

# *Metaphor, ad hoc concepts and word meaning - more questions than answers\**

ROBYN CARSTON

---

## **Abstract**

Recent work in relevance-theoretic pragmatics develops the idea that understanding verbal utterances involves processes of ad hoc concept construction. The resulting concepts may be narrower or looser than the lexical concepts which provide the input to the process. Two of the many issues that arise are considered in this paper: (a) the applicability of the idea to the understanding of metaphor, and (b) the extent to which lexical forms are appropriately thought of as encoding concepts.

## **1 Introduction**

A well-established feature of the relevance-theoretic view of verbal comprehension is that concepts lexically encoded in the linguistic expression uttered may be pragmatically strengthened (or narrowed) as part of the process of deriving the intended explicit content of the utterance. Possible cases are given in (1), focussing on the italicised words:

- (1) a. Ann is *happy*.  
b. I want to meet some *bachelors*.  
c. The *birds* wheeled above the waves.

In the case of (1a), the idea is that the encoded concept HAPPY is quite general, covering a wide range of positive states of mind, while the concept communicated and understood in a particular context, represented as HAPPY\*, is considerably narrower and denotes a much more specific mental state; possibilities include a brief intense feeling of joy, a more lasting low-key feeling of general well-being, a sense of satisfaction after the successful completion of a task, and so on. (Note that these descriptions are necessarily circumlocutionary ways of indicating in natural language

---

\*This paper is part of a much longer chapter, on the pragmatics of on-line concept construction, in Carston (2002). Many thanks to Corinne Iten, Rosa Vega-Moreno, Deirdre Wilson and Vladimir Ōegarac for enjoyable discussions on the issues raised in this paper.

the nature of the *atomic* concept at issue.) In the case of (1b), suppose the context is one in which the speaker has made it clear that she wants to settle down and have children, then the relevant concept, BACHELOR\*, is narrower than the encoded concept whose extension is the set of unmarried adult males; ‘eligibility for marriage’ is a crucial component of the derived concept. In the case of (1c), the BIRD\* concept inferred would be confined to sea birds, and so would exclude sparrows, robins, owls, woodpeckers, etc. In each case, the process involves accessing a relevant subset of the information (logical and encyclopaedic) made available by the lexical concept and using this to construct the intended concept, a concept whose extension is a proper subset of the extension of the lexical concept.

In more recent work, it has been recognised that essentially the same process is involved in understanding loose uses of words (see Carston 1996/97, 2002; Sperber & Wilson 1998; Wilson & Sperber 2000/02). That is, relevance-driven ad hoc concept-building may result in a concept whose extension includes entities not in the denotation of the lexicalised concept; that is, the concept may be broader than the encoded one. Again, it is this pragmatically inferred concept that is taken to be a constituent of the proposition the speaker intended to communicate on the particular occasion of use, even though the consequence in a case of loosening is the loss of an element of encoded meaning. Some possible examples are given in (2):

- (2)    a. There is a large *square* of lawn at the back.  
       b.     This steak is *raw*.  
       c. Ken’s a (real) *bachelor*.            [where Ken is legally married]

The area of lawn referred to in (2a) is very unlikely to be truly a square (with four sides equal in length and at right angles to each other); rather it is approximately square, and this holds for many other uses of geometrical terms: a ‘round’ lake, a group of people standing in a ‘straight line’, a person with a ‘triangular’ face, etc. In (2b), the steak, perhaps served in a restaurant, is not really raw but is much less cooked than the speaker wishes; in (2c), Ken is a married man who behaves like a bachelor (of a certain sort), and so on. In each case, the extension of the pragmatically inferred concept, SQUARE\*, RAW\*, BACHELOR\*, is more inclusive in certain respects than that of the lexical concept, SQUARE, RAW, BACHELOR, from which it was derived.

In this paper, I discuss two of the many issues that are raised by this idea of pragmatic concept building. The first concerns the understanding of metaphorical uses of language, which are treated within relevance theory as a kind of loose use, involving no special mechanisms or processes dedicated to the interpretation of metaphors alone or even to figurative language more generally. It follows from the loose use account that the result of understanding a metaphor is that the logical and encyclopaedic information attached to the loosely-used encoded concept is effectively

sorted into two disjoint sets of properties: those which are taken to be part of the intended interpretation and those which are not, and in the set of excluded properties is some logical or defining feature of the encoded concept. The denotation of the new concept (the metaphorical meaning) includes, or at least overlaps with, the denotation of the encoded concept. Given the unified account of narrowing and loosening, briefly discussed above, it also follows that the ad hoc concept, pragmatically fashioned from the encoded one, contributes to the explicit content of the utterance. In the next section, I consider the extent to which this gives us a fully satisfactory account of how we understand metaphors. The second issue, taken up in section 3, concerns some questions about word meaning which are prompted by this relevance-theoretic account of lexical pragmatics, specifically whether or not what is termed ‘conceptual’ lexical encoding in the theory really is conceptual in any sense that meshes with established ideas about the nature of concepts.

## 2 Metaphor: loose use and ad hoc concepts

### 2.1 Where does metaphorical meaning come from?

The account of metaphor in terms of loose use and relevance-driven processing has been, and continues to be, something of a breakthrough in the understanding of metaphor. However, my concern here is to suggest that, as so far expounded, it may be incomplete, perhaps needing to be supplemented by a further cognitive component in order to be fully explanatory.

I will largely confine myself to the simplest of cases, those involving positive attributions of a property which, if taken literally, is inapplicable to its subject, e.g. ‘Mary is a steamroller/robot/icicle/gazelle, etc’. In so doing I avoid certain extra complications, which are irrelevant here, that arise in thinking about cases of ‘metaphorical’ negations, such as ‘Tom isn’t a human being/man/surgeon, etc.’ (see discussion in Carston 2002, chapter 5). So let’s consider how the loose talk story works for the fairly standard, but still mildly evocative, metaphor in (3):

(3) Robert is a bulldozer.

Sperber & Wilson (1986/95, 236) say the following about this example: ‘[it is] a fairly conventional metaphor whose interpretation involves bringing together the encyclopaedic entries for *Robert* and *bulldozer*, which do not normally come together in a subject-predicate relationship. The result will be a wide array of contextual implications, many of which, being contradictory, can be automatically discarded. The relevance of [(3)] will be established by finding a range of contextual effects which can be retained as weak or strong implicatures. Here there is no single strong

implicature that automatically comes to mind, but rather a slightly weaker, less determinate range having to do with Robert's persistence, obstinacy, insensitivity and refusal to be deflected.' On the updated account, on which an ad hoc concept BULLDOZER\* is constructed, the idea is that certain elements of the encyclopaedic entry of BULLDOZER which are highly accessible in the particular context are carried over into the construction of the new concept which, as a constituent of the explicature, warrants the particular implicatures derived.

