Semantics for Textual Inference

Cleo Condoravdi
Annie Zaenen

NSF funded Workshop
LSA Institute
University of Colorado at Boulder
July 9-10, 2011

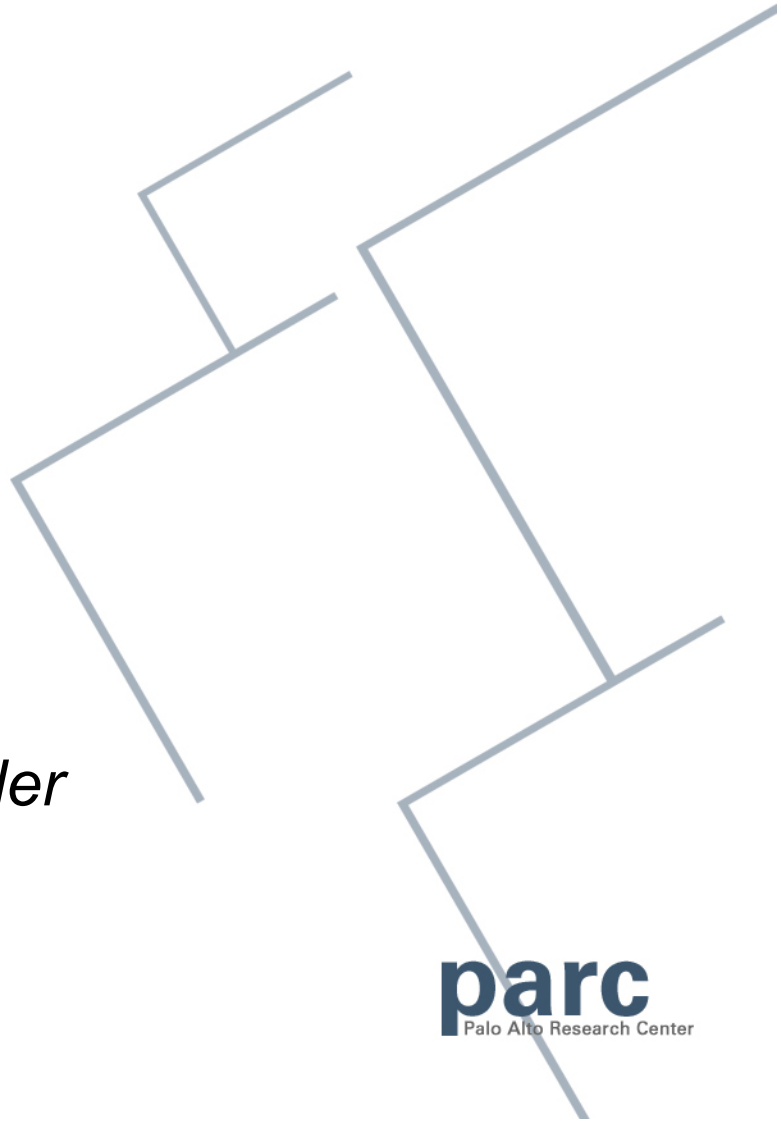

parc
Palo Alto Research Center

Motivation

Within the linguistic computational world a common perspective has emerged on what is common to these natural language understanding tasks under the heading “textual inferencing”. The aim is to develop systems that can decide, when given two natural language statements, what the inferential relation between the two is. Textual inference simplifies the general language understanding problem by limiting its interest to direct inferences avoiding complicated chains of inferences and specialized world knowledge. Semantics as practiced by linguists could play a role in the development of textual inference systems, but most of current work in linguistic semantics has a very different focus. This workshop aims to bring together researchers interested in semantics and in computational textual inference to discuss the virtues and drawbacks of various semantic approaches. The aim of the workshop is to make the community of semanticists more aware of the computational issues in natural language understanding and to expose computer scientists to a variety of semantic approaches.

parc
Palo Alto Research Center

Toward NL Understanding

A measure of understanding a text is the ability to make inferences based on the information conveyed by it.

parc
Palo Alto Research Center

Access to content: existential claims

What happened? Who did what to whom?

Microsoft managed to buy Powerset.

⇒ *Microsoft acquired Powerset.*

Shackleton failed to get to the South Pole.

⇒ *Shackleton did not reach the South Pole.*

The destruction of the file was not illegal.

⇒ *The file was destroyed.*

The destruction of the file was averted.

⇒ *The file was not destroyed.*

parc
Palo Alto Research Center

Access to content: monotonicity

What happened? Who did what to whom?

Every boy managed to buy a small toy.

⇒ *Every small boy acquired a toy.*

Every explorer failed to get to the South Pole.

⇒ *No experienced explorer reached the South Pole.*

No file was destroyed.

⇒ *No sensitive file was destroyed.*

The destruction of a sensitive file was averted.

