COGNITIVE LINGUISTICS AN INTRODUCTION

Vyvyan Evans and Melanie Green

EDINBURGH UNIVERSITY PRESS

© Vyvyan Evans and Melanie Green, 2006

Edinburgh University Press Ltd 22 George Square, Edinburgh

Typeset in Sabon and Gill Sans by Servis Filmsetting Ltd, Manchester, and printed and bound in Great Britain by Antony Rowe Ltd, Chippenham, Wilts

A CIP record for this book is available from the British Library

ISBN 0 7486 1831 7 (hardback) ISBN 0 7486 1832 5 (paperback)

The right of Vyvyan Evans and Melanie Green to be identified as authors of this work has been asserted in accordance with the Copyright, Designs and Patents Act 1988.

What does it mean to know a language?

Cognitive linguists, like other linguists, study language for its own sake; they attempt to describe and account for its systematicity, its structure, the functions it serves and how these functions are realised by the language system. However, an important reason behind why cognitive linguists study language stems from the assumption that language reflects patterns of thought. Therefore, to study language from this perspective is to study patterns of conceptualisation. Language offers a window into cognitive function, providing insights into the nature, structure and organisation of thoughts and ideas. The most important way in which cognitive linguistics differs from other approaches to the study of language, then, is that language is assumed to reflect certain fundamental properties and design features of the human mind. As we will see throughout this book, this assumption has far-reaching implications for the scope, methodology and models developed within the cognitive linguistic enterprise. Not least, an important criterion for judging a model of language is whether the model is psychologically plausible.

Cognitive linguistics is a relatively new school of linguistics, and one of the most innovative and exciting approaches to the study of language and thought that has emerged within the modern field of interdisciplinary study known as cognitive science. In this chapter we will begin to get a feel for the issues and concerns of practising cognitive linguists. We will do so by attempting to answer the following question: what does it mean to know a language? The way we approach the question and the answer we come up with will reveal a lot about the approach, perspective and assumptions of cognitive linguists. Moreover, the view of language that we will finish with is quite different from the view suggested by other linguistic frameworks. As we will see throughout this book, particularly in the comparative chapters at the ends of Part II and Part III, the

answer to the title of this chapter will provide a significant challenge to some of these approaches. The cognitive approach also offers exciting glimpses into hitherto hidden aspects of the human mind, human experience and, consequently, what it is to be human.

I.I What is language for?

We take language for granted, yet we rely upon it throughout our lives in order to perform a range of functions. Imagine how you would accomplish all the things you might do, even in a single day, without language: buying an item in a shop, providing or requesting information, passing the time of day, expressing an opinion, declaring undying love, agreeing or disagreeing, signalling displeasure or happiness, arguing, insulting someone, and so on. Imagine how other forms of behaviour would be accomplished in the absence of language: rituals like marriage, business meetings, using the Internet, the telephone, and so forth. While we could conceivably accomplish some of these things without language (a marriage ceremony, perhaps?), it is less clear how, in the absence of telepathy, making a telephone call or sending an e-mail could be achieved.

In almost all the situations in which we find ourselves, language allows quick and effective expression, and provides a well developed means of **encoding** and **transmitting** complex and subtle ideas. In fact, these notions of encoding and transmitting turn out to be important, as they relate to two key functions associated with language, the **symbolic function** and the **interactive function**.

1.1.1 The symbolic function of language

One crucial function of language is to express thoughts and ideas. That is, language encodes and externalises our thoughts. The way language does this is by using **symbols**. Symbols are 'bits of language'. These might be meaningful subparts of words (for example, *dis-* as in *distaste*), whole words (for example, *cat*, *run*, *tomorrow*), or 'strings' of words (for example, *He couldn't write a pop jingle let alone a whole musical*). These symbols consist of **forms**, which may be spoken, written or signed, and meanings with which the forms are conventionally paired. In fact, a symbol is better referred to as a **symbolic assembly**, as it consists of two parts that are conventionally associated (Langacker 1987). In other words, this symbolic assembly is a **form-meaning pairing**.

