

Oxford
LINGUISTICS

Meaning in Language

An Introduction to
Semantics and Pragmatics

THIRD
EDITION

Alan Cruse

OXFORD TEXTBOOKS IN LINGUISTICS

Meaning in Language

An Introduction to
Semantics and
Pragmatics

Third Edition

Alan Cruse

OXFORD
UNIVERSITY PRESS

Contents

<i>Typographic conventions</i>	ix
<i>From the preface to the first edition</i>	x
<i>Preface to the second edition</i>	xi
<i>Preface to the third edition</i>	xii
Part 1 Fundamental Notions	1
1 Introduction	3
2 Logic and meaning	21
3 Concepts and meaning	51
Part 2 Words and Meanings	71
4 Lexical units	73
5 Contextual variability of word meaning	97
6 Paradigmatic relations of inclusion and identity	127
7 Paradigmatic relations of exclusion and opposition	149
8 Lexical hierarchies	165
9 Syntagmatic semantic relations	177
10 Describing lexical senses 1: dimensions and structures	193
11 Describing lexical senses 2: approaches to the specification of word meaning	211
12 Extensions of meaning	237

Part 3 Grammatical Meaning	263
13 Grammatical meaning; nouns and noun phrases	265
14 Argument structure	279
15 Verbs and adjectives	295
16 Prepositions	319
17 Derivational affixes	343
 Part 4 Pragmatics	359
18 Speech acts	361
19 Reference and deixis	379
20 Conversational implicatures	411
 <i>Epilogue</i>	447
<i>Answers to questions</i>	451
<i>References</i>	469
<i>Author index</i>	479
<i>Subject index</i>	482

Typographic conventions

Small capitals

For concepts; occasionally for lexical roots

Small capitals in square brackets

For semantic features

Angled brackets

For selectional restrictions

Bold type

For technical terms when first introduced

Italics

For citation forms when not set as displayed examples

Bold italics

For emphasis and stress placement

Single quotation marks

For quotations from other authors; 'scare quotes'

Double quotation marks

For meanings

Question marks

For semantic abnormality

Asterisks

For ungrammaticality or extreme semantic abnormality

From the preface to the first edition

The aim of this book is not to present a unified theory of meaning in language (I am not even sure that that would be a worthwhile project), but to survey the full range of semantic phenomena, in all their richness and variety, in such a way that the reader will feel, on completing the book, that he or she has made face-to-face contact with the undeniably messy 'real world' of meaning. At the same time, it aims to show that even the messy bits can, at least to some extent, be tamed by the application of disciplined thinking. As far as semantic theories are concerned, I have been unashamedly eclectic, adopting whatever approach to a particular problem seems genuinely to shed light on it. If there is a theoretical bias, it is in favour of approaches which, like the cognitive linguistic approach, embrace the continuity and non-finiteness of meaning.

This is not intended to be a 'baptismal' text; it would probably not be suitable for absolute beginners. The sort of readership I had in mind is second- or third-year undergraduates and beginning postgraduates who have completed at least an introductory course in linguistics, and who require an overview of meaning in language, either as preparation for a more detailed study of some particular area, or as background for other studies. I would hope it would be found useful, not only by students of linguistics, but also students of ancient and modern languages, translation, psychology, perhaps even literature.

Most of the material in the book has grown out of courses in general semantics, lexical semantics, and pragmatics, given to second- and third-year undergraduates and postgraduates at Manchester University over a number of years. I owe a debt to generations of students in more than one way: their undisguised puzzlement at some of my explanations of certain topics led to greater clarity and better exemplification; critical questions and comments not infrequently exposed weaknesses in the underlying arguments; and very occasionally, a genuine flash of insight emerged during a classroom discussion.

The final form of the text was significantly influenced by constructive comments on a draft by Jim Miller of the University of Edinburgh, an anonymous American reviewer, and John Davies of Oxford University Press, although, of course, full responsibility for remaining imperfections lies with myself.

Preface to the second edition

In this revised edition, the original overall plan, purpose, and character of the book have not been changed. The revisions are of three principal kinds. The first involves updating. Although only three years have elapsed since the publication of the first edition, there have been significant theoretical advances, particularly in lexical semantics. Accordingly, a new chapter has been added (Chapter 14), which outlines new thinking on word meaning. References have also been brought up to date. The second type of revision is the filling in of perceived gaps. There are of course constraints of space, and some omissions are unavoidable, but I have become aware of missing topics which really deserve a mention, such as frames, generalized conversational implicatures, and constraints on relevance, to name but three. In the end, of course, the selection is a personal one. The third sort of revision has involved the improvement of existing treatments of topics, sometimes by reorganizing whole chapters, as in the case of Chapters 2 and 18 (originally Chapter 17), in other cases by rewriting short passages to improve clarity and/or consistency.

In preparing this edition, I have benefited from the comments and suggestions of two anonymous referees, although I of course accept full responsibility for the final result.

Preface to the third edition

This edition contains a number of changes vis-à-vis the second edition. The most obvious is that Part Three, on grammatical semantics, has been divided into separate chapters. In addition to the material covered in the second edition, there are two new chapters, one on the semantics of prepositions and one on the semantics of derivational affixes. There has also been a good deal of reorganization, in an attempt to achieve a more logical grouping of material. For instance, quantifiers are now dealt with in Chapter 2, on logical matters, where they naturally belong. Readers familiar with the second edition may need to consult the table of contents and the index to locate topics. A number of chapters have been considerably expanded and updated.

In preparing the new edition, I have benefited greatly from detailed questions and comments on the second edition by Ernie Davis and Jim Feist, and from conversations on many topics with Delia Bentley. I am, of course, entirely responsible for the final text.

Part 1

Fundamental Notions

In this first section of the book, a number of fundamental, but fairly general, notions are introduced, which need to be grasped before the more detailed discussions in later sections can be properly appreciated. Chapter 1 has a scene-setting function, identifying the place of linguistic signs and linguistic communication in the broader domains of semiotics and communication in general. Chapter 2 introduces a number of vital conceptual tools drawn from the field of logic. Chapter 3 outlines the nature and properties of concepts, which underlie a conceptual approach to meaning.

CHAPTER 1

Introduction

1.1 Communication	5	1.4 The linguistic study of meaning	15
1.2 Language as a sign system: some basic notions of semiotics	10	1.5 Plan of the book	18
1.3 Approaches to the study of meaning	14		

CHAPTER 1

Introduction

1.1 Communication

Meaning makes little sense except in the context of communication: the notion of communication therefore provides as good a place as any to start an exploration of meaning. Communication can be conceived very broadly, including within its scope such matters as the transfer of information between biological generations via the genetic code, the interaction between a driver and his car, and indeed any sort of stimulus-response situation. Here we shall take a narrower perspective on communication and confine ourselves principally to what is surely the paradigm communicative scenario, namely, linguistic interaction between human beings.

1.1.1 Linguistic communication: an overview

The following components or stages can be identified in a typical act of linguistic communication involving a speaker and an addressee:

- (i) The speaker normally has a purpose in communicating.
- (ii) The speaker constructs a message to be communicated.
- (iii) The speaker constructs an utterance with which to convey the message.
- (iv) The speaker transforms the utterance into a physical signal.
- (v) The speaker transmits the signal.
- (vi) The addressee receives the signal.
- (vii) The addressee decodes the signal to recover the utterance.
- (viii) The addressee reconstructs the message from the utterance.
- (ix) The addressee infers the purpose of the communication.

Let us now look at the individual stages in more detail.

1.1.2 Purpose

It is probably not controversial to say that every act of communication (or attempted act) is driven by some desire or purpose, or complex of desires or purposes on the part of the communicator. But we can say more. Communication

is a social act and presupposes an audience. While the purpose of the communication is the fulfilment of a desire or achievement of an aim, the means of achieving this is via the effect of the communicative act on the audience. We may start, then, with a speaker (confining our attention, for the sake of simplicity, to spoken communication) who is motivated to attempt to achieve some goal by producing an utterance. Obviously, some utterances will be better adapted than others for the achievement of the speaker's goal, so a major task for the speaker is to construct an appropriate utterance. This is no simple matter, even though as speakers we are generally not aware of the complications.

1.1.3 Message

An essential component of the vast majority of messages that humans transmit linguistically consists of one or more propositions. (As we shall see, however, it is not possible for a message to consist *only* of propositions.) We shall not enquire deeply into the nature of propositions here (for more detailed discussion, see Chapter 2): it is enough for present purposes to say that propositions may be viewed as mental entities that correspond to potential events or states of affairs in the world, and are true or false with respect to some actual events or states of affairs. (The reader is warned that the topic of propositions is a philosophical minefield.)

Two important points must be made at this juncture. The first is that propositions are not to be confused with sentences. Take the sentence *John loves Mary*. Is this true or false? The question does not make sense, even though we understand the sentence perfectly well, unless and until we can identify some specific individual designated by *John*, a second individual designated by *Mary*, and the time at which the sentence was uttered. Only with this information to hand can we identify the proposition expressed by the utterance. This leads to a second point, namely, that a proposition or set of propositions does not exhaust the content of a message. Two types of non-propositional meaning are especially important, what is sometimes called **propositional attitude**, and **expressive meaning**.

Identifying a proposition is only part of the process of interpreting a message. A proposition on its own does not actually communicate anything. The addressee needs to know, for instance, whether the speaker is presenting the proposition as true (as opposed to expressing it, say, for the purpose of ridiculing it), and if so, with what strength of commitment to its truth; whether it forms part of a question or command; whether it is intended as a promise, warning, prediction, or complaint, and so on. (Matters of this sort are discussed in Chapter 18 under the heading of 'speech acts'.) In addition to this type of psychological orientation towards the proposition expressed, a speaker may also wish to convey attitudes or emotions with regard to the things, people, or events, and so on referred to in the utterance (for more on this aspect of meaning see Chapter 10). Think, for instance, of the role of the (currently ubiquitous) expletive in *He always forgets to lock the*

fucking door: including it or omitting it has no effect on the proposition expressed, but it is nonetheless an important part of the message.

1.1.4 Constructing the utterance

Let us now put ourselves in the position of a speaker who has a purpose in mind, who knows what message they wish to convey in order to achieve that purpose, and now has to formulate an utterance to transmit to the hearer. (The sequence of events described here must not be taken too literally: there is undoubtedly a lot of 'parallel processing' going on; also, a lot of the decision-making occurs at an unconscious level.) The message, let us assume for the sake of simplicity, consists of a proposition and a propositional attitude. The linguistic resources at the disposal of the speaker consist of words, grammar, and prosodic features such as intonation and stress. Again, for the sake of simplicity, we shall put aside prosodic matters, and we shall also ignore paralinguistic features such as facial expressions and body language (without, however, implying that these are unimportant).

There are two important basic facts about the relationship between messages and utterances which must be confronted at this point. The first concerns the relation between **type** information and **token** information. Take the utterance *The cat is sitting on the windowsill*. Normally we can take it that the proposition this sentence is intended to express relates to a particular cat sitting on a particular windowsill on a particular occasion. Consider the word *cat*. There is nothing in the meaning of *cat* which points to any particular cat: *cat* gives only type information, namely, that the individual entity referred to belongs to the category CAT. However, in order to understand the message, an addressee will need to be able to identify which cat the utterance is about, that is to say, will need token information. This information cannot be encoded directly in an utterance: language encodes only type information. Even mentioning the cat's name does not do this: *Socks is sitting on the windowsill*. This indicates to the addressee that a specific cat must be identified, as does the definite article in *the cat*, but the actual identity of the cat is not encoded in the name, since many cats have the same name. The same applies to *the windowsill*. (For more details on definiteness and reference, see Chapter 19.) It is a general feature of language that token information cannot be directly encoded.

The second important point is that utterances rarely directly encode all of the type information necessary to understand the message. Speakers normally expect addressees to work out some parts of a message for themselves. The following is a simple example:

- (1) A: Have you cleared the table and washed the dishes?
 B: I've cleared the table.

The normal interpretation of B's answer is that B has not washed the dishes. This type of inference is called an **implicature** (for fuller treatment see Chapter 20).

But this is not encoded in the utterance (although it could have been—we are not talking here about what language inherently cannot express). Not needing to spell everything out is a valuable economy feature in language.

These two basic points have important consequences for a speaker constructing an utterance. First of all, the speaker must directly encode sufficient information to enable the hearer to be able to identify easily the things, people, and so on referred to in the utterance. Crucially, exactly what information is necessary for this purpose varies enormously with the context of the communicative act. For present purposes, we can consider context under three headings: physical context, human context, and discourse context.

Physical context: In a conversation, this refers to things that the participants can see, hear, and so on, and each participant can assume is accessible to the other(s). So, for instance, if there is a knock at the door in the middle of a face-to-face conversation, one of the participants can say *That'll be John*. But if it is a telephone conversation, then more would be needed: *There's someone at the door. It's probably John*.

Human context: This refers basically to the speaker's assessment of characteristics of the addressee(s) which are liable to affect their ability to comprehend the message. A speaker will formulate an utterance differently according to their assessment of such matters as the addressee's age, intelligence, and relevant knowledge. For instance, the same message is likely to be conveyed by different utterances when speaking to a child than when speaking to an adult, or to a layman and an expert, or to someone with a limited command of a language and a native speaker.

Discourse context: The discourse context of an utterance consists of the utterance(s) which preceded it in the same discourse, whether a conversation or a text. For instance, *Will you get one for me, too, while you're there* is unlikely to be intelligible to someone unaware of what has been said previously.

The above parameters of context are obviously also relevant to a speaker's decision regarding the balance between explicit and implicit type information—how much (and what) to encode in the utterance and how much to leave to the addressee to figure out. If the speaker relies too heavily on the addressee's powers of inference, there is a danger that the wrong message will be conveyed.

So far, we have been concerned only with the contents of messages and the efficiency of utterances to convey them. But there is rather more to communication than the conveyance of messages (as understood here). For instance, an utterance might not achieve the speaker's purpose, in spite of conveying a clear message, if it is not appropriate to the context. There are various ways of being inappropriate to context. For instance, if a speaker pays no attention to the norms of politeness or cultural sensibilities, they risk offending the addressee; 'speaking like a textbook' in the context of an informal conversation, or conversely, using racy colloquialisms to a patient during a medical consultation, or sermon, or job interview can all

be counter-productive with respect to a speaker's purpose (of course the effect might be deliberate, in which case it should be considered part of the message).

1.1.5 Encoding the utterance as a signal

Since utterances in their initial form as mental representations cannot be transmitted directly (telepathy, if it exists at all, does not appear to be reliable), they must be converted into a form that can be transmitted, namely, a **signal**. This involves translating the linguistic form into a set of instructions to the speech organs, which, when executed, result (in the case of a spoken communication) in an acoustic signal.

1.1.6 Transmission and reception of the signal

Every mode of communication has a **channel** through which the signal travels: for speech, we have the auditory channel, for normal writing and sign language, the visual channel, for Braille, the tactile channel, and so on. As the signal travels from sender to receiver, it alters in various ways, through distortion, interference from irrelevant stimuli, or loss through fading. These changes are referred to collectively as **noise**. As a result, the signal picked up by the receiver (the **received signal**) is never precisely the same as the transmitted signal. If every detail of the transmitted signal was crucial for the message being transmitted, communication would be a chancy business. However, efficient communicating systems like language compensate for this loss of information by building a degree of redundancy into the signal. Essentially this means that the information in a signal is given more than once, or is at least partially predictable from other parts of the signal, so that the entire signal can be reconstructed even if there is significant loss of detail during transmission. This is helped by the discrete nature (see below, Section 1.2.3) of linguistic signs. Think of a printed word: it generally takes quite a lot of physical degradation to render it unidentifiable. It is said that language is roughly 50 per cent redundant.

1.1.7 Decoding the signal

Once the signal is received, the addressee's first job is to decode it in order to recover the original form of the utterance as intended by the speaker. This involves recognizing the words and the prosodic cues, retrieving meanings, and assembling these into a representation of a full grammatical structure. There is, however, a complication, even at this initial stage of decoding: it is not always possible to construct a full representation of an utterance purely on the basis of the acoustic signal transmitted. This is because of the phenomenon of ambiguity—the same sound representing two (or more) different structures. For instance, (2) and (3) can sound exactly the same, even though they differ both lexically and grammatically:

- (2) I've scent for my mother.
- (3) I've sent for my mother.

In such cases, disambiguation—deciding which structure was intended by the speaker—requires context to be taken into account.

1.1.8 Construal of the full message

Arguably the central task for the addressee is to work out the speaker's message on the basis of the utterance, taken together with relevant features of the context. As we have seen, by no means all the message is encoded in the utterance. The major aspects of the message which are not encoded in the utterance are the identity of referents, and (by definition) implicatures. A wide range of message components can be conveyed via implicatures. Many of them are propositional in nature, but many are not. Just to take one example, politeness can be directly encoded by such expressions as *thank you* and *please*, but is often conveyed by means of implicatures. In (4) and (5), for instance, politeness is not correlated with form, and is conveyed by implicature:

- (4) (i) Wash the dishes. (less polite)
- (ii) Would you mind washing the dishes? (more polite)
- (5) (i) Have some more cake. (more polite)
- (ii) Would you mind having some more cake? (less polite)

A fuller treatment of politeness can be found in Chapter 20.

