ISSUES IN PRAGMATICS (PLIN 3001) 2006-07

LEXICAL PRAGMATICS

6. Word meaning and the semantics/pragmatics distinction: (a) Lexical narrowing

1. Introduction

We started this term by looking at a range of data which suggest that words are often used in ways that depart from their linguistically-encoded 'literal' meanings. In particular, these literal meanings can be **narrowed**, as in (1), used **approximately**, as in (2) or **metaphorically**, as in (3):

(1a) All doctors *drink.* ('drink liquid', 'drink alcohol', 'drink a lot of alcohol')
(1b) Mary is a *working mother.* (Lakoff: 'stereotypical working mother')
(1c) I have a *temperature*. ('higher than normal temperature')

(2a) This coat cost *1,000* dollars. ('about 1,000 dollars') (round numbers)
(2b) The stones form a *circle, an oval, a pyramid*. ('approximately a circle') (geometric)
(2c) This injection will be *painless*. ('nearly painless') (negatively defined terms)

(3a) Mary is *a rose, a lily, a daisy, a violet; a jewel, a diamond, a ruby, a pearl.*
(3b) The leaves *danced* in the breeze.
(3c) That book *puts me to sleep.* (literal? approximation? hyperbole? metaphor?)

Our main aim was to find an account of these examples that might also extend to deal with **category extensions** such as (4), and **neologisms** such as (5):

(4a) Roger Federer is the next *Sampras*. (used at Wimbledon 2003)
(4b) Iraq is this generation's *Vietnam*.
(4c) Brown is the new *black.* (statement in a fashion magazine)

(5a) The newspaper boy *porched* the newspaper.
(5b) They *Learjetted* off to Miami.
(5c) He *Houdinied* his way out of the closet.

I suggested that all of (1)-(5) might be reducible to two main types of case: (a) **lexical narrowing**, where a word is used in a more **specific** sense than the linguistically encoded one, with consequent narrowing of the linguistic denotation, and (b) **lexical broadening**, where a word is used in a more **general** sense, with a broader denotation.

In the next few lectures, we looked at a range of broader issues about how word meanings are represented and acquired, and more generally about the nature and acquisition of concepts and their role in categorisation and inference. Here are some of the conclusions we drew:

(6a) Most lexical concepts are simple and unanalysable, hence **triggered** rather than learned. The innateness involved is of **mechanisms** for concept construction rather than actual concepts.

(6b) Concepts give access to various types of mentally-represented information, **lexical**, **logical** and **encyclopaedic**, which may affect their processing and use.

(6c) Word use creates a variety of **prototypicality effects**. These may be explained in **pragmatic** rather than semantic terms, based on the choice of different subsets of assumptions from the encyclopaedic entries of the encoded concepts.

At the end of the lecture on prototypicality effects, we looked at Barsalou's claim that comprehension often involves the construction of ad hoc concepts, using information derived from the encyclopaedic entries of existing concepts as outlined in (6c). He suggested that the concept construction process was constrained by considerations of **discourse context, accessibility** and **relevance**. For the next few lectures, I want to pursue this idea and try to show how the lexical-pragmatic processes in (1)-(5) may be explained using ad hoc concept construction under relevance-based constraints. The assumption is that every time a word is used, it may convey a slightly different ad hoc concept, which is not encoded but pragmatically derived.

This account (if successful) would have broader implications of at least three types. First, there are implications about the relation between lexical semantics and pragmatics (or word meaning and word use). Before Grice, ordinary-language philosophers tended to defend a **use** theory of meaning on which facts about word use are seen as giving direct insight into word meaning, and the lexicon would consist of a huge range of descriptive facts about word use. Grice (in 'Logic & conversation', lecture 1) was one of the first to argue that many facts about word use do not give us direct insight into word meaning, but follow from more general pragmatic principles. On this approach, rather than trying to record every fact about word use in the lexicon, we would combine a relatively simple semantics with a general account of pragmatics or language use. Notice that prototype semantics is a return to the old pre-Gricean position on which all facts about word use give direct insight into word meaning. I want to show how the Gricean argument can be taken further, using a relatively simple semantics plus the pragmatic creation of ad hoc concepts to account for prototypicality effects. (The argument goes beyond Grice's because it treats lexical-pragmatic processes as contributing not only to **implicatures** but also to explicit truth-conditional content, i.e. the **proposition expressed**.)

