# Ad hoc concepts: atomic or decompositional?*

*Alison Hall*

**Abstract**

The meanings of words are frequently pragmatically modulated in context, so that the concept a word is used to communicate is different than what it encodes. In this short paper, I begin to consider the nature of these communicated, 'ad hoc' concepts. I start from the assumption that the concepts encoded by monomorphemic lexical items are atomic – or, at least, that the stable concepts these lexical items are often used to communicate are – and assess whether it is plausible to maintain that genuinely ad hoc concepts are also atomic, or whether some variety of decompositional account is to be preferred.
*Keywords:* ad hoc concepts, relevance theory, concept atomism, decompositionalism

## 1 Introduction

According to recent work in lexical pragmatics (Blutner 1998, Carston 2002, Recanati 2004, Sperber and Wilson 2008, Wilson and Carston 2007), the meanings of words are frequently pragmatically modulated in context, so that their contribution to the proposition explicitly expressed by an utterance is different from their linguistically encoded meaning. Research so far has concentrated on the relation between the extensions of encoded and communicated concept, and on the pragmatic factors driving the modulation. The question I aim to begin to address here, which has so far been largely ignored, is what modulation involves from the internalist point of view and what the nature of the resulting concepts is.

I follow Fodor and relevance theorists in assuming that lexical concepts (for monomorphemic lexical items) are atomic: that is, unstructured – not constituted by either more primitive concepts or by inferential roles (see Fodor et al 1980, Fodor 1998, Sperber and Wilson 1986/1995, Carston 2002). In the next section, I present briefly some details of atomism and why I take it to be a plausible view, but the main focus of this paper will be the ad hoc concepts that result from lexical modulation.

Accounts of modulation see it as coming about through a sorting or rearrangement, in terms of accessibility, of logical and encyclopaedic information associated with the lexical concept. For example, Carston (2002: 326) discusses lexical narrowing, illustrated by an utterance of "I want to meet some bachelors", where "bachelor" would be used to denote not all those entities denoted by the encoded concept, but a smaller set. The idea is that the encoded concept BACHELOR[1] gives access to various sort of information: lexical (which can be set aside for the purposes of this paper), logical, and encyclopaedic. The logical entry for BACHELOR would consist of inference rules specifying that anything that is a bachelor is male, human, adult and unmarried; encyclopaedic information would include information about 'certain sorts of bachelor, the irresponsible, ... sort, the elderly, solitary, misogynous sort, and those who are youngish, heterosexual, and capable of long-term commitment' (ibid). In the context of utterance, the hearer is aware that the speaker is interested in marriage, and

*This is work in progress which will be revised into a longer paper on lexical modulation. I acknowledge funding from a British Academy Postdoctoral Fellowship and a Research in Paris fellowship from Paris City Hall.

[1] Small caps indicate mental representations.

so it is likely that assumptions about this third sort of bachelor (young, eligible, etc) will be more highly activated than information about the others, so it is that subset of assumptions that is accessed and, as Carston puts it, 'used to construct the ad hoc concept' (ibid). This narrowing 'seems to involve elevating an encyclopaedic property of [the lexical concept] to a logical (or content-constitutive) status': in this example of "bachelor", the property in question might be ELIGIBLE FOR MARRIAGE, for instance (ibid: 339). The result is a different concept, BACHELOR*, which would also differ from the encoded concept in having a much smaller encyclopaedic entry, consisting of the set of activated assumptions selected from the encyclopaedic entry of the encoded concept, and would have a narrower denotation. The complementary process of loosening seems to involve dropping at least one of the logical (defining) properties of the lexical concept: for example, when "raw" is used of an undercooked steak, the property HAS RECEIVED NO COOKING would be dropped, resulting in an ad hoc concept that entails 'undercooked'.

The question I want to address in this paper is whether this account of modulation as presented so far by the relevance theorist can be reconciled with the concept atomism that they are committed to, and specifically, whether atomism can be maintained for genuinely ad hoc concepts. Some authors have recently suggested that this is incompatible with the account of modulation roughly sketched above, and have argued that, even assuming lexical concepts are atomic, ad hoc ones must have some kind of decompositional structure (though are not classic definitions).

In the next section I present some background on concept atomism and why it is a tenable view. After that, I will assume that atomism makes sense for lexical concepts, and discuss recent work that grants this much but suggests that ad hoc concepts are not atomic. I argue against that conclusion and that a case can be made for the view that ad hoc concepts are indeed atomic.

