CONCEPTS, OR: WHAT IS IT THAT A WORD DESIGNATES?

by
John R. Taylor

Saussure asserted that the linguistic sign associates an 'acoustic image' with a 'concept'. Although many linguists speak of 'concepts' with reference to the semantic pole of the linguistic sign, others have been scathing in their critique. While defending a 'conceptualist' semantics, I point to a major problem that a concept-based semantics must address – that of compositionality. I argue that it could be an error to suppose that the concepts that words designate are fixed, stable, and invariant. Concepts are 'emergent'.

1. Problems with Concepts

Probably, anyone with even a rudimentary background in linguistics will be familiar with the little diagram in the *Cours de linguistique générale*, which illustrates Saussure's notion of the 'linguistic sign' (Saussure 1964:99). The diagram shows the association of an 'acoustic image' (what we should nowadays call a 'phonological representation') with a 'concept'. Thus, the sign *tree* associates the phonological representation /tri/ with the concept 'tree' (represented, in the diagram, by a little picture of a tree). The linguistic sign (named 'symbolic unit') is fundamental in Langacker's Cognitive Grammar. Indeed, for Langacker (1987:11), a language is nothing other than a set of symbolic units, each of which associates a phonological representation with a semantic representation.

While broadly subscribing to the Saussurian view of the linguistic sign, Langacker has scrupulously avoided using the word 'concept'. (He prefers the term 'conceptualisation', implying, thereby, a dynamic process, rather than a fixed, stable entity; cf Langacker 1987: 100. The relevance of this distinction will become clear in due course.) Others, however, have been less squeamish. Like Saussure, Sapir, and, more recently, Pinker and Chomsky (amongst many others), have happily used the word 'concept' with reference to the semantic pole of the linguistic sign. Thus, for Sapir, the word *house*

> is the symbol, first and foremost, not of a single perception, nor even of the notion of a particular object, but of a 'concept', in

> other words, of a convenient capsule of thought that embraces thousands of distinct experiences and that is ready to take in thousands more (Sapir 1921:3; my emphasis).

On the meanings of complex expressions, Sapir has this to say:

> If the single significant elements of speech are the symbols of concepts, the actual flow of speech may be interpreted as the setting of these concepts into mutual relations (Sapir 1921:3).

Later in his book *Language*, in a chapter entitled 'Grammatical Concepts', Sapir analyses (1):

(1) The farmer kills the duckling.

as involving three 'concrete' concepts – 'farmer', 'duckling', and the activity of 'killing' – and several 'relational' concepts, such as 'definiteness of reference' (expressed by *the*), 'declarative modality' (expressed by subject-verb order and by *-s*), and concepts of number and time. Sentence (1) is taken to express, in all, no less than thirteen 'convenient capsules of thought', i.e. 'concepts', all set into the appropriate 'mutual relations'.

Almost three quarters of a century later, Pinker expressed a position that, once due allowance is made for the changed role of syntax in grammatical theory, is not radically different from the Sapirean (and Saussurian) view:

> The way language works, then, is that each person's brain contains a lexicon of words and the concepts they stand for (a mental dictionary) and a set of rules that combine the words to convey relationships amongst concepts (a mental grammar) (Pinker 1994: 85).

Chomsky also equated the meanings of lexical items with concepts; indeed, words are but 'labels for concepts'. Moreover, he suggests that the concepts in question might exist prior to the learning of the labels. Consider the following remarks, from the Managua Lectures:

> The speed and precision of vocabulary acquisition leaves no real alternative to the conclusion that the child somehow has the concepts available before experience with language and is basically learning labels for concepts (Chomsky 1988:28; my emphasis).

Having discussed some semantic properties of English *persuade* and Spanish *persuadir*, Chomsky remarks:

> Someone who knows no Spanish at all will know these facts about the word *persuadir*, and the same is true, in essence, of the child learning Spanish – or English or other human languages. The child must have enough information to determine that the form *persuadir* is the one that corresponds to the preexisting concept (Chomsky 1988:32).

This is a remarkable passage, for several reasons. Chomsky's statement about the 'speed and precision' of vocabulary acquisition will certainly astonish anyone familiar with the child language literature. There is ample documentation that children acquiring their mother tongue do not instantaneously determine the semantic content of the words that they take up into their vocabulary (Barrett 1997). Equally, the idea that languages provide labels for pre-existing concepts (something that Saussure vigorously denied) ignores the extraordinary diversity of the semantic values of lexical items in different languages. Indeed, it is rather hard to find convincing examples (*persuade*/*persuadir* could be one of the few!) of words in one language which have exact translation equivalents in another, i.e. which designate exactly the same concepts.

Many might also wish to raise a more basic issue, and query the very notion of 'concept', which linguists from Saussure and Sapir to Chomsky and Pinker, have appealed to in their discussion of linguistic semantics. Kempson (1977) sums up a common view. Taking up Sapir's idea of a concept as a 'convenient capsule of thought', she states that to discuss linguistic meaning in terms of such entities is to put forward claims that are 'entirely untestable and hence vacuous' (pp. 16-7). Since we have no independent means – outside of the linguistic evidence – for studying the properties of these 'capsules of thought', to talk about meaning in terms of concepts:

> does no more that substitute for the problem term *meaning* the equally opaque term *concept*... If meaning is to be explained in

> terms of concepts, it is essential that the term *concept* itself be given a rigorous definition (Kempson 1977:17).