This seems fine as far as it goes. But consider the sort of properties that are taken to have been communicated as pertaining to Robert: obstinacy, insensitivity, refusal to be deflected, not listening to other people's views, ruthlessness in pursuing his own interests, etc. The questions that need consideration concern where these properties have come from, and how they have been accessed. As indicated above, according to the general account of loose use, they are recovered via the encyclopaedic entry of the concept BULLDOZER. But how? These properties do not actually feature in that entry, since bulldozers (i.e. those large tractor-like machines used for moving earth, rocks, etc) are not persistent and obstinate, nor do they ignore other people's views, refuse to be deflected, or ruthlessly pursue their own goals. Only human beings seem to have psychological properties such as these (taking them literally, as we must). It is difficult to see how any encyclopaedic sorting process can, by itself at least, effect the transition from the property BULLDOZER, which is literally inapplicable to Robert, to a set of attributes that may well be true of him, because none of those attributes are found in the encyclopaedic entry of BULLDOZER. Note that, in the quote above, Sperber & Wilson say the interpretation process involves bringing together the encyclopaedic entries for ROBERT and BULLDOZER. The idea here, I think, is that the first of these plays an important role in constraining the information selected from the second one because it exerts a considerable influence on the accessibility ranking of the information in the BULLDOZER entry. But this does not help with the problem at hand, which is that the required information is just not in there, so is not available to be promoted in accessibility. Nor, presumably, is it (yet) in the encyclopaedic entry for Robert since, if it were, the predication would be pointless and the utterance uninformative. Intuitively, what goes on is that properties of Robert interact somehow with properties of bulldozers thereby making highly accessible some different properties which can be plausibly predicated of Robert; it is the 'somehow' in this description that remains to be cashed out.

Searle (1979/91) and Martinich (1984/91) have made similar observations about the example in (4). None of the properties that we take to be attributed to Sally by this utterance (for instance, extreme emotional reserve, lack of generosity towards other people, etc) are properties of blocks of ice.

(4) Sally is a block of ice.

As Martinich (1984/91, 511) observes, an inferential process from the coldness of blocks of ice to the coldness of Sally trades on an equivocation in the meaning of ‘cold’ and so is not in fact a valid inference. Another way of putting it is to say that the relevant concepts in the encyclopaedic entries of ‘bulldozer’ and ‘block of ice’ have themselves to be taken as used metaphorically, so that what we have are metaphors within metaphors (loose uses within loose uses). However, that doesn’t get us any closer to an explicit account of how the process works, since extending the context further by exploring the encyclopaedic entries of these metaphorically used concepts, at the next level as it were, does not break through the metaphorical web, into the realm of those properties that are being literally attributed to the human subjects in (3) and (4). For instance, it is not the case that the encyclopaedic information associated with the property of efficient land-clearing provides us with the sort of properties that we can understand as being literally attributed to Robert in (3), nor that the information associated with the concepts of coldness and hardness to the physical touch includes the sort of psychological traits that we understand as attributed to Sally by a speaker of (4). Pugmire (1998, 99) puts the point well when he says: ‘a predicate does not project unmodified from a non-metaphorical into a metaphorical context. Iron cannot, *except metaphorically*, be stubborn, persistent, or headstrong’.

The properties literally predicated in (3) and (4) are not only false of their subjects, they are necessarily false; that is, according to our (naive) metaphysical understanding of the universe, the entity denoted by the subject just isn’t eligible, in any situation, for the property denoted by the predicate. By contrast, there are cases, such as those in (5), where the encoded concept being metaphorically used *could* be literally true of the subject, whether it actually is or not. If it is in fact false, then this is merely a contingent matter, there is no category mistake (no violation of our naive metaphysics).

- (5) a. My son is a baby.  
 b. Suzannah is a princess.  
 c. Bob is a soldier/surgeon/butcher/artist/magician.

A person’s son might well be a baby; a female human called Suzannah could be a princess; a male person could be a soldier or a surgeon, etc. Let us suppose, though, that these sentences are uttered in a situation where it is mutually manifest to the speaker and hearer that the property denoted by the concept encoded in the predicate is not true of the person referred to at the time of utterance. In these cases, the general account of interpreting loose uses, in terms of a relevance-driven process of encyclopaedic sorting, is perfectly adequate; for instance, in (5a), properties represented in the encyclopaedic entry for BABY, such as their total dependence on

their carers, their demands for immediate attention, their inability to consider other people's needs, etc. can be understood as literally true of a particular kind of adult person; similarly, certain (stereotypical) properties of actual princesses, such as that they are rich and privileged, and get treated with great care and deference, can be understood as attributed literally to Suzannah in (5b), and so on, *mutatis mutandis*, for the range of predicates in (5c).

So, while it is a general feature of metaphors that the literal meaning of a predicate is not what the speaker intends to communicate, there are two distinct cases: those where the property could (in some circumstance or other) hold of the entity it is predicated of and those where it simply could not because there is a crossing of fundamental type or category boundaries (machines and humans, inorganic matter and humans, material artefacts and emotions). The latter kind seems to raise a problem for an account which turns solely on the recovery of relevant properties from the encyclopaedic information attached to the encoded concept(s). In so far as this is a genuine problem for a loose use account of the metaphors in (3) and (4), it must be all the more so for truly creative cases such as the following:


- (6) a. The fog comes on little cat feet.  
(from Carl Sandberg's poem *Fog*)
- b. Love is the lighthouse and the rescued mariners.  
(from Oskar Davi...o's poem *Hannah*)
- c. Life's but a walking shadow, a poor player that struts and frets his hour upon the stage, and then is heard no more;  
(Shakespeare: *MacBeth* V, v, 24-26)

These are highly evocative metaphors which fulfil the presumption of relevance, not through the communication of a few strong implicatures, but, rather, through a very wide range of weakly communicated implicatures. It is this that gives them their poetic quality (see Sperber & Wilson 1986/95, 199-200, 224, 235-237). Different hearers/readers may entertain different specific implications in accordance with the particularities of their own encyclopaedic knowledge systems (and their imaginative capacities). The question, again, is how we go from information plausibly stored in memory about cat feet to thoughts about fog, or from our knowledge about lighthouses to thoughts about the nature of love.