To end this introductory section, two qualifications are in order. First, although the term 'ad hoc concept' is used in the literature to refer to concepts communicated by the use of a word that are different than what that word encodes, it is likely that many of these communicated concepts are not, in fact, ad hoc. Our stock of stable (atomic) concepts may well far exceed our stock of simple lexical items (see Sperber and Wilson 1998). For example, for verbs such as "open" or "cut" that seem to have very general meanings, it is likely that we have several different established concepts, and which one is retrieved on a particular occasion depends on (among other contextual factors) what is being opened or cut: we are likely to employ different concepts when we think about opening a window, versus a wound, versus one's mouth, or when thinking about cutting a cake, versus someone's hair, versus one's finger, and so on. Likewise, we may already have two or three different concepts stored that denote different categories that we would express using the word "bachelor", as discussed above. What I am interested in here is genuinely ad hoc concepts – that is, not already established elements of the subject's conceptual repertoire, but new concepts constructed in the comprehension process, to underwrite the occasion-specific inferences that are intended but would not all be derived from the encoded concept[2]. Henceforth, when I refer to 'ad hoc concept', I mean the genuinely novel ones.

Second, an alternative view on word meanings is that words encode not concepts but something more like rules for use, or pointers to an area of encyclopaedic memory, or a 'grab bag' of associated information (Rayo 2011), containing the materials from which the concept communicated is to be constructed or retrieved. This idea has been suggested in different

[2] A further qualification: To the extent that they are really ad hoc, they are not, strictly speaking, concepts, as concepts are generally taken to be stable mental entities that serve to organize information in memory. Nevertheless, I continue to employ the term 'concept' here for these temporary constructions.

forms by several authors, including Carston (2002, 2011), Pietroski (2005), Recanati (2004), and Schiffer (2003). While I find this idea appealing, I continue here to assume that word meanings are concepts, but I think that the contents of the paper could be straightforwardly recast on most of these 'word meanings as something else' views, particularly when one takes into account the fact that there is a strong intuition that many words, at least, have denotations, suggesting that one particular concept has a kind of priority, even if it is not actually the word meaning (this is, after all, why it made sense to talk of modulation as involving narrowing or loosening).

## 2 Fodorian and relevance-theoretic concept atomism

Concept atomism is the view that basic concepts are unstructured: they do not have constituent concepts, nor are they constituted by their inferential roles. In RT, a concept is taken to consist of an 'address' or 'label', which various types of information (lexical, logical, encyclopaedic) are stored under, and which can appear as a constituent of a thought. Fodor adopts the 'mental file' metaphor for concepts, on which the information associated with a concept is the file, and the 'file name' would be what RT calls the conceptual address, so is what appears as the thought constituent. When it comes to the structure of the file contents, though, there are some differences between Fodor and RT. According to RT (Sperber and Wilson 1986/1995, chapter 2), abstracting away from lexical information, the concept gives access to two types of information, logical and encyclopaedic. As illustrated in section 1 for BACHELOR, the logical entry of a concept consists of a set of inference rules capturing analytic implications of the concept, but usually falling well short of being definitional. For example, CAT would have a logical entry consisting of an inference rule stating that anything that is a cat is an animal; perhaps also a small number of other rules, e.g. that a cat is a mammal. These seem generally to be treated as content-constitutive by relevance theorists (Sperber and Wilson 1986/1995, Carston 2002: chapter 5).

Fodor, at least since the 1990s (e.g. his 1998), rejects the idea that any such inference rules are content-constitutive, his reason for this being the impossibility of drawing a principled analytic-synthetic distinction. However, the familiar Quinean objections to an analytic-synthetic distinction can easily be avoided if one adopts a psychological notion of analyticity. Within relevance theory, Horsey (2006), drawing on Boghossian (1994, 1997), defends the idea that an inference is analytic, hence constitutive of a concept's content, just in case the inference is associated as a meaning postulate with the concept. The traditional philosophical notion of analyticity is dropped: these meaning postulates need not be veridical. The way that Horsey cashes out the difference between such 'psychosemantically analytic' and synthetic information is, briefly, as follows. They are stored in different formats, with the analytic (logical) information taking the form of inference rules that support spontaneous inference and are reliably activated whenever a representation containing the relevant concept is tokened. In contrast, contingent information is represented in the form of propositional encyclopaedic knowledge: it may not be reliably activated by the tokening of the concept, so does not support spontaneous inference (which is why, as illustrated by Carston's discussion of BACHELOR, summarized in section 1, the encyclopaedic entry of a concept may well contain many internal conflicts and inconsistencies).

