

Lexical pragmatics, *ad hoc* concepts and metaphor: A Relevance Theory perspective

Robyn Carston

Ostensive communication, the paradigm case of which is verbal communication, is the domain of a dedicated cognitive system, according to Relevance Theory (RT). This 'pragmatics' module is responsible for inferring the content or meaning that the communicator intends by his/her ostensive stimulus. An important sub-process within the system is the adjustment or modulation of lexically-encoded meaning, which makes it possible for speakers to communicate a vastly greater range of concepts than those that are stably encoded in their linguistic system. This includes the meaning communicated by at least some cases of metaphorically-used language. Taking a broadly Fodorian view that lexical concepts are atomic (unstructured), this paper looks at some issues raised by the idea that addressees infer *ad hoc* concepts as part of the on-line comprehension process. As a cognitive-scientific theory, RT is open to evidence from a range of sources, including native speaker-hearer intuitions, recorded instances of linguistic communication (corpus data) from both communicatively typical and atypical populations, results from relevant psychological and psycholinguistic experiments, and findings in cognitive neuroscience on brain activation during both utterance production and comprehension. At the end of the paper, some tentative suggestions are made of questions arising from work in theoretical lexical pragmatics which might be amenable to investigation at the neuropragmatic level.*

Keywords: utterance understanding, Relevance Theory, lexical semantics, lexical pragmatics, explicature, *ad hoc* concepts

1. Introduction: Relevance Theory as cognitive science

Pragmatics within the Relevance Theory framework is fundamentally cognitive-scientific as distinct from pragmatic theories whose basis is philosophical, sociological or linguistic. The account of communication and comprehension that it aims to provide focuses on the on-line processes of utterance interpretation and the nature of the mental system(s) responsible for them (Sperber & Wilson 1986/1995; Wilson & Sperber 2004).

Relevance Theory (**RT**) has developed in tandem with two important and closely connected ideas in cognitive science: (1) that the mind is modular, and (2) that many mental processes are performed by ‘fast and frugal’ heuristics. Arguments from evolutionary psychology suggest that the mind is massively modular, in the sense that a great many distinct dedicated procedures and processes have evolved to solve specific cognitive problems (Cosmides & Tooby 1994; Sperber 2002).¹ And it seems that, in solving a wide variety of everyday problems, we employ rather simple, albeit ecologically rational, heuristics rather than foolproof algorithms or explicit reasoning processes (Gigerenzer et al. 1999). These mechanisms making up the mind’s ‘adaptive toolbox’ are fast and frugal in that they carry out limited computations and consult just a small salient subset of all the available information, rather than working through myriad possibilities and comparing candidate solutions. A process of this sort can only achieve a high degree of accuracy in a domain to whose particular characteristics it is specifically tailored or adapted.

The relevance-theoretic approach to communication situates pragmatics within this sort of cognitive framework, that is, one which consists of largely domain-specific capacities, each with the function of solving a specific long-standing problem in human mental life and employing quick, relatively cheap computations to do so (Sperber & Wilson 2002; Wilson 2005). The human pragmatic capacity is such a dedicated system: its specific domain is ostensive stimuli (verbal utterances and other acts of ostensive communication) and the comprehension procedure it employs is a fast and frugal heuristic. The procedure can be stated simply and informally as follows: for any ostensive stimulus, assess hypotheses about its interpretation in order of their accessibility and accept the first one that meets current expectations of relevance (where occasion-specific expectations of relevance are regulated by a general presumption of ‘optimal relevance’ conveyed by all instances of ostensive communication). This heuristic is, of course, underpinned by an account of the nature of cognitive relevance (which is, in brief, a positive function of cognitive effects and a negative function of processing effort) and of what it is for an ostensive stimulus to be optimally relevant. For the full story, see Wilson & Sperber (2004).

While ideas in cognitive science have had a formative influence on relevance-theoretic pragmatics from its beginnings, a more recent development is the use of experimental techniques from cognitive psychology and psycholinguistics, which are increasingly bringing rich empirical evidence to bear on aspects of the theory (see, for

instance, the papers in Noveck & Sperber 2004). Reaction time data derived from experiments employing a range of task types, including self-paced reading, sentence verification, acceptability judgements, lexical decision or naming within a cross-modal priming paradigm, among others, have provided valuable evidence about the nature and time-course of utterance comprehension processes. For instance, there is now extensive experimental work testing (and confirming) the RT view of scalar implicatures as context-sensitive and thus effortful inferences versus an alternative view of them as default and automatic (see, in particular, Breheny et al. 2006; Noveck & Sperber 2007). There is also experimental work, using a variety of priming techniques, which shows that the activation and deactivation of conceptual information is significantly different in the pragmatic processes of meaning selection (disambiguation) and meaning construction (accessing metaphorical meaning and enriching lexical concepts). This provides some initial encouragement for the current RT account of metaphor comprehension and of lexical meaning adjustment more generally (Rubio Fernández 2007; 2008), although further detailed testing is needed.

More recently still, insights and techniques from cognitive neuroscience have begun to yield new kinds of data, with the potential for providing a further source of empirical evidence bearing on processing predictions coming from relevance theory and some other pragmatic theories (for useful overviews, see Bambini 2005; Bambini & Bara *forthcoming*). Neuroscientific methodologies, including electrophysiological recordings of neural activity in the brain and a range of brain imaging techniques, are now being employed in studies of language processing, in particular at the level of lexical/syntactic decoding, but also at the level of pragmatic inference.² For instance, ERP studies, which have the advantage of being time-locked to the onset of the stimulus under scrutiny, have already yielded results which are consistent with the RT view that the derivation of scalar implicatures is not an automatic default but an effort-demanding inferential process (see Noveck & Posada 2003 and Chevallier et al. *this issue*). This evidence converges pleasingly with the psychological evidence from the timing studies referred to above.

A property of ERP studies which is potentially illuminating for pragmatics is their ability to reveal the occurrence of distinct kinds of neurophysiological responses during the time-course of a particular cognitive process, thereby providing insights into processing load which are inaccessible to the standard reaction time measures made within many of the psycholinguistic paradigms adopted into experimental

pragmatics. As Coulson (2004: 199) points out, it is possible for two cognitive processes or tasks to take the same amount of time, but for one to recruit more neural processing resources and thus to be more costly and complex than the other. Coulson goes on to outline how this sort of evidence from ERP studies has already been employed in choosing between different accounts of metaphor comprehension versus literal processing, finding in favour of those which deny a qualitative difference between them with regard to the neural resources they deploy. More generally for relevance theory, since processing effort (together with cognitive effects) is a key component in assessing the relevance, hence the 'correctness', of an utterance interpretation, it seems that ERP graphs could take us one step further than temporal measurement of overt behaviour in understanding why certain interpretations of utterances win out over others. Thus, albeit with important caveats (see note 2), we can look forward to the new evidence that *neuropsychopragmatics*, alongside *psychopragmatics*, may provide towards confirming or disconfirming our theoretical accounts of pragmatic phenomena.

In this paper, which is pitched entirely at the level of cognitive theorising and a rather speculative variety at that, I set out some issues that arise within the relevance-theoretic account of how words, as tools in ostensive communication, are understood in context to convey a range of different meanings (including cases of metaphorical meaning), which are pragmatically inferred for the specific occasion of use. At the end of the paper, I briefly outline some questions whose ultimate resolution might be aided by the kind of empirical evidence attainable from studies using techniques from both psychopragmatics and neuropsychopragmatics.

2. *Explicature and lexical pragmatics*

It is a basic assumption of RT pragmatics that the meaning encoded in the linguistic expression type that a speaker utters inevitably underdetermines the content that she communicates, not only her implicatures, but also the propositional content she communicates explicitly (the 'explicature' of the utterance). In other words, there is no proposition communicated whose content is recoverable by linguistic decoding processes alone; pragmatic inference inevitably plays a role in the recovery of speaker meaning, that is, her overtly intended content (Carston 2002: ch. 1; 2009).

