

1. Some Preliminary Remarks about Compositionality, Entailment, etc.

Semantics = study of the meaning end of the sound-meaning correspondence, where *meaning* is intended to denote something about the way the world is (e-semantics) or the world as we mentally represent it (i-semantics).

Formal semantics is the study of the composition of complex concepts, like "John loves Mary" or "Newton and the great array of astronomers who have succeeded him have proved that, within planetary distances, matter attracts with a force varying inversely as the square of the distance," or "There was a young lady from Worcester, who used to crow like a roosester."

It's clear, to take the easy case, that if you know the meaning of the string "John loves Mary" you also know the meaning of the string, "Mary loves John".

1. a. John loves Mary.
- b. Mary loves John.

This is the beauty of compositionality. Since the difference between the meanings of a. and b. lie entirely in the positions the elements "John" and "Mary" occupy in the structure, formal semantics tends to take the meanings of the lexical items somewhat for granted, treating them as set-theoretic entities (as sets), without worrying too much about how we decide what belongs (e.g.) in the set of cats vs. what belongs in the set of dogs. Formal semantics focuses on how the way that sentences are put together results in the intuition that "John loves Mary" entails the concept JOHNLOVESMARY and "Mary loves John" entails MARYLOVESJOHN. (Here we encounter the tradition that the mental concepts we entertain are Capital English, an entertainingly ethnocentric notion).

This has enabled formal semanticists to explain a lot of things. They can explain why

2. *Logical Entailment*
 - a. Mary and Sue love John.
 Sue loves John.
 - b. The chair by the door
 The thing which is a chair and is also by the door
 - c. Bossy is a brown cow.
 Bossy is brown and Bossy is a cow.
 - d. Someone loves everyone
 Everyone didn't eat their dinner.
 - e. Socrates is a man.
 All men are mortal.
 Socrates is mortal.

"John and Bill love Mary" entails "John loves Mary". They can explain why, in a given context, the two NPs in 2b designate the same real-world object. They can explain why if Bossy is a brown cow, Bossy must be brown and Bossy must be a cow. They can explain (usually) why sentences like 2d have two different meanings. And they are extremely proud of how they can explain that, given the truth of the first two sentences of 2e, the truth of the third then follows.

This is called "entailment", and simply means that the truth of some proposition P ensures the truth of some proposition Q, purely on formal grounds. That is, it doesn't matter what the individual content of the lexical items "brown" and "cow" are in 2c, e.g., provided they each pick out some set of entities, the predicate "brown cow" refers to the intersection of those two sets and the sentence situates the individual Bossy within that intersection. Hence their disinterest in the specific characteristics of the meanings of lexical items; as long as they pick out sets, they do their job, and compositionality will do the rest.

Now, it is worth pausing to consider, as a philosopher would say, what the *explanandum* is here. Really, what the formal system is accounting for is our firm intuition that if A is true and B is true then C must be true. (Once the formal laws have been elucidated, of course, they can be used to demonstrate the truth of a lot of entailments that don't seem to have the weight of intuition behind them - nicely predictive, in that respect.) But it's the intuitive firmness or seeming conceptual necessity of the argument that motivated people to try and explain it in the first place.

3. *Intra-lexical entailment?*

- a. Mary killed John.
John is dead.
Mary caused John to become dead.
- b. John is a bachelor.
John is unmarried
John is male
John is an unmarried male
- c. I painted the wall.
I caused the wall to become covered with paint.
- d. Toot is a cat.
Toot is an animal
Toot is a mammal
Toot is furry.
Toot is a carnivore.
Toot may be a pet.
Toot's tail twitches involuntarily when she's stalking something...

If all of the above are entailment relations with the same status as the one about Socrates' mortality, then there's more to say about the "meaning" of, say, "cat" or "paint" than that the concept CAT picks out the set of cats, or that the concept PAINT picks out the set of pairs <x,y> such that x painted y, etc.

