Computational Semantics and Pragmatics

Autumn 2011

Raquel Fernández

Institute for Logic, Language & Computation
University of Amsterdam

What is this course about?

About the semantics and pragmatics of natural language — about meaning and interpretation in context, and about language use in interaction. Some general key questions we will address are:

- how can we model the meaning of words?
- what kind of inferences can we draw from sentences or discourses?
- how do we use and interpret language in dialogue?

The course is also about using computational and empirical methods to explore semantic/pragmatic phenomena

- computational resources such as linguistic corpora and databases
- algorithms and automatic tools

Related Courses

- This is a new course at the interface of the *Logic & Language* and the *Language & Computation* groups at the ILLC.
- (Mildly) related courses within the MoL:
  - Structures for Semantics (*Maria Aloni / Robert van Rooij*)
  - Meaning, Reference and Modality (*Paul Dekker*)
  - Inquisitive Semantics (*Jeroen Groenendijk / Floris Roelofsen*)
  - Language and Optimality (*Reinhard Blutner / Henk Zeevat*)
  - Mechanisms of Meaning (*Henk Zeevat*)
  - Elements of Language Processing and Learning (*Khalil Sima'an*)
  - Cognitive Models of Language and Beyond (*Rens Bod*)

Prerequisites

No formal prerequisites are required to follow the course. However, some basic things are expected from you:

- an interest in natural language, particularly in semantics/ pragmatics - in meaning, interpretation, and interaction.
- an empirical orientation: an interest in the empirical evidence (or lack thereof) behind theoretical claims; and in working with data.
- a computational inclination: an interest in computational methods of enquiry and evaluation
  - ∗ does this mean you need to know how to program? No! but if you do, then you'll have the chance to use your programming skills.

⇒ Please fill in the **student questionnaire** on the website to let me know about your background and interests.

Practical Matters

- **Lecturer:** Raquel Fernández, <raquel.fernandez@uva.nl>
- **Website:** Slides, references, and other important information will be posted on the course's website: http://www.illc.uva.nl/~raquel/teaching/cosp2011/
- **Timetable:** Thursday 15-17 in D1.168 till 27 Oct, then G2.04
  * no class in the following two weeks; need to find a different slot in the week of 26 Sept. [*we'll discuss this later*]
- **Seminars:** There may be talks at the ILLC that are relevant to the course and that you are welcome/encouraged to attend:
  * Computational Linguistics Seminar (CLS)
  * DIP (discourse processing) Colloquium

  Check the ILLC Events webpage for details.

Evaluation

- Homework exercises involving any of the following:
  - ∗ analytical thinking
  - ∗ use of online corpora and web interfaces to examine data
  - ∗ running algorithms to obtain results
- Reading relevant research papers and presenting or discussing them (to be made more concrete later on)
- Individual final paper to be presented at the end of the course
  - ∗ on-topic philosophical/theoretical essays are in principle OK, but
  - ∗ ideally, your project should include an empirical/computational component, e.g. analysis of real data or some sort of implementation

Plan for today

1. Overview of the main topics of the course
2. Introduction to Textual Entailment

Overview of Course Topics

Meaning and Understanding

How can we characterise what understanding natural language is?
This is a tough question and there are plenty of proposals...

- to know the meaning of a (declarative) sentence is to know what the world would have to be like for the sentence to be true
- ... to know how it changes the context (by adding knowledge, by making relevant follow-up expressions,etc.)
- ... to be able to use an expression appropriately given the conventions of a linguistic community
- ... to be able to (re)act according to what is expected ...

All these *takes* on meaning and understanding can be seen as complementary: all possibly necessary but none sufficient.

Meaning and Inference

Another necessary condition for natural language understanding is the ability to recognise entailment and contradiction.

- If you understand these sentences, you can recognise that (1) and (2) are contradictory ...

(1) No civilians were killed in the Najaf suicide bombing.
(2) Two civilians died in the Najaf suicide bombing.

- ... and that if (3) is true then (4) is true as well.

(3) Apple filed a lawsuit against Samsung for patent violation.
(4) Samsung has been sued by Apple.

Recognising whether entailment holds is a core aspect of our ability to understand language.

Recognising Textual Entailment

Textual Entailment is a notion broader than logical entailment defined by the computational linguistics community as follows:

> Textual entailment is a relation that holds between a pair $\langle T, H \rangle$ of natural language expressions (a *text* and a *hypothesis*), such that a human who reads (and trusts) T would infer that H is most likely true.