It might be supposed that these literary metaphors are of a very different nature from the conventional cases in (3) and (4), that they call for a special sort of measured effortful processing (with commensurately greater effects). A distinction is often made between time-limited comprehension and leisurely comprehension (see Gerrig 1989, Gibbs 1994, chapter 3, Pilkington 2000, chapter 4). While this seems a valid distinction (though perhaps better viewed as a continuum), I see no reason to suppose


that different interpretive mechanisms are involved; rather, the expectation of relevance can be set at a higher level when temporal constraints are relaxed, so that more effort may be invested in exploring connections and going beyond the first few accessible levels of information, and the result should be an ever wider array of cognitive effects. The process is not different in kind from that involved in understanding the quite ordinary cases of the bulldozer and the block of ice, in (3) and (4). A complete explanation of how it is that we access the intended properties in these cases, and form the intended but unencoded concepts BULLDOZER\* and BLOCK OF ICE\*, should provide us with all the ingredients needed for a complete explanation of the creative cases.<sup>1</sup>

As discussed elsewhere (Carston 1996/97, 2002), there are two subcases of the loose use of concepts, the ‘pure’ broadening case resulting in a proper subset relation between the denotations of the lexical concept and the ad hoc concept, shown here in (7a), and the combination broadening/narrowing case resulting in a denotational overlap of the lexical concept and the ad hoc concept, shown in (7b). A third possibility, which arises when we look at certain cases of metaphor, is that the denotations of the two concepts do not intersect at all, as shown in (7c).




---

<sup>1</sup> It is now a commonplace in the psychological literature that understanding metaphorically used predicates takes no longer, so is no more effort-demanding, than understanding literal uses. See, for instance, Gerrig (1989) and Gibbs (1994, chapter 3), where many of the experimental findings supporting this position are cited. Gibbs criticises the relevance-theoretic view of ‘metaphor-as-loose-talk’ for ‘incorrectly assuming that metaphors, and other tropes ... obligatorily demand additional cognitive effort to be understood’ (1994, 232). Whatever the validity of this criticism at the time it was made, it is not a correct characterisation of current relevance-theoretic thinking: the interpretation of literal and of loose (including metaphorical) utterances proceeds in the same way (implications are considered in their order of accessibility and the process stops once the expectation of relevance is fulfilled), so the account does not predict that loose (including metaphorical) uses will generally require more processing effort than literal uses. Indeed, it is to be expected on this view that, in appropriate contexts, a metaphorical interpretation of an utterance may be more easily derived than a literal one.



The contingent falsehood examples in (5) are most likely cases of (7b). Take, for instance, an utterance of ‘Bob is a magician’, in a context in which Bob, who is a surveyor by profession, has just rapidly produced a very nice meal for his flatmates out of a few unpromising old cans of food. What is communicated about Bob is that he has done something surprising, that he is quick and dexterous in his actions, that his performance belies the onlookers’ perception, etc. These are all properties that a professional magician is expected to have, so the ad hoc concept *MAGICIAN\** which is constructed in understanding this utterance includes in its denotation quite a few professional magicians, though probably not all (for instance, old magicians who’ve lost their touch), as well as non-magicians like Bob who have the properties in question.

However, when we turn back to the categorial falsehoods, such as (3), (4) and (6), the appropriate picture seems to be that given by (7c). Support for this comes from the earlier observations that, in these cases, the properties that make up the ad hoc concept cannot be recovered from the encyclopaedic (or, indeed, logical) entry of the encoded concept. Consider again *BULLDOZER* and *BULLDOZER\**. The logical entry for the lexicalised concept *BULLDOZER* may contain an inference rule with the output *HEAVY MACHINERY OF A CERTAIN SORT*, or something along those lines, and its encyclopaedic entry includes information about its land-clearing function, its effectiveness in mechanically pushing away large volumes of earth, rocks and other debris, and about its physical appearance, perhaps represented by a visual image of some sort. The components of the ad hoc concept *BULLDOZER\** include representations of the properties mentioned earlier: obstinacy and persistence, insensitivity to other people’s feelings and views, single-mindedness in pursuing personal interests, etc. It’s not clear to me whether a one-off ad hoc concept is usefully thought of as having a logical entry, but for those



people for whom BULLDOZER\* has become a stable element of their conceptual systems, the likely logical entry is HUMAN TEMPERAMENT OF A CERTAIN SORT (this is certainly reflected in the many dictionaries that give a second sense to the word 'bulldozer'). The entities in the world that fall under these two concepts comprise disjoint sets. Another possible source of support for this third picture comes from linguistic semantic change. Among the processes which bring about such change the following three are often distinguished: lexical narrowing, lexical broadening and metaphorical transfer or shift (see Campbell 1998, chapter 10). For instance, the word 'insult' in English originally meant 'to jump on', but, through its repeated metaphorical application to a certain kind of uncomfortable psychological effect, the physical action meaning has been lost. This would seem to indicate that, at a particular time in the history of the language, the two concepts, INSULT (= jump on), encoded by the lexical form 'insult', and INSULT\* (= offend), an ad hoc concept pragmatically inferred from it, referred to distinct categories of behaviour.

The question remains: how is the move from the lexically encoded concept to the ad hoc concept effected in these fundamental category-crossing cases? I have said that the crucial properties of the new concept in such cases are not to be found in the encyclopaedic entries of the lexically encoded concepts, BULLDOZER, ICE, LIGHTHOUSE, etc., but, of course, this information plays an important role in their recovery. Exactly what that role is and what else is required has not yet been made explicit in the relevance-theoretic account of metaphor (or in any other account of on-line metaphor interpretation that I am aware of). I have nothing new to offer here, my primary aim being the modest, albeit rather negative, one of placing the issue on the table as one that must ultimately be addressed. I'll simply mention briefly two lines of thought in the wider literature on metaphor which might be worth considering in this regard.

The first comes from psychological investigations of the general human capacity for making analogies between what may be quite disparate subject domains, a capacity that manifests itself most strikingly in creative thinking, both artistic and scientific, but which is also prevalent in everyday cognitive activity (see, for instance, Holyoak & Thagard 1995, Gentner, Holyoak, & Kokinov 2001). This capacity is not well understood, but one point of agreement seems to be that a crucial element is the (partial) mapping or alignment of the *structures* of the different domains, that is, a matching of dimensions and relations (as opposed to a search for attributes shared by entities in the domains). For example, the scientific analogies between the atom and the solar system, or between sound transmission and waves of water, are based on similarities in the structures of the two systems rather than similarities in the entities themselves (electrons and planets, etc). At the more mundane level, young children can readily perceive relational

correspondences between, say, the domain of bird life and that of human life, aligning bird with person, nest with house, tree with backyard, singing with talking, etc. Once a structural alignment of two domains is assumed, correspondences can be read off from a position in the structure of one domain (the 'source' domain) to the corresponding position in the other domain (the 'target' domain). Perhaps, the capacity to perform this sort of relational mapping is a component of understanding certain verbal metaphors. So, for example, the metaphorical utterance 'Ariel Sharon is a hawk' might involve a structural mapping from the domain of birds to the domain of politicians, with Sharon occupying the same position in the politician domain as the hawk does in the bird domain, so that characteristics such as political aggressiveness, readiness to make preemptive strikes, etc. might be read off from correspondences with the hawk's ferocity (among birds), its quickness in finding a prey and going for the kill. Exactly how such a story might go for our bulldozer case, I am less sure about. I take it that we want a mapping from, in the source domain, the movement of a bulldozer and its effects on the environment it goes through, and, in the target domain, a particular kind of human social behaviour and its effects on the people at whom it is directed. Conceivably, then, particular correspondences can be read off, for instance, between a bulldozer's indiscriminate clearing of material in its path to a person's insensitive overriding of other people's views and feelings.