1.1.9 Inference of purpose

The final stage in the comprehension of a communication is to infer some sort of answer to the question "What was the speaker's purpose in sending me this message?" and to decide how to respond. It is clear even from the simplified picture of linguistic communication presented above that although it is mostly experienced as rapid and effortless, it is in fact an astonishingly complex process. It is not surprising, therefore, that utterances on occasion fail to achieve their purpose.

1.2 Language as a sign system: some basic notions of semiotics

Any natural human language is a complex sign system, 'designed' to ensure infinite expressive capacity, that is to say, there is nothing that is thinkable which represents a 'no-go area' for language (provided no limit is placed on the complexity of utterances). Each elementary sign is a stable symbolic association between a meaning and a form (phonetic or graphic); elementary signs may combine together in a rule-governed way to form complex signs which convey correspondingly complex meanings. Below we consider how certain notions drawn from semiotics (the study of signs and sign systems) apply to language.

1.2.1 Iconicity

Signs can generally be classified as iconic or arbitrary. **Iconic** signs are those whose forms mirror their meanings in some respect; signs with no natural analogical correspondences between their forms and their meanings are called **arbitrary**. A simple example is provided by the arabic and roman numerals for "three": 3 and III. The arabic form gives no clue to its meaning; the roman version, on the other hand, incorporates "threeness" into its shape, and is thus iconic. Iconicity is a matter of degree, and usually coexists with some degree of arbitrariness. Three horizontal lines would be just as iconic as the roman III: the fact that in the roman symbol the lines are vertical is arbitrary.

Iconicity enters language in several guises. The majority of words in a natural language are arbitrary: the form of the word *dog*, for instance, does not mirror its meaning in any respect. However, the so-called **onomatopoeic** words display a degree of iconicity, in that their sounds are suggestive (to varying degrees) of their meanings:

bang clank tinkle miaow splash cuckoo peewit curlew
whoosh thud crack ring wheeze howl rumble (etc.)

The predominance of arbitrariness in the vocabulary is not an accidental feature, but is a crucial **design feature** of natural language. There is a limited stock of dimensions of formal variation in linguistic signs; if all signs were iconic, it is difficult to see how universal expressivity could be achieved.

Some iconicity is also apparent in grammar. For instance, words which belong together tend to occur together. In *The tall boy kissed the young girl* we know that *tall* modifies *boy* and not *girl* because *tall* and *boy* come next to each other in the sentence. In some languages this relationship might be shown by grammatical agreement, which is a kind of resemblance, and therefore also iconic. Another way in which iconicity appears in the grammar is that grammatical complexity by and large mirrors semantic complexity.

1.2.2 Conventionality

Many of the signs used by humans in communication are **natural** in the sense that they are part of genetically inherited biological make-up and do not have to be learned, although a maturational period may be necessary before they appear in an individual, and they may be moulded in various ways to fit particular cultural styles. The sort of signs which are natural in this sense will presumably include facial expressions like smiling, frowning, indications of fear and surprise, and so on, perhaps many of the postural and proxemic signs (involving physical distance between interlocutors) that constitute the so-called 'body language', certain types of gesture, vocal indications of excitement, desire, etc. (whether or not linguistic), and many more. Natural signs are likely to be the most cross-culturally interpretable.

Other signs have conventionally assigned meanings; they have to be specifically learned, and are likely to differ in different communities. Linguistic signs are the prototypical **conventional** signs. Even onomatopoeic words usually have a significant conventional element; often the iconic nature of the word can only be appreciated, as it were, with hindsight. Take the Turkish word *bülbül*. What does it refer to? A baby's babbling? The noise of a mountain spring? In fact, it means "nightingale". Looking back, one can make the connection. It is not only linguistic signs that are conventional. Obscene or offensive gestures, for instance, can vary quite remarkably cross-culturally: I was once reprimanded for pointing the soles of my feet at the Prime Minister of Iraq (in Arab culture this is disrespectful: my disrespect was entirely inadvertent). Even in Europe, conventional gestures can differ: Greeks are famously—and slightly inaccurately—said to shake their heads to say "Yes", and nod to say "No".

1.2.3 Discreteness

Some signs can vary gradually in their form, and their meaning varies in parallel with the change of form, like the fisherman's indication of the size of 'the one that got away'; these are called **continuous signs**. Other signs have fixed shapes, and must be chosen from a limited inventory: intermediate forms are not allowed, the range of possibilities is 'chunked'; such signs are described as **discrete**. Linguistic signs are virtually all of the discrete variety. Again, this is not an accidental feature, but has a close connection with iconicity and arbitrariness: continuous signs are necessarily iconic; arbitrary signs are necessarily discrete.

1.2.4 Double articulation

Simple signs can be combined to form complex signs which signal complex meanings. But simple signs themselves typically display complexity in their formal aspect. Thus, the formal aspect of the sign *cat* (as pronounced) is built up from the simpler phonological units /k + æ + t/, which are not individually associated with grammatical or semantic properties. This feature is usually termed **double articulation**. It allows a language to have a very large inventory of elementary signs without overstressing the discriminatory and memory capacities of human language users (imagine if every word in a language had to be realized as a single phoneme).

1.2.5 Simple and complex signs

Linguistic signs may be simple or complex. This does not just mean that they can occur singly or in groups of various sizes: the occurrence of two or more signs together does not necessarily result in a complex sign. Take the case of someone who answers a question with the word *Yes*, at a higher than usual pitch, and at the same time smiling. This person has not produced a complex sign with three constituents, only three simple signs simultaneously. The meanings of the three

signs are simply added together: there is no interaction between the signs other than additivity. Contrast this with a minimally complex sign such as *red wine*: to obtain the meaning of this sign, we do not simply add the meaning of *red* to the meaning of *wine* (that would give us something like “wine plus redness”). What happens is that the meaning of *red* interacts with the meaning of *wine* by restricting it.

The syntax of a language is essentially a system of rules for combining simple meanings together in a systematic way to form more complex meanings, through combining simple signs to form complex signs. There is no theoretical upper limit to the complexity of linguistic signs. This is rendered possible by the **recursive** nature of syntax, that is, the existence of rules which can be applied indefinitely many times (like the one which yields *This is the dog that worried the cat that killed the rat that ate the corn that*). Such rules are an essential prerequisite for the ‘universal expressivity’ of language—the fact that anything thinkable is expressible, or at least can be approximated to any given degree of accuracy.

1.2.6 Linguistic, paralinguistic, and non-linguistic signs

The prototypical scenario for linguistic communication is two people engaged in face-to-face conversation. Of course, in such an encounter, language signals are exchanged; but normally so are many other types of signal, and these modify and/or supplement the linguistically encoded message. Let us, then, briefly look at the semiotic environment of language in a typical conversation.

The signs that accompany language can be divided into two major types—**paralinguistic** and **non-linguistic**. The defining characteristic of paralinguistic signs will be taken here to be an extreme dependence on the accompanying language. Either they cannot be produced except during speech (because they are carried on the voice), or they cannot be interpreted except in conjunction with accompanying language. Examples of the first variety are abnormal volume, tempo, pitch, and voice quality. For these to function as signs, there must be a departure from some (personal) baseline or norm. For instance, abnormally high volume, fast tempo, or high pitch typically signal a heightened emotional state. Examples of the second variety include pausing, emphatic gestures, and gestures which metaphorically depict, for instance, direction of motion.

The functions of paralinguistic signs can be conveniently classified under three headings:

- (i) **punctuation**: there are signs which have functions parallel to those of punctuation in written language, mainly to segment the stream of speech so as to facilitate processing.
- (ii) **modulation**: this involves the addition of an emotive or attitudinal colouring to the linguistically encoded message.
- (iii) **illustration**: some signs ‘depict’ a key element in the message, such as a direction of movement, or a shape; the depiction may be relatively literal, like the hand movements of someone describing the climbing of a spiral

staircase, or metaphorical, as when vertical and parallel hands accompany the setting of limits of some kind.

Not all the signs that occur alongside language are paralinguistic in the sense defined. For instance, one may smile or frown while speaking, and this may well modulate the message. But smiles and frowns (and many other signs) are perfectly interpretable and producible in the absence of any accompanying language. These are therefore to be considered as **non-linguistic**.

1.3 Approaches to the study of meaning

Meaning may be studied as a part of various academic disciplines. There is of course a significant degree of overlap between disciplines, but characteristically all have something idiosyncratic and unique in their approach (the following remarks are merely illustrative).

1.3.1 Philosophy

Linguists typically take the existence of meaning for granted and accept it as an intuitively accessible 'natural kind'. They do not ask questions like *How is it possible for anything to mean something?* or *What sort of relation must hold between X and Y for it to be the case that X means Y?*. These kinds of questions are the province of the philosopher, particularly the philosopher of language.

1.3.2 Psychology

Meaning is a major concern of the psychology of language and psycholinguistics (I shall not attempt to distinguish these). A distinctive feature here is the experimental study of how meanings are represented in the mind, and what mechanisms are involved in encoding and decoding messages. An example of a fact that could only emerge within a psycholinguistic framework is that in the lexical decision task, where experimental subjects observe strings of letters flashed on a screen and must indicate by pressing the appropriate button whether the string represents a word or not, responses are faster to words with concrete meanings than to words with abstract meanings, even when extraneous factors like length and frequency are rigorously controlled. This observation presumably provides a clue to the role of meaning in word recognition (to the best of my knowledge it is still a puzzle).

1.3.3 Sociology

The major concern of sociological approaches to language and meaning (sociolinguistics and the sociology of language) is the role that language plays in society and in the creation and maintenance of social relations.

1.3.4 Neurology

Psychologists take a 'macro' view of mental states and processes. Neurologists, on the other hand, want to know how these states and processes are implemented at the neuronal level. A psychologist might be broadly compared with a computer programmer, and a neurologist to the designer of computer chips. Meaning, like everything else in mental life (at least if one is a physicalist) must boil down ultimately to connections between neurons.

1.3.5 Semiotics

Semioticians view language as one sign system amongst many, and seek out those features which render it so successful. They are also likely to give emphasis to marginal aspects of linguistic signification. The recent strong interest in iconicity in language represents a significant overlap between the linguistic and semiotic approaches to meaning.

1.3.6 Statistics

The availability of large bodies of data concerning the actual use of language (corpora), both written and spoken, in a form which allows for sophisticated statistical analysis, has contributed much to the study of language. Of particular interest to semanticists are the patterns of co-occurrence of words in natural texts.

1.3.7 Linguistics

It is not easy to encapsulate the linguistic approach to meaning in language succinctly. There are perhaps three key aspects. The first is that native speakers' semantic intuitions are centre-stage, in all their subtlety and nuances: they constitute the main source of primary data. The second is the importance of relating meaning to the manifold surface forms of language. The third is the respect paid not just to language, but to languages in all their variety.

1.4 The linguistic study of meaning

1.4.1 What are we trying to achieve?

1.4.1.1 Describing meanings

It is usually assumed that linguistic expressions can be assigned some sort of context-independent semantic value, although there is much disagreement regarding exactly what this is. A very important task is to discover a way of specifying or describing such meanings. Among other things, this involves recognizing different types of meaning, such as the distinction mentioned earlier between propositional and non-propositional meanings, and within the latter, the distinction between

propositional attitude and expressive meaning. The position taken in this book is that in general, meanings are not finitely describable, so this task boils down to finding the best way to approximate meanings as closely as is necessary for current purposes (lexicographers have long had to confront this problem for words).

I.4.1.2 How meaning varies with context

The meanings of all linguistic expressions vary with the context in which they occur. For instance, the shade of colour indicated by (natural) *red hair* and *red wine* are markedly different; the periods of time denoted by *month* in (6) and (7) are quite likely to be different:

- (6) He's here for a month. (could be four weeks; not dependent on time of utterance)
- (7) He's here for the month. (will depend on time of utterance, but could be 31 days)

Some variations, like the sex of the doctor in *Our doctor has just married a policeman* and *Our doctor has just married an actress* can be predicted by general principles; other variants are less, or not at all predictable. The generation of implicatures is another important aspect of the influence of context on meaning. Scholars concerned with linguistic meaning seek a revealing account of contextual variation.

I.4.1.3 What happens when meanings combine?

Another vital aspect of semantics is how simple(r) meanings combine to form more complex meanings. To some extent this is a function of grammatical structure: for instance, the way *red* and *hat* combine in *a red hat* is not the same as the way *turn* and *red* combine in *to turn red*. But differences occur even within the same grammatical construction: the mode of combination of *red* and *hat* in *a red hat* is different from that of *long* and *eyelash* in *long eyelashes* (compare the import of *long* in *long eyelashes* and *a long river*).

I.4.1.4 New meanings from old

A striking feature of linguistic expressions is their semantic flexibility: beyond their normal contextual variability, they can be bent to semantic ends far removed from their conventional value, witness *She swallowed it hook, line, and sinker* or *You'll find her in the telephone book*. The study of such extensions of meaning is an important task for semantics.

I.4.1.5 Systematicity and structure: the possibility of formalization

All semanticists are to some extent looking for regularities and system in the way meanings behave, as this leads to maximally economical descriptions. The most dedicated to this aspect of semantics are those who attempt to model the semantic

behaviour of natural language expressions by means of a strict logical or quasi-mathematical formalism. This route will not be followed in this book.

1.4.2 Branches of the study of linguistic meaning

1.4.2.1 Lexical semantics

Lexical semantics studies the meanings of words; the focus here is on 'content' words like *tiger*, *daffodil*, *inconsiderate*, and *woo*, rather than 'form'/'grammatical' words like *the*, *of*, *than*, and so on. To a non-specialist, the notion of meaning probably has a stronger link with the idea of the word than with any other linguistic unit: words are, after all, what are listed in dictionaries, and the main function of a dictionary is to tell us what the listed words mean. For this reason, lexical semantics perhaps provides the easiest access route into the mysteries of semantics in general, and this is one reason why it has been given a prominent place in this book.

1.4.2.2 Grammatical semantics

Grammatical semantics studies aspects of meaning which have direct relevance to syntax. This has many manifestations, which can only be briefly illustrated here. One problem is the meaning of syntactic categories (problematic, because not everyone believes they can be assigned meanings). Consider, for instance, the differences in the meaning of *yellow* in the following:

- (8) She wore a yellow hat. (adjective)
- (9) They painted the room a glowing yellow. (noun)
- (10) The leaves yellow rapidly once the frosts arrive. (verb)

Another aspect of grammatical semantics is the meaning of grammatical morphemes like the *-ed* of *walked*, the *-er* of *longer*, the *re-* and the *-al* of *retrial*, and so on. Clearly grammatical semantics overlaps with lexical semantics, partly because some grammatical elements are words (like *the* and *of*), but more particularly because some aspects of the meanings of full lexical items determine to some degree their grammatical behaviour (for instance, the fact that *I am studying that question* is grammatical, but not *I am knowing the answer to that question*).

1.4.2.3 Formal semantics

Formal semantics studies the relations between natural language and formal logical systems such as the propositional and predicate calculi. Such studies usually aim at modelling natural language as closely as possible using a tightly controlled, maximally austere logical formalism. It is arguable that sometimes such studies shed more light on the formalism used than on the language being modelled; nonetheless, valuable insights have come from this approach. To date, most (but not all) such studies have concentrated on the propositional/sentential level of meaning, without attempting to delve into the meanings of words.

1.4.2.4 Linguistic pragmatics

For present purposes, pragmatics can be taken to be concerned with aspects of information (in the widest sense) conveyed through language which (a) are not encoded by generally accepted convention in the linguistic forms used, but which (b) nonetheless arise naturally out of and depend on the meanings conventionally encoded in the linguistic forms used, taken in conjunction with the context in which the forms are used. This rather cumbersome formulation is intended to allow into pragmatics things like the identity of the individual referred to by *John* in *I saw John today*, and the assumption that the room in question had several lights in *John entered the room; all the lights were on*, at the same time excluding, for instance, the possibility that the person saying *I saw John today* had a private ad hoc arrangement with the hearer that whenever he said *John*, he should be taken to mean "Mary" (since it does not arise naturally out of the normal meaning of *John*), and excluding also the possibility of someone's inferring from a speaker's slurred speech that s/he was drunk (since this does not depend on the conventional meanings of the words uttered). Pragmatics is usually contrasted with semantics, the latter dealing with conventionalized meaning. The distinction is drawn in different ways by different scholars, but since the range of topics covered by semantics + pragmatics is largely the same for everyone, not much will be made of this distinction in this book.

1.4.2.5 The approach adopted in this book

We are not yet in a position to rule out any approaches which yield insights, even if some such approaches appear at first sight incompatible. This book therefore takes an ecumenical position on many issues. Insofar as there is a theoretical bias, it is towards the cognitive semantic position. This means, in particular, that the meaning of a linguistic expression is taken to arise from the fact that the latter gives access to a particular conceptual content. This may be of indeterminate extent: no distinction is made between linguistic meaning and encyclopaedic knowledge.

Since this book is not intended to propound a body of theory, but to acquaint non-specialists with the range of semantic phenomena in language, there is a bias towards descriptive coverage at the expense of theoretical rigour.