However, there's another possibility. The entailment relations above might not have the status of the one about Socrates' mortality. Rather, they might simply be pieces

of world knowledge, not arguments that words are made up of compositional elements that enter into true entailment relations, as I mean to suggest with my final list in d, perhaps more on a par with "bachelors are unlikely to have children" or whatever. Of course, the study of such world knowledge and how we come by it, and how it's organized, etc., is an interesting and important subject, but it's likely that it falls more into the domain of psychology than linguistics. Linguistics will be concerned with how a given string gets hooked up with a given batch of conceptual knowledge, perhaps, but not be concerned with the structure of the world knowledge per se. If there are elements of lexical meaning that seem to constrain linguistic possibilities, that's the appropriate subject of lexical semantics. The big methodological problem is figuring out if/when/where to draw the line.

But wait. Let's be a little more clear about the terms of the debate. It all basically hinges on what concepts are linguistic atoms participate in compositional structures.

2. Jackendoff: Semantics and Cognition

2.1 E-semantics vs. I-semantics

4. E semantics

utterances \leftrightarrow things in the world
Platonism

I semantics

utterances \leftarrow \rightarrow knowledge of language \leftarrow \rightarrow grasp of the world \leftarrow \rightarrow world
phonologymorphologysyntax.....semantics.....perception

5. intension vs. intention

Either I'm confused about this distinction, or Jackendoff is (note that the former is probably more likely than the latter). I **think** he's intending to use *intension* where he says *intention*. And it's not clear to me that intensional semantics -- a version of set-theoretic semantics that can make reference to the world as instantiated in the head of the speaker -- doesn't count as an I-semantic theory. However, hopefully I'll get more educated about this over the course of the next week and be able to clarify the problem somewhat. Intensional semantics is still an E-semantic theory in that it constructs alternative possible worlds *with reference to* the real one (e.g. " w_2 is a world in all respects identical to the real world except in that Mary went home in w_2 ", in a treatment of the sentence "John believes that Mary went home").

(Note: anyone who wants to tackle model-theoretic I-semantics for their short paper as cited by Jackendoff (Zwarts and Verkuyl) is strongly encouraged to do so)

6. a. "String S is a grammatical sentence of language L"

is shorthand for "A speaker of language L judges string S grammatical"

Similarly, Jackendoff claims,

b. "A string S of L is true if conditions XYZ... obtain in the real world"

is shorthand for

"A string S of L is true if conditions XYZ... obtain in the speaker's construal of the real world."

Now, this seems *prima facie* unobjectionable, and certainly the spirit behind it is understood; we don't manipulate objects in the real world with our mind, but rather concepts that correlate, sometimes, to objects in the real world. But it seems to miss an important fact about concepts: they are *public*. We'll see more of this when we read Fodor, but consider Humpty Dumpty and Alice, ("There's glory for you!" "I don't understand what you mean by *glory*". "I mean, there's a nice knock-down argument for you!" "But "glory" doesn't mean "a nice knock-down argument"!" etc.), or the status of someone in a small claims court who says *he* thought the sign "No Shirt, No Shoes, No Service" meant, "Take off your shirt and shoes *or else* you won't get any service", not "You must be wearing a shirt and shoes *in order* to get any service".

7. Jackendoff says:

"It is *no longer possible* simply to characterize 'the world' logically or set-theoretically... Rather, it is an empirical problem to determine what sorts of entities inhabit the world as humans construe it. These entities may or may not be characterizable in set-theoretic terms"

This does seem *prima facie* unreasonable. it seems hard to conceive of an entity, mental or otherwise, that cannot receive a set-theoretic characterization; sets are just collections of entities, so mental entities should be just as set-theoretically characterizable as anything else. Of course, what he means is, a set-theoretic characterization is *missing* something important (hence the "simply" above).

He goes on to say, "Furthermore, truth-conditions must be states in the vocabulary of human construal of the world". This isn't objectionable, in fact, they have usually been stated in English, which is of course one vocabulary of human construal of the world...

2.2 *Compositionality*

Any kind of generativist is committed to the programme of composing sentence meanings from their constituent meanings. Lexical semanticists, further, are usually committed to composing lexical meanings from smaller meanings. (Fodor doesn't; we'll see more about this later.)

Jackendoff notes that this helps with the acquisition problem: because meaning is compositional, it's a formal system; a formal system needs primitives and combinatory rules; the learner can't make the jump from nothing to some kind of "unit", hence there must be an innate vocabulary of semantic primitives from which more complex meanings can be composed/on which other meanings can be built. On one construal, the subject of lexical semantics is the investigation of the nature of the primitive semantic units.