There is a range of pragmatic tasks involved in determining the proposition explicitly communicated, including disambiguation,

assignment of referents to indexicals and filling in missing constituents, as illustrated in (1), and various other enrichments or adjustments of encoded content, as indicated in the examples in (2) (where the particular propositional meaning given here is, of course, just one of indefinitely many possibilities):

- (1) He has taken enough from her.

Communicating: JIM HAS ENDURED ENOUGH ABUSIVE TREATMENT FROM MARY

- (2) a. I've eaten.

Communicating: I'VE EATEN DINNER TONIGHT

- b. He told her the combination and she opened the safe

Communicating: TOM TOLD SUE THE COMBINATION FOR THE LOCK AND SUE OPENED THE SAFE USING THE COMBINATION TOM HAD GIVEN HER

- c. The water is boiling.

Communicating: THE WATER IS VERY HOT [*not necessarily at boiling point*]

The proposition explicitly communicated in (2a) contains constituents of content that don't appear in the surface form of the sentence uttered and, arguably, not at any deeper level of linguistic structure either. Such constituents are known as linguistically unarticulated constituents (UCs) and their existence is a source of considerable contention. A similar point could be made about (2b) since the predicate *opened the safe* has been considerably enriched on the basis of the presumed relevance of the preceding clause (which itself requires disambiguation and enrichment of 'combination'). Another possible analysis here, though, is that the general concept OPEN encoded by the verb *open* is narrowed down to a more specific concept OPEN* with the approximate meaning of 'open x using a particular combination of letters and numbers'. Such a pragmatic adjustment of a lexically encoded concept seems more evidently to have taken place in (2c) where BOILING has been used loosely, perhaps hyperbolically, for a broader concept BOILING* which encompasses both the normal boiling point of water (99.97 degrees Celsius) and a range of somewhat lower temperatures.

As this brief survey of pragmatic contributions to the level of explicature indicates, some of the pragmatic processes at work here are linguistically indicated or required (indexical saturation, disambiguation) while others seem to be free from such linguistic mandate and motivated entirely by the presumption that the utterance will meet a certain standard of communicative relevance. In this paper, I will concentrate on the processes of lexical adjustment which underlie

the widespread observation that a word which seems to have a single stable meaning in the linguistic system can nevertheless be used to express a wide range of distinct concepts on different occasions of utterance. This is less controversial than the idea that there may be constituents of explicature content that are wholly unarticulated by the linguistic form used (as claimed for (2a) above) and a promising suggestion is that, in the long run, many instances put forward as UCs will turn out instead to be cases of pragmatic adjustment of linguistically encoded meaning (as suggested for *open* in (2b) above).

As discussed in recent relevance-theoretic work, modulation or adjustment of the meaning encoded by a linguistic constituent involves an interaction among the lexically encoded concept, the other concepts encoded by the utterance and contextual information, constrained by the hearer's expectation of relevance (see, for example, Wilson & Carston 2007). The outcome of this process is what is known as an *ad hoc* concept ('ad hoc' in that it has to be inferentially derived on, and for, the particular occasion of use) and, as in the examples just discussed, it is marked with an asterisk (*BOILING**, *OPEN**, etc.) to distinguish it from the context-independent lexically-encoded concept (*BOILING*, *OPEN*, etc.).³ The pragmatically derived concept may be more specific or more general than the encoded concept; that is, its denotation may be either a proper subset or a superset of the denotation of the linguistically encoded concept, or it may be a combination, both extending the lexical denotation and excluding some part of it.

In order to appreciate how productive this lexical pragmatic process is, consider the following simple example:

(3) Let's dance.

It's not too difficult to think of a range of scenarios in each of which *dance* would be understood somewhat differently: suppose speaker and addressee are (a) at a ball where the orchestra has just started playing a waltz, (b) at a Scottish céilidh where a six-person round is about to begin, (c) at a party where people are moving about individualistically, apparently in response to blaring rock music, or (d) suppose the speaker is Rudolf Nureyev addressing Margot Fonteyn. Although the word *dance* is used literally in all these cases, the particular concept expressed is likely to be distinct in the four situations, each one a more specific (narrower) concept than the lexically encoded concept *DANCE*. Suppose next that the interlocutors are, in fact, already dancing but in a somewhat unenergetic lacklustre way when one of them, inspired by a change in the music, utters (3), thereby commu-

nicating a concept that might be roughly paraphrased as ‘dance in an intense, focussed, lively way’, or ‘do our flamboyant party piece’, hence further narrowings of the lexical concept. Equally, the opposite sort of concept adjustment, loosening or broadening, could result in any of a range of concepts, from those that involve extensions to the range of bodily movements included in the denotation, for instance, what might be roughly paraphrased as ‘walk together in a light, rhythmic way, keeping in step with each other’, through to those of a more clearly metaphorical nature, like ‘spend our life harmoniously together, attuned and responsive to one another, never moving far apart’. Further variations in the ‘dance’ concept expressed will come with a change of subject: consider, for instance, *Bees dance to tell their conspecifics where nectar is located*, or *See how the daffodils dance in the breeze*. In short, the unambiguous verb *dance* might be used to communicate any of an indefinite range of related concepts (DANCE*, DANCE**, DANCE***, ...). See Carston (2002), Wilson & Carston (2007; 2008) for more examples and more detailed analysis and explanation.

As noted above, this is a ‘free’ pragmatic process in that nothing in the linguistic form indicates that it must be carried out; it is therefore optional, that is, there are contexts where the encoded lexical concept would suffice (an utterance of *Children in most cultures dance spontaneously* might be an example where the encoded concept DANCE is communicated). The consensus is that these pragmatic adjustments contribute to explicature (hence to the truth-conditional content of the utterance) rather than merely being implicated (for arguments supporting this point, see the references immediately above).

Although contextualist philosophers of language have for quite some time been pressing the point that word meaning is irremediably context-sensitive and occasion-specific (see, in particular, Travis 2008), the attempt to give a cognitive account of the mechanisms and processes involved is quite recent. Thus there are many intriguing questions yet to be pursued, in particular concerning the nature of *ad hoc* concepts in the mind, how they are ‘constructed’ or accessed in the course of utterance comprehension, how they can be progressively conventionalised and lose their ‘ad hoc’ status, and so on. In this paper, I’ll take up two central issues (or sets of issues). The first concerns the nature of the linguistic semantic input to the pragmatic modulation processes, specifically whether or not lexically encoded concepts are atomic or structured/decompositional, a question which then arises in turn for the output, that is, the derived *ad hoc* concepts themselves. The second array of issues concerns certain figurative uses of words (and phrases), including metaphor, simile and meton-

ymy. Each of these gets a different treatment within RT: metaphor is claimed to be a case of loose use (on a continuum with approximations, category extensions and hyperboles) and so is accounted for by the same inferential mechanism of concept adjustment as the cases just discussed; similes have received little attention, but it has generally been assumed that they are not subject to the kind of conceptual adjustment, in particular broadening, that their corresponding metaphors undergo; metonymic uses present an interesting challenge since they seem to result in an *ad hoc* concept which contributes to explication (truth-conditional content) but they are not instances of narrowing or broadening of the encoded lexical concept.

3. *Lexical semantics according to Relevance Theory*

Starting with the first set of issues, let's consider the relevance-theoretic stance on lexical meaning. The first thing to say is that, as regards their encoded semantics, words are taken to be a heterogeneous lot: some of them encode full-fledged concepts; some encode 'pro-concepts' or conceptually incomplete information (for discussion and examples, see Sperber & Wilson 1998: 185); some encode procedural meaning (constraints on pragmatic inference) and this category is itself very heterogeneous, allegedly including pronouns, discourse connectives, tense, aspect and mood indicators, particles and interjections. I will confine my attention here to those words that, arguably, encode full-fledged concepts.⁴ The central claim about these is that they encode *atomic concepts* rather than molecular (structured) concepts. Abstracting away from the important formal linguistic information (phonological and syntactic) stored in lexical entries, what this means is that there is a simple mapping from lexical form to mental concept; the concept is completely unstructured and the lexical entry does not specify any further information about its content or semantic behaviour. In short, the position is essentially the same as that of Fodor's "disquotational lexicon": the word *house* means HOUSE, *miserable* means MISERABLE, *keep* means KEEP, and so on (see Fodor 1998; Fodor & Lepore 2002). (However, for an important difference between RT and Fodor on conceptual content, see note 6 below). Given the heterogeneity point above, it follows that the conceptual atomism claim applies to only a subset of the vocabulary, albeit a sizable one, consisting of many of those lexical items that fall into what are informally termed 'open word classes', specifically nouns, verbs, adjectives and at least some adverbs.

