Preliminaries

Semantics I Simon Charlow

simon.charlow@rutgers.edu

September 2, 2015

1 Semantics in generative grammar

There are famous and compelling arguments for **generativity** in syntax. We have extremely fine-grained judgments about the grammaticality of an infinity of sentences (in principle), including ones we've never heard. To an astonishing degree, these judgments are shared among speakers of the same language.

A similar argument for **generativity in semantics**: we unconsciously, effortlessly pair linguistic form with meaning, even for forms we've never encountered. To an astonishing degree, speakers of the same language associate the same meaning with a given linguistic form (hard to see how linguistic communication could be managed otherwise).

Since there's an infinite number of interpretable sentences, we can't be idiomatically pairing form and meaning. We're finite beings, and so there must be a finite system that can interpret any of the infinite grammatical sentences of a language.

Today: how is this possible? Beginnings of a theory of **semantic competence**.

2 Compositional semantics

Compositionality: the meaning of a complex expression is a function of the meaning of its parts, and the way those parts are put together. A bedrock methodological assumption in formal semantics.

Assuming compositionality has some important consequences:

i. Given a finite number of ways to combine expressions and a finite lexicon, interpreting a natural language should be tractable in principle. Gives us traction on the problem of pairing form with meaning.

ii. If interpretation is compositional, it matters how structures are put together. Therefore, semantics must take syntax seriously.

iii. Conversely, semantics can be brought to bear on syntax: it's conceivable that a syntactically plausible structure turns out to be uninterpretable or gives rise to unattested interpretations.

iv. Compositionality goes all the way down. We must be able to assign meanings to all the (interpreted) constituents of complex expressions, down to the morpheme.

Our central goal: pairing structured syntactic representations with meanings in a compositional way. In another (more formal) manner of speaking, elucidating the compositional nature of the **interpretation function** $[\![\cdot]\!]$. That is, if $\mathscr{T}$ is some structure, $[\![\mathscr{T}]\!]$ is its meaning.

Other varieties of semantics: lexical, feature-based ("Markerese"), proof-theoretic.

Though we will see some lexical semantics in this course, we'll mostly be concerned with **compositional semantics**, i.e. the ways in which languages build meanings for complex expressions from the meanings of their parts.

3 What is meaning?

3.1 Meaning as truth-conditions

One of the reasons semantics is so slippery is that the object of study is (in part) meaning, and it's not at all obvious to introspection what meaning *is*.

Proposal: **meaning is truth-conditions**. To know the meaning of a sentence is to know what the world has to be like for the sentence to be true. E.g. "snow is white" is true iff (if and only if) snow is white. More formally:

$$[\![\text{snow is white}]\!] = \text{true iff snow is white}$$

This (known as the disquotational principle) sounds really trivial. Let me reassure you with a couple of appeals to authority:

> To know the meaning of a sentence is to know its truth-conditions.
>
> (Heim & Kratzer 1998)

> Semantics with no treatment of truth conditions is not semantics.
>
> (Lewis 1970)

More seriously, it *would* be trivial, if that was all there was to our theory of meaning. But it isn't. It just places a substantive constraint (given, e.g., that we understand what it means for snow to be white) on our theory of semantic competence. Any such theory that tells us that "snow is white" can be true when snow isn't white is not a good theory.

3.2 Interlude: entailment and its varieties

A **entails** B ('$A \Rightarrow B$') iff whenever A is true, B has to be true too, i.e. if $[\![A]\!] = \text{true}$ implies $[\![B]\!] = \text{true}$.

Example: from an utterance of "snow is white", we can conclude that something is white (it's impossible for snow to be white without *something* being white). On the other hand, we *cannot* conclude something is black (it's possible for snow to be white without something being black.).

(1) Snow is white.
$\Rightarrow$ Something is white.

That sort of entailment is commonly known as **asserted** or **at-issue**. Importantly, it's destroyed by negation:

(2) Snow isn't white.
$\not\Rightarrow$ Something is white.