1.5 Plan of the book

The organization of the book is as follows. It is in four parts. Part 1 discusses a range of basic notions that underlie virtually all discussions of meaning within linguistics; Part 2 concentrates on aspects of the meanings of words; Part 3 deals with semantic aspects of grammar; Part 4 introduces the core areas of pragmatics, and highlights the relations between meaning and context.

Within Part 1, Chapter 1 provides a very general introduction to questions of meaning, locating the linguistic study of meaning within the wider context of the study of signs and communication in general. Chapter 2 introduces a set of fundamental conceptual tools, mostly drawn from the field of logic, which, because of their wide currency in discussions of semantic matters, constitute indispensable background knowledge for a study of meaning in language. Chapter 3 outlines the nature of concepts and provides a basis for the conceptual approach to meaning adopted in the rest of the book.

In Part 2, Chapter 4 provides a general introduction to the study of lexical meanings, first discussing whether there are any restrictions on what sort of meanings lexical items can bear, then discussing the relationship between lexical items—the basic units of composition—and words. In Chapter 5, the focus is on the range of variation observable in a single word form in different contexts, ranging from the major differences displayed by homonyms to subtle modulations of sense. Chapters 6 and 7 deal with relations of sense between lexical items which can occupy the same syntactic position—in other words, paradigmatic sense relations, such as hyponymy, meronymy, incompatibility, synonymy, antonymy, complementarity, reversivity, and converseness. Chapter 8 looks at the most important paradigmatic structures in the vocabulary of a language, namely, lexical hierarchies. In Chapter 9, meaning relations between words in the same syntactic construction, that is, syntagmatic sense relations, are examined. Topics discussed include the nature of normal and abnormal collocations, reasons for a tendency for certain types of words to co-occur, and the nature and consequences of selectional pressures of words on their partners in a string. Chapters 10 and 11 deal with the problems of describing lexical senses, Chapter 10 concentrating on basic dimensions and structures of meaning, and Chapter 11 on theoretical approaches. Chapter 12 describes the main types of process, such as metaphor and metonymy, which enable new meanings to be produced from old ones.

Part 3 deals with grammatical, as opposed to lexical meaning. Chapter 13 seeks to clarify the notion of grammatical meaning before going on to examine meanings associated with nouns and noun phrases, such as number and gender. The subsequent chapters in the section illustrate the sorts of grammatical meaning associated with different levels of grammar. Chapter 14 looks at the clause-level features of argument structure and transitivity, Chapter 15 at meanings associated with verbs and verb phrases, such as tense and aspect, and with adjectives, such as degree. Chapter 16 considers the sorts of meanings typically carried by prepositions, and Chapter 17 examines the semantics of a representative selection of derivational affixes.

Part 4 covers topics which are usually considered to fall under pragmatics, in that either they involve aspects of meaning which cannot be satisfactorily treated unless context is taken into account, or they are not propositional in nature (or both). Chapter 18 provides an outline of speech act theory, mainly following Austin and Searle. It discusses the acts that people perform when they are speaking—acts such as stating, requesting, warning, congratulating,

commanding, and so on. The range of different types of such act is surveyed and their nature examined. Chapter 19 is concerned with reference, i.e. establishing connections between utterances and the extralinguistic world. Reference is portrayed as the assigning of values to variables, the variables being signalled by definite expressions and the values being items in the extralinguistic world. Various strategies for indicating (on the part of the speaker) and determining (on the part of the hearer) correct referents are discussed, including the use and interpretation of deictic elements, names and descriptions. Chapter 20 deals with conversational implicatures, that is, those aspects of the intended meaning of an utterance which are not encoded in its linguistic structure, but are, as it were, 'read between the lines'. Different types of conversational implicature are described and some proposed explanations of how they arise are considered.

Suggestions for further reading

Much fuller accounts of the semiotic environment of spoken language can be found in Argyle (1972), Beattie (1983), and Ellis and Beattie (1986). For a stimulating coverage of broadly the same range of topics as the present volume, but from a different perspective, see Seuren (2009). Allan (2001) has similar coverage, but is more advanced.

CHAPTER 2

Logic and meaning

2.1 Introduction	23	2.5 The logic of classes	40
2.2 Propositions	23	2.6 Presuppositions	42
2.3 Quantifiers and quantification	33	2.7 Possible world semantics	45
2.4 The logic of relations	39	2.8 Some useful distinctions	46

CHAPTER 2

Logic and meaning

2.1 Introduction

This chapter introduces a set of fundamental conceptual tools, mostly drawn from the field of logic, which, because of their wide currency in discussions of semantic matters, constitute indispensable background knowledge for a study of meaning in language. The level of treatment here is fairly elementary, but hopefully will be sufficient to encourage clarity of thought.

2.2 Propositions

Propositions were briefly introduced in the previous chapter. It is now time to examine them in a little more detail. Some people find the notion of a proposition hard to grasp, and indeed, there is a great deal of controversy in philosophy as to their true nature. For our purposes, however, we can sidestep most of the subtleties. A simple proposition attributes some property to an entity, or a relation between two or more entities. It is either true or false (even if it is not practically or even physically possible to ascertain which): truth and falsehood are essentially properties of propositions. The truth or falsehood of an utterance is a sign that at least one proposition has been expressed. A proposition is not a specifically linguistic entity (although we shall not dwell too long on the problem of what sort of entity it is, and in what Platonic or other realm it subsists). The same proposition may be expressed by an indefinitely large number of sentences:

John saw Mary.
John saw his sister.
Mary was seen by Peter's uncle.
John saw Jane's best friend.
etc.

These can all express the same proposition, provided, of course, that Mary is John's sister, and that John is Peter's uncle, and so on. So, for our purposes, we can

regard a proposition as a unit of meaning. But it is not merely a unit of meaning, it is a (potential) unit of knowledge (it is only a unit of knowledge if it is true). The item *cat* in *The cat sat on the mat* is the carrier of a unit of meaning, but it is not an item of knowledge.

2.2.1 Propositional logic: propositions as unanalysed units

We shall first of all consider propositions as elementary units, without enquiring into their internal structure. This perspective is known as 'first-order logic', and the system of expressing the relations between propositions in terms of formulae and theorems is known as the propositional calculus.

Before we go any further, it will be useful to consider how propositions are related to sentences, statements, and utterances.

Sentences. Let's start with sentences (confining ourselves to declarative sentences). Obviously we use sentences to express propositions. But sentences are not propositions. A sentence is basically a string of words constructed in accordance with the grammatical rules of some language and does not have a truth value. Take the sentence *The cat sat on the mat*. As it stands, it is neither true nor false. It becomes true or false when it is asserted of some specific cat and some specific mat on a specific occasion. Until definite referring expressions in a sentence have been assigned referents, it does not express a specific proposition. A sentence like *The cat sat on the mat* can be used to express an indefinitely large number of different propositions (i.e. with reference to different cats and different mats) on different occasions of use. Although the sentence *The cat sat on the mat*, as a linguistic structure abstracted from particular uses, does not have a truth value, it does nonetheless have a meaning that is different from that of, say, *The rat sat on the mat*. That is because although it does not express a proposition, it places limits on the range of propositions it can be used to express. Obviously, *The cat sat on the mat* can express propositions that *The rat sat on the mat* cannot, and vice versa. The range of propositions a sentence can be used to express is called its **propositional content**, and this is an important component of the semantic potential of the sentence.

Statements. Let us now consider what a statement is. First of all, we shall take it that a statement is a sentence that is used to express a proposition. But a proposition may be 'entertained', without any stance being adopted towards its truth or falsehood, as, for instance, in a logic class, where propositions are entertained, and their interrelationships studied, for example:

- (I) All writers on semantics are clairvoyant.
- Alan Cruse is a writer on semantics.
- Hence, Alan Cruse is clairvoyant.

These are not statements, although they express propositions. A proposition on its own cannot actually communicate anything: it is not an item of knowledge. To communicate, it must be energized with some kind of illocutionary force. For

instance, the proposition expressed by *The earth is a polyhedron* is of no interest until someone claims it is true: a true proposition is an item of knowledge. The illocutionary force in such a case would be that of assertion. We may define a (minimum) statement, therefore, as a sentence expressing a proposition, uttered, as Lyons puts it, "with epistemic commitment", that is, the proposition is presented as being true. A statement therefore has a truth value: the proposition being asserted may in fact be true or false. (There are other possibilities for illocutionary force, to be studied in due course.)

We have so far assumed that a statement must be a full grammatical sentence. But this is not necessary. Consider the following exchange:

- (2) A: What's the time?
B: Half past four.

Is B expressing a proposition with epistemic commitment? Obviously, yes. Then what proposition is expressed? Again the answer is obvious: "The time is 4.30". But this is not what is actually *said*. In such cases, it is up to the hearer to reconstruct the full form of the proposition on the basis of contextual clues, but there is no doubt that it is the full form of the proposition that the speaker intends to convey, and they are therefore making a statement.

Utterances. We shall take an utterance to be a piece of language produced on a particular occasion with a particular intent. It may or may not express a proposition, it may or may not carry any illocutionary force, it may or may not be a full sentence. An utterance is a token-level entity; a sentence is a type-level entity. If John says *The cat sat on the mat* and Mary says *The cat sat on the mat*, two utterances have been produced, but only one sentence. If John's utterance and Mary's utterance refer to different cats, then two different propositions have been expressed: the proposition, like the utterance, is a token-level entity.

The meaning carried by a statement *qua* statement (i.e. a proposition + illocutionary force) may not adequately represent everything the speaker intended to convey by uttering the sentence to which it applies. As an example, consider the following, mentioned in the previous chapter:

- (3) A: Have you cleared the table and washed the dishes?
B: I've cleared the table.

In normal circumstances, it would clearly be part of B's intended message that she had not washed the dishes. Yet this cannot be obtained by elaborating or completing the proposition expressed. This is even clearer in the following case where it is obviously the speaker's intention to convey the proposition that A is too late for supper:

- (4) A: Am I in time for supper?
B: I've cleared the table.

These extra (i.e. covert) propositions are expected to be inferred by the hearer on the basis of contextual information, but they go well beyond the mere filling out of

missing bits in what is actually said. We shall give the name **utterance meaning** to the totality of what the speaker intends to convey by making an utterance, within certain necessary limits.

Two people might, for instance, have an arrangement such that *How was your day?* means "How was your day?", but *How was the day today?* means "Mum and Dad will be out late this evening—we'll be on our own". We would wish to exclude this sort of meaning from utterance meaning. I think the key point is that such a use requires a special ad hoc stipulation on the part of speaker and hearer: the hearer's understanding of the clue does not arise entirely from his general knowledge of the rules and conventions governing the use of the forms in question.

2.2.1.1 Logical operators

First-order logic is concerned with relations between propositions and how these relations affect the truth values of compound propositions. Logical operators are elements of logical formulae which allow the expression of these relations. This will become clearer when we look at some examples.

The negative operator. The application of the negative operator, symbolized by ' $\sim$ ', to a proposition has the effect of switching its truth value. So, if some proposition P is true, then $\sim P$ is false, and if P is false, then $\sim P$ is true. Let P = "Cats are mammals". We can take it that this is true. Therefore $\sim P$ ("It is not the case that cats are mammals") is false. If Q is the clearly false "All bachelors are married", then $\sim Q$ ("It is not the case that all bachelors are married") is true.

Logical connectives are operators that serve to combine elementary propositions together to form compound propositions. The different connectives are defined by truth tables which show how the truth value of a compound proposition relates to the truth values of the component propositions. The principal logical connectives are implication, conjunction, disjunction, and equivalence.

Implication. Implication is signalled by the symbol ' $\rightarrow$ '. The truth-table definition is as follows:

P	Q	$P \rightarrow Q$
T	T	T
T	F	F
F	T	T
F	F	T

This can be translated as "It is never the case that P is true and Q is false". It is sometimes translated as "If P then Q ", but this is misleading, because, as the table shows, if P is false, then $P \rightarrow Q$ is true whatever the truth value of Q . (See the comments below on entailment.) Notice that the order of the propositions is significant for this relation.

Conjunction. Conjunction is signalled by the symbol '&'. Its truth-table definition is as follows:

P	Q	P&Q
T	T	T
T	F	F
F	T	F
F	F	F

This needs no comment: P&Q is true only if both P and Q are true. In this case the order of the propositions is irrelevant.

Disjunction. Disjunction is signalled by the symbol 'v', and is defined as follows:

P	Q	PvQ
T	T	T
T	F	T
F	T	T
F	F	F

This is often translated as "P and/or Q": PvQ is true as long as at least one of the two propositions P and Q is true. The order of the propositions is irrelevant.

Equivalence. Equivalence can be regarded as bi-directional implication, and is signalled by the symbol ' $\leftrightarrow$ '. It is defined as follows:

P	Q	P $\leftrightarrow$ Q
T	T	T
T	F	F
F	T	F
F	F	T

P $\leftrightarrow$ Q is true only if both P and Q have the same truth value. The order of the propositions is irrelevant.

2.2.1.2 Application to natural language

In their strict logical versions, the above connectives are not very useful for the semantics of natural language. But in a modified form they are very useful. The problem with them as they stand is that they do not depend in any way on the meanings of the conjoined propositions, but only on relations between their truth values. This gives rise to counter-intuitive results when applied to ordinary language. For instance, P implies Q provided that it is never the case that P is true and Q is false. This means that a necessarily false proposition implies any other proposition whatsoever, hence "All bachelors are married" implies "John is an

orthodontist". Intuitively, this is not what we mean by "If P then Q". It is sometimes said that a natural language equivalent of implication can be found in sentences of the form: *If John can speak Cree, then I'm a Dutchman* (meaning that since it is manifestly false that John can speak Cree, according to the speaker, then all propositions whatsoever are implied). What we need is a relation that holds between, say, "Fido is a dog" and "Fido is an animal", but not between "All bachelors are married" and "John is an orthodontist". We shall therefore replace implication with a stricter relation, entailment.

Entailment. Entailment is the relation which holds between the propositions listed under P and the corresponding propositions under Q in the following:

P	Q
Fido is a dog	Fido is an animal
John killed the wasp	The wasp died
All dogs are purple	My dog is purple

To say that proposition P **entails** proposition Q means that the truth of Q follows logically and inescapably from the truth of P, and the falsity of P follows likewise from the falsity of Q. So, in the first pair of propositions above, if it is true that Fido is a dog, then it follows ineluctably that Fido is an animal, and if Fido is not an animal, then there is no way Fido can be a dog. Similarly, in the second pair, if John killed some wasp, then we cannot avoid the conclusion that the wasp died, and if the wasp did not die, then it cannot be the case that John killed it. Unlike implication, entailment is a relation between meanings. Notice that if P entails Q, then it necessarily follows that P implies Q, but if P implies Q, then it does not necessarily follow that P entails Q.

Two other properties of entailment must be emphasized. The first is that the relation is not determined by context: it is context-independent, since it depends entirely on the meanings of the constituents of the sentences. Consider a case where John has in front of him a box of coloured disks, in which all the red disks are round in shape, and all the green disks are square. In such circumstances, the truth of "John picked a square disk from the box" follows inescapably from the truth of "John picked a green disk from the box". But clearly this relation of truth values does not arise from relations between *green* and *square*, but from the context: it would in principle have been just as easy to have all the red disks square, and the green disks round. On the other hand, the relation between "It's a dog" and "It's an animal" is independent of any particular contexts.

The second property is that the truth of the entailed proposition must follow inescapably from the truth of the entailing proposition. It is not enough for it to be usually true, or even almost always true; it has to be unthinkable that it might be false. Consider the relation between "It's a dog" and (i) "It's a pet" and (ii) "It can bark". Most dogs that most people encounter are pets, but there are such things as wild dogs, so the relationship is merely one of expectation. This is not entailment.

Likewise in the case of (ii), most dogs can bark, but a dog with a defective larynx does not thereby cease to be a dog, so the relationship is not logically necessary. Only logically necessary, context-independent relationships count as entailment.

The notion of entailment can be used to define a useful set of propositional relations. Where necessary, we shall use ' $\Rightarrow$ ' to indicate entailment.

Equivalence. Semantic equivalence can be straightforwardly defined as mutual entailment, which we can represent, following the pattern used for logical equivalence, as ' $P \Leftrightarrow Q$ '. The following are examples of equivalence:

John killed the wasp	The wasp was killed by John
The wasp is dead	The wasp is not alive
It began at 10 o'clock	It commenced at 10 o'clock

If it is true that John killed the wasp, then it is also true that the wasp was killed by John and if it is true that the wasp was killed by John, then it is also necessarily true that John killed the wasp; a parallel two-way entailment holds between the members of the other two pairs.

Contrariety. Contrary propositions may not be simultaneously true, although they may be simultaneously false. We can represent contrariety as ' $P \Rightarrow \sim Q$ '. The following are examples:

John killed the wasp	The wasp is alive
John killed the wasp	Mary killed the wasp
This paint is red	This paint is green

It cannot be simultaneously true of some wasp both that John killed it and that it is still alive; on the other hand, if the wasp is actually dead, but it was Bill who killed it, then both "John killed the wasp" and "The wasp is alive" are false. Notice that the entailment relation only goes from P to Q, not from Q to P. Thus, "This paint is red" entails "This paint is not green", but "This paint is not green" does not entail "This paint is red", since it might well be, for instance, yellow. Notice that this is not identical to the logical relation of disjunction $P \vee Q$, since this allows both P and Q to be true.