Lyons (1977:113) is also sceptical. While not denying that 'concepts' might be useful constructs in psychological theory, he argues that the use of the term in semantic theory is essentially vacuous. Suppose I say that in order to use the word *table* correctly (i.e. of tables), a person needs to 'have the concept' of table. The evidence that a person does have the concept can only be the person's ability to use the word appropriately. But if I am capable of judging whether or not a person does use the word appropriately, there is no need for me to bring in the concept as a criterion for correct use. Concepts become redundant. In fact, 'having the concept', and being able to use the word appropriately, might turn out to be one and the same thing.

Roy Harris goes further in his critique. According to Harris, Saussure bequeathed to linguistics two 'interconnected errors' – the 'fixed-code fallacy', and the 'telementation fallacy'. Saussurian conceptualism dictates that linguistic knowledge 'is essentially a matter of knowing which words stand for which ideas', i.e., knowing a language involves knowing the set of linguistic signs that constitute the language; furthermore, in virtue of this knowledge, speakers are able to 'transfer thoughts from one mind to another' (Harris 1981: 9). Concepts, however, are invisible to public scrutiny; consequently, there can be no assurance that the concepts one person associates with the words of a language are at all commensurate with the concepts that another person associates with the words. This being so, there can be no 'transfer' of thought from one person to another, in the manner suggested by Saussure's illustration (Saussure 1964:27) of the verbal communication process.

Pursuing these lines of criticism leads inexorably to the neo-behaviourism of Quine, and to a purely observational theory of language. The meaning of a word is its use:

> there is no more to the meaning of an expression than the overt use that we make of the expression. Language is a skill that each of us acquires from his fellows through mutual observation, emulation, and correction in jointly observable circumstances. When we learn the meaning of an expression we learn only what is observable in overt verbal behavior and its circumstances (Quine 1987:130).

Quine, being a philosopher, never actually got down to describe, in any detail, the facts of use that were supposed to fully exhaust the meaning of a word (and which, moreover, would predict the likely uses of the word on future occasions, and speakers' intuitions vis-à-vis hypothetical uses). This is the job of linguists, not philosophers. How, then, have linguistic semanticists proceeded, especially those who reject a conceptualist semantics? Lyons was one such. In place of a conceptual semantics, he proposed a strictly structuralist approach. The meaning of a word was nothing other than the set of (paradigmatic and syntagmatic) relations that a word contracts with other words (Lyons 1968:443; for a similar approach, see also Cruse 1986:16). But this approach inevitably leads to a view of language as a vast calculus of language-internal relations. It ignores the relation of words to world, and of words to a person's conceptualisation of the world. In brief, it rejects the symbolic nature of language.

In order to reassert the symbolic thesis against its critics, we first need to remove a number of red herrings. Harris contends that a symbolic system, if it is to function for the purpose of the communication of thoughts, must be represented identically in the minds of its users. There is clearly no such requirement, as everyday experience shows. In writing this paragraph, I am attempting to get into the consciousness of my reader certain thoughts that I entertain (specifically, about Roy Harris's thoughts on the fixed-code fallacy). Whether I am successful or not depends on many factors, including, obviously, my reader's proficiency in English. But there is certainly no requirement that my reader's knowledge of English exactly coincides with mine, and no linguist, to my knowledge, has ever asserted that this is a requirement to successful communication. (And if my reader is even reasonably proficient in English, this becomes the least important aspect to the success of the communication process!) So much, then, for the fixed-code fallacy. The supposed invisibility of concepts to public verification could also be a red herring. Of course, it is trivially true that I cannot have direct access to another person's inner mental life. But I can certainly believe that another person's mental experience is of the same kind as mine. This is because I have a theory (Fodor 1980) of human beings, their emotions, sentiments, and mental processes. The theory tells me that other humans are constructed in much the same manner as I am, and have minds much the same as mine. This entitles me to attribute to other humans an inner life of concepts and experience,

on the evidence of even minimal overt behaviour (including linguistic behaviour). Lieberman (1991) has argued that the 'uniquely human' ability to empathise was a precondition for the evolution of language. It is empathy that enables us to disregard the 'telementation fallacy', and to consider a symbolic account of language a not implausible undertaking. Also taking an evolutionary perspective, Noble and Davidson (1996) claim that the emergence of symbolic ability in higher mammals was a precondition for the evolution of human language. And with respect to language ontogeny, Ingram (1989: Ch 6) argues that the development of language in the individual builds on an emerging symbolic ability in the young child.

It is less easy to respond to Lyons' objection. Certainly, we evaluate a person's knowledge of a word (and, indeed, of a language) by examining the use they make of the word (and of the language). Irrespective of what a person might tell us about how they understand the concepts designated by the words, provided they use the words appropriately, we do not further enquire into the concepts. (Wittgenstein (1978:182) made just this point.) To this, I would respond that, if we are engaged in merely describing usage, there could indeed be no reason to bring in concepts. However, a conceptualist semantics – as pursued by Langacker, for example – aims to go beyond a description of usage. The aim is to explain (or at least, to motivate) the facts of usage in terms of the conceptualisations that linguistic expressions symbolise. A driving assumption of Cognitive Grammar, is that facts of usage are not arbitrary patterns of learned behaviour, they are motivated (some more transparently than others) by the conceptualisations that they symbolise.

Even if we are inclined to be sympathetic to a symbolic, conceptualist account of word meanings (as I am), there are clearly some major problems that need to be addressed. Two stand out, in particular.

a) A common objection (cf. Kempson 1977:16) is that, even if we accept that some words (*house, table, toothache*) might designate concepts, there are surely many words for which this seems not to be the case at all. What could be the 'concepts' associated with *and, the, of*?