Horsey's idea receives support from experimental work (Rubio Fernandez 2007), which found that in metaphor interpretation, which is taken to involve loosening, certain 'core features' of word meaning (hence logical or content-constitutive features of lexical concepts) are always activated upon tokening of the word, even in contexts that heavily favour the metaphorical interpretation. Despite the fact that loosening involves dropping certain of these

core features – for example the feature PLANT in the case of a metaphorical use of "cactus" to describe a person's character – Rubio Fernandez found that they remained highly activated well beyond the point at which the metaphorical interpretation had been recovered, requiring more active suppression than is involved in disambiguation[3].

The upshot of this is that the psychological logical-encyclopaedic distinction adopted by RT survives the Quinean objections. A remaining point of difference between Fodor and RT is whether logical information is content-constitutive. Horsey argues that it must be, and Carston (e.g. 2002: 326, 339) and Sperber and Wilson (1986/1995) generally seem to speak as though they take it to be, though Carston (2010) suggests she is not committed to this idea. For Fodor, none of the information associated with a concept is constitutive; the content is constituted entirely by nomological mind-world relations. Good evidence for the idea that none of the information is constitutive is provided by the fact that in principle, none of it is necessary to the identity of the concept, nor does it determine the concept's denotation. To take the standard philosopher's example, if cats were discovered to be robots, this would require a significant change in the logical information associated with CAT, but it would remain the same concept, with the same denotation. A realistic example is provided by "whale": most people assume initially that whales are fish; when they learn that whales are in fact mammals, they may substitute the new logical property for the old one in their mental file contents, but the concept's identity and denotation do not change. The position I adopt here, which will be the background for the discussion of atomic versus decompositional ad hoc concepts in section 4, is that a logical-encyclopaedic distinction can be drawn, but that none of the information – logical or encyclopaedic – is constitutive of the concept.

Finally in this section, I'll briefly discuss the main problem perceived with atomism, which is that it appears to entail radical concept nativism: that is, the thesis that all atomic concepts, including COMPUTER, CARBURETTOR, BOOK, GIRAFFE are innate, because, being atomic, it is difficult to see how they could be acquired inferentially. Inferential acquisition seems far more plausible than either innateness (defended by Fodor 1975, 1990) or some kind of brute-causal acquisition (e.g. that defended by Fodor 2008), and many researchers have thought that only by positing an at least partly definitional or otherwise decompositional structure to concepts would it be possible to explain how they are inferentially acquired. Nevertheless, several millennia of trying to define lexical concepts have produced plausible, complete definitions for little more than BACHELOR, suggesting strongly that the premise that lexical concepts have definitions is false. Fodor and others have given several other well-known arguments against definitions and against more recent 'prototype' views of concepts having internal structure (see Fodor et al 1980 and papers in Margolis and Laurence 2000)[4].

---

[3] For example, when the word "bug" is uttered in a context heavily biased towards the 'hidden microphone' sense, core features of the other linguistically encoded sense, such as INSECT, would not remain so highly activated after disambiguation has occurred (Swinney 1979).

[4] Some recent versions of decompositionalism recognize the problems with classical definitions: Jackendoff (2002) rejects the assumption that definitions must involve word meanings being built by combining other word meanings, and this opens the way to a more sophisticated decompositional theory on which word meanings could be composed not only of linguistically expressed definitions, but also of imagistic elements, spatial components, and other features. Such an account promises to handle some of the cases Fodor used to argue against classical definitions, such as how to define "red", which presumably would have to decompose into the primitive COLOURED, plus something else that distinguishes red from other colours but does not itself contain COLOURED as a component of its meaning. The problem lies in defining the "completer" linguistically, but, if Jackendoff is correct, then a visual representation (probably of a prototypical shade of the colour) could play this role. What remains to be established, though, is whether this is really a comprehensive solution to the problem of specifying completers, and whether it avoids the problem encountered with linguistic definitions: that the completer (whether linguistic or not) makes redundant some of what, it is agreed, ought to be primitives.