⇒ *A file was not destroyed.*

parc
Palo Alto Research Center

Access to content: temporal domain

What happened when?

Ed visited us every day last week.

⇒ *Ed visited us on Monday last week.*

Ed has been living in Athens for 3 years.

Mary visited Athens in the last 2 years.

⇒ *Mary visited Athens while Ed lived in Athens.*

The deal lasted through August, until just before the government took over Freddie. (NYT, Oct. 5, 2008)

⇒ *The government took over Freddie after August.*

parc
Palo Alto Research Center

Toward NL Understanding

Local Textual Inference

- "Direct inferences": no complicated chains of inferences, no (specialized) world knowledge

Veridicality reasoning

- Did an event mentioned in the text actually occur?

Temporal reasoning

- When did an event happen? How are events ordered in time?

Spatial reasoning

- Where are entities located and along which paths do they move?

Causality reasoning

- Enablement, causation, prevention relations between events

Knowledge about words for access to content

The verb “acquire” is a hypernym of the verb “buy”
The verbs “get to” and “reach” are synonyms

Inferential properties of “manage”, “fail”, “avert”, “not”

Monotonicity properties of “every”, “a”, “no”, “not”

Every (↓) (↑), A (↑) (↑), No(↓) (↓), Not (↓)

Restrictive behavior of adjectival modifiers “small”, “experienced”, “sensitive”

The type of temporal modifiers associated with prepositional phrases headed by “in”, “for”, “through”, or even nothing (e.g. “last week”, “every day”)

Construction of intervals and qualitative relationships between intervals and events based on the meaning of temporal expressions

parc
Palo Alto Research Center

Types of Approaches

"Shallow" approaches: many ways to approximate

- String-based (n-grams) vs. structure-based (phrases)
- Syntax: partial syntactic structures
- Semantics: relations (e.g. triples), semantic networks

➳ Confounded by negation, syntactic and semantic embedding, long-distance dependencies, quantifiers, *etc*.

"Deep(er)" approaches

- Syntax: full syntactic analysis
- Semantics: a spectrum of meaning representations depending on aspects of meaning required for the task at hand

➳ Scalability

parc
Palo Alto Research Center

Recognizing Textual Inferences

parc
Palo Alto Research Center

Textual Inference Task

Does premise *P* lead to conclusion *C*?

Does text *T* support the hypothesis *H*?

Does text *T* answer the question *H*?

… without any additional assumptions

P*: Every explorer failed to get to the South Pole.*

C*: No experienced explorer reached the South Pole.*

Yes

parc
Palo Alto Research Center

Local Textual Inference Initiatives

PASCAL RTE Challenge (Ido Dagan, Oren Glickman) 2005, 2006

PREMISE

CONCLUSION

TRUE/FALSE

Rome is in Lazio province and Naples is in Campania.

Rome is located in Lazio province.

TRUE (= entailed by the premise)

Romano Prodi will meet the US President George Bush in his capacity as the president of the European commission.

George Bush is the president of the European commission.

FALSE (= not entailed by the premise)

parc
Palo Alto Research Center

Entailment and Contradiction Detection

Text:		Kim hopped.
Hypothesis:		Someone moved.
Answer:	TRUE	

Text:		Sandy touched Kim.
Hypothesis:		Sandy kissed Kim.
Answer:	UNKNOWN	

Text:		Sandy kissed Kim.
Hypothesis:		No one touched Kim.
Answer:	NO	

parc
Palo Alto Research Center

Entailment and Contradiction Detection

Text:	Kim hopped.
Hypothesis:	Someone moved.
Answer:	TRUE

Text:	Sandy touched Kim.
Hypothesis:	Sandy kissed Kim.
Answer:	UNKNOWN

Text:	Sandy kissed Kim.
Hypothesis:	No one touched Kim.
Answer:	NO

Text:	Sandy didn't wait to kiss Kim.
Hypothesis:	Sandy kissed Kim.
Answer:	AMBIGUOUS

parc
Palo Alto Research Center

World knowledge intrusion

Romano Prodi will meet the US President George Bush in his capacity as the president of the European commission.

George Bush is the president of the European commission.

FALSE

Romano Prodi will meet the US President George Bush in his capacity as the president of the European commission.

Romano Prodi is the president of the European commission.

TRUE

G. Karas will meet F. Rakas in his capacity as the president of the European commission.

F. Rakas is the president of the European commission.

TRUE

parc
Palo Alto Research Center

Inference under a particular construal

Romano Prodi will meet the US President George Bush in his capacity as the president of the European commission.

George Bush is the president of the European commission.

FALSE/UNKNOWN (= not entailed by the premise on the correct anaphoric resolution)

G. Karas will meet F. Rakas in his capacity as the president of the European commission.

F. Rakas is the president of the European commission.

TRUE (= entailed by the premise on one anaphoric resolution)

parc
Palo Alto Research Center

Conventional meaning vs. speaker meaning

Not a pre-theoretic but rather a theory-dependent distinction

Multiple readings

- ambiguity of meaning?
- single meaning plus pragmatic factors?