Other sorts of entailments exist. **Presuppositions** are entailments that both a speaker and their interlocutor must share at the time an utterance is made for the utterance to make sense. A presupposes B iff both A and its negation entail B. In another manner of speaking, presuppositions *project* out from under negation. (Non-veridical contexts more generally, i.e. attitude verbs, conditionals, questions, etc.)

(3) They know my Uber was late. $\Rightarrow$ My Uber was late
They don't know my Uber was late. $\Rightarrow$ My Uber was late.

Worth thinking about the correctness of this definition, in particular how well it works for, say, **expressive** adjectives (*damn*, *fucking*, etc.).

Implicatures are defeasible entailments that arise from communicative norms. Contra assertion and presupposition, the classical view of implicatures is that they're not part of what is conventionally encoded (recent approaches disagree).

(4) A: How many kids does Steve have?
B: Steve has two kids. ($\rightsquigarrow$ Steve has *exactly* two kids)

(5) If Steve has two kids, he can claim this tax rebate.
(What if Steve has three kids? Can he claim the rebate?)

3.3 Model-theoretic semantics

Meaning as reference. A sentence *refers* to true if the sentence is true, and to false if the sentence is false.

But how does this come about? How does the theory look inside of a sentence, look at the world, and then determine what the sentence's reference is?

The way we'll practice semantics (and the way it's usually practiced) is **model-theoretic**. That is, we ask whether a sentence is true with respect to a formal (i.e. mathematical) model, rather than with respect to the actual world.

Whether a sentence is true or false in a given model tells us something about what the sentence means. Evaluating a sentence's truth or falsity relative to a class of models tells us even more.

In **extensional** models we have individuals, truth values, and some facts about what properties individuals have and what relations the individuals bear to each other (i.e. who the men are, whether John likes syntax, etc.). Expressions are *synonymous* iff they have the same extension (i.e. reference).

Extensional models are not adequate for all cases. Suppose John came and Bill came. Relative to this model, both "John came" and "Bill came" have the extension (reference) true. Yet it's conceivable to think one but not the other.

(6) Mary thinks John came.
$\not\Rightarrow$ Mary thinks Bill came.

Intensional models: Possible worlds brought into the mix. Meanings have different extensions depending on which possible world they're evaluated at. Expressions are *synonymous* iff they have the same extension in every possible world. By those lights, "John came" and "Bill came" are not synonymous.

3.4 The world is not enough?

Truth conditions seem less suited to other cases.

Hyper-intensionality: the following entailment does not go through, even though "two and two is four" and "Fermat's last theorem is true" are logically equivalent (to wit, there's no possible worlds where either is false).

(7) John knows that two and two is four.
$\not\Rightarrow$ John knows that Fermat's last theorem is true.

Questions do not seem to lend themselves to being analyzed in terms of truth conditions (when is a question true??). Instead, a notion of *answerhood conditions* seems potentially more appropriate.

(8) A: Which student read a book?
B: BILL.

(9) A: Which student read which book?
B: BILL read ASPECTS, and SUE read GB.

Something similar goes for **imperatives** and **directive language** more generally. Don't really seem like a good fit for truth conditions (when is an imperative true??). A notion of *compliance conditions* seems more appropriate.

(10) (I order you to) take your seat at once!

Indefinites seem to be special in a way that truth conditions don't capture. The following two sentences are truth-conditionally identical (that is, the first entails the second and vice versa). Nevertheless, only the first can be continued with *he sat down*.

(11) A man walked into the room.

(12) Not every man didn't walk into the room.

Intriguingly, **disjunctions** also pose problems. Ignoring tense, "John comes" has the same truth conditions as "John or John and Bill comes" (this may not be obvious; try to think of a scenario that would make one true and the other false). But the entailment doesn't go through. The first sentence can be true even though the second is false (e.g. maybe John and Bill are sworn enemies).