RTE can be seen as an abstract generic ability that captures inferential/semantic capabilities required by many tasks involving understanding.

⇒ How can we model this ability computationally? Challenges include:

- characterising the *sources* of the entailment (syntactic, semantic,...)
- background knowledge
- ambiguity

Lexical Semantics

Next, we will move on to lexical semantics (meaning of words).

Formal compositional semantics employs a rather crude notion of lexical meaning:

$$
\begin{array}{lll}
[\![dolphin]\!] = \{x \mid x \text{ is a dolphin}\} & f : D \rightarrow \{1,0\} & \langle e,t\rangle \\
[\![envy]\!] = \{\langle x,y\rangle \mid x \text{ envies } y\} & f : D \rightarrow (D \rightarrow \{1,0\}) & \langle e,\langle e,t\rangle\rangle
\end{array}
$$

How can we model word senses and the relations that hold between them in a more fine-grained manner?

- Hyponymy and Hypernymy: relation of semantic inclusion that holds between a more general term such as *'bird'* and a more specific term such as *'robin'*
- Synonymy: relation of semantic identity between senses, e.g. *'aurora/dawn/sunrise'*, *'whore/prostitute'*
- Antonymy: relation of semantic oppositeness between senses, e.g. *'tall/short'*, *'dead/alive'*
- Meronymy: part-whole relation between senses, e.g. *'elbow/arm'*, *'keyboard/computer'*

Distributional Semantic Models

We will focus on Distributional Semantic or Vector Space Models.

- These models take a *usage-based* view of word meaning.
- Their basic underlying idea is that word meaning depends on the contexts in which words are used.
- An example by Stefan Evert: what's the meaning of *'bardiwac'*?

∗ He handed her her glass of bardiwac.
∗ Beef dishes are made to complement the bardiwacs.
∗ Nigel staggered to his feet, face flushed from too much bardiwac.
∗ Malbec, one of the lesser-known bardiwac grapes, responds well to Australia's sunshine.
∗ I dined on bread and cheese and this excellent bardiwac.
∗ The drinks were delicious: blood-red bardiwac as well as light, sweet Rhenish.

$\Rightarrow$ *'bardiwac'* is a heavy red alcoholic beverage made from grapes

The Distributional Hypothesis

- DH: The degree of semantic similarity between two linguistic expressions A and B is a function of the similarity of the linguistic contexts in which A and B can appear (Harris, 1954)
- DSMs make use of mathematical and computational techniques to turn the informal DH into empirically testable semantic models.
- They build contextual semantic representations from data about language usage.
- These representations are defined as an abstraction over the linguistic contexts in which a word is encountered.

⇒ We will study the philosophical ideas behind these models and the computational techniques currently used to build them.

Conversational Implicature

Then we'll move on to more typically pragmatic issues...
Entailments are not the only inferences we are able to make when we understand language in context:

A: Which room is the seminar in next week?
B: It's in the G building.
$\rightsquigarrow$ A does not know in which room the seminar is.

A: Where can I get gas around here?
B: There is a garage around the corner.
$\rightsquigarrow$ A can get gas at a garage around the corner.

According to the philosopher Paul Grice, we are able to make inferences like the ones above, called conversational implicatures, because we follow general principles of cooperation.

The Cooperative Principle: Make your contribution such as it is required, at the stage at which it occurs, by the accepted purpose or direction of the talk exchange in which you are engaged. (Grice 1975)

The Gricean Maxims of Conversation

- Maxim of Quality: be truthful
  - Do not say what you believe to be false.
  - Do not say that for which you lack adequate evidence.
- Maxim of Quantity:
  - Make your contribution as informative as is required (for the current purposes of the exchange).
  - Do not make your contribution more informative than is required.
- Maxim of Relevance: be relevant
- Maxim of Manner: be perspicuous.
  - Avoid obscurity of expression / Avoid ambiguity.
  - Be brief / Be orderly.

H.P. Grice (1975) Logic and conversation.In *Syntax and Semantics, Vol. 3, Speech Acts*, ed. by Peter Cole and Jerry L. Morgan. New York: Academic Press.

Computational Exploration of Implicature

We'll look into different computational explorations of implicature, focusing on two issues (time permitting):

- Interpretation: indirect answers. Polar questions are not always answered with a plain 'yes'/'no'. The intended answer is often *implicated*.

A: Do you want to go ahead and start?	A: Is the new judge really this good?
B: I was hoping that you would.	B: He is great.