Second, there are implications about the relative sizes of our stock of words and our stock of concepts. Recall that according to the classical empiricist view, there are many more words than concepts: a few simple, unanalysable concepts combine and recombine to yield the meanings of a

much larger stock of words. On Fodor's view, by contrast, there are roughly as many words as concepts: the meanings of most morphologically simple words are simple, unanalysable concepts, so there is roughly a one-to-one correspondence between words and concepts. The relevance-theoretic approach suggests that both Fodor and the classical theory are wrong. In fact, there are vastly many **more** concepts than words, because new ad hoc concepts can be created and used in the interpretation of almost any word.

This in turn has implications for innateness. If new ad hoc concepts can be created at will, it seems obvious that they can't themselves be innate in Fodor's strong sense. In the last two lectures, I'll return to issues of innateness and acquisition, and continue the argument that many of Fodor's insights can be preserved by adopting a rather weaker position, on which it is not the concepts themselves but the mechanisms for constructing them that are innate. What makes this approach interesting, and links it to the literature on modularity and innateness, is the increasing evidence that there are different **kinds** of concepts, and innate **domain-specific** concept construction mechanisms which go well beyond what any general-purpose learning theory would provide.

2. Existing accounts of narrowing

Many approaches to lexical pragmatics take for granted that narrowing, approximation and metaphorical extension are distinct pragmatic processes, which lack common descriptions or explanations and need not be studied together. For example, narrowing is often analysed as a case of **default inference** to a **stereotypical,** or prototypical, interpretation. I'll look at the account in Levinson (2000:37-8; 112-34) to illustrate (see also Blutner 1998; 2002; 2004).

Levinson (2000:37-8, 112-34) is a 'neo-Gricean' pragmatist who uses a modified set of maxims to deal with 'generalised implicatures', i.e. implicatures carried in 'all normal contexts'. The idea is to treat these generalised implicatures as **default inferences**, i.e. inferences which are automatically generated, but may be cancelled in context. This allows for a relatively code-like approach to lexical pragmatics. Levinson deals with various enrichment and narrowing processes using an Informativeness Principle, which gives rise to a more specific Informativeness Heuristic:

Informativeness Principle
Amplify the informational content of the speaker's utterance, by finding the most **specific** interpretation, up to what you judge to be the speaker's m-intended point ... [ibid: 114]
[the speaker's 'm-intended point' is the speaker's overtly intended meaning]

Informativeness heuristic
What is expressed simply is stereotypically exemplified.

The I-heuristic accounts for stereotypical narrowings such as (1b) above (*Mary is a working mother*), while the more general I-Principle deals with other types of narrowing, e.g. (1c) (*I have a temperature*).

This account seems to me to run into serious problems. First, narrowing is a highly flexible and context-dependent process. For example, there may be several possible degrees or directions of narrowing, as in (1a) (*All doctors drink*), where *drink* may be narrowed to 'drink alcohol' or 'drink a lot of alcohol'), or in (8), where *bird* is narrowed in different ways in different contexts:

(8a) As I worked in the garden, a *bird* perched on my spade.
(8b) *Birds* wheeled above the waves.
(8c) A *bird*, high in the sky, invisible, sang its pure song.
(8d) At Christmas, the *bird* was delicious.
(8e) John opened the birdcage, and the *bird* flew across the room.

