

ISSUES IN PRAGMATICS (PLIN 3001) 2006-07

LEXICAL PRAGMATICS


2. Approaches to lexical semantics

1. Introduction

Last week, we looked at a variety of ways in which the concept **communicated** by use of a word may differ from the concept **encoded**. I suggested as a hypothesis that most of these departures can be reduced to two main types of lexical pragmatic process: (a) **narrowing** of the linguistically encoded word meaning, and (b) **broadening** of the linguistically-encoded word meaning. In the second half of term, we'll consider how this hypothesis might be developed and tested. First, though, we need to think a little more about lexical semantics, and about the nature of concepts. What do word meanings look like, and how are they acquired?

Let's start with the simple model of lexical semantics that I presented last week:

Simple model of lexical semantics:



The claim tacitly made in this model is that word meanings are **concepts**, and that concepts represent **categories** of objects, events or properties in the world (or in alternative possible worlds). On this account, the concept BIRD may be activated in your mind in two different ways: by hearing the word 'bird', or by seeing a bird. The study of concepts can therefore be approached from two different directions: (a) we can start from the linguistic end, and investigate what properties concepts would have to have to function adequately as word meanings, or (b) we can start from the cognitive end, and investigate the role of concepts in categorisation and cognition. Approach (a) is taken by lexical semantics. Approach (b) is taken by cognitive psychologists working on problems of categorisation. The best treatment of concepts will be one that responds to both types of evidence, taking both linguistic and cognitive factors into account.

For a long time, it seemed that a single theory of concepts could explain both their

role as word meanings and their role in cognition and categorisation. This was the **classical** (or **empiricist**, or **decompositional**) theory, which claims that a small set of basic concepts, typically acquired via the senses, may be combined by logical operations such as conjunction and negation to define an infinite variety of more complex concepts. On this approach, the concept of a bachelor would be a conjunction of more basic concepts such as ADULT & HUMAN & MALE & NOT-MARRIED (which might in turn decompose into more basic, sensory concepts). This complex concept would function, on the one hand, as the meaning of the word ‘bachelor’, and on the other, as a set of necessary and sufficient conditions enabling us to recognise bachelors as distinct from other things. How do you understand the word ‘bachelor’? You recover the definition ADULT & HUMAN & MALE & NOT-MARRIED. How do you know that someone is a bachelor? You tick off the defining features ADULT & HUMAN & MALE & NOT-MARRIED. On this classical (empiricist) approach, the meanings of most words are **definitions**, and objects, events or properties are categorised by checking them against their definitions.

The empiricist version of the classical account of concepts is also a theory of **concept acquisition**. It claims that simple concepts are acquired through the senses, and the more complex concepts built up out of simple concepts are therefore indirectly acquired through the senses too. Thus, the concepts CHAIR, HORSE, BUILDING, UNIVERSITY, etc. are built up out of simpler concepts representing perceptual features such as SHAPE, SIZE, TEXTURE, COLOUR, SOUND and so on, combined to form a definition of the objects, events or properties they represent. The empiricist doctrine is, essentially, that ‘nothing is in the mind that was not first in the senses’, so on this version of the classical account, all conceptual knowledge is ultimately based in perceptual representations.

In the last 25 years or so, this classical account of concepts has come under attack from at least three directions: as a theory of word meaning, as a theory of categorisation, and as a theory of concept acquisition. In semantics, Fodor has argued that word meanings are not definitions; in cognitive psychology, there is accumulating evidence that we do not categorise objects by checking them against their definitions; and in concept acquisition, there is increasing evidence that children are born with a considerable amount of conceptual knowledge that is not acquired through the senses. We need to consider some of these arguments, and discuss what alternative accounts we might put in its place. This week we’ll look at Fodor’s critique of the classical account, and the alternative he proposes.

2. Arguments against the classical account of word meanings

In 1975, a book and an article were published that substantially undermined the classical account of word meaning. The book was Jerry Fodor's *The Language of Thought*, and the paper was Fodor, Fodor & Garrett's 'The psychological unreality of semantic representations'. Both these works concluded that the meanings of most words are not complex definitions but simple, unanalysable concepts: for example, the meaning of the word 'bachelor' is the simple concept BACHELOR, the meaning of the word 'telephone' is the simple concept TELEPHONE, and so on. (This is also known as the 'atomic' concept view.) This conclusion has much more general consequences for our ideas about concept acquisition, and I'll return to these at the end of the lecture. For the moment, though, I'll concentrate on Fodor's arguments against the classical approach to word meaning itself, and his alternative account; see Lawrence & Margolis: 8-24 and 52-59 for criticisms of the classical approach, and Lawrence & Margolis: 59-71 and Margolis 1998 for summary of Fodor's full views, which we'll just make a start on today.).

