Semantics

LING 553

Sept. 10, 2008

1. SEMANTICS: PUTTING THIS ALL TO USE

> It is astonishing what language accomplishes. With a few syllables it expresses a countless number of thoughts, and even for a thought grasped for the first time by a human it provides a clothing in which it can be recognized by another to whom it its entirely new. This would not be possible if we could not distinguish parts in the thought that correspond to parts of the sentence, so that the construction of the sentence can be taken to mirror the construction of the thought. [...] If we thus view thoughts as composed of simple parts and take these, in turn, to correspond to simple sentence-parts, we can understand how a few sentence-parts can go to make up a great multitude of sentences to which, in turn, there correspond a great multitude of thoughts. The question now arises how the construction of the thought proceeds, and by what means the parts are put together so that the whole is something more than the isolated parts.
>
> Gottlob Frege, "Logische Untersuchungen. Dritter Teil: Gedankengefuege"

So then...

- What are the basic elements of our system? (DWP (1981): "This is not an easy question to answer since it asks, in effect, What is there in the world?")

- What system do we use to combine them?

General goal: we want to know what a sentence "means". e.g., our theory of semantics should tell us that the sentence **snow is white**[1] means that snow is white.

"Snow is white" means that snow is white.
↑object language ↑metalanguage

This isn't *stupid*, though it looks stupid when the object language is the metalanguage...

"Der schnee ist weiß" means that snow is white.

But mostly it isn't quite helpful yet, since we don't know what it means to "mean".

[1] Typographical note: similar to H&K, I'll use **this bold sans-serif font** to indicate a linguistic string of the object language, i.e. the language we're studying. (Properly: I'll use it as such everywhere other than in the previous sentence, where it's used in an odd self-referential way.) Thus, **Julie Legate** = "Julie Legate" = the linguistic object pronounced, roughly, as [ʤulilɛgeit].

1.1. Tarski's definition of truth

The foundation of truth conditional semantics:

> Let us consider an arbitrary sentence; we shall replace it by the letter *p*. We form the name of this sentence and we replace it by another letter, say *X*. We ask now what is the logical relation between the two sentences "*X is true*" and '*p*.' It is clear that from the point of view of our basic conception of truth these sentences are equivalent. In other words, the following equivalence holds:
>
> (T) *X is true if, and only if, p.*
>
> ....Now at last we are able to put into a precise form the conditions under which we will consider the usage and the definition of the term "*true*" as adequate from the material point of view: we wish to use the term "*true*" in such a way that all equivalences of the form (T) can be asserted, and *we shall call a definition of truth "adequate" if all these equivalences follow from it.*

An example of (T), from Tarski:

"Snow is white" is true if, and only if, snow is white.
or, again,
"Der schnee ist weiß" is true if, and only if, snow is white.

Ramifications:

- Sentences have truth values: they "are" true or false.
- A mapping exists from sentences (i.e., linguistic expressions) to truth values (that maps, e.g., **Snow is white** to TRUE and **Snow is black** to FALSE).

AN IMPORTANT NOTE ABOUT THE TARSKI SENTENCE: the truth conditional semantics will generally have the format given. We know that our semantics has succeeded if we get to, e.g.,

(1) **Snow is white** is true iff snow is white.

Remember, however, that the name of the sentence is arbitrary; (T) need not take the form "**X** is true iff X". For instance, if we're investigating the meaning of "the", the sentence in (2) isn't especially helpful; we may want (3) instead.

(2) **The book is white** is true iff the book is white.
(3) **The book is white** is true iff there is a single book, and that book is white.

When writing derivations, resist the temptation to take metalanguage results and rephrase them to be identical to the object language.

1.2. A rough start

Ontological fact: Truth values exist.[2]

Definition: The interpretation function, ⟦...⟧, is a function from linguistic expressions to their meanings. (So far, we know that it maps sentences to truth values.)

We want sentences to have truth values, i.e. "true" or "false". So we'll set up the set of truth values, which we'll call D_t and which is {TRUE, FALSE}. Then ⟦**Julie snores**⟧ = TRUE (if indeed she snores); the ⟦...⟧ function maps the linguistic string **Julie snores** to the truth value TRUE.