Since the opposite view, that lexical meanings are compositional, is widely supported, some brief discussion is in order. It seems to me that the Fodorian arguments against *definitional* lexical decomposition are unassailable (see, for instance, J. D. Fodor et al. 1975; Fodor et al. 1980; Fodor 1998). The most compelling of these, perhaps, is that no-one has been able, despite centuries of trying, to give adequate definitions for any but a tiny group of words (for example, *bachelor*, *mother*, and their ilk). There seem to be principled reasons why, for instance, natural kind terms cannot be analysed beyond the observation that a whole bunch of them entail ANIMAL or COLOUR or PLANT or METAL: what concept is to be added to COLOUR in order to give us RED, other than RED itself, what concept(s) are to be composed together with ANIMAL to give us HORSE, and so on? Furthermore, children's conceptual acquisition seems to proceed from the allegedly more conceptually complex concept (MOTHER, RED, HORSE) to the allegedly more basic or primitive one (PARENT, COLOUR, ANIMAL). I won't rehearse the arguments in any more detail here. Most advocates of complex lexical meanings nowadays favour some kind of *non-definitional* form of decomposition (for example, Pustejovsky 1995; Jackendoff 2002: ch. 11; Vicente & Martínez Manrique *forthcoming*). I touch on some of the problems with this seemingly more reasonable decompositionalist position below.

On the relevance-theoretic view, what the encoded atomic concept amounts to is an address in memory or, viewed from a different perspective, a basic element of the language of thought (a monomorphemic 'word' in Mentalese). The content or semantics of this entity is its denotation, what it refers to in the world, and the lexical form that encodes it, in effect, inherits its denotational semantics. This conceptual address (or "file name")⁵ gives access to a repository of mentally represented information about the concept's denotation, some of which is general and some of which, such as stereotypes, applies only to particular subsets of the denotation. This information includes conceptually represented assumptions and beliefs, held with varying degrees of strength, and also, in some cases at least, imagistic and/or sensory-perceptual representations. A distinction is standardly made in the theory between this kind of information, which is stored in the 'encyclopaedic entry' associated with the concept, and the 'logical entry' for the concept. Logical entries consist of inference rules (rather than propositional representations) which are, crucially, taken to be content-constitutive (Sperber & Wilson 1986/1995; Horsey 2006). This logical/encyclopaedic distinction is rather controversial and needs a lot more consideration than I can give it here.⁶

Returning to the *dance* example in (3) above, the idea would be that the decoded atomic concept DANCE gives access to a range of knowledge and beliefs about the activity of dancing, including general information (conceptual and imagistic) about the kinds of bodily movements it involves and about its expressive and social functions, information about specific kinds of dancing, and more idiosyncratic information (episodic memories) based on one's own observations and experiences of particular instances of the activity. When Tom and Mary, a couple of long-standing, are walking along a busy town street, feeling happy and relaxed, and Mary, in romantic mood, says *Let's dance*, she is most likely not suggesting that they break into an unaccompanied waltz or tango, or perform a balletic pas-de-deux, but rather that they walk in a more mutually attuned way, closer together, taking lighter steps, rhythmically, in time with each other, and so on. The denotation of the concept expressed, DANCE*, is certainly broader than that of the encoded DANCE and possibly also narrower (it might exclude very elaborate dance-movements that require years of training). On the relevance-theoretic account of how this kind of word meaning adjustment takes place in on-line utterance interpretation, it is simply one case of a more general process of *mutual parallel adjustment* in which tentative hypotheses about contextual assumptions, explicatures and contextual implications are incrementally modified so as to yield an overall interpretation which is both inferentially sound and satisfies the hearer's expectations of relevance. So, in the example under discussion, the explicature LET'S DANCE* is, at least partially, the result of backwards inference processes responsive to Tom's on-line hypotheses about the relevance (the intended contextual implications) of Mary's utterance, implications along the lines of 'we are in special harmony with each other tonight; let's enhance further this feeling of closeness and mutual accord; we can walk together in a more coordinated and graceful way despite all the people around us, etc'. These play a shaping role in the derivation of the non-lexicalised, probably ineffable, *ad hoc* concept DANCE*. For much more fully realised accounts of the way in which the lexical adjustment process works, see Wilson & Sperber (2002), Vega Moreno (2007), Wilson & Carston (2007; 2008).

Vicente & Martínez Manrique (*forthcoming*) take the position that the kind of 'rampant polysemy' entailed by the free pragmatic process of *ad hoc* concept formation just outlined is incompatible with the view that lexically encoded concepts are atomic and they advocate a decompositional view of lexical concepts. I take it that what lies behind this allegation of incompatibility is the thought that, if

a word's standing meaning can be adjusted/modified so that different meanings/senses are communicated on different occasions of use, there must be an array of meaning components that can be played about with, highlighted, backgrounded, dropped, or otherwise rearranged. An unstructured monolithic atom does not provide us with the distinct parts needed for the job. But, if this is the thinking, it is quite wide of the mark since the account of *ad hoc* concept formation is not semantic, not internal to the linguistic system, but wholly pragmatic; that is, the kind of information that does the work is, for the most part, general encyclopaedic knowledge/beliefs about the world (including, in particular, information about the entities and/or properties denoted by the lexical concept). So, even if word meanings were decompositional, the component features or subconcepts would seldom, if ever, be sufficient on their own to account for the (more or less indefinite) range of concepts that can be communicated by the use of a particular word form on different occasions of utterance. As described above, *ad hoc* concepts are an outcome of the process of finding the interpretation of an utterance that meets one's expectations of (optimal) relevance, that is, the interpretation which has a satisfactory range of cognitive implications and requires no gratuitous processing effort. A requirement here is that these implications are properly inferentially warranted and a major source of premises for deriving such implications is the logical and encyclopaedic information activated by the decoded lexical concept. In short, the *lexical* decomposition issue is really not relevant at the level (conceptual, non-linguistic) at which this is going on.

One of the current non-definitional decompositionalist positions is James Pustejovsky's generative lexicon hypothesis. According to this view, the lexicon consists of quite complex lexical entries; for instance, the entries for nouns like *potato*, *cake*, *book*, *knife* include information about both the origin (natural or artefactual) and the purpose (the telic role) of the entities in their denotations. The problems with this approach are legion (see Blutner 2002; Fodor & Lepore 2002; de Almeida 2004; Bosch 2007; Lossius Falkum 2007). Here are two of them: (a) the lexical entries posited include an arbitrary subset of general world knowledge (for example, that books are written for the purpose of being read, that windows consist of a frame (often made of wood) and a pane of glass, that cakes come into existence through human action, that knives are for cutting, and so on), and (b) the approach can account for only a very restricted range of cases of meaning modulation (not just in practice, but in principle!), leaving the vast bulk of context-specific senses to be explained by a pragmatic

account, which, once provided, will, of course, also apply to the few that are allegedly resolved by intra-lexical means.