The second line of thought, sometimes occurring in combination with the first one, is the currently very popular view that there is a large number of preexisting metaphorical schemes, which play a fundamental role in structuring some of our more abstract concepts (such as LIFE, LOVE, MIND, TIME) and which are readily available to the processes of utterance interpretation. Commonly discussed examples are LIFE IS A JOURNEY, THE MIND IS A CONTAINER, TIME IS A MOVING ENTITY, ARGUMENT IS WAR (see Lakoff 1993, Gibbs 1994). One very general scheme proposed by these authors is PSYCHOLOGICAL FORCE IS PHYSICAL FORCE, which enables us to understand psychological phenomena (traits of character, mental processes) in terms of physical (perceptible) phenomena. It may be this that underpins the 'bulldozer' example, or it may be that a much more specific scheme such as PEOPLE ARE MACHINES is involved. Similarly, there might be a metaphorical scheme mapping emotions aroused by other humans, such as love, anger or grief, onto visceral experiences of the physical world (so perhaps underpinning the 'lighthouse and rescued mariners' case). If we really do have such schemes as part of our conceptual makeup, then part of the role of the relevant encyclopaedic information accompanying a metaphorically used lexical concept is to provide a connection with the appropriate metaphorical scheme. For instance, information about the physical movements of bulldozers and their effects instantiates one

side of the mapping between physical forces and psychological characteristics, so that (perhaps) the bulldozer property of indiscriminately clearing away anything in its path maps to the human psychological trait of ignoring other people in the pursuit of one's own goals. Whether such schemes are really employed in this way in understanding metaphorical utterances remains to be seen; the idea has its critics (see, for instance, Glucksberg & McGlone 1999, Keysar et al. 2000), and its wider framework, according to which thought is fundamentally metaphorical, raises many more questions.

A further issue, which should be acknowledged, though it lies well beyond my scope here, is whether an approach in terms of propositional *conceptual representations* (explicatures and implicatures) can ever do full justice to the processes and results of comprehending a metaphor. From a phenomenological perspective, what is striking about so many metaphors is their imagistic quality; for instance, the cat feet of the fog, the lighthouse and the rescue of the shipwrecked sailors. Even in the utterly banal bulldozer case, people report having a mental image of a bulldozer, perhaps two images that merge in a certain way, the one of Robert, the other of a bulldozer, aspects of the two being combined so that we *see* Robert *as* a bulldozer, or, in the case of 'that surgeon was a butcher', we may have a mental image of a figure dressed in surgeon's garb and in an operating theatre, but who is raising a cleaver to hack at the flesh on the table. What relationship there is, if any, between these apparently imagistic mental representations and the sort of representations I have been concentrating on - conceptual, propositional, syntactic - is something that needs to be explicated. For instance, could it be that we derive conceptual representations ('that surgeon was rough/brutal/behaved without appropriate care and skill', etc) through scrutinising the internal image, rather as we might form thoughts through looking at an external picture? If so, this would provide a ready explanation for the open-endedness and variability of metaphor interpretations. See Davies (1983) and Moran (1997) for discussion of 'proposition' theories and 'image' theories of metaphor and possible relations between them.

In many good poetic metaphors, there are other apparently nonpropositional effects that are achieved: qualitative states of mind, such as sensations and feelings, are evoked and these, rather than anything conceptual, may be precisely what the poet is striving for through his creative use of words. Sperber & Wilson (1986/95, 224) suggest that (at least some of) the *affective* effects of poetic metaphors may turn out to be explainable in cognitive conceptual terms, through their account of weak communication which involves a slight increase in the manifestness of a very wide array of weakly manifest assumptions. See Pugmire (1998, chapter 7) for a discussion of whether or not metaphors might 'come down to nascent, inchoate thoughts or to devices for suggesting these' (p.98), and Pilkington (2000, chapters 6-8) for a searching discussion of the extent to

which poetic effects can be adequately captured in entirely cognitive terms.

## 2.2 Metaphors, similes and their explicatures

In the previous subsection, I dwelt on the question of how we arrive at an interpretation of a metaphorical utterance. This question arises both for the original relevance-theoretic account (Sperber & Wilson 1985/86, 1986/95) and for the slightly revised one involving ad hoc concepts. I'll focus now on this relatively recent idea that an ad hoc concept is constructed and functions as a constituent of what is explicitly communicated by a metaphor<sup>2</sup>. Recall that this position simply follows from the idea that metaphor is a kind of loose use, together with the unitary account of the effects of enrichment and loose use. An interesting consequence is that it seems to provide a nice way of reflecting the often discussed difference between a metaphor and its corresponding simile:

- (8) a. Mary is a bulldozer.  
       b. Mary is like a bulldozer.
- (9) a. Love is the lighthouse and the rescued mariners.  
       b. Love is like the lighthouse and the rescued mariners.