Of course, we need to revise this fairly quickly, because:

- Linguistic strings other than sentences—e.g., **my coffee mug**—have meanings, but not apparently truth values. (Is "my coffee mug" true or false?)

Well, that's OK. ⟦α⟧ will be a truth value when α is a sentence, but other linguistic expressions will be mapped to other kinds of things.

- Some sentences don't seem to have truth values. For instance, ⟦**Is it raining?**⟧ doesn't seem to correspond to TRUE or FALSE the way ⟦**Julie snores**⟧ does; and that's even more apparent for ⟦**Who's enrolled in LING 553?**⟧ or ⟦**Close the door.**⟧

Rather than saying "sentences...have truth values", it's more accurate to say "declarative sentences have truth values". The interpretation function will map other kinds of sentences to other kinds of meanings.

- Some strings seem to lack interpretation, like the underlined part of **my <u>mug is on the</u> table**, or the sentence **my mug is on the gostak**.

To some extent, we're going to have to deal with these things as we reach them. At least part of the answer is syntactic...

[2] So do linguistic expressions, but we won't really need to include them in our theory.

1.3. Syntactic excursus

Recall that many sentences are ambiguous, and at least some of those ambiguities are syntactic.

(4) I saw her duck.

(5) I saw the man with the binoculars.
  a. ...and Mary did so with the telescope.
  b. ...and Mary saw the one with the telescope.

Clearly ⟦<string of words>⟧ isn't going to be sufficient to derive a meaning. We have to have some sense of the structure.

(6)

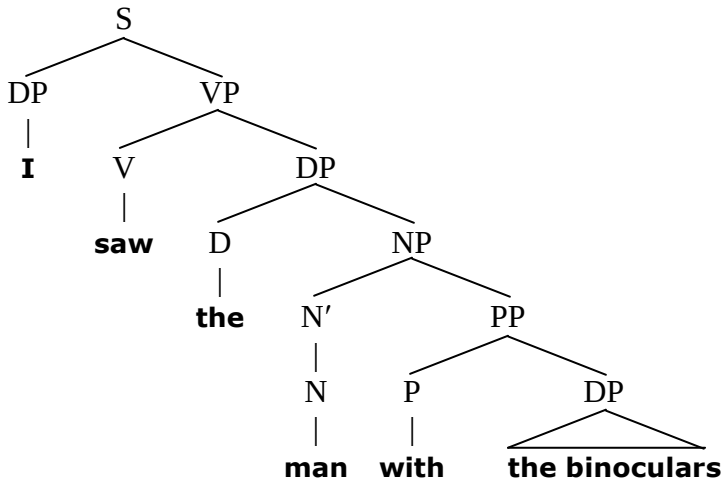

(7)

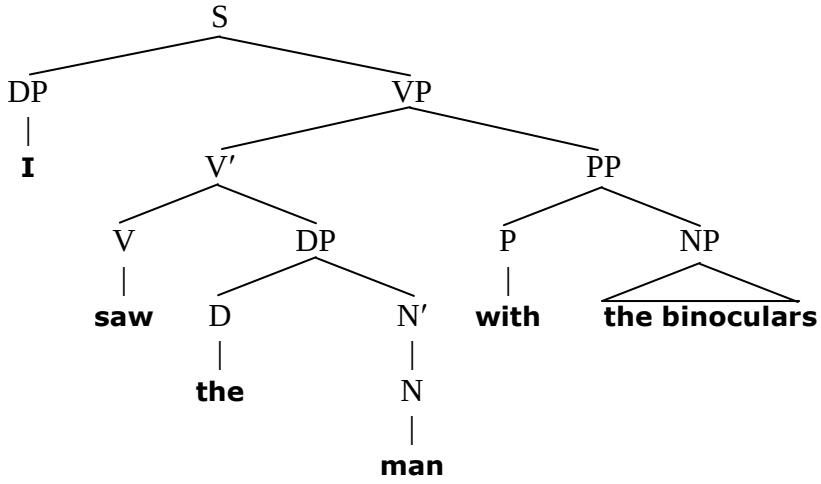

...or, more to the point, we don't want to interpret strings like ⟦**Julie snores**⟧, but instead structures like:

(8)

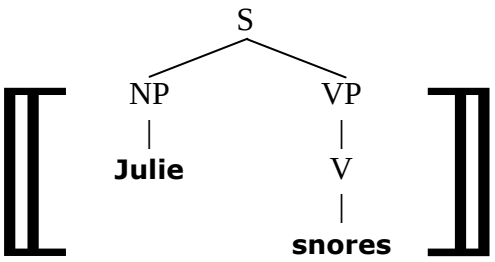

When it comes to **mug is on the**, we have something that isn't in the domain of the interpretation function, because it's not a legitimate syntactic structure, but a weird sort of partial, broken one:

(9)

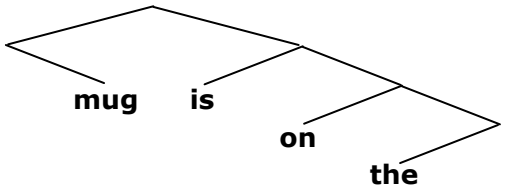

And with **my mug is on the gostak**, the structure gives us enough to more or less form a meaning with a small hole in it where the details of **gostak** go.[3]

Two more notes before leaving the syntactic excursus. First, we're going to work as much as possible with geometry and not with actual syntactic node labels. (Which is good, because in minimalism, there are no node labels.) So don't stress over what the nodes are labeled, or indeed if they're labeled at all.

Second, we've revised the function ⟦...⟧ so that, instead of being a function from linguistic strings to their meanings, it's a function from linguistic strings-and-structures to their meanings. Resist the temptation to put something in the brackets that is part of the metalanguage and not part of the object language. (That's obvious now; it'll be more tempting later. We'll see why.)

[3] Recall the original example that uses **gostak**: if you're told "The gostak distims the doshes" and are then asked "What does the gostak do to the doshes?", you can answer "It distims them". The syntactic structure goes a long way towards providing semantic information.

1.4. Rules of interpretation

To interpret (8), we'll need rules like the following (cf p. 16, H&K):[4]

(R1) $[\![\overset{\text{NP}}{\underset{\alpha}{|}}]\!] = [\![\alpha]\!]$

(R2) $[\![\overset{\text{VP}}{\underset{\alpha}{|}}]\!] = [\![\alpha]\!]$

(R3) $[\![\overset{\text{V}}{\underset{\alpha}{|}}]\!] = [\![\alpha]\!]$

(R4) $[\![\overset{\text{S}}{\underset{\alpha \quad \beta}{\wedge}}]\!] = [\![\alpha]\!] \,..^{?}.. \, [\![\beta]\!]$

...where [[S]] is some combination of [[α]] and [[β]]. Then, working somewhat informally up through the tree, we end up finding that (8) ≡ [[**Julie snores**]], which depends on [[**Julie**]], [[**snores**]], and the exact nature of R4. What seems likely (or indeed possible)?

Given that [[**Julie**]] is, by definition, the meaning of the linguistic string "Julie", we might ask what we mean by a name like "Julie". The most likely answer seems to be that we mean Julie—the person, the actual object in the world. (Note that this involves an ontological commitment to there being actual objects in the world. That's an assumption I feel OK about making.)

Let's suppose, then, that there are things in the world. We can put them in a set and name it for convenience—we had D_t earlier, which was the set of truth values {TRUE, FALSE}, so named because *t* designates "truth values". So we can call things in the world "entities" and name the set of them D_e. Julie and my coffee mug and the University of Pennsylvania are all things in this set.

Now what? Lydia suggests:

> [[**Julie**]] = Julie (an entity, *e*)
> [[**snores**]] = an act of snoring (an action, *a*)
> (R4): [[S]] = TRUE if and only if the entity does the action
> = TRUE if and only if [[NP]] engages in [[VP]]

[4] For simplicity, I've cut out one of the nodes, i.e. the one in which the proper name **Julie** is dominated by an N node, rather than directly by an NP.