⇒ Can we use real data to automatically predict whether the answer intended to convey "yes" or "no"?
To investigate this we'll use data and code developed by Chris Potts.

- Production: generation of referring expressions

⇒ how can we automatically generate expressions that refer to entities in the context obeying the maxims of conversation, i.e. without generating unintended implicatures?

We'll also see that implicature may also play a role in lexical semantics – in computing the relevant sense of words.

Dialogue Interaction

Finally we'll put the interpretation (hearer) and the production (speaker) perspective together to look into dialogue.

- Dialogue or conversation is the most basic setting for language use.
- Dialogue is a form of interaction and brings in extra challenges.
- Crucially, it involves multiple participants, which requires coordination.
  - ∗ content coordination: utterances in a dialogue are connected to form a coherent discourse; speakers need to avoid misunderstanding.
  - ∗ interaction coordination: turn-taking (who speaks when) and integration of language with other modalities (gestures, gaze, . . .)

From the British National Corpus (KP5):

```
A: Did you get your tickets for Crowded House?
B: No!
   There is not one ticket left in the entire planet!
   So annoying!
C: Where for?
B: Crowded House.
   My brother is going and he doesn't even like them.
A: Why doesn't he sell you his ticket?
B: Cos he's going with his work.
   And Sharon.
A: Oh, his girlfriend?
B: Yes.
   They are gonna come and see me next week.
A: Not Sharon from Essex?
B: No, she's Sharon from <laughing> Australia.
A: Oh, alright then.
B: That's the only reason I forgive him.
   <laughing> Cos she's not born in this country!
```

Burnard (2000) *Reference Guide for the British National Corpus (World Edition)*, Oxford Univ. Computing Services.

Dialogue Issues

We'll look into the following dialogue phenomena, drawing on data from dialogue corpora and discussing how these phenomena are treated in current computational dialogue research.

- Dialogue acts ($\sim$ *speech acts*): what DA types are there? how can we identify them? what are they useful for, what do they tell us about dialogue?
- Grounding, the process by which DPs establish mutual understanding, including feedback strategies such as clarification requests.
- Alignment processes by which DPs converge in their choice of linguistic forms

Raquel Fernández (2011 draft) *Dialogue*, The Oxford Handbook of Computational Linguistics (2nd Edition). To appear.

Break

Textual Entailment

Textual entailment is a relation that holds between a pair $\langle T, H \rangle$ of natural language expressions (a *text* and a *hypothesis*), such that a human who reads (and trusts) T would infer that H is most likely true.

T	H	TE
Eyeing the huge market potential, currently led by Google, Yahoo took over search company Overture Services Inc last year.	Yahoo bought Overture.	✓
Since its formation in 1948, Israel fought many wars with neighboring Arab countries.	Israel was established in 1948.	✓
The National Institute for Psychobiology in Israel was established in May 1971 as the Israel Center for Psychobiology by Prof. Joel.	Israel was established in May 1971.	×
Arabic is used densely across North Africa and from the Eastern Mediterranean to the Philippines, as the key language of the Arab world and the primary vehicle of Islam.	Arabic is the primary language of the Philippines.	×

General information (references, resources, etc.) can be found on the Textual Entailment Portal of the ACL:

http://aclweb.org/aclwiki/index.php?title=Textual_Entailment

Applications of Textual Entailment (1)

RTE can be seen as an abstract generic task that captures inferential/semantic capabilities required by many applications.

Automatic summarization: creation of a *reduced version* of a text, document, or set of documents.

- In general, a summary should be entailed by the source text.
- Two main methods: extraction and abstraction
- Abstraction requires natural language generation
- Extraction should avoid redundancy: dismiss sentences that are entailed by previously selected sentences.

Applications of Textual Entailment (2)

Question-answering (QA): automatically answering a question posed in natural language.

Basic architecture of a QA system:

- question classifier: determines the type of answer required (who/what/why/how/...)
- retrieval module: selects relevant documents (search engine)
- filter: selects candidate expressions from the retrieved documents that match the answer type
- answer extractor: selects an answer among the candidates

- Q: Who painted Guernica? $\rightsquigarrow$ A: X painted Guernica
- T: Guernica is grey, black and white, 3.5 metres tall and 7.8 metres wide, a mural-size canvas painted in oil. Picasso's purpose in painting it was to bring the world's attention to the bombing of the Basque town of Guernica by German bombers, who were supporting the Nationalist forces of General Franco during the Spanish Civil War.
- H_1: Picasso / H_2: German bombers / H_3: General Franco painted Guernica
- **Answer extraction via entailment**: Does T entail H_1 / H_2 / H_3 ?