A form can be a sound, as in [kæt]. (Here, the speech sounds are represented by symbols from the International Phonetic Alphabet.) A form might be the orthographic representation that we see on the written page: *cat*, or a signed gesture in a sign language. A **meaning** is the conventional ideational or semantic content associated with the symbol. A symbolic assembly of form and meaning is represented in Figure 1.1.



Figure 1.1 A symbolic assembly of form and meaning



Figure 1.2 Levels of representation

It is important to make it clear that the image of the cat in Figure 1.1 is intended to represent not a particular referent in the world, but the idea of a cat. That is, the image represents the meaning conventionally paired with the form pronounced in English as [kæt]. The meaning associated with a linguistic symbol is linked to a particular mental representation termed a **concept**. Concepts, in turn, derive from **percepts**. For instance, consider a piece of fruit like a pear. Different parts of the brain perceive its shape, colour, texture, taste, smell and so on. This diverse range of perceptual information deriving from the world 'out there' is integrated into a single **mental image** (a representation available to consciousness), which gives rise to the concept of PEAR. When we use language and utter the form *pear*, this symbol corresponds to a conventional meaning, and therefore 'connects' to a concept rather than directly to a physical object in the external world (see Figure 1.2).

Our cognitive abilities integrate raw perceptual information into a coherent and well defined mental image. The meanings encoded by linguistic symbols then, refer to our **projected reality** (Jackendoff 1983): a mental representation of reality, as construed by the human mind, mediated by our unique perceptual and conceptual systems.

We stated above that the symbolic function of language serves to encode and externalise our thoughts. We are now in a position to qualify this view. While our **conceptualisations** are seemingly unlimited in scope, language represents a limited and indeed limiting system for the expression of thought; we've all experienced the frustration of being unable to 'put an idea into words'. There is, after all, a finite number of words, with a delimited set of conventional meanings. From this perspective then, language merely provides **prompts** for the construction of a conceptualisation which is far richer and more elaborate than the minimal meanings provided by language (Fauconnier 1997; Turner 1991). Accordingly, what language encodes is not thought in its complex entirety, but instead rudimentary instructions to the conceptual system to access or create rich and elaborate ideas. To illustrate this point, consider the following illustration adapted from Tyler and Evans (2003):

(1) The cat jumped over the wall.

This sentence describes a jump undertaken by a cat. Before reading on, select the diagram in Figure 1.3 that best captures, in your view, the trajectory of the jump.

We anticipate that you selected the fourth diagram, Figure 1.3(d). After all, the conventional interpretation of the sentence is that the cat begins the jump on one side of the wall, moves through an arc-like trajectory, and lands on the other side of the wall. Figure 1.3(d) best captures this interpretation. On first inspection, this exercise seems straightforward. However, even a simple sentence like (1) raises a number of puzzling issues. After all, how do we know that the trajectory of the cat's jump is of the kind represented in Figure 1.3(d)? What information is there in the sentence that provides this interpretation and excludes the trajectories represented in Figures 1.3(a–c)?

Even though the sentence in (1) would typically be judged as unambiguous, it contains a number of words that have a range of interpretations. The behaviour described by *jump* has the potential to involve a variety of trajectory shapes. For instance, jumping from the ground to the table involves the trajectory represented in Figure 1.3(a). Jumping on a trampoline relates to the trajectory represented in 1.3(b). Bungee jumping involves the trajectory represented in 1.3(c), in which the bungee jumper stops just prior to contact with the surface. Finally, jumping over a puddle, hurdle, wall and so on involves an arc-like trajectory as in 1.3(d).



Figure 1.3 Possible trajectories for The cat jumped over the wall

If the lexical item *jump* does not in itself specify an arc-like trajectory, but is vague with respect to the shape of the trajectory, then perhaps the preposition *over* is responsible. However, *over* can also have a range of possible interpretations. For instance, it might mean 'across', when we walk *over* a bridge (a horizontal trajectory). It might mean 'above', when an entity like a hummingbird is *over* a flower (higher than but in close proximity to). Equally, *over* could mean 'above' when a plane flies *over* a city (much higher and lacking close proximity). These are just a few of the possibilities. The point to emerge from this brief discussion is that *over* can be used when different kinds or amounts of space are involved, and with a number of different trajectories or paths of motion.