There are clearly great similarities in how the members of each of these pairs are understood and in the result of the interpretation process; in fact, the implicatures are probably the same, so, for both (8a) and (8b), Mary is understood to be obstinate, single-minded, insensitive to other people's feelings, etc. However, their overall impact is not identical, the metaphor usually being experienced as somehow more direct and forceful than the simile. Various observations can be made about what might account for this felt

---

<sup>2</sup> Over the past twenty odd years, Sam Glucksberg and his colleagues have been developing their 'class-inclusion' account of nominal predicate metaphors such as 'My job is a jail' (see Glucksberg (2001) for a summing up of the ideas and of the experimental work supporting them). What is asserted by an utterance of such an example is that the entity designated by the subject term (e.g. my job) is a member of a superordinate category referred to by the predicate expression (e.g. the category of things which are confining, externally imposed, difficult to escape from, etc.), a category which includes both prototypical jails and the speaker's job. The category (or concept) thereby created has no established lexical form; it is an ad hoc category, which can be represented as JAIL\*, reflecting the considerable similarity between this view and the relevance-theoretic one. A detailed comparison of the two accounts remains to be done, but see Rubio-Fernandez (2002) for interesting discussion and some suggestive experimental results bearing on the two approaches.

difference. For instance, while the metaphor is literally false, similes are standardly said to be trivially true (since everything is like everything else in some respect or other). I don't think this strikes at the heart of the issue, though, since the corresponding negatives are just as much metaphors and similes, although their truth/falsity status is reversed. It is sometimes said that while a simile invites the addressee to make a comparison between two unlike things, a metaphor requires the addressee to conceive of one thing as actually *being* another (unlike) kind of thing. Whether or not this is right, an account of metaphor understanding which involves the construction of an ad hoc concept from an encoded concept may reflect the felt difference, since similes do not seem to undergo such a process. Compare (10a) and (10b) as possible propositions expressed by the simile in (8b):

- (10) a. MARY IS LIKE A BULLDOZER\*  
 b. MARY IS LIKE A BULLDOZER  
 c. Mary is like a human being  
 d. A pear is like a fruit.

While it makes sense to say that Mary belongs to a particular category of BULLDOZERS\*, there is no more sense in claiming that she is *like* a BULLDOZER\* than that she is like a human being (given that she *is* one). That is, (10a) is odd in much the same way as it would be odd to say that a pear is like a fruit, that a sparrow is like a bird, or that a dog is like a mammal. The oddity lies in saying of an entity which is a member of some category that it is (merely) *like* that category. So while a metaphor and its corresponding simile may communicate the same set of implicatures, the difference between them may be captured by the fact that an ad hoc concept is constructed as part of the explicit content of the metaphor, while the lexically encoded concept is preserved in the simile, as in (10b).

Incorporation of an ad hoc concept in the proposition explicitly communicated by a metaphorical utterance has another interesting consequence. Recall that, according to the relevance-theoretic account, many metaphors involve weak communication; that is, the intention made mutually manifest by the speaker is not (or not only) to make strongly manifest some small number of specific assumptions but rather to make weakly manifest a wide range of assumptions (a conceptual space). In such a case, the implicatures of the utterance are indeterminate: the speaker has not singled out and endorsed any particular assumptions, but has rather encouraged the hearer to explore within the range of activated assumptions. Precisely which ones the hearer does in fact derive as implicatures of the utterance is, in large measure, a matter of his own choice and responsibility. The

more creative or unusual a metaphor, the wider the range of possibilities and the weaker the speaker's endorsement of any specific implicated propositional form. This much has been part of the account since its beginnings in the 1980s. The point here is that, given the view that the hearer constructs an ad hoc concept that replaces the encoded concept in the explicitly communicated propositional form, the characteristic of indeterminacy must carry over from the implicatures of many metaphorical utterances to their explicatures.

All the implicatures derived by the hearer have to be inferentially warranted and the ad hoc concept plays a crucial role in this, since most of the implicated properties are features of its encyclopaedic entry. What we have here is multiple instances of the process of mutual parallel adjustment of explicitly and implicitly communicated assumptions (for discussion and exemplification of this process, see Wilson and Sperber 2000/02). It may be that the time course of interpretation is such that, before he has a complete explicature, the hearer accesses a variety of particular assumptions from the, possibly vast, range activated by the utterance and treats them as potential implicatures of the utterance; if so, there is then a good deal of backwards inference involved in shaping the new concept which will figure in the explicature and ultimately warrant the set of implicated conclusions. The indeterminacy that pertains to implicatures, therefore, pertains equally to the ad hoc concept; that is, the relation between the concept constructed by the hearer and that in the speaker's own thought is one of sufficiently close resemblance rather than identity. There may be quite a range of subtly different concepts licensed by an utterance of, for instance, 'Robert is a bulldozer': BULLDOZER\*, BULLDOZER\*\*, BULLDOZER\*\*\*, etc. No specific one is strongly communicated and the hearer's construction of any one of them is good enough for the communication to have succeeded.

A final question, which I merely raise without attempting to answer, concerns the extent of the process of ad hoc concept formation in metaphor understanding. Most of the examples above have involved single constituents, either a predicate or a referring expression, but, clearly, whole sentences can be used metaphorically, as in the following examples, where (11a) might describe a bad-tempered boss and his cowered employees, (11b) might be a report on the state of a failing institution or company, and (11c) a characterisation of the course of a destructive marriage or friendship:

- (11) a. When the old lion wakes up and starts roaring again we had all better run for cover.
- b. Despite lavish nursing, the patient has yet to leave his sick-bed and take a few tottering steps out of doors.

- c. The buds of hope and love called out by a day or two of sunshine are frozen again and again till the tree is killed.

The proposition expressed by (11a) might involve several ad hoc concepts which compose together with some encoded concepts along the lines in (12):

- (12) WHEN THE OLD [LION]\* [WAKES UP]\* AND STARTS [ROARING]\* AGAIN WE HAD ALL BETTER [RUN FOR COVER]\*

Another possibility is that longer stretches of the encoded conceptual structure, phrases or even the whole logical form, are to be taken as used loosely (metaphorically) and a complex (structured) ad hoc concept pragmatically constructed on that basis:

- (13) [[THE PATIENT] [HAS YET TO [[LEAVE HIS SICK-BED] AND [TAKE A FEW TOTTERING STEPS OUT OF DOORS]]]]\*

Neither of these possibilities seems sustainable in the case of certain wholly metaphorical poems and stories (allegories or parables); rather, they would seem to be given an initial full interpretation (pragmatically disambiguated, enriched, etc) within a frame or (fictional) domain, from which, subsequently or in parallel, their metaphorical interpretation is projected. This is, clearly, an issue that needs a lot more thought.

### 3 Word meaning and concepts

In this section, I leave the issues concerning metaphor interpretation and move to a rather different set of concerns. These are questions about the nature of encoded word meaning which are raised by the general account of lexical pragmatics in terms of on-line concept construction, whether resulting in a narrower or broader concept or some combination of the two. The discussion so far has proceeded on the assumption that lexical items like ‘cat’, ‘sing’, ‘open’, ‘raw’ and ‘happy’ encode (atomic) concepts, where, on a mental representational construal, concepts are ‘words of Mentalese’, that is, constituents of sentences in the language of thought (Fodor 1975, 1998). What I want to consider now is the possibility that this view of word meaning is not right, that such ‘conceptual encodings’ are (in many instances, at least) not really full-fledged concepts, but rather concept schemas, or pointers to a conceptual space, on the basis of which, on *every* occasion of their use, an actual concept (an ingredient of a thought) is pragmatically

inferred. The discussion is, I'm afraid, highly speculative, appealing to intuition rather than providing hard argument.

