependence of form and meaning, since the structures of syntax are as they are because of the meanings that they express.[j]

3.3 Saussure: Paradigmatic and syntagmatic relations in language

However, Saussure also made a third major distinction between two pairs of concepts: that between **paradigmatic** and **syntagmatic** relations in language. It is a particularly important one for systemic functional linguists, because it introduces the concept that is the most central of all in understanding how Systemic Functional Grammar differs from other functional approaches to language.[k] Interestingly, most theories of language other than Systemic Functional Grammar concentrate almost entirely on syntagmatic relations. The consequence is that they pay very little attention to paradigmatic relations - even though it is clear that for Saussure the two were equally important, and indeed complementary. I believe that Saussure was right, and the model presented here reflects this belief.[l]

Paradigmatic relations are relations of contrast. Unlike syntagmatic relations, paradigmatic relations exist only in the **potential** and never in an **instance**. From the viewpoint of the text analyst, they express a contrast between the meaning (and so the form) that WAS chosen for use in the text and the one or more meanings (and so forms) that MIGHT have been chosen (but were not). In

other words, they exist only in the **language** that is used to produce a **text-sentence** - and not in the text-sentence itself. ile there are relations of contrast at the level of **form**, in a systemic functional grammar it is the contrasts between **meanings** that are seen as the generative base of the lexicogrammar. The result is that the purely formal contrasts in a language play no role in how the grammar operates in the generation of a sentence. (The term 'sentence' is used throughout this *Handbook* as a short form for 'text-sentence'; i.e. a sentence is always to be viewed as an instance at the level of form, unless it is specified that it is being considered at the level of meaning.)

We must now ask the question: "Are these paradigmatic relations at the level of **meaning** or at the level of **form**, or both?". The answer is there are clearly CONTRASTS at the level of form, e.g. as between the forms of *student* and *students.* But in a systemic functional grammar it is the contrasts between **meanings** that are seen as the generative base of the lexicogrammar. The result is that the purely formal contrasts in a language play no role in how the grammar operates. The next section illustrates the way in which the concept of **choice between meanings** is the key concept in a systemic functional grammar.

The nearest we can get to having a paradigmatic relationship in an **instance** of language is to have a contrasting PAIR of examples - a 'minimal pair', to introduce a term that is more frequently used in discussing the sounds of a language (i.e. in phonology).[5] So in a pair of examples such as *I've been discussing that student with Peter* and *I've been discussing those students with Peter* there is a paradigmatic relationship - and so a difference of meaning - between *that student* and *those students.* Specifically, the contrast occurs within a **nominal group** (this being roughly equivalent to a 'noun group' or 'noun phrase' in other terminologies). The key point is that in any given nominal group you cannot have both *that* and *those,* or both *student* and *students* - i.e. you must choose between them. Here the difference affects more than one word, and it is the difference between the two meanings of 'singular' and 'plural'.

Typically, contrasts between **meanings** are realized as contrasts between **forms**. But the fact is that there is often not a straightforward relationship between a set of meanings and the equivalent set of forms. And this in turn means that you cannot rely completely on the contrasts between the forms of a language as a way of thinking about the contrasts between meanings. The contrasts between forms are merely a rough guide to the contrasts between meanings, and in a systemic functional approach it is the contrasts between meanings that are the heart of the lexicogrammar.

Why should we give this priority to paradigmatic relations between meanings? It is because, when someone is planning a sentence, it is clearly the contrasts between alternative meanings between which he or she chooses - rather than the contrasts between forms. For example, if two outputs display a contrast between forms, as in *that student* and *those students*, the only significance of the contrast in

5. For example, we know that English makes a distinction between the phonemes /b/ and /p/, because the difference between the two is needed to show the difference of meaning between the 'minimal pair' of words /bet/ and /pet/.

form is that it expresses a contrast in meaning. In other words, the difference between 'singular' and 'plural' is ultimately a difference of meaning rather than form.

A model of language that is built around choices between **meanings** is therefore inherently more insightful than one that is built around contrasts between mere **forms**. Indeed, it is one of the most important characteristics of natural languages such as English (in contrast with artificial languages such as computer 'languages' and the 'languages' of logic) that meanings and forms do **not** correspond exactly. If they did there would be no ambiguity, and modelling language would be much easier - and also much less fascinating. A systemic functional grammar, then, is one that is built around the core concept that the meaning potential of a language is best modelled as a NETWORK OF CHOICES BETWEEN MEANINGS. (We will come back to the implications of the word 'network' shortly.)

We will now look at the concept of **syntagmatic relations**. There are in fact two aspects to syntagmatic relations in language: **part-whole relations** and **sequential relations**. The more fundamental concept is that of part-whole relations. A widely used term with a similar meaning is 'structural relations', and the everyday meaning of the word 'structure' conveys very effectively the core concept of syntagmatic relations. This is the relationship of a number of 'parts' to some 'whole' of which those part are components (e.g. Subject, Main Verb and Complement in a clause, and determiner and head in a nominal group, to mention just some of the elements of those two units). In broad terms, this is the relationship known traditionally in linguistics as 'constituency'.[m]

Given the framework of concepts that we established in the last section, it is natural to ask: "Are syntagmatic relations important at the level of **meaning**, or at the level of **form**, or both?" And then, for each of the two levels: "Are they important in a model of language as **potential** or in text as **instance** - or both?" The answer is that part-whole relations are relevant at both levels and in both the potential and instances - though they are more complex and more dominant at the level of form. And the relationship of sequence is, in principle, only relevant in instances at the level of form.[n]

The part-whole relationship is found at the level of meaning in the way that 'things' function as components of 'situations'.[o] But at the level of form we find both the part-whole relationship and the relationship of sequence.[p] You can see this most clearly in terms of a specific example, such as *those students* - i.e. in an **instance** at the level of **form.** This is because this gives us a physical object that we can actually see, and whose parts we can examine. **Forms** of language are always easier to think about than **meanings**, because forms are one stage more concrete than meanings.