Consider a further complication. Figure 1.3(d) crucially represents the cat's motion ending at a point on the opposite side of the wall relative to the starting position of the jump. Yet no linguistic element in the sentence explicitly provides us with this information.

Example (1) therefore illustrates the following point: even in a mundane sentence, the words themselves, while providing meanings, are only partially responsible for the conceptualisation that these meanings give rise to. Thought relies on a rich array of encyclopaedic knowledge (Langacker 1987). For example, when constructing an interpretation based on the sentence in (1), this involves at the very least the following knowledge: (1) that the kind of jumping cats perform involves traversing obstacles rather than bungee jumping; (2) that if a cat begins a jump at a point on one side of an obstacle, and passes through a point above that obstacle, then gravity will ensure that the cat comes to rest on the other side of the obstacle; (3) that walls are impenetrable barriers to forward motion; (4) that cats know this, and therefore attempt to circumnavigate the obstacle by going over it. We use all this information (and much more), in constructing the rich conceptualisation associated with the sentence in (1). The words themselves are merely prompts for the construction process.

So far, then, we have established that one of the functions of language is to represent or symbolise concepts. Linguistic symbols, or more precisely symbolic assemblies, enable this by serving as prompts for the construction of much richer conceptualisations. Now let's turn to the second function of language.

1.1.2 The interactive function of language

In our everyday social encounters, language serves an interactive function. It is not sufficient that language merely pairs forms and meanings. These formmeaning pairings must be recognised by, and be accessible to, others in our community. After all, we use language in order to 'get our ideas across', in other words to **communicate**. This involves a process of transmission by the speaker, and decoding and interpretation by the hearer, processes that involve the construction of rich conceptualisations (see Figure 1.4).



Figure 1.4 The interactive function

The messages we choose to communicate can perform various interactive and **social functions**. For example, we can use language to change the way the world is, or to make things happen:

- (2) a. I now pronounce you man and wife.
 - b. Shut the door on your way out!

The utterance in (2a), spoken by a suitably qualified person (such as a member of the clergy licensed to perform marriages), in an appropriate setting (like a church), in the presence of two unmarried adults who consent to be joined in matrimony, has the effect of irrevocably altering the social, legal and even spiritual relationship between the two people. That is, language itself can serve as a **speech act** that forever alters an aspect of our reality.

Similarly, in the example in (2b), the utterance represents a command, which is also a type of speech act. Language provides a means of communication, allowing us to share our wishes and desires. Moreover, the way in which these wishes and desires are expressed signals who we are, and what kind of relationship we have with our addressee. We would be unlikely to issue a command like (2b) to the Queen of England, for example.

Another way in which language fulfils the interactive function relates to the notion of **expressivity**. Language is 'loaded', allowing us to express our thoughts and feelings about the world; consider the different mental images evoked by the following expressions, which might be used by different speakers to refer to the same individual:

- (3) a. the eminent linguist
 - b. the blonde bombshell

While the example in (3a) focuses on the profession of the individual and her relative standing in that profession, the example in (3b) focuses on her physical appearance. Moreover, although both these sentences relate to a female linguist, the person's gender cannot be inferred from the sentence in (3a) while it can from the second sentence due to normative patterns of linguistic behaviour and social stereoptypes. That is, we typically use the expression *blonde bomb-shell* to describe the physical attributes of women rather than men.

Language also plays a role in how we affect other people in the world, and how we make others feel by our choice of words. That is, language can provide information about **affect** (emotional response):

- (4) a. Shut up!
 - b. I'm terribly sorry to interrupt you, but . . .

These examples also illustrate the way in which we present our public selves through language. The language we choose to use conveys information about our attitudes concerning others, ourselves and the situations in which we find ourselves.

Language can be used to create scenes or **frames** of experience, indexing and even constructing a particular context (Fillmore 1982). In other words, language use can invoke frames that summon rich knowledge structures, which serve to call up and fill in background knowledge.

- (5) a. How do you do?
 - b. Once upon a time . . .

