Categorisation and idealised cognitive models

In this chapter, we continue our exploration of the human conceptual system by focusing on categorisation: our ability to identify perceived similarities (and differences) between entities and thus group them together. Categorisation both relies upon and gives rise to concepts. Thus categorisation is central to the conceptual system, because it accounts, in part, for the organisation of concepts within the network of encyclopaedic knowledge. Categorisation is of fundamental importance for both cognitive psychologists and semanticists, since both disciplines require a theory of categorisation in order to account for knowledge representation and indeed for linguistic meaning. Central to this chapter is the discussion of findings that emerged from the work of cognitive psychologist Eleanor Rosch and her colleagues in the 1970s, and the impact of these findings on the development of cognitive semantics. In particular, we will be concerned with the work of George Lakoff, who addressed findings relating to prototype structure and basic level categories revealed by research in cognitive psychology, and who developed a cognitive semantic theory of idealised cognitive models (ICMs) in order to account for these phenomena. The influence of Lakoff's research, and of his book Women, Fire and Dangerous Things (1987), was important for the development of cognitive semantics. In particular, this book set the scene for cognitive semantics approaches to conceptual metaphor and metonymy, lexical semantics (word meaning) and grammatical structure. In this chapter, then, we set out the theoretical background of Chapters 9 and 10 where we will address Lakoff's theory of conceptual metaphor and metonymy and his theory of word meaning in detail.

We begin the chapter by explaining how Rosch's research on categorisation was important in the development of cognitive semantics, setting this discussion against the context of the classical view of categorisation that was superseded by Rosch's findings. We then look in detail at the findings to emerge from Rosch's research (section 8.2) and explore the development of Lakoff's theory of cognitive models that was developed in response to this research (section 8.3). Finally, we briefly explore the issue of linguistic categorisation in the light of the empirical findings and theoretical explanations presented in this chapter (section 8.4).

8.1 Categorisation and cognitive semantics

In the 1970s the **definitional** or **classical theory** of human categorisation – so called because it had endured since the time of the ancient Greek philosophers over 2,000 years ago – was finally called into question. The new ideas that contributed most significantly to this development are grouped together under the term **prototype theory**, which emerged from the research of Eleanor Rosch and her colleagues. In fact, 'Prototype Theory' was less a theory of knowledge representation than a series of findings that provided startling new insights into human categorisation. In so far as the findings led to a theory, Rosch proposed in her early work that humans categorise not by means of the **necessary and sufficient conditions** assumed by the classical theory (to which we return below), but with reference to a **prototype**: a relatively abstract mental representation that assembles the key attributes or features that best represent instances of a given category. The prototype was therefore conceived as a schematic representation of the most salient or central characteristics associated with members of the category in question.

A problem that later emerged was that the view of prototypes as mental representations failed to model the relational knowledge that humans appear to have access to (recall from the last chapter that relational knowledge is one of the properties of encyclopaedic knowledge addressed by Frame Semantics). These criticisms led to further developments in prototype theory. Some scholars argued for a revised view of the prototype, suggesting that the mental representation might correspond to an exemplar: a specific category member or 'best example' of a category, rather than a schematic group of attributes that characterise the category as a whole. However, these exemplar-based models of knowledge representation were also problematic because they failed to represent the generic information that humans have access to when they use concepts in order to perform a host of conceptual operations, including categorisation. Indeed, the most recent theories of categorisation assert that a key aspect of knowledge representation is the dynamic ability to form simulations, an idea that was introduced in the previous chapter. Thus, in a number of respects, prototype theory has been superseded by more recent empirical findings and theories. Despite this, there are a number of reasons why a chapter on categorisation in general, and

prototype theory in particular, is essential for a thorough understanding of cognitive semantics.

Firstly, an investigation of prototype theory provides a picture of the historical context against which cognitive linguistics emerged as a discipline. The development of prototype theory in the 1970s resonated in important ways with linguists whose research would eventually contribute to defining the field of cognitive semantics. Charles Fillmore and George Lakoff were both members of faculty at the University of California at Berkeley where Eleanor Rosch was also conducting her research, and both were influenced by this new approach to categorisation. For Lakoff in particular, Rosch's discovery that psychological categories did not have clearly definable boundaries but could instead be described as having 'fuzzy' boundaries reflected his own views about language: Lakoff thought that lexical and grammatical categories might also be most insightfully conceived as categories with rather fluid membership. This led Lakoff to apply this new view of psychological categories to linguistic categories (such as word meanings). In this way, 'Prototype Theory' inspired some of the early research in cognitive semantics.

Secondly, and perhaps more importantly, although it now seems that prototype theory cannot be straightforwardly interpreted as a theory of knowledge representation, the empirical findings that emerged from this research demand to be accounted for by any theory of categorisation. In other words, the **prototype effects** or **typicality effects** that Rosch discovered are psychologically real, even if the early theories of knowledge representation that were proposed to account for these effects have been shown to be problematic. Indeed, a central concern in Lakoff's (1987) book was to address the problems that early prototype theory entailed, and to propose in its place a theory of cognitive models.

Thirdly, as we mentioned above, Lakoff's (1987) book set the scene for the development of three important strands of research within cognitive linguistics: (1) Conceptual Metaphor Theory (Chapter 9); (2) cognitive lexical semantics (Chapter 10); and (3) a cognitive approach to grammar that influenced the well-known constructional approach developed by his student Adele Goldberg (to which we return in Part III of this book).

Finally, *Women, Fire and Dangerous Things*, despite its rather meandering presentation, in many ways defines the two key commitments of cognitive linguistics: the 'Generalisation Commitment' and the 'Cognitive Commitment'. Lakoff's book took what was then a relatively new set of findings from cognitive psychology and sought to develop a model of language that was compatible with these findings. In attempting to model principles of language in terms of findings from cognitive psychology, Lakoff found himself devising and applying principles that were common both to linguistic and conceptual phenomena, which thus laid important foundations for the cognitive approach to language.

8.1.1 The classical theory

Before presenting Rosch's findings concerning categorisation, it is important to set her research in some historical context. The 'classical theory' of categorisation was the prevalent model since the time of Aristotle and holds that conceptual and linguistic categories have definitional structure. This means that an entity represents a category member by virtue of fulfilling a set of necessary and (jointly) sufficient conditions for category membership. These conditions are called 'necessary and sufficient' because they are individually necessary but only collectively sufficient to define a category. Traditionally, the conditions were thought to be sensory or perceptual in nature. To illustrate, consider once more the familiar lexical concept BACHELOR. For an entity to belong to this category, it must adhere to the following conditions: 'is not married'; 'is male'; 'is an adult'. Each of these conditions is necessary for defining the category, but none of them is individually sufficient because 'is not marrried' could equally hold for SPINSTER, while 'is male' could equally hold for HUSBAND, and so on. In theories of linguistic meaning, necessary and sufficient conditions have taken the form of semantic primitives or componential features, an idea that we have mentioned in previous chapters (recall our discussion of semantic universals in Chapter 3 and our discussion of the dictionary view of linguistic meaning in Chapter 7). As we have seen, the idea of semantic primitives has been influential in semantic theories that adopt the formal 'mentalist' view proposed by Chomsky, which is primarily concerned with modelling an innate and specialised system of linguistic knowledge. This is because, in principle at least, semantic primitives suggest the possibility of a set of universal semantic features that can be combined and recombined in order to give rise to an infinite number of complex units (word meanings). This approach is reminiscent of the characterisation of human speech sounds in phonetics and phonology, where a bundle of articulatory features makes up each speech sound. It is also reminiscent of the characterisation of sentence structure in terms of strings of words that combine to make phrases, which then combine to make sentences. In other words, the influence of the semantic decomposition approach reflects the influence of structural approaches to sound and grammar upon the development of theories of word meaning. This kind of approach is attractive for a formal theory because it enables the formulation of precise statements which are crucial to the workings of the 'algorithmic' or 'computational' model favoured by these approaches. For example, Katz (1972) argued that the English noun *chair* names a category that can be decomposed into the set of semantic features or markers shown in Table 8.1.

However, while many (usually formal) linguists would argue that 'decompositional' approaches have worked rather well for modelling the structural aspects of language such as phonology or syntax, many linguists (both formal

| BJECT | |
|----------------------|--|
| HYSICAL | |
| ON-LIVING | |
| RTEFACT | |
| URNITURE | |
| ORTABLE | |
| OMETHING WITH LEGS | |
| OMETHING WITH A BACK | |
| OMETHING WITH A SEAT | |
| EAT FOR ONE | |

| Table 8.1 | Semantic | features or | markers | for the | category | CHAIR |
|-----------|----------|-------------|---------|---------|----------|-------|
|-----------|----------|-------------|---------|---------|----------|-------|

and cognitive) also recognise that the classical decompositional theory of word meaning suffers from a number of problems. We discuss here three of the most serious problems with this approach.