Applications of Textual Entailment (3)

Machine translation: automatically translating a text into a different language.

- Automatic translations are evaluated against *gold standard* human translations.
  - ∗ an automatic translation should be semantically equivalent to a human one: both translations should entail each other.
- Entailment may also be used to find alternatives that are easier to translate by the system
  - ∗ the source text T may be paraphrased into T', where T entails T' (assuming the system is able to translate T' but not T)

T: Apple files a lawsuit against Samsung for patent violation.
T': Apple accuses Samsung of patent violation.

Other applications: e.g. text simplification or automatic scoring of student answers

Textual Entailment and Logic

We may want to think of TE in terms of logical entailment:

Let the logical meaning representations of T and H be ϕ_T and ϕ_H, and B be a conjunction of axioms or knowledge base.
If $(\phi_T \wedge B) \models \phi_H$, then $\langle T, H \rangle$ is a correct textual entailment pair.

Obvious challenges:

- assigning ϕ_T and ϕ_H to natural language expressions T and H
- defining B
- checking whether $(\phi_T \wedge B) \models \phi_H$ holds

What do these challenges involve and how can we address them with computational tools?

Ambiguity

Natural language is highly ambiguous. Ambiguity may be syntactic (multiple structural groupings), lexical (multiple parts of speech), semantic (multiple word senses or compositional interpretations), pragmatic (multiple available referents), ...

(5) Two sisters reunited after 18 years in checkout counter
two sisters [reunited [after 18 years] [in checkout counter]]
two sisters [reunited [after 18 years [in checkout counter]]]

(6) Squad helps dog bite victim [helps [[dog bite] victim]] / [helps [dog] [bite victim]]

(7) Teacher strikes idle kids — strikes: V/N idle: V/A

(8) Iraqi head seeks arms — head: body part/leader arms: body part/weapons

(9) The French agreement. — French: by the French/in France/...

(10) If the baby doesn't thrive on cows' milk, boil it. — it = the baby / the milk

Humans are usually able to quickly select one reading in a given context. But ambiguity resolution is a huge problem for computational systems that aim at natural language understanding.

Ambiguity and Textual Entailment

Assigning logical representations ϕ_T and ϕ_H to T and H requires disambiguation - selecting particular readings of these expressions.

T: A bomb exploded near the French bank.
H: A bomb exploded near a building.

In the above example, T entails H only if the word *'bank'* is used with the sense "office or quarters of a financial institution".

Thus, disambiguation, such as word sense disambiguation, is needed to assign logical representations that allow us to check if the logical entailment holds.

Background Knowledge

What kind of knowledge B is required to check whether $(\phi_T \wedge B) \models \phi_H$ holds and where can we get it from?

- we may construct a knowledge base by hand ...
- we may extract it from online resources, such as Wikipedia, or for lexical knowledge bases such as WordNet

  (see homework)

Automated Reasoning

Assuming we have been able to assign ϕ_T and ϕ_H and define B, how do we check whether $(\phi_T \wedge B) \models \phi_H$ holds?

- $(\phi_T \wedge B) \models \phi_H$ is true iff whenever $(\phi_T \wedge B)$ is true, ϕ_H is true.
- that amounts to checking whether the following formula is valid, i.e. true in all possible models

$$(\phi_T \wedge B) \wedge \phi_H$$

- Can we check validity computationally? Not for first-order logic: FOL is undecidable, that is:
  - ∗ there is no algorithm capable of checking validity in finite time for all possible input formulas.

  Fortunately there are some *partial solutions* we can exploit: we'll see how to use theorem provers and model builders in combination to tackle the problem in the next class.

Readings

- We'll discuss the following paper in the next class:

Johan Bos & Katja Markert (2005) Recognising Textual Entailment with Logical Inference. In *Proc. of HLT/EMNLP*.

- For a general overview on Textual Entailment see:

Ion Androutsopoulos & Prodromos Malakasiotis (2010) A Survey of Paraphrasing and Textual Entailment Methods. *Journal of Artificial Intelligence Research*, vol. 38, pp. 135-187.

What's next?

- Check the website of the course; it will be updated tonight with a link to the slides and other material, including an "overview bibliography" which you can browse to get a better impression of what the course will be about.
- Please fill in the student questionnaire!
- Attend the CLS seminar next week
- On the website you'll find a link to Homework # 1, which needs to be submitted by September 22.
- Need to fix day and time of next class ...