The example in (5a) creates a greeting frame, signalling an acknowledgement of another person and a recognition that this is the first time they have met. It also signals a degree of formality, which expressions like *hey*, *what's up?* or *hi* would not. Analogously, the utterance in (5b) signals the beginning of a fairy-tale. In other words, just by hearing or reading the expression in (5b) an entire frame is invoked, which guides how we should respond to what follows, what our expectations should be and so forth.

In summary, we've seen that not only does language encode particular meanings, but also that, by virtue of these meanings and the forms employed to symbolise these meanings which constitute part of shared knowledge in a particular speech community, language can serve an interactive function, facilitating and enriching communication in a number of ways.

1.2 The systematic structure of language

Having seen some examples of what language is used for, let's now consider how language is structured. Language is a system for the expression of meaning and

for carrying out its symbolic and interactive functions. So, what evidence is there for the systematicity of language?

I.2.1 Evidence for a system

Language consists of symbolic assemblies that are combined in various ways to perform the functions we described in section 1.1. A symbolic assembly is a conventional **linguistic unit**, which means that it is a piece of language that speakers recognise and 'agree' about in terms of what it means and how it is used. As we will see later in the book, particularly in Part III, one of the prominent concerns in cognitive approaches to grammar is how to model the inventory of linguistic units that make up a language. For example, speakers of Modern English 'agree' that the form *cat* is used to refer to a certain kind of meaning which we illustrated in Figure 1.2. A conventional unit can be a meaningful subpart of a word, which linguists call a **morpheme** (*anti-dis-establish* . . .), a whole word, a string of words that 'belong' together (a **phrase**) or a whole sentence. Now let's consider another example:

(6) He kicked the bucket

This utterance consists of a sentence that has an **idiomatic meaning** in English. That is, its meaning is not predictable from the integrated meanings of the individual words. A non-native speaker of English who has not learnt the 'special' idiomatic meaning will only be able to interpret example (6) literally. Native speakers of English, on the other hand, while also being able to interpret the sentence literally, often cannot avoid the idiomatic meaning 'he died'. Of course, whether a literal versus an idiomatic interpretation is accessed depends on the situation or **context** in which the utterance occurs.

Focusing for now on the idiomatic interpretation, we can view this utterance as a unit that has a particular meaning associated with it. Therefore, it counts as a symbolic assembly. Another term for symbolic assembly that is employed by some cognitive linguists is **construction** (e.g. Goldberg 1995). We will look in detail at the notion of symbolic assemblies and constructions in Part III of the book.

When we change certain aspects of the sentence in (6), the meaning is affected. For example, if we change the object (the thing being kicked), as in (7), we lose the idiomatic meaning and are left with a **literal** utterance:

(7) He kicked the mop.

For many cognitive linguists, what makes example (7) 'literal' is that this sentence 'as a whole' does not represent a construction. Instead, the meaning of (7)

is interpreted by **unifying** the smaller units, the words. In contrast, example (6) is interpreted as a whole single unit: a construction. One way of expressing this idea in more intuitive terms is to use the metaphor of 'storage': suppose we store our knowledge of words, phrases and complex constructions in a mental 'box'. The behaviour of larger constructions, like *kick the bucket*, suggests that these are stored as 'chunks' or single units, just like words. The meanings of sentences like (7) on the other hand are 'built' by unifying the individual words that make them up.

Now consider another example. If we change the structure of example (6) in the following way, we also lose the idiomatic meaning:

(8) The bucket was kicked by him.

This example shows that, in addition to meaning, constructions (formmeaning pairings) have particular formal grammatical patterns associated with them. In other words, the properties of the construction relate not only to the individual words that make it up, as in (6), but also to the grammatical form, or word order. The passive construction in (8), in which *the bucket* is placed in subject position, fails to provide the idiomatic meaning associated with the sentence in (6). We can conclude from this that the linear arrangement of the words in the sentence constitutes part of an individual's knowledge of idiomatic constructions like (6).

