Introduction to natural language semantics
Class 1: what is formal semantics?

Ivano Ciardelli

October 16, 2018

1 What is semantics?

- From Steven Pinker's *The language instinct*:

> As you are reading these words, you are taking part in one of the wonders of the natural world. For you and I belong to a species with a remarkable ability: we can shape events in each other's brains with exquisite precision. I am not referring to telepathy or mind control or the other obsessions of fringe science; even in the depictions of believers these are blunt instruments compared to an ability that is uncontroversially present in every one of us. That ability is language. **Simply by making noises with our mouths, we can reliably cause precise new combinations of ideas to arise in each other's minds.** The ability comes so naturally that we are apt to forget what a miracle it is.

- An example (from Reuters last Thursday):

(1) The two-man U.S.-Russian crew of a Soyuz spacecraft en route to the International Space Station survived a dramatic emergency landing in Kazakhstan on Thursday when their rocket failed in mid-air.

- This piece does not just cause you to think of a spacecraft. It leads you to think of a complex and very specific series of events involving various participants.
- If you trust the source of the sentence, and thus accept (1) as a true description of the world, the content of (1) will be integrated into your worldview.
- In this way, we can learn a great deal about events that we did not personally witness, and collectively build a body of shared knowledge.
- It can be used not just to convey information, as in (1) but also:
  - to impart instructions

(2) Bring water and salt to a boil in a large saucepan; pour polenta slowly into boiling water, whisking constantly until all polenta is stirred in [...]

– request information

(3) Where did the spacecraft land?

– bring new concepts into existence

(4) I have called this principle, by which each slight variation, if useful, is preserved, by the term of Natural Selection.
(Darwin, On the origin of species)

– suggest ideas

(5) The keys might be in the car.

– make proposals

(6) Shall we go for a walk?

- Language can be used to cause humans to think ideas that they have never thought before—hopefully the essence of what is going on in a classroom.
- **How is this achieved?**
- Crucially, not by learning conventions about the meaning of each sentence.
- Sentence (1) has never been used before in the history of the universe, yet the write was able to produce it, and we can understand perfectly well what it means.
- How do we manage to produce pieces of language that encode **new** thoughts, and conversely, how can we decode pieces of language that we have never encountered before?
- Frege was the first to fully realize the importance of the question, and suggest the core of (what we think must be) the answer:

> 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 is 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. (Frege, *Logical Investigations*)

- In more modern terminology, the crucial ideas would read:
  - meanings are composed from parts which correspond to sentence parts;
  - the construction of the meaning mirrors the construction of the sentence;
  - by combining sentence parts in infinitely many different ways we can thus build an infinite number of corresponding meanings.
- Two main ideas: first, human language is a **discrete combinatorial system**; infinitely many signals are obtained by combining a finite number of basic blocks in infinitely many different ways.
- This strategy seems to be unique in the animal world. In other species, it seems that two other strategies are used:
  - A finite, fixed repertoire of signals: e.g., vervet monkeys have specific alarm calls for eagles, snakes, and leopards, which elicit different responses.
  - A signal modulated in strength to convey the degree of a relevant quantity: e.g. bees convey the position of a source of food by performing a waggle dance; the richer a source of food, the more rapidly a bee will waggle.
- Second, at the heart of the working of human language lies the following principle:

  Compositionality principle:
  the meaning of a complex expression is determined by the meanings of its constituent expressions and the rules used to combine them.
- Summing up, we understand (1) because we understand the words that occur in it, and we understand the grammar of the sentence.
- A human language comes with:
  - a lexicon, specifying certain words and corresponding meanings;[1]
  - a grammar, specifying rules to assemble words into sentences, and of corresponding rules to combine the meanings of words to yield the meaning of the sentence.
- We will be concerned with both components: meanings of lexical items, and semantic composition rules.

[1] Actually the atomic constituent of sentences, called a *morpheme*, is often smaller than a whole word, but I'll set this aside here for simplicity.