8.1.2 The definitional problem

While the classical theory holds that categories have definitional structure, in practice it is remarkably difficult to identify a precise set of conditions that are necessary and sufficient to define a category. This requires the identification of all those features that are shared by all members of a category (necessary features) and that together are sufficient to define that category (no more features are required). The following famous passage from the philosopher Wittgenstein's discussion of the category GAME illustrates the difficulty inherent in this approach:

Consider for example the proceedings that we call 'games'. I mean board-games, card-games, ball-games, Olympic games and so on. What is common to them all? - Don't say: 'There must be something common, or they would not be called "games"' - but look and see whether there is anything common to all. - For if you look at them you will not see something that is common to all, but similarities, relationships, and a whole series of them at that. To repeat: don't think, but look! - For example at board-games, with their multifarious relationships. Now pass to card-games; here you find many correspondences with the first group, but many common features drop out, and others appear. When we pass next to ball-games, much that is common is retained, but much is lost. - Are they all 'amusing'? Compare chess with noughts and crosses. Or is there always winning and losing, or competition between players? Think of patience. In ball-games there is winning and losing; but when a child throws his ball at the wall and catches it again, this feature has disappeared. Look at the parts played

by skill and luck; and at the difference between skill in chess and skill in tennis. Think now of games like ring-a-ring-a-roses; here is the element of amusement, but how many other characteristic features have disappeared! And we can go through the many, many other groups in the same way; we see how similarities crop up and disappear. (Wittgenstein 1958: 66)

This passage reveals that there is no single set of conditions that is shared by every member of the category GAME. While some games are characterised by AMUSEMENT, like tiddlywinks, others are characterised by LUCK, like dice games, still others by SKILL or by COMPETITION, like chess. In other words, it appears to be impossible to identify a definitional structure that neatly defines this category. To present a simpler example, consider the category CAT. We might define this category as follows: 'is a mammal'; 'has four legs'; 'is furry'; 'has a long tail'; 'has pointy ears'. What happens if your cat gets into a fight and loses an ear? Or gets ill and loses its fur? Does it then stop being a member of the category CAT? The definitional approach therefore suffers not only from the problem that the definitions are often impossible to identify in the first place, but also from the problem that definitions are, in reality, subject to exceptions. A three-legged one-eared hairless cat is still a cat. It seems, then, that a category need not have a set of conditions shared by all members in order to 'count' as a meaningful category in the human mind. It is important to emphasise here that we are not dealing with scientific categories, but with the everyday process of categorisation that takes place in the human mind on the basis of perceptual features. While a biologist could explain why a three-legged one-eared hairless cat still 'counts' as a member of that species from a scientific perspective, what cognitive psychologists and linguists want to explain is how the human mind goes about making these kinds of everyday judgements in the absence of scientific knowledge.

8.1.3 The problem of conceptual fuzziness

A second problem with the classical view is that definitional structure entails that categories have definite and distinct boundaries. In other words, an entity either will or will not possess the 'right' properties for category membership. Indeed, this appears to be the case for many categories. Consider the category ODD NUMBER. As we learn at school, members of this category are all those numbers that cannot be divided by 2 without leaving a remainder: 1, 3, 5, 7, 9 and so on. This category has clearly defined boundaries, because number is either odd or even: there is no point in between. However, many categories are not so clearly defined but instead have 'fuzzy' boundaries. Consider the category, it is less clear whether CARPET should be considered a member. Consider the

COGNITIVE LINGUISTICS: AN INTRODUCTION

| Table 8.2 Problems for the classical theory of cate | tegorisation |
|---|--------------|
|---|--------------|

Definitional problem: difficult or impossible to identify the set of necessary and sufficient conditions to define a category

The problem of conceptual fuzziness: not all categories have clear boundaries

category BIRD. While it is obvious that birds like ROBIN and SPARROW belong to this category, it is less obvious that animals like PENGUINS and OSTRICHES do, neither of which can fly. The difficulty in deciding to set the boundary for certain categories is the problem of conceptual 'fuzziness'. If the classical theory of categorisation is correct, this problem should not arise.

8.1.4 The problem of prototypicality

The third problem with the definitional view of categories is related to the problem of conceptual fuzziness, but while the problem of conceptual fuzziness concerns what happens at the **boundaries** of a category, the problem of prototypicality concerns what happens at the centre of a category. As we will see in the next section, findings from experimental cognitive psychology reveal that categories give rise to prototype or typicality effects. For example, while people judge TABLE or CHAIR as 'good examples' or 'typical examples' of the category FURNITURE, CARPET is judged as a less good example. These asymmetries between category members are called typicality effects. While we might expect this to happen in the case of categories that have fuzzy boundaries, experiments have revealed that categories with distinct boundaries also show typicality effects. For example, Armstrong et al. (1983) found that the category EVEN NUMBERS exhibits typicality effects: participants in their experiments consistently rated certain members of the category including '2', '4', '6', and '8' as 'better' examples of the category than, say, '98' or '10,002'. Categories that exhibit typicality effects are called graded categories. Typicality effects represent a serious challenge for the classical theory, because if each member of a category shares the same definitional structure, then each member should be equally 'typical'. These problems with the classical theory of categorisation are summarised in Table 8.2.

8.1.5 Further problems

Laurence and Margolis (1999) discuss further problems with this approach which we mention only briefly here. These are what they call the problem of psychological reality and the problem of ignorance and error.

The problem of typicality: many categories, including some with clear boundaries, exhibit typicality effects

The problem of psychological reality relates to the fact that there is no evidence for definitional structure in psychological experiments. For example, we might expect words with a relatively 'simple' definitional structure or small set of features (like, say, *man*) to be recognised more rapidly in word-recognition experiments than words with a more 'complex' definitional structure or greater number of features (like, say, *cousin*). This expectation is not borne out by experimental evidence. The problem of ignorance and error relates to the fact that it is possible to possess a concept without knowing what its properties are. In other words, possessing a concept is not dependent upon knowing its definition. For example, it is possible to have the concept WHALE while mistakenly believing that it belongs to the category FISH rather than the category MAMMAL.

8.2 Prototype theory

Prototype theory is most closely associated with the experimental research of cognitive psychologist Eleanor Rosch and her colleagues. In this section, we present an overview and discussion of Rosch's research, which is largely based on experimental findings.

8.2.1 Principles of categorisation

Prototype theory posits that there are two basic principles that guide the formation of categories in the human mind: (1) the **principle of cognitive economy**, and (2) the **principle of perceived world structure**. These principles together give rise to the human **categorisation system**.

Principle of cognitive economy

This principle states that an organism, like a human being, attempts to gain as much information as possible about its environment while minimising cognitive effort and resources. This cost-benefit balance drives **category formation**. In other words, rather than storing separate information about every individual stimulus experienced, humans can group similar stimuli into categories, which maintains **economy** in cognitive representation.

Principle of perceived world structure

The world around us has **correlational structure**. For instance, it is a fact about the world that wings most frequently co-occur with feathers and the ability to fly (as in birds), rather than with fur or the ability to breathe underwater. This principle states that humans rely upon correlational structure of this kind in order to form and organise categories.

8.2.2 The categorisation system

These two principles give rise to the human categorisation system. While the principle of cognitive economy has implications for the level of detail or level of inclusiveness with which categories are formed, the principle of correlational structure has implications for the representativeness or prototype structure of the categories formed (Rosch 1977, 1978). Rosch (1978) suggests that this gives rise to a categorisation system that has two dimensions: a horizontal and a vertical dimension. This idea is represented in Figure 8.1.

The vertical dimension relates to the level of inclusiveness of a particular category: the higher up the vertical axis a particular category is, the more inclusive it is. Consider the category DOG in Figure 8.1. Relative to this category, the category MAMMAL is higher up the vertical axis and includes more members than the category DOG. The category MAMMAL is therefore more inclusive than the category DOG. The category COLLIE, however, is lower on the vertical axis and has fewer members; this category is less inclusive than the category DOG. In contrast, the horizontal dimension relates to the category distinctions at the same level of inclusiveness. Hence, while DOG and CAR are distinct categories, they operate at the same level of detail. In the next two subsections, we look in more detail at the evidence for these two dimensions of categorisation.

8.2.3 The vertical dimension

The vertical dimension derives from the discovery by Rosch and her colleagues (Rosch *et al.* 1976) that categories can be distinguished according to **level of inclusiveness**. Inclusiveness relates to what is subsumed within a particular category. As we have seen, the category FURNITURE is more inclusive than the category CHAIR because it includes entities like DESK and TABLE in addition to CHAIR. In turn, CHAIR is more inclusive than ROCKING CHAIR because it includes other types of chairs in addition to rocking chairs. The category ROCKING CHAIR



Figure 8.1 The human categorisation system

| Superordinate level | Basic level | Subordinate level |
|---------------------|-------------|-------------------|
| | CHAIR | KITCHEN CHAIR |
| | | LIVING-ROOM CHAIR |
| FURNITURE | TABLE | KITCHEN TABLE |
| | | DINING-ROOM TABLE |
| | LAMP | FLOOR LAMP |
| | | DESK LAMP |

Table 8.3 Example of a taxonomy used by Rosch *et al.* (1976) in basic-level category research

only includes rocking chairs, and therefore represents the least inclusive level of this category. Rosch and her colleagues found that there is a level of inclusiveness that is **optimal** for human beings in terms of providing optimum cognitive economy. This level of inclusiveness was found to be at the mid-level of detail, between the most inclusive and least inclusive levels: the level associated with categories like CAR, DOG and CHAIR. This level of inclusiveness is called the **basic level**, and categories at this level are called **basic-level categories**. Categories higher up the vertical axis, which provide less detail, are called **superordinate** categories. Those lower down the vertical axis, which provide more detail, are called **subordinate** categories. This is illustrated in Table 8.3.

