# Foundations of Natural Language Processing
# Lecture 18a
# Introduction to Semantics

Alex Lascarides

# Meaning

**Last Time:** Syntactic parsing

**Now:** How can syntax inform meaning?

- Desiderata for formal semantic representations of sentences
- Method for constructing semantic representations from a syntax tree.

We will start with:

- What is meaning?

- What should a representation of meaning capture?

# Meaning

- One grand goal of artificial intelligence is to *understand* what people mean when they talk.

- But how do we know if we succeeded?

# Literal Meaning

Semantics is concerned with how expressions relate to 'the world'. This includes both their

- truth conditions and denotation (literal meaning)
- connotation (other associations)

For now, we're interested only in literal meaning.

# Eliza

Joseph Weizenbaum wrote the computer program **Eliza** in 1969 to demonstrate how easily people can be fooled into thinking that machines have some deep understanding.

**young woman:** Men are all alike.
**eliza:** In what way?
**young woman:** They're always bugging us about something specific or other.
**eliza:** Can you think of a specific example?
**young woman:** Well, my boyfriend made me come here.
**eliza:** Your boyfriend made you come here?
**young woman:** He says I'm depressed much of the time.
**eliza:** I am sorry to hear you are depressed.
**young woman:** It's true. I'm unhappy.
**eliza:** Do you think coming here will help you not to be unhappy?

Online demo: `http://www.manifestation.com/neurotoys/eliza.php3`

# What is meaning? What is understanding?

- These are deep philosophical questions

- NLP usually takes a more pragmatic view: can the computer behave as though it understands (in order to do what we want)?

  – Dialogue systems (e.g., Eliza)
  – Machine translation
  – Question answering

- What issues will we face in building such systems?

# A Concrete Goal

- We would like to build
  - a machine that answers questions in natural language.
  - may have access to knowledge bases
  - may have access to vast quantities of English text
- Basically, think Alexa!
- This is typically called **Question Answering**

# Semantics

- To build our QA system we will need to deal with issues in **semantics**, i.e., meaning.

- Sentential semantics: how word meanings combine (study now)
  - Who did what to whom; when, how, why. . .

    John loves Mary $\neq$ Mary loves John

    $\Rightarrow$ Someone loves Mary

- Lexical semantics: the meanings of individual words (study after that)
  E.g., John is male, Mary is female,
  loves is more closely related to like than to sees, antonym of hate. . .

# What we've learned so far about sentential syntax

- It captures linguistic generalisations about grammaticality (*substitutability*)
- It generates an unbounded set of grammatical sentences via a finite lexicon and finite rules (*recursion*)
- We can induce probabilistic grammars from a treebank, and so tackle (pervasive) syntactic ambiguity.

# What about Meaning? Spoiler Alert!

- Sentential syntax reveals information about sentence meaning

| | | |
|---|---|---|
| John loves Mary | $\mapsto$ | $love(j, m)$ |
| Mary is loved by John | | |
| Mary loves John | $\mapsto$ | $love(m, j)$ |

- Decisions about how to resolve syntactic ambiguity are tied up with decisions about (intended) meaning.

- Syntactic ambiguity (almost) always yields a semantic ambiguity.

- Resolving syntactic ambiguities does *not*, however, resolve all semantic ambiguities

  - word sense, semantic scope, anaphoric expressions
    (*all to be studied later in this course*)

  so reasoning about context is also very important
  (*also studied later in this course*)

# Summary

- Meaning representations are important for many NLP tasks because they captures:
  - Who did what to whom in a way that abstracts away from (irrelevant) syntactic details
  - Valid inference, and information about reference

- Linguistic syntax informs meaning (and *vice versa*)

# What we'll study now. . .

- Principle of Compositionality

- Exploit compositionality to augment a grammar with a semantic component, which *deterministically derives* the logical form of a sentence from its syntax tree.