(13) If John had come, the party would have been fun.
$\not\Rightarrow$ If John or John and Bill had come, the party would have been fun.

So we might need something *richer* than truth conditions (indeed, each of these cases has been used at one time or another to argue in favor of an enrichment to the truth-conditional paradigm).

Still, many things can be usefully analyzed extensionally, even more intensionally. In this class, we will spend a good deal of time doing extensional truth-conditional semantics. Then we will learn a bit about some extensions.

4 Kinds of ambiguity

Compositionality means: every node in an interpreted structure is itself associated with an interpretation. The interpretation of the larger structure is assembled from the meanings of the smaller structures.

Complicated by certain factors. Among them ambiguity.

Lexical ambiguity: when one or more words in a sentence has more than one dictionary definition. Not going to focus on in this class.

(14) John went to the bank.

Structural ambiguity: Phrase structure rules for a toy grammar:

$$\text{NP} \rightarrow \text{N} \qquad \text{NP} \rightarrow \text{N PP}$$
$$\text{VP} \rightarrow \text{V DP} \qquad \text{VP} \rightarrow \text{VP PP}$$

These rules (along with the obvious additions) generate two possible trees for e.g. *the boy saw the elk with the binoculars*:

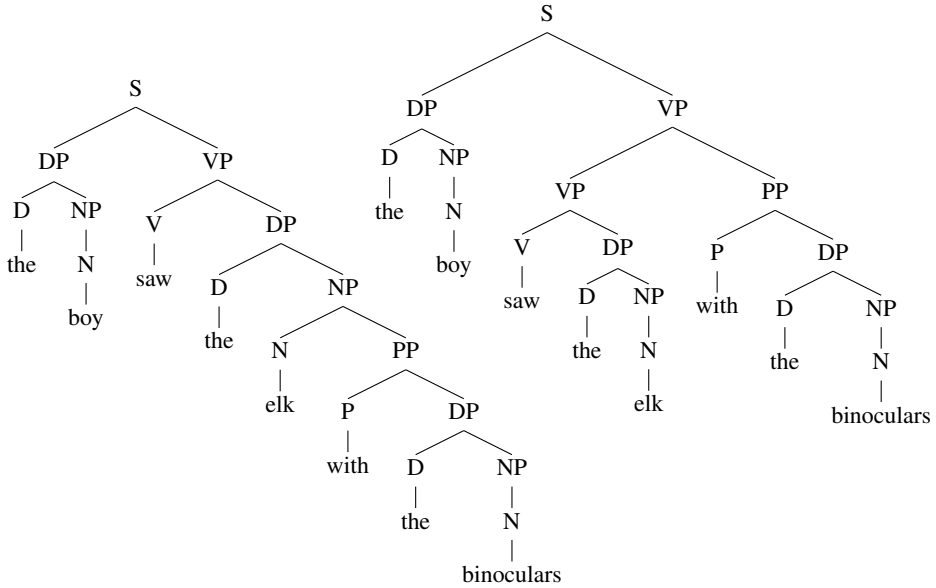

It's extremely natural to think of these two structures as corresponding to two different interpretations of the sentence (you might not have noticed that the

sentence is ambiguous since the other interpretation is fairly implausible, but the ambiguity is real).

Syntax determines the **order of semantic combination**. This means different trees for the same string of words can receive different interpretations.

Notice how this is bound up with constituency: in the first tree *saw the elk* is a constituent (VP), whereas in the second it's not. In the first tree, *the elk with the binoculars* isn't a constituent, whereas in the second, it is (DP). Compositionality means that all and only the constituents in a tree receive interpretations.

Another example: say you saw the sentence "the dish comes with eggs and ham or bacon" on a menu. There's two ways to interpret this: in the first, you're given an option of meat; in the second, you're given a choice between breakfast and bacon.