Another long-standing supporter of lexical decomposition is Ray Jackendoff. In Jackendoff (2002), he points out that all the standard arguments against decomposition assume that it involves other lexicalised concepts (words) and suggests that a more reasonable hypothesis is that the ingredients that make up a word's meaning are some other kind of element altogether, which cannot be captured by using natural language forms. He makes an analogy with the physical decomposition of substances like oxygen and sulphur into elementary particles that are different in kind, such as electrons, protons, neutrons, and so on. But, while the general point seems sound, there is very little in the way of a concrete proposal about the nature of the (non-definitional) subatomic components that might make up lexical meaning. Jackendoff's conceptual decompositions employ features like CAUSE, PATH, OBJECT, EVENT, STATE, which may, but more often do not, mean the same as the apparently corresponding English words, so it is difficult to assess the content of the proposed (partial) conceptual analyses. His further claim is that these non-definitional lexical decompositions are 'completed' by abstract image structures, so, for instance, while *red* would share the component concept COLOUR with all other colour terms, it would be distinguished from them by its own particular imagistic component (Jackendoff 2002: 345-350). This proposal is interesting and deserves more sustained scrutiny than I can give it here. As with Pustejovsky's (1995) proposed *qualia* structures for lexical items, I think it could be reinterpreted in RT terms as providing formulations of some of the material to be found in the encyclopaedic entries of lexicalised atomic concepts.⁷

4. What are ad hoc concepts?

The questions in the domain of relevance-theoretic lexical pragmatics that strike me as most in need of some long hard thought concern the nature of *ad hoc* concepts. Are *ad hoc* concepts the same kind of entity as lexical concepts (apart from not being lexicalised)? Are they atomic or decompositional (perhaps even definitional)? Do they have logical and/or encyclopaedic entries? What do they look like as mental representations (that is, what lies behind the stand-in asterisked notation, TIRED*, DANCE*, and so on)? How stable and/or long-lasting are they as components of our thinking apparatus? This is a research programme with most of the work yet to be done and

I don't have much to offer here beyond a few hunches, hopes and intuitions.

Let's consider the question of whether these pragmatically derived concepts are atomic or complex (decompositional). In line with their decompositional stance on lexical meaning, Vicente & Martínez Manrique (*forthcoming*) advocate a decompositional view of *ad hoc* concepts, so that, for instance, ANGEL* (as in *My girlfriend is an angel*) is made up of KIND, GOOD and a few other atomic concepts. They claim that this is "simpler and more intelligible" than an atomistic view. I can see the intuitive appeal of the view (component elements get removed or added in), but the lack of evidence for the initial lexical input having the required internal structure undermines the apparent simplicity and intelligibility. On the basis of the paraphrases often provided in the RT literature for *ad hoc* concepts, it might look as if they are being construed as decompositional; for example, DANCE* has been glossed as 'dance in an intense, focussed, lively way', TIRED* as 'tired to the extent that one does not want to go out', RAW* as 'so grossly undercooked as to be virtually inedible', and so on. But the idea behind this is that *ad hoc* concepts are, generally, ineffable, in the sense that, as well as not being lexicalised, there isn't a linguistic phrase that fully encodes them either, and the paraphrases are intended as just a rough indication to aid readers in understanding what we have in mind in particular cases.

A decompositional view might also seem to have been implied by my talk (Carston 2002: 239) of the dropping of logical properties (in the case of loose uses) and the promoting of encyclopaedic properties (in the case of narrowing), though it doesn't strictly follow, since these properties are clearly not internal components of the lexical concepts themselves and need not be taken that way for *ad hoc* concepts either. In fact, it was my aim then, as now, to maintain a consistently atomic view of concepts if at all possible.

It might be useful at this point to remind ourselves what is meant by 'narrowing' and 'broadening' in the theory (as so far developed): these are descriptions of the outcomes of pragmatic adjustment processes rather than of the processes themselves. The *denotation* of the pragmatically inferred concept is narrower or broader (or both)⁸ than the denotation of the lexical concept which provided the evidential input to its derivation. The idea is not that there are two distinct processes – of making narrower and making broader – but rather a single overall pragmatic adjustment/modulation process with these various possible results. The perspective is essentially an externalist semantic one. There are important questions about the internal proc-

esses and representations involved in the shift from (atomic) lexical concept to (atomic) *ad hoc* concept – what they are and how exactly they work – and these remain to be answered. Accepting that this is the case, I shall try now to flesh out a little the notion of an *ad hoc* (that is, pragmatically inferred) concept.

Let us suppose that Fodor has successfully dispatched the old empiricist idea that lexical concepts are complexes built out of a relatively small set of primitive atomic concepts. The next reasonable assumption to make is that in our thinking we are employing, as well as lexicalised atomic concepts, a range of atomic concepts that are not encoded in our particular linguistic systems (Carston 1996; Sperber & Wilson 1998). Given the unequivocal differences among languages with regard to the concepts that are lexicalised in them (the different ways they carve up ‘semantic space’, as it is sometimes put), this seems pretty uncontroversial. (Cases of cross-linguistic lexical differences are familiar enough not to need extensive exemplification: for instance, the concept of ‘aunt or uncle’ is lexicalised in some languages but not in English; the concept of ‘grandmother or grandfather’ is lexicalised in English but not in Serbian; English has the two words *foot* and *leg* where Japanese has a single word *ashi*, and so on). Other considerations that weigh in favour of the claim that our atomic concepts exceed our lexicalised concepts are: (i) the (apparently still contentious) view that we are born with at least some innate concepts, (ii) the view that children’s word learning often involves matching a piece of linguistic form to a concept (antecedently acquired), and (iii) the fairly widely accepted idea that at least some animals have some concepts (but, of course, no lexical items).

So the idea here is that, although, for example, there is a range of English words describing states of tiredness (*tired*, *weary*, *sleepy*, *bored*, *exhausted*), it seems likely that the concepts featuring in the thoughts of even the monoglot English speaker concerning such states come in a much finer grain (that is, there are lots more concepts than words in this domain). So also for our ‘happy’ words and our ‘HAPPY’ concepts, our words for colours such as *green* and *red* versus our finer-grained concepts for kinds of green and red, our words for different manners of walking, cutting, opening, eating, smiling, versus the range of concepts we possess in each of these domains, and so on for the rest of our open class vocabulary. I see no reason to suppose that these stable non-lexicalised concepts, regularly employed in our thinking, are not essentially the same in kind as lexical concepts; that is, they are atomic and just as likely to come with logical and encyclopaedic entries as are lexical concepts.

The implication of this for the pragmatic process of inferring *ad hoc* concepts in utterance interpretation is that it may result in a tokening of one of these non-lexicalised atomic concepts which is already an established component of the hearer's conceptual system. In such a case, the pragmatic process is more one of retrieving or locating the relevant concept than of constructing it.

But, of course, other pragmatically inferred concepts may be quite new or, at least, have made such an infrequent appearance in the hearer's thinking repertoire that there is no established conceptual address for them, hence no logical or encyclopaedic entry (or, to put it another way, no mental file has yet been opened for them). Strictly speaking, these new, possibly one-off, *ad hoc* entities are not concepts, although they have the potential to become concepts, that is, stable, enduring components of Mentalese. Nevertheless, even in their pre-conceptual manifestation, they can make a contribution to structured propositional states, specifically explicatures, alongside fully-fledged concepts (whether lexical or *ad hoc*) and play a role in warranting certain implications of the utterance. They are best thought of as metarepresentational or interpretive, where what is metarepresented or interpreted is whatever concept the speaker intended by her particular use of a word, and their conceptual potential is partially grasped by the hearer insofar as he takes them to contribute to the grounding of intended implications of the utterance.

To illustrate this sort of case, consider the following example: suppose you and I have been discussing one of our colleagues, Lionel, who is a prolific producer of academic papers, but, in our view, hasn't written anything new or interesting for years. Knowing that you've just read his latest offering, I ask you what it's like and you reply:

- (4) He's taken ideas from several different theories and stewed them together.

I interpret you as meaning (implicating) that Lionel's new paper doesn't contain any original thought, that it is a mixture of already existing ideas, that he has attempted to make something that looks like a new theory but what he has produced is, in fact, an unstructured mess. However, your use here of the verb *stew* is new to me; the *ad hoc* component of the explicature which I construct, based on the encoded concept STEW, is not a concept already active in my conceptual system. It is "STEW*", which is both *ad hoc* and metarepresentational (hence the quotation marks). Based on what I know already about 'stewing', both as a value-neutral culinary term

and as a term with negative connotations in *The tea was stewed* and *She got into a stew over her lost keys*, together with the specifics of the current context, I can figure out that as used here it is a somewhat negative expression denoting the process of putting together various disparate bits and pieces and failing to produce anything from them that is fresh or clearly articulated. Although this 'concept' is new to me and I might have to think a bit before trying to employ it myself, it does the job quite adequately in this particular one-off communicative exchange where the intended implications are plain enough.⁹

This is not a particularly distinctive new use of the verb *stew* and at least some of the clues for how to construe it are readily available from more well-established uses. There are, no doubt, more creative new uses than this one, which are more strikingly distinct from existing concepts associated with the word form. Because they will be less closely related to any concept already established in one's Mentalese repertoire, they are likely to be harder for a speaker to come up with and for a hearer to grasp, and are more likely to be found in carefully crafted literary texts.