This illustrates an important fact about the study of language. This is the fact that it is easier to think about **instances** than about the **potential**, and that it is easier to think about **forms** than **meanings**. So when we are working at the level of form we can usually produce an example consisting of a normal string of words to illustrate whatever point we are making, but it is more of a challenge to have to produce a representation of an example at the level of meaning, i.e. the set of semantic features that must have been chosen in generating that string.

This handbook and its companion the *Functional Semantics Handbook* are about the analysis of texts, and it will now be clear that the natural way into the study of language is to start from the study of **instances**, rather than the study of the **potential**, and to start with instances at the level of **form** rather than instances at the level of meaning. The goal of this *Handbook* is therefore to show how to describe instances of language at the level of form, and the goal of the *Functional Semantics Handbook* is to build on these to show how to analyze instances at the level of **meaning**.

In modern linguistics, the structure of sentences is usually shown by 'tree diagrams' (though other notations are occasionally used too). This use of a 'tree' as a way of modelling the structure of a sentence is by now so well established that we have largely forgotten that the origin of the metaphor is a living thing in the natural world - i.e. an oak tree or a fir tree. So it may be worth pointing out that it is in fact a rather odd model to use for analyzing sentences, in that tree diagrams are typically drawn in such a way that the 'root' of the tree is at the top of the diagram and the 'leaves' are at the bottom! However, even though the term 'tree' is an odd description, TREE DIAGRAMS are by far the most insightful of the various visual representation of the structural relations in sentences that have been used. The rest of this handbook contains numerous examples of such diagrams.¶

We can summarize this section by saying that, in terms of Figure 2, the place of **syntax** in an overall model of language is in the **syntagmatic** relations within an **instance** at the level of **form**.

4 Some other major characteristics of this model of language

4.1 A lexicogrammar as a 'sentence-planner'

We have seen that this theory of language gives theoretical priority to choices between meanings - rather than choices between forms. From the viewpoint of modelling language in use, there is a great advantage that follows from this. It is that the choices in the system networks of language can be seen as corresponding to the last major stage of the several stages of decision-making through which a person must go - completely unconsciously, typically - WHEN PLANNING A SENTENCE. This is what gives this type of model of language its explanatory power.

The corollary of making choices between meanings the generative base of the lexicogrammar is, as I have pointed out earlier, that there is often not a simple one-to-one relationship between a meaning and the form in which it is expressed. In practical terms, this means that the **realization rules** that handle these relations are sometimes more complicated than they would be if the contrasts were simply at the level of form. (Examples of simple cases of such rules are given in Section 5.) This is the price to be paid for having a model of language in which the system network expresses choices between meanings - and all systemic functional grammarians would agree that this is a small price that is well worth paying for the greater insights that it makes possible. In a systemic functional grammar, then, it is the **meaning potential** - which is modelled as a great network of semantic features - that constitute the core of the model of language. It is the work of the

realization component to apply its rules automatically, in order to ensure that the output is a natural text.[r]

The main visible difference between Figure 1 and Figure 2 is the addition of the distinctions between the **potential** and the **instances**. One important effect of this addition is to make it clearer that the systemic functional model of language is inherently oriented to the **production** (or **generation**) of texts, rather than to **understanding** them. This expresses the view - which I first set out a few pages back - that a language is, essentially, a program for turning 'concepts' (or, more accurately, a communicative purpose and its associated propositional content) into a string of words, punctuation and spaces.[6]

4.2 The role of the lexicogrammar in understanding text

If the model set out in Figure 2 is inherently a model of language from the viewpoint of GENERATION, this question arises: 'How can it be used WHEN THE ADDRESSEE IS INTERPRETING AN INCOMING TEXT?' The answer is that a program with a different set of components from those in Figure 2 is used - but they are components that are systematically DERIVED FROM the components in Figure 2. The reason why a different program is needed is that the set of problems that have to be solved when a person (or a computer) interpreting an incoming text is very different from the set of problems that arise when a person (or computer) is generating a text. Assuming a written input, the first stage is to turn a string of words into a tree representation of each text-sentence.[7] The component that performs this task is the **syntax parser**, whose task is to discover the syntactic structure that relates the string of words that is the input to the parser to each other. Here, interestingly, the equivalent component in generation is so trivial that it is usually ignored: it is the process of stripping away the syntax that has been used to get the words into the right sequence. (In the diagram in Figure 2, the output at the level of form is shown as 'richly labelled tree structures'; the syntax stripper is applied after this, and it simply strips away everything except the words and the punctuation (or intonation, in spoken discourse). Getting rid of the syntax is easy; the hard part is the Addressee's task of reconstructing it from the scanty evidence of the sequence of items and spaces that is the input to the parser. Essentially, the working of the parser rests on a strategy for discovering the likely structural relations between items as the Addressee works his/her way through each text-sentence from the beginning to the end. And then we need to discover, from the tree structure, what semantic features have been chosen in generating it - and this is the task of the **semantic interpreter**. Essentially, this is the operation of the realization rules and the traversal of the system network in reverse. In other words, understanding a text involves understanding not only its meanings but also how those meanings were converted into words.[s]

6. Here I am assuming that the text is in the written mode; the spoken mode is more complex.

7. While the stages required to process a spoken and a written input are in principle the same, in practice a spoken input requires an extra stage. One of the reasons is that spoken text brings with it the additional problem of segmenting the speech flow into words.