In a remarkable series of experiments, Rosch found that basic-level categories provided the most inclusive level of detail at which members of a particular category share features in common. In other words, while the superordinate level (e.g. MAMMAL) is the most inclusive level, members of categories at this level of inclusiveness share relatively little in common when compared to members of categories located at the basic level of inclusiveness (e.g. DOG).

Attributes

Rosch *et al.* (1976) found that the basic level is the level at which humans are best able to list a cluster of **common attributes** for a category. To investigate this, Rosch and her colleagues gave subjects 90 seconds to list all the attributes they could think of for each of the individual items listed in a particular taxonomy. Six of the taxonomies used by Rosch *et al.* are presented in Table 8.4. (It is worth pointing out to British English readers that because Rosch's experiments were carried out in the United States, some of the American English expressions may be unfamiliar.)

Table 8.5 lists common attributes found for three of these taxonomies. In the table, lower levels are assumed to have all the attributes listed for higher levels and are therefore not repeated. Table 8.5 illustrates the fact that subjects were only able to provide a minimal number of shared attributes for superordinate categories. In contrast, a large number of attributes were listed as being shared

COGNITIVE LINGUISTICS: AN INTRODUCTION

| Superordinate | Basic level | Subordinates | |
|--------------------|-------------|----------------------|------------------------|
| MUSICAL INSTRUMENT | GUITAR | FOLK GUITAR | CLASSICAL GUITAR |
| | PIANO | GRAND PIANO | UPRIGHT PIANO |
| | DRUM | KETTLE DRUM | BASE DRUM |
| FRUIT | APPLE | DELICIOUS APPLE | MACKINTOSH APPLE |
| | PEACH | FREESTONE PEACH | CLING PEACH |
| | GRAPES | CONCORD GRAPES | GREEN SEEDLESS GRAPES |
| TOOL | HAMMER | BALL-PEEN HAMMER | CLAW HAMMER |
| | SAW | HACK HAND SAW | CROSS-CUTTING HAND SAW |
| | SCREWDRIVER | PHILLIPS SCREWDRIVER | REGULAR SCREWDRIVER |
| CLOTHING | PANTS | LEVIS | DOUBLE KNIT PANTS |
| | SOCKS | KNEE SOCKS | ANKLE SOCKS |
| | SHIRT | DRESS SHIRT | KNIT SHIRT |
| FURNITURE | TABLE | KITCHEN TABLE | DINING-ROOM TABLE |
| | LAMP | FLOOR LAMP | DESK LAMP |
| | CHAIR | KITCHEN CHAIR | LIVING ROOM CHAIR |
| VEHICLE | CAR | SPORTS CAR | FOUR-DOOR SEDAN CAR |
| | BUS | CITY BUS | CROSS-COUNTRY BUS |
| | TRUCK | PICK-UP TRUCK | TRACTOR-TRAILER TRUCK |

Table 8.4 Six of the taxonomies used by Rosch et al. (1976) as stimuli

| Table 8.5 | Examples | of attribute | lists (based | l on Rosch et al. | 1976: appendix I) |
|-----------|----------|--------------|--------------|-------------------|-------------------|
| | 1 | | ` | | 11 / |

| tool | clothing | furniture |
|----------------------|----------------|------------------------------|
| make things | you wear it | no attributes |
| fix things | keeps you warm | CHAIR |
| metal | PANTS | legs |
| SAW | legs | seat |
| handle | buttons | back |
| teeth | belt loops | arms |
| blade | pockets | comfortable |
| sharp | cloth | four legs |
| cuts | two legs | wood |
| edge | LEVIS | holds people – you sit on it |
| wooden handle | blue | KITCHEN CHAIR |
| CROSS-CUTTING | DOUBLE-KNIT | no additional |
| HAND SAW | PANTS | LIVING-ROOM CHAIR |
| used in construction | comfortable | large |
| HACK HAND SAW | stretchy | soft |
| no additional | - | cushion |

by basic-level categories, while just one or two more specific attributes were added for subordinate categories. Hence, while subordinate categories have slightly more attributes, the basic level is the most inclusive level at which there is a cluster of shared attributes.

Motor movements

In this experiment, Rosch *et al.* set out to establish the most inclusive level at which properties of human physical interaction with a category are found to cluster. This experiment also revealed that basic level categories were the most inclusive level at which members of categories share **motor movements**. To demonstrate this, subjects were asked to describe the nature of their physical interaction with the objects listed. It was found that while there are few motor movements common to members of a superordinate category, there are several specific motor movements listed for entities at the basic level, while entities at the subordinate level make use of essentially the same motor movements. This provides further evidence that the basic level is the most inclusive level, this time with respect to common interactional experiences. This is illustrated in Table 8.6.

Similarity of shapes

For this experiment, Rosch *et al.* sought to establish the most inclusive level of categorisation at which shapes of objects in a given category are most similar. In order to investigate this, the researchers collected around 100 images from sources like magazines and books representing each object at each level in the taxonomies listed in Table 8.4. The shapes were scaled to the same size and then superimposed upon one another. Areas of overlap ratios were then measured, which allowed the experimenters to determine the degree of similarity in shape. While objects at the superordinate level are not very similar in terms of shape (compare the outline shapes of car, bus and motorcycle, for example, as instances

| scan |
|-----------------------------|
| |
| turn |
| turn, move back position |
| bend |
| extend-touch |
| bend |
| touch |
| release weight |
| straighten, lean back |
| |
| sink |
| |

Table 8.6 Motor movements for categories at three levels of inclusiveness (based on Rosch *et al.* 1976: appendix II)

of the category VEHICLE), and while objects at the subordinate level are extremely similar, the basic level was shown to the most inclusive level at which object shapes are similar. In other words, the basic level includes a much greater number of instances of a category than the superordinate level (for example, DOG versus COLLIE) that can be identified on the basis of shape similarity.

Identification based on averaged shapes

In a fourth experiment, Rosch and her team devised averaged shapes of particular objects. They did this by overlapping outlines of entities belonging to a particular category. For all points where the two outlines did not coincide, the central point between the two lines was taken. Subjects were then shown the shapes and provided with superordinate, basic-level and subordinate terms to which they were asked to match the shapes. The success rate of matching shapes with superordinate terms was no better than chance, while subjects proved to be equally successful in matching averaged shapes with basic-level and subordinate terms. For example, the superordinate category VEHICLE consisted of overlapped shapes for car, bus and motorcycle, which are significantly different in shape and therefore less recognisable. On the other hand, the basic-level category CAR, represented by overlapping shapes of different types of cars, did not involve significant differences in shape, and was easily identifiable. Again, although there is a greater degree of similarity at the subordinate level, the basic level is more inclusive. The absence of shape similarity at the superordinate level compared to the evident shape similarity at the basic level goes some way towards explaining why the basic level is the optimum categorisation level for the human categorisation system, which is based, among other things, on perceptual similarity.

Cognitive economy versus level of detail

The major finding to emerge from Rosch's research on basic-level categorisation is that this level of categorisation is the most important level for human categorisation because it is the most inclusive and thus most informative level. It is worth emphasising why this should be the case. After all, Rosch *et al.*'s findings seem to show that the subordinate level is at least as informative as the basic level, if not more so, given that it provides more detailed information in addition to the information represented at the basic level. Recall that, when asked to list attributes of CAR and SPORTS CAR, subjects typically listed more attributes for SPORTS CAR than for CAR. This is because the subordinate category SPORTS CAR is likely to be identified with the same attributes as CAR, plus some extra attributes specific to SPORTS CAR.

The reason why the basic level is the most salient level of categorisation relates to the tension between similarity of members of a category and the principle of cognitive economy. While entities at the subordinate level are most alike (rocking chairs have most in common with other rocking chairs), different categories at the subordinate level are also very similar (rocking chairs are pretty similar to kitchen chairs). At the basic level, on the other hand, while there are also similarities within a particular category (all chairs are pretty similar to one another), there are far fewer between-category similarities (a chair is not that similar to a table). To illustrate this point, let's compare and contrast the basic-level and subordinate level categories given in Table 8.7.

Crucially, for a category to achieve cognitive economy (to provide the greatest amount of information at the lowest processing cost), it must share as many common within-category attributes as possible, while maintaining the highest possible level of between-category difference. In intuitive terms, it is easier to spot the differences between a chair and a lamp than between a desk lamp and a floor lamp. This demonstrates why the basic level of categorisation is 'special': it is the level which best reconciles the conflicting demands of cognitive economy. Therefore the basic level is the most informative level of categorisation.

This notion of cognitive economy has been described in terms of **cue valid**ity. According to Rosch (1977: 29) 'cue validity is a probabilistic concept' which predicts that a particular cue – or attribute – becomes more valid or relevant to a given category the more frequently it is associated with members of that category. Conversely, a particular attribute becomes less valid or relevant to a category the more frequently it is associated with members of other categories. Thus 'is used for sitting on' has 'high cue validity' for the category CHAIR, but 'is found in the home' has low cue validity for the category CHAIR because many other different categories of object can be found in the home in addition to chairs.