Compare the words *house* and *anyway*. It seems quite natural (*pace* Quine and others) to say that *house* designates a concept, and it seems not unreasonable to attempt a characterisation of this concept. On the other hand, it would be odd to frame an account of the meaning of *house* in terms of how the word is used. (What is the use of the word *house*, if not to designate the house concept?) It is just the other way round with *anyway*. It sounds weird to talk about the 'concept' of *anyway*. On the other hand, it would be quite natural to describe the use of *anyway*. For example, we could say that the word is typically used when a speaker wishes to return to the main topic after a brief digression: *Well, anyway, as I was saying...* . Indeed, the Quinean doctrine, that the meaning of a word is its use, that an account of the use of a word exhausts its semantic description, seems particularly apt for a word like *anyway*.

Clearly, if we are to pursue a full-fledged conceptualist semantics, we will need an account of concepts that is able to accommodate the meanings of all kinds of words, including words like *the, of, and*, and so on.

b) Second, there is the question of compositionality. Wittgenstein (1978:181) asked: 'How do the meanings of the individual words make up the sense of the sentence "I still haven't seen him yet"'? The point of Wittgenstein's question was to suggest that we cannot get at the meaning of a complex expression simply by stringing together, one after the other, the concepts designated by the individual words. The meaning of a complex expression is not a procession of concepts. In order to pursue a conceptualist semantics, we need a syntax of concepts.

How will the 'syntax of concepts' relate to the 'syntax of forms' (i.e. to 'syntax' in the standard linguistic sense)? There is a long tradition in philosophical semantics, which aims to 'translate' the meaning of a sentence into its logical, propositional form. The exercise has been thought useful, to the extent that logical form rarely coincides with syntactic form. The sentences in (2) all express, in different ways, the universally quantified proposition in (3).

(2) a. All cats are carnivores.
    b. Cats are carnivores.

c. Every cat is a carnivore.
d. The cat is a carnivore.
e. A cat is a carnivore.
f. If that's a cat, it's a carnivore.
etc.

(3) ∀x, x(cat) and x(carnivore)

However, the sentences in (2) are by no means synonymous. They each 'construe' the universal assertion in different ways. A basic assumption of conceptualist semantics (as pursued, e.g., in Langacker's Cognitive Grammar), is that the different construals map rather directly onto the surface syntax of the English sentences. That is to say, the words, and their alignment, rather directly symbolise the complex concepts that the sentences in (2) designate.

2. *Defending Concepts*

A first task for a conceptualist semantics, is to clarify the notion of 'concept'.

As a working definition, we can accept what I take to be the prevailing view amongst psychologists, i.e. that a concept is a principle of categorisation (Komatsu 1992). Roughly: To 'have' a concept, is to know what range of entities and states of affairs fall under the concept. There are, to be sure, different and competing models of categorisation, e.g. categorisation by necessary and sufficient conditions, or categorisation by similarity to a prototype, where 'prototype' could be the memory of a specific instance, the characterisation of a good example, or a statistical centre of a category (Smith and Medin 1981). These differences can be ignored for our present purposes. What may be relevant to our topic (though there is insufficient space here to pursue the matter), is the suggestion, developed in Murphy and Medin (1985), that 'concepts' are not just arbitrary, ad hoc collections of entities, but are based in theories ('folk theories') of how the world works. It is our theories that give coherence to our categories, and to our concepts.

It is worth noting that the idea of a concept as a principle of categorisation is not biased towards any word class. Certainly, we tend to think of concepts as essentially nominal; recall Saussure's example

of *tree* as illustration of the linguistic sign. But the proper use of relational terms also entails categorisation, and therefore concepts. Bowerman (1996) has elegantly described the categorisations of spatial relations symbolised by the English prepositions *in/on*, in contrast to Dutch *in/aan/op/om*, Spanish *en*, and Finnish inessive /adessive cases.

Still, the two problems outlined above remain. How can a conceptualist semantics (a) accommodate all the words and morphemes in a language, and (b) account for the combinatorial properties of words and morphemes?

I believe Langacker's Cognitive Grammar offers a framework for satisfying these requirements. Lack of space precludes a full justification of this thesis. (The reader is referred to Langacker's two-volume *Foundations of Cognitive Grammar* (1987, 1991).) But a brief illustration may be in order. Let's take, as a test case, a particularly problematic example. It is not at all obvious that *the* designates a concept. So, how can the concept designated by the complex expression *the house* be derived by the combination of the 'concept' symbolised by *the* and the concept symbolised by *house*?

I need to briefly introduce some aspects of the theory of Cognitive Grammar.

1) profile vs. base. What an expression designates is called the expression's profile. Designation (profiling) takes place against background assumptions; these constitute the expression's base. Thus, *uncle* designates a male human; the kinship relation between the designated person and the nephew(s)/niece(s) is contained within the expression's base. *Uncle* does not designate the kinship relation, even less, the nephew(s) and/or niece(s); the word profiles a male human understood against the base of the (unprofiled) kinship relation.

Can an expression's profile be identified with the concept that the expression symbolises? Clearly not. There is more to the uncle concept than the notion of a male human. The concept resides in the profiling of a male human against the base.

2) domains. Typically, various (more or less coherent) knowledge configurations (or domains) will come together in the base (and constitute an expression's domain matrix). Obviously, the domain of kinship relations is central to an understanding of *uncle*; but kinship

notions, in turn, are based in domains of sexuality, reproduction, marital relations, etc. Human beings are also conceptualised with respect to domains of time and space, bodily functioning, social structures, etc. Very likely, different occasions on which a word is used will differentially activate various aspects of domain-based knowledge, alhough some domains (e.g. kinship in the case of *uncle*) might be so intrinsic that they will be relevant to just about all uses of a word.