Or consider *Englishman* in (9), which would evoke different stereotypes in a discussion of cooking, cricket, sailing, seduction, etc:

(9) John is an *Englishman.*

More generally, Barsalou's experiments on prototypicality effects show that narrowing **can't** be accounted for by appeal to ready-made stereotypes, because there are subtle differences in interpretation and typicality rankings in different contexts. Now the main appeal of default inference is that it **isn't** flexible, and can therefore be approached in code-like terms, and seen as yielding '**generalised** implicatures' which are carried 'in all normal contexts'. But there is a huge amount of evidence (e.g. the examples above) to show that narrowing **is** highly context-dependent and flexible, and **can't** therefore be approached in code-like terms.

Second, Barsalou's experiments also showed that not all narrowing is stereotypical (for an interesting survey of the various types of narrowing, see Lakoff 1987). For example, in (10), *man* might be narrowed to an **idealised** rather than a stereotypical interpretation, indicating that Churchill is a man worthy of the name rather than a typical man:

(10) Churchill was a *man*.

An adequate account of narrowing should explain how hearers decide on the appropriate degree, direction and type of narrowing on each occasion of use. But what does the I-Principle tell us? It

tells us that the hearer of (1) and (8)-(10) should find an appropriate narrowing using his judgement about 'the speaker's m-intended point'. The 'm-intended point' is another word for the speaker's meaning, so all the I-principle really tells us is 'narrow until you think you've got the intended interpretation'. In other words, the I-principle doesn't explain how hearers **find** the speaker's meaning. It **presupposes** that they have some other way of working out the speaker's meaning, which is not given in this account (or anywhere else in Levinson's work).

Levinson (ibid.:118) acknowledges the context dependence of I-implicatures, but still claims that 'at a sufficient level of abstraction' they are default inferences which 'hold as preferred interpretations across contexts, and indeed across languages.' But this claim also runs into problems. As we've seen in (2)-(5) above, at the lexical level, broadening appears to be just as strong a tendency as narrowing, and it is not clear why narrowing rather than broadening should be seen as the default. Nor is it clear how an approach based on the I-principle and default inference could generalise to approximations or metaphors (which Levinson treats as blatant violations of a maxim of truthfulness in the regular Gricean way)

A final problem with the accounts of narrowing by people like Levinson and Blutner is that they retain Grice's assumption that the only way that pragmatic processes can contribute to utterance interpretation is at the level of **implicature**. Thus, Blutner (1998: 140) defines the goals of Lexical Pragmatics as involving 'the combination of a compositional semantics with a general mechanism of conversational implicature'. In the next section, I'll argue that many cases of concept narrowing can't be analysed in this standard Gricean way, because they affect the truth conditions of utterances: that is, they fall on the **explicit** rather than the implicit side.

3. <u>A relevance-theoretic approach to narrowing</u>

Let's look first at a clear case where narrowing affects the truth conditions of utterances, and must therefore be dealt with at the level of the proposition expressed (in Grice's terms, **what is said** rather than what is **implicated**). (In fact it's a case of reference resolution rather than lexical narrowing, but it has the same logical structure as narrowing, with a set of logically related interpretations among which the hearer has to choose. Suppose we meet in the College Library, and I say:

(11) I've been here all week.

Here in general means 'the place where the speaker is' – but where is that? In uttering (11), I might mean that I've been *in this room*, *in this library*, *in this building*, *in this town*, *in this country*, etc. all week. The range of possible interpretations has a subset structure: the cases in

which I've been *in this room* all week are a subset of those in which I've been *in this library, this college, this town*, etc., all week. It's particularly clear in this case that the hearer's choice of interpretation affects the truth conditions of the utterance: (11) has different truth conditions depending on whether the speaker is taken to mean *in this room, this town, this country*, etc.