Contradiction. Contradictory propositions must have opposite truth values in every circumstance: that is, they cannot be either both true or both false. In any particular circumstance, one member of a contradictory pair must be true and the other false. Contradiction is bi-directional, and can be represented as ' $P \Leftrightarrow \sim Q$ '. The following sentences exemplify contradictory pairs:

The wasp is dead	The wasp is alive
John is still singing	John is no longer singing
No dogs are brown	At least some dogs are brown

If John is still singing, then it is false that he is no longer singing; to this extent, this is like contrariety. However, there is a crucial difference: if it is false that John is still singing, then it must be the case that he is no longer singing, and if it is false that he is no longer singing, then he must be still singing.

Compatibility. For some pairs of propositions, the truth values vary independently of one another: they may be both true, both false, or one true and the other false:

John is retired John is married
It is Tuesday today Christmas day falls on a Wednesday this year

This is equivalent to saying that there is no logical relation between the propositions. However, this is not without significance as a semantic relation.

The relations described in this section have an important role in the analysis of meaning relations between words (see Chapters 6–9).

Strictly speaking, the relations described above hold between propositions, not between sentences, because sentences do not have truth values. However, one frequently encounters in semantic texts statements to the effect that such-and-such a sentence entails (for example) some other sentence. This can be taken as a kind of shorthand for something slightly more complex. Saying that sentence S^1 entails sentence S^2 means that in any context where S^1 expresses a true proposition, S^2 also necessarily expresses a true proposition, provided that corresponding definite referring expressions in the two sentences are co-referential. Obviously, in the case of *It's a dog* and *It's an animal*, the two occurrences of *it* must refer to the same entity for the logical relation to hold, and in the case of *John killed the wasp* and *The wasp died*, we must be talking about the same wasp, and the time references must be the same. The use of the shorthand version is harmless as long as the distinction between propositions and sentences is kept in mind, and as long as it is obvious which is intended on any particular occasion. In what follows and in the rest of the book, propositions and sentences will generally not be typographically or otherwise distinguished, except where there is a danger of confusion.

2.2.2 Predicate logic: the inner structure of propositions

In this section, propositions are, as it were, opened up to reveal something of their inner structure. This perspective is known as **second-order logic**, and the system of representing the structure of propositions is known as the **predicate calculus**.

2.2.2.1 Arguments and predicates

A closely linked pair of concepts which are absolutely fundamental to both logic and semantics are argument and predicate. No attempt will be made here to explore the philosophical background and underpinning of these notions: the basic notions are fairly accessible and they will be treated in an elementary fashion. Put simply, an argument designates some entity or group of entities, while

a predicate attributes some property to the entity denoted by the argument, if there is only one, or a relation between the entities denoted by the arguments, if there is more than one. The following are some examples:

Proposition: John is tall
 Argument: John
 Predicate: is tall
 Proposition: John likes Mary
 Arguments: John, Mary
 Predicate: likes
 Proposition: John gave Mary a rose
 Arguments: John, a rose, Mary
 Predicate: gave

An argument may have a more or less complex internal structure. For instance, a whole proposition, itself possessing argument(s) and predicate, may constitute an argument, as in the following example:

Proposition: John was surprised that the man was tall
 Arguments: John, that P
 Predicate: was surprised
 Proposition P: the man was tall
 Argument: the man
 Predicate: was tall

There are various ways of incorporating propositions as constituents of complex arguments, and there is no limit to the resulting degree of complexity. An account of this is beyond the scope of this book.

An atomic proposition therefore consists of a predicate with one or more arguments. One conventional way of representing a proposition is as follows:

TALL(John)	(= "John is tall")
POLICEMAN(John)	(= "John is a policeman")
LIKE(John, Mary)	(= "John likes Mary")
AFRAID(John, Mary)	(= "John is afraid of Mary")
GIVE(John, rose, Mary)	(= "John gave Mary a rose")

Notice that *John gave Mary a rose* and *John gave a rose to Mary* express the same proposition. For the logical representation, the order in *John gave a rose to Mary* is preferred because it is more general: not all three-argument verbs allow the ditransitive construction (*John donated his library to the school*, **John donated the school his library*).

2.2.2.2 Valency/adicity

The number of arguments a predicate takes is known as its **valency** (sometimes **adicity**). In this book a distinction will be made between the logical valency of a predicate and the syntactic valency of a verb or adjective. The logical valency of a predicate is determined by the number of arguments for it to be logically complete, irrespective of the number of complements the verb or whatever takes in natural language. This is a difficult notion to define, but it is intuitively fairly easy to grasp. Take the case of *John bought a car*. Syntactically, there are only two participants, *John* and *a car*. But logically, an act of buying involves a buyer, a seller, the thing bought, and a price. If any of these is absent, then whatever happened was not a case of buying. In some cases, the language does not allow the specification of all the arguments. For instance, *The ring fetched £2,000* describes a situation which necessarily included a seller and a buyer, but there is no way to include these in the sentence, at least not in the sentence nucleus. Some cases are more tricky. Consider *Mary ate the spaghetti* and *Mary has eaten*. If it is true that Mary has eaten, then it follows that she ate something, and one might be tempted to say that the logical structure of *Mary has eaten* is EAT(Mary, something). However, it can plausibly be argued that the *eat* in *Mary ate the spaghetti* and the *eat* in *Mary has eaten*, although they are obviously related in meaning, are nonetheless different verbs. If Mary ate a cherry, we would not say that she had eaten *tout court*: the monovalent *eat* may be glossed “have a meal”.

2.2.2.3 Analytic, paradoxical, and synthetic propositions

Analytic propositions. These are propositions that are necessarily true by virtue of their internal make-up:

John's uncle is a man.

The proposition “John is tall” is either true or false.

Paradoxical propositions. These are propositions that are necessarily false by virtue of their internal make-up:

All bachelors are married.

John's sister is a man.

Synthetic propositions. These are propositions whose truth values are a matter not of their internal make-up, but of their correspondence or otherwise with the facts:

John's sister is married.

This paint is green.

All dogs are brown.

(Notice that the last example is factually false, but it is not logically false—the world could have been different.)

The terms analytic, paradoxical, and synthetic can be adapted to apply to sentences:

Analytic. Analytic sentences are sentences which automatically express true propositions in any context by virtue of the meanings of their constituent words and their arrangement.

Paradoxical. Paradoxical sentences automatically express false propositions.

Synthetic. Synthetic sentences are those which express true propositions in some (conceivable) contexts (although they may be false of the world as we know it) and false ones in others (this is the normal kind of sentence used in communication).

2.3 Quantifiers and quantification

Quantification is concerned with expressions like:

No Albanians came to the party.
Some of my best friends are troglodytes.
All aardvarks can sing the Marseillaise.

The subject NPs in the above are **quantified NPs**; the sentences express a quantification.

2.3.1 Quantifiers in classical predicate logic

A quantification requires a **quantifier**, (e.g. *no*, *some*, *many*, *all*, etc.), a **restriction** (which indicates the sort of things being quantified, e.g. *Albanians*), and a **scope**, which expresses what is true of the items designated by the quantified NP.

Classical predicate logic recognizes just two quantifiers, (i) the **existential quantifier** (usually symbolized as $\exists$), which in its quantificational properties corresponds to such English expressions as *somebody*, *a cat*, *some book*, etc. and (ii) the **universal quantifier** (symbolised as $\forall$), which corresponds to expressions like *all men*, *every aardvark*, *everybody*, and so on. Some idea of the nature of quantifiers can be gained by a closer examination of these.

Consider the English sentences below:

- (5) Everybody saw Mary.
- (6) Somebody saw Mary.
- (7) Mary saw somebody.
- (8) Mary saw everybody.

These would be translated into predicate calculus by means of formulae with roughly the structure of the following:

- i. For all (x), (x) saw Mary.
- ii. For some (x), (x) saw Mary.
- iii. For some (x), Mary saw (x).
- iv. For all (x), Mary saw (x).

Here, (x) is called a **variable**, because it does not have a fixed reference, and the quantifier is said to **bind** the variable. *Mary saw (x)* is called a **propositional function** and when (x) is given a referential value, it forms/expresses a proposition. For a sentence like *All aardvarks are left-pawed* a more complex representation is required:

For all (x) , if (x) is an aardvark then (x) is left-pawed

And for existential quantification, take *An aardvark sang*:

For some (x) , (x) is an aardvark and (x) sang

Now let us look at sentences with two quantifiers, like:

(9) Every aardvark saw a springbok.

This sentence is ambiguous: either all the aardvarks saw a particular springbok, or every aardvark had a springbok-viewing experience, but not necessarily of the same springbok. This ambiguity can be captured by placing one quantifier within the scope of the other in two different ways. Suppose we start with the existential quantifier in the outer position, with the universal quantifier in its scope. The resultant sentence has a 'formal' translation as follows:

There exists some (x) , such that (x) is a springbok and for all (y) , if (y) is an aardvark then (y) saw (x) .

Reversing the order of the quantifiers yields:

For all (y) , if (y) is an aardvark then there exists some (x) such that (x) is a springbok and (y) saw (x) .

Quantifiers interact in regular ways with negatives, and similar sorts of ambiguities can arise as with two quantifiers. Take the sentence *Alf the aardvark didn't see a springbok*. The most natural interpretation of this would be:

It is not the case that there existed a (y) such that (y) was a springbok and Alf saw (y) .

Here, the existential quantifier is within the scope of the negative operator (translated as *it is not the case that...*). But there is another possible interpretation for this type of structure, as in *John did not see a sniper, and was shot as he crossed the road*. Here the negative is inside the scope of the existential operator:

There existed an (x) such that (x) was a sniper and it is not the case that John saw (x) .

There is a similar interaction between a negative and universal quantifier. Consider the sentence *All the aardvarks did not see Pik*. The most natural interpretation of this is once again with the negative having widest scope:

It is not the case that for all (x) , if (x) is an aardvark, (x) saw Pik.

It is less natural to read this sentence with the quantifier having the widest scope:

For all (x) , if (x) is an aardvark then it is not the case that (x) saw Pik.

A more natural encoding of this meaning is *None of the aardvarks saw Pik*. If we have two quantifiers and a negative, as in *All the aardvarks did not see a springbok*, there are in theory six possible interpretations, although some of them are somewhat unnatural:

For all (x) if (x) is an aardvark, then it is not the case that there exists a (y) such that (y) is a springbok and (x) saw (y).

(No aardvark saw a springbok)

For all (x) if (x) is an aardvark, then there exists a (y) such that (y) is a springbok and it is not the case that (x) saw it.

(For every aardvark there was a springbok that it did not see)

It is not the case that for all (x) if (x) is an aardvark, then there exists a (y) such that (y) is a springbok and (x) saw (y).

(Not all the aardvarks saw any springbok)

It is not the case that there exists a (y) such that (y) is a springbok and for all (x) if (x) is an aardvark then (x) saw (y).

(No springbok was seen by all the aardvarks)

There exists a (y) such that (y) is a springbok and for all (x) if (x) is an aardvark then it is not the case that (x) saw (y).

(There is a springbok that none of the aardvarks saw)

There exists a (y) such that it is not the case that for all (x) if (x) is an aardvark then (x) saw (y).

(There is a springbok that was not seen by all the aardvarks)

When one quantifier is within the scope of another, the including quantifier is said to have **wider scope** (this applies not only to the 'classical' quantifiers just dealt with, but to *most*, *many*, *a few*, etc. It is possible to arrange quantifiers in order of their preferences for wide scope; this at least partially determines the preferred readings of propositions with more than one quantifier. One suggestion for the order of preference is as follows:

each > every > all > most > many > several > some > a few

The effect of these different degrees of inherent tendency to have wide scope can be seen in the following:

(I0) A springbok was seen by many aardvarks.

(II) A springbok was seen by each aardvark.

In the preferred interpretation of (I0), *many* is within the scope of *a*, and we take it that a single springbok is involved. In (II), however, the scopes are reversed, and we assume a plurality of springboks. This is a consequence of the fact that *each* has the stronger tendency to wide scope, strong enough to override the tendency of a subject to take wide scope; *many*, on the other hand, is overridden by *a* in subject position.

2.3.2 Generalized quantifiers

The trouble with the classical quantifiers of predicate logic is that firstly, there are quantifying expressions that intuitively belong together with *every* and *some/a*, but which cannot be expressed in the predicate calculus (for instance *many*, *few*, *more than half*, etc.) and secondly, there are many whose expression is clumsy and counter-intuitive. Ronnie Cann, a formal linguist, (Cann 1993) gives the following as a translation of *At least two students laughed*:

There exists an (x) and a (y) such that (x) is a student and (y) is a student and (x) is not the same as (y) and (x) laughed and (y) laughed

This can be adapted for any specific number, but at some cost in plausibility.

A more fruitful way of looking at quantifiers is to say that they express relations of quantity between sets of elements. These are relations which are not concerned with the identity of any of the elements in the sets that they relate, but only with their numbers. The relevant sets in a quantified sentence such as *Every aardvark sneezed* are (i) the set of things which satisfy the subject nominal, i.e. the set of aardvarks and (ii) the set of things which satisfy the predicate, i.e. the set of sneezers. One way of accounting for the relation between these sets which is expressed by the sentence is to say that the set of aardvarks is a subset of the set of sneezers. A more general way which allows a uniform treatment of a wider range of quantifiers is in terms of cardinality, i.e. the number of elements in a set, together with operators such as '=', '>', '<', etc. This gives the following interpretations:

- (I2) *Every X is Y*: "The number of elements in the set of things that are X but not in the set of things that are Y is zero"
- (I3) *Some X are Y*: "The number of elements in the set of things that are both X and Y is greater than zero"
- (I4) *Five X are Y*: "The number of elements in the set of things that are both X and Y is five"
- (I5) *Most X are Y*: "The number of elements in the set of things that are both X and Y is greater than the number of elements in the set of things that are X but not Y"
- (I6) *Neither X is Y*: "The number of elements in the set of things that are both X and Y is zero, and the number of elements in X is two"

Some of these analyses are straightforward, but some merit further discussion. Take the analysis of *most*: as it stands, it would also serve for *the majority of* and *more than half*. Is this an entirely satisfactory account of *most*? Cann suggests that for at least some speakers *most* requires a greater proportion than *more than half*: that is to say, if out of 100 aardvarks, 51 sneezed and 49 did not, this would not justify the use of *most*, but would satisfy the formula given above (and would justify *more than half*). Cann's solution is to say that the required proportion for *most* is pragmatically determined by reference to context (this should be taken to include the identity of X and Y); he incorporates a contextual

proportional factor c in the formula. His account of *most* can be expressed verbally as follows:

- (17) *Most X are Y*: "The number of elements in the set of things that are both X and Y is greater than the number of elements in the set of things that are X but not Y, and is greater than the number of elements in the set of things that are X multiplied by the contextual proportional factor c "

This allows us to set a proportion of things that are X—say, 80 per cent—that the number of things that are both X and Y must exceed to justify the use of *most*.

A contextual factor is also needed for *many*. Cann points out that the proportion which would justify the use of *many* need not be as much as *more than half*; for instance, *Many civil servants receive knighthoods* may mean no more than that the proportion of knighted civil servants is greater than that of other comparable professions, and may still be quite a small percentage. Cann's formula for *many* (modified) is:

- (18) *Many X are Y*: "The number of elements in the set of things that are both X and Y is greater than the number of elements in the set of things that are X multiplied by the contextual proportional factor c "

It should be noted, however, that the c which appears in (17) is not the same as that which appears in (18); they should therefore be distinguished as c_{most} and c_{many} .

2.3.3 Directional entailment

From sentence (19):

- (19) Every dog barked

we can validly infer *Every spaniel barked*, but not *Every animal barked*; we can also infer *Every animal made a noise*, but not *Every dog barked loudly*. Notice that the valid entailment goes from less specific to more specific for the subject term, but from more specific to less specific for the predicate term; more technically, the subject in (19) is **downward entailing**, and the predicate is **upward entailing**. The directional pattern of entailment is characteristic of the quantifier *every*, which creates a downward-entailing environment for its subject and an upward-entailing environment for its predicate. The full range of patterns is as follows:

(i) subject upward entailing; predicate upward entailing

- (20) If some dogs barked, then some animals barked.
 (21) If some animals barked, then some animals made a noise.
 (provided *some* is interpreted as "at least one")
 (22) *If some dogs barked, then some spaniels barked.
 (23) *If some dogs barked, then some dogs barked loudly.

(ii) subject and predicate downward entailing

- (24) If no dogs barked, then no spaniels barked.
- (25) If no dogs barked, then no dogs barked loudly.
- (26) *If no dogs barked, then no animals barked.
- (27) *If no dogs barked, then no dogs made a noise.

(iii) subject downward entailing, predicate upward entailing

- (28) If every dog barked, then every spaniel barked.
- (29) If every dog barked, then every dog made a noise.
- (30) *If every dog barked, then every animal barked.
- (31) *If every dog barked, then every dog barked loudly.