2 What is *formal* semantics?

- Target phenomenon: meaning composition in natural language.
- We want a concrete, mathematical model of the phenomenon, capable of making predictions that can be tested.
- For this we will use tools from logic, in particular:
  - a formal notion of models
  - a recursive definition of expressions
  - a recursive definition of semantic value of expressions in a model
- Showing that such an approach is possible was the groundbreaking contribution of Richard Montague (1970, 1973).
- After him, a logical semantics of this kind for a fragment of natural language is called a **Montague grammar**.
- Two styles of doing formal semantics:
  - **Direct interpretation**: natural language expressions are mapped directly to semantic objects in the model.
  - **Indirect interpretation**: natural language expressions are translated to expressions in a logical language, which are then mapped to semantic objects in the model.
- In this course, we will use an indirect interpretation.
- That is, we use a logical language with a fixed semantics, and we will interpret natural language by giving:
  - translations of NL words into logical language
  - rules for deriving the translation of a complex expression from the translations of its constituents
- In the end, given a sentence and a specific syntactic analysis of it, the theory will output a formula encoding the meaning of the sentence.
- For instance, to anticipate, we will want our step-by-step process to generate the usual translations for the following sentences:
  - Alice loves Bob $\rightsquigarrow L(a, b)$
  - every professor snores $\rightsquigarrow \forall x(P(x) \rightarrow S(x))$
  - some professor snores $\rightsquigarrow \exists x(P(x) \wedge S(x))$
- Usually, in a first-order logic course, these translations are done by hand.
- By contrast, now we want them to arise in a systematic way given the meanings of the constituents and the composition rules.

- It is natural to assume that, in these sentences:
  - Alice $\rightsquigarrow a$
  - professor $\rightsquigarrow P$
  - snores $\rightsquigarrow S$
- But what about *every*? And what about the compounds *loves Bob*?
- And by what rules are the meanings assembled?
- In the next few classes we will be concerned with filling in the question marks in structures such as:

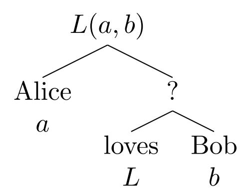

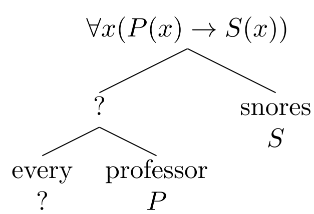

3 What must the theory deliver?

- We want to give lexical entries and rules which jointly can be used to associate each sentence with its meaning.
- We face two important questions:
  1. What is the meaning of a sentence?
     That is, what kind of object should our semantic theory deliver?
  2. How can we tell if the meaning produced by the theory is correct?
     That is, how can our theory be tested empirically?

3.1 What notion of sentence meaning?

- With respect to the first question, the standard answer in semantics is to follow logic: what we want to obtain are the sentence's **truth-conditions**.
- More explicitly: semantics should determine what things have to be like in order for the sentence to be true.
- For some purposes, this view is to narrow:
  - First, language contains not only declarative sentences like (7), but also interrogative sentences like (8) and imperatives like (9), to which the notion of truth does not naturally apply.

    (7) Alice went to Paris.

    (8) Where did Alice go?

    (9) Go to Paris!

    We thus need a separate view about the semantics of interrogatives and imperatives.
  - Second, even for declaratives, we seem to need something more fine-grained for some purposes. For instance, consider:

    (10) Every dog is a dog.

    (11) It rains or it doesn't rain.

    They have the same truth-conditions (always true) but intuitively different meanings (the former is about dogs, the latter about rain).
- Nevertheless, truth-conditions capture important aspects of the meaning of a (declarative) sentence.
  - They determine the sentence's **informative content**:
    by asserting a sentence, a speaker conveys the information that the state of affairs is one of those where the sentence is true.
  - They allow us to characterize the relation of **entailment**:
    α entails β if in all circumstances where α is true, β is true as well

3.2 How to assess the predictions of a theory?

- The truth-conditional perspective on meaning naturally suggests two ways to assess the predictions of a theory.
- **Truth-value judgments**: check whether speakers of the language judge the sentence as true/false in the circumstances predicted by the theory.
  - Ex: "Some cats are sleeping" should be predicted false in a situation where no cat is sleeping.
- **Entailment judgments**: check whether speakers judge a sentence as following from another in accordance with the theory's entailment predictions.
  - Ex: "Some cats are sleeping" should be predicted to entail "Some animals are sleeping".
- However, the picture is complicated by two important linguistic phenomena which need to be factored in:
  - implicatures
  - presuppositions

3.3 Implicatures

- Consider (12):

(12) a. Alice can speak English, French, and Spanish.
b. $\rightsquigarrow$ Alice doesn't speak German.

- But is this an entailment?
- It seems not: in a situation where Alice speaks English, French, Spanish, and German, (12-a) seems true and (12-b) false.
- It seems that (12-b) does not follow from the literal meaning of (12-a), but rather from the assumption that the speaker is giving us a complete description of the situation.
- Inferences like the one in (12) do not involve only the literal meaning of sentences, but also the assumptions that speakers make about one another's behavior when using sentences.
- Via these assumptions, speakers manage to convey more than what they literally say:

(13) a. Alice went to London or to Paris.
b. $\rightsquigarrow$ speaker believes Alice went to London or Paris.
c. $\rightsquigarrow$ speaker doesn't know which one.