Taken together with what was said above on the profile-base relationship, these remarks entail that the concept symbolised by a word need not be a fixed, invariant entity. I take up this notion later.

3) thing vs relation. Nominal expressions designate things; expressions of all other syntactic categories (verbs, prepositions, etc.) designate relations. The related entities may be either things, or themselves relations.

4) schematic vs. contentful. The concept designated by *house* is fairly contentful, in contrast to the concepts designated by *thing, it, he, do,* etc. These are highly schematic in content. Also highly schematic in content are bound morphemes. The plural morpheme, for example, merely designates a plurality of entities, whose identity is fixed outside of the morpheme itself.

5) autonomous vs. dependent. Some concepts can, as it were, stand by themselves, in that one can conceptualise the designated entity without making necessary reference to other entities. Other concepts are inherently dependent, i.e. they necessarily invoke other entities. *House* is (relatively) autonomous, *in* (which invokes a relation between entities) is conceptually dependent. (If there is an in-relation, there has got to be a container, however schematic the container concept may be.) Autonomy vs. dependence also applies to the phonological pole of the linguistic sign. The plural morpheme (in English) is phonologically dependent, in that it cannot be pronounced by itself, but must always attach to a noun stem.

6) valence. Linguistic units can combine, to the extent that a valence relation can be established between them. That is, some component in the structure of the one can be 'unified' with some component in

the structure of the other. The composite expression usually 'inherits' the profile of one of its constituents (traditionally, this is the 'head' of the expression).

The conceptual syntax of *the house* can now be explicated. *House* designates the concept 'house', understood as a type of entity (much as Sapir suggested). *The* designates a specific instance of a type, uniquely identifiable by both speaker and hearer. Its designation is highly schematic; the word is also phonologically dependent. Combining the two symbolic units *the* and *house*, we get a complex unit that designates a specific instance of the type 'house'. (Note that the complex expression inherits the profile of *the*; that is to say, the determiner is the head of the complex expression.)

Kempson had argued that words such as *the*, though meaningful, obviously do not designate concepts, hence that a full-blown conceptual semantics is implausible. In order to rebut the argument, we need a fully articulated conceptual-semantic theory. The brief discussion above has shown, I trust, how Cognitive Grammar sets out to meet this challenge.

3. Compositionality

If words designate concepts, complex expressions (phrases, sentences, texts), we may say, designate 'complex' concepts. The above example showed how the complex concept 'the house' can be assembled from its component parts. Compositionality – the computation of the meanings of complex expressions from the meanings of their components – would appear to be a sine-qua-non of linguistic creativity. The very fact that speakers can, and do, construct complex expressions, many of them never before encountered, and the fact that hearers are well able to fully comprehend these utterances, absolutely requires, or so it would seem, that compositionality be an inherent 'design feature' of human language. The individual words have got to make a predictable contribution to the meaning of the whole. On the other hand, it is a notable fact that the meanings of complex expressions are rarely, if ever, strictly compositional. Typically, the meaning of a complex expression is much more specific than (or even at variance with) the meaning that would result from combining the meanings of its parts.

Take the expression *the football under the table* (cf. Langacker 1987:281). In terms of strict compositionality, the expression locates the football in a region in space that is lower than the table. Very many spatial configurations match this description. Perhaps the football is squashed underneath an overturned table, or the football is suspended from a leg of the table, which in turn is suspended from a hovering helicopter. We would almost certainly not entertain these possibilities. We would probably interpret the expression to mean that the table was in its canonical position, standing on its legs on the floor, and that the football was located on the floor within the area circumscribed by the table's legs. This image is not linguistically encoded by the expression. Indeed, strictly speaking, the preferred interpretation of *the football under the table* actually conflicts with the compositional meaning, in that the football is not actually 'under' the table at all – no part of the football is located 'under' the lowest part of the table, i.e. under the bottom of the table's legs.

Many scholars have drawn attention to this state of affairs, and have been troubled by it. In their different ways, Relevance Theory, and Bierwisch's 'two-level' model of semantics, are largely driven by problems which arise from (the absence of strict) compositionality.

Consider, first, Relevance Theory (Sperber and Wilson 1986). Fundamental to Relevance Theory is the notion that a person's representation of the world is in the form of propositions. Propositions have to be 'semantically complete'; only thus can a proposition be true or false (p. 72). However, many (most?) sentences that speakers utter express logical forms that are 'non-propositional', i.e. are semantically incomplete. These sentences can be compatible with 'an indefinite range of propositional forms' (p. 188). Consequently, their semantics needs to be 'enriched'. Enrichment goes beyond matters of referent assignment and disambiguation. Even after *bat* in *Peter's bat is grey* has been disambiguated ('mammal' vs. 'hitting instrument'), the precise nature of the possessive relation (is it the bat owned by Peter?, the one he killed?, the one he chose?) still needs to be fixed. The whole point of Relevance Theory, as I read it, is to propose mechanisms for 'closing the gap' between 'semantic representations and propositional forms' (p. 188).

The 'two-level' theory was anticipated by John Searle, and focuses more on the meanings of individual words. (For a summary and a critique of the theory, see Taylor 1994.) Consider the verb *cut* in the following examples (Searle 1980:221).

(4) a. Bill cut the grass.
b. The barber cut Tom's hair.
c. Sally cut the cake.
d. I just cut my skin.
e. The tailor cut the cloth.