Fodor gives three main arguments against the classical approach, which I'll give in ascending order of importance:

(a) Experimental evidence (Fodor, Fodor & Garrett section 2.2.) This argument is designed to show that the meaning of 'bachelor' should *not* be decomposed into the complex concept UNMARRIED & ADULT & HUMAN & MALE. It thus strikes right at the heart of the classical view, since 'bachelor' is one of the strongest arguments for the classical account: if it doesn't decompose into a definition, then surely nothing does.

The argument goes as follows. There are well-known psychological tests for the presence of negation in an utterance. Negation causes processing difficulties when it interacts with quantifiers like 'all' or 'some', or with other negative items such as 'false' or 'deny'. For example, if you have to decide whether a particular utterance is true, or whether a particular argument is valid, you'll take longer with (2) and (4) than with the corresponding positive utterances (1) and (3):

- (1) It's true that somebody came
- (2) It's *not* true that *nobody* came.
- (3) It's true that they *said* that he was ill.
- (4) It's *false* that they *denied* that he was ill.

Fodor, Fodor and Garrett did experiments to show that not only the lexical item 'not', but also the negative morpheme 'un' (e.g. in 'unexciting') cause processing difficulties when

they interact with quantifiers and negation. They ran the same tests on implicitly negative words like 'doubt' and 'deny', which contain no obvious negative element but are semantically negative ('doubt' means BELIEVE NOT; 'deny' means SAY NOT), and showed that these implicitly negative words also create the predicted processing difficulties. They then turned their attention to 'bachelor'. Their hypothesis was that if the classical account of concepts is right, 'bachelor' should behave semantically like a negative item, because its definition includes the concept UNMARRIED (or NOT-MARRIED). We should therefore test the interaction of 'bachelor' with quantifiers and negative items, to see whether it causes the predicted processing difficulties. If it does, the classical approach to word meaning is confirmed; if not, it is disconfirmed.

In the tests, Fodor, Fodor and Garrett compared performance on sentences like (5) and (6), where (5) contains the word 'bachelor' and (6) contains the word 'unmarried'.

- (5) If practically *all* of the men in the room are *bachelors*, then *few* of the men in the room have wives.
 (6) If practically *all* of the men in the room are *unmarried*, then *few* of the men in the room have wives,

If the classical approach to concepts is right, these sentences should cause similar processing difficulties, because their semantic representations will both contain the negative concept NOT-MARRIED. However, when asked to evaluate the validity of arguments like (5) and (6), subjects performed significantly worse with (6), which contains 'unmarried', than they did with (5), which contains 'bachelor' (and so on for other similar examples).

Conclusion: "We take this result to suggest strongly that [words like 'bachelor'] do not act as though they contain a negative element in their linguistic representation." (F, F & G: 522).

(b) Absence of adequate definitions. (Fodor, 'The present status of the innateness controversy', pp. 283-88) Fodor claims that 'bachelor' is quite exceptional in being easily defined. He points out that not only have philosophers failed over many years to define philosophically important words such as 'true', 'know' and 'good', but linguists, despite many years of effort, have failed to define even such simple verbs as 'kill' and 'paint'. For example, it is standardly claimed that 'kill' can be defined as CAUSE TO DIE. However, though all killings may count as cases of causing to die, not all cases of causing to die seem to count as killings. Compare (7) (an alleged consequence of chaos theory) with (8):

- (7) A butterfly, by flapping its wings in Japan, can cause someone to die in Brazil.
 (8) ?A butterfly, by flapping its wings in Japan, can kill someone in Brazil.

That is, 'kill', unlike 'cause to die', carries a connotation of intention, or agency. Similarly, 'paint X' is standardly defined as COVER THE SURFACE OF X WITH PAINT, but as Fodor points out, there are many cases which would satisfy the proposed definition without being cases of painting. For example, when I dip my brush in the paintbox I cover its surface with paint, but I don't *paint* the brush:

- (9) By dipping the paintbrush into the jar, John covered its surface with paint.
 (10) ?By dipping the paintbrush into the jar, John painted it.

Generally, however much you add to them, so-called 'definitions' such as CAUSE TO DIE or COVER WITH PAINT turn out to be not really definitions at all, however much you refine them: they are *necessary conditions* (one-way entailments) but not *necessary and sufficient conditions* (two-way entailments), which is what definitions should be.