Several authors have taken the arguments against decompositionalism and prototype theory to show that atomism is the only remaining option, in which case, atomic concepts must be learnable, and there has been much work in the last 15 years on this issue: see, for example, Margolis (1998), Laurence and Margolis (2011), Weiskopf (2008). I cannot assess this literature here, but these authors provide several options for accounts of concept acquisition that meet the rationality requirement for counting as cases of learning, raising the prospects of a solution to the major obstacle to atomism.

## 3 Ad hoc concepts and decompositionalism

Assuming that lexical concepts are atomic, I now turn to the nature of the ad hoc concepts that are inferred and contribute to explicit utterance content. It has largely been assumed in the relevance-theoretic literature that the results of lexical modulation, like the input, are atomic concepts. To reflect the fact that they are different from but related to the lexical concept, the notational convention is to use stars to indicate the ad hoc nature of the concept. So, on different occasions of use, the word "bachelor", which encodes BACHELOR, might express different ad hoc concepts BACHELOR*, BACHELOR**, BACHELOR***, etc, any of which might be narrower, looser, or both narrower and looser than BACHELOR (where the number of stars are merely to distinguish different concepts – not to indicate degree of similarity to the lexical concept).

I've said, in section 1, that many so-called "ad hoc concepts" are not really ad hoc, in that they are likely to be already stable entities stored in the subject's memory. This would be the case with many of the examples discussed in the literature, including many metaphorical uses that are conventionalized to some extent:

(1) a. She is <u>happy</u>/<u>upset</u>/<u>tired</u>.
b. a <u>flat</u> country/a <u>raw</u> steak/a <u>square</u> face.
c. My neighbour's an <u>angel</u>/a <u>saint</u>/<u>dragon</u>/<u>witch</u>/<u>battleaxe</u>/<u>sergeant major</u>.

It seems likely that most of us will have previously encountered a range of related but different concepts that we would express with the word "happy", that correspond to various states of mind, and that are more specific than any abstract, general concept that "happy" might encode as its lexical meaning. Such concepts may well be used frequently enough to have become stable inhabitants of our memories. The same goes for the other underlined terms in (a), for which a more specific concept of upsetness/tiredness, would generally be retrieved on any given occasion of utterance, and for (b), where the underlined words would be interpreted loosely. Several, if not all, of the metaphors in (c) are fairly frequent and conventional: interpreting them might, in effect, be more a disambiguation process than a process of constructing a novel concept. Many, and perhaps the overwhelming majority of, cases of ad hoc concepts will be stable mental entities: that is, they are concepts proper, but not actually ad hoc – they are just different from the lexical concept. For these concepts, then, there is no reason to assume that they are any different in nature from lexical concepts: if atomism is a plausible view for lexical concepts, then these pragmatically inferred concepts can also be assumed to be atomic. The atomism versus decompositionalism question here is, then, relevant only to genuinely ad hoc concepts – that is, one-off (or perhaps just so infrequent as never to stabilize in the subject's memory) entities that are constructed to underwrite a set of inferences (implicatures) that is highly specific to a particular occasion of utterance.

Examples of such genuinely ad hoc concepts would include certain novel metaphors, or even quite conventionalized metaphors that the subject encounters for the first time, such as "cactus" or "prickly" or "spiky" used of a person (the first being a novel metaphor used in Rubio Fernandez's 2007 study, while the latter two are fairly conventionalized). In each case, whether a hearer will form a new ad hoc concept will depend on whether he has previously encountered the characteristic in question enough that it will have become a stable concept available for retrieval from memory; even if it has, encyclopaedic assumptions activated via tokening of the lexical concept CACTUS might cause him to conceptualize the denoted property differently, hence result in the formation of a genuinely ad hoc concept. Other cases might be novel category extensions. Wilson (2004) mentions cases such as "Federer is the new Sampras" and "Iraq is this generation's Vietnam", where the idea is that "Sampras" and "Vietnam" are used to denote not the individual and the particular war, but broader categories, paraphrasable respectively by descriptions such as 'supremely skilled tennis player who dominated his generation' and 'disastrous military intervention'. Those particular broader categories may well be familiar ones, corresponding to stable concepts, but this highly productive "X is the new Y" paradigm regularly offers fresh candidates for genuinely ad hoc concepts. For example, some hearers may not previously have found use for concepts denoting the categories 'fashionable hot alcoholic drinks for winter' or 'pretty Hollywood it-boy who specializes in playing troubled young men on the big screen'[5], which seem to be approximately what the final nouns, respectively, express in the following two cases (both are thanks to Guardian Weekend magazine's 'The Measure' series, October-November 2011):

(2) Winter Pimm's. With hot apple juice. The new mulled wine.
(3) Is Michael Fassbender the new Ryan Gosling?