Again, it's natural to think of this ambiguity as having to do with different possible structures. Let's add the following rules to our grammar:

$$S \rightarrow S \text{ and } S$$
$$S \rightarrow S \text{ or } S$$

We now generate two trees for the relevant structure (using e.g. "eggs" as shorthand for "the dish comes with eggs"):

(15)

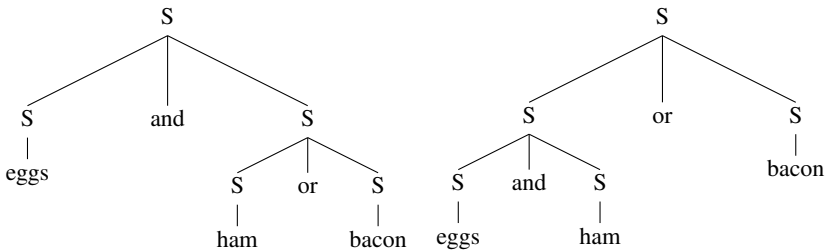

Suppose we have the following two interpretation rules:

(16) $[\![P \text{ and } Q]\!] = \text{true iff } [\![P]\!] = \text{true and } [\![Q]\!] = \text{true}$

(17) $[\![P \text{ or } Q]\!] = \text{true iff } [\![P]\!] = \text{true or } [\![Q]\!] = \text{true}$

Look at the tree on the left. *And* guarantees that you get to eat both the *eggs* branch and the *ham or bacon* branch. Eating the *ham or bacon* branch means you eat either the *ham* branch or the *bacon* branch. All together, you're guaranteed a proper (albeit non-kosher) breakfast.

Now look at the tree on the right. *Or* lets you eat either the *eggs and ham* branch or the *bacon* branch. If you choose the former, *and* guarantees that you eat both the *eggs* branch and the *ham* branch. If you choose the later, it's bacon for breakfast.

Again we see that structure matters. And again, the reason it matters is that it dictates the order of semantic combination.

Other cases are less clear. For instance **pronominal** "ambiguity" and **scope** ambiguity (the scare quotes reflect current consensus in the field, but as recently as 1979, some formal semanticists denied there was a real thing as scope ambiguity):

(18) No boy$_i$ should insult his$_{i/j}$ mother.

(19) A guard was standing in front of every building.

Are these structural ambiguities? *Lexical* (worth thinking about how that might be possible for scope ambiguity)? Something else?

You might even think (and some do) that the cases of implicature we looked at reflect some sort of ambiguity.

5 What linguistic phenomena are semantic? Where is the division?

It is in general not a priori whether a given phenomenon reflects syntactic processes, semantic processes, or even both.

There-insertion seems sensitive to the properties of a quantified DP. But are the relevant properties syntactic or semantic (or both)?

(20) There were {some, no, many, more than three, *the, *most, *all, ...} linguists at the party.

NPI licensing:

(21) I *(didn't) ever think I'd be here.

(22) {Everyone, *someone} who ever walked through that door regretted it.

Question-answer congruence:

(23) Who met Sue?
{BILL, *Bill} met {Sue, *SUE}.

Superlatives vis à vis extraction:

(24) *Who$_i$ did you take the picture of e_i?

(25) Who$_i$ did you take the best picture of e_i?

"Sloppy" readings vis à vis identity in ellipsis:

(26) John went to the bank, and then BILL did __.

(27) John likes his mom, and BILL does __ too.

Extra argumentation has to be brought to bear on these cases. Do syntactic or semantic explanations offer a better fit for the data? Is one or the other more explanatory?

6 For next class

i. Truth-conditions and compositionality: Heim & Kratzer Ch. 1 [12pp].

ii. The semantics-pragmatics interface: Grice: "Logic and conversation" [18pp], Stalnaker: "Pragmatic presuppositions" [12pp].

iii. Start having a look through the Allwood et al. textbook, if you have some extra time.