This point is also illustrated by an **ungrammatical** sentence, a sentence that does not correspond to any of the formal patterns associated with the constructions of English, as in (9), and consequently does not have a conventional meaning associated with it. Ungrammaticality is indicated by an asterisk:

(9) *Bucket kicked he the

As we noted above, the sentence in (6) qualifies as a construction because it consists of particular words arranged in a particular order, and these words are conventionally associated with a particular (idiomatic) meaning. However, we have suggested that constructions can also give rise to 'literal' meanings. To illustrate this, we will examine another sentence that has both idiomatic and literal meanings. For instance, consider the following linguistic joke:

- (10) A: Waiter, what is this fly doing in my soup?
 - B: I think that's the breaststroke, sir!

This joke turns on the ambiguity between the regular interrogative construction, in which a speaker is enquiring after the intention or purpose of something or someone (*What's that seagull doing on the roof? What's that woman* doing over there?), and the 'What's X doing Y construction', studied in detail by cognitive linguists Paul Kay and Charles Fillmore (1999), in which the speaker is indicating that a particular situation is incongruous or unacceptable (*What are you doing mearing those bunny ears? What are those clothes doing on the floor?*). Notice that each of these interpretations requires a different kind of response. For the regular interrogative construction, the response should consist minimally of a piece of information corresponding to the question word (*building a nest; maiting for a bus*). For the 'what's X doing Y' construction, on the other hand, the expected response is typically an explanation, excuse or apology (*I'm going to a fancy-dress party; I've been busy*).

Crucially, for example (10), these two very different meanings are conventionally associated with exactly the same words arranged in the same sequence. The humorous effect of the waiter's reply rests on the fact that he has chosen to respond to the 'wrong' interpretation. While the diner is employing the 'what's X doing Y' construction, the waiter prefers to respond to the interrogative construction.

The examples in this section illustrate the fact that there is a systematic relationship between words, their meanings and how they are arranged in conventional patterns. In other words, language has a systematic structure.

1.2.2 The systematic structure of thought

Does the systematic structure found in language reflect a systematic structure within our conceptual system? Cognitive linguists certainly think so. Cognitive linguists explore the hypothesis that certain kinds of linguistic expressions provide evidence that the structure of our conceptual systems is reflected in the patterns of language. Moreover, as we will see throughout this book, the way the mind is structured can be seen as a reflection, in part, of the way the world (including our sociocultural experience) is structured and organised. Consider the examples in (11).

- (11) a. Christmas is fast approaching.
 - b. The number of shares we own has gone up.
 - c. Those two have a very close friendship.

These examples relate to the abstract **conceptual domains** of TIME (11a), QUANTITY (11b) and AFFECTION (11c). A conceptual domain is a body of knowledge within our conceptual system that contains and organises related ideas and experiences. For example, the conceptual domain of TIME might relate a range of temporal concepts including *Christmas*, which is a temporal event. Notice that in each sentence in (11) the more abstract concepts *Christmas*, *number (of shares)* and *friendship* are understood in terms of conceptual domains relating to *concrete* physical experience. For instance, Christmas is conceptualised in terms of the domain of physical MOTION, which is evident in the use of the word *approaching* in (11a). Clearly *Christmas* (and other temporal concepts) cannot literally be said to undergo motion. Similarly, the notion of *number of shares* is conceptualised in terms of VERTICAL ELEVATION, which is clear from the use of the phrase *gone up* in (11b). Finally, *friendship* is conceptualised in terms of PHYSICAL PROXIMITY in (11c), which is shown by the use of the word *close*.

One of the major findings to have emerged from studies into the human conceptual system is that abstract concepts are systematically structured in terms of conceptual domains deriving from our experience of the behaviour of physical objects, involving properties like motion, vertical elevation and physical proximity (Lakoff and Johnson 1980, 1999). It seems that the language we use to talk about temporal ideas such as *Christmas* provides powerful evidence that our conceptual system 'organises' abstract concepts in terms of more concrete kinds of experiences, which helps to make the abstract concepts more readily accessible.