These sentences exemplify quite ordinary, literal uses of the verb *cut*. Note, though, that the cutting activity is different in each case. (You do not, for example, cut the cake by running the lawn mower over it; and there is a big difference between cutting your hair and cutting your finger). In other words, *cut* makes different contributions to the truth conditions of the five sentences. Assuming strict compositionality, we should have to say that the verb is five-ways ambiguous. For Searle, this would be absurd, for we could end up with as many different meanings of *cut* as there are things to cut, and ways to cut them. Searle therefore assumes that *cut* makes exactly the same semantic contribution to each of the sentences. The sentences are then interpreted relative to knowledge of how the world is, and how people interact with the world. (This knowledge – which need not be 'propositional' – constitutes what Searle calls 'the Background'.) Searle is therefore proposing two levels of semantic representation: a strictly compositional, 'linguistic' level, at which words (barring idiomatic usage) contribute fixed, invariant meanings to the meaning of a complex expression, and a level of interpretation, relative to the Background. It is this second level at which truth conditions are established.

The idea of two levels of meaning – a level of semantic form (the 'linguistic' meaning) and a level of conceptual interpretation, was taken up by Bierwisch and his associates. An early statement of the two-level approach may be found in Bierwisch (1981); a more recent account is in Bierwisch and Schreuder (1992). In contrast to Searle's informal account, the two-level theorists have attempted to formalise the interpretation process.

Consider the case of prepositions. Locational prepositions (in English, German, etc.) have the general semantic form shown in (5) (cf. Lang 1991:129).

(5) $\lambda y\, \lambda x$ [LOC{x. REG(y)}...]

(5) locates a figure object x in a Region determined with reference to a ground object y. ('...' in (5) allows for further specifications, if needed.) Prepositions differ with respect to the characterisation of the function REG. Thus, the semantic form of *in* is as follows. (6) states that a figure object x is located in the Place of the ground object y.

(6) λy λx [LOC{x. PLACE(y)}]

Now consider the following expressions.

(7) a. the water in the vase
b. the crack in the vase
c. the flowers in the vase

Although in each case, it can be said that the figure object (the water, the crack, the flowers) is located in the Place that the vase occupies, the *in*-relation is interpreted differently. In (a), the water is in the hollow internal region of the vase (the vase, moreover, is probably in its canonical, upright position); in (b) the crack is in the material substance of the vase; while in (c) the flowers are held in position by virtue of their stems being 'in' the vase (which is, again, in its upright position). These interpretations emerge on the basis of conceptual knowledge about containers, and about water, cracks, and flowers; they are not actually part of the 'semantic form' of the expressions.

Note that logical properties of *in*-expressions are determined at the level of conceptual interpretation, not at the semantic level. Given (6), *in* should express a transitive relation. Sometimes, this is indeed the case:

The letter is in my briefcase
and
My briefcase is in my car
therefore
The letter is in my car

The following, though, is distinctly odd.

The letter is in my car
and

My car is in the parking lot
therefore
The letter is in the parking lot

Sometimes, transitivity does not obtain at all:

There's a hole in my umbrella
and
My umbrella is in my hand
therefore
There's a hole in my hand

Entailment is normally taken to be a quintessentially linguistic-semantic relation. Yet the above examples show that the entailments of complex expressions may be determined, not at the linguistic-semantic level, but at the level of conceptual interpretation. For some, this could be good enough reason to query the validity of splitting off conceptual interpretation from semantics proper. But if we conflate the linguistic-semantic and the conceptual levels, we must give up, it seems, on compositionality. On the other hand, if we wish to preserve compositionality, we have to reckon with compositionally derived meanings which may be semantically deficient in important respects.

It must be emphasised that the above examples are not rare quirks of usage; the phenomenon investigated under the two-level rubric is endemic in natural language. Practically any word, selected at random from the dictionary, can be associated with different sets of truth conditions, according to its context of use. *Under the tree* looks like a fairly innocuous expression (and *under* and *tree* do not appear to present any inherent problems). Yet if I have a picnic *under the tree*, and if I bury some money *under the tree*, the locations designated by *under the tree* are different. If I say that the underground railway passes *under the tree*, the location is different yet again. In this case, the different interpretations arguably have to do with different interpretations of *tree* rather than *under*. A picnic under the tree takes place on the ground, under the branches and foliage of the tree. To bury something under the tree conceptualises 'tree' as that part that is visible above the ground. The railway that passes under the tree, passes under the tree conceptualised in its entirety – branches, trunk, roots, and all.

4. The Building-Block metaphor

The principle of semantic compositionality rests upon the assumption that the entities that get combined ('concepts', or 'word meanings') are fixed, stable, and invariant. This assumption almost has the status of a logical necessity. (How can we have 'complex concepts' unless there are invariant building blocks out of which these complex concepts are formed?) It is an assumption that also underlies quite a lot of work in lexical semantics. The idea is that not only are the meanings of complex expressions assembled from the meanings of their component words, but word meanings themselves are constructed out of simpler meanings, i.e. 'semantic features', or 'semantic primitives'. One of the best-known of such attempts was the feature theory of Katz and Postal (1964).

As Fodor (1980:289) has pointed out, there are two parts to any programme of 'conceptual reduction': one must (a) 'exhibit' the reductions, i.e. show that they are feasible, and (b) show that the reductions 'play a role in the psychological processes that mediate the use of language'. Fodor claims that for the vast majority of words, definitions in terms of simpler elements are just not possible. (There typically remains a hard core of counterexamples to even the most elaborate definitions.) And even in cases where semantic decomposition does seem to be doable, there are compelling psycholinguistic reasons for rejecting such accounts. For example, there is no evidence that words whose meanings are allegedly 'more complex' take longer to process than words that are featurally less complex (Fodor et al. 1980). Moreover, children do not acquire concepts from the most general, i.e. from those which contain the fewest features (such as 'animate creature', 'inanimate object'), down to the most particular, which contain the largest number of features. Instead of working through the Porphyrian tree from top to bottom, so to speak, they first zero in on names for 'basic level' objects – *ball, dog, baby*, or even names for specific examples of basic level categories. Featurally, these concepts have to be amongst the most complex (Taylor 1995:252-3).