The overall picture, then, is one of pragmatically inferred (retrieved or constructed) *ad hoc* concepts that range from those that already have a firm presence in the hearer's cognitive system and so, in that sense at least, are not 'ad hoc', to those that are entirely *ad hoc* (new, occasion-specific) but don't qualify (yet) as full 'concepts', with probably various intermediate cases (involving degrees of *ad hocness* and/or degrees of conceptualness).

5. *Ad hoc concepts and figurative language use*

The second set of issues that falls within this general area of a free pragmatic process of *ad hoc* concept formation concerns whether and, if so, how this construct plays a role in an account of our understanding of certain figurative uses of language such as hyperbole, metaphor, simile, metonymy, synecdoche and epizeuxis (immediate word repetitions). Here I will focus on metaphorical uses, with some consideration of corresponding similes. As is well-known, within current relevance theory, comprehension of a metaphorical use is a case of *ad hoc* concept formation where, crucially, the concept inferred is much broader in its denotation than the lexical concept from which it was derived.¹⁰ Corresponding similes, on the other hand, are assumed to work rather differently and it is the literal lexical concept, rather

than the broadened *ad hoc* concept, that appears in their explicature, as in (5d), communicated by the simile in (5c):

- (5) a. John is a shark
b. JOHN IS A SHARK*
c. John is like a shark
d. JOHN IS LIKE A SHARK

The reason for this seems clear enough: it would make little sense to say of someone who is a member of a certain category (here SHARK*) that he is (merely) *like* things in that category; that would be comparable to saying that an apple is (merely) like a fruit or a robin is (merely) like a bird (Carston 2002: 357-358).

Nevertheless, one might feel there is something amiss here in that what is communicated by (5a) and (5c) is surely very similar, if not identical (perhaps just differing in directness or force) and yet the key concept in the explicature in each case is quite different: the denotation of SHARK* has a radically broader denotation than SHARK since, as well as actual sharks, it includes some human beings. The apparent closeness in the meaning of metaphors and corresponding similes has been captured in other theories in various ways; for instance by treating metaphors as ellipsed similes or by treating similes as hedged metaphors.

In fact, contrary to the 'same meaning' intuition, there is empirical evidence that people interpret metaphors and similes rather differently. Glucksberg & Haught (2006) found that experimental participants made different judgements about the acceptability/aptness of corresponding metaphors and similes and that their interpretations of metaphors were more likely to involve emergent properties (that is, properties that are not directly associated with either the metaphor topic or vehicle) than their interpretations of the corresponding similes. O'Donoghue (2009) points out that there are instances of similes which simply have no direct metaphorical counterpart (and vice versa) and she makes a persuasive case for there being certain contexts in which similes are a more effective communicative device than their corresponding metaphors. Both of these studies support accounts, such as the RT one, which take the concept explicitly communicated by metaphors and similes to be different. Finally, it's worth noting that autistic people who find metaphorical uses difficult to understand seem to be much less troubled by similes (see Happé 1993), although exactly what they understand as communicated by the simile cases is something that needs closer investigation.

Focusing now on metaphor alone, in recent years the *ad hoc* concept account of metaphor comprehension has been developed in some detail (Carston 2002; Vega Moreno 2007; Wilson & Carston 2008; Sperber & Wilson 2008) and has led to new questions and, of course, criticisms. One interesting question concerns how emergent properties are to be accounted for. For example, understanding *Robert is a bulldozer* might well include deriving the implication that Robert is insensitive, but INSENSITIVE is not likely to occur in the hearer's encyclopaedic entry for bulldozing machines (the metaphor vehicle). Pilkington (*forthcoming*) considers this question and criticises some existing accounts that try to explain property emergence in wholly conceptual inferential terms. He argues that mental imagery (across a range of sensory modalities) plays a central role in accounting for emergent properties in metaphor comprehension and, following McGinn (2004), he insists that imagery is a distinctive type of mental category which is not reducible to the conceptual. Taking the case of *Robert is a bulldozer*, the way properties emerge is through imagining or (mentally) seeing Robert *as* a bulldozer and then, by internal scrutiny of that mental image, 'reading off' properties which can be represented conceptually as INSENSITIVE, OVERBEARING, UNSWERVING (or as related non-lexicalised concepts, such as UNSWERVING*). I agree with the general direction of these remarks (see Carston 2002: 349-356) and believe that future work within RT on the pragmatics of various kinds of 'figurative' language use, including metaphor, should look more closely at the role of imagistic representation.

A second question, one that has interested me for some time, concerns just how far we should or can take the *ad hoc* concept approach, what range of cases it applies to. While it provides a neat and convincing account of how we understand spontaneous conversational (often somewhat conventionalised) cases of metaphor, such as *John is a shark*, *That surgeon is a butcher*, *She bulldozed the entire committee into acquiescence*, and so on, it is not obvious that it carries over to more innovative cases or to those that are extended and developed over a stretch of discourse/text (perhaps a whole poem). Consider the following familiar example (truncated in the interests of space):

- (6) All the world's a stage,
And all the men and women merely players:
They have their exits and their entrances;
And one man in his time plays many parts,
(Shakespeare: *As You Like It*, 2/7)

In this developed metaphor (or metaphorical ‘conceit’), it seems unlikely to me that comprehension involves the formation of a series of (radically broadened) *ad hoc* concepts (STAGE*, PLAYERS*, EXITS*, ENTRANCES*, PARTS*, and so on). Rather, what seems to go on is that a literal interpretation of the whole text is maintained and is metarepresented as a whole, so that what we have is a representation of an imaginary state of affairs in which human life takes place on a large theatre stage, and the phases of each person’s life and the activities he or she takes part in are a matter of acting out a pre-existing script. Our mental representations of this non-actual, imagined world are compartmentalised and sealed off from our factual beliefs (our representations of the actual world), as with games of make-believe or pretence and other surreal or fantastical conceptions that we recognise as such. Processing of the (strictly false) literal interpretation within the metarepresentation will yield a range of implications and other effects, some of which will be judged to apply to the actual world, that is, to be true (for instance, ‘The course of human life is largely predetermined’, ‘We are powerless against the passing of time’, ‘Most of our activities and concerns are of only momentary significance’, and so on). The hearer/interpreter may disembed these implications from the metarepresentational frame, taking them as speaker-meant, and carry them over into his descriptive mental representation of the actual world.

This, clearly, is a very different sort of interpretive process from the lexical pragmatic mechanism of *ad hoc* concept construction, so it might look as if I’m claiming that there are two kinds of metaphor (the lexical and the extended). Rather, I see it as a matter of processing load or threshold: there’s a point up to which interpreters can and do adjust or modulate the literal encoded meaning (that is, construct *ad hoc* concepts to fit the world as they know it) and beyond which they don’t/can’t. When this point is reached, the literal meaning is maintained but, given that it is clearly not speaker-meant, it is metarepresented and held, as it were, for further processing. There may be individual differences as regards the tipping point for moving from the one mode of processing to the other. Clearly, a lot more needs to be said about how this second kind of interpretation works, the kind of effort it requires and the effects it achieves, and whether the processes involved are to be thought of as more controlled and reflective than the fast, automatic pragmatic processes engaged in the comprehension of ordinary conversational lexical metaphors. For more detail, see Carston (2010).¹¹