1.3 What do linguists do?

As we have begun to see, cognitive linguists form hypotheses about the nature of language, and about the conceptual system that it is thought to reflect. These hypotheses are based on observing patterns in the way language is structured and organised. It follows that a theory of language and mind based on linguistic observation must first describe the linguistic facts in a systematic and rigorous manner, and in such a way that the description provides a plausible basis for a speaker's tacit knowledge of language. This foundation for theorising is termed **descriptive adequacy** (Chomsky 1965; Langacker 1987, 1999a). This concern is one that cognitive linguists share with linguists working in other traditions. Below, we provide an outline of what it is that linguists do and how they go about it.

I.3.1 What?

Linguists try to uncover the systems behind language, to describe these systems and to **model** them. Linguistic models consist of theories about language. Linguists can approach the study of language from various perspectives. Linguists may choose to concentrate on exploring the systems within and between sound, meaning and grammar, or to focus on more applied areas, such as the evolution of language, the acquisition of language by children, language disorders, the questions of how and why language changes over time, or the relationship between language, culture and society. For cognitive linguists, the emphasis is upon relating the systematicity exhibited by language directly to the way the mind is patterned and structured, and in particular to conceptual structure and organisation. It follows that there is a close relationship between cognitive linguistics and aspects of cognitive psychology. In addition to this, applied linguistics also informs and is informed by the cognitive linguistics research agenda in various ways (see Chapters 3 and 4 for further discussion of this point).

1.3.2 Why?

Linguists are motivated to explore the issues we outlined above by the drive to understand human cognition, or how the human mind works. Language is a uniquely human capacity. Linguistics is therefore one of the **cognitive sciences**, alongside philosophy, psychology, neuroscience and artificial intelligence. Each of these disciplines seeks to explain different (and frequently overlapping) aspects of human cognition. In particular, as we have begun to see, cognitive linguists view language as a system that directly reflects conceptual organisation.

1.3.3 How?

As linguists, we rely upon what language tells us about itself. In other words, it is ordinary language, spoken every day by ordinary people, that makes up the 'raw data' that linguists use to build their theories. Linguists describe language, and on the basis of its properties, formulate hypotheses about how language is represented in the mind. These hypotheses can be tested in a number of ways.

1.3.4 Speaker intuitions

Native speakers of any given human language will have strong **intuitions** about what combinations of sounds or words are possible in their language, and which interpretations can be paired with which combinations. For example, native speakers of English will agree that example (6), repeated here, is a well-formed sentence, and that it may have two possible meanings:

(6) He kicked the bucket.

They will also agree that (7) and (8), repeated here, are both well-formed sentences, but that each has only one possible meaning:

- (7) He kicked the mop.
- (8) The bucket was kicked by him.

Finally, and perhaps most strikingly, speakers will agree that all of the following examples are impossible in English:

- (12) a. *bucket kicked he the
 - b. *kicked bucket the he
 - c. *bucket the kicked he
 - d. *kicked he bucket the

Facts like these show that language, and speakers' intuitions about language, can be seen as a 'window' to the underlying system. On the basis of the patterns that emerge from the description of language, linguists can begin to build theoretical 'models' of language. A model of language is a set of statements that is designed to capture everything we know about this hidden cognitive system in a way that is principled, based on empirical evidence and psychologically plausible.

1.3.5 Converging evidence

How do cognitive linguists evaluate the adequacy of their models? One way is to consider **converging evidence** (Langacker 1999a). This means that a model must not only explain linguistic knowledge, but must also be consistent with what cognitive scientists know about other areas of cognition, reflecting the view that linguistic structure and organisation are a relatively imprecise but nevertheless indicative reflection of cognitive structure and organisation. By way of illustration, consider the scene in Figure 1.5.

How might we use language to describe a scene like this? Most English speakers will agree that (13a) is an appropriate description but that (13b) is 'odd':

- (13) a. The cat is on the chair.
 - b. ?The chair is under the cat.