Although few linguists, nowadays, would subscribe to semantic analysis à la Katz and Postal, the explication of complex concepts in terms of simpler ones has driven much of Wierzbicka's work over the past decades. It is axiomatic, in Wierzbicka's approach, that complex

meanings can be defined in terms of a small, finite set of simple meanings:

> One cannot define all words, because the very idea of 'defining' implies that there is not only something to be defined (a definiendum) but also something to define it with (a definiens, or rather, a set of 'definienses'.
>
> The elements which can be used to define the meaning of words (or any other meanings) cannot be defined themselves; rather, they must be accepted as 'indefinibilia', that is, as semantic primes, *in terms of which all complex meanings can be coherently represented.* (Wierzbicka 1996:9-10; emphasis added).

On Wierzbicka's programme, the identity of the primitives is open to empirical discovery. Primitives are arrived at, not through introspection, but by trial-and-error definitions. She makes the strong claim that all words *can* – as a matter of empirical fact – be defined in terms of about 55 primitives. She also hypothesises that these primitives will be lexicalised in all languages.

Arguably, therefore, Wierzbicka has gone some way towards satisfying the first of Fodor's requirements, i.e. that conceptual reduction be 'exhibited'. What is open to question is whether her definitions are psychologically relevant to language users, i.e. whether speakers of a language really do construct complex concepts out of simpler components.

Acquisition is again a relevant testing ground. Wierzbicka's theory predicts that complex concepts can only be acquired on the back of the innate and universal semantic primitives (cf. Wierzbicka 1996:17); were it to be otherwise, the very rationale for definitions in terms of simpler concepts falls away. If a person lacks the building blocks – the 'definienses' – they obviously cannot construct the complex meanings – the 'definienda'. Wierzbicka concedes the point, but in a (for her!) uncharacteristically non-assertive way. (Observe the hedges: *seems reasonable to conjecture, on the whole.*)

> It seems reasonable to conjecture that children absorb the semantic universe of their native language gradually, moving, on the whole, from simpler concepts to more complex ones. (Wierzbicka 1996:221)

The fact is, the first words that children acquire are nothing at all like the semantic primitives that Wierzbicka proposes. Children's first words are not at all words like *think, know, after, I, under, because, if.* On the contrary, the first words (e.g. words for basic level concepts) are words that, on Wierzbicka's account, are semantically very complex. For example, her definition of *mouse* (a word that is likely to be acquired quite early) extends over almost two pages (pp. 340-41). (If the complex concepts used in the definition were themselves to be explicated in terms of the primitives, the definition would be many times longer!) Consider, as another example, the definition of *cloud* (p. 220). 'Cloud' includes, in its definition, the concepts 'sun' and 'sky'. The definition of *sun*, in turn, includes 'sky'. This means, that the 'simpler' concept 'sky' has got to be acquired before 'sun', and 'sun' before 'cloud'. Likewise, 'head' must come before 'face', and 'face' before 'eyes' (p. 219). Surprisingly, Wierzbicka rejects the implications of her analyses for language acquisition:

> [I]f we define 'eyes' in terms of 'face', or 'sun' in terms of 'sky', this does not mean that we expect children to learn the word *face* before the word *eyes*, or the word *sky* before the word *sun*, because the acquisition of concepts is one thing, and the acquisition of words, another (pp. 221).

If this passage is intended to mean that the mere use, by a child, of a word of the adult language does not entitle us to infer that the child has acquired the adult concept, we can certainly go along with the distinction that Wierzbicka proposes. Certainly, the acquisition of a word may be a lengthy process, associated along the way with all manner of overextensions (which are easy to detect) and underextensions (which are not so easy to detect!). However, on a conceptualist semantics, under- and overextensions of use would be taken to be symptomatic of (and even, the consequence of) conceptual under- and overextensions (vis-à-vis adult norms). But if a child does use a word (such as *sun, eyes,* or *cloud*) appropriately, in accordance with adult norms, there can be no reason to query that an appropriate concept has been acquired. If Wierzbicka's semantic analyses are correct, a child who uses *sun, eyes,* etc. appropriately, has got to have acquired the component concepts. There is no way, it seems to me, of getting round the very clear, and very strong predictions made by the theory of conceptual primitives. On the

other hand, Wierzbicka's distinction between the plane of concepts and the plane of word meanings, could be little more than a ploy to protect the theory of semantic primitives against potential counter-evidence. And the counterevidence is, I suggest, massive. Acquisition sequence fails to correlate even remotely with the supposed complexity of concepts. This must cast doubt on the theoretical status of semantic reductionism, and on the very notion of compositionality at the level of word meanings.

5. Emergent Concepts

The building-block metaphor is so dominant in our thinking about language that many people believe that it's got to be correct, and are sceptical of any attempts to dethrone it. The alternative metaphor that I want to propose is that of a network; linguistic expressions provide access to a network of knowledge, which is selectively accessed as the occasion requires. Acquisition – a topic I touched on in the foregoing section – can now be seen in a different light. Acquisition is not a process of building up a concept from its constituent parts; it consists in the gradual elaboration of a knowledge network.