Fodor concludes that it is a mistake to model our semantics on the behaviour of the word 'bachelor', which is very much the exception in being easily defined. The classical approach claims that the vast majority of words – e.g. 'kill', 'paint' – are definable. But this remains an article of faith unless the definitions can actually be provided.

(c) The knock-down argument. This is given in Fodor, Fodor & Garrett: 527, though rather briefly, so I'll expand on it here. One of the claims made by linguists who defend the classical approach to word meaning is that by analysing word meanings into definitions, we can capture intuitions of semantic similarity. For example, by decomposing 'bachelor' into ADULT & UNMARRIED & HUMAN & MALE, 'man' into ADULT & HUMAN & MALE, 'boy' into NOT-ADULT & MALE, and so on, we can capture the intuition that words like 'bachelor', 'man', 'boy', etc. are semantically related, because they share certain features of their definition. If you believe in the classical approach to word meaning, then, you will want to decompose words like 'red' and 'horse', to account for the meaning relations between 'red', 'blue' and 'coloured', and between 'horse', 'cow' and 'animal'. For example, you'll assume that 'red' decomposes into a complex concept consisting of COLOURED, plus some further concept X which will distinguish it from the meaning of 'blue'; and similarly for 'horse' and 'cow'. The problem then arises over the nature of this missing concept X.

What we need is a concept that picks out whatever you have to add to COLOURED to get the meaning of 'red'. X should not have COLOURED as part of its meaning: first, because if it did it would then *mean* 'red' all on its own, and it would be superfluous to decompose it further; and second, because if it did, it would then have to be decomposed further to bring

out its relation to other colour concepts. The conclusion is that X must be a concept which means [RED BUT NOT COLOURED]. But this is incoherent. Precisely because RED entails COLOURED, a concept which means 'red but not coloured' would be internally contradictory. Hence, no such concept can exist.

Conclusion: there is certainly a one-way entailment between 'red' and 'coloured': that is, it's true that IF something is red, THEN it's coloured (though the entailment doesn't run the other way: it's not true that IF something is coloured, THEN it's red). However, this one-way entailment cannot be captured by decomposing 'red' into [COLOURED PLUS X], as the classical approach requires, because there is no 'X' to complete the definition. Parallel arguments apply to natural objects (or 'kinds') such as 'horse/animal', 'flower/plant', 'gold/metal', etc. Thus, intuitions of semantic relatedness must be explained some other way.

3. Fodor's alternative approach

Fodor's conclusion, which he still holds (Fodor 1998), is that the meanings of most morphologically simple words (e.g. 'bachelor') are simple, unanalysable (atomic) concepts, which do not decompose into definitions. Thus, the meaning of 'bachelor' is BACHELOR, the meaning of 'man' is MAN, and there is no way to tell, just by looking at their associated concepts, that there is a semantic relation between 'bachelor' and 'man'. That is why, in our simple model of lexical semantics, I claimed that the word 'bird' encodes the concept BIRD, rather than a complex definition such as FEATHERED CREATURE WITH WINGS.

One of the apparent advantages of the classical model of word meanings was that it captured our intuitions of semantic relatedness. Fodor does not deny our intuitions of semantic relatedness, but simply proposes to capture them in a different way. He assumes (as is clearly true) that humans have not only a **grammar** but also a **logic**. A grammar is a set of rules that relate phonological representations to semantic representations. A logic is a set of **inference** rules that relate **concepts** (or larger conceptual representations) to **other concepts** (or larger conceptual representations). These are like the rules you study in your logic course, e.g. the AND-elimination rules which enable you to make valid inferences such as (9) and (10):

(11) $P \ \& \ Q \ \rightarrow \ P$

(12) $P \ \& \ Q \ \rightarrow \ Q$

Fodor's claim is simply that there are many more inference rules than logicians have thought. There is, for example, a BACHELOR rule, that accounts for the inference in (13), and a

MAN rule, that accounts for the inference in (14):


(13) x is a BACHELOR \rightarrow x is a MAN

(14) x is a MAN \rightarrow x is MALE

These additional inference rules he calls *meaning postulates*. His claim is that intuitions of semantic relatedness are explained by appeal to inference rules showing the logical relations among concepts. So what would be done in the classical approach by appealing to definitions, as illustrated above, would be done in Fodor's approach by appeal to logical inference rules.