6. Conclusion: Lexical semantics/pragmatics and neuropragmatics

As will be amply evident from the previous sections, there are many open questions about the nature of *ad hoc* concepts and the role they play in ostensive communication, and one might even question whether there really are any such entities. Debate about this at the theoretical level has scarcely begun and few questions have been formulated precisely enough yet for serious experimental investigation. There are almost as many unresolved issues concerning the nature of lexically encoded word meanings, though the debate in this area has been alive and kicking for many years, especially regarding the issue of the internal structure of word meanings. Whether current brain scanning techniques or event related potential (ERP) measures or any of the other investigative methodologies of neuroscience could cast any evidential light on this question seems unlikely to me. For instance, suppose that certain brain areas or sections of neural circuitry 'light up' when a particular word is accessed (in a particular context) or that the N400 component of an ERP is larger for one word than for another (in a particular discourse context). Suppose also that this can be taken to be indicative of something about the kind of 'meaning' (information, conceptual features) that has been activated by the word or about the ease/difficulty of integrating it with already represented meaning. Although interesting in itself, it's hard to see how this sort of finding could help us with the more fine-grained question of whether that information (or conceptual feature) is a component of the internal semantics of the word (its lexically encoded meaning) or, rather, is an element of the encyclopaedic information associated with the (possibly atomic, internally unstructured) concept encoded by the word.¹²

However, there are some broader-brush issues concerning 'explicature' and lexical pragmatics generally, and others concerning metaphorically-used language specifically, which might be more neuroscientifically tractable. Some of these questions have been around for a while and others arise from distinctions made in this paper:

1. Do the various pragmatic modulations of word meaning (*ad hoc* concepts) contribute to explicature (often equated with the truth-conditional content of the utterance), as claimed (and argued for) in Relevance Theory, or are they, following a more Gricean approach, merely implicated (implicitly communicated, non-truth-conditional)?

2. Is there a single process of pragmatic modulation or lexical adjustment which, as claimed in RT, can result in a narrower concept, a broader concept, or one that is some kind of combination of the two, or are narrowing/enriching and broadening/loosening really distinct processes, with possibly distinct representational effects (only one of them contributing to explicature perhaps)?

3. More specifically, does metaphorical meaning contribute to explicature, as claimed in RT, or should it, differently perhaps from other kinds of lexical adjustments, be understood as having merely implicated effects?

4. Is there a continuum of loose uses between banal cases of concept broadening (such as approximations), category extensions, hyperboles and metaphors, with no clear cut-off points between them, as claimed in RT (see note 10), or is metaphor a distinct category of language use with its own special properties, as claimed in many other theories of language use?

5. Does imagistic representation and imagery processing play a more central or dominant role in the comprehension of metaphor than in literal language use?

6. Are there distinctive modes of processing for different kinds of metaphor, perhaps distinguishing spontaneous and relatively conventional instances, on the one hand, and more creative or developed cases, on the other, as suggested in this paper?

7. Widening out again, are the pragmatic processes involved in the recovery of the explicature of an utterance, including lexical adjustment processes, properly inferential and just as sensitive to the speaker's communicative intentions as those deployed in implicature derivation, as claimed in RT, or are they local 'dumb' associative processes, sensitive only to least-effort contextual fit considerations, as suggested by Recanati (2004)?

Any empirical evidence from cognitive neuroscience (or anywhere else for that matter) bearing on these questions would be very welcome.

Address of the Author

Robyn Carston: Linguistics, University College London, Chandler House,
2 Wakefield Street, London WC1N 1PF, and Centre for the Study of
Mind in Nature, University of Oslo
robyn.carston@ucl.ac.uk

Notes

* I would like to thank Deirdre Wilson, Catherine Wearing, Vladimir Žegarac and Richard Breheny for helpful discussions about issues in this paper. Many thanks also to Alessandro Lenci for challenging comments on the penultimate version of the paper and to Valentina Bambini both for the initial invitation to contribute to this special issue on neuropragmatics and for ongoing support and advice throughout the writing process. I gratefully acknowledge funding from the Centre for the Study of Mind in Nature, University of Oslo.

¹ What characterises a modular cognitive system from the perspective of the evolutionary psychologists is not so much the set of properties that Fodor (1983) ascribed to perceptual input systems (being local, fast, mandatory and encapsulated), but “the presence of dedicated mechanisms, (typically biological adaptations to regularities in some domain) which cannot be seen as special cases of more general mechanisms operating in broader domains” (Wilson 2005: 1131).

² I note some rather negative assessments of what brain imaging can reveal about cognitive processes, from the strong position of Fodor (1999), who argues that brain maps are, in principle, incapable of telling us anything of interest about mental functions, to the milder report of Coltheart (2006), who claims that they haven’t told us anything about the mind *so far* (but see rejoinders in the same issue of *Cortex* 42), to Geurts (2008) who, focusing specifically on language processing (syntactic, semantic and pragmatic), finds existing fMRI studies in this area to be mostly ill-conceived, poorly designed and over-interpreted. Geurts does, however, hold out greater hope for ERP studies due to the high temporal resolution of the data they provide.

³ This notion of *ad hoc* concepts should be distinguished from Lawrence Barsalou’s notion of *ad hoc* categories. The latter are categories that people construct in particular situations to achieve particular goals (e.g., the category of ‘things I need to take with me on a short trip’, ‘ways to keep stress to a minimum when writing a conference presentation’), as distinct from common established categories, such as ‘birds’, ‘furniture’, ‘games’ (see, e.g., Barsalou 1991). As regards concepts, Barsalou’s view is that they are quite generally not stored in long-term memory but are temporary, highly context-sensitive entities, constructed in working memory: “[...] people have the ability to construct a wide range of concepts in working memory for the same category. Depending on the context, people incorporate different information from long-term memory into the current concept that they construct for a category” (Barsalou 1987: 118). On this account, then, *every concept is ad hoc* and so the meanings of words cannot be concepts. In RT, on the other hand, concepts are seen as relatively stable enduring mental structures which *can* function as word meanings (see, for instance, Sperber & Wilson 1998).

⁴ Elsewhere, I have speculated that encoded word meaning (of the allegedly conceptual sort) might be better construed as, quite generally, not a matter of full-fledged concepts but something much more schematic and abstract, not a component of thought as concepts are (Carston 2002: 359–364). Clearly, any such position would have significant ramifications for pragmatics – for instance, the process of *ad hoc* concept formation would become perfectly general and obligatory in comprehending words in context, and it would not be a ‘free’ pragmatic process because it would no longer be entirely pragmatically motivated. I set the idea aside for the purposes of this paper, but it has a great deal going for it, in my view, and seems to be gaining ground, albeit in various guises (see, for instance, Bosch 2007; Pietroski 2008; and Pritchard 2009).

⁵ Fodor (2008) adopts the metaphor of mental “files” which seems to be pretty much equivalent to relevance theorists’ talk of conceptual addresses, which give

access to various kinds of information. For instance, he says: "When you are introduced to John [...] you assign him a Mentalese name and you open a mental file, and the same Mentalese expression (M(John)) serves both as John's Mentalese name and as the name of the file that contains your information about John; [...] according to this story, we think in file names; tokens of file names serve both as the constituents of our thoughts and as the Mentalese expressions that we use to refer to things we think about" (Fodor 2008: 94-95, emphasis in the original).

⁶ It is on this point that RT and Fodor part company. Although early Fodor (as in Fodor 1975; J. D. Fodor et al. 1975; and Fodor et al. 1980) advocated conceptual inference rules (or meaning postulates) for capturing the validity of arguments such as 'X kill Y \square Y die' and 'X red \square X coloured', later Fodor (since at least the early 1990s) renounced these, along with any other kind of conceptual/inferential role semantics (see, in particular, Fodor 1998: 108-112). On this later view, content is constituted wholly by nomological relations between mental symbols and the world (entities or properties), and, as shown by Quine, there is no principled analytic/synthetic distinction. Within relevance theory, on the other hand, the logical/ encyclopaedic distinction has been robustly defended by Richard Horsey (2006), using both evolutionary considerations and developmental work on concept acquisition (following ideas set out by Dan Sperber 1994; 1997). Thus, this is a psychological distinction and does not coincide with the (probably untenable) philosophical analytic/synthetic distinction.