The idea that natural language *sentences* do not encode (do not translate into) sentences in the language of thought but rather provide a template or schema for constructing such language of thought sentences has been fundamental to relevance theory since its beginnings: 'Linguistically encoded semantic representations are abstract mental structures which must be inferentially enriched before they can be taken to represent anything of interest' (Sperber & Wilson 1986/95, 174). At the same time, however, it has been assumed that, with the exception of indexicals and a few other expressions that encode procedures, most words encode (atomic) concepts, that is, constituents of thought. The suggestion here is that the abstract schematic nature of a decoded 'semantic' representation (or logical form) is manifest not only at the global sentence level (by variables, gaps and missing sections of propositional structure) but also at the local level of the basic elements of the representation. While sentences encode thought/proposition templates, words encode concept templates; it's linguistic underdeterminacy all the way down.

Fodor (1998) takes the position that 'English [and public language systems quite generally] inherits its semantics from the contents of the beliefs, desires, intentions, and so forth that it's used to express' (1998, 9). In Carston (2002, chapter 1), I argued that this does not work for natural language *sentences*, which cannot be assigned a 'real' truth-conditional semantics, even by inheritance, since they drastically underdetermine propositional content. However, I left it open that the view might be true of the *lexical* level: words that encode concepts may inherit a referential semantics from the concepts (CAT, SING, OPEN, RAW, HAPPY, etc.) which they encode (and so activate in an addressee's mind when uttered). One upshot of questioning the position that lexical items encode anything recognisable as a fully fledged concept will be to cast doubt on even this rather meagre claim that word types in natural language can be thought of as having a truth-based semantics by inheritance.

Focussing on the word 'happy', let's consider the concept that it is supposed to encode, a concept which is to provide communicative access to a wide range of other more specific concepts, including one for a steady state of well-being, another for a momentary experience of intense joy, another for the sense of fulfilment that accompanies a successful negotiation, and so on. The idea is that the lexically encoded concept HAPPY is distinct from all of these; it is more general and abstract than any of them, but provides the basis, in appropriate contexts, for processes of pragmatic enrichment so that addressees can come to grasp one of the more specific concepts and incorporate it into their representation of the speaker's thought. But what is not at all clear is whether we



ever actually have (hence sometimes try to communicate) thoughts in which this very general lexicalised concept features as a constituent, or indeed what the property of being HAPPY is, as opposed to being HAPPY\* or HAPPY\*\*, etc. Could it be that the word ‘happy’ does not encode a concept, but rather ‘points’ to a conceptual region, or maps to an address (or node, or gateway, or whatever) in memory? This pointing or mapping provides access to certain bundles of information from which the relevance-constrained processes of pragmatic inference extract or construct the conceptual unit which features in the speaker’s thought.

Sometimes it seems that Sperber & Wilson are envisaging something along these lines: ‘Quite generally, the occurrence of a word in an utterance provides a piece of evidence, a pointer to a concept involved in the speaker’s meaning. ... A verb like ‘open’ acts as a pointer to indefinitely many notions or concepts ...’ Sperber & Wilson (1997/98, 196-97). However, this is selective quoting; in between these two sentences they say: ‘It may happen that *the intended concept is the very one encoded by the word*, which is therefore used in its strictly literal sense’ (my highlighting). But when we try to think about the general concept OPEN and to have a thought in which such a general concept features, as opposed to any of the more specific concepts that we grasp in understanding ‘open one’s mouth’, ‘open the window’, ‘open a can’, ‘open a discussion’, etc. the experience is an odd one, as we seem to have no definite thought at all.

Searle (1983) discusses the fact that many truth-conditionally distinct meanings can be communicated by perfectly literal uses of the unambiguous verb ‘open’:

- (14) a. Jane opened the window.  
 b. Bill opened his mouth.  
 c. Sally opened her book to page 56.  
 d. Mike opened his briefcase.  
 e. Pat opened the curtains.  
 f. The child opened the package.  
 g. The carpenter opened the wall.  
 h. The surgeon opened the wound.

What constitutes opening a book is very different from what constitutes opening one’s mouth, which is quite different again from what constitutes opening a package, etc., so that there are different constituents corresponding to ‘open’ in the proposition expressed

in each of these examples<sup>3</sup>. Furthermore, and more directly relevant to my point here, given our current background assumptions, we are unable to fully understand (assign a determinate truth-conditional content to) the following:

- (15) a. Bob opened the grass.  
       b. Chris opened the fork.  
       c. Jane opened a hair.

According to Searle, although we can grasp the literal meaning of each word in these sentences and their syntax is unproblematic, we do not have any know-how concerning the opening of grass or forks or hairs. One way of construing this in the current context is that the verb ‘open’ points us to a particular region in encyclopaedic memory at which all manner of information about kinds of opening is stored, or at least made accessible, but it does not include what would be needed for us to construct the kind of full-fledged concept that could feature in a thought involving a relation between Bob and the grass, Chris and the fork, etc., a thought which could then be evaluated against some situation or event in the world involving Bob and the grass, Chris and the fork, etc., and judged true or false. Assuming we have referents for ‘Bob’ and ‘the grass’ but no other knowledge about a specific context, then all we get from an utterance of (15a) is that Bob performed some act involving the grass and that that act can be described as one of ‘opening’. Whatever mental representation we form from this contains not a descriptive concept, that is, one with a determinate denotation, but an interpretive, metarepresentational one. But surely, if the word ‘open’ encodes a concept, we should be able to have thoughts which include that concept, and there should be no particular difficulty in grasping the proposition expressed by ‘Bob opened the grass’.

At the least, it seems that the ‘concepts’ encoded by bits of lexical phonological material are a rather different kind of thing from the concepts that feature as ingredients in our thoughts and in interpretations of our utterances. That there is a distinction of this sort to be made is further supported by the general account of how ad hoc concepts are constructed: the lexical form maps to a ‘conceptual’ address in memory to which is

---

<sup>3</sup> Searle’s explanation of these differences is given in terms of what he calls the Background, that is, a set of (mostly unconscious) assumptions, dispositions, and know-how, without which the encoded meaning of a sentence does not determine any definite truth condition. Relative to one set of background assumptions and practices, a sentence may determine one truth condition, relative to another set, it may determine a different truth condition, and, relative to some sets, it may not determine any definite truth condition.

attached a package (or packages) of information; we dip into this package and take out just a part of it. The process is always selective, there is always some subset of the activated information which is left behind or discarded, whether the ultimate upshot is one of a narrowing or a broadening, in the sense discussed earlier. The result is a concept, a constituent of the thought the speaker is communicating; however, the building materials for this process made available by the lexical item (perhaps including several conflicting properties or stereotypes, as in the case of 'bachelor') do not constitute a concept. If this is right, there is a sense in which all concepts are ad hoc, that is, temporary constructs arising for specific purposes at particular times, as Barsalou (1987, 1993) has claimed.