8. a. primitives of lexical representation in the child primitives in the adult
 "Different languages may appear semantically incommensurate"

Analogy: range of syntactic and phonological possibilities allowed by UG; some are selected by some languages, some by others. The same array+selection process may well happen in lexical semantics. For example, English lexicalizes MANNER elements with MOTION elements; by all appearances, Romance languages do not. (Or my claim that some languages lexicalize HAVE and some do not).

- b. Semantic primitives words
- there's the possibility that "bound sememes" exist

Hence, the paraphrase argument in general is bad (c.f. Fodor on kill vs. cause to die). But this is a difficult point: if it's entailment relations in general that we're trying to explain, how else can we construe "entailment" in a principled way? General insight, though, is right.

The big methodological problem in lexical semantics, alluded to above, can be restated as, "how can we justify a semantic primitive?" I.e. what aspects of meaning are linguistically relevant and what are not?

One might think that part of the answer to resolving that question is to posit a level of semantic representation separate from the real-world concepts that mediates between the syntax and real-world knowledge.

- 9. a. *3-stage mapping*

Syntax $\leftarrow \rightarrow$ Lexical semantic representation $\leftarrow \rightarrow$ Encyclopedia

- b. *2-stage mapping*

Syntax $\leftarrow \rightarrow$ Encyclopedia

Jackendoff asserts that there's no predictive distinction to be made between a. and b., as long as b. has a principled list of what elements of Encyclopedic meaning are relevant for the syntax, or as long as the syntax somehow knows what's relevant and what's not. (another possible short paper here: Choi and Bowerman). This may be true, but then the list is essentially an intermediate level of representation, or whatever characterizes the syntax's knowledge. The 3-stage mapping gets around a potential problem: that of characterizing the difference between languages. If the syntax magically knows what's relevant in language 1 and that list is different in language 2, the situation is considerably different than in his phonological analogy, where he claims that the syntax doesn't care about phonological distinctions between, e.g. the string "dog" and "cat", but it does care about word boundaries, phrasal intonation, and so on. However, *no* language's syntax cares about the phonological distinctions between "dog"-equivalents and "cat"-equivalents, and *all* languages care about intonation and word boundaries, but it's Jackendoff's claim (and probably an empirical fact) that different languages care about different meaning elements. The 3-stage mapping adds a mechanism that could

possibly account for such differences, I think. Maybe the 2-stage one can as well; J. certainly thinks so.

10. *Why categories are not sets*

Formal semantics: "Toot is a cat" = 1 iff $\exists x$ (Toot is x & x is a cat)

I-semantics: categories aren't sharp, like sets. They are often:

- a. *fuzzy* e.g. "red"
- b. *dependent* e.g. "small", "near"
- c. *indeterminate* e.g. "game", "climb", "chair"
(i.e.. definitional criteria are not subject to a requirement of logical conjunction, as is true for, e.g., "brown cow")
- d. *imagistic* e.g. "duck" vs. "goose"
- e. *spatial relations*

11. *Spatial representations in more detail:*

locations vs. locatums

- a. The bike is next to the house.
- b. The house is next to the bike.

spatial "metaphors"

- c. John kept the messenger in Istanbul Location
(is Constantinople)
- the money Possession
- the light red Property
- the meeting on Monday Scheduling (time?)

So Jackendoff wants to point out that "keep"'s meaning, unlike its sort of knee-jerk definition, doesn't necessarily have a possession component involved, and hence wants to argue that it decomposes (in other work). Fodor points out that the differences between kinds of keeping in 11c seem to boil down to differences between the kinds of things being kept, rather than being located in *keep* itself. It's not clear to me that they're making different points; I've made a similar argument w/r to the interpretation of "have" elsewhere.

Anyway, J's main point is that it seems like spatial relations are part of our basic cognitive machinery, and that is reflected in our basic linguistic machinery. And that set-theoretic machinery is not enough to capture these sorts of generalizations.