Another possible case, offered by Carston (2010), is "She asked questions frantically and buried all my answers". This use of the word "bury" to denote the action of dismissing or refusing to process someone's comments, may be new to hearers, in which case it could express a genuinely ad hoc concept BURY*, even though it may be modelled on similar uses such as "This is a good day to bury bad news" (which itself may have given rise to a genuinely ad hoc concept BURY** the first time a hearer encountered such a use). There is reason to think, then, that not only are stable non-lexicalized concepts retrieved in utterance comprehension to replace the lexical meanings in the proposition expressed, but that genuinely novel ad hoc concepts are sometimes formed for specific purposes and then discarded. It is these genuinely ad hoc concepts, and the process of inferring them – 'modulation' – that I will be referring to from now on.

Some recent work questions whether the account of modulation, illustrated in section 1, as manipulation of logical/encyclopaedic properties is reconcilable with relevance theory's commitment to concept atomism. Vicente and Martinez-Manrique (2010) and Allott and Textor (2011) suggest that, even assuming lexical and other stable concepts are atomic, ad hoc concepts must be a very different kind of entity. To start with, they are not, strictly, concepts, as explained in footnote 2. Furthermore, the logical-encyclopaedic distinction as seen within RT, where logical information is spontaneously inferred but encyclopaedic information is not, clearly only applies to re-uses of already stored, hence stable concepts. In the case of stable concepts, none of the associated information (logical or encyclopaedic) need be constitutive or reference-determining. This does not seem to make sense for ad hoc

[5] As with all such paraphrases provided in the RT ad hoc concepts literature, these are intended merely to roughly indicate to the reader what the author has in mind. They should not be taken as suggesting that the ad hoc concepts decompose into the features listed. (See Carston 2010, section 4.)

concepts which, in contrast, would simply be constituted by the 'clusters' of properties that are activated and underwrite inference on the particular occasion: a different cluster of properties would result in a different concept. This form of decompositionalism seems to be, for those authors, the only way that RT talk of 'dropping logical properties' (in loosening) and 'promoting a property to content-constitutive (logical) status' (in narrowing) can be made intelligible. As Vicente and Martinez-Manrique say, 'The atomist owes an explanation of modulation processes consistent with her atomism' (2010).

## 4 Atomic ad hoc concepts?

While decompositionalism may seem a plausible approach to ad hoc concepts, what I want to do in this section is suggest that atomism cannot be dismissed so easily. For one thing, talk of properties being 'promoted' to or 'dropped' from the logical entry does not necessarily imply manipulating the internal structure of the concept, since these logical and encyclopaedic properties are not part of the internal structure of the atomic concept (the lexical concept) that is the starting point of the process. What actually makes two atomic concepts different concepts (whether stable or ad hoc) is simply that they are different symbols in the language of thought, and this is independent of the associated logical and encyclopaedic information.

I agree that the advocates of a decompositional 'cluster' structure must be correct that the ad hoc concept is formed from the properties that are activated and licence inference to implications that fulfil the speaker's communicative intention. Allott and Textor, and Vicente and Martinez-Manrique take this to mean that the activated properties constitute the concept, and this decompositionalism is endorsed by Carston (2011), at least for temporary, one-off entities. This is not the only option, however. What is needed here is an account of how atomic ad hoc concepts can arise inferentially in the comprehension process, and below, I discuss why such an account cannot yet be ruled out.

