Semantics and Pragmatics. Lecture 2.

Matthew McKeever (lecturer, mckeever@hku.hk / mipmckeever@gmail.com)

Xindi Ye (TA, xindiye@connect.hku.hk)

- **You should have Moodle access**
- **Sign up**
- **Last week**

This week we're going to introduce the goals and methods of semantics, the study of meaning.

Idea 1: Deep Structures

Compare and contrast

1) You see Herman. Who do you see?
2) Nǐ kànjiàn Herman. Nǐ kànjiàn shéi?

Notice the difference, particularly with regards to the behaviour of the question words ("who", "shéi"). In English, to turn a statement into a question, you have to (normally) rearrange the sentence. Not so in Mandarin and other languages (the fancy linguistics term is 'wh- in situ' languages).

A universal grammar rule. All languages are wh in situ, at some level.

This is not an actual universal grammar rule (is there are any). It's just a potential one. That means, at some level, the English is actually

3) You see who?

But what does 'at some level' mean?

It means that modern linguistics proposes a distinction between different forms one and the same sentence can take. There are thousands of papers on this (which I haven't read), so we'll just propose a sort of toy example, using outdated terminology and theories but ones that can still be found, in some version, in today's syntactic theory textbooks. So let us define:

Deep structure in some sense represents the underlying form of the sentence, the one in accordance with universal grammatical rules

And:

Surface structure is the result of applying certain operations to the deep structure to make it a suitable sentence for the given language.

So we would have, for English:

Deep structure: You see who?

Surface: Who do you see?

Moreover, there are putative rules that take one from deep to surface structures. One is movement. Movement has been very important in linguistic theory (again, syntacticians spent their whole lives debating this issue, and I'm not a syntactician, so this is just for the sake of example and not meant to be an accurate report on what people think in 2022):

Step 1: Original sentence. You see who?

Step 2. Delete the 'who'. You see ~~who~~?

Step 3. Move 'who'. Who you see ~~who~~?

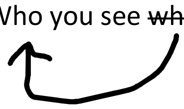

Step 4. Final sentence. Who do you see ~~who~~?

The core idea: a given sentence can have more than one form, and deriving the surface from the deep form might involve a series of steps

That language has hidden structure is obvious once you see it. Consider:

4) Second-hand books and laptops are available here for a good price

And consider disambiguations:

5) [Second-hand books] and laptops are available here for a good price
6) [Second-hand books and laptops] are available here for a good price

If we want to be fancy, we can draw phrase structure trees, although the above crude brackets our point for us: there is hidden structure to sentences.

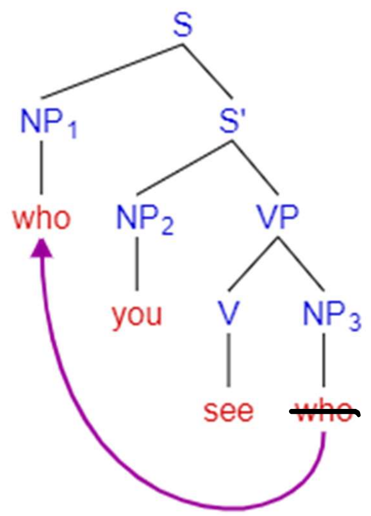

The key idea is that using ideas from syntax, we take sentences, which are after all just linear sequences of words, and impose (sometimes hidden) structure on them. It's these more structured entities that are the input to the sort of semantics we'll consider.

Idea 2: Functions

Let's change examples for a more simple case. Imagine a very simple sentence with no weird features

7) Datri sings

No ambiguity, but we still have to bracket or draw a tree:

{Datri} {Sings}

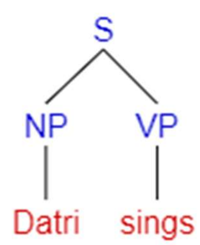

Our goal, *the* goal of semantics: assign *some thing* as meaning to each node on the tree (or to each bracketed expression) that determines the meaning of the whole.

Math to the rescue?

Why *some thing*? Well, think about maths. Consider "1", "I", "一". There is surely something these expressions have in common. Here's Frege:

> There is at present a very widespread tendency not to recognize as an object anything that cannot be perceived by means of the senses; this leads here to numerals' being taken to be numbers, the proper objects of our discussion… But such a conception is untenable, for we | cannot speak of any arithmetical properties of numbers whatsoever without going back to the [meaning] of the signs. For example, the property belonging to 1, of being the result of multiplying itself by itself, would be a mere myth; for no microscopical or chemical investigation, however far it was carried, could ever detect this property in the possession of the innocent character that we call a figure one... The characters we call numerals have, on the other hand, physical and chemical properties depending on the writing material. One could imagine the introduction some day of quite new numerals, just as, e.g., the Arabic numerals superseded the Roman. Nobody is seriously going to suppose that in this way we should get quite new numbers, quite new arithmetical objects, with properties still to be investigated. Thus we must distinguish between numerals and their [meanings]. ("Function and Concept", available here.)

This is very important. When doing semantics, not only of maths, always distinguish numerals and their meanings. '1' is a very different thing from 1. '1' has a shape, and a sound; 1 has neither.

But the key idea is: there is some thing that "1", "I", "一" all have in common. Frege thinks that that thing is their meaning, in some sense. But what is that? Well, that's an open question, one for metaphysicians! But we know it's shared by "1", "I", "一", and has properties like being the result of multiplying itself by itself.

This idea—that words stand in relations to things in the world--is at the heart of the semantic theories we'll be considering.

But we need more than that. Consider:

- sin(16)

Let's grant that '16' means something (the same as 'XVI', for example). What about 'sin'? It's slightly different, it seems. 'sin' is something that, when applied to a number gives a number between -1 and 1 (details, which I'm embarrassingly shaky on, don't matter.)