Wierzbicka's definition of *mouse* is instructive in this respect. The definition, as mentioned, extends over almost two pages, and includes such information as the characteristic size, shape, and colour of mice, their habitat, their manner of moving, and the sounds they make. Also included is the fact that mice are (or are thought of as being) timid, quiet, and inconspicuous; that cats chase them; that they are fond of eating cheese; that they live near humans; that they are regarded as pests; and that people try to get rid of them. (Surprisingly, the fact that some people – stereotypically women – have a phobic terror of mice, is not included in the definition.) Wierzbicka motivates the contents of her definition largely on linguistic grounds. A cat can be 'a good mouser'; poor quality Cheddar cheese can be called (or used to be called) 'mousetrap cheese'; a shy, timid, and inconspicuous person (usually female) can be called 'a mouse' (or 'a grey mouse'); and so on. Idioms ('as poor as a church mouse') and nursery rhymes ('Three blind mice') are also called in evidence. Excluded from the definition are 'encyclopedic' facts about mice that are not reflected in everyday linguistic usage, such as their

geographical distribution, the length of the gestation period of the female mouse, the size of the mouse litter, and such like.

The definition brings together a good deal of folk knowledge about mice. But we need to ask, what is the theoretical (and practical) status of the definition? It is clearly not the case that the complete definition can be slotted, *in toto*, into any sentence that contains the word *mouse*, as a replacement for the word. This would be to suppose that all parts of the definition are equally relevant, and equally salient, on any occasion on which the word *mouse* is used. If I describe a person as a 'grey mouse', I do not wish to convey that they are partial to Cheddar cheese, or that they are chased by cats! Rather, what Wierzbicka's definition has captured, is a network of knowledge about mice, which may be selectively accessed on any occasion on which the word is used. Rather than designating a stable, invariant mouse-concept, the word *mouse* provides access to this knowledge network.

Moreover, the 'mouse-network' is not a clearly demarcated region in conceptual space, distinct from all others. The mouse network overlaps and intersects with other networks. If, for Wierzbicka, 'cat' goes into the mouse concept, then 'mouse' surely must go into the cat concept! If the mouse network contains knowledge of how mice typically move, then 'mouse' must be part of the network associated with the verb *scuttle*. If part of the mouse network is knowledge that mice inhabit places where humans live, and that they are unwelcome there, then *infested* has 'mice' as part of its semantic network.

It is important, then, to emphasise that Wierzbicka's definition of *mouse* is an artefact, obtained by summing over very many conventionalised uses of the word. (Other styles of definitions – e.g. minimalist definitions – are also artefacts, of course, obtained by abstracting what is identified as common to a range of uses.) The definition summarises a body of knowledge that is shared (or believed to be shared) by speakers of English, which enables a person to construct a mouse concept that is appropriate to any given use of the word. This is not to say that *mouse* is multiply polysemous. Nor is this a recipe for complete indeterminacy of meaning. On the contrary, we may suppose that certain concepts, or clusters of very similar concepts, have been so frequently evoked in association with certain words, that they constitute ready-made conceptual gestalts, available as 'default concepts', e.g. in a zero context, or in a context which is

minimally constraining. Still, the possibility exists for the construction of ad hoc concepts, appropriate to novel contexts of use.

There have been attempts to formalise the 'emergent' aspect of word meaning – most notably, by Pustejovsky (1991). Pustejovsky proposed that a semantic representation of object names incorporates information that 'structures our basic knowledge about the object' (p. 427), such as what the object is made of, its purpose or function, how it is created, or comes about. This knowledge is differentially activated by context, especially in interaction with a verb meaning. *Read a novel, write a novel, buy a novel, print a novel,* each activates different components of the noun's semantic structure, foregrounding some, and backgrounding others. Pustejovsky's proposals for semantic structure bear some similarity to Langacker's notion of domain matrix; as pointed out, the domain matrix pertains to various kinds of background knowledge against which profiling takes place. One important difference between Pustejovsky's theory and Langacker's is that Pustejovsky's formalisations rigidly circumscribe the range of interpretations that a noun can elicit. (The 'closed' nature of Wierzbicka's definitions, which I address below, invites a similar conclusion.) Langacker's model, in contrast, emphasises the potential open-endedness of the domain matrix (any bit of idiosyncratic knowledge could in principle be relevant to an expression's interpretation), while allowing at the same time that certain domains might be more intrinsic to the conceptualisation than others.

Wierzbicka's definition, I have suggested, sums over a range of conventionalised uses of the word *mouse*; these include idioms, fixed expressions, typical collocations, standard metaphorical uses, and so on. The set of conventionalised uses is not a fixed quantity, either for an individual speaker (speakers are always encountering new uses of old words, and uses, once learned, may gradually be forgotten), nor, even less, for the speech community. (The notion of a speech community is also, of course, an artefact, which can either sum over the knowledge of individual speakers, or abstract what is presumed to be common to them.) Although Wierzbicka (1996:348) presents her mouse-definition as 'discrete and finite', in contrast to encyclopedic knowledge, which is 'provisional', 'cumulative' and 'inexhaustible', we have to doubt whether there exists a discrete, finite, 'mouse network' at all.

There is another aspect to the 'conventionalised uses' of a word that I want to mention in conclusion. A word not only provides access to a network of conceptual knowledge, it is also an entry point into a network of linguistic knowledge, i.e., knowledge about linguistic expressions. Linguistic knowledge includes knowledge about the syntactic configurations ('constructions') that a word can enter into, as well as knowledge about its collocational preferences. Recall the earlier discussion of *cut*. I doubt whether any speaker of English ever feels the need to compute the meaning of *cut the grass* from the meanings of its parts. *Cut the grass* is an established locution, familiar to just about all users of English, and which names an equally familiar 'practice'. Likewise with *cut the cake, cut one's finger, cut the cloth*, and so on. To 'know' a word, is not only to have access to a conceptual network, it is also to know the kinds of linguistic contexts in which the word is conventionally used. Any attempt to come up with a definition of *cut* that abstracts away from the particularities of its use in context, is just as artefactual as Wierzbicka's cumulative definition of *mouse*.

