

Textual Entailment

Part 1: Introduction

Sebastian Pado

Institut für Computerlinguistik

Universität Heidelberg, Germany

Rui Wang

Language Technology

DFKI, Saarbrücken, Germany

AAAI 2013, Bellevue, WA

Thanks to Ido Dagan and Dan Roth for permission to use slides

About Us

- Sebastian Pado

Professor of Computational Linguistics

Heidelberg University,
Heidelberg, Germany



- Rui Wang

Researcher in Language Technology

German Research Center for Artificial Intelligence,
Saarbrücken, Germany



Structure of the Tutorial

- Part 1 [SP]: Introduction and Basics
- Part 2 [RW]: Classes of Strategies and Learning
 - * BREAK*
- Part 3 [SP]: Knowledge and Knowledge Acquisition
- Part 4 [SP]: Applications
- Part 5 [RW]: Multilingual, Component-based System Building

3

Part 1: Overview

- Language Processing
 - Variability in Language
- Textual Entailment
 - What is it and what is it good for?
- The Textual Entailment ecosystem
 - The “Recognizing Textual Entailment” Challenges

4

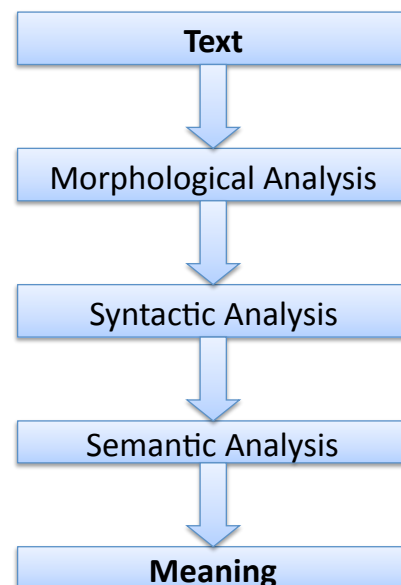
Natural Language Processing

- Text is the dominant modality to represent **knowledge** in many fields (science, industry, ...)
- Text is the dominant modality in which users **interact** with computers
- We (and our computers) need to be able to
 - **extract** knowledge from texts and
 - **draw inferences**

5

Language Processing as Analysis

- Input: Text
- Output: Formal meaning representation
 - E.g. predicate logics, description logics, modal logics, ...
- Inference: Logical calculus defined by meaning representation



6

Logical Entailment

- “A hypothesis H is entailed by a premise P ($P \models H$) iff in every model where P holds, H holds as well”
 - Relevant devices: Theorem provers, model checkers, deduction systems, ...

7

Problems of Representation

- The analysis approach formalizes language meaning **as precisely as possible: complete disambiguation**
- Language is **imprecise** and **incomplete**
 - **Ambiguity:**
*Yesterday, Peter passed by the **bank***
*I saw the man **with the telescope***
 - **Deictic expressions:**
you, he, yesterday
- Full analysis difficult and often highly ambiguous

8

Problems of Inference

- People are willing to accept “loose” inferences [Norvig 1987]:
 1. The cobbler sold a pair of sturdy boots to the alpinist.
 2. The cobbler made the sturdy boots
- People use “loose speak” [Fan & Porter 2004] to formulate search queries

9

Is All Disambiguation Necessary?

- Consider concrete instances of inference
 1. Obama addressed the general assembly yesterday
 2. The president gave a speech at the UN
- To decide whether (1) implies (2), we do NOT care whether...
 - ... “address” also has other senses
 - ... there are other referents for “the president”
 - ... what the exact date of “yesterday” is

10

Application-specific Processing


- Current dominant paradigm in language processing
 - Build task-specific models for semantic processing: Only treat **relevant** phenomena for given task
 - Semantic similarity → Distributional Methods
 - Semantic types → Named Entity Recognition
 - ...
- Robust, often accurate, models for individual tasks
- BUT huge no generalization / consolidation

Fragmentation of processing, no “theory”

11

Reimagining Semantic Processing

- The goal of processing is **not** to analyze individual texts
- Instead: determine the **relationships** that hold among texts
- Most important relationship: **Entailment**
 - Does Text A imply Text B?
(including common sense cases)



12

What Is Textual Entailment?

- TE is a **framework** for semantic language processing
 - **Not a concrete model!**
- Components:
 1. Concept of entailment (and its properties)
 2. Perspective on language processing centered around **variability**
 3. Body of research, community

13

Entailment

- A *directional* relation between two text fragments:
Text (t) and Hypothesis (h):

t **entails** h ($t \Rightarrow h$) if humans reading t will infer that h is most likely true [Dagan & Glickman 2004]

14

Textual vs. Logical Entailment

- Logical Entailment:
 - Define formal representation language
 - Define translation into formal language
 - **Entailment is what the representations say it is**
- Textual Entailment:
 - Collect entailment judgments for text pairs
 - Develop processing methods that can reproduce these judgments
 - **Entailment is what the speakers say it is**

15

Textual vs. Logical Entailment

“**Loose**” entailment: Textual but not logical

T: The technological triumph known as GPS was incubated in the mind of Ivan Getting.

H: Ivan Getting invented the GPS.

“**Uninformative**” entailment :Logical but not textual

T: The technological triumph known as GPS was incubated in the mind of Ivan Getting.