⁷ In his review of this paper, Alessandro Lenci objected to my atomistic stance on lexical concepts and pointed out that most *psychological models* of concepts and thus of word meanings are decompositional (that is, they represent lexical concepts as made up of clusters of features), citing Vigliocco & Vinson (2007). In fact, what that paper makes clear is that most psychologists simply *presuppose* that lexical concepts are decompositional and develop their models on that basis. There are quite different kinds of features proposed in different psychological models, ranging from those that are highly abstract (e.g., THING, EVENT, STATE, PATH) to those that are 'embodied' (fundamental aspects of our sensory-motor functioning) and some of them seem to be clearly contingent properties of the real world entities denoted by the concept/word (e.g., <domesticated> vs. <wild>, <having a tail>, <having four legs>, for various animal terms).

It seems that there is still no end in sight to the atomism/decompositionalism debate(s). But what strikes me is how different the concerns and criteria are that different theorists bring to their accounts of concepts (and word meaning). For philosophers like Jerry Fodor, key constraints on an adequate account of concepts include their interpersonal shareability (the 'publicity' constraint) and the compositionality of their content (see Fodor 1998). For many psychologists, on the other hand, models of conceptual structure (and of word meaning) are shaped by (i) 'feature norms' derived by canvassing speakers' reflective assessments of the central features of the meaning of words, and (ii) more direct (on-line) measures (behavioural and neuroscientific) of feature activation in participants' minds when particular words are presented to them. It is far from obvious, then, that the label 'concept' refers to the same entity for both parties (and very clear that the term 'semantic' does not) so that little conciliatory progress is likely to be made until these differences are mapped out and resolved.

⁸ Instances of lexical meaning modulation that result in an *ad hoc* concept which is both broader in some respects and narrower in others than the linguistically encoded meaning are not unusual (see discussion in Carston 2002: 334) and may well be the standard outcome for metaphorical uses (see Carston *forthcoming*).

⁹ These embryonic thoughts on metarepresentational *ad hoc* proto-concepts obviously need a lot of development. However, there is a nice link, yet to be pur-

sued, with a discussion in Sperber (1985: 50-53) on what he calls “semi-propositional” representations, that is, conceptual representations which are not fully propositional in that they include an “unanalysed or incompletely analyzed term”. Such a term is a metarepresented component of the conceptual representation. One of his examples concerns a (communication-derived) belief held by Bob that *stagflation* is a problem in Western economies, where Bob is not quite sure what the word *stagflation* means. Here we have metarepresentation of a word whose lexically-encoded meaning is yet to be pinned down, while in my example in the text the encoded meaning (of *stew*) is known, but what is not fully grasped is the *ad hoc* concept the speaker intended by her particular use of the word. There are clear similarities between the two phenomena. See also Sperber (1997) on reflective concepts (which are metarepresentational) as distinct from intuitive concepts.

¹⁰ A further distinctive RT claim is that while metaphorical use involves a quite radical broadening of denotation, it is on a continuum with other cases of loose use, including approximations and hyperbolic uses, and there are no sharp cut-off points between these uses of language (consider *Her husband is a saint* – is it hyperbolic, metaphoric or both?). Thus metaphors are not to be thought of as a natural kind or as having any special distinctive properties. This continuity view is distinct from the idea that metaphorical use involves *ad hoc* concept construction and it is quite possible to advocate either one of them without the other. Unlike the relatively recent *ad hoc* concept account of metaphor, the continuity view has been around since the early days of RT in the 1980s; for a recent vigorous defence of the claim, see Sperber & Wilson (2008).

¹¹ Metonymic uses of language raise new issues since they are plainly not cases of lexical/phrasal narrowing or broadening (for example, *The twinset and pearls seems to be offended*, *No comment from Buckingham Palace*, *She married a free ticket to the opera*). There seems to be a fairly regular pragmatic process of using a highly salient characteristic of a person or persons as an abbreviatory means of referring to them. Many of these are familiar and routine (*Downing Street*, *The White House*, *The Guardian*, and so on), but other, more novel, cases can have a range of more or less striking effects. Exactly how (and even whether) they are to be analysed within the general lexical pragmatic approach presented here remains to be seen.

¹² The discussion in Van Berkum (2009 and *this issue*), which is much more alert to the distinction between linguistic and pragmatic processing than most neuro-cognitive work in this area, indicates clearly what a coarse-grained index of utterance processing the N400 is. It seems that N400 amplitude is affected by degree of linguistic semantic coherence/anomaly, by degree of coherence with wider discourse context, by the extent to which the utterance achieves optimal relevance for the hearer, and by a range of other features which might cause a temporary processing blip but which are strictly irrelevant to the retrieval of the communicatively-intended content (explicatures and implicatures), such as whether the linguistic content meshes with stereotypes based on the speaker’s sex or social class or age. In fact, the N400 seems to be affected by just about any interpretive ‘surprise’ (e.g., seeing someone who is apparently about to shave himself reach for a rolling pin) and so is not even specific to pragmatic processing. As Valentina Bambini (p.c.) says, the general point here is also true of the P600 and probably of all electrophysiological components, so we should not expect specificity of these components. She believes that we can derive useful results from ERP but only if we ask it ‘the right questions’ (as for all experimental techniques).

Bibliographical References

- BAMBINI Valentina 2005. A survey in neurocognitive pragmatics: with insights for the study of lexical combination. In BERTUCCELLI PAPI Marcella (ed.). *Studies in the Semantics of Lexical Combinatory Patterns*. Pisa: Edizioni Plus. 327-364.
- BAMBINI Valentina & Bruno G. BARA *forthcoming*. Neuropragmatics. In VERSCHUEREN Jef & Jan-Ola ÖSTMAN (eds.). *Handbook of Pragmatics Online*. Amsterdam: Benjamins.
- BARSALOU Lawrence 1987. The instability of graded structure in concepts. In NEISSER Ulric (ed.). *Concepts and Conceptual Development: Ecological and Intellectual Factors in Categorization*. New York: Cambridge University Press. 101-140.
- BARSALOU Lawrence 1991. Deriving categories to achieve goals. In BOWER Gordon H. (ed.). *The Psychology of Learning and Motivation: Advances in Research and Theory* 27. San Diego, CA: Academic Press. 1-64.
- BLUTNER Reinhard 2002. Lexical semantics and pragmatics. *Linguistische berichte* 10. 27-58.
- BOSCH Peter 2007. Productivity, polysemy, and predicate indexcality. In TEN CATE Balder & Henk ZEEVAT (eds.). *Proceedings of the Sixth International Tbilisi Symposium on Language, Logic and Computation*. Berlin: Springer. 58-71.
- BREHENY Richard, Napoleon KATSOS & John WILLIAMS 2006. Are scalar implicatures generated by default? *Cognition* 100. 434-463.
- BURTON-ROBERTS Noel (ed.) 2007. *Pragmatics*. Basingstoke: Palgrave Macmillan.
- CARSTON Robyn 1996. 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 Robyn 2002. *Thoughts and Utterances: the Pragmatics of Explicit Communication*. Oxford: Blackwell.
- CARSTON Robyn 2009. The explicit/implicit distinction in pragmatics and the limits of explicit communication. *International review of pragmatics* 1. 35-62.
- CARSTON Robyn 2010. Metaphor: *ad hoc* concepts, literal meaning and mental images. *Proceedings of the Aristotelian Society* 110.
- CARSTON Robyn *forthcoming*. Metaphor: old words, new concepts, imagined worlds. *Language and cognition* 3.
- CHEVALLIER Coralie, Mathilde BONNEFOND, Jean-Baptiste VAN DER HENST, Ira A. NOVECK *this issue*. Using ERPs to capture inferential processes guided by prosodic cues.
- COLTHEART Max 2006. What has functional neuroimaging told us about the mind (so far)? *Cortex* 42. 323-331.
- COSMIDES Leda & John TOOBY 1994. Origins of domain-specificity: the evolution of functional organization. In HIRSCHFELD & GELMAN 1994. 85-116.
- COULSON Seana 2004. Electrophysiology and pragmatic language comprehension. In NOVECK & SPERBER 2004. 187-206.