4.3 The unity of the model

Another significant characteristic of the model set out in Figure 2 is that it has dispensed completely with the 'two-book' model of language. There is no separate 'lexicon' - and in this it is unlike most models of language that are not systemic functional models. This difference arises because the model of language presented here is not a 'grammar', in the strict sense of the word, but a **lexicogrammar**. The concept of a 'lexicogrammar' is one that is only found in systemic functional linguistics. So when the term 'grammar' is used in this handbook it is used as a short form for 'lexicogrammar' (and not, therefore, to refer to the 'grammar component' rather than the 'lexis component' of the model). This follows naturally from taking 'choice between meanings' as the fundamental organizing principle of language. The old 'two book' model of language reflects a difference at the level of form - not meaning.

One great advantage of the present approach is that it reflects the fact that there is no clear boundary in English - or indeed in any other language, so far as I am aware - between items that are 'grammatical' (which traditionally belong in the 'grammar') and items that are 'lexical' (which are traditionally placed in the 'lexicon'). For example, in the 'two book' approach some of the words that would appear in the context of *She is ... clever*, would be said to be 'grammatical' and some to be 'lexical' - but it is simply a distinction that it is misleading to make. Some of the possible words are: *as, so, too, less, very, quite, fairly, really, relatively, terrifically, impressively, amazingly*, etc. What these words have in common is that they are **items** rather than **structures** (as opposed to being 'lexical' rather than 'grammatical') - as well as the fact that they all function as the same element of structure as each other in this example. In a systemic functional grammar, then, all items are integrated into one great lexicogrammar, so that there is no need to try to decide for each word whether it is 'grammatical' or 'lexical', and so whether it should be dealt with in 'the grammar' or 'the lexicon'. Once again, then, we can see the unfortunate influence of the 'two book' model of language.

Moreover, the same integrative principle applies to meanings that are expressed in **intonation** or **punctuation** - two additional types of 'form' that occur with spoken and written texts respectively. Indeed, different languages may use either an item, a syntactic structure or an intonational marker to express what is essentially the same meaning. Any model of language that omits intonation and punctuation is therefore deficient. In the full systemic functional grammar on which this handbook is based, the system networks that model the meanings of intonation and punctuation are integrated fully with those that are expressed in structures and items - and realization rules of the same type integrate their forms with the rest of the output from the lexicogrammar.[t]

4.4 One pass through the system network generates one syntactic unit

Finally, I have to point out that the simplified diagram in Figure 2 has one major limitation. This is that it shows ONLY ONE PASS THROUGH THE GRAMMAR - and

so the generation of only one unit. What the diagram does not show is that the cycle from meaning potential to the instance in form must be repeated FOR EACH SEMANTIC UNIT - AND SO EACH SYNTACTIC UNIT - that is to be generated. You can add this to Figure 2 - either mentally or on the diagram itself - by drawing a line that runs out from the left side of the 'realization rules' box and up into the left side of the system networks box. Then add an arrowhead and label the line 're-entry'. In other words, there are TWO main types of realization: (1) the direct generation of structures, items and intonation or punctuation, and (2) the re-entry to the system network to make more choices between meanings, and so to generate more forms, for a lower unit in the tree structure.

The next section provides a small working example of how a lexicogrammar operates. It shows the generation of a simple nominal group. What it does NOT show is the clause in which it might typically occur. To provide a context for it, let's say that two overworked lecturers in a university meet briefly in a corridor, and the following conversation takes place:

A: *Where have you been for the last half hour? I've been looking for you everywhere.*

B: *I've been discussing .... with Peter.*

The row of dots gives us the context, in a simple clause, of our example.

5 A fragment of a working systemic functional grammar

5.1 The system network

The little example of how a systemic functional grammar works that is presented here is taken from a much fuller lexicogrammar for 'things' in English. A **thing** is a semantic unit that is typically realized at the level of form in a **nominal group**. ile the source grammar for 'things' has over 150 systems that are realized grammatically (and many hundreds more that are realized lexically), the present grammar has just four systems that are realized in grammar and two that are realized in lexis. The highly simplified model to be presented here therefore includes only a few of the many meanings and structures that are expressed in the nominal group. Specifically, they will be restricted to a few that use the to elements of the **head** and the **deictic determiner**. We will meet the other types of 'determiner' in Chapters 6 and 14.

The key concept of a systemic functional grammar is a **system** - this term being used here in a technical sense where it means a 'choice between two or more semantic features'. The heart of the grammar is therefore the system network of semantic features. Consider the little example shown in Figure 3:

SYSTEM NETWORK OF MEANINGS

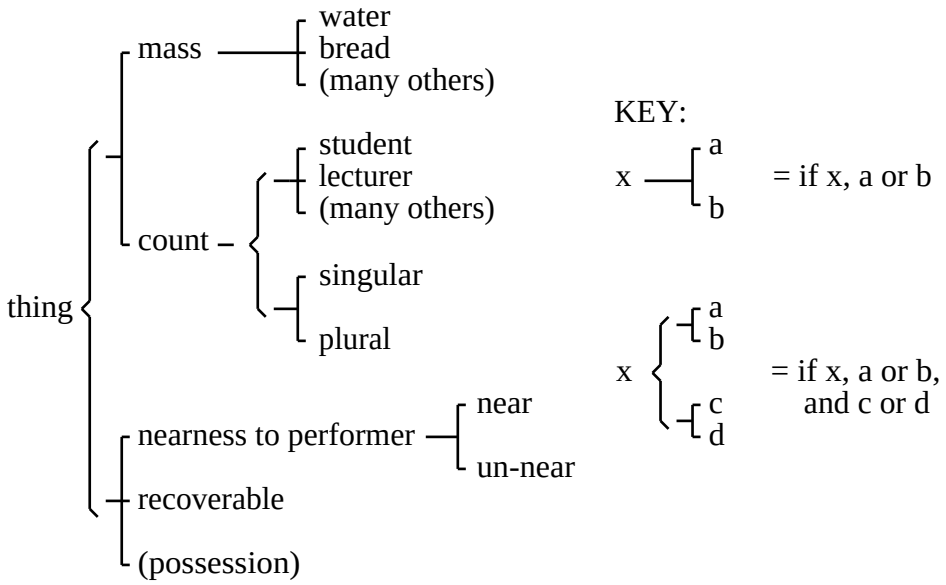

Figure 3: A highly simplified system network for the 'thing' in English

Each system in the system network has an **entry condition** and, since this is typically a feature in another system, a number of systems of related features typically combine to create a **system network**. Furthermore, as Figure 3 shows, there can be parallel entries to more than one system. IT IS THE SYSTEM NETWORK OF SEMANTIC FEATURES THAT MODELS THE MEANING POTENTIAL OF A LANGUAGE.

The system network corresponds to the top left box in Figure 2. The way to use it is to 'traverse' it, starting with the leftmost feature. Whenever an 'and' bracket is encountered all the systems to its right must be entered, so that the pathway through the network typically becomes a set of branching pathways.

When you have completed a traversal of the network, you will have collected a **selection expression** of semantic features such as:

[thing, count, plural, student, nearness to performer, un-near].

Notice that features are typically written in square brackets, to show their status as features. This **output** from the network corresponds to the top right box in Figure 2 - i.e. it is one **instance** of this little lexicogrammar's **meaning potential**. This small network of six systems will generate eighteen different selection expressions, each of which constitutes a representation at the level of meaning of a different nominal group. This is not very many - but clearly it could be quickly scaled up by the addition of other systems. The full network for 'things' in the Cardiff Grammar (from which this is derived) currently has around 120 semantic features - most of which have an associated realization rule - and it generates many millions of different nominal group structures.

5.2 The realization component

This selection expression of features becomes the **input** to the **realization** component. This is the bottom left box in Figure 2, and it contains two main types of statement: (1) **realization rules**, as given in Figure 4, and (2) **potential structures**, which simply show the sequence in which those elements that are fixed in sequence must appear (such as those in the nominal group) Here we will introduce an oversimplified potential structure, which can be summarized as:

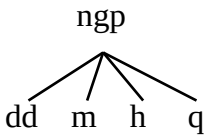

This simply means that when any of the four elements of a **deictic determiner (dd)**, a **modifier (m)**, a **head (h)** and a **qualifier (q)** occur in a **nominal group (ngp)**, they come in this sequence. An example would be: *those* [dd] *very fresh* [m] *oranges* [h] *from Spain* [q]. A fuller nominal group may have many more elements. These facts about the nominal group apply potentially to all nominal groups, and the realization rules given in Figure 2 use the potential structure to ensure that the elements occur in the right sequence.[8]

The first rule in Figure 4 inserts the nominal group itself into the structure being built. (In our example of *I've been discussing ... with Peter*, we would at this point be generating the nominal group to fill the Complement following the Main Verb, i.e. *discussing*.)

8. It may be of interest to point out that there is no equivalent potential structure for the clause, because very few elements occur at fixed places in sequence.

FEATURE	CONDITIONAL FEATURE(S)	REALIZATION
thing		insert ngp
water bread student lecturer		h < "water" h < "bread" h < "student" h < "lecturer"
near	singular or mass plural	dd < "this" dd < "these"
un-near	singular or mass plural	dd < "that" dd < "those"
recoverable		dd < "the"
plural		h < "+s"

Figure 4: Some simplified realization rules for the 'thing' in English

Each of the next four rules assigns one of the four representative examples of nouns given here to the 'head'. Specifically, the rule for [water] says: 'The head is *water*' - or, more technically and so more precisely, 'The head is expounded by the item *water*'.[9]

The two rules below these four illustrate the important fact that THERE IS NOT ALWAYS A SIMPLE ONE-TO-ONE RELATIONSHIP BETWEEN MEANING AND FORM. The first rule says that, if the 'thing' is either a 'singular' thing' or a 'mass' thing (such as 'water' or 'bread', each of which is viewed in English as an uncountable 'stuff'), then the deictic determiner will be *this* - but if it is a 'plural' thing, the determiner will be *these.* And there is a similar pattern for the meaning of 'un-near'. However the meaning 'recoverable' in the next rule is simpler, in that it is always realized by *the.* Notice that the meaning of *the* is something like 'If you inspect the rest of this nominal group, you will know which one I am referring to', and it is this that the feature [recoverable] summarizes.

The final realization rule is another simple one, and it says 'If the 'thing' is 'plural' the head is given the morpheme *s* as a suffix'. Notice that there is no need to specify that this doesn't apply to 'mass' things such as 'bread', because the system network ensures that 'mass' things do not enter the system in which 'singular' or 'plural' is chosen.

Clearly, this little grammar leaves out a rather large proportion of the many

9. I could have used an abbreviated definition or an expression such as 'H_2O' to characterize the 'core' meaning of the item *water* (i.e. as a 'clear, tasteless liquid essential for life'), but here I have followed the simpler method of using the form of the word, placed in single quotation marks, to indicate that this is a meaning (i.e. a 'sense') of the word and not its form.

complex meanings that can be expressed through the nominal group in English. Equally clearly, it ignores various problems, such as the plurals of words like *box* and the irregular plurals of *men* and *women*, etc). All of these matters are covered in the full network from which this simplified one has been taken. But the fact that this little lexicogrammar is quite limited in its coverage of English nominal groups doesn't matter, because our purpose here is simply to illustrate the BASIC PRINCIPLES of how a grammar that is founded on the concept of 'choice between meanings' actually works. The key concept, then, is that the system network of a language (or any other sign system) defines the 'meaning potential' of that language, and the realization component defines the 'form potential'. So a lexicogrammar also specifies, when it is set to work, the **instances** that are possible at the levels of both **meaning** (in the selection expression of semantic features) and **form** (in the string of word forms).