Given the rich range and finely varied nature of the concepts that can be communicated by a single lexical form, it seems we have (at least) the following two options in giving an account of the encoded meaning of a word: multiple ambiguity, with each of the (related) meanings a full-fledged concept, or the very abstract, attenuated, schematic meanings, which come from taking the 'univocality' or 'monosemy' line (see Grice 1978/89, Atlas 1989, Ruhl 1989) and which, I am suggesting, are not concepts. Even setting aside possible economy considerations, reflected in the well known methodological principle of Modified Occam's Razor, the ambiguity approach is ruled out as any sort of general solution since what we're dealing with here is not a fixed set of (related) senses but an apparently indefinite range of possibilities (see Sperber & Wilson 1998, 186). So, even if we decided that 'open' encodes, say, four distinct (but related) literal senses, the pragmatic account of ad hoc concept construction would be required for the myriad other meanings that arise in particular contexts. But once we have an effective account of these cases, the supposedly encoded meanings might as well be derived in the same general way, given an initial pointer to the right area of information in memory. Of course, individual concepts within the (indefinitely large) set of those communicable by 'open' differ greatly in their frequency of occurrence in thought and communication (compare a concept of the usual mode of opening one's mouth with one of an action of prising lips and teeth apart with one's fingers, etc.) so that some are retrievable more or less ready-made while others require a more constructive process. This conception of word meaning and concept communication goes at least some way towards accounting for the phenomenon of polysemy (the multiple related senses of a word).

However, while this view of word meaning as very abstract and schematic is perhaps plausible for a range of cases, such as the verbal and adjectival encodings just discussed, it might be thought to be much less so when we turn to natural kind terms, like 'cat', 'lion', 'water', 'tree'. There is a strong intuition that 'cat' encodes a concept CAT, which

features in thoughts, and not just some abstract schema for constructing CAT\* concepts or some pointer to knowledge about cats. Utterances such as those in (16) do not seem to be susceptible to the indeterminacy of those in (15), that is, we have no difficulty grasping a propositional content, even though we may judge it false and/or bizarre in other ways:

- (16) a. My neighbour's cat likes to read Shakespeare.  
 b. Mr Jones was transformed into a cat for a few hours.

It is noticeable that natural kind terms do not figure much in discussions on polysemy, perhaps a further indication of their stable conceptual content (which is not to deny their high susceptibility to figurative use). So perhaps there are different kinds of lexical meanings, with some words encoding full-fledged concepts, others encoding a schema or a pro-concept (see Sperber & Wilson 1998 on the latter) and others a procedure or inferential constraint.

Suppose it is right that there is a sizeable class of words that do not encode particular concepts (senses), but rather concept schemas, or pointers, or addresses in memory (which of these is the best metaphor remains unclear), what effect does this have on our depiction of the output of the linguistic processing system, the decoded logical form or 'semantic' representation? I think it doesn't make much difference; we just alter our construal of such representations as HAPPY, OPEN, etc., from being concepts, constituents of thought, to being labels for pointers or schemas. This is not a change that brings with it any new problems for a principle of semantic compositionality, conceived of as a principle that holds at the level of encoded linguistic meaning (expression-type meaning) rather than at the level of the proposition expressed by an utterance (see discussion in Carston 2002, chapter 1). However, not too surprisingly, this way of thinking about word meaning does initiate new questions and problems. For instance, the introduction of a whole additional population of mental entities, which are distinct from concepts and don't seem to have any function in mental life except to mediate the word/concept relation, is not to be taken on lightly. On the other hand, it is simply a counterpart, at the lexical semantic level, to what we've already taken on at the sentence semantic level: proposition schemas, which mediate the sentence/thought relation.

One particularly challenging question that arises concerns acquisition: if word meanings are these abstract schematic entities that do not feature in our thinking about the world, how do we ever manage to acquire (learn) them? Consider how the word meaning acquisition story goes on a Fodorian sort of account, according to which learning a natural language is 'learning how to associate its sentences with the

corresponding thoughts. To know English is to know, for example that the form of words ‘there are cats’ is standardly used to express the thought that there are cats; ...’ (Fodor 1998, 9). So learning a word, say ‘cat’ or ‘open’, is learning to associate it with the corresponding concept, CAT or OPEN, on the basis of which the word inherits the semantics of the concept, namely that it refers to the property *cat* or the property *open*, respectively. There are two acquisitional processes here; the acquiring of the concept and the acquiring of the link between lexical form and concept, and giving a satisfactory explanation of these is no simple matter. However, on an account of word meaning according to which it is not a concept but something more like a general schema for building concepts, or a pointer to a range of concepts, the task of explaining acquisition appears to become much more difficult.

How does a child acquire such an abstract entity as a schema or a pointer, something that doesn’t play any role in her thoughts? For instance, suppose she already has several concepts of the OPEN\* variety, which do figure in her thoughts about the world, say one for the action of opening her mouth, another for opening a door, another for opening a carton of milk. On the current hypothesis, none of these concepts is the meaning of the word ‘open’, so she can’t just map the phonological form /open/ onto one of those and thereby acquire the (adult) meaning of ‘open’. Of course, in another sense, these are all (context-dependent) meanings of the word ‘open’, meanings that particular *tokenings* of the word ‘open’ have. So, at different points along the route to establishing a more or less stable meaning for the lexical form, various different form-concept correspondences may be made. The question, which won’t receive any answer here, is how the more general schema or indicator arises and how it comes to be the meaning of the lexical expression *type*. There must be some process of abstraction, or extraction, from the particular concepts associated with the phonological form /open/ to the more general ‘meaning’, which then functions as a gateway both to the existing concepts of opening and to the materials needed to make new OPEN\* concepts which may arise in the understanding of subsequent utterances.<sup>4</sup>

Recanati (1998) refers to work by Hintzman (1986) on a multiple-trace memory model of word meaning acquisition, according to which each experience of a token (an utterance) of a word leaves an ‘episodic’ memory trace. On a new encounter with the