Cue validity is maximised at the basic level, because basic level categories share the largest number of attributes possible while minimising the extent to which these features are shared by other categories. This means that basic-level categories simultaneously maximise their **inclusiveness** (the vertical dimension) and their **distinctiveness** (the horizontal dimension) which results in optimal cognitive economy by providing a maximally efficient way of representing information about frequently encountered objects.

| Basic level | Subordinate level | |
|-------------|-------------------------------|--|
| TABLE | DINING TABLE KITCHEN TABLE | |
| CHAIR | DINING CHAIR | |
| | LOUNGE CHAIR | |

Table 8.7 Comparison between levels of categorisation

Perceptual salience

It is clear from Rosch's findings that categorisation arises from perceptual stimuli. When we categorise objects, we do so according to various types of sensoryperceptual input, including shape, size, colour and texture, as well as kinaesthetic input representing how we interact physically with objects. Another way of describing the importance of the basic level, then, is by relating it to **perceptual salience**. There are a number of additional lines of evidence that support the position that the basic level represents the most salient level of categorisation.

The basic level appears to be the most abstract (that is, the most inclusive and thus the least specific) level at which it is possible to form a mental **image**. After all, we are unable to form an image of the category FURNITURE without imagining a specific item like a chair or a table: a basic-level object. This is consistent with the finding that averaged shapes cannot be identified at the superordinate level as there are insufficient similarities between entities at this very high level of inclusiveness. This is also consistent with the fact that Rosch's subjects often struggled to list attributes for the superordinate level. You can try this experiment yourself: if you ask a friend to draw you a picture of 'fruit' or 'furniture' they will draw you apples and bananas or tables and chairs. These are all basic-level categories. There is no recognisable or meaningful shape that represents the superordinate level of categorisation.

Based on a picture verification task, Rosch *et al.* (1976) also found that objects are **perceived** as members of basic-level categories more rapidly than as members of superordinate or subordinate categories. In this experiment, subjects heard a word like *chair*. Immediately afterwards, they were presented with a visual image. If the word matched the image, subjects pressed a 'match' response key. If the word did not match the image, they pressed a different response key. This enabled experimenters to measure the reaction times of the subjects. It emerged that subjects were consistently faster at identifying whether an object matched or failed to match a basic level word than they were when verifying images against a superordinate or subordinate level word. This suggests that in terms of perceptual verification, objects are recognised more rapidly as members of basic-level categories than other sorts of categories.

Language acquisition

Rosch *et al.* (1976) found that basic-level terms are among the first concrete nouns to emerge in child language. This investigation was based on a case study of a single child, consisting of weekly two-hour recordings dating from the initial period of language production. All relevant utterances were independently rated by two assessors in order to determine whether they were superor-dinate, basic or subordinate level terms. The study revealed that the individual

noun-like utterances were overwhelmingly situated at the basic level. Rosch *et al.* argued that this finding provided further support for the primacy of the basic level of categorisation.

Basic-level terms in language

The language system itself also reveals the primacy of the basic level in a number of ways. Firstly, basic-level terms are typically monolexemic: comprised of a single word-like unit. This contrasts with terms for subordinate level categories which are often comprised of two or more lexemes - compare chair (basic-level object) with rocking chair (subordinate-level object). Secondly, basic-level terms appear to occur more frequently in language use than superordinate or subordinate level expressions. More speculatively, Rosch (1978) has even suggested basic-level terms may have emerged prior to superordinate- and subordinate-level terms in the process of language evolution. Of course, given that evidence for the primacy of the basic level is so overwhelming, we might wonder why we need the other levels of categorisation at all. In fact, the superordinate and subordinate levels, while they may not be cognitively salient, have extremely useful functions. As Ungerer and Schmid (1996) explain, the superordinate level (for example, VEHICLE) highlights the functional attributes of the category (vehicles are for moving people around), while also performing a collecting function (grouping together categories that are closely linked in our knowledge representation system). Subordinate categories, on the other hand, fulfil a specificity function.

Are basic-level categories universal?

Of course, if we can find evidence for basic-level categories among English speakers, two questions naturally arise. Firstly, do members of all cultures or speech communities categorise in this way? Given that all humans share the same cognitive apparatus, it would be surprising if the answer to this question were 'no'. This being so, the second question that arises is whether the same basic-level categories are evident in all cultures or speech communities. Clearly, this question relates to 'the extent to which structure is "given" by the world versus created by the perceiving organism' (Rosch *et al.* 1976: 429). Put another way:

[B]asic objects for an individual, subculture, or culture must result from *interaction* between potential structure provided by the world and the particular emphases and state of knowledge of the people who are categorizing. However, the environment places constraints on categorizations. (Rosch *et al.* 1976: 430)

It follows that while the environment partly delimits and thus determines the nature of the categories we create, these categories are also partly determined by the nature of the interaction between human experiencers and their environment. This finding, of course, is consonant with the thesis of embodied cognition.

This view of categorisation entails that while the organisation of conceptual categories into basic, superordinate and subordinate levels may be universal, the level at which particular categories appear may not be. This relates not only to cross-linguistic or cross-cultural variation in the broader sense, but is also reflected within a single speech community or culture where acquired specialist knowledge may influence an individual's taxonomy of categories. For instance, Rosch et al. (1976) found that for most of their North American subjects the category AIRPLANE was situated at the basic level. However, for one of their subjects, a former aircraft mechanic, this category was situated at the superordinate level, with specific models of aircraft being situated at the basic level. This reveals how specialist knowledge in a particular field may influence an individual's categorisation system. At the cross-cultural level, the cultural salience of certain objects may result in taxonomic differences. For example, the anthropologist Berlin and his colleagues (1974) investigated plant naming within the Mayan-speaking Tzeltal community in Southern Mexico. They found that in basic naming tasks members of this community most frequently named plants and trees at the (scientific) level of genus or kind (for example, *pine* versus *millom*) rather than at the (scientific) level of class (for example, tree versus grass). When Rosch et al. (1976) asked their North American students to list attributes for TREE, FISH and BIRD as well as subordinate instances of these categories, they found that, on average, the same number of attributes were listed for TREE, FISH and BIRD as for the subordinate examples, suggesting that for many speakers TREE, FISH and BIRD may be recognised as a basic-level category. The differences between the Tzeltal and North American speakers indicates that aspects of culture (for example, familiarity with the natural environment) can affect what 'counts' as the basic level of categorisation from one speech community to another. However, it does not follow from this kind of variation that any category can be located at any level. While our interaction with the world is one determinant of level of categorisation, the world itself provides structure that also partly determines categorisation, an issue to which we now turn.

8.2.4 The horizontal dimension

The horizontal dimension of the categorisation system (recall Figure 8.1) relates in particular to the principle of perceived world structure which we introduced earlier. This principle states that the world is not unstructured, but possesses correlational structure. As Rosch points out, 'wings correlate with

feathers more than fur' (Rosch 1978: 253). In other words, the world does not consist of sets of attributes with an equally probable chance of co-occurring. Instead, the world itself has structure, which provides constraints on the kinds of categories that humans represent within the cognitive system.

One consequence of the existence of correlational structure in the world is that cognitive categories themselves reflect this structure: the category prototype reflects the greater number of correlational features. Recall that categories often exhibit typicality effects, where certain members of the category are judged as 'better' or more representative examples of that category than other members. Members of a category that are judged as highly prototypical (most representative of that category) can be described as category prototypes. This feature of category structure was investigated in a series of experiments reported in Rosch (1975), which established that prototypical members of a category were found to exhibit a large number of attributes common to many members in the category, while less prototypical members were found to exhibit fewer attributes common to other members of the category. In other words, not only do categories exhibit typicality effects (having more or less prototypical members), category members also exhibit family resemblance relations. While for many categories there are no attributes common to all members (not all members of a family are identical in appearance), there is sufficient similarity between members that they can be said to resemble one another to varving degrees (each having some, but not all, features in common).

Goodness-of-example ratings

In order to investigate the prototype structure of categories, Rosch (1975) conducted a series of experiments in which subjects were asked to provide **goodness-of-example** ratings for between fifty and sixty members of each category, based on the extent to which each member was representative of the category. Typically, subjects were provided with a seven-point scale. They were asked to rate a particular member of the category along this scale, with a rating of 1 indicating that the member is highly representative, and a rating of 7 indicating that the entity was not very representative. Presented in Table 8.8 are the highest- and lowest-ranked ten examples for some of the categories rated by American undergraduate students. It is worth observing that the experiments Rosch employed in order to obtain goodness-of-example rating were 'linguistic' experiments. That is, subjects were presented with word lists rather than visual images.