What's interesting about this set of interpretations from the relevance-theoretic point of view is that the hearer's choice of interpretation affects not only the proposition expressed by the utterance, but also its cognitive effects. The narrower the interpretation, the greater the cognitive effects. For example, the interpretation on which the speaker of (11) is saying that she's been *in this room* all week entails the interpretation on which she's been *in this building, in this town, in this country*, etc., all week. It will thus achieve all the cognitive effects obtainable from these less specific interpretations, and more besides. We might thus diagram the range of possible interpretations of (11) in the following way:

As we narrow the interpretation of *here*, we obtain successively greater sets of cognitive effects.

Notice, now, that this is exactly the logical structure we find in cases of lexical narrowing. We have a range of increasingly narrower interpretations (as in (1a-c)), and as we narrow the 'literal' meaning we get greater cognitive effects. Which of these logically related interpretations should the hearer choose? Recall that, according to Levinson's I-Principle, the hearer should narrow 'up to what he takes to be the speaker's m-intended point', but no guidance is given on how he should recognise this when he finds it. It seems that relevance theory can do better than that.

According to relevance theory, utterances raise **expectations of relevance**, and the

comprehension process involves **mutual adjustment** of explicit content, context and cognitive effects in order to satisfy the particular expectations of relevance raised by the utterance. More precisely, the hearer has a **presumption of optimal relevance:**

Presumption of optimal relevance
(12a) The utterance will be **at least relevant enough** to be worth processing.
(12b) It will be the **most relevant** one compatible with the speaker's abilities and preferences.

and can use the following comprehension procedure to find an appropriate interpretation:

Relevance-theoretic comprehension procedure
Follow a path of least effort in looking for cognitive effects.
(13a) Consider interpretations (e.g. **narrowings**, **contexts**, **implications**) in order of accessibility;
(13b) Stop when your **expectation of relevance** is satisfied.

An adequate account of narrowing should explain what triggers it, what direction it takes, and when it stops. Here are the answers given by relevance theory. First, narrowing is triggered by the search for relevance. Second, it follows a path of least effort, choosing the most accessible set of contextual assumptions, the most accessible narrowing, the most accessible set of contextual implications. Third, it involves **mutual adjustment** of contextual assumptions, explicit content and cognitive effects in order to satisfy the particular expectations of relevance raised by the utterance; and fourth, it stops when these expectations of relevance are satisfied, or abandoned. (Mutual adjustment involves choosing from a range of possible contextual assumptions, implicatures and implicatures in such a way that (a) the chosen contextual assumption and explicatures together logically imply the implicatures, and (b) there are enough implicatures to satisfy the hearer's expectation of relevance.)

Let's illustrate how this approach might apply to the case of (1c). Imagine Mary is responding to Peter's suggestion that they visit his aunt in hospital:

(1c) I have a temperature.

Strictly speaking, (1c) would be true if Mary had a temperature of normal or below. However, it would typically not satisfy Peter's expectation of relevance on this interpretation: he will be expecting a response to his suggestion that they visit his aunt in hospital. Now in the first place, we usually take for granted that someone's temperature is at or near normal, so this information would not be relevant enough to be worth our attention. In the second place, it's hard to see what implications this response would have in a context Peter has easily available. Assuming a

spreading activation model of memory, in this particular situation Peter's assumptions about temperatures, hospital visits and the connections between them would be most highly activated at this point. Thus, by following a path of least effort in mutually adjusting contextual assumptions, explicatures and implicatures, he is likely to interpret (1c) as communicating that Mary has a temperature high enough to make it inadvisable for her to visit his aunt in hospital. This does not, of course, mean that the word *temperature* **literally means** 'above normal temperature'. It just shows that a word which encodes one concept can be used to express another. More generally, what triggers narrowing is the search for relevance. Why do we narrow in this case by going up the temperature scale rather than down? Because we have easily accessible encyclopaedic assumptions about the implications of raised temperatures, and much less easily accessible assumptions about the implications of low temperatures. How do we narrow? By mutually adjusting context, explicatures and implicatures in order to satisfy expectations of relevance. When do we stop? When our expectations of relevance are satisfied (or abandoned).