(iv) subject no entailment, predicate upward entailing

- (32) If most dogs barked, then most dogs made a noise.
(provided 'most' is interpreted not to exclude 'all')
- (33) *If most dogs barked, then most spaniels barked.
- (34) *If most dogs barked, then most animals barked.
- (35) *If most dogs barked, then most dogs barked loudly.

2.3.4 Negative polarity items

Directional entailment properties correlate in an interesting way with so-called **negative polarity items (negpol)**. These are expressions which are only normal in certain types of environment, typically containing a negative element of some kind. Typical examples are *anyone*, *anything*, *ever*:

- (36) He never says anything.
He rarely says anything.
I haven't seen anyone.
No one has ever reached the top.
Few people have ever reached the top.

Compare the normality of these with those in (37):

- (37) *A man has ever reached the top.
*All men have ever reached the top.
*Some men have ever reached the top.
*Most men have ever reached the top.
*Many men have ever reached the top.
*He always says anything.
*He sometimes says anything.
*He usually says anything.

The correlation with direction of entailment is that negpol are licensed in downward-entailing environments, but not in upward-entailing environments or environments where there is no entailment. This correlation is interesting, but it

does not really constitute an explanation, and there are other unexplained properties of negpols, such as the fact that some occur happily with questions and conditionals, while others do not :

- (38) I didn't say anything.
 Did you say anything?
 If you say anything, you're finished.
 I didn't say a word.
 If you say a word, you're finished.
 *Did you say a word?
 I couldn't care less.
 *Could you care less?
 *If I could care less, I would do something about it.

2.4 The logic of relations

Another useful set of concepts borrowed from logic are to do with relations between individual entities. The entities may be anything at all: objects, people, places, ideas, propositions. So can the relations be anything: "brother of", "smaller than", "has played string quartets with", "entails". The logical properties of such relations can be grouped under four headings: transitivity, symmetry, reflexivity, and converseness.

2.4.1 Transitivity

A relation that is **transitive** is one such that if A is related in this specific way to B and B to C, then it follows inescapably that A stands in the relation to C. Suppose A, B, and C are people, and the relation is "is taller than". Then if A is taller than B and B is taller than C, then A is necessarily taller than C. If a relation is **intransitive**, then if A stands in the relation to B and B to C, then it is logically impossible for A to stand in the relation to C. This is the case with "is the mother of": if A is the mother of B and B the mother of C, then A cannot be the mother of C. A relation may be neither transitive nor intransitive; we shall call such a relation **non-transitive**. If John has played duets with Bill, and Bill has played duets with Tom, then we are not in a position to conclude anything regarding John's musical relations with Tom.

2.4.2 Symmetry

A **symmetric** relation is one such that if A stands in a particular relation to B, then B necessarily stands in that same relation to A. For instance, if A is near to B, then B is near to A. If a relation is **asymmetric**, then if A stands in the relation to B, B cannot stand in the same relation to A. An example of an asymmetric relation is "is taller than". Once again, it is useful to have a designation for relations that are

neither symmetric nor asymmetric, such as “is sexually attracted by”; we shall label these **non-symmetric**.

2.4.3 Reflexivity

This property is not of great usefulness in semantic analysis: it is included for the sake of completeness. A relation is **reflexive** if something necessarily stands in that relation to itself. This includes most types of identity relation such as “has the same name as”, “is the same age as”, etc. Again, we can recognize **irreflexive** relations like “is taller than”, and **non-reflexive** relations such as “knows the weight of”.

2.4.4 Converseness

Converseness is a relation between relations. Two relations are **converses** if one yields the same proposition as the other when the arguments are reversed. By this criterion, “above” and “below” are converses (perhaps more strictly “is above” and “is below”), because *A is below B*, assuming constancy of A and B, expresses the same proposition as *B is above A*. Other examples of (sentences expressing) converse relations are: *A is B's offspring/B is A's parent*, *A saw B/B was seen by A*, *A sold B to C/C bought B from A* (it is usual to disregard ‘automatic’ adjustments in grammatical realization, such as the change from *to* to *from* in the case of *buy* and *sell*). We have defined converseness in terms of two-way entailment between two sentences (e.g. *A is taller than B* and *B is shorter than A*). It is useful to have a relation defined on a one-way entailment. For instance, *A is B's doctor* entails *B is A's patient*, but the reverse entailment does not hold because other practitioners in the medical and paramedical field, such as dentists and speech therapists, also have patients. We shall say that *doctor* is a **semi-converse** of *patient*.

2.5 The logic of classes

Another useful set of concepts is drawn from the logic of classes.

2.5.1 Identity

Two classes C^1 and C^2 are said to be **identical** if everything that belongs to C^1 also belongs to C^2 , and vice versa. Thus, the class of fathers and the class of male parents are identical, as are the class of pairs of spectacles and the class of pairs of glasses (on the relevant interpretation of *glasses*)

2.5.2 Inclusion

Class C^1 is said to include class C^2 if everything that is a member of C^2 is also a member of C^1 , but not vice versa. Thus, for instance, the class of animals includes the class of dogs, the class of aardvarks, etc. The set of dogs is described as a

subclass of the set of animals, and the set of animals as a superclass of the set of dogs. (Inclusion is defined here so as to exclude identity; it can be defined so as to include identity.)

2.5.3 Disjunction

Classes C^1 and C^2 are said to be disjunct if no member of C^1 is also a member of C^2 . The class of cats and the class of aardvarks are disjunct in this sense, as are the class of red things and the class of green things.

2.5.4 Intersection

Classes C^1 and C^2 are said to intersect if they have some members in common, but each has members which do not belong to the other (i.e. complete intersection, or identity, is excluded here). The class of red things and the class of round things intersect in this fashion, as do the class of architects and the class of amateur musicians. The set of common members to two (or more) overlapping classes is often referred to as the intersection of the two (or more) classes.

2.5.5 Union

The combined set of members belonging to either of two (or more) classes (including overlapping classes) is called the union of the two (or more) classes. Thus the union of the class of dogs and the class of cats is constituted by the class of all entities which are either cats or dogs.

2.5.6 Class relations and propositional relations

There are obvious connections between the class relations described above and the propositional relations described earlier. For instance, the fact that *It's a dog* entails *It's an animal* is not unconnected to the fact that the class of dogs is a subclass of the class of animals, and scarlet things being a subclass of red things relates to the entailment from *It's scarlet* to *It's red*. Some relations emerge more naturally within one approach than another. For instance, we gave no propositional parallel for class overlap. This could be called propositional independence, since the truth of either proposition in such a pair imposes no logical restraint on the truth value of the other, as in the case of *John is an architect* and *John is an amateur musician*. The relation of contradiction, on the other hand, emerges more naturally within the propositional approach. To find an equivalent of the relation between, say, *dead* and *alive* within the class approach, we would have to say that the class of dead things and the class of alive things were (i) disjunct and (ii) exhaustive of a superclass of animate things.

2.5.7 Mapping

It sometimes happens that the members of one class have a relation of correspondence of some kind with one or more members of a parallel class. This type of correspondence is known as **mapping**. An example will make this clear. It is a well-known fact that a person's fingerprints are uniquely distinctive. If, therefore, we think of the class of persons and the class of fingerprints, there is a straightforward mapping relation between the two classes, in that each member of one class corresponds to a specific member of the other set. This is known as **one-to-one mapping**. Contrast this situation with the two classes FATHERS and CHILDREN. Every member of the FATHER class corresponds to one or more members of the CHILDREN class, but every member of the CHILDREN class corresponds to a single specific member of the FATHER class. Here we have **one-to-many** mapping between fathers and children, but **many-to-one** mapping between children and fathers. Yet another elementary mapping relation holds between the class of word forms and the class of meanings. If we allow the possibility of synonymy, then some word forms (e.g. perhaps *begin* and *commence*) will map onto the same meaning, whereas other word forms (e.g. *bank*) will map onto more than one meaning. This is known as **many-to-many** mapping (see Chapter 5 for a more detailed consideration of mapping between words and meanings).

2.6 Presuppositions

If John says to Mary *I'm glad you got the job*, he is not normally telling her that she got the job. It is likely that Mary would already be in possession of that information. The function of the clause *you got the job* is to specify what it is that the speaker is glad about. The truth of the expressed proposition is assumed to be common knowledge shared by speaker and addressee. This is a typical example of presupposition: the speaker presupposes that the hearer got the job and expresses their pleasure at the fact.

It might seem at first sight that we have here a case of entailment: if it is true that X is glad that Y got the job, then it is true that Y got the job. However, there are crucial differences between presupposition and entailment. Consider the example immortalized by Bertrand Russell:

(39) The present king of France is bald.

This presupposes (40):

(40) There is a king of France.

There are two things to note about (39). First, intuitively, it in some way implies (40). But so does (41):

(41) The present king of France is not bald.

In other words, the presupposition is not affected by ordinary sentence negation. (However, this is only in the 'default' interpretation.) We can also interpret (41) as negating, or cancelling the presupposition, especially if it is appropriately stressed:

The present king of France is *not* bald because there *is* no present king of France.

Second, if France currently had a king, then (39) would be either true or false (ignoring the question of how much hair one has to lose to qualify as bald). But as it happens, there is no current king of France (although there are pretenders to the French throne). What, then, is the truth value of (39)? It cannot be true, but on the other hand it does not seem right to describe it as false, either. We can say that a consequence of the falsity of (40) is that (39) and (41) lack a truth value. These two peculiarities serve to distinguish presupposition from entailment. Take a straightforward case of entailment. Sentence (42) entails sentence (43):

(42) Ashley is a man.

(43) Ashley is male.

However, the two characteristics of presupposition are absent here. First, sentence (44), the negation of (42), does not imply (43):

(44) Ashley is not a man.

Second, suppose that (43) is false. The consequence of this is not that (42) and (44) lack truth values: on the contrary, (42) is false and (44) is true.

Expressions that carry presuppositions are known as presupposition triggers or inducers. There are many different types of these. The following is a selection. In each case, the 'a' sentence is a presupposition trigger, and the 'b' sentence one of its presuppositions (a sentence may have several presuppositions):

(45) a. I want to meet the winner of yesterday's egg-and-spoon race.

b. There was an egg-and-spoon race yesterday.

(46) a. I was amazed that he turned up.

b. He turned up.

(47) a. He managed to escape.

b. He acted with the intention of escaping.

(48) a. It wasn't Mary who stole the caviar.

b. Somebody stole the caviar.

(49) a. If only we had listened to him, he would still be alive.

b. We did not listen to him.

(50) a. John drank it.

b. It was a liquid.

In this section we shall (briefly) look at ways of accounting for presuppositions in terms of logic. Two logical accounts will be examined.

2.6.1 Russell

Russell believed that the presupposition of a sentence should be explicitly represented in a representation of the sentence's meaning. His suggestion for the logical form of (39) is as follows:

(51) The king of France is bald = $\exists x (\text{king}(x) \ \& \ \sim \exists y ((y \neq x) \ \& \ \text{king}(y)) \ \& \ \text{bald}(x))$

This translates roughly as "There exists an entity x such that x is the king of France and there does not exist an entity y such that y is the king of France and x is not y [this is Russell's way of capturing the meaning of the definite article] and x is bald".

Russell's account has the virtue of being explicit. It can also explain the dual interpretation of (41) in terms of the scope of the negative operator. The operator can have wide scope including the whole compound proposition, thus denying the truth of the proposition:

(52) The king of France is not bald = $\sim(\exists x (\text{king}(x) \ \& \ \sim \exists y ((y \neq x) \ \& \ \text{king}(y)) \ \& \ \text{bald}(x)))$

This translates as "It is not the case that there exists an entity x such that x is the king of France and there does not exist an entity y such that y is the king of France and x is not y and x is bald". A compound proposition is false if any of its component simple propositions is false, so this formulation allows an interpretation in which the king of France does not exist. The default interpretation of (41) has the following logical structure:

(53) The king of France is not bald = $\exists x (\text{king}(x) \ \& \ \sim \exists y ((y \neq x) \ \& \ \text{king}(y)) \ \& \ \sim \text{bald}(x))$

This translates as "there exists an entity x such that x is the king of France and there does not exist an entity y such that y is the king of France and x is not y and x is not bald".

There are problems with Russell's formulation. First, it gives the presupposition and the attribution of baldness exactly the same status, even though intuitively, the sentence does not actually state that there is a king of France, but does state that he is bald. Second, the falsity of the presupposition entails the falsity of the compound proposition, that is to say, if (40) is false, then both (39) and (41) are false.

2.6.2 Strawson

Strawson took up a suggestion by Frege and made it more explicit. In Strawson's version, presupposition is a relation between propositions:

(54) A proposition P is a presupposition of another proposition Q if and only if:
 If P is true then Q is true
 If P is false then Q is true

Notice that it follows from this that if *Q* is false, then *P* can be neither true nor false, that is to say, the truth of *Q* is a precondition for *P* having a truth value. This succeeds in capturing the constancy of presupposition under negation and the absence of a truth value when the presupposition is false. However, unlike Russell's formulation, it does not account for the interpretation of sentences like (41) when the presupposition is cancelled. It also fails to account for the intuition that sentences like (55) are false, rather than lacking truth value:

(55) The king of France visited me yesterday.

There is an extensive literature on presupposition and the various approaches to it, which it is not possible to go into here. For more details, see the sources recommended in the suggestions for further reading at the end of this chapter.

2.7 Possible world semantics

This is an approach to the semantic interpretation of logical formulae in certain systems of formal semantics, in particular, those like the so-called **Montague semantics**, that are based on **intensional logic**.

The basic idea is that the actual world (in the broadest sense of everything existing) is only one of an infinite number of conceivable alternative worlds which differ in at least one respect from the actual world. Some possible worlds are very close to the actual world. For instance, there is a possible world identical to the actual world except that I did not mistype the current sentence. Others differ in major respects: for instance, there is one in which Napoleon was victorious at Waterloo. Others are hugely different, where, for instance, our galaxy does not exist. The following examples give a very brief idea of the sort of uses made of the concept of possible worlds (for a fuller understanding, more advanced works need to be consulted).

Extension:

The extension of a noun like *dog* is the set of all dogs in the actual world.

The extension of a sentence like *Dogs are animals* is its truth value in the actual world.

Intension:

The intension of *dog* is the set of all dogs in all possible worlds.

The intension of a sentence is its truth condition(s), interpreted as the set of all possible worlds in which it is true.

Analytic, paradoxical, and synthetic propositions:

An analytic proposition is one which is true in all possible worlds.

A paradoxical proposition is one which is false in all possible worlds.

A synthetic proposition is one which is true in at least one possible world.

Entailment:

P entails Q is true if in all worlds in which *P* is true, *Q* is also true.

Rigid designator:

A rigid designator is a term which refers to the same individual in all worlds in which that individual occurs. For instance, the term *eleven* designates the same number in all worlds where it has the same meaning as in the actual world. This may seem a tautology, but contrast *eleven* with *the number of gold medals won by Britain at the 2004 Olympics*, which also designates a number, but clearly not the same number in all possible worlds. Proper names in general are held to be rigid designators, but **definite descriptions** are in principle capable of designating different individuals in different worlds, without change of meaning.

2.8 Some useful distinctions

2.8.1 Reference, denotation, and sense

Language is used to communicate about things, happenings, and states of affairs in the world, and one way of approaching the study of meaning is to attempt to correlate expressions in language with aspects of the world. This is known as the **extensional** approach to meaning.

The thing or things in the world referred to by a particular expression is its **referent(s)**: in saying *The cat's hungry*, I am (normally) referring to a particular cat, and that cat is the referent of the expression *the cat*. The whole utterance attributes a particular state to the cat in question. We can also consider the whole class of potential referents of the word *cat*, namely, the class of cats. This, too, is sometimes called the **reference** of the word *cat*. But this is clearly different from the designation of particular individuals as in the case of *The cat's hungry*, so, to avoid confusion, we shall follow Lyons and say that the class of cats constitutes the **denotation** of the word *cat*. So, in the case of *The cat's hungry*, the word *cat* **denotes** the class of cats, but *the cat* **refers** to a particular cat.

The alternative to an extensional approach to meaning is an **intensional** approach. Take the word *cat*. Why do we use it to refer to cats, rather than, say, to platypuses or aardvarks or spiny anteaters? One answer is that the word is associated with some kind of mental representation of the type of thing that it can be used to refer to, and aardvarks do not fit the description associated with the word *cat*. This representation constitutes what is called the **sense** of the word (or at least part of it). We shall assume in this book that the main function of linguistic expressions is to mobilize concepts, that concepts are the main constituents of sense, and that sense (and hence concepts) constrains (even if it does not completely determine) reference. (It should be noted that some authors, for instance Lyons, understand *sense* in a different way. For them, sense is a matter of the relations between a word and other words in a language. So, for instance, the sense of *cat*

would be constituted by its relations with other words such as *dog* (a cat is necessarily not a dog), *animal* (a cat is an animal), *miaow* (*The cat miaowed* is normal but *?The dog miaowed* is not.)