Family resemblance

Rosch argues that prototype structure, as exhibited by goodness-of-example ratings, serves to maximise shared information contained within a category. As

| Rank | BIRD | FRUIT | VEHICLE | FURNITURE | WEAPON |
|---------|-----------------|-------------------|---------------|---------------|-------------|
| Top eig | ht (from more t | o less representa | tive) | | |
| 1 | Robin | Orange | Automobile | Chair | Gun |
| 2 | Sparrow | Apple | Station wagon | Sofa | Pistol |
| 3 | Bluejay | Banana | Truck | Couch | Revolver |
| 4 | Bluebird | Peach | Car | Table | Machine |
| | | | | | gun |
| 5 | Canary | Pear | Bus | Easy chair | Rifle |
| 6 | Blackbird | Apricot | Taxi | Dresser | Switchblade |
| 7 | Dove | Tangerine | Jeep | Rocking chair | Knife |
| 8 | Lark | Plum | Ambulance | Coffee table | Dagger |
| 9 | Swallow | Grapes | Motorcycle | Rocker | Shotgun |
| 10 | Parakeet | Nectarine | Streetcar | Love seat | Sword |
| Bottom | ten (from more | to less represen | tative) | | |
| 10 | Duck | Pawpaw | Rocket | Counter | Words |
| 9 | Peacock | Coconut | Blimp | Clock | Hand |
| 8 | Egret | Avocado | Skates | Drapes | Pipe |
| 7 | Chicken | Pumpkin | Camel | Refrigerator | Rope |
| 6 | Turkey | Tomato | Feet | Picture | Airplane |
| 5 | Ostrich | Nut | Skis | Closet | Foot |
| 4 | Titmouse | Gourd | Skateboard | Vase | Car |
| 3 | Emu | Olive | Wheelbarrow | Ashtray | Screwdriver |
| 2 | Penguin | Pickle | Surfboard | Fan | Glass |
| 1 | Bat | Squash | Elevator | Telephone | Shoes |

Table 8.8 A selection of goodness-of-example ratings (based on Rosch 1975: appendix)

Rosch puts it, 'prototypes appear to be those members of a category that most reflect the redundancy structure of the category as a whole' (Rosch 1978: 260). In other words, the more frequent a particular attribute is among members of a particular category, the more representative it is. The prototype structure of the category reflects this 'redundancy' in terms of repeated attributes across distinct members, or exemplars. This entails that another way of assessing prototype structure is by establishing the set of attributes that a particular entity has (Rosch and Mervis 1975). The more category-relevant attributes a particular entity has, the more representative it is.

In order to investigate this idea, Rosch and Mervis (1975) presented twenty subjects with six categories: FURNITURE, VEHICLE, FRUIT, WEAPON, VEGETABLE and CLOTHING. For each category, the experimenters collected twenty items that were selected to represent the full goodness-of-example scale for each category, from most to least representative. The subjects were each given six items from each category and asked to list all the attributes they could think of for each item. Each attribute then received a score on a scale of 1–20, depending

on how many items in a category that attribute had been listed for: the attributes that were listed most frequently were allocated more points than those listed less frequently. The degree of family resemblance of a particular item (for example, CHAIR in the category FURNITURE) was the sum of the score for each of the attributes listed for that item: the higher the total score, the greater the family resemblance. Rosch and Mervis's findings showed a high degree of correlation between items that received a high score and their goodness-ofexample ratings. Table 8.9 illustrates these ideas by comparing some of the attributes common across the category BIRD against two members of the category: ROBIN (judged to be highly representative) and OSTRICH (judged to be much less representative).

This table illustrates that the number of relevant attributes possessed by a particular category member correlates with how representative that member is judged to be. Robins are judged to be highly prototypical: they possess a large number of attributes found across other members of the BIRD category. Conversely, ostriches, which are judged not to be very good examples of the category BIRD, are found to have considerably fewer of the common attributes found among members of the category. Therefore, while OSTRICH and ROBIN are representative to different degrees, they nonetheless share a number of attributes and thus exhibit a degree of family resemblance. The claim that category members are related by family resemblance relations rather than by necessary and sufficient conditions entails that categories are predicted to have fuzzy boundaries. In other words, we expect to reach a point at which, due to the absence of a significant number of shared characteristics, it becomes unclear whether a given entity can be judged as a member of a given category or not.

| Attributes | ROBIN | OSTRICH | |
|---------------------------|-------|---------|--|
| lays eggs | yes | yes | |
| beak | yes | yes | |
| two wings | yes | yes | |
| two legs | yes | yes | |
| feathers | yes | yes | |
| small | yes | no | |
| can fly | yes | no | |
| chirps/sings | yes | no | |
| thin/short legs | yes | no | |
| short tail | yes | no | |
| short neck | yes | no | |
| moves on the ground by | yes | no | |
| hopping | | | |

Table 8.9 Comparison of some attributes for ROBIN and OSTRICH

8.2.5 Problems with prototype theory

As we noted at the outset of this chapter, it has been argued that prototype theory is inadequate as a theory of knowledge representation. In this section, we briefly review some of the objections, as well as consider whether Rosch and her colleagues intended their findings to be interpreted directly as a model of knowledge representation.

We begin with a number of criticisms discussed by Laurence and Margolis (1999), who present a survey of the criticisms that have been levelled against prototype theory in the literature. The first criticism, which Laurence and Margolis describe as **the problem of prototypical primes**, concerns the study of ODD NUMBERS that we discussed earlier (Amstrong *et al.* 1983). Recall that this study found that even a 'classical category' of this nature exhibits typicality effects. Armstrong *et al.* argue that this poses potentially serious problems for Prototype Theory since such effects are not predicted for classical categories.

The second criticism that Laurence and Margolis identify is that, like the classical theory, prototype theory also suffers from the problem of ignorance and error: it fails to explain how we can possess a concept while not knowing or being mistaken about its properties. The basis of this criticism is that a concept with prototype structure might incorrectly include an instance that is not in fact a member of that category. The example that Laurence and Margolis use to illustrate this point is that of a prototypical GRANDMOTHER, who is elderly with grey hair and glasses. According to this model, any elderly grey-haired woman with glasses might be incorrectly predicted to be a member of this category. Conversely, concepts with a prototype structure may incorrectly exclude instances that fail to display any of the attributes that characterise the prototype (for example, a cat is still a cat without having any of the prototypical attributes of a cat).

The third criticism that Laurence and Margolis discuss is called **the missing prototypes problem**: the fact that it is not possible to describe a prototype for some categories. These categories include 'unsubstantiated' (non-existent) categories like US MONARCH and heterogeneous categories like OBJECTS THAT WEIGH MORE THAN A GRAM. In other words, the fact that we can describe and understand such categories suggests that they have meaning, yet prototype theory as a model of knowledge representation fails to account for such categories.

Finally, Laurence and Margolis describe the problem of compositionality, which was put forward by Fodor and Lepore (1996). This is the criticism that prototype theory provides no adequate explanation for the fact that complex categories do not reflect prototypical features of the concepts that contribute to them. To illustrate this point, Laurence and Margolis cite Fodor and Lepore's example of PET FISH. If a prototypical PET is fluffy and affectionate and a prototypical FISH is grey in colour and medium-sized (like a mackerel), this does not predict that a prototypical PET FISH is small and orange rather than medium, grey, fluffy and affectionate.

As this brief discussion of the criticisms levelled against prototype theory indicates, Rosch's findings have often been interpreted directly as a theory of knowledge representation (a theory about the structure of categories as they are represented in our minds). Indeed, Rosch explored this idea in her early work (albeit rather speculatively). Consider the following passage:

> [A prototype can be thought of] as the abstract representation of a category, or as those category members to which subjects compare items when judging category membership, or as the internal structure of the category defined by subjects' judgments of the degree to which members fit their 'idea' or 'image' of the category. (Rosch and Mervis 1975: 575)

Rosch retreats from this position in her later writings. As she later makes explicit, 'The fact that prototypicality is reliably rated and is correlated with category structure does not have clear implications for particular processing models nor for a theory of cognitive representations of categories' (Rosch 1978: 261). In other words, while typicality effects are 'real' in the sense that they are empirical findings, it does not follow that these findings can be directly 'translated' into a theory of how categories are represented in the human mind. In other words, experiments that investigate typicality effects only investigate the categorisation judgements that people make rather than the cognitive representations that give rise to these judgements.

This point is central to Lakoff's (1987) discussion of Rosch's findings. Lakoff argues that it is mistaken to equate prototype or typicality effects with cognitive representations. Rather, typicality effects are 'surface phenomena'. This means that they are a consequence of complex mental models that combine to give rise to typicality effects in a number of ways. Typicality effects might therefore be described in intuitive terms as a superficial 'symptom' of the way our minds work, rather than a direct reflection of cognitive organisation. Lakoff (1987) therefore attempts to develop a theory of cognitive models that might plausibly explain the typicality effects uncovered by Rosch and her colleagues. As we will see in the next section, Lakoff's theory of cognitive models avoids the problems that we summarised above which follow from assuming Prototype Theory as a model of knowledge representation.

8.3 The theory of idealised cognitive models

In his book, *Women, Fire And Dangerous Things* (1987), George Lakoff set out to develop a theory of category structure at the cognitive level that could

account for the empirical findings presented by Rosch and her colleagues. This theory was called the **theory of idealised cognitive models**, and represented one of the early frameworks that helped define cognitive semantics as a research programme.