---

<sup>4</sup> This acquisition question arises for anyone who takes seriously the underdeterminacy thesis at the lexical level, such as, for instance, Blutner (2002), who says: ‘Given a polysemous lexeme, its meaning representation may either refer to a primary conceptual variant (representing its base sense), or it may be a more abstract unit referring to some form of underspecified structure’. He favours the latter, more schematic, option.

word, all those accumulated traces are activated, but with different weightings (accessibility rankings) depending on the particular context, so that some subset consisting of the most strongly activated traces gives the meaning of the word on that specific occasion. This leads to the radical claim that there is no lexical meaning in the sense of a stable encoding: ‘Words, as expression-types, do not have ‘meanings’, over and above the collection of token-experiences they are associated with. The only meaning which words have is that which emerges in context’ (Recanati 1998, 630). Whether one takes this extreme line or the more conservative one I am taking (on which words do encode something, albeit something very schematic, which simply sends the system off to a particular region in long-term memory), the multiple-trace model offers a suggestive approach to the question of how we acquire word meanings. On the more conservative view, it would need to be supplemented by an account of how the encoded schema or pointer emerges, or is abstracted, over time from the accumulation of traces of particular cognitive experiences of tokens of the word.

## References

- Atlas, J. (1989) *Philosophy Without Ambiguity*. Oxford: Clarendon Press.
- Barsalou, L. (1987) The instability of graded structure in concepts. In Neisser, U. (ed.) *Concepts and Conceptual Development: Ecological and Intellectual Factors in Categorization*, 101-140. New York: Cambridge University Press.
- Barsalou, L. (1993) Flexibility, structure, and linguistic vagary in concepts: Manifestations of a compositional system of perceptual symbols. In Collins, A. C., Gathercole, S. E. & Conway, M. A. (eds.) *Theories of Memory*, 29-101. London: Lawrence Erlbaum.
- Blutner, R. (2002) Lexical semantics and pragmatics. *Linguistische Berichte Sonderheft 10*, 27-58.
- Campbell, L. (1998) *Historical Linguistics*. Edinburgh: Edinburgh University Press.
- Carston, R. (1996/97) Enrichment and loosening: complementary processes in deriving the proposition expressed? *UCL Working Papers in Linguistics 8*, 205-232. Reprinted 1997 in *Linguistische Berichte 8*, Special Issue on Pragmatics, 103-127.
- Carston, R. (2002) *Thoughts and Utterances: The Pragmatics of Explicit Communication*. Oxford: Blackwell.
- Davies, M. (1983) Idiom and metaphor. *Proceedings of the Aristotelian Society* 83, 67-85.
- Davis, S. (ed.) (1991) *Pragmatics: A Reader*. Oxford: Oxford University Press.
- Fodor, J. (1975) *The Language of Thought*. New York: Thomas Crowell Co.
- Fodor, J. (1998) *Concepts: Where Cognitive Science Went Wrong*. Oxford: Clarendon Press.
- Gentner, D., Holyoak, K. & Kokinov, B. (eds.) (2001) *The Analogical Mind: Perspectives from Cognitive Science*. Cambridge, Mass.: MIT Press.
- Gerrig, R. (1989) Empirical constraints on computational theories of metaphor. *Cognitive Science* 13, 235-241.

- Gibbs, R. (1994) *The Poetics of Mind*. Cambridge: Cambridge University Press.
- Glucksberg, S. (2001) *Understanding Figurative Language: From Metaphors to Idioms*. Oxford University Press.
- Glucksberg, S. & McGlone, M. (1999) When love is not a journey: What metaphors mean. *Journal of Pragmatics* 31, 1541-1558.
- Grice, H.P. (1978/89) Further notes on logic and conversation. In Cole, P. (ed.) *Syntax and Semantics 9: Pragmatics*, 113-127. New York: Academic Press. Reprinted in Grice, H.P. 1989 *Studies in the Way of Words*, 41-57. Cambridge, Mass.: Harvard University Press.
- Hintzman, D. (1986) Schema abstraction, in a multiple trace memory model. *Psychological Review* 93, 411-28.
- Holyoak, K. & Thagard, P. (1995) *Mental Leaps: Analogy in Creative Thought*. Cambridge, Mass: MIT Press.
- Keysar, B., Shen, Y., Glucksberg, S. & Horton, W. (2000) Conventional language; how metaphorical is it? *Journal of Memory and Language* 43, 576-593.
- Lakoff, G. (1993) The contemporary theory of metaphor. In Ortony, A. (ed.) 1993, 202-251.
- Martinich, A. (1984/91) A theory for metaphor. *Journal of Literary Semantics* 13, 35-56. Reprinted in Davis, S. (ed.) 1991, 507-518.
- Moran, R. (1997) Metaphor. In Hale, B. & Wright, C. (eds.) *A Companion to the Philosophy of Language*, 248-68. Oxford: Blackwell.
- Ortony, A. (ed.) (1993) *Metaphor and Thought*, second edition. New York: Cambridge University Press.
- Pilkington, A. (2000) *Poetic Effects: A Relevance Theory Perspective*. Amsterdam: John Benjamins.
- Pugmire, D. (1998) *Rediscovering Emotion*. Edinburgh: Edinburgh University Press.
- Recanati, F. (1998) Pragmatics. *Routledge Encyclopedia of Philosophy*, vol. 7, 620-633.
- Rubio-Fernandez, P. (2002) Understanding nominal metaphors: beyond class inclusion. Typescript. RCEAL, University of Cambridge.
- Ruhl, C. (1989) *On Monosemy: a Study in Linguistic Semantics*. Albany: State University of New York Press.
- Searle, J. (1979/91) Metaphor. In Searle, J. 1979. *Expression and Meaning*, 76-116. Reprinted in Davis, S. (ed.) 1991, 519-537.
- Searle, J. (1983) *Intentionality*. Cambridge: Cambridge University Press.
- Sperber, D. & Wilson, D. (1985/86) Loose talk. *Proceedings of the Aristotelian Society* 86, 153-171. Reprinted in Davis, S. (ed.) 1991, 540-549.
- Sperber, D. & Wilson, D. (1986/95) *Relevance: Communication and Cognition*. Oxford: Blackwell; Cambridge, Mass.: Harvard University Press. Second edition (with postface) 1995.
- Sperber, D. & Wilson, D. (1998) The mapping between the mental and the public lexicon. In Carruthers, P. & Boucher, J. (eds.) 1998. *Language and Thought: Interdisciplinary Themes*, 184-200. Cambridge: Cambridge University Press.
- Wilson, D. & Sperber, D. 2000/02. Truthfulness and relevance. *UCL Working Papers in Linguistics* 12, 215-254. Revised version 2002 in *Mind* 111 (443), 583-632.