First, consider the process of concept acquisition, and specifically, acquiring a concept for natural kinds. A rational learning account for atomic concepts must involve the subject inferring a concept from a cluster of perceptible properties. For example, if he observes that the following cluster of properties are manifested by some objects, then he would infer that they belong to the same category: [ANIMAL, CERTAIN SIZE, HAIRY, BARKS, .....]. Given humans' built-in assumption that natural kinds have unknown 'essences' that explain observable properties and that these essences are criterial for category membership (cf Medin and Ortony 1989), learning a concept will involve inferring from these 'surface' properties to the presence of some essential property, and assigning an arbitrary label (e.g. DOG) to the category of things that are assumed to have these properties. It is this label that is the concept – that is, the constituent of thought – and serves as a conceptual address under which the associated information can be organized and added to, or, if some of it is discovered to be wrong, replaced (e.g. FISH replaced with MAMMAL in WHALE's logical information. This mental label, itself unstructured and not constituted by the associated properties, is an atomic concept. Supposing that this process is not restricted to inferring a natural kind category from observable properties, then acquisition of other kinds of concepts would work in roughly the same way: acquiring a concept would involve, initially, assigning a mental label to a category of objects or properties that share certain other properties (which may be derived via perception or inference).

The construction of atomic ad hoc concepts would be similar to this initial stage of concept acquisition via inference. A cluster of properties is activated (largely from the logical and encyclopaedic entry of a lexical concept); if this cluster is not similar enough to a stored cluster of properties that correspond to an existing category, the hearer infers the presence of

a novel category, whose members share a particular essential feature, and assigns it a new mental label – i.e. an atomic concept. This concept is what forms part of the proposition expressed by the speaker. The formation of the concept would occur because of the subject's representation of the cluster of properties that are activated, and which would likely all have the same status (a logical-encyclopaedic distinction does not seem to apply here). But, as Rupert (2001: 524) says, speaking of the acquisition of stable concepts, that 'does *not* imply that said cluster defines the newly introduced term'. The same consideration applies to ad hoc concepts.

Within various frameworks, suggestions exist as to how the mind coins such terms: one that is perhaps the most compatible with LOT is that of Changeux (1983/85), who describes how 'prerepresentations' (or, the more standard term used by Rupert 2001, 'protorepresentations') briefly cohere; many pass away, but any further experience or use causes them to begin to stabilize. Among ad hoc concepts, there is likely to be a continuum of degrees of stability, with some clusters of properties being highly unstable and temporary, and perhaps conforming to the decompositional model suggested by Allot and Textor, Vicente and Martinez Manrique, and Carston. The creation of an atomic ad hoc concept would require the system going beyond merely representing these properties, to coin a new primitive in the language of thought (a file label or conceptual address) prompted by the inference that the activated properties are the features of a new category for which the hearer does not already possess a concept.

Why think that the inference system should coin these new, highly temporary LOT terms? The pragmaticists mentioned above seem to dismiss the idea, but there are several reasons why it perhaps should not be ruled out so quickly.

First, this process, or something very similar, occurs all the time anyway. Not only must it be very common in development, when the child is forming many new protorepresentations many of which will stabilize into concepts proper, but we continually coin new demonstrative-like terms in LOT to keep track of objects in our environments. Some of these concepts will be highly temporary, existing only for the duration of our attention to the object. Still, each tokening of a new mental demonstrative constitutes opening a new mental file on an object, the label of which is the concept, and the contents the associated information – consisting, initially, of just whatever group of properties is activated on that occasion. These singular concepts and the kind of ad hoc concepts that I've been discussing in this paper are merely different species of newly created, usually temporary, LOT terms (see Recanati 1993), allowing us to think of the referents or denotations *as such*, rather than under any particular description.

Why should a temporary mental term be coined (and associated file opened) for these novel categories, which are only hypothesized, and do not need to be tracked like the referents of singular concepts? The inferential comprehension process may require some, albeit short-lived, stability. As illustrated in the discussion of many examples by Carston (2002) and Wilson and Sperber (2002), it is often the case that not only does the ad hoc concept form part of the explicature, but the hearer will be required to construct additional premises in order to warrant the intended implications. One of the examples given is "I'm tired", in response to the question "What do you feel like doing today?" This is an example of an utterance that would communicate a range of relatively weak implicatures (that is, an indeterminate set of implicatures, none of which may determinately be something that the speaker had in mind). Deriving these requires the construction of the ad hoc concept TIRED*, which is more specific than the lexical meaning of "tired", and its incorporation in a range of premises, such as IF X IS TIRED* THEN X DOES NOT WANT TO PLAY SQUASH TODAY, and so on. Like this one, the examples tend to involve concepts that are non-lexicalised but already possessed by speaker and hearer, but the point is that some stability – i.e. tokening a primitive

as a mental label for the activated cluster of properties – may be required for the hearer to scrutinize the proposition expressed for the range of implications that will fulfil the communicative intention. Here is where atomism provides the advantage of not requiring that the category be thought of under a particular description: the decompositional story seems to require that it be thought of, rather, as satisfying what will in some cases be an extensive list of descriptive information corresponding to the activated properties.