Lakoff argued that categories relate to idealised cognitive models (ICMs). These are relatively stable mental representations that represent theories about the world. In this respect, ICMs are similar to Fillmore's notion of frames, since both relate to relatively complex knowledge structures. While ICMs are rich in detail, they are 'idealised' because they abstract across a range of experiences rather than representing specific instances of a given experience. In Lakoff's theory, ICMs guide cognitive processes like categorisation and reasoning. For example, Barsalou (1983) argues that 'ad hoc' categories like WHAT TO TAKE FROM ONE'S HOME DURING A FIRE also exhibit typicality effects. Lakoff argues that categories of this kind, which are constructed 'online' for local reasoning, are constructed on the basis of pre-existing ICMs. In other words, faced with a house fire, our ability to construct a category of items to be saved relies on pre-existing knowledge relating to the monetary and sentimental value attached to various entities, together with knowledge of the whereabouts in the house they are, the amount of time likely to be available and so on. In the next two subsections, we look in more detail at the properties of ICMs.

8.3.1 Sources of typicality effects

Lakoff argues that typicality effects can arise in a range of ways from a number of different sources. In this section, we present some of the ICMs proposed by Lakoff, and show how these are argued to give rise to typicality effects.

The simplest type of typicality effects

Typicality effects can arise due to mismatches between ICMs against which particular concepts are understood. To illustrate, consider the ICM to which the concept BACHELOR relates. This ICM is likely to include information relating to a monogamous society, the institution of marriage and a standard marriageable age. It is with respect to this ICM, Lakoff argues, that the notion of BACHELOR is understood. Furthermore, because the background frame defined by an ICM is idealised, it may only partially match up with other cognitive models, and this is what gives rise to typicality effects. Consider the Pope, who is judged to be a poor example of the category BACHELOR. An individual's status as a bachelor is an 'all or nothing' affair, because this notion is understood with respect to the legal institution of MARRIAGE: the moment the marriage vows have been taken, a bachelor ceases to be a bachelor. The concept POPE, on the other hand, is primarily understood with respect to the ICM of the CATHOLIC CHURCH whose clergy are unable to marry. Clearly, there is a mismatch between these two cognitive models: in the ICM against which BACHELOR is understood, the Pope is 'strictly speaking' a bachelor because he is unmarried. However, the Pope is not a prototypical bachelor precisely because the Pope is understood with respect to a CATHOLIC CHURCH ICM in which marriage of Catholic clergy is prohibited.

Typicality effects due to cluster models

According to Lakoff, there is a second way in which typicality effects can arise. This relates to **cluster models**, which are models consisting of a number of converging ICMs. The converging models collectively give rise to a complex cluster, which 'is psychologically more complex than the models taken individually' (Lakoff 1987: 74). Lakoff illustrates this type of cognitive model with the example of the category MOTHER, which he suggests is structured by a cluster model consisting of a number of different MOTHER **subcategories**. These are listed below.

- 1. THE BIRTH MODEL: a mother is the person who gives birth to the child.
- 2. THE GENETIC MODEL: a mother is the person who provides the genetic material for the child.
- 3. THE NURTURANCE MODEL: a mother is the person who brings up and looks after the child.
- 4. THE MARITAL MODEL: a mother is married to the child's father.
- 5. THE GENEALOGICAL MODEL: a mother is a particular female ancestor.

While the category MOTHER is a composite of these distinct sub-models, Lakoff argues that we can, and often do, invoke the individual models that contribute to the larger cluster model. The following examples reveal that we can employ different models for MOTHER in stipulating what counts as a 'real mother' (Lakoff 1987: 75).

- (1) a. BIRTH MODEL I was adopted and I don't know who my real mother is.
 - b. NURTURANCE MODEL I am not a nurturant person, so I don't think I could ever be a real mother to my child.
 - c. GENETIC MODEL My real mother died when I was an embryo, and I was later frozen and implanted in the womb of the woman who gave birth to me.

d. BIRTH MODEL

I had a genetic mother who contributed the egg that was planted in the womb of my real mother, who gave birth to me and raised me.

e. BIRTH MODEL

By genetic engineering, the genes in the egg my father's sperm fertilised were spliced together from genes in the eggs of twenty different women. I wouldn't call any of them my real mother. My real mother is the woman who bore me, even though I don't have any single genetic mother.

Lakoff argues that cluster models give rise to typicality effects when one of the ICMs that contributes to the cluster is viewed as primary. This results in the other subcategories being ranked as less important: 'When the cluster of models that jointly characterize a concept diverge, there is still a strong pull to view one as the most important' (Lakoff 1987: 75). This is reflected in dictionary definitions, for example, which often privilege one of the MOTHER sub-models over the others. Although many dictionaries treat the BIRTH MODEL as primary, Lakoff found that *Funk and Wagnall's Standard Dictionary* selected the NURTURANCE MODEL while the *American College Dictionary* chose the GENEALOGICAL MODEL.

Typicality effects due to metonymy

Lakoff argues that a third kind of typicality effect arises when an **exemplar** (an individual instance) stands for an entire category. The phenomenon whereby one conceptual entity stands for another is called **metonymy** and is explored in much more detail in the next chapter. To illustrate metonymy consider example (2):

(2) Downing Street refused comment.

In this example, the official residence of the British Prime Minister stands for the Prime Minister. In other words, it is the Prime Minister (or his or her press officer) who refuses to comment. Similarly, in example (3) it is the vehicle owner who is standing for the car.

(3) I'm parked out the back.

A metonymic ICM can be a subcategory, as in the case of one of the subcategories of a cluster model, or an individual member of a category that comes to stand for the category as a whole. An important consequence of this is that the metonymic model, by standing for the whole category, serves as a **cognitive** **reference point**, setting up norms and expectations against which other members of the category are evaluated and assessed. It follows that metonymic ICMs give rise to typicality effects, as other members of the category are judged as atypical relative to the metonymic model.

An example of a metonymic ICM is the cultural stereotype HOUSEWIFE-MOTHER, in which a married woman does not have paid work but stays at home and looks after the house and family. The HOUSEWIFE-MOTHER stereotype can give rise to typicality effects when it stands for, or represents, the category MOTHER as a whole. Typicality effects arise from resulting expectations associated with members of the category MOTHER. According to the HOUSEWIFE-MOTHER stereotype, mothers nurture their children, and in order to do this they stay at home and take care of them. A WORKING MOTHER, by contrast, is not simply a mother who has a job, but also one who does not stay at home to look after her children. Hence the HOUSEWIFE-MOTHER model, by metonymically representing the category MOTHER as a whole, serves in part to define other instances of the category such as WORKING MOTHER, which thus emerges as a non-prototypical member of the category.

Lakoff proposes a number of different kinds of metonymic models, any of which can in principle serve as a cognitive reference point and can thus give rise to typicality effects. We briefly outline some of these below.

Social stereotypes

The HOUSEWIFE-MOTHER model is an example of a social stereotype. These are conscious ICMs which emerge from public discussion. Against this background, we can re-evaluate the category BACHELOR. The stereotypical bachelor in our culture is a womaniser who lacks domestic skills. Typicality effects can arise if a particular bachelor contrasts with this stereotype. For instance, an unmarried man with one sexual partner who enjoys staying at home cooking and takes pride in his housework may be judged atypical with respect to the social stereotype for bachelors. This shows how the social stereotype BACHE-LOR, which represents one element in the category BACHELOR, can stand for the category as a whole thus giving rise to typicality effects.

Typical examples

Typicality effects can also arise in relation to **typical examples** of a particular category. For instance, in some cultures ROBIN and SPARROW are typical members of the category BIRD. This is because in some parts of the world these birds are very common. In this respect, our environment has consequences for what we judge as good examples of a category. Furthermore, Lakoff argues that we may evaluate a member of the category bird with respect to a typical example. In this way, typicality effects arise when the typical example stands for the entire category.

Ideals

Lakoff suggests that some categories are understood in terms of **ideals**, which may contrast with typical or stereotypical instances. For example, we might have an ideal for the category POLITICIAN: someone who is public-spirited, altruistic, hardworking and so on. This may contrast with our stereotype of politicians as egotistical, power-hungry and obsessed with 'spin'. Once more, typicality effects occur when the ideal stands metonymically for the entire category. For instance, with respect to our ideal the utterance *He's a great politician* might be interpreted as a positive evaluation. However, with respect to our social stereotype, the same utterance would be interpreted as a negative evaluation.

Paragons

Individual category members that represent ideals are **paragons**. For instance, David Beckham, arguably the world's best-known soccer star, is good-looking, a committed father, glamorous, married to a pop star and captain of the England team, as well as being one of the world's most successful footballers. For many people around the world, Beckham represents a FOOTBALL paragon. Similarly, Rolls-Royce represents a paragon in terms of LUXURY CARS, Nelson Mandela represents a paragon in terms of POLITICAL LEADERS, Winston Churchill in terms of WAR LEADERS, Noam Chomsky in terms of GENERATIVE LINGUISTS, and so on. Because paragons stand for an entire category, they set up norms and expectations against which other members of the category may be evaluated. For instance, the comment, 'He's no Nelson Mandela' about a particular political leader may represent a negative assessment as to the leader's altruism and so forth. In this way, paragons give rise to typicality effects.