- The extra content which is conveyed in this way is called an **implicature**.
- Another (real) example: American movie producer Samuel Bronston was being questioned under oath during a bankruptcy hearing.

  Q: Do you have any bank accounts in Swiss banks, Mr. Bronston?

  A: No, sir.

  Q: Have you ever?

  A: The company had an account there for about six months, in Zurich.

- We hear Bronston's response as conveying that he himself never had an account in Switzerland. It was later discovered that, in fact, Bronston had had a Swiss bank account for five years.
- Is this a case of perjury?
- The US Supreme court ruled that it wasn't.
- Indeed, Bronston's answer was literally true; however, by means of it Bronston implicated something false.
- The implicature comes about because as listeners we interpret B.'s reply as a full answer to the question.

  If he had had a bank account himself, it would be relevant to say so; since he didn't, we take it that *only* the company had a Swiss bank account, not him.

- Notice that in normal situations, where interlocutors are cooperative, such a conclusion would indeed be sensible; we make such inferences all the time:

  A: Did you read *War and Peace*?

  B: I read a summary of it.

  $\rightsquigarrow$ B didn't read *War and Peace.*

  A: Are you coming to the party?

  B: I have to work.

  $\rightsquigarrow$ B is not coming.

- Implicatures arise from the way in which language is *used* by rational agents.
- This aspect of language is is the subject of a discipline called **pragmatics**, which originated with Grice (1975).
- How can we distinguish entailments from implicatures?
  - Direct truth-value intuitions. E.g., in a situation where Alice went to London (14) seems true even though the speaker is not uncertain. This indicates that the ignorance associated with *or* is an implicature.

(14) Alice went to London or to Paris.

- Truth-conditional content, unlike implicatures, can be embedded under operators like negation. Thus, by viewing what content is negated we can distinguish entailments from implicatures.

(15) a. Alice did not go to London or to Paris.
b. $\approx$ she went to neither
c. $\not\approx$ either she went to neither or speaker is informed

- Nevertheless, in some cases the situation is not clear-cut.
- One examples are so-called *upper-bounding inferences* with numerals.

(16) a. Bea has three children.
b. $\rightsquigarrow$ Bea does not have four children.

- Is (16-a) false if Bea has four children?
- Consider the negation:

(17) a. Bea does not have three children.
b. $\overset{?}{\approx}$ Bea has less than three children.
c. $\overset{?}{\approx}$ Bea has a number of children different from 3.

- In cases such as these, the same inferences are explained either semantically (as entailments) or pragmatically (as implicatures) by different theories.
- In this course we will not deal with pragmatics; still it is important to be aware of the distinction between entailments and implicatures, and of the fact that when a statement is made, we "hear" more into it than its literal content.

3.4 Presuppositions

- Sometimes, truth-value intuitions are not sharp either.

(18) The largest natural number is prime.

(19) Munich is no longer the capital of Germany.

- We are reluctant to judge such sentences as true or false.
- The complication is that (18) and (19) are associated with *presuppositions*.
- Presuppositions can be characterized as entailments which persist under negation (and under the operation of turning the sentence into a question).

(20) a. The largest natural number is prime.
b. $\rightsquigarrow$ There is a largest natural number.

(21) a. The largest natural number is not prime.
b. $\rightsquigarrow$ There is a largest natural number.

(22) a. Is the largest natural number prime?
b. $\rightsquigarrow$ There is a largest natural number.

- Notice that this is not the same for ordinary entailments.

(23) a. Alice loves Bob.
b. $\rightsquigarrow$ Someone loves Bob.

(24) a. Alice does not love Bob.
b. $\not\rightsquigarrow$ Someone loves Bob.

(25) a. Does Alice love Bob?
b. $\not\rightsquigarrow$ Someone loves Bob.

- How should our semantics account for presuppositions?
- A simple option: view the truth-conditions of a sentence as a **partial function**, defined only when the presuppositions of the sentence are true.
- So, presuppositions can be dealt with in truth-conditional semantics.
- Nevertheless, there are cases where intuitions are not very clear-cut.

(26) The king of France is sitting in this chair.