Paradoxically, we come back to a position not unlike that advocated by Quine, and other sceptics of a conceptualist semantics. The position is not quite the same, however. For Quine, the meaning of a word is its use; conceptualisations associated with its use are irrelevant, because unobservable and unverifiable. I think it is still possible to maintain the symbolic thesis, i.e. the idea that linguistic expressions associate a phonological form with a concept, the latter understood as a principle of categorisation (what counts as a valid instance of this concept?). What is clear, is that concepts need not be compositionally built up from simpler constituents, either at the level of the word, or at the level of the phrase. Concepts 'emerge' against the background of previous experience with the language (which entails knowledge of typical collocations and co-occurrence patterns), against accumulated encyclopedic ('folk') knowledge, in interaction with the exigencies of the communicative situation.

Linguistics Section
University of Otago
Dunedin
New Zealand

References

Barrett, Martyn. 1997. Early lexical development. In: P. Fletcher and B. MacWhinney (eds.), The Handbook of Child Language. Oxford: Blackwell.

Bierwisch, Manfred. 1981. Basic issues in the development of word meaning. In: W. Deutsch (ed.), The Child's Construction of Language. London: Academic Press, 341-387.

Bierwisch, Manfred and R. Schreuder. 1992. From concepts to lexical items. Cognition 42.23-60.

Bowerman, Melissa. 1996. Learning how to structure space for language: a crosslinguistic perspective. In: Paul Bloom, Mary A. Peterson, Lynn Nadel and Merill F. Garrett (eds.), Language and Space. Cambridge, Mass.: MIT Press, 385-436.

Chomsky, Noam. 1988. Language and Problems of Knowledge: The Managua Lectures. Cambridge, Mass.: MIT Press.

Cruse, D. A. 1986. Lexical Semantics. Cambridge: Cambridge University Press.

Fodor, Jerry. 1980. The present status of the innateness controversy. In: Representations: Philosophical Essays on the Foundations of Cognitive Science. Cambridge, Mass.: MIT Press, 257-316.

Fodor, Jerry, M. F. Garrett, E. C. Walker, and C. H. Parkes. 1980. Against definitions. Cognition 8.263-367.

Harris, Roy. 1981. The Language Myth. London: Duckworth.

Ingram, David. 1989. First Language Acquisition: Method, Description and Explanation. Cambridge: Cambridge University Press.

Katz, Jerold and Paul Postal. 1964. An Integrated Theory of Linguistic Descriptions. Cambridge, Mass.: MIT Press.

Kempson, Ruth. 1977. Semantic Theory. Cambridge: Cambridge University Press.

Komatsu, Lloyd. 1992. Recent views of conceptual structure. Psychological Bulletin 112.500-526.

Lang, Ewald. 1991. A two-level approach to projective prepositions. In: Gisa Rauh (ed.), Approaches to Prepositions. Tübingen: Narr, 127-167.

Langacker, Ronald. 1987. Foundations of Cognitive Grammar. Vol 1. Stanford: Stanford University Press.

Langacker, Ronald. 1991. Foundations of Cognitive Grammar. Vol 2. Stanford: Stanford University Press.

Lieberman, Philip. 1991. Uniquely Human: The Evolution of Speech, Thought, and Selfless Behavior. Cambridge, Mass.: Harvard University Press.

Lyons, John. 1968. Introduction to Theoretical Linguistics. Cambridge: Cambridge University Press.

Lyons, John 1977. Semantics. 2 Vols. Cambridge: Cambridge University Press.
Murphy, G. and D. Medin. 1985, The role of theories in conceptual coherence. Psychological Review 92.289-316.
Noble, William and Iain Davidson. 1996. Human Evolution, Language and Mind. A Psychological and Archaeological Inquiry. Cambridge: Cambridge University Press.
Pinker, Stephen. 1994. The Language Instinct: How the Mind creates Language. New York: William Morrow.
Pustejovsky, James. 1991. The generative lexicon. Computational Linguistics 17.409-441.
Quine, Willard Van Orman. 1987. Quiddities: An Intermittently Philosophical Dictionary. Cambridge, Mass.: Harvard University Press.
Sapir, Edward. 1921. Language. New York: Harcourt, Brace and World.
Saussure, Ferdinand de. 1964. Cours de linguistique générale. Paris: Payot.
Searle, John. 1980. The background of meaning. In: J. R. Searle, F. Kiefer and M. Bierwisch (eds.), Speech Act Theory and Pragmatics. Dordrecht: Reidel, 221-232.
Smith, E. E. and D. L. Medin. 1981. Categories and Concepts. Cambridge, Mass.: Harvard University Press.
Sperber, Dan and Deirdre Wilson. 1986. Relevance: Communication and Cognition. Oxford: Blackwell.
Taylor, John R. 1994. The two-level approach to meaning. Linguistische Berichte 149.3-26.
Taylor, John R. 1995. Linguistic Categorization: Prototypes in Linguistic Theory. Oxford: Oxford University Press.
Wierzbicka, Anna. 1996. Semantics: Primes and Universals. Oxford: Oxford University Press.
Wittgenstein, Ludwig. 1978. Philosophical Investigations. Trans. G. E. M. Anscombe. Oxford: Blackwell.