Generators

According to Lakoff, members of some categories are 'generated' by a core subset of category members called **generators**. These generators are judged to be more prototypical than the other category members that they generate. For example, the natural numbers are represented by the set of integers between zero and nine, which are combined in various ways in order to produce higher natural numbers. For instance, the number 10 combines the integers 1 and 0. Thus the entire category NATURAL NUMBERS is generated from a small subset of single-digit integers. Lakoff argues that this is why the numbers 1 to 9 are judged as prototypical members of the category NATURAL NUMBERS than much larger numbers. Another example of a set of generators is Morse Code. In this system the generators are the 'dot' and the 'dash'. While the 'dot' represents the letter 'E', the 'dash' represents the letter 'T'. Because all other letters represent combinations of dots and/or dashes, the 'letters' 'E' and 'T' are likely to be more prototypical than the others for regular Morse Code users.

| Stereotypes | represent cultural norms and expectations regarding instances of the category |
|------------------|---|
| Typical examples | represent the most frequent or commonly encountered instances of the category |
| Ideals | combine the ideal properties of the category |
| Paragons | represent actual instances of an ideal |
| Generators | members of a category are 'generated' by a core subset of members |
| Salient examples | represent memorable or well-known actual instances of a category |

Table 8.10 Summary of some metonymic ICMs

Salient examples

Finally, memorable or salient examples can also give rise to a type of metonymic ICM. For instance, Oxford University is a salient example of a university, in part due to its history (it received its royal charter in the thirteenth century), in part due to the esteem in which its teaching and scholarship have traditionally been held and in part due to the nature of the colleges that make up the university, both in terms of the structure of the institution and its architecture. Although in many ways atypical in terms of British and other international higher education institutions, people, particularly in the United Kingdom, often rely upon Oxford as a point of comparison for other universities. Typicality effects occur when Oxford serves to establish a means of evaluating and assessing another university.

In other words, salient examples, like prototypes in general, provide cognitive reference points that not only structure a category metonymically, but can influence the decisions we make, for instance whether we decide to go to a particular university based on how similar it is to a salient example like Oxford. Table 8.10 provides a summary of some of the types of metonymic ICMs proposed by Lakoff.

In sum, Lakoff argues that cluster models and metonymic ICMs can give rise to typicality effects in different ways. While the cluster model provides a converging cluster of cognitive models which gives rise to typicality effects by ranking one of the subcategories as more important than the others in the cluster, a metonymic model can stand for the category as a whole and gives rise to typicality effects by defining cultural expectations relating to this category. We will look in more detail at metonymy in Chapter 9.

8.3.2 Radial categories as a further source of typicality effects

Lakoff proposes that the cluster model for MOTHER and the metonymic HOUSE-WIFE-MOTHER stereotype taken together contribute to a **composite prototype** for MOTHER: a prototype derived from two models. This prototype provides **representative structure** for the category. For example, the composite prototype for the category MOTHER includes a female who gave birth to the child, was supplier of 50 per cent of the genetic material, stayed at home in order to nurture the child, is married to the child's father, is one generation older than the child and is also the child's legal guardian. In other words, the composite prototype draws upon information from the BIRTH MODEL, the GENETIC MODEL, the NUR-TURANCE MODEL, the MARITAL MODEL, the GENEALOGICAL MODEL and the HOUSEWIFE MODEL, which is a social stereotype. This type of prototype is an idealisation which provides schematic information. Importantly, further models can be derived from this composite prototype. These models include ADOPTIVE MOTHER, FOSTER MOTHER, BIRTH MOTHER and SURROGATE MOTHER. As Lakoff points out:

These variants are not generated from the central model by general rules; instead, they are extended by convention and must be learned one by one. But the extensions are by no means random. The central model determines the possibilities for extensions, together with the possible relations between the central model and the extension models. (Lakoff 1987: 91)

A composite prototype and extensions of this kind are modelled in terms of a radiating lattice structure. The composite prototype is positioned centrally with other subcategories represented as extending from the **central case** (see Figure 8.2).

Crucially, the non-central cases in such radial categories are not strictly predictable from the central case but are cultural products. For instance, the



Figure 8.2 Radial network for the category MOTHER

subcategories of MOTHER listed below are all understood in terms of how they diverge from the central case.

- 1. STEPMOTHER married to the father but didn't supply genetic material or give birth.
- 2. ADOPTIVE MOTHER provides nurturance and is the legal guardian.
- 3. BIRTH MOTHER gave birth and supplied genetic material but put the child up for adoption hence does not nurture the child and has no legal responsibilities.
- 4. FOSTER MOTHER charged by the state to nurture the child but is not the child's legal guardian.
- 5. SURROGATE MOTHER gives birth to the child, typically does not supply the genetic material and has no other obligations to the child.

Thus radial categories of this kind provide a fourth way in which typicality effects can arise. These effects occur when the subcategories are seen to deviate from the composite prototype. Moreover, as particular categories can become more conventionalised than others, different subcategories in a radial category can develop different degrees of prototypicality.

Importantly, radial categories are not 'generators'. The central case does not productively generate new subcategories of the MOTHER category. While the subcategories are **motivated** in the sense that they are licensed by the prototype, this is a consequence of our cultural experience. For instance, the subcategory SURROGATE MOTHER is a consequence of recent achievements in medicine and cultural trends and has appeared in the second half of the twentieth century. In sum, radial categories are motivated, but knowing a prototype does not predict what subcategories will become conventionally adopted in the culture. We will have more to say about radial categories and how they apply to word meaning in Chapter 11.

To summarise this section, we have seen that there are four ways in which Lakoff accounts for typicality effects. The first kind of typicality effect arises from mismatches between ICMs. The second kind of typicality effect arises from more complex cognitive models which Lakoff calls cluster models. These consist of a number of distinct subcategory models. Typicality effects occur when one subcategory is deemed to be more salient than the others. The third kind of typicality effect relates to metonymic ICMs. These are essentially exemplar-based cognitive models in which a particular member of a given category stands for the category as a whole. Assessed with respect to the metonymic models, other members of a category may be evaluated as being atypical. The fourth kind of typicality effect arises from radial categories, in which members of a radial category exhibit degrees of typicality depending on how close to the composite prototype they are.

8.3.3 Addressing the problems with prototype theory

In section 8.2.5, we reviewed a number of problems that have been claimed to undermine the validity of prototype theory as a model of knowledge representation. In this section, we look at how Lakoff's theory of ICMs addresses these problems.

The first problem we saw was the problem of prototypical primes, which relates to the unexpected typicality effects exhibited by 'classical' categories. Lakoff argues that this finding is not problematic for a prototype-based theory of cognitive models, because these effects can be explained by the nature of the cognitive model that underlies them. Recall that the integers 0-9 are generators: they have a privileged place in the category REAL NUMBER precisely because they form the basis of the category. Within this set, there is a submodel EVEN NUMBERS, which consists of numbers that can be divided by 2, and a submodel ODD NUMBERS for those that cannot. Lakoff argues that because a set of generators can metonymically stand for the category or model as a whole, then the generators included in the submodel ODD NUMBERS (the numbers 1, 3, 5, 7, 9) can stand for the entire category. Against this metonymic model, other odd numbers appear to be less representative of the category, resulting in typicality effects. Although the category ODD NUMBER remains a 'classical' category in the sense that it has definite rather than fuzzy boundaries, it still exhibits typicality effects, which Lakoff argues can be accounted for by the theory of cognitive models. Of course, if typicality effects were interpreted as a direct reflection of cognitive representation of categories, the findings of Armstrong et al.'s study would certainly be unexpected. This example goes some way towards explaining why prototype theory cannot be straightforwardly translated into a model of cognitive representation.

The second problem we saw was the problem of ignorance and error. This relates to the idea that it is possible to possess a concept while not knowing or being mistaken about its properties. For example, a concept with prototype structure might incorrectly include an instance that is not in fact a member of that category, or incorrectly exclude instances that are a member of the category but fail to display any of the attributes that characterise the prototype. However, this problem only arises on the assumption that typicality effects are equivalent to cognitive representation. In other words, tendencies to categorise elderly women with grey hair and spectacles as members of the category GRANDMOTHER (when they might not be) or the failure to categorise sprightly blonde women as members of the category GRANDMOTHER (when they might not be) arise from the social stereotype for GRANDMOTHER which can stand for the category as a whole. In Lakoff's model, this is only one ICM among several for the category GRANDMOTHER, which means that both 'correct' and 'incorrect' instances of categorisation can be accounted for. Equally, it is possible to

possess the concept WHALE while believing it is an instance of the category FISH rather than MAMMAL. Again, this can be accounted for on the basis of metonymic models. A typical property of fish is that they have fins and live in the sea while a typical property of mammals is that they have legs and live on land. Thus, based on the typicality of attributes within the ICM, a whale might be 'miscategorised' as a fish.

The third problem we saw relates to 'missing prototypes'. According to this criticism, it should be possible to describe a prototype for any category we can conceive, yet it is not possible to describe a prototype for 'unsubstantiated' (non-existent) categories like US MONARCH and heterogeneous categories like OBJECTS THAT WEIGH MORE THAN A GRAM. Once more, this problem only arises on the assumption that typicality effects equate to cognitive representation. According to the theory of idealised cognitive models, categories like these are constructed 'on-line' from pre-existing cognitive models, like the 'ad hoc' categories we discussed earlier. Recall that ICMs are relatively stable knowledge structures that are built up on the basis of repeated experience: it is the non-conventional status of non-existent and heterogeneous categories that predicts that such categories would be unlikely to exhibit typicality effects.