The little grammar of the nominal group that is set out here shows only what happens in a single **unit** - the nominal group. If the output is to be a sentence with more than one layer of structure - as in any natural test-sentence - then the network will first be entered to generate a **clause** and its elements, and then RE-ENTERED to generate any 'lower' unit that is required. This may be a **nominal group**, as here, or one of the other types of group that we will meet. The important point is that each such re-entry to the network will add a new unit to the structure, each with its own internal structure. And, as the later chapters of Part B show, language has an amazing ability to embed UNITS WITHIN UNITS - groups within groups, clauses within clauses, clauses within groups, and so on. Yet - remarkably - the simple principles for generating structure that have been illustrated here apply in all of these cases.

6 Should we start from form or meaning?

As the Preface and Chapter 1 have made clear, the relationship between this and its sister volume *The Functional Semantics Handbook* is that we begin, in the present volume, with the level of form - and so with the words of English and their relationships to each other through syntax (to oversimplify a little) - and we then go on to ask what these forms mean.

But in a functional approach to understanding the nature of language you might think that the natural approach to explaining it would be to start from the level of meaning and then to say how the meanings are realized at the level of form. This is certainly one possible approach - and if this book had been intended to explain the use of Systemic Functional Grammar to GENERATE Text-sentences, that would have been the only possible approach.[u]

So why did I decide to take the approach that I have?

The reason is that, if our goal is to be able to ANALYZE text-sentences, we need to start from the level of FORM. The reason is that, when we start to analyze a text, our starting point is a string of spoken or written symbols - i.e. the most explicitly formal part of the level of FORM of the language. In starting from the level of form, then, we are engaging in THE CONSCIOUS EQUIVALENT OF THE UNCONSCIOUS ACTIVITY OF ANALYZING THE INCOMING TEXT-SENTENCES THAT WE HEAR AND READ THROUGHOUT OUR

EVERYDAY LIVES. It is the natural way to approach language.

However, because form and meaning are so closely interwoven in language, it is helpful to have as good a picture as possible of the MEANINGS that are being expressed by the forms that we are analyzing. This is especially true WHEN A SINGLE FORM MAY STAND FOR TWO OR MORE MEANINGS. This is the concept of ambiguity, or, more accurately, **polysemy**. We will meet examples of this in Section 7.9 of Chapter 3, when we will note that certain **lexical verbs** have more than one meaning, and we will meet other examples when we look ate **modal verbs** (in Section 4 of Chapter 4, and also in the forms and meanings of MOOD in Chapter 5 - and indeed at many other points in this book.

7 The multi-strand approach to meaning

In showing how syntax works in English I will need to bring into the picture, in a systematic way, the various **functions** that syntax serves. This means that I will introduce at the appropriate points of the book the several **strands of meaning** that run through every text - most clearly in the unit of the clause. ile each language of the world differs from each other in how it uses syntax to express meaning - in large or small ways - all languages have in common the fact that they use syntax to express many different strands of meaning.

To make a complete analysis of a text, we must attend to at least EIGHT different major strands of meaning. These are: **experiential** (including **temporal**) meaning, **interpersonal** meaning, the meaning of **polarity**, the meaning of **validity**, **affective** meaning, the meaning of **logical relations**, and the **thematic** and **informational** strands of meaning.

I should add, however, that Halliday recognizes just FOUR major strands of meaning (or 'metafunctions'), these being the 'experiential', the 'logical', the 'interpersonal' and the 'textual'. Roughly speaking, my eight strands of meaning are sub-divisions of Halliday's four. In this approach my last two correspond to Halliday's 'textual', and my' interpersonal', 'validity' and 'affective' strands are all included in his 'interpersonal'.[v]

As this book will demonstrate, it is important to give each of the eight major strands of meaning the status of being a strand of meaning in its own right, because each reflects a very different range of semantic choices. Figure 5 shows the eight **strands of meaning** mentioned above, together with the names of the main **system networks** within each that are expressed in the clause.

However, this is not the moment to try to spell out what these strands of meaning are. Instead, we will gradually build up an understanding of what each contains as the book progresses - especially in Chapter 4. And by the end of Chapter 5 you will have met all of the eight strands of meaning shown in Figure 5.

8 Summary

Section 5 provides an introduction to the way in which the model of language on which this *Handbook* is based actually works. But it covers only a tiny corner

of language. As a correction to this, Figure 5 provides an overview of the major central types of meaning in English. As we will see by the end of Chapter 6, it is the simultaneous realization of specific meanings selected from within each of these eight broad strands of meaning that constitutes syntax. Figure 5 can therefore be seen as an outline map of the contents of the top left box in Figure 2 (in Section 3 of this chapter).