Moving on to the cases that particularly interest those who argue against the classical view of concepts and in favour of fuzzy or prototypical meanings, consider (1b):

(1b) Mary is a working mother.

As Lakoff (1987: 80-82) points out, (1b) would generally be understood as communicating more about Mary than that she is a female parent who works: it might suggest, for example, that Mary's children are not grown up and living away from home, that she did not give them up for adoption at birth; that she not only works but works for money, for more than an hour or two a week, and so on. In other words, (1b) is typically understood as applying to some subset of the people who satisfy the definition 'female parent who works'. As we've seen from the literature on prototype effects, similar cases of concept narrowing occur in the interpretation of virtually every utterance. In the light of relevance theory, it's easy to see why. If the concept encoded by an expression (which we're here assuming denotes the set of 'female parents who work') does not yield enough effects to satisfy the hearer's expectation of relevance, he should narrow it down, thus increasing its effects, to a point where it does yield enough effects. Narrowing in the direction of a stereotype is a particularly easy way to go. However, this is a fact about pragmatics, or language use, and need have no implications for lexical semantics: *working mother* simply denotes the set of female parents who work, and will be narrowed in different ways on different occasions.

More generally, when you look at language, you find scales and nested sets of interpretations everywhere: e.g. *It's raining* has different implications depending on how hard it is

raining, how long it rains, etc. Sapir claims that this scalarity pervades the whole vocabulary:

‘Any two houses selected at random offer the contrast of *more* or *less* on hundreds of features which are constitutive of the concept “house”. (Sapir 1944: 94)

If this is true, relevance theory has something to offer that other pragmatic theories aren’t even attempting: it offers a genuine account of how we decide when to narrow, in what direction to narrow, and when to stop.

4. Relevance and ad hoc concepts

How exactly does concept narrowing work? I’d like to end by looking a little more closely at the idea that humans have the ability to construct ad hoc concepts on the model of existing concepts, using contextual assumptions derived from encyclopaedic entries of these existing concepts, with different denotations and different cognitive effects.

As noted above, the assumption made in much of the semantic and psychological literature is that humans have a finite stock of concepts, and that each of them is the encoded meaning of some word. This assumption seems to be a hangover from the code model of communication, which claims that the only way to communicate a concept is to encode it. Although the inferential model has now largely replaced the code model in the areas of implicature, disambiguation and reference assignment, its implications for the study of lexical pragmatics are still largely unexplored.

As far as I can see, once we abandon the code model and move to the inferential model, there is no reason (apart from Fodor’s argument from innateness) to assume that every human has a fixed stock of concepts, each of which is the encoded meaning of some word. Dan Sperber and Deirdre – following psychologists like Barsalou – have been exploring the idea that the mind is much more flexible and creative: it can construct and use new concepts at a moment's notice (generally on the model of existing concepts) – say, to represent a particular blade of glass, a particular sensation, a particular cloud. And one of the ways this flexibility is exercised is in communication and understanding. (Sperber & Wilson 1998, Carston 2002; Wilson & Sperber 2002)

Such flexibility is only to be expected on an inferential account of communication. On an inferential account, the concept encoded by a word should merely act as a starting point for inferential comprehension: it can be narrowed, broadened, and played around with in other ways – as long as we have rich enough pragmatic theory to explain (and constrain) the resulting interpretations. This is what the relevance-theoretic comprehension procedure should provide.

Here's a suggestion about how the machinery for constructing interpretations based on ad hoc concepts might work. In interpreting (8d), say, the hearer might open a new conceptual address BIRD* whose logical and encyclopaedic entries have still to be determined:

(8d) At Christmas, the bird was delicious.