The final problem we saw related to compositionality: the criticism that prototype theory fails to provide an adequate explanation for the fact that complex categories do not reflect prototypical features of the concepts that contribute to them. For example, we saw that the category PET FISH does not represent prototypical attributes of the categories PET and FISH. Observe, however, that this criticism assumes that PET FISH is a straightforward composite of the meanings of the two conceptual categories PET and FISH. According to the cognitive model this concept has category structure independently of the two categories to which it is related. In other words, although a pet fish is a type of pet and a type of fish, experience of pet fish gives rise to an independently structured cognitive model in which the prototypical pet fish is the goldfish. The experiential basis of the cognitive model therefore explains why the attributes of this category differ from those of PET and FISH.

8.4 The structure of ICMs

In this section, we explore in more detail the structure of ICMs. So far, we have likened the ICM to Fillmore's notion of a frame and have shown how ICMs can give rise to typicality effects of various kinds. However, we will show that Lakoff's ICMs encompass a wider range of conceptual phenomena than frames and that frames are just one kind of ICM. In Lakoff's theory, ICMs are complex structured systems of knowledge. ICMs structure **mental spaces**: conceptual 'packets' of knowledge constructed during ongoing meaning construction (see Chapter 12). As Lakoff observes, '[a] mental space is a medium

for conceptualization and thought. Thus any fixed or ongoing state of affairs as we conceptualize it is represented by a mental space' (Lakoff 1987: 281). Examples include our understanding of our immediate reality, a hypothetical situation or a past event. In particular, language prompts for the construction of mental spaces in ongoing discourse. The role of ICMs is to provide the background knowledge that can be recruited in order to structure mental spaces. We referred to this process as **schema mapping** in Chapter 5, a process that is also called **schema induction**. According to Lakoff, ICMs depend upon (at least) five sorts of structuring principles for their composition: (1) image schemas; (2) propositions; (3) metaphor; (4) metonymy; and (5) symbolism. We briefly consider each of these structuring principles in turn.

Image schematic ICMs

For Lakoff, a fundamental 'building-block' of conceptual structure is the image schema (recall Chapter 6). Lakoff argues that, in many respects, image schemas serve as the foundation for conceptual structure. He argues that our experience and concepts of SPACE are structured in large part by image schemas like CON-TAINER, SOURCE-PATH-GOAL, PART-WHOLE, UP-DOWN, FRONT-BACK and so on. This means that image schemas like these structure our ICM (or mental model) for SPACE.

Propositional ICMs

Lakoff uses the term 'propositional' in the sense that ICMs of this kind are not structured by 'imaginative devices' (1987: 285) like metaphor and metonymy. Instead, propositional ICMs consist of elements with properties and relations that hold between those elements. An ICM of this kind consists of propositional (or factual) knowledge. For example, our knowledge of the 'rules' involved in requesting a table and ordering food in a restaurant emerges from a propositional ICM. Another sort of propositional ICM might be a taxonomic classification system, for example the biological systems that classify plants and animals.

Metaphoric ICMs

Metaphoric ICMs are structured by the projection or mapping of structure from a source domain to a target domain. For example, when the domain or ICM of LOVE is metaphorically structured in terms of a JOURNEY, as illustrated by expressions like *Their relationship has come a long way*, the ICM for LOVE is metaphorically structured. We return to this subject in more detail in the next chapter.

Metonymic ICMs

We have already examined metonymic ICMs in some detail. As we saw above, ICMs like stereotypes, paragons and ideals are metonymic in the sense that a single type or individual stands for the entire category. We also examine metonymy in more detail in the next chapter.

Symbolic ICMs

ICMs of this kind represent the knowledge structures that Fillmore described in terms of semantic frames. Semantic frames involve lexical items (and grammatical constructions), which cannot be understood independently of the other lexical items relative to which they are understood. Recall the examples of *buy*, *sell* and so on which are understood with respect to the COMMERCIAL EVENT frame that we discussed in the previous chapter. Because this kind of ICM (or semantic frame) is explicitly structured by language (rather than providing a purely conceptual structure that underlies language), its structure contains symbolic units; this is why Lakoff describes it as symbolic.

8.5 Summary

In this chapter we outlined the classical theory of categorisation, which assumes necessary and sufficient conditions, and identified the problems inherent in this approach. We then looked in some detail at prototype theory, the model of categorisation that emerged from research carried out by cognitive psychologist Eleanor Rosch and her colleagues. This research revealed that many categories have prototype structure rather than definitional structure. In addition, Rosch found that categories for concrete objects are most informative at the basic level. However, we saw that assumptions concerning the direct 'translation' of Rosch's findings into a model of knowledge representation gave rise to a number of problems. We then looked at how the empirical findings from this research inspired the development of Lakoff's theory of idealised cognitive models (ICMs). The main claim to emerge from this research was that typicality effects are surface phenomena, arising from underlying ICMs of various kinds. Lakoff argues that prototype structure is not to be directly equated with conceptual structure and organisation, but that typicality effects emerge from three sources: mismatches between ICMs; one subcategory becoming primary in a cluster model; and metonymic ICMs. The latter two types of ICM additionally give rise to radial categories which give rise to a fourth source of typicality effect. Finally, we examined the nature of ICMs in more detail and looked at the various ways in which they are structured. Lakoff argues that ICMs structure mental spaces (entities that serve

as the locus for on-line conceptualisation) by providing the background knowledge that structures these mental spaces. ICMs can be structured in a range of ways. We considered **image schematic ICMs**, **propositional ICMs**, **metaphoric ICMs**, metonymic ICMs and **symbolic ICMs**. We will return immediately to metaphor and metonymy in the next chapter. We return to radial categories in Chapter 10 and to mental spaces in Chapter 11.

Further reading

Prototypes and basic-level categories

- Rosch (1975)
- Rosch (1977)
- Rosch (1978)
- Rosch and Mervis (1975)
- Rosch et al. (1976)

These are among the key articles by Rosch and her collaborators which present their findings concerning prototypes and basic-level categories. The two 1975 papers deal with experimental evidence for prototype effects. The 1976 paper is concerned with basic level categories. The 1977 and 1978 papers provide summaries and overviews of key developments based on the earlier findings. The 1978 paper is particularly important because Rosch explicitly distances herself from earlier suggestions that experimental findings can be considered a direct reflection of cognitive organisation of category structure.

The theory of idealised cognitive models

- Lakoff (1987). While long and sometimes meandering, this book is one of the seminal volumes that sets out the cognitive semantics framework. It introduces and develops the theory of ICMs.
- Taylor (2003). Taylor's book, first published in 1989 and now in its third edition, is an excellent introduction to Rosch's research and the interpretation of these findings within cognitive semantics. Moreover, Taylor elaborates on and extends many of the issues first addressed by Lakoff, particularly as they apply to language.

Other views of categorisation and conceptual organisation

• Komatsu (1992); Laurence and Margolis (1999). Both these articles provide overviews of different approaches to categorization, including prototype theory. These articles are of particular interest

because prototype theory is compared and contrasted with other approaches. The Komatsu article is shorter and more accessible. The Laurence and Margolis volume consists of collected papers by the foremost researchers in the field, including cognitive linguists, formal linguists, philosophers and psychologists.

Exercises

8.1 The classical theory

What are the main claims associated with the classical theory of categorisation? What kinds of problems are inherent in this approach?

8.2 Prototype theory

How is the theory of prototypes and basic level categories different from the classical theory? What do the principles of cognitive economy and perceived world structure contribute to this theory?

8.3 Prototype structure

Try Rosch's experiments for yourself.

(i) List as many attributes as you can for each level of the following taxonomy. What do your findings show?



(ii) List all the motor movements relating to each level of the following taxonomy. What does this experiment reveal?



(iii) Collect judgements from three non-linguists for the following members of the category KITCHEN UTENSIL. Ask them to rank the members on a 1 (good example) to 7 (bad example) scale. Discuss your findings in the light of Rosch's claims.

| pepper-mill |
|---------------|
| plate |
| sink-plunger |
| rolling-pin |
| salad spinner |
| saucepan |
| saucer |
| scales |
| spatula |
| spoon |
| teacup |
| teapot |
| toaster |
| whisk |
| wooden spoon |
| sink plug |
| |

8.4 Idealised cognitive models (ICMs)

What are the ICMs against which the following terms are understood: *bachelor, spinster, boy, girl*? How do these distinct ICMs contribute to the quite different connotations associated with the pairs *bachelor-spinster* and *boy-girl*? (You will need to state first what the common connotations associated with each of these words are.)

8.5 The theory of ICMs

In view of the theory of ICMs, give a detailed account of why the following concepts might be judged as non-prototypical with respect to their corresponding categories. You will first need to state your assumptions about the prototypical attributes associated with the categories in question.

| (a) | STEPFATHER | [category: FATHER] |
|-----|-------------|------------------------------|
| (b) | 977 | [category: CARDINAL NUMBERS] |
| (c) | OSTRICH | [category: BIRD] |
| (d) | TARZAN | [category: BACHELOR] |
| (e) | NORTH KOREA | [category: NATION] |
| | | |

8.6 Radial categories

Consider the category KNIFE. What are the various subcategories associated with this category? What is the prototype? Explain your reasoning.