2.3 *What sort of entities are in the human construal of the world?*

"Nothing in the world demands a sharp boundary between red and orange, nothing in the world distinguishes climbing from other kinds of motion. It is the human need to sort the particulars of the world into categories that creates these divisions"

Seems slightly misleading here: we've already established that we're **not** sorting particulars of the world, but our internal representations of them mediated by our

perceptions. And it seems like he's recanting his anti-set-theoretic argument above, in some sense: if the boundary between red and orange is so fuzzy that it denies set-theoretic characterization, then it's not "sharp". But the basic point is clear: the world doesn't overtly indicate possession, Mondays, gems vs. rocks, etc; these are divisions that are purely relative to our human categorizations. But it's clear anyway that a theory of human categorizations is necessary, not a theory of the world (that's the job for other kinds of scientists).

3 Levin: Introduction: The Theoretical Perspective

12. *Linking principles*: generalizations (presumably robust in the mind) about the relationship between a lexeme's meaning and its possible syntactic environments

Now, this is slipperier than you might imagine, as you will see when we get to Gleitman's syntactic bootstrapping, or theories of measuring-out. Certainly it's the case that certain lexemes don't belong in certain frames. But is it because of a linking principle between elements of lexeme meaning and syntactic frame constituency, or is it because the frame itself carries a certain meaning? The tension between these two approaches is central to some of my own work.

The way to isolate elements of meaning which are relevant for the linking principles, asserts Levin, is to identify some syntactic behaviors that cluster together for certain semantically similar groups of verbs, and then identify the meaning element (usually in contrast with other *prima facie* similar groups of verbs) which triggers the grouping of behaviors. She goes on to give several examples (I've given some examples of hypotheses about what meaning elements might be relevant in these cases; we'll look extensively at argumentation concerning these issues. The "meaning elements" are **not** meant to be conclusions, rather, suggestions about where an investigator might look for relevant data):

- spray/load relevant meaning elements: telicity, fore/background?
13. a. John sprayed the wall with paint.
 John sprayed paint on the wall.
- b. John filled the glass with milk.
 *John filled milk in the glass.
- c. *John poured the glass with milk.
 John poured milk into the glass.
- d. John covered the wall with paint.
 *John covered paint on the wall.
- causative/inchoative relevant meaning elements: change of state, causation?
14. a. The window broke.
- b. John broke the window.
- c. The rabbit suddenly appeared.
- d. *The magician appeared the rabbit.

- benefactive alternation relevant meaning elements: product vs. material?
15. a. Sue carved a toy out of wood for Hansel.
 - b. Sue carved Hansel a toy out of wood.
 - c. Sue carved some wood into a toy for Hansel.
 - d. *Sue carved Hansel some wood into a toy.

- middle formation relevant meaning elements: affectedness of object?
16. a. *The whale sees easily.
 - b. The whale frightens easily.

Conative/transitivity/body part ascension and transitive verbs:

- relevant meaning elements: contact, change of state
17. a. Sue broke the vase.
 - b. Bill cut the bread.
 - c. Mary touched the cat.
 - d. Joe kicked the tire.
18. a. The vase broke
 - b. *The bread cut.
 - c. *The cat touched.
 - d. *The tire kicked.
19. a. The vase broke easily.
 - b. The bread cut easily.
 - c. *The cat touched easily.
 - d. *The tire kicked easily.
20. a. Sue broke the boy's arm /*Sue broke the boy on/in the arm.
 - b. Bill cut the man's arm/Bill cut the man on the arm.
 - c. Mary touched the cat's head/Mary touched the cat on the head.
 - d. Joe kicked the girl's shin/Joe kicked the girl on the shin.
21. a. *Sue broke at the vase.
 - b. Bill cut at the bread.
 - c. *Mary touched at the cat.
 - d. Joe kicked at the tire.

Different languages might group different verbs differently, but the generalization of meaning-association and syntactic positions should hold for some classes, whatever the particular instantiations of the classes are. So, e.g., "break" is a causative/inchoative alternator in English, but not, apparently, in Italian, while "cook" is in both.

22. The unaccusative hypothesis

The subject of certain single-argument verbs originates in object position, and moves to subject position in the course of the derivation.

23. Derived subjects

The train arrived.

The tree fell.

The sun rose.

The book is black

...

Base-generated subjects

The boy ran.

The dog barked.

The baby laughed.

The boy ate.

...