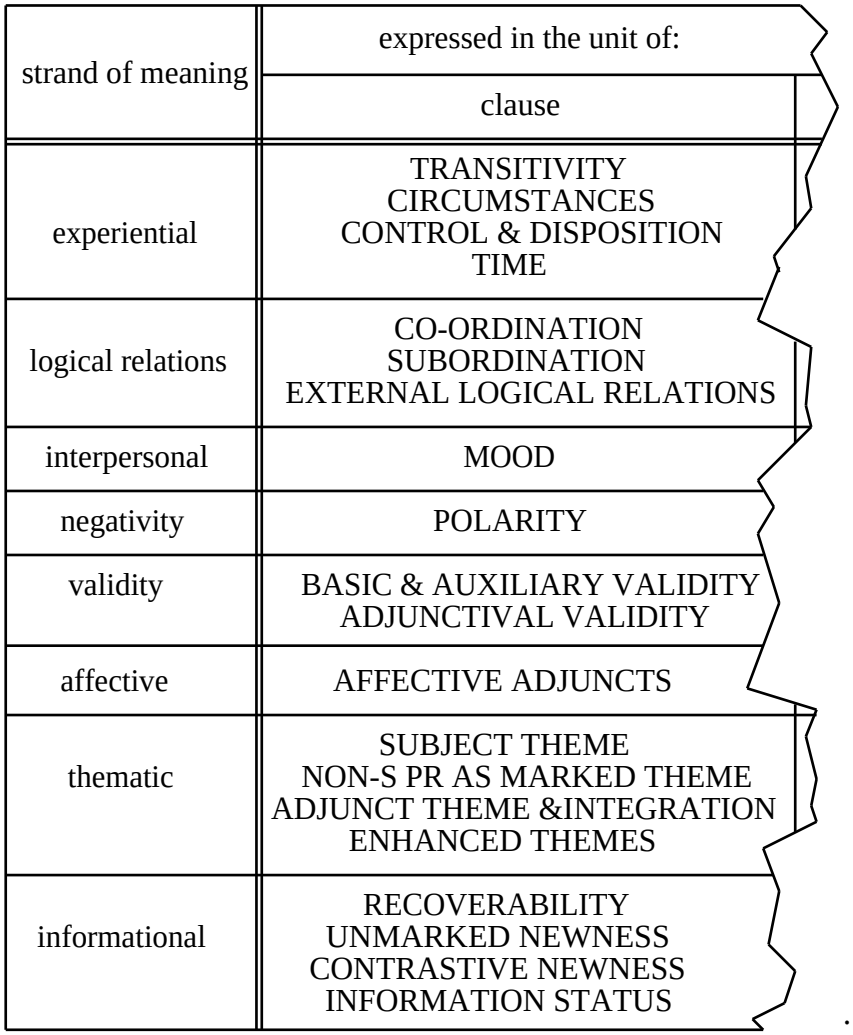

Figure 5:
The major strands of meaning and system networks for the English clause

If you want a single diagram that summarizes the contents of this chapter, it is that diagram. Firstly, it shows how the choices between **semantic features**, which are related to each other in a **system network**, generate a **selection expression**. Then it shows how the features of the selection expression become the input to the **realization component** - and this component then automatically generates the actual **syntax** and **items** that are the output of the lexicogrammar. Figure 2 therefore summarises both what the overall shape of a language is, and how a language works. And Section 5 provides a simple but fully explicit example of how the lexicogrammar works, for a small part of the English nominal group.

All this provides a **model** of the lexicogrammar of English. We considered various types of model, and concluded that a model of language must be abstract (rather than concrete) and that it must be a functioning model (rather than static or behaving). I rejected the traditional 'two books' model (which still seems to be the dominant model for most linguists), and I suggested that the concept of a **computer program** gives us a better model of what a language is like than any other.

Endnotes

These endnotes provide 'follow-up' comments and references for readers with prior knowledge and experience of linguistics. They include occasional brief comparisons with other approaches - including the 'traditional' approach to grammar and, within systemic functional grammar, comparisons with the proposals of Halliday and others.

a. By the end of this chapter the relationship between a 'language' and a 'sentence', as these terms are used in the present model, will be clear. Note that it will **NOT** be the relationship found in some types of generative-formal linguistics, where a language is defined as a set of sentences, and a sentence is thought of as a bracketed string of word-forms.

b. Perhaps the major advantage of having an observable picture (or model) is that this relieves your conscious mind of the work of trying to remember, as you think about it, what all the various bits of it are like and how they are related to each other. By placing some parts of the model inside other parts we can temporarily forget about the details of what is inside the superordinate parts. Computers, of course, have no such memory limitations, and this is why the judicious use of a combination of the strengths of the computer and the strengths of the human brains may free us - at least to some extent - from this limitation. Indeed, we have found this to be one of the major advantages of using a computer as the tool for modelling language in the COMMUNAL Project at the University of Wales, Cardiff. For further details, see Endnote f of Chapter 1.

c. Two obvious reasons for our extremely strong preference for two-dimensional pictures over three-dimensional models are (1) that two-dimensional pictures are cheaper and quicker to produce, and (2) that they have enormous advantages of economy in storage, whether on paper (in books, journals and files), or on film. However, new ways of storing images electronically are now beginning to make us re-evaluate this assumption of many centuries. Nonetheless we seem to be further than ever - for very good reasons - from making books obsolete. Hence the continuation in the present work of the tradition of two-dimensional models.

d. However, in the case of a tpy car the child's imagination is able supply the missing elements, and the failure of the model to function appropriately is central to its purpose as a toy. In other words, the role of a toy is less that of explaining things than stimulating the child's imagination. It is true that this can also be one of the roles of a model in science, but it remains the case that, if we are looking for a model to help us the understand the nature of language, it will help if we can find one that is a functioning model.

e. However, it must be admitted that, within the limitations of the 'two book' model, the Harris-Chomsky concept of a 'generative grammar' - e.g. the classic 'transformational generative grammar decribed in Chomsky's Aspects of the theory of syntax (1965) - is in a sense a 'functioning' grammar. Such grammars have been adapted for use in computer models for parsing the syntax of a string of words, or for the computer 'generation' of sentences. However, Chomsky himself goes out of his way to emphasize that the 'generative' capacity of his grammars is not in any sense a model of how people produce or understand sentences; he offers his grammars as an elegant way of capturing what a speaker of a language 'knows' about his or her language. It is thus essentially a