Inference rules are one-way entailment rules. Thus, we could capture the intuition of relatedness between 'red' and 'coloured' by setting up an inference rule allowing us to argue from X IS RED to X IS COLOURED. The advantage of this approach is that we can then capture whatever meaning relations exist, without being committed, as on the classical approach, to finding some set of concepts which *exhaust* the meaning of 'red' by assigning it a definition. The meaning of 'red' is simply the concept RED, which is logically related to the concept COLOURED. The upshot is that our intuitions of meaning-relatedness are explained not in the grammar, as on the classical account, but in the logic:

Fodor's model of lexical semantics and logic:



I find these arguments against the classical view largely convincing, and Fodor's approach to word meaning is essentially the one adopted in *Relevance* (Sperber & Wilson 1986/1995, chap 2, sect 4: 83-93). However, there have been many attempts to defend a classical approach to word meaning against Fodor's arguments, and indeed in his 1998 book *Concepts*, Fodor still sees himself as a lone voice defending the atomistic view against the massed ranks of cognitive science. As I mentioned earlier, the classical view of concepts has also been criticised by cognitive psychologists working on categorisation, and some of these psychologists have seen their account of concepts as providing an alternative to both the classical and Fodorian accounts of lexical semantics. Next week we'll start looking at some arguments based on categorisation and their implications for both lexical semantics and

lexical pragmatics. To end this week, I'd like to start looking at an immediate question raised by Fodor's rejection of the classical approach to word meaning: if the meanings of words like 'horse', 'red', 'telephone', etc. aren't built up, as the empiricists claim, from bundles of simpler sensory concepts, how *are* they acquired?

4. Implications of Fodor's approach for concept acquisition

The debate between Fodor and the classical theorists about the nature of word meanings is linked to a long-standing debate in the history of philosophy about how concepts are acquired. This debate is generally seen as one between empiricists and rationalists (or nativists). Empiricists (e.g. Locke, Berkeley, Hume) claimed that all our knowledge ultimately derives from the senses plus a few very general reasoning processes: we have a small stock of innate (or innately determined) sensory concepts – e.g. for shape, colour, texture, smell – which, once 'triggered' or activated by exposure to appropriate experiences, may be combined and recombined to form more complex concepts, by processes of generalisation, conjunction, disjunction, negation, abstraction, etc, which are themselves innate. Our concept of a horse, for example, would be a complex construction out of sensory concepts for shape, colour, texture, etc., via mental processes of abstraction or generalisation, as in the classical account.

Rationalists (or nativists, e.g. Descartes, Kant) claimed that our stock of innate (or innately determined) concepts is much greater, and goes well beyond those delivered to us by the senses: it might include, for example, abstract concepts of space, time, causation, etc., and mathematical concepts such as SQUARE and CIRCLE, as well as general concepts such as MAN and HORSE. Notice that the debate between Empiricists and Rationalists is not about *whether* anything is innate (they all concede that simple, unanalysable concepts must be innate, or innately determined), but about *how much* is innate.

Fodor takes an extreme rationalist position which commits him to the claim that there are a vast number of innate concepts. Here are the bare bones of his argument:

Step 1: A concept must be either learned or innate.

Step 2: Only complex concepts can be learned. All simple, unanalysable concepts are innate.

Step 3: The classical view claims that our stock of simple, unanalysable concepts is quite small, and that the meanings of most words, e.g. 'telephone', 'electron', 'cabbage', 'car', are complex and decomposable (hence capable of being learned).

Step 4: The classical view is wrong. All (morphologically simple) words express simple, unanalysable concepts. For example, the concepts TELEPHONE, ELECTRON, CABBAGE,

CAR, etc. are simple and unanalysable.

Step 5: Hence, the concepts TELEPHONE, ELECTRON, CABBAGE, CAR, etc. are innate.

In their commentary on Fodor, Samet & Flanagan (1989) call this position 'Radical Concept Nativism' (Fodor himself calls it 'mad dog nativism'), and add:

"Most philosophers and cognitive scientists think radical concept nativism is outrageous.... George Lakoff probably speaks for most cognitive scientists when he says of Fodor's radical concept nativism, 'I find such an idea too bizarre to take seriously'." (Samet & Flanagan 1989: 189).

Before considering whether this view is outrageous or unacceptable, I'd like to mention a piece of evidence from language acquisition that seems to confirm Fodor's anti-empiricist stance. This is discussed towards the end of 'The present status of the innateness controversy' (around p. 311), and involves apparently well-established facts about the order of vocabulary acquisition in children. Consider the following hierarchies of concepts:

ANIMAL, DOG, POODLE;
 PLANT, FLOWER, TULIP;
 BUILDING, HOUSE, BUNGALOW;
 VEHICLE, CAR, VOLKSWAGEN.