Figure 1.5 The cat is on the chair

Why should (13b) be 'odd'? It's a perfectly grammatical English sentence. From what psychology has revealed about how the human mind works, we know that we have a tendency to focus our attention on certain aspects of a visual scene. The aspect we focus on is something about which we can make certain predictions. For example, in Figure 1.5 we focus on the cat rather than the chair, because our knowledge of the world tells us that the cat is more likely than the chair to move, to make a noise or to perform some other act. We call this prominent entity the figure and the remainder of the scene the ground, which is another way of saying 'background' (see Chapter 3). Notice that this fact about human psychology provides us with an explanation for why language 'packages' information in certain ways. In (13a) the cat has a prominent position in the sentence; any theory of language will tell you that sentence initial position is a 'special' position in many of the world's languages. This accords with the prominence of the corresponding entity in the visual scene. This explanation, based on the figure-ground distinction, also provides us with an explanation for why (13b) is 'odd'. This is an example of how converging evidence works to strengthen or confirm theories of language. Can you think of a situation in which (13b) would not be odd?

1.4 What it means to know a language

Let's look more closely now at some of the claims made by cognitive linguists about how language is represented in the mind. We have established that the linguist's task is to uncover the systematicity behind and within language. What kinds of systems might there be within language? We'll begin to answer this question by introducing one fundamental distinction based on the foundational work of pioneering cognitive linguist Leonard Talmy. Talmy suggests that the **cognitive representation** provided by language can be divided into **lexical** and **grammatical** subsystems. Consider the following example:

(14) The hunter tracked the tigers.

Notice that certain parts of the sentence in (14) – either whole words (free morphemes), or meaningful subparts of words (bound morphemes) – have been marked in boldtype. What happens when we alter those parts of the sentence?

- (15) a. Which hunter tracked the tigers?
 - b. The hunter tracks the tigers.
 - c. Those hunters track a tiger.

All the sentences in (15) are still about some kind of tracking event involving one or more hunter(s) and one or more tiger(s). What happens when we change

the 'little' words like *a*, *the* and *those* and the bound morphemes like *-ed* or *-s* is that we then interpret the event in different ways, relating to information about number (how many hunters or tigers are/were there?), tense (did this event happen before now or is it happening now?), old/new information (does the hearer know which hunters or tigers we're talking about?) and whether the sentence should be interpreted as a statement or a question.

These linguistic elements and morphemes are known as **closed-class** elements and relate to the grammatical subsystem. The term *closed-class* refers to the fact that it is typically more difficult for a language to add new members to this set of elements. This contrasts with the non-boldtype 'lexical' words which are referred to as **open-class**. These relate to the lexical subsystem. The term *open-class* refers to the fact that languages typically find it much easier to add new elements to this subsystem and do so on a regular basis.

In terms of the meaning contributed by each of these two subsystems, while 'lexical' words provide 'rich' meaning and thus have a **content func-tion**, 'grammatical' elements perform a **structuring function** in the sentence. They contribute to the interpretation in important but rather more subtle ways, providing a kind of 'scaffolding' which supports and structures the rich content provided by open-class elements. In other words, the elements associated with the grammatical subsystem are constructions that contribute **schematic meaning** rather than rich contentful meaning. This becomes clearer when we alter the other parts of the sentence. Compare (14) with (16):

- (16) a. The movie star kissed the directors.
 - b. The sunbeam illuminated the rooftops.
 - c. The textbook delighted the students.

What all the sentences in (16) have in common with (14) is the 'grammatical' elements. In other words, the grammatical structure of all the sentences in (16) is identical to that of (15). We know that both participants in the event can easily be identified by the hearer. We know that the event took place before now. We know that there's only one movie star/sunbeam/textbook, but more than one director/rooftop/student. Notice that the sentences differ in rather a dramatic way, though. They no longer describe the same kind of event at all. This is because the 'lexical' elements prompt for certain kinds of concepts that are richer and less schematic in nature than those prompted for by 'grammatical' elements. The lexical subsystem relates to things, people, places, events, properties of things and so on. The grammatical subsystem on the other hand relates to concepts having to do with number, time reference, whether a piece of information is old or new, whether the speaker is providing information or requesting information, and so on.