The diplomat talked to most victims

The diplomat did not talk to all victims

UNKNOWN / YES

You can have the cake or the fruit. I don't know which.

You can have the fruit

UNKNOWN / YES

parc
Palo Alto Research Center

Carving out well-behaved aspects of the problem

parc
Palo Alto Research Center

Between shallow and logic-based approaches

Parc's BRIDGE and Stanford's NatLog system are somewhere between shallow, similarity-based approaches and deep, logic-based approaches

In BRIDGE reasoning is done with a particular type of logical forms derived from parsed text

In NatLog reasoning is done with surface forms organized into phrases

No full translation to a logical formalism

No disambiguation

Special reasoning modules, no theorem proving

parc
Palo Alto Research Center

Recognizing textual entailments

Lexical categories and relations

Subsumption and monotonicity calculus

Polarity propagation

Semantic relations

Temporal Modification

parc
Palo Alto Research Center

Semantic relations

Presupposition (Factive predicates)

It is surprising that there are no WMDs in Iraq.

It is not surprising that there are no WMDs in Iraq.

Is it surprising that there are no WMDs in Iraq?

If it is surprising that there are no WMDs in Iraq, it is because we had good reasons to think otherwise.

Entailment (Implicative predicates)

It has been shown that there are no WMDs in Iraq.

It has not been shown that there are no WMDs in Iraq.

Has it been shown that there are no WMDs in Iraq?

If it has been shown that there are no WMDs in Iraq, the war has turned out to be a mistake.

Inferences in the temporal domain

In 2008 Ed visited us every month.

$\Rightarrow$ *Ed visited us in July 2008.*

Last year, in July, Ed visited us every day.

$!\Rightarrow$ *Last year Ed visited us every day.*

Ed has been living in Athens for 3 years.
Mary visited Athens in the last 2 years.

$\Rightarrow$ *Mary visited Athens while Ed lived in Athens.*

Ed has been living in Athens for 2 years.
Mary visited Athens in the last 3 years.

$!\Rightarrow$ *Mary visited Athens while Ed lived in Athens.*

parc
Palo Alto Research Center

Temporal modification under negation and quantification

Temporal modifiers affect monotonicity-based inferences

Everyone arrived in the first week of July 2000.

⇒ *Everyone arrived in July 2000.*

No one arrived in July 2000.

⇒ *No one arrived in the first week of July 2000.*

Everyone stayed throughout the concert.

⇒ *Everyone stayed throughout the first part of the concert.*

No one stayed throughout the concert.

⇒ *No one stayed throughout the first part of the concert.*

parc
Palo Alto Research Center

Quantified modifiers and monotonicity

Modifier dropping

Every boy bought a toy from Ed.

⇒ *Every boy bought a toy.*

Last year, in July, he visited us every day.

!⇒ *Last year he visited us every day.*

Modifier adding

Every boy bought a toy.

!⇒ *Every boy bought a toy from Ed.*

Last year he visited us every day.

⇒ *Last year he visited us every day in July.*

parc
Palo Alto Research Center

Textual inferences and meaning

What do theories of linguistic meaning tell us about textual inference?

"But the semantics literature, it almost never gives a full account of any inferences whatsoever. It is seriously concerned with truth conditions and figuring out how semantics should work in a general way. But it rarely goes back and figures out, for various fragments, what the overall complete stock of inferences should be." Moss (2009)

parc
Palo Alto Research Center

Textual inferences and logic

Interest in “lower order” logics and proof theory

Interest in logics and proof theory

Traditional (syllogistic) logic reasoned with regimented natural language sentences

parc
Palo Alto Research Center

Natural Logic

parc
Palo Alto Research Center

Two views on Natural Logic

Johan van Benthem's pioneering work on natural logic in the 1980's and 1990's

The proposed ingredients of a logical system to satisfy his goals would consist of

(a) Monotonicity reasoning, i.e., predicate replacement
(b) Conservativity, i.e., predicate restriction
(c) Algebraic laws for inferential features of particular lexical items

parc
Palo Alto Research Center

Two views on Natural Logic

Larry Moss' aims in recent work:

Show that significant parts of natural language inference can be carried out in decidable logical systems.

Obtain complete axiomatizations, whenever possible, because the resulting logical systems are likely to be interesting.

Make connections to fields like complexity theory, (finite) model theory, proof theory, decidable fragments of first-order logic, and algebraic logic.

Thank you

parc
Palo Alto Research Center

Some Familiar NLP Tasks

- Document Classification
- Information Retrieval and Extraction
- Topic detection and tracking
- Machine Translation
- Question Answering
- NL interfaces to databases
- Dialogue Systems