

Natural Language Processing for Knowledge Representation and Reasoning

Michaël Thomazo

April 14th, 2014
Dresden

A few words about me and the course

Me:

- ▶ member of the Computational Logic Group
- ▶ office 2036 if needed (door open = you are welcome, door closed = do not enter)
- ▶ mail: michael.thomazo@tu-dresden.de

The course:

- ▶ DS3-DS4
- ▶ INF-E005
- ▶ distinction between course and tutorial will not respect the schedule on the Webpage. In particular today.

Origin of the course

- ▶ my research focuses on so-called “ontological query answering”
- ▶ an often-cited application area is the Semantic Web, through the use of standards

Problem: real world applications and real-world data

Benchmarks, and in particular coming from real-world, are missing.
However:

- ▶ there is a lot of data on the Web, in textual form
- ▶ there is knowledge expressed in textual form
- ▶ how much of “my work” could be applied there?”

My leitmotiv during this course

- ▶ look at what are the interactions between NLP and logic
- ▶ what has been done, what is missing, what succeeded, what failed, what is currently tried, what should be tried?

What NLP could bring to KRR

KRR usually uses highly normalized data

- ▶ relational database
- ▶ RDF triple stores
- ▶ other technologies

All these formats are not directly created by humans, as a sentence would → no direct communication.

Can KRR bring anything to NLP?

- ▶ in the realization of “NLP tasks”?
- ▶ in the more general setting of artificial intelligence? Towards the creation of “intelligent systems?”

Expectations/Goal of the course/Evaluation (1)

Expectations:

- ▶ questions are more than welcome
- ▶ given the schedule, eating during the course is no problem...
- ▶ ... but nothing too smelly, please!
- ▶ or we can arrange a slightly different schedule?
- ▶ you are expected to be at ease with technical parts of the course

Goals:

- ▶ to motivate you to consider approaches that are at the frontier between the two fields
- ▶ to give you a KRR-based view on (a small part of) NLP

Expectations/Goal of the course/Evaluation (1)

Evaluation:

- ▶ high-level knowledge of the material of the course
- ▶ technical ease

Starting point for this first lecture

- ▶ Levesque's invited talk at IJCAI'13

What is an intelligent system?

I mentioned “intelligent systems”. What is it?

What is an intelligent system?

I mentioned “intelligent systems”. What is it?

the Turing test

Is this a good test? (Hint: there are