COGNITIVE LINGUISTICS: AN INTRODUCTION

Lexical subsystem	Grammatical subsystem
Open-class words/morphemes	Closed-class words/morphemes
Content function	Structuring function
Larger set; constantly changing	Smaller set; more resistant to change
Prompts for 'rich' concepts, e.g. people, things, places, properties, etc.	Prompts for schematic concepts, e.g. number, time reference, old vs. new, statement vs. question, etc.

Table 1.1 Properties of the lexical and grammatical subsystems

A further important distinction between these two subsystems concerns the way that language changes over time. The elements that comprise the lexical (open-class) subsystem make up a large and constantly changing set in any given human language; over a period of time, words that are no longer 'needed' disappear and new ones appear. The 'grammatical' (closed-class) elements that make up the grammatical subsystem, on the other hand, constitute a smaller set, relatively speaking, and are much more stable. Consequently, they tend to be more resistant to change. However, even 'grammatical' elements do change over time. This is a subject we'll come back to in more detail later in the book when we discuss the process known as grammaticalisation (see Chapter 21).

Table 1.1 provides a summary of these important differences between the lexical and grammatical subsystems. Together, these two subsystems allow language to present a cognitive representation, encoding and externalising thoughts and ideas.

Having provided a sketch of what it means to know a language from the perspective of cognitive linguistics, we will now begin to examine the cognitive linguistics enterprise in more detail. In particular, we must consider the assumptions and commitments that underlie the cognitive linguistics enterprise, and begin to examine this approach to language in terms of its perspective, assumptions, the cognitive and linguistic phenomena it considers, its methodologies and its approach to theory construction. We turn to these issues in the next chapter.

1.5 Summary

We began this chapter by stating that cognitive linguists, like other linguists, attempt to describe and account for linguistic **systematicity**, **structure** and **function**. However, for cognitive linguists, language reflects patterns of thought; therefore, to study language is to study patterns of **conceptualisa-tion**. In order to explore these ideas in more detail we looked first at the functions of language. Language provides a means of **encoding** and **transmitting**

ideas: it has a symbolic function and an interactive function. Language encodes and externalises our thoughts by using symbols. Linguistic symbols consist of form-meaning pairings termed symbolic assemblies. The meaning associated with a linguistic symbol relates to a mental representation termed a concept. Concepts derive from percepts; the range of perceptual information deriving from the world is integrated into a mental image. The meanings encoded by linguistic symbols refer to our projected reality: a mental representation of reality as construed by the human mind. While our conceptualisations are unlimited in scope, language merely provides prompts for the construction of conceptualisations. Language also serves an interactive function; we use it to communicate. Language allows us to perform speech acts, or to exhibit expressivity and affect. Language can also be used to create scenes or contexts; hence, language has the ability to invoke experiential frames.

Secondly, we examined the evidence for a linguistic system, introducing the notion of a conventional **linguistic unit**, which may be a **morpheme**, a **word**, a string of words or a sentence. We introduced the notion of **idiomatic meaning** which is available in certain **contexts** and which can be associated with **constructions**. This contrasts with **literal** meaning, which may be derived by **unifying** smaller constructions like individual words. Word **order** constitutes part of an individual's knowledge of particular constructions, a point illustrated by **ungrammatical** sentences. We also related linguistic structure to the systematic structure of thought. **Conceptual domains** reflected in language contain and organise related ideas and experiences.

Next, we outlined the task of the cognitive linguist: to form hypotheses about the nature of language and about the conceptual system that it reflects. These hypotheses must achieve **descriptive adequacy** by describing linguistic facts in a systematic and rigorous manner. Linguists try to uncover, describe and **model** linguistic systems, motivated by the drive to understand human cognition. Linguistics is therefore one of the **cognitive sciences**. Cognitive linguists carry out this task by examining linguistic **data** and by relying on native speaker **intuitions** and **converging evidence**. As an example of converging evidence, we explored the linguistic reflex of the distinction made in psychology between **figure** and **ground**.

Finally, we looked at what it means to know a language, and introduced an important distinction between kinds of linguistic knowledge: the **cognitive** representation provided by language can be divided into lexical and grammatical subsystems. The lexical subsystem contains open-class elements which perform a content function. The grammatical subsystem contains closed-class elements, which perform a structuring function providing schematic meaning.