2.8.2 Use and mention

Consider the difference between (56):

- (56) Snow has four letters
 Snow is a noun
 Snow is a natural kind term
 Snow is an English word
 Snow is easy to pronounce

and (57):

- (57) Snow is white
 Snow damages crops
 Snow is frozen water

The difference between these two sets is usually designated as a difference between **use** (here, of the word *snow*), as in (57), and **mention** (of the word *snow*), as in (56). In the sentences in (56) we are using the word form *snow* to identify a word of the language, and we then proceed to say something about that word; in the sentences in (57) we are using the word form to identify a substance in the world, prior to predicating something of it. A simple way of distinguishing the two is to apostrophize (or italicize) the language unit in question. If this makes a negligible effect on the meaning, then it is a case of mention:

- 'Snow' has four letters
 'Snow' is an English word
 *'Snow' is white
 *'Snow' damages crops

Mention may involve any stretch of language:

- 'Go to home' is ungrammatical
 *Go to home is ungrammatical

Discussion questions and exercises

1. Arguments and predicates

Considering the following as logical predicates, mark them as one-, two-, three-, or four-place (1, 2, 3, 4) (think in terms of the intuitive maximum number of arguments):

yawn, steal, thank, pay, be tall, be taller than, meet, put, imagine, daydream, cost, understand, explain.

2. Sentence, statement, utterance, and proposition

In the light of the definitions given in this chapter, of which of the above can the following be said? Assume we are talking about declaratives, and that statements and utterances inherit the properties of entities they embody (for instance, a statement will be in the form of a sentence and will inherit its properties):

- X was inaudible.
- X was uninformative.
- X was false.
- X was in a foreign accent.
- X was ungrammatical.
- X was insincere.

3. Propositional and non-propositional meaning

In which of the following pairs of sentences do the members, on their most likely construal, have the same propositional content? Look for situations in which they would make different statements, ask different questions, or issue different commands:

- (a) (i) Take your hands off me (said by a woman to a man)!
- (ii) Take your filthy paws off me (ditto)!
- (b) (i) I always get my bread from Gregg's, because it's cheaper.
- (ii) I always buy my bread from Gregg's, because it's cheaper.
- (c) (i) Don't you find him rather skinny?
- (ii) Don't you find him rather thin?
- (d) (i) Have you read the stuff he wrote about telepathy?
- (ii) Have you read the garbage he wrote about telepathy?
- (e) (i) She was there at the start of the race.
- (ii) She was there at the beginning of the race.
- (f) (i) John hasn't turned up.
- (ii) John hasn't turned up yet.
- (g) (i) Old Joshua Hobbblethwaite died last week.
- (ii) Old Joshua Hobbblethwaite passed away last week.

4. Entailments

In which of the following does the (a) sentence entail the (b) sentence? Are there any problems?

- | | |
|---------------------------|--------------------------------|
| (a) (i) X is a cat. | (g) (i) X is a pet. |
| (ii) X has four legs. | (ii) X is alive. |
| (b) (i) X is a cat. | (h) (i) X is not dead. |
| (ii) X is an animal. | (ii) X is alive. |
| (c) (i) X is a cat. | (i) (i) X has stopped smoking. |
| (ii) X is a quadruped. | (ii) X doesn't smoke any more. |
| (d) (i) X is a quadruped. | (j) (i) X taught Y Z. |
| (ii) X has four legs. | (ii) Y learnt Z. |
| (e) (i) X is a quadruped. | (k) (i) X killed Y. |
| (ii) X is an animal. | (ii) Y is not alive. |
| (f) (i) X is a pet. | (l) (i) X watched Y. |
| (ii) X is an animal. | (ii) Y was doing something. |

5. Relations between propositions

Mark the propositional relationship between the members of the following pairs of sentences as either EQUIVALENCE, CONTRARIETY, CONTRADICTION, or CONVERSENESS:

- (a) (i) Proposition P is true.
(ii) Proposition P is false.
- (b) (i) John likes Mary.
(ii) John dislikes Mary.
- (c) (i) Mary agrees with the statement.
(ii) Mary disagrees with the statement.
- (d) (i) Mary borrowed the book from John.
(ii) John lent the book to Mary.
- (e) (i) John killed the wasp.
(ii) The wasp is still alive.
- (f) (i) John is not married.
(ii) John is a bachelor.

6. Logic of relations

Classify the following relations with regard to their TRANSITIVITY (i.e. either TRANSITIVE, INTRANSITIVE, or NON-TRANSITIVE) and their SYMMETRY (i.e. as SYMMETRIC, ASYMMETRIC, or NON-SYMMETRIC):

parent of; ancestor of; brother of; related to; sibling of; friend of; near to; to the right of; far from; resembles.

7. What are the presuppositions of the following?

- (a) Lesley is a lesbian.
- (b) Lesley plays the clarinet brilliantly.
- (c) Lesley will graduate next year.
- (d) Lesley is sorry for all the trouble she has caused.

- (e) It was Lesley who wrote the letter.
- (f) When Lesley was ill, Jane deputized for her on the committee.

8. Which of the following are implicitly negative? Examine their collocability with negpols.

hardly	often	seldom	occasionally
mostly	a few	far (from)	near
free (from)	beware of	take care to	avoid

Suggestions for further reading

The treatment here has been very informal. A similar elementary treatment, but with more practical exercises will be found in Hurford and Heasley (1983). Lyons (1995) develops the philosophical background more fully, but still at an elementary level. Those requiring initiation into logical formalization will find an accessible introduction in Allwood, Anderson, and Dahl (1977). Lyons (1977) gives a more detailed treatment of many of the topics touched on here. Cann (1993), McCawley (1981), and Larson and Segal (1995) are only for those who are really serious about the application of logic to language. Presupposition is dealt with briefly in this chapter; Cruse (1992d) gives a fuller, but still introductory survey of different theoretical approaches; a much more detailed account can be found in Levinson (1983).


CHAPTER 3

Concepts and meaning

3.1 The importance of concepts

53

3.2 Theoretical approaches to concepts 54



CHAPTER 3

Concepts and meaning

3.1 The importance of concepts

A concept is a mental construct that stands in a relation of correspondence to a coherent category of things in some world, prototypically the real world, but potentially also imaginary, fictional, or virtual worlds. The notion 'conceptual category' encompasses these two aspects, the mental construct and the category of entities. The view taken in this book is that all meaning, whether conventionally associated with a linguistic expression or expressions, or whether it arises by pragmatic construal, is conceptual in nature.

Concepts are absolutely vital to the efficient functioning of human activities and human society.

Codification of experience. The essence of a concept is that it enables to treat a range of objectively different experiences as being instances (or 'tokens') of one and the same type of experience. Most of us have a wide range of experiences with a number of different cats, with different colours and markings, different sizes, different characteristic behaviours in different circumstances, and so on. But we also have a concept CAT, which enables us to see these as all belonging to a single type. This is extremely important for intellectual and social activity.

Learning. Without conceptual categories, learning would be piecemeal, chaotic, confined to acquiring bodies of knowledge about vast numbers of individual entities. Concepts provide a locus for accumulating information relevant for a class of individuals or events. This gives an efficient way of learning from past experience, since, as is well known, history never repeats itself exactly. Also, new knowledge gained on the basis of interaction with one or more individuals can be easily generalized to other members of a category. Conversely, recognizing that an individual belongs to a particular category, perhaps on the basis of just a few criteria, can make available a much wider range of information about that individual.

Communication. Conceptual categories are essential for linguistic communication. As we saw in the opening chapter, language does not operate in terms of

particularities, that is, of individual people, things, events, and so on. A linguistic message is coded in terms of conceptual categories, and individual referents intended by a speaker must be inferred from these.

3.2 Theoretical approaches to concepts

3.2.1 The classical approach

The classical approach to categorization, which goes back at least to Aristotle, but is still often taken for granted, defines a category in terms of a set of **necessary and sufficient criteria** (or conditions, or features) for membership. So, for instance, the criteria for some X to qualify for inclusion in the category GIRL are:

X is human
X is female
X is young

If any of these criteria are not satisfied, then X is not a girl (i.e. the criteria are **individually necessary**); if all the criteria are satisfied, then X is a girl (i.e. the criteria are **jointly sufficient**). (The above set of criteria can be taken as a definition of the meaning of *girl*.)

3.2.1.1 Some problems of the classical approach

There is a certain undeniable obviousness about this way of defining categories. However, it has a number of shortcomings.

Lack of plausible analyses. The superficial plausibility of the Aristotelian analysis of *girl* (and similar words) is misleading. The words like *girl*, which apparently can be satisfactorily defined by means of a set of necessary and sufficient features constitute a relatively small proportion of the vocabulary at large, and are confined to certain semantic areas, such as kinship, and specialized terms for animals specifying age and sex, and so on. There are many everyday words whose meanings cannot be captured by means of a set of necessary and sufficient features. Wittgenstein's famous example is *game*. He argued that it was impossible to draw up a list of features possessed by all games which jointly distinguish games from non-games. One might suggest the following as possible criteria:

- (i) involves winning and losing
- (ii) involves more than one person
- (iii) has arbitrary rules
- (iv) done purely for enjoyment

However, none of these is necessary:

- (i) There are many games which do not involve winning and losing: party games, such as charades, Matthew, Mark, Luke, and John, kissing games; children's games such as leapfrog, hallalevo, and hopscotch, etc.
- (ii) Solitaire is a game for one person.
- (iii) Many children's games, such as dressing-up games, and cowboys and Indians, have no statable rules.
- (iv) Many games are played professionally.

In spite of this lack of necessary and sufficient criteria, we communicate using the word *game* perfectly successfully, and without any sense of linguistic imperfection. Such examples can be multiplied almost indefinitely: *apple, dog, table, water, house, flower, dance, violin*, etc.

Fuzzy boundaries. An Aristotelian definition of a category implies a sharp, fixed boundary. However, much empirical research on category structure has been interpreted as showing that the boundaries of natural categories are fuzzy and contextually flexible. For instance, the American scholars Brent Berlin and Paul Kay (1969), who studied colour categories from a psycholinguistic and anthropological linguistic perspective, found that while judgements of certain examples of colours were relatively constant across subjects and reliable within subjects on different occasions, judgements of borderline instances, for instance between red and orange, or between blue and purple, showed neither agreement amongst subjects nor reliability with a single subject on different occasions. William Labov, an American sociolinguist, (1973) studied subjects' naming of line drawings illustrating cups, mugs, vases, bowls, and the like which systematically varied parameters such as ratio of height to width, curved or straight sides, presence or absence of a handle. Again, the finding was that certain items received reliable assignation to a particular category, while others were uncertain. He also found that contextual conditions could alter subjects' responses, so that, for instance, an instruction to imagine all the items as containing rice extended the boundaries of the BOWL category, while a similar instruction to imagine coffee as contents extended the CUP category. Such results receive no natural explanation within the classical (Aristotelian) picture. It is possible that the fuzziness of natural categories has been overstated—see the section below on the dynamic construal approach to categories.

Internal structure of categories. As far as the classical conception of categories goes, everything that satisfies the criteria has the same status, that is to say, something is either in the category, or not in it, and that is all there is to say about the matter. However, language users have clear intuitions about differences of status of items within a category: some members are felt to be 'better' examples of the category than others. For instance, an apple is a better example of a fruit than is a date, or an olive. In other words, categories have internal structure: there are central members, less central members, and borderline cases. No account of these facts can be given using the classical approach.

3.2.2 Concepts and kinds

Aristotle distinguished two kinds of substance: primary substances, that is to say, individual entities, like Madonna and Barack Obama's favourite tie, and secondary substances, like butter, human beings, computers, and aardvarks. The focus of interest in what follows will be on secondary substances, because these have the most direct relevance to language, especially words. I shall follow the philosopher Ruth Millikan and characterize substances (of both sorts) as entities in the world about which a coherent body of knowledge can be assembled as a result of repeated acquaintance. Secondary substances, which I shall henceforward refer to as 'kinds', are not identical to categories. Take the category of persons whose surname begins with the letter Y. This does not constitute a kind, because the definition exhausts what is to be known about the category—there is no relevant body of knowledge to be assembled. Compare this with, for instance, DOG, POLITICIAN, ELECTRON, which are true kinds in the relevant sense.

It is hardly necessary to point out the advantages of acquiring a store of knowledge about a range of substances, but in order to do this we have to have at least two things, (i) some sort of mental location for each substance, and a way of linking new knowledge to it and (b) a way of identifying recurrences of substances. Millikan identifies one fundamental function of a concept as that of identifying recurrences of a kind. I shall call this aspect of the function of a concept as a 'kind-detector'. An important feature of a kind-detector is that it has to carry out its identifying function in a wide range of contextual conditions, and will typically have to rely on a variety of interacting perceptual and other clues. It follows from this that a kind-detector will not in general be characterizable in terms of a set of necessary and sufficient features, or even a prototype representation. Notice that although words are in a sense public property, kind-detectors are personal property: there is no logical requirement for my kind-detector DOG to be the same as anyone else's. What is necessary for successful communication is that it identifies the same kind in the world on a sufficient number of occasions.

The kind-detector has logical properties. It is not error-free—it may identify correctly or incorrectly on any particular occasion. It is responsible for the truth values of utterances in which it is involved. Although a kind-detector is relatively stable, it is not unchangeable, and is designed to improve with experience. Besides being responsible for truth values, a kind-detector governs important logical relations between a concept and other concepts, the most important being inclusion and exclusion, which underlie sense relations such as hyponymy and incompatibility (see Chapters 6 and 7).

The purpose of identifying a kind is twofold: to give access to a stored body of knowledge, and to modify that body of knowledge by addition of new material, or by strengthening, weakening, or deleting existing material. This body of knowledge is called by Millikan the 'conception', and I shall retain her term. (What is

here called a concept is a combination of a kind-detector and a conception.) It is of course vitally important for reasoning, planning of action, and so on. Identifying a kind gives immediate access to this body of knowledge; conversely, features of the identified kind may be added to the general store, and thus be made available to new identifications.

Like kind-detectors, people's conceptions are not public, but are private property. But while different persons' kind-detectors of, for example, DOG are constrained by the need to correctly identify the 'public' substance DOG, conceptions are under no such constraints, or at least the constraints are somewhat weaker. Although no two persons' conceptions of a particular substance will be identical, a smooth flow of communication does of course depend on a significant degree of overlap.

It is doubtful whether conceptions as such, as opposed to kind-detectors, have logical properties; they are nonetheless deeply implicated in language processing, reasoning, and so on. Many non-logical effects, such as prototype effects, basic-level status, kind-of relations, and the emergent meaning arising in a conceptual blend, are all conception-driven, rather than kind-detector-driven.

3.2.3 Prototype theory

We shall first of all describe what might be called the 'standard' approach to prototype theory, deriving from the work of Eleanor Rosch and her co-workers. The main thrust of Rosch's work has been to argue that natural conceptual categories are structured around the 'best' examples, or **prototypes** of the categories, and that other items are assimilated to a category according to whether they sufficiently resemble the prototype or not.

3.2.3.1 Degrees of centrality

Rosch's most basic experimental technique is the elicitation of subjects' **Goodness-of-Exemplar ratings** (henceforward GOE ratings). Subjects are asked to give a numerical value to their estimate of how good an example something is of a given category. The rating scale typically goes something like this:

- 1 very good example
- 2 good example
- 3 fairly good example
- 4 moderately good example
- 5 fairly poor example
- 6 bad example
- 7 very bad example/not an example at all

So, for instance, if the category was VEGETABLE, the ratings of various items might be as follows:

POTATO, CARROT	1
TURNIP, CABBAGE	2
CELERY, BEETROOT	3
AUBERGINE, COURGETTE	4
PARSLEY, BASIL	5
RHUBARB	6
LEMON	7

Significantly, subjects do not find this to be a totally meaningless task. While there is of course a great deal of variation between subjects, statistically, the results within a culturally and linguistically homogeneous population cluster strongly round particular values. The prototypes of categories are determined by selecting the item with the lowest average numerical score.

GOE ratings may be strongly culture-dependent. (Familiarity is undoubtedly a factor influencing GOE scores, but the scores cannot be reduced to familiarity.) For instance, in a British context (say, a typical class of undergraduates), DATE typically receives a GOE score of 3–5 relative to the category of FRUIT, but an audience of Jordanians accorded it an almost unanimous 1.

Wittgenstein described the instances of the category GAME as manifesting a relationship of **family resemblance**: the members of a human family typically resemble one another, but there may well not be any set of features that they all possess, and it may be possible to find two members who have no features in common. However, they will be linked by a chain of intermediate members with whom they do share features. So, for example, A may have no features in common with C, but has the same nose as B, who in turn has the same eyes as C. Prototype theory embraces Wittgenstein's notion that family resemblance unites the members of a category, but adds to it the vital idea of central and peripheral members.

3.2.3.2 Prototype effects

Taken in isolation, the existence of stable GOE scores might be thought to be of minor cognitive significance. However, there is abundant evidence that prototypicality, as measured by GOE scores, correlates strongly with important aspects of cognitive behaviour. Such correlations are usually referred to as **prototype effects**. The principal prototype effects are as follows:

Order of mention. When subjects are asked to list the members of a category, and especially if they are put under time pressure, the order of listing correlates with GOE ratings, with the prototypical member showing a strong tendency to appear early in the list.

Overall frequency. The overall frequency of mention in such lists also correlates with GOE score.

Order of acquisition. Prototypical members of categories tend to be acquired first, and order of acquisition correlates with GOE rating.