'static' model. See Fawcett (1993a) for a discussion of this issue for linguists, and Fawcett (1994a) for a similar discussion for computational linguists.

f. The work on implementing a reasonably full model of language as a computer program has led to many modifications to the standard systemic functional model, including (1) the widespread use of probabilities to replace, in many cases, absolute rules of 'grammaticality', and (2) the concept that the probabilities on featurres in systems - and so, if a feature is assigned 0%, the system itself - can be changed in the course of the generation of a sentence. See Fawcett, Tucker and Lin (1995) for a fairly full account of this work.

g. Ultimately, however, a model of language as an abstract functioning phenomenon must also be set within a model of a human being as a physical, behaving phenomenon - and moreover one in which the human being is a member of many social groups, from vary large to very small ones, permanent and temporary ones, etc.

h. I am indeed offering you here a 'two books' view of language - and the two books are*The Functional Syntax Handbook* and *The Functional Semantics Handbook.* But these two books divide up language in a very different way from the old division into 'grammar' and 'vocabulary'.

i. Clearly, formal grammarians who work in the 'autonomous syntax' approach would not accept this concept as wholeheartedly as most functional grammarians would - if they do at all.

j. There is a difference between the Sydney Grammar and the Cardiff Grammar on the question of the level or levels of 'meaning'. From one viewpoint this difference may appear to be relatively small, but in the long run it has major implications for the theory and description of languages. See Fawcett (1999) for a discussion of the differences between the two models of language. Butler (2002b:471), in his comparison of the two versions of SFG, writes that 'in my view the Cardiff model represents a considerable improvement on the Sydney account', his principal reason being that 'this model of the relationship between syntactic and semantic phenomena offers .... a much clearer programme of explanation of how semantics motivates syntax'.

k. Saussure in fact used the term 'associative' instead of 'paradigmatic' - or rather, the French equivalent.

l. All systemic functional linguists would agree with Halliday when he says that 'one of the things that distinguishes systemic functional grammar [from various other types of functional 'grammar'] is that it gives priority to paradigmatic relations' (Halliday 1994:15). This is because the system networks of meanings are the generative base of a systemic functional grammar. But I think that he goes a step too far when he continues that systemic functional grammar 'interprets language not as a set of structures but as a network of **SYSTEMS**, or interrelated sets of options for making meaning'. Here the word 'not' is misleading - certainly in terms of the Cardiff Grammar that is presented here but also, I believe, for Halliday's own version of SFG. This is because in SFG we also represent language as a structure - most obviously in instances at the level of form, but also, less obviously, at the level of meaning.

m. In the Cardiff Grammar, however, we break the relationship of 'constituency' down into alternating relations of **componence** and **filling**. Chapter 24 summarises the theory of syntax within which these concepts are to be understood.

n. In principle there are no sequential relationships at the level of semantics, sequence being the result of expressing part-whole relations at the level of form, However, it could be argued that the way in which in the relationship of **co-ordination** between units at the level of semantics is modelled in the grammar is 'sequential', because the first of any two co-ordinated units is generated before the grammar goes on to generate the next.

o. In the Cardiff Grammar, the technical name for the type of structural relationship that is present at both levels and in both the potential and the instances is the relationship of **componence** - i.e. constituency without sequence (for a fuller discussion see Chapter 24).

In the semantic system network this part-whole relationship is represented as if it was a choice (for reasons that we need not go into here), but in fact there is not a choice between, let us say, a 'situation' and a 'thing' - because the grammar always pre-selects the appropriate one with a 100%

probability.) However, the relationship is expressed overtly in the instances. First the grammar generates a selection expression for a 'situation', and then, for a role in its structure such as Agent, it generates another selection expression - such as the bundle of features that is realized in a nominal group (e.g. *those students* (which is used as an example later in this chapter). See Fawcett, Tucker and Lin (1993) for a fully specific account of how this is done in a systemic functional grammar, and Fawcett (1996) for an exemplification of the complex type involved in 'complementation'.

p. The **form potential** of a language contains a simple component containing the **potential structures** of those units whose sequence of elements is largely fixed - i.e. all units except the clause. Indeed it is perhaps the single most important fact about the syntax of English that there is no point in trying to use a potential structure for the clause. This is because the enormous number of potential variations in the sequence of its elements make this an extremely wasteful way to try to characterize the variations in sequence. This fact is illustrated and discussed in the final sections of each of Chapters 3 and 4, in which I present small systemic functional grannars that summarize the meanings and the syntax covered so far. Those sections illustrate the reason why those who try to capture syntactic relations in re-write rules find their task so hard (once they get beyond 'toy' grammars).

q. Halliday uses box diagrams, rather than tree diagrams. However, he does not explain why. Box diagrams are notationally equivalent to tree diagrams, in that they too model constituency relationships. Like most linguists other than those who work in Halliday's 'Sydney Grammar' version of SFG, I consider that tree diagrams provide a more flexible and more insightful representation of syntactic structure. Whenever it is necessary they can be enriched through the concept of **conflating** elements (as is done in the Cardiff Grammar), in order to enable them to give expression to the various strands of meaning that are realized in the clause.

A major problem with box diagrams is that they can only show clearly the structure of a single layer of structure. When you are trying to analyze a sentence containing several layers of structure, tree diagrams of the type used here are the clearest and most economical way to diagram the relationships. The 'depth' of units in a sentence is in fact often up to seven or more, and tree diagrams, as all publications in the Cardiff Grammar framework show, are able to handle all the syntagmatic complexity of language. Tree diagrams are also quicker and easier to draw, because they use fewer lines.