- DE ALMEIDA Roberto G. 2004. The effect of context on the processing of type-shifting verbs. *Brain and language* 90. 249-261.
- FODOR Jerry A. 1975. *The Language of Thought*. Sussex: Harvester Press.
- FODOR Jerry A. 1983. *Modularity of Mind*. Cambridge, MA: MIT Press.
- FODOR Jerry A. 1998. *Concepts. Where Cognitive Science Went Wrong*. Oxford: Clarendon Press.
- FODOR Jerry A. 1999. Diary: Why the brain. *London review of books* 21. 68-69. Available at: http://www.lrb.co.uk/v21/n19/fodo01_.html
- FODOR Jerry A. 2008. *LOT 2: The Language of Thought Revisited*. Oxford: Clarendon Press.
- FODOR Jerry A., MERRILL F. GARRETT, EDWARD C. T. WALKER & CORNELIA H. PARKES 1980. Against definitions. *Cognition* 8. 263-367.
- FODOR Jerry A. & ERNIE LEPORE 2002. The emptiness of the lexicon: reflections on Pustejovsky. In FODOR Jerry & ERNIE LEPORE *The Compositionality Papers*. Oxford: Clarendon Press. 89-119.
- FODOR Janet D., JERRY A. FODOR & MERRILL F. GARRETT 1975. The psychological unreality of semantic representations. *Linguistic inquiry* 6. 515-531.
- GEURTS Bart 2008. Neurocognition of language: good, bad, and bogus. Available at: <http://ncs.ruhosting.nl/bart/papers/ncl.pdf>
- GIGERENZER Gerd, PETER M. TODD & the ABC Research Group 1999. *Simple Heuristics that Make Us Smart*. Oxford: Oxford University Press.
- GLUCKSBERG Sam & CATRINEL HAUGHT 2006. On the relation between metaphor and simile: when comparison fails. *Mind and language* 21. 360-378.
- HAPPÉ Francesca 1993. Communicative competence and theory of mind in autism: a test of Relevance Theory. *Cognition* 48. 101-119.
- HIRSCHFELD Lawrence A. & SUSAN A. GELMAN (eds.) 1994. *Mapping the Mind: Domain Specificity in Cognition and Culture*. New York: Cambridge University Press.
- HORN Laurence & GREGORY WARD (eds.) 2004. *The Handbook of Pragmatics*. Oxford: Blackwell.
- HORSEY Richard 2006. *The Content and Acquisition of Lexical Concepts*. London: University of London. PhD dissertation.
- JACKENDOFF Ray 2002. *Foundations of Language: Brain, Meaning, Grammar, Evolution*. Oxford: Oxford University Press.
- LOSSIUS FALKUM Ingrid 2007. Generativity, relevance and the problem of polysemy. *UCL working papers in linguistics* 19. 205-234.
- MCGINN Colin 2004. *Mindsight*. Cambridge, MA: Harvard University Press.
- NOVECK Ira A. & ANDRES POSADA 2003. Characterising the time course of an implicature: an evoked potentials study. *Brain and language* 85. 203-210.
- NOVECK Ira A. & DAN SPERBER (eds.) 2004. *Experimental Pragmatics*. Basingstoke: Palgrave Macmillan.
- NOVECK Ira A. & DAN SPERBER 2007. The why and how of experimental pragmatics: the case of 'scalar inferences'. In BURTON-ROBERTS 2007. 184-212.
- O'DONOGHUE Josie 2009. Is a metaphor (like) a simile? Differences in meaning, effect and processing. *UCL working papers in linguistics* 21. 125-149.
- PIETROSKI Paul 2008. Minimalist meaning, internalist interpretation. *Biolinguistics* 2. 317-341.

- PILKINGTON Adrian *forthcoming*. Metaphor comprehension: Some questions for current accounts in Relevance Theory. In SORIA & ROMERO *forthcoming*.
- PRITCHARD Timothy 2009. *The Pattern of Word Meaning: How Words Enable us to Express Content*. London: Kings College London. PhD dissertation.
- PUSTEJOVSKY James 1995. *The Generative Lexicon*. Cambridge, MA: MIT Press.
- RECANATI François 2004. *Literal Meaning*. Cambridge: Cambridge University Press.
- RUBIO FERNANDEZ Paula 2007. Suppression in metaphor interpretation: Differences between meaning selection and meaning construction. *Journal of semantics* 24. Special issue on processing meaning. 345-371.
- RUBIO FERNÁNDEZ Paula 2008. Concept narrowing: The role of context-independent information in concepts. *Journal of semantics* 25. 381-409.
- SORIA Belén & Esther ROMERO *forthcoming*. *Explicit Communication: Robyn Carston's Pragmatics*. Basingstoke: Palgrave Macmillan.
- SPERBER Dan 1985. Apparently irrational beliefs. In SPERBER Dan. *On Anthropological Knowledge*. Cambridge: Cambridge University Press. 35-63.
- SPERBER Dan 1994. The modularity of thought and the epidemiology of representations. In HIRSCHFELD & GELMAN 1994. 39-67.
- SPERBER Dan 1997. Intuitive and reflective beliefs. *Mind and language* 12. 67-83.
- SPERBER Dan 2002. In defense of massive modularity. In DUPOUX Emmanuel (ed.). *Language, Brain and Cognitive Development: Essays in Honor of Jacques Mehler*. Cambridge, MA: MIT Press. 47-57.
- SPERBER Dan & Deirdre WILSON 1986/1995. *Relevance: Communication and Cognition*. Oxford: Blackwell.
- SPERBER Dan & Deirdre WILSON 1998. The mapping between the mental and the public lexicon. In CARRUTHERS Peter & Jill BOUCHER (eds.). *Language and Thought: Interdisciplinary Themes*. Cambridge: Cambridge University Press. 184-200.
- SPERBER Dan & Deirdre WILSON 2002. Pragmatics, modularity and mind-reading. *Mind and language* 17. Special Issue on Pragmatics and Cognitive Science. 3-23.
- SPERBER Dan & Deirdre WILSON 2008. A deflationary account of metaphors. In GIBBS Ray (ed.). *The Cambridge Handbook of Metaphor and Thought*. Cambridge: Cambridge University Press. 84-105.
- TRAVIS Charles 2008. *Occasion-Sensitivity: Selected Essays*. Oxford: Oxford University Press.
- VAN BERKUM Jos J. A. 2009. The neuropragmatics of 'simple' utterance comprehension: an ERP review. In SAUERLAND Uli & Kazuko YATSUSHIRO (eds.). *Semantics and Pragmatics: From Experiment to Theory*. Basingstoke: Palgrave Macmillan. 276-316.
- VAN BERKUM Jos J. A. *this issue*. The brain is a prediction machine that cares about good and bad. Any implications for neuropragmatics?
- VEGA MORENO Rosa E. 2007. *Creativity and Convention: The Pragmatics of Everyday Figurative Speech*. Amsterdam: John Benjamins.

- VICENTE Augustin & Fernando MARTÍNEZ MANRIQUE *forthcoming*. On Relevance Theory's atomistic commitments. In SORIA & ROMERO *forthcoming*.
- VIGLIOCCO Gabriella & David VINSON 2007. Semantic representation. In GASKELL Gareth (ed.). *Oxford Handbook of Psycholinguistics*. Oxford: Oxford University Press. 195-216.
- WILSON Deirdre 2005. New directions for research on pragmatics and modularity. *Lingua* 115. 1129-1146.
- WILSON Deirdre & Robyn CARSTON 2007. A unitary approach to lexical pragmatics: relevance, inference and *ad hoc* concepts. In BURTON-ROBERTS 2007. 230-260.
- WILSON Deirdre & Robyn CARSTON 2008. Metaphor and the 'emergent property' problem: a relevance-theoretic treatment'. *The Baltic International Yearbook of Cognition, Logic and Communication*, vol. 3: *A Figure of Speech*. New Prairie Press. Available at: <http://www.thebalticyearbook.org/>
- WILSON Deirdre & Dan SPERBER 2002. Truthfulness and relevance. *Mind* 111. 583-632.
- WILSON Deirdre & Dan SPERBER 2004. Relevance theory. In HORN & WARD 2004. 607-632.