H: Two plus two equals four.

16

Entailment and Variability

- Variability is a central fact of language
 - TE can be seen as the task of distinguishing **meaning-preserving** from **meaning-changing** variability

The **Global Positioning System** **was incubated in the mind of** an American physicist, **Ivan Getting**. \Rightarrow **Ivan Getting invented GPS.**

Abbreviations, Paraphrases, Change of Voice, Apposition, ...

17

Variability and Inference

- Variability is important in, but not all of, inference:
 - Inferences about language variability
 - I **bought** a watch => I **purchased** a watch
 - Inferences about the extra-linguistic world
 - it **rained** yesterday => it **was wet** yesterday
- Most (Text, Hypothesis) pairs involve both
 - No definite boundary between the two
- Crucial role of both kinds of knowledge (**cf. Part 3**)


18

Recognizing Textual Entailment

- “Common ground” for processing approaches
 - Contrast to analysis-centered approach
 - No abstract gold standard
- Allows direct comparison of different processing approaches (**cf. Part 2**)
 - “Depth of analysis” up to each approach
- Mid-term goal: Identification and combination of best strategies from various approaches (**cf. Part 5**)

19

“Easy-first processing”



- Perform as many inferences over natural language representations as possible
- Resort to formal meaning representation when necessary

20

Why Work With Textual Entailment?

- Conceptual benefits:
 - A concept of “common sense” inference
 - Alternatively, framework to address language variability
 - Novel perspective on the needs of language processing
- Practical benefits:
 - An attractive “meta framework” for language processing
 - A unified perspective on many research questions at the boundary of language processing, machine learning, and knowledge representation

21

Textual Inference in Applications

QA:

Question: What **affects** blood pressure?

“**Salt** **causes an increase** in blood pressure”

IR:

Query: **symptoms** of IBS

“**IBS is characterized by vomiting**”

22

Story Comprehension

(ENGLAND, June, 1989) - Christopher Robin is alive and well. He lives in England. He is the same person that you read about in the book Winnie the Pooh. As a boy, Chris lived in a pretty home called Cotchfield Farm. When Chris was three years old, his father wrote a poem about him. [...]

1. Christopher Robin was born in England.
2. Winnie the Pooh is a title of a book.
3. Christopher Robin's dad was a magician

cf. also Part 4

23

Practical Role of Textual Entailment

- Young task: Introduced about 10 years ago
- A prominent concept in semantic processing
 - 20000 Google Scholar hits for “Textual Entailment”
- Important role: The “Recognizing Textual Entailment” Challenges (PASCAL/NIST)
 - Yearly preparation of new datasets
 - Created utilizing (or simulating) reductions from real systems' output
 - Shared task: Practical and conceptual advances

24

RTE Data

	TEXT	HYPOTHESIS	TASK	ENTAILMENT
1	<i>Regan attended a ceremony in Washington to commemorate the landings in Normandy.</i>	<i>Washington is located in Normandy.</i>	IE	False
2	<i>Google files for its long awaited IPO.</i>	<i>Google goes public.</i>	IR	True
3	<i>...: a shootout at the Guadalajara airport in May, 1993, that killed Cardinal Juan Jesus Posadas Ocampo and six others.</i>	<i>Cardinal Juan Jesus Posadas Ocampo died in 1993.</i>	QA	True

25

Developments of the Task

- RTE 1, 2: Single-sentence T-H pairs
- RTE 3+: Longer texts
- RTE 4: Contradiction
 - Generalization to more relations
- RTE 5: Search Task (single H, multiple Ts)
- RTE 6+: Application-specific datasets
 - RTE 8 (2013): Student Response Analysis

26

Development of Methods



27

Remainder of this Tutorial

- Part 2 [RW]: Classes of Strategies and Learning
 - Which methods can be used to decide entailment?
- Part 3 [SP]: Knowledge and Knowledge Acquisition
 - What kinds of knowledge are necessary? Where can we find them or how can we learn them?
- Part 4 [SP]: Applications
 - How can language processing applications use entailment?
- Part 5 [RW]: Multilingual, Component-based System Building
 - How can we develop sustainable entailment systems?

28

Reference List

- I. Dagan and O. Glickman (2004). Probabilistic textual entailment: Generic applied modeling of language variability. Proceedings of the PASCAL workshop on Learning Methods for Text Understanding and Mining.
- J. Fan and B. Porter (2004). Interpreting Loosely Encoded Questions. Proceedings of AAI, 399-405.
- Heilman, M. and N. Smith (2010). Tree edit models for recognizing textual entailments, paraphrases, and answers to questions. Proceedings of NAACL, 1011–1019.
- P. Norvig (1987). Inference in text understanding. Proceedings of AAI, 561–565.

29

Reference List

- Shnarch, E., J. Goldberger, and I. Dagan (2011). A probabilistic modeling framework for lexical entailment. Proceedings of ACL, 558–563.
- Stern, A. and I. Dagan (2011). A confidence model for syntactically-motivated entailment proofs. Proceedings of RANLP, 455–462.
- Stern, A. , R. Stern, I. Dagan, and A. Felner (2012). Efficient search for transformation-based inference. In Proceedings of ACL, 283-291.

30