Vocabulary learning. Children at later stages of language acquisition, when vocabulary enlargement can be greatly influenced by explicit teaching, learn new words more readily if they are provided with definitions that focus on prototypical instantiations than if they are given an abstract definition that more accurately reflects the total range of the word's meaning.

Speed of verification. In psycholinguistic experiments, subjects who are required to respond as quickly as they can to a categorization task produce faster responses if the tasks involve a prototypical member. In a typical set-up, subjects see a pair of words, say FRUIT: BANANA, flashed up on a screen, and they are to respond as quickly as possible by pressing one of two buttons, the one labelled *Yes* if the second named item belongs to the category indicated by the first item and *No* otherwise. Results show that responses to, for instance, FRUIT: APPLE, where the second item is a prototypical member of the class denoted by the first, are faster than, say, FRUIT: DATE (for average British subjects).

Priming. Another psycholinguistic technique involves the phenomenon of priming. In a typical set-up, subjects see strings of letters flashed onto a screen and their task is to respond *Yes* (by pressing the appropriate button) if the string of letters makes a word of (say) English, and *No* if it does not. Responses are timed electronically. It is a well-established experimental fact that if a word is preceded by a semantically related word, response to it will be speeded up. So, for instance, a *Yes* response to DOCTOR will be faster if NURSE has been just previously presented. It is found that the presentation of a category name has the greatest speeding-up effect on the prototype of a category, and the effect is proportionately less as we move away from the centre of the category to the periphery (as measured by GOE scores).

3.2.3.3 Intuitive unity, definitional polyvalence

Most of the work on prototypes has been carried out by psychologists, and the nature of the experiments reflects this. A purely linguistic characterization of categories with a prototypic organization (it is not necessary to assume that **all** categories have this sort of structure) is that they show intuitive unity, but are definitionally polyvalent. That is to say, they cannot be captured by means of a single definition, but require a set of definitions. For instance, the semantic field covered by the term *game* can be quite well described by means of a restricted set of definitions, but no satisfactory unitary definition exists.

3.2.3.4 Fuzzy boundaries

A common position is to maintain that only the prototype has 100 per cent membership of a category, the degree of membership of other items being dependent on their degree of resemblance to the prototype, this, in turn, being reflected by their GOE score (it has sometimes been claimed—wrongly, in my opinion—that when subjects give GOE ratings, they are actually judging degree of membership). From this one would have to conclude that a natural category has

no real boundaries, and indeed this has been explicitly claimed by, for instance, Langacker, a leading figure in the field of cognitive linguistics (1991b: 266):

There is no fixed limit on how far something can depart from the prototype and still be assimilated to the class, if the categorizer is perceptive or clever enough to find some point of resemblance to typical instances.

Not all scholars belonging to the cognitive linguistics fraternity agree that GOE and DOM (degree of membership) should be equated. However, there is general agreement that category boundaries are typically fuzzy. (Arguments against equating GOE and DOM will be detailed below.)

3.2.3.5 The mental representation of conceptual categories

The earliest hypotheses regarding the mental representation of categories suggested that there was some sort of portrait of the prototypical member, against which the similarity of other items could be computed and their status in the category determined. This idea fell out of favour when it was realized that many 'portraits' would have to be three-dimensional and would have to incorporate characteristic behaviour (although Jackendoff still envisages all these possibilities for his 3-D representation of conceptual categories). Many prototype theorists (following the American cognitive linguist George Lakoff) speak only of 'prototype effects', and remain uncommitted on the subject of the form of mental representations.

More recently, feature-based treatments of prototype structure have appeared. With these, categories with a prototype structure are represented by a set of features. However, unlike the classical features, these do not constitute a set of necessary and sufficient criteria, except perhaps for the prototype itself. Rather, the features are such that the more of them that are manifested in some particular instantiation, the higher the GOE score the item in question will obtain (note that in GOE terms, a score of 1 is high and 7 low). In such systems, features may be differentially weighted, that is to say, some features will have a greater effect on determining centrality in the category than others (there is nothing in principle to prevent some features being necessary). It is possible that no actual member of a category has all the features in the prototype specification. In that case, the core of the category may be represented by a group of the highest scoring members.

The general idea can be illustrated using the category *VEHICLE*. The features listed in (1) would seem to be plausible (note that these have not been subjected to empirical testing, they are based on my intuitions: the list is illustrative, not necessarily exhaustive):

- (1) a. Designed to go on roads
- b. Has own propulsive power
- c. Can go faster than unaided human
- d. Can carry persons/goods in addition to driver
- e. Has 4 wheels
- f. Metallic construction
- g. Persons/goods enclosed
- h. Manoeuvrable

Clearly a central example of the category of vehicle, such as CAR, will have all these features. If they are correct, it ought to be possible, for items judged not to be central, to pinpoint features they do not possess. For instance, a typical class of students will mark the following items as non-prototypical in the class of VEHICLE. For each of them, there are features from the above list which are missing:

TRAIN	not designed to go on roads not manoeuvrable
TRACTOR	not designed to go on roads driver not always covered
BICYCLE	doesn't have own propulsive power does not carry persons/goods in addition to driver

(The category VEHICLE, like GAME, is one for which it is not possible to draw up an adequate set of necessary and sufficient features; notice, however, that there may be features—[CONCRETE] is a possible example—which are necessary.)

3.2.3.6 Levels of categorization

Categories occur at different levels of inclusiveness, with more specific ones nested within more inclusive ones:

- (2) a. vehicle—*car*—hatchback
- b. fruit—*apple*—Granny Smith
- c. living thing—creature—animal—*cat*—Manx cat
- d. object—implement—cutlery—*spoon*—teaspoon

Normally, one level of specificity in each set, called the **basic** (Rosch et al. 1976) or **generic** (Berlin et al. 1973) level of specificity, has a special status, and importance. (The basic level items in (2) are printed in bold italic.) In addition to the basic level, two further levels of specificity are commonly identified: **superordinate level** and **subordinate level**. These are not defined merely by their position in the chain—there are other characteristics which distinguish one level from another.

Basic-level categories. The main characteristics of basic-level items are as follows:

- (i) It is the most inclusive level at which there are characteristic patterns of behavioural interaction. Imagine being asked to mime how one would behave with an animal. This is rather difficult without knowing whether the animal in question is a crocodile or a hamster. Likewise with, say, an item of furniture. However, the assignment is relatively easy if it involves a cat, horse, mouse, or chair.
- (ii) It is the most inclusive level for which a clear visual image can be formed. This is similar in principle to the previous characteristic: try to visualize an item of cutlery or a vehicle, without it being any specific type. A fork or a lorry, however, are easy to visualize.

- (iii) It is the most inclusive level at which part-whole information, including relations between parts, is relevant. This is particularly marked in the case of categories of artefacts above the basic level, such as CROCKERY, FURNITURE, and TOOL. But many biological categories such as ANIMAL, INSECT, and FLOWER also encompass a variety of forms.
- (iv) It is the level of most rapid categorization. Individual items can be assessed for their membership of basic-level categories more rapidly than their membership of categories which depart from basic level. Thus, confronted by, say, an Alsatian, we can decide more quickly that it belongs to the category DOG than that it belongs to either ALSATIAN or ANIMAL.
- (v) It is the level used for neutral, everyday reference. They are often felt by speakers to be the 'real' name of the referent. Suppose A and B are sitting at home; A hears a noise outside and says *What's that?* B looks out of the window and sees an Alsatian in the garden. How does B reply? Out of the following choices, normally (b) will be chosen:
 - (a) It's an animal.
 - (b) It's a dog.
 - (c) It's an Alsatian.

The other two responses would require special contextual conditions.

- (vi) The names of basic-level categories tend to be morphologically simple, and 'original', in the sense of not being metaphorical extensions from other categories. Take the case of *spoon*, which is a basic-level term; all the more specific categories have more complex names: *teaspoon*, *tablespoon*, *soup spoon*, *coffee spoon*, etc.
- (vii) The basic level is the level at which the best categories are created. Good categories are those which maximize the following characteristics:
 - a. distinctness from neighbouring categories
 - b. internal homogeneity
 - c. informativeness.

Generally speaking, categories which are more inclusive than the basic level (e.g. ANIMAL) have less internal homogeneity, while narrower categories (e.g. ALSATIAN) show less distinctness from neighbouring categories. Consider a division of animals into male and female: this would yield two clear categories which might have utility in certain circumstances. But they would not be good categories by the above criteria because (i) distinctness from neighbouring categories is restricted to one feature, (ii) internal homogeneity is likewise restricted: as a result, a female mouse resembles a male mouse far more than it resembles a female elephant (and the same is true for all animals), even though it falls into a different category. The third characteristic involves the amount of information made available by the fact that one knows what category something belongs to.

Superordinate-level categories. The superordinate level comprises categories that are more inclusive than basic-level categories. Each superordinate category typically covers a number of basic-level categories. The main characteristics of the superordinate level are as follows:

- (i) In terms of the 'good category principle', superordinate categories such as ANIMAL, BIRD, FISH, and INSECT score highly on distinctness from sister categories, but have relatively low internal homogeneity compared with basic-level categories.
- (ii) Superordinate-level categories have fewer defining attributes than do basic-level categories, and as a result, family resemblance relations are less marked.
- (iii) Names for superordinate categories, especially for categories of artefacts, often differ from basic-level names in respect of the count noun/mass noun distinction (see Chapter 13 for a more detailed account of this distinction). For instance, in the case of *cup*, *plate*, *bowl* (all count nouns), the superordinate is *crockery*, a mass noun. Other examples are: *cutlery* (mass noun), *knife*, *fork*, *spoon* (count nouns); *furniture* (mass noun), *table*, *chair*, *cupboard* (count nouns); *footwear* (mass noun), *boot*, *shoe*, *sandal* (count nouns). The reverse situation is illustrated by, for instance, *iron* and *copper* (mass nouns), which are both *metals* (count noun); similarly *beer*, *wine*, *lemonade* (mass nouns), which are *beverages* (count noun), and *pepper*, *pimento*, *coriander* (mass nouns) which are all *spices* (count noun).

Subordinate-level categories. Subordinate level categories differ from basic-level categories in the following main ways:

- (i) By the good category principle, they are less good than basic-level categories, because although members under the same superordinate resemble one another to a high degree, they have low distinctiveness from members of sister categories.
- (ii) They do not add a great deal of information to what is provided by the basic-level category which includes them. When subjects are asked to list distinctive attributes for a subordinate category, their lists differ very little from the lists they give for the corresponding basic-level item.
- (iii) Names for subordinate-level categories are frequently morphologically complex, the most common pattern being a modifier-head structure (for instance, *teaspoon*, *bread knife*, *coffee cup*, *blue tit*, *herring gull*). This is sometimes taken as an indication that they are distinguished from their corresponding basic-level category by a single property or feature. However, this is misleading. The 'single property' is a feature of the name, but one cannot automatically assume that the conceptual content is also distinguished by a single feature. Usually the property signalled by the name has an identifying role, but the item typically differs from the including basic-level category (and from its sisters) in respect of a range of properties.

3.2.3.7 Problematic aspects of the prototype model

While the standard prototype-theoretical approach undoubtedly sheds light on the nature of natural conceptual categories, it is not without its problematic aspects.

The 'guppy effect'. The categories which result from the combination basic categories have often been claimed to be problematical for prototype theory. The most famous example is the category PET FISH. It appears that what emerges as the best example in experimental studies of this category (at least in an American setting) is GUPPY. The problem is that a guppy is judged to be neither a prototypical pet (CAT and DOG, for instance, are given higher GOE ratings) nor a prototypical fish (TROUT gets a higher rating). It is not entirely clear why one should expect the prototype of a compound category XY to have prototype status in x and y separately. But one might expect the prototype approach to provide a way of predicting the prototype of a compound category from the representations of the constituent categories. Attempts to do this have been inconclusive (for an example, see Hampton (1991)).

The basis of GOE ratings and their interpretation. The first point to make is that although subjects readily enough make GOE judgements on the basis of two words (category name and item name), this is surely rather unnatural: it would presumably be more revealing to produce GOE ratings for actual objects or events, etc. Furthermore, this would be likely to highlight the fact that the GOE scale is a conflation of several more basic scales. One of these is undoubtedly familiarity, although it can be shown that GOE ratings cannot be reduced to familiarity ratings. Another is well-formedness: APPLE may well receive a high rating in the category FRUIT if only the words are presented, but what if an actual apple were presented, and it happened to be rotten? Well-formedness does not necessarily correlate with familiarity. Most mushrooms are at least slightly deformed in one way or another. Yet there seems little doubt that a perfectly formed specimen would receive the highest GOE rating (other things being equal). Another factor is important, which in Cruse (1990) is called 'quality'. Think of an emerald. Most emeralds are pale in colour and have faults in the form of tiny cracks, etc. The best emeralds are deep in colour, but these are rare, and are even more susceptible to faults. An emerald with a deep glowing green colour would be voted the prototype on the basis of its 'quality', which is distinct from frequency and well-formedness. Here, then, we have at least three independent strands potentially making up a GOE score, and there may be more.

Interaction among features. A major criticism of the prototype model of category structure is that a simple feature list is far too simplistic. It is true that the move from the classical 'necessary and sufficient features' model to a 'more features = more central' model represents an improvement. And further refinements are possible, for instance, assigning different 'weights' to different features, so that some have a greater effect on GOE rating than others. But an important aspect of category structure remains unaccounted for, and that is the

interaction between features. This has several aspects. One concerns GOE ratings. There is not a straightforward relationship between number of features and GOE rating, even if weighted features are used. The reason is that the effect on GOE rating of a single feature can depend on the presence or absence of other features. To take a simple example, informants asked to rate spoons for GOE will judge a wooden spoon to be less good than an otherwise similar metal one if it is small, but equally good if it is large.

A perhaps more important aspect of the interrelatedness of features is functional relatedness. It is obvious that central examples of the category BIRD will possess the features [FLIES] and [HAS WINGS]. What is missing from a list of features is that wings are necessary for flight.

Finally, consider the relation between the features [HAS WINGS] and [HAS A BEAK] in connection with the category BIRD. We will probably want to say that aeroplanes and butterflies also have wings, but they are very different from the wings of birds. All these interactions detract from the adequacy of a feature list as a characterization of a natural category.

Context sensitivity. Another shortcoming of the prototype model is that it is not able to account for context sensitivity. Many experimental studies have demonstrated that judgements as to which are the best members of a category can be influenced by providing a context for the judgement. But it is intuitively obvious that judgements of the 'best' examples of, say, the category CAR are going to depend on whether one has in mind a racing context, a context of town use, or long-distance travel. It seems likely that if none of these are made explicit, then the word *car* evokes some sort of 'default' context; it is unlikely that we make our judgements in a genuine zero context.

Category boundaries. One of the most serious shortcomings of the standard prototype view is that no category boundary is recognized. The few scholars who do admit that a boundary exists, evince little interest in it (e.g. Lakoff). Yet a category without a boundary is virtually useless: a primary function of a category is to discriminate between things which are in it and things which are not in it. The classical view of categories, with necessary and sufficient features, set a boundary (albeit an unnaturally sharp one) but allowed for no internal structure. In getting rid of necessary and sufficient features, prototype theory has thrown the baby out with the proverbial bath water: by providing an account of gradable centrality, it has lost the ability to set a boundary. The view taken here is that a fully satisfactory description of a category must specify both the internal structure and the location of a boundary or boundary area.

Category boundaries are important for a number of reasons. They crop up in everyday language use. If someone declares of some entity A *A is an X* and someone else disagrees and says *No, it isn't*, then either they disagree as to the nature of A, or they disagree as to the location of the boundary of the category X. Certain adjectives, such as *imitation* as in *imitation leather*, or *fake* as in *a fake*

Picasso, have an implicit negating effect—imitation leather does not fall into the category LEATHER, and a fake Picasso falls outside the category WORK BY PICASSO.

Prototype theory focuses on category inclusion and neglects category exclusion. Hence there is no explanation for the mutual exclusion relation between, for instance, the categories CAT and DOG, or TULIP and DAFFODIL. This is easy to account for in terms of the location of category boundaries. But the only way to build exclusion into a prototype representation would be to have features with negative weighting. For instance, [BARKS], [WAGS TAIL WHEN PLEASED], [HAS FEATHERS], [LAYS EGGS] would all have negative weighting for the category CAT.