In each series, there is a general or superordinate concept (ANIMAL), a basic, or middle-level concept (DOG) and a specific or subordinate concept (POODLE). Here, the classical view of word meaning seems to predict a certain order of acquisition: children should not be able to learn the meaning of a middle-level word like 'dog' unless they already have the concepts ANIMAL, etc., out of which its meaning is composed. Thus, the order of acquisition should be, first ANIMAL, then DOG, and finally POODLE. The evidence is, however, that children generally acquire the meaning of 'dog' *before* they acquire the meaning of 'animal'. More generally, they acquire the meanings of the so-called basic or middle-level words before they have the superordinate vocabulary out of which, on the decompositionalist approach, the meanings of these basic-level words should be composed. (Notice, incidentally, that the terms involved in this order-of-acquisition argument are typically natural-kind terms, the ones that figured in what I earlier called the 'knock-down argument' against the classical view of word meaning. We therefore now have two arguments against the classical approach to natural-kind terms: one based on the lack of adequate definitions, and one based on order of acquisition.)

On Fodor's view, by contrast, DOG, POODLE, ANIMAL, etc. are all simple, unanalysable concepts, and there is no reason why DOG should not be acquired (or 'activated') before either POODLE or ANIMAL. Around these pages, Fodor makes some interesting speculations about why the concept DOG is generally activated before ANIMAL or POODLE. To understand these arguments, we'll have to look more closely at the relation Fodor sees between innateness and learning, and at the role he attributes to experience in learning, on the one hand, and in activating innate concepts on the other.

We'll look more closely at some of these arguments in a later lecture, but support for Fodor's position comes from the repeated failure of empiricist accounts of concept acquisition throughout the history of philosophy. For example, the logical positivists in the middle of the twentieth century took seriously the project of showing how, say, the concept of a CHAIR could be constructed out of 'sense data', i.e. sensory representations, along the lines the empiricists proposed. This programme failed because our concept of a chair involves notions of persistence through time and independence of an observer's presence, which are simply incompatible with the claim that this concept is reducible to a collection of sense data – which by definition cannot exist independently of a perceiver. Returning to the classical theory of concepts, it follows that the concept CHAIR cannot be decomposed into a small set of sensory concepts (for shape, colour, texture, etc.) which constitute a definition. This again supports Fodor's anti-classical position.

There have been many attempts to show that Fodor's arguments for the innateness of concepts are unsound, or to develop alternative accounts of concept acquisition that would avoid them. Two lines of investigation are worth exploring. One concedes Fodor's argument that concepts do not have the classical structure of definitions, but denies his conclusion that they therefore have no internal structure at all. In particular, so-called 'prototype' theories of concepts, which claim that concepts have the internal structure of a prototype, have been seen as providing an alternative to Fodor's radical concept nativism, and we'll start looking at them next week. A second line of investigation is to look more carefully at Fodor's one-step argument from 'This concept is not learned' to 'This concept is innate', and to find a third position which avoids the more extreme consequences of Fodor's position – say by expanding our notion of learning or our notion of innateness. We'll do this later in the term.

Homework

1. Fodor claims that DOG is acquired before ANIMAL and POODLE, and that this is an argument against the classical view of concepts. How would he analyse the meanings of 'dog', 'animal' and 'poodle', and how would he show the semantic relations between them?

2. How might one account for the order of acquisition of DOG, ANIMAL and POODLE? (If you want to read Fodor's views on this, see Fodor 1981: 298-end.)

Reading

- Fodor, J. 1981 'The present status of the innateness controversy'. Read pp 257-292 and pp 298-end), in J. Fodor *Representations*. MIT Press.
- Laurence, S. & Margolis, E. 1999 Concepts and cognitive science (read pp 8-24, 52-59 on the classical view; 59-72 on Fodor's view). In E. Margolis & S. Laurence (eds) *Concepts: Core Readings*. MIT Press.

Background references

- Samet & Flanagan 1989 Innate representations, in S. Silvers (ed.) *Rerepresentation*: 189-210. (critique of Fodor).
- Margolis, E. 1998 'How to acquire a concept'. *Mind & Language* 13: 347-69 (for arguments that Fodor's atomistic concepts can be learned).
- Laurence, S. & Margolis, E. 2002 Radical concept nativism. *Cognition* 86: 25-55.