Another notation that is occasionally encountered is the mainland European tradition of 'sister dependency'. This was adopted for rhetorical structure relations at the level of discourse (Mann and Thompson 1987) - though so far as I am aware the reason for this decision has never been explained. And within the lexicogrammar itself, Halliday's concept of 'hypotaxis' (dependency without constituency) is essentially the same, and could be shown by that notation as well as in box diagram form. I have shown (in Fawcett and Davies 1992) that, with respect to the rhetorical structure relationships, the sister dependency relations can be re-expressed without loss of insight as part-whole relations. This is implemented in the computer model of discourse grammar in the COMMUNAL Project, and it makes possible the integration of the rhetorical structure with the existing implementation of a grammar of exchange structure - which is already implemented in a tree diagram form. Withe regard to 'hypotaxis' within sentences, I show at various points in the present book that all of the relations for which Halliday uses 'hypotaxis' can be handled naturally - and far more naturally, in my view - through other structural relations.

r. It is an important point of theory that the 'form potential' is seen here as being specified by the realization rules - and so ultimately by the meaning potential (together with the simple componence and sequencing rules of the 'potential structures'). Thus the 'form potential' is not expressed as a 'lower' level of network. This saves the description of the language from having to have another layer of networks. See Fawcett (1999) for a full discussion of the alternative positions within Systemic Fungtional Linguistics on levels and system networks in relation to 'meaning'.

s. Many linguists will at first have difficulty in accepting the unconventional viewpoint that language is inherently oriented to the production of texts rather than to understanding them. This may in part be because there is such a strong tradition in linguistics that a model of language itself should be neutral between the production and understanding of language - essentially on the self-evident grounds that a language has to be able to be used for both. But this deceptively attractive line of reasoning is mistaken, as I have argued in Fawcett (1994a). It is logically possible that

language might be inherently structured for either production or reception, and then that a procedure for drawing upon that model for the converse process might be developed. And this is what I believe to be the case. The meanings that are expressed in natural languages are the Performer's meanings (e.g. the deixis of time and place, thematic meanings, affective meanings, etc). It is in fact perfectly feasible to build a text-understanding system which includes a **syntax parser** and a **semantic interpreter**, each of these being constructed around the acceptance of the position that it relies, directly or indirectly, on understanding how the process of generation works. We know that it is possible to build such models because we have done it, in the framework of the COMMUNAL Project. For a full presentation of the argument for this position, see Fawcett (1994a), for descriptions of the computer implementations of these components, see Weerasinghe and Fawcett (1993) and (Weerasinghe (1994) for the parser and for the semantic interpreter see O'Donoghue (1991 & 1994). O'Donnell (1994) describes a partially similar approach.

The suggestion that a 'parser' has no equivalent in the process of generation may also surprise some readers. Again, I would ask you to read Fawcett (1994a). But for now consider this fact. In understanding, one major task is to turn a string of words into a string of words with a structured tree above it. What is the converse in generation? Clearly, the task of the equivalent component in generation is to remove the syntax tree from a representation consisting of a syntax tree whose terminal leaves are a string of words. Such a **syntax stripper** is a trivially simple component to model, when compared with the difficulty of building a successful parser.

t. Again, see Fawcett, Tucker and Lin (1993) for an overview of how the various components fit together. Fawcett (1990) gives a fairly full picture of how the intonation fits into the overall model.

u. In Fawcett, Tucker and Lin (1993) we do indeed take the generative approach, and the generative apparatus starts with the the system networks of meanings. However, the version of the system networks that can be used for describing texts - and that therefore correspond to the syntactic structures introduced in this handbook - are not to be found there but in the sister volume to this handbook. the *Functional Semantics Handbook* (Fawcett in press).

v. The multifunctional principle was first made explicit in Halliday (1968), and it was developed in Halliday (1970b) and many subsequent works, including Halliday (1985) and (1994). For an examination of the grounds for dividing the meaning potential of a language into 'functional components' (or 'metafunctions', to introduce their more technical name), and for my proposals for further distinctions among them, see Fawcett (1980:34f.) For discussions of the relationship between Halliday's and my alternative proposals, see Ellis (1987) and Gregory (1987). For a general critical overview of the concept of 'metafunction', see Butler (1985). Halliday himself has sometimes presented a picture of three and sometimes four main 'metafunctions', and he has changed his mind about the location of one system (POLARITY) at least twice (Fawcett 1980:30) - as indeed I have too. So the question of which system networks of meaning belong in which 'metafunctions' is not one with simple answers. Moreover, in his more recent work (e.g. Halliday (1985) and (1994), in contrast with Halliday (1973)) he has treated logical relations as 'above' the clause, so the picture of what rank of unit each 'metafunction' is relevant to has also changed over the years.

The task of assigning 'types of meaning' - i.e. **system networks** covering specific areas of meaning - to their appropriate 'metafunctions' seems to be a matter of considerable importance for some systemic functional linguists. I must admit that I helped to trigger that debate by my proposals in Fawcett (1973/81 and 1980) that Halliday's FOUR (as there were then) 'functional components' (or 'metafunctions') for the English clause should be expanded to the EIGHT main 'functional components' mentioned in the main text (plus three 'minor' and 'accessory' functions). I still consider that it is helpful to think in terms of about this number of different functional components when one is analyzing a text - or 'strands of meaning' as I prefer to call them - but I have now changed my view in one respect. I now consider that it is **NOT** a matter of great importance that we should have criteria that enable us to assign a particular system network to one or other of Halliday's three (or four) 'metafunctions' - or indeed to one of my eight. In other words, I do not think it is a major theoretical issue. I discuss the value of the concepts of 'metafunction' and 'strand of meaning' in Section 14 of Chapter 4.