4 What logic?

- What kind of logical language can we used ?
- How about first-order logic?
- Problem: unclear how we can fill the slots below:

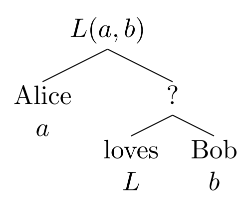

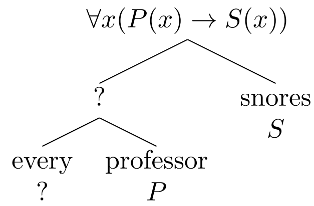

- Again Frege comes to help:

> In my essay "Negation", I considered the case of a thought that appears to be composed of one part which is in need of completion or, as one might say, unsaturated, and whose linguistic correlate is the negative particle, and another part which is a thought. We cannot negate without negating something, and this something is a thought. Because this thought saturates the unsaturated part or, as one might say, completes what is in need of completion, the whole hangs together. And it is a natural conjecture that logical combination of parts into a whole is always a matter of saturating something unsaturated. (Frege, *Logical Investigations*)

- Some expressions in natural language are unsaturated; a complete content arises when their slots are saturated by the content of every expressions with which they combine:
  - loves Bob $\approx$ $L(\_, b)$
  - every professor $\approx$ $\forall x(P(x) \rightarrow \_(x))$
- But what does this "saturating" talk mean precisely?
- Frege proposed to model "saturation" in terms of the mathematical notion of **application of a function to an argument**.
- Modern formal semantics has built on this insight: the main (though not the only) rule underlying semantic composition is taken to be **function application**.
- Accordingly, natural language expressions should be translated in a framework which has **expressions for functions**.
- Moreover, these functions should be **typed**, i.e., they expect an argument of a specific semantic type: loves Bob expects in input an entity, while every professor expects a property.
- For these reasons, the main framework used for formal semantics is the **typed λ-calculus** (Church, 1940), a formalism developed to deal with functions with restrictions on the semantic types of their arguments.
- In typed λ-calculus, we will be able to translate our "unsaturated expressions" as follows:
  - loves Bob $\rightsquigarrow$ $\lambda x.L(x, b)$
  - every professor $\rightsquigarrow$ $\lambda Q.\forall x(P(x) \rightarrow Q(x))$

5 Typed λ-calculus

5.1 General architecture

- Syntax:
  - a universe T of types
  - for each type $\sigma \in T$, a set of expressions $\alpha : \sigma$ of that type
- Semantics:
  - for each type σ, a domain D_σ^M of objects of type σ
  - for each expression $\alpha : \sigma$, an object $[\![\alpha]\!]^{M,g} \in D_\sigma^M$

	Syntax	Semantics
Types	σ	D_σ^M
Expressions	$\alpha : \sigma$	$[\![\alpha]\!]^{M,g} \in D_\sigma^M$

5.2 Types

- The set of types is given by the following inductive definition:
  - Basic types: (this repertoire will be enriched later)
    - e: the type of individuals
    - t: the type of truth-values
  - Inductive clause:
    - if σ, τ are types, then $\langle \sigma, \tau \rangle$ is a type:
      the type of functions from objects of type σ to objects of type τ
- Particular cases:
  - $\langle e, t \rangle$: the type of functions from individuals to truth-values (predicates)
  - $\langle e, \langle e, t \rangle \rangle$: the type of functions from pairs of individuals to truth-values (binary relations)

5.3 Domains for each type

- A model M will supply a domain D of entities.
- We then recursively associate a set of objects to each type.
  - $D_e^M := D$
  - $D_t^M := \{0, 1\}$
  - $D_{\langle \sigma, \tau \rangle}^M := (D_\sigma^M)^{D_\tau^M} = \{f \mid f : D_\sigma^M \to D_\tau^M\}$
- In particular:

- $D^M_{\langle e,t\rangle} := \{f \mid f : D \to \{0,1\}\}$, which can be identified with $\wp(D)$ via:
  * if $f : D \to \{0,1$ then $f \mapsto \{d \in D \mid f(d) = 1\}$
  * if $S \subseteq D$ then $S \mapsto f(x) =$
- $D^M_{\langle e,t\rangle} := \{0,1\}^D \approx \wp(D)$
- $D^M_{\langle e,t\rangle} := (\{0,1\}^D)^D \approx \{0,1\}^{D\times D} \approx \wp(D \times D)$

• When a model is implicit, we omit reference to it.

References

Alonso Church. A formulation of the simple theory of types. *Journal of Symbolic Logic*, 5:56–68, 1940.

H.P. Grice. Logic and conversation. In Peter Cole and Jerry Morgan, editors, *Syntax and Semantics*, volume 3, pages 41–58. Harvard University Press, 1975.

Richard Montague. Universal grammar. *Theoria*, 36(3):373–398, 1970. ISSN 1755-2567.

Richard Montague. The proper treatment of quantification in ordinary English. In J. Hintikka, J.M.E. Moravcsik, and P. Suppes, editors, *Approaches to natural language*, pages 221–242. Reidel, Dordrecht, 1973.