The position of a category boundary has psycholinguistic consequences. In an experimental task in which subjects have to say as quickly as possible whether or not an item belongs to a given category, it is found that the speed of response is slowest for items close to the category boundary. Hence, in (3), the **bold** items will be responded to slowest:

- (3) FRUIT *apple* ***tomato*** *potato*
 VEHICLE *car* ***bicycle*** *chair*
 WORD *hand* ***malk*** *pkhq*

The location of the boundary of a category is not determined by the location of its prototype core. Two categories may have more or less identical prototypes but their boundaries may be different; equally, two categories may have the same boundaries but differ in their prototypes. The first case can be illustrated by the French word *corde* and its English equivalent *rope*. Native informants suggest the same sort of examples as the best in the respective categories, from which we infer that the respective prototypes are at least very similar. However, there is evidence that the boundaries of CORDE and ROPE are not the same. The French dictionary *Le Petit Larousse* defines *ficelle* ("string") as "une corde mince" (= "a thin rope"); it would never occur to an English lexicographer to define *string* as "thin rope". It seems, therefore, that although FICELLE falls within the boundary of the category CORDE, STRING falls outside the boundary of the category ROPE, in other words, CORDE is the broader category. The converse case, of two categories with the same boundary but differing in the location of their cores, is harder to find, but perhaps COURAGE and BRAVERY are examples. There is something contradictory about ?*The men showed great bravery, but no courage*. But there is evidence that their core regions are not the same. Student informants in an informal test rated (4) as a better example of a brave act, and (5) as a better example of a courageous act:

- (4) A person jumps into a fast-flowing river in an attempt to save someone who has fallen in.
 (5) A person risks his/her career and livelihood by exposing malpractice and injustice at the heart of government.

Degrees of membership and fuzzy boundaries. There has been some dispute among prototype theorists concerning the relationship between the GOE of an item with respect to a category and its **degree of membership** (henceforward DOM). Some say that in giving GOE scores, subjects are in fact giving DOM scores. The claim is that only the prototype of a category has 100 per cent membership. However, one must bear in mind that subjects in the standard experiments were not actually asked to rate degree of membership. What they were asked to do was to rate items as to how good they were as members of particular categories. Saying that they were giving DOM ratings is a subsequent interpretation, requiring a construal of membership as a gradable property. While this is possible, it renders one of the measures, GOE and DOM, redundant. This would be a pity, firstly because there is intuitive resistance to the proposition that an ostrich, say, is only 30 per cent a bird. Secondly, it is possible to define DOM in such a way that it designates a different property from GOE. There are two ways of doing this.

The first way requires us to think of a category boundary as relatively sharp, and to construe the category as a container. Everything wholly on one side of the boundary will have 100 per cent membership, whatever its location relative to the prototype; anything that is wholly on the other side of the boundary will have zero membership. However, thinking of an ordinary container, it is possible for something to be partly in it and partly outside it, i.e. to be fractionally contained. This gives us a way of interpreting the notion of partial membership of a category. When we say, for instance, that a priest is *to some extent a social worker*, we are effectively placing him partly in the category and partly not in the category. That is to say, we are construing the categories PRIEST and SOCIAL WORKER as partially overlapping.

An alternative way of looking at partial membership of a category is to think of a category as having a borderline area, rather than a sharp boundary; that is to say, the boundary has a definite location, but is to a greater or lesser degree fuzzy. On this view, any particular item will be a full member of a particular category, or not a member at all, or it will be a borderline example. It is amongst borderline examples that the notion of partial membership applies. For instance, most people would probably judge BICYCLE and SKATEBOARD to be borderline instances of the category VEHICLE. Here, the notion of degree of membership becomes operational, and I myself, for instance, would judge BICYCLE to have a higher degree of membership than SKATEBOARD.

In Croft and Cruse (2004: 95) it is argued that there is no need for the notion of 'fuzzy boundary', and that the experimental results usually taken to point to the fuzziness of category boundaries can be explained by different construals of (sharp) boundaries on different occasions and in different contexts. On this view, only the first interpretation of partial membership would be valid. The situation with regard to BICYCLE and SKATEBOARD would be that on some occasions the boundary of VEHICLE would be construed to include BICYCLE, but exclude SKATEBOARD, but on other occasions would be construed to include (or exclude) both.

3.2.4 Dynamic construal

Underlying most work on conceptual categories there has been an implicit assumption that mature, normal human beings possess an inventory of discrete and more or less stable categories with permanent mental representations, which are activated in appropriate circumstances. (Of course, new concepts can be learned, but these will have the same stable character as those already in the inventory.) The goal of conceptual studies, it is assumed, is to elucidate the nature of these mental entities. Conceptually oriented studies of word meaning typically assume that a word is associated with a specific concept (or, in the case of ambiguity, with a specific set of such concepts). However, the notion of fixed categories has not gone unchallenged. For instance, experiments by Barsalou, a cognitive psychologist, have shown that people are capable of forming ad hoc categories with no conventional names, such as THINGS TO TAKE ON A PICNIC, THINGS ON A DESK THAT COULD BE USED TO POUND IN A NAIL, SUITABLE BIRTHDAY PRESENTS FOR A TEN-YEAR-OLD, and so on. These categories showed the same characteristics as established categories, with graded centrality and characteristic features. These and similar results suggest that categories are generated online as and when required. However, all is not flux and chaos—that would be fatal for communication. The suggestion is that previous experience has a permanent effect on our tendency to construe concepts in particular ways. The psychologists Smith and Samuelson have this to say:

Critically, the accrual of these long-term changes provides a source of stability in a continually changing system. If there are statistical regularities, patterns, in our experiences that recur over and over again, then as each moment of knowing is laid on the preceding moments, weak tendencies to behave and to think in certain ways will become strong tendencies—sometimes so strong that they will not be easily perturbed and thus might seem fixed. (Smith and Samuelson 1997: 175–6)

In other words, concepts may have any degree of fixedness/fluidity. The consequences of this view for word meaning will be taken up in Chapter 5.

Discussion questions and exercises

1. Prototype features

Suggest a set of prototype features for one or more of the following conceptual categories (or select your own example(s)):

CLOTHES FRUIT MUSICAL INSTRUMENT HOBBY BUILDING HOUSEHOLD APPLIANCE

For each category, draw up a list of possible members, including some marginal cases, and ask another person to assign GOE ratings. Consider to what extent the ratings can be accounted for in terms of your suggested features.

2. Category levels

Assign the following categories to superordinate, basic, or subordinate level:

BIRO TEASPOON SANDAL UNDERWEAR SEAGULL DAISY GRASS BULLDOZER BUS
MOUNTAIN BIKE SELF-RAISING FLOUR WALNUT SUGAR ARMCHAIR DELICATESSEN SUPERMARKET
PETROL STATION TOWN HALL PARK MOTORWAY ROAD CANAL POLICE STATION BUILDING
GROCERIES WINE CHAMPAGNE BEVERAGE MILK

Suggestions for further reading

That meaning is essentially conceptual in nature is one of the central tenets of **cognitive linguistics**. For elementary introductions to cognitive linguistics see Ungerer and Schmid (1996). Ultimately, a reader interested in this approach will eventually want to tackle the foundational texts. The 'bible' of the cognitive approach is Langacker's two-volume *Foundations of Cognitive Grammar* (1987 and 1991a). However, this is not an easy read; fortunately, many of the basic topics are expounded in a much more accessible form in Langacker (1991b). The interested reader will also find articles on a wide range of cognitive linguistic topics in the journal *Cognitive Linguistics*. See Millikan (2000) for a thorough discussion of the notion of 'kinds'.

An alternative 'conceptual' approach to meaning can be found in the works of Jackendoff; Jackendoff (1983) provides a good introduction. An interesting comparison between Jackendoff's approach and the cognitive linguistic approach (including a contribution from Jackendoff himself) can be found in Volume 7.1 of *Cognitive Linguistics*, which also gives a fairly full bibliography of Jackendoff's later work.

Cruse (1990) provides an introduction to prototype theory as applied to lexical semantics. (The volume which includes this article also contains many other articles on the topic.) A fuller account is to be found in Taylor (1989); Cruse (1992c) is a critical review of this. Ungerer and Schmid (1996) has an interesting chapter on categorization. For a more psychological view of the prototype approach to meaning, see the articles in Schwanenflugel (1994), especially those by Hampton and Murphy. Cruse (1995) attempts to apply prototype theory to lexical relations.

Part 2

Words and Meanings

To the layman, words are the bearers of meaning par excellence in language. While it is in danger of understating the importance of other linguistic structures and phenomena in the elaboration of meaning, this view is not entirely unjustified: words do have a central role to play in the coding of meaning, and are responsible for much of the richness and subtlety of messages conveyed linguistically. Hence it is no accident that this section of the book is the most substantial. Here, after the introductory Chapter 4, we discuss how word meanings vary with context (Chapter 5), paradigmatic sense relations (Chapters 6 and 7), the nature of lexical hierarchies (Chapter 8), how the meanings of words in the same grammatical sequence interact (Chapter 9), the problems of describing word senses (Chapters 10 and 11), and finally, how new meanings grow out of old ones (Chapter 12).

CHAPTER 4

Lexical units

4.1 What is a word?	75	4.4 Lexification	82
4.2 The notion "possible word meaning"	77	4.5 Delexification	91
4.3 Words vs lexical units	81		

Lexical units

4.1 What is a word?

In a descriptive introduction to meaning such as this, it is inevitable that the meanings of words should loom large, even though in more formally oriented accounts, word meanings are left largely unanalysed, or are reduced to mere skeletons of their true selves. There are, of course, more or less reputable justifications for such neglect. However, most (linguistically innocent) people have an intuition that meaning is intimately bound up with individual words; indeed, this, *par excellence*, is what words are for. While such an intuition seriously underestimates other aspects of meaning, it is not, in itself, wrong, and an adequate introduction to meaning should not shrink from the slipperiness and complexity of word meaning simply because it cannot be neatly corralled into a favoured formalization. Hence, the present and the following eight chapters will be devoted to various aspects of lexical semantics. We begin by looking at what sort of thing a word is.

There has been a great deal of discussion of the nature of the word as a grammatical unit, too much even to summarize here. Most of it, anyway, is not relevant to our concerns. But it is as well to have some idea of what we are dealing with. The notion has notoriously resisted precise definition. Probably the best approach is a prototypic one: what is a prototypical word like? Well, for our purposes, the classical characterization as 'a minimal permutable element' will serve. This attributes two features to a prototypical word:

- (i) It can be moved about in the sentence, or at least its position relative to other constituents can be altered by inserting new material.
- (ii) It cannot be interrupted or its parts reordered.

In other words, in modifying a particular sentence, we are by and large obliged to treat words as structurally inviolable wholes. Let's see briefly how this works. Take a sentence like (1):

(1) The government is strongly opposed to denationalization.

Reordering appears in such examples as (2)–(4):

(2) The government is opposed to denationalization—strongly.

(3) What the government is strongly opposed to is denationalization.

(4) It is denationalization that the government is strongly opposed to.

And the possibilities for the insertion of new material are as follows:

(5) The (present) government, (apparently), is (very) strongly (and implacably) opposed (not only) to (creeping) denationalization, butetc.

Notice that the only possible insertion points are between words. Words, of course, are separated by spaces in writing, although not usually by silences in speech. They also have a characteristic internal structure, in that they prototypically have no more than one **lexical root**. This notion will become clearer below, but, for instance, the lexical roots of the following words are shown in capitals:

GOVERNment	reORDERing	STRONGly	deNATIONalization
OPPOSed	TYPically	CLEARer	LEXical

Some words, such as HEDGE-HOG, BUTTER-FLY, and BLACK-BOARD seem to have more than one lexical root. These, however, are atypical, and for many of them it is possible to argue that the apparent roots are not fully autonomous, semantically, but form a fused root. Other words have no lexical roots at all: these are the so-called **grammatical words** like *the*, *and*, and *of*.

At this point it is necessary to be somewhat more precise about what we mean by a word. In one sense, *obey*, *obeys*, *obeying*, and *obeyed* are different words (e.g. for crossword purposes), in another sense, they are merely different forms of the same word (and one would not, generally speaking, expect them to have separate entries in a dictionary); on the other hand, *obey* and *disobey* are different words in both senses, but *bank* (river) and *bank* (money) are the same word for crossword purposes, but we would expect them to have separate dictionary entries and they are therefore different words in the second sense. Finer distinctions are possible, but for our purposes it will be sufficient to distinguish **word forms** and **lexemes**. Word forms, as the name suggests, are individuated by their form, whether phonological or graphic (most of our examples will be both); lexemes can be regarded as groupings of one or more word forms which are individuated by their roots and/or derivational affixes. So, *run*, *runs*, *running*, and *ran* are word forms belonging to the same lexeme **run**, while *walk*, *walks*, *walking*, *walked* belong to a different lexeme, **walk**, distinguished from the former by its root; likewise, *obey*, *obeys*, *obeying*, and *obeyed* belong to a single lexeme and *disobey*, *disobeys*, *disobeying*, and *disobeyed* despite having the same root as the first set, belong to a different lexeme, distinguished this time by the possession of the derivational affix *dis-*. A simple test for derivational affixes (the matter is in reality, however, not quite as clear-cut as this) is that they are never grammatically

obligatory. For instance, in *John is disobeying me*, *disobey* can be substituted by *watch*, without giving an ungrammatical sentence, which shows that *dis-* is not essential to the grammatical structure of the sentence. This is true of all occurrences of *dis-*. On the other hand, any verb which will fit grammatically into the frame *John is — me* must bear the affix *-ing*, showing that it is not a **derivational**, but an **inflectional** affix: word forms that differ only in respect of inflectional affixes belong to the same lexeme. It is the word-as-lexeme which is the significant unit for lexical semantics.

4.2 The notion “possible word meaning”

It is worthwhile to pose the question of whether there are any restrictions on possible meanings for words. We may approach this in two stages. We can first of all ask whether there are any universal restrictions; and then we can inquire as to the existence of language-specific restrictions.

Let us take the first question first. Is there anything conceivable that could never be the meaning of a word? It will be as well to restrict ourselves to notions that can be expressed by a combination of words, otherwise we shall be in really deep water. One line of thinking can, I think, be disposed of relatively quickly. It may be thought that no language could possibly have a word meaning, for instance, “to face west on a sunny morning while doing something quickly”. I confess that I would be astonished to find such a word. But the reason is not that it is theoretically impossible, but that it would be of such limited utility. Languages have words, at least partly, because in the cultures they serve, the meanings such words carry need to be communicated. (Of course, cultural evolution can leave words stranded, as it were, but this does not invalidate the basic point that words at some stage must be motivated in terms of possible use.) This means that if some culture had a use for the notion expressed, then it would not be surprising if there were a word for it. In the case in question, for instance, maybe the word could designate a specific sort of act of disrespect towards the Sun God, which carried specific penalties. If we take into account the possibility of outlandish (to us) religious beliefs, it is clear that the scope for improbable word meanings of this sort is (almost) unlimited.

Now let us look at a different sort of case. Take the sentence (6):

(6) The woman drank the wine slowly.

The notion “drank slowly” could easily be lexicalized (i.e. expressed by a single word): we have in English, after all, verbs such as *quaff* and *sip*, which combine the meaning of “drink” with some adverbial manner component. Similarly, a verb meaning “drink wine” is not at all implausible, as one of the senses of *drink* in English is specifically “drink alcoholic beverage” (as in *Mary doesn’t drink, she’ll just have an orange juice*). In contrast to these more-or-less plausible

word meanings, consider next the possibility of having a word meaning "The woman drank" (*blisk*), or "the wine slowly" (*blenk*). On this system, *Blisk wine* would mean "The woman drank wine", and *The woman drank blenk* would mean "The woman drank the wine slowly". It seems clear that here we are in the realms not of implausibility, but of impossibility. As a further example, consider the phrase *very sweet coffee*. It is perfectly within the bounds of possibility that there should be a single word meaning "sweet coffee", or "very sweet", even "very sweet coffee", but it is not conceivable that there should be a word meaning "very — coffee" (i.e. any adjective applied to *coffee* would be automatically intensified). What is the difference between the possible and the impossible cases? There seem to be two parts to the answer. First, a word meaning is not allowed to straddle the vital subject-predicate divide. Second, possible word meanings are constrained in a strange way by semantic dependencies. It is first necessary to distinguish **dependent** and **independent** components of a semantic combination. The independent component is the one which determines the semantic relations of the combination as a whole with external items. So, for instance, in *very large*, it is *large* which governs the combinability of the phrase *very large* with other items. Thus the oddness of, say, *?a very large wind* is attributable to a semantic incompatibility between *large* and *wind*—there is no inherent clash between *very* and *wind*, as the normality of *a very hot wind* demonstrates. By similar reasoning, the independent item in *warm milk* is *milk*, and in *drink warm milk* is *drink*. By following this line of reasoning, we can establish chains of semantic dependencies. For instance, the chain for *very young boy* is:

"very" → "young" → "boy"

and that for *drink warm milk* is:

"warm" → "milk" → "drink"

The constraint that we are looking at says that the elements that constitute the meaning of a word must form a continuous dependency chain. This means, first, that there must be a relation of dependency between elements. This rules out "wine slowly" as a possible word meaning, because there is no dependency between "wine" and "slowly" in "Drink wine slowly". Second, there must be no gaps in the chain which need to be filled by semantic elements from outside the word. This rules out cases like "very — milk", where the dependency chain would have to be completed by an external item such as "hot".

Two further suggestions regarding possible word meanings are worth examining. The first is from the famous American linguist Noam Chomsky (1965: 201):

... there are no logical grounds for the apparent non-existence of words such as *LIMB*, similar to *limb* except that it designates the single object consisting of a dog's four legs so that *its LIMB is brown* ... would mean that the object consisting of its four legs is brown. Similarly, there is no a priori reason why a natural language could not contain a word *HERD*, like the collective *herd* except that it denotes a single scattered object with cows as parts, so that *a cow lost a leg* implies *the HERD lost a leg*, etc.