'Sin' stands for a function, something is something that takes a number and returns a number. It's different from 16 which simply *is* a number. And 'sin(16)' the whole thing, is also a number (about -0.29, google tells me).

So, in elementary mathematics, it seems we need to recognize at least two things. Numbers, and things that take numbers and give back numbers. The latter things, functions, have a sort of incompleteness: they look for an input value to yield an output value. Numbers don't look for anything.

Functions In Natural Language

Frege thought that natural language had the same structure as mathematics: there were objects, and functions. An easy one is things like

- "The mother of Obama"

Frege thought that: 'Obama' stands for Obama; 'the mother of' stands for a function that maps an object to its mother, and the whole phrase, accordingly, stood for the mother of Obama. Part of natural language is just like 'sin(16)'—we apply functions to objects to get other values.

Frege thinks it holds generally:

> "Statements in general, just like equations or inequalities or expressions in Analysis, can be imagined to be split up into two parts; one complete in itself, and the other in need of supplementation, or unsaturated. Thus, e.g., we split up the sentence
>
> 'Caesar conquered Gaul'
>
> into 'Caesar' and 'conquered Gaul'. The second part is unsaturated - It contains an empty place; only when this place is filled up with a proper name, or with an expression that replaces a proper name, does a complete sense appear." (Same reference as above)

So here's the idea: expressions like 'conquered Gaul' are functions. Expressions like 'Caesar' denote objects. But that encourages an important question. Sin is a function that outputs numbers. What does the function 'conquered Gaul' output?

Well, let's think about it. Let's assume we're on board with claims like that '1' stands for a number, and 'sin' for a function. What do we say about something like:

- 7+6

Well, here’s a claim: “+” is a function that yields a number from a pair of numbers.

And now consider

- 10+3
- 14-1
- 26/2

There’s some sense in which these all stand for the same object. Different functions applied to different arguments can yield the same number.

But now consider these *other* two-place functions:

- 13>14
- 10=10
- 20<=20

Note:

i) Functions take numbers as arguments to yield … something. So the whole expression should denote the output of a function (given, respectively, 13 and 14, 10 and 10, and 20 and 20, as inputs)

ii) The above expressions clearly don’t yield numbers: the output isn’t a number.

iii) There’s some sense in which the last two are more similar to each other than either is to the first one—they both share the property of *being true*.

Now here’s the important but weird idea. From i), we need to find something for the function, when given values, to output. It isn’t a number. And on analogy with the fact that we take 10+3 and 14-1 to stand for the same thing, there’s maybe some reason to think the second and third sentences stand for the same thing.

Yes, says Frege. While functions like + or sin take numbers to numbers, functions like =, < and <= take numbers to truth values. There are two truth values, True and False. They are weird abstract objects, that Frege is led to posit because he needs to find something for functions denoted by things like =, < and <= to stand for, with the requirement, or at least possibility, that different function-object(s) pairs can lead to the same object.

We’ve come quite far. We’ve introduced the ideas of objects (numbers), functions (like sin), and truth values. We’re now ready to introduce fully half of the main part of formal semantics.

Idea 3. Compositionality

What, if anything, does all this philosophy of mathematics have to do with language? Well, remember the principle we introduced last week:

Compositionality. The meaning of a complex expression is determined by the meaning of its parts and how they are combined.

What Frege has begun to do for us is show how the meaning of a complex expression, such as “7+5”, is determined by the meaning of its parts. In particular, he’s told us what “7” means; what “5” means, and what “+” means. It makes sense, I hope you agree. But what that means is that Frege has shown us how to capture an important aspect of the compositionality principle, a principle that, to repeat, seems true and something we want to capture.

The great power of the Fregean idea is that if we allow ourselves objects, functions, and truth-values as possible meanings for expressions, we can go quite a long way in explaining compositionality. Let me just specify some meanings this gives:

- Datri sings
- ‘Datri’ =[[Datri]]
- ‘sings’= [the function that takes an x as input and maps to True just in case x sings]]

And a rule:

- If you have a sentence in which—at deep structure—an expression standing for an object and one standing for a function are beside each other, the meaning of the complex expression is the result of applying the function to the object

This rule, function application, can yield a compositional theory of meaning for quite a big chunk of language. Using only these ideas, we can determine the meaning of complex sentences like:

- If everybody comes to the party, we’ll run out of beer and the party will be bad

A couple more examples:

- Datri sings and Datri dances
- ‘and’= [[the function that takes two truth values are inputs and returns True just in case both the inputs are true]]

Now here's a challenging and cool one:

- Everybody sings.
- 'Everybody'= [[the function that takes a function f and returns T provided every person is mapped by f to True]]

This connection of ideas, then, seems to offer promise for capturing one part of meaning. We work out the **deep structure** of a sentence and assign **functions** and objects to its parts, and we see that the result is a system that obeys the intuitive principle of **compositionality**. Great!

Idea 4: Cognitive Significance

Or, maybe not so great. In the framework we're developing, "7+5" and "13-1" are going to mean the same thing. That's because that thing—the number 12—is the thing that results from composing the numerals and function words in the sentence, and we've so far been concentrating on that compositional aspect of meaning.

But if it's an interesting fact that meaning is compositional, it's also an interesting fact that meaning is something that we humans **understand.** And the way we understand meaning places constraints on what our theory of meaning can be.

In particular, here's another fundamental fact:

Fact. Someone can understand "7+5" without understanding "13-1"

But surely

Cognitive Significance. If there are two expressions a and b, and someone can understand a without understanding b, then a and b don't mean the same thing.

Making compositionality and cognitive significance fit together is an open challenge. Next week's reading will begin to show us how to try to do it, by dropping the strong and strange assumption that all true and all false sentences stand for the same thing.