Following the relevance-theoretic comprehension procedure, and using contextual assumptions based on the encyclopaedic entry of the existing concept BIRD, together with expectations about the level and type of cognitive effects to be achieved, he would start deriving cognitive effects, and when he has enough to satisfy his expectations of relevance, he would stop. The result of the mutual adjustment process would be an ad hoc concept BIRD* which differs from BIRD in two respects: (a) it has a much smaller encyclopaedic entry, consisting, say, of something like the typical assumptions about birds we eat at Christmas, and (b) it has a narrower denotation (picking out only the set of birds which satisfy those assumptions and enable these cognitive effects to be derived). Such a concept, once constructed, might be forgotten, or stored in memory for future use, so that two people could share and use a stable narrowed concept BIRD* which is not the linguistically-encoded meaning of any word of English.

The relevance-theoretic framework, with its notion of cognitive effect and processing effort, allows us to add a little more detail to how this construction process might go. (I sketched an account briefly with the example of *temperature* above.) First, which subset of encyclopaedic assumptions should the hearer test first? Answer: the most accessible ones. Which encyclopaedic assumptions will be most accessible? Let’s assume a spreading-activation model of memory, in which the mention of CHRISTMAS will add a layer of activation to assumptions in the encyclopaedic entry of BIRD which contain some reference to Christmas (e.g. robin, turkey, goose), and the mention of DELICIOUS adds an extra layer of activation to edible birds (e.g. turkey, goose, grouse). The result will be a double dose of activation for the assumptions involving birds traditionally eaten at Christmas – which will therefore be the most accessible assumptions to use in constructing the ad hoc concept. (Another source of activation would be particular expectations about the type of cognitive effects to be achieved, which would activate assumptions likely to yield those effects). Next, how does the hearer test this interpretation? Answer: by seeing whether, by adding its conceptual address to the proposition expressed, and adding its encyclopaedic entry to the context, he achieves enough cognitive effects to satisfy his expectation of relevance. If so, he accepts it; if not, he looks again.

The claim is, then, that a word which linguistically encodes a certain concept doesn't

necessarily **express** that concept. Instead, it may express another, more specific concept, constructed ad hoc as part of the interpretation process, which is determined by considerations of relevance. On this account, many prototypicality effects that we looked at in the lectures on prototypes are explainable as the result of lexical narrowing. Next week, we'll look at some further effects that result from lexical broadening, as in (2)-(5).

Homework

1. Come prepared to criticise or defend Levinson's account of narrowing.

2. How would *tired* in (b) be understood on a mutual adjustment account?
   (a) Peter: Would you like to come to the cinema?
   (b) Mary: I'm *tired.*

Reading

Levinson, S. 2000. *Presumptive Meanings.* CUP: 37-8, 112-34.
Carston, R. 2002. *Thoughts and Utterances.* Chapter 5, pp. 320-328.
Wilson, D. 2003. Relevance and lexical pragmatics. *Italian Journal of Linguistics* 15.2: 273-291

Background references

Lakoff, George 1987. *Women, Fire and Dangerous Things,* Chicago University Press.
Blutner, R. 1998. Lexical pragmatics. *Journal of Semantics* 15: 115-162.
Blutner, R. 2004. Pragmatics and the lexicon. *Handbook of Pragmatics.* Blackwell.
Wilson, D. 1998. Linguistic structure and inferential communication. *Proceedings of the 16th International Congress of Linguists,* Paris.
Wilson, D. & D. Sperber. 1998. Pragmatics and time. In R. Carston & S. Uchida (eds) *Relevance theory: Applications and implications.* Benjamins
Sapir, E. 1944. Grading: A study in semantics. *Philosophy of Science* 11: 93-116.
Sperber, D. & Wilson, D. 1998. The mapping between the mental and the public lexicon. In P. Carruthers & J. Boucher (eds) *Language and Thought.* CUP, Cambridge: 184-200.
Carston, R. 1996. Enrichment and loosening: Complementary processes in deriving the proposition expressed? *University College London Working Papers in Linguistics* 8.