A closely related issue concerns the overall inferential nature of the comprehension process. The hearer develops hypotheses about explicatures (proposition expressed), contextual assumptions, and implicated conclusions in parallel (Wilson and Sperber 2002), and what makes the process inferential, according to Sperber and Wilson, is that the final interpretation is a valid argument, with explicatures and contextual assumptions warranting implicatures. Carston (2006) argues against Recanati's (2004) suggestions that the derivation of explicatures can be non-inferential, achieved purely by associative processes. Recanati (2006) concedes that, at least at the stage of 'externalizing' the explicature, which appears to mean confirming the hypothesized explicature as part of the overall warranted interpretation, explicature derivation requires mindreading and inference. A possible worry with the decompositional cluster account – but one that needs much more working out – is that it regresses to a non-inferential, merely associationist, account of modulation. This is because, as far as the inferential pragmatic processor is concerned, the 'cluster' is merely a random collection of activated properties. So the thought is that confirming the hypothesis about the explicature, for the modulation to count as an inferential process, requires more than just the presence of a set of properties, but also the inference that they are connected, hence inferring the existence of a novel category, for which a new mental name is coined.

Finally, one of the factors in favour of, or at least making possible, the cluster decomposition position, was that it appears not to make sense to talk of a logical-encyclopaedic distinction for genuinely ad hoc concepts. The idea is that the concept does not yet have a logical or encyclopaedic entry; all that is there are the activated properties, and they have not yet been organized under a conceptual address. So all the properties in the cluster can be considered constitutive. However, it is not obvious to me that a logical-encyclopaedic distinction of some kind is ruled out for ad hoc concepts. Consider the cases of category extensions discussed in section 3 (examples (2) and (3) and the preceding text). It was suggested there that some such cases may result in the formation of ad hoc concepts rather than merely retrieval of already stored concepts. A potential problem for the cluster account is that these are cases where the activated information plausibly includes a prototype. For example, in the "Federer is the new Sampras" example, the hearer's stored concept of the individual, Pete Sampras, is inevitably activated, and its role here would be as a prototype of the broader category, SAMPRAS*, of which he is, after all, a prototypical member. A prototype can, clearly, not be a constitutive feature of a concept alongside other defining information. So some structuring of the associated information is necessary, and the question for the cluster account is whether this is possible without the creation of a temporary mental name under which to arrange this information[6].

To conclude this short note, then, while all the arguments sketched above need much more thought, I hope to have provided some reasons why atomism should still be considered, alongside decompositionalism, as a live option for newly created ad hoc concepts.

[6] This may also be necessary for the suppression of unwanted logical inferences that are spontaneously activated when a stored concept is tokened (Rubio Fernandez 2007, discussed in section 2). If an inference from the relevant activated properties to a new category and concept is not required for this suppression, and the unwanted properties can be filtered out without it, then it is no longer clear what grounds there are for supposing that anything other than the lexical concept appears in the proposition expressed.

## References


Allott, N. and Textor, M. 2011. Lexical pragmatic adjustment and ad hoc concepts. Paper presented at Theoretical Pragmatics workshop, Berlin.

Blutner, R. 1998. Lexical pragmatics. *Journal of Semantics* 15: 115-162.

Boghossian, P. 1994. Inferential role semantics and the analytic/synthetic distinction. *Philosophical Studies* 73: 109-122.

Boghossian, P. 1997. Analyticity. In B. Hale and C. Wright (eds), *A Companion to the Philosophy of Language*, 331-368. Oxford: Blackwell.

Carston, R. 2002. *Thoughts and Utterances: The Pragmatics of Explicit Communication*. Oxford: Blackwell.

Carston, R. 2006. How many pragmatics systems are there? In M. Frapolli (ed.), *Saying, Meaning and Referring: Essays on Francois Recanati's Philosophy of Language*, 18-48. Basingstoke: Palgrave.

Carston, R. 2010. Explicit communication and 'free' pragmatic enrichment. In B. Soria and E. Romero (eds), *Explicit Communication: Robyn Carston's Pragmatics*, 217-287. Basingstoke: Palgrave.

Carston, R. 2011. Lexical meaning and concept communicated. Paper presented at Workshop on Semantics, Pragmatics and Rhetoric, ILCLI, Donostia.

Changeux, J. P. 1983/1985. *Neuronal Man*. Princeton NJ: Princeton University Press.

Fodor, J. 1975. *The Language of Thought*. New York: Thomas Crowell.

Fodor, J. 1990. *A Theory of Content and Other Essays*. Cambridge, MA: MIT Press.

Fodor, J. 1998. *Concepts. Where Cognitive Science Went Wrong.* Oxford: Clarendon Press.

Fodor, J. 2008. *LOT 2*. Oxford: OUP.

Fodor, J., Garrett, M., Walker, E. and Parkes, C. 1980. Against definitions. *Cognition* 8: 263-367.

Horsey, R. 2006. *The Content and Acquisition of Lexical Concepts*. PhD thesis, UCL.

Jackendoff, R. 2002. *Foundations of Language*. Oxford: OUP.

Margolis, E. 1998. How to acquire a concept. *Mind & Language* 13: 347-69.

Margolis, E. and Laurence, S. 2000. *Concepts: Core Readings*. Cambridge, MA: MIT Press.

Margolis, E. and Laurence, S. 2011. Learning matters: the role of learning in concept acquisition. *Mind & Language* 26: 507-539.

Medin, D. and Ortony, A. 1989. Psychological essentialism. In S. Vosniadou and A. Ortony (eds), *Similarity and analogical reasoning*, 179-195. New York: CUP.

Pietroski, P. 2005. Meaning before truth. In G. Preyer and G. Peters (eds) *Contextualism in Philosophy: Knowledge, Meaning and Truth*, 255-302. Oxford: OUP.

Rayo, A. 2011/forthcoming. A plea for semantic localism. *Noûs*.

Recanati, F. 1993. *Direct Reference: From Language to Thought*. Oxford: Blackwell.

Recanati, F. 2004. *Literal Meaning*. Cambridge: CUP.

Recanati, F. 2006. Reply to Carston. In M. Frapolli (ed.), *Saying, Meaning and Referring: Essays on Francois Recanati's Philosophy of Language*, 49-54. Basingstoke: Palgrave

Rubio-Fernandez, P. 2007. Suppression in metaphor interpretation: Differences between meaning selection and meaning construction. *Journal of Semantics* 24: 345-71.

Rupert, R. 2001. Coining terms in the language of thought. *Journal of Philosophy* 98: 499-530.

Schiffer, S. 2003. *The Things We Mean*. Oxford: OUP.

Sperber, D. and Wilson, D. 1986/1995. *Relevance: Communication and Cognition*. Oxford: Blackwell.

Sperber, D. and Wilson, D. 1998. The mapping between the mental and the public lexicon. In P. Carruthers and J. Boucher (eds) *Language and Thought: Interdisciplinary Themes*, 184-200. Cambridge: CUP.

Sperber, D. and Wilson, D. 2008. A deflationary account of metaphor. In R. Gibbs (ed.), *The Cambridge Handbook of Metaphor and Thought*, 84-105. Cambridge: CUP.

Swinney, D. 1979. Lexical access during sentence comprehension: (Re)consideration of context effects. *Journal of Verbal Learning and Verbal Behavior* 18: 645–660.

Vicente, A. and Martinez-Manrique, F. 2010. Relevance theory's atomistic commitments. In B. Soria and E. Romero (eds), *Explicit Communication: Robyn Carston's Pragmatics*, 42-58. Basingstoke: Palgrave.

Weiskopf, D. 2008. The origins of concepts. *Philosophical Studies* 140: 359-384

Wilson, D. 2004. Relevance and lexical pragmatics. *UCL Working Papers in Linguistics* 16: 343-360.

Wilson, D. and Carston, R. 2007. A unitary approach to lexical pragmatics: relevance, inference and ad hoc concepts. In N. Burton-Roberts (ed.) *Pragmatics*, 230-259. Basingstoke: Palgrave.

Wilson, D. and Sperber, D. 2002. Truthfulness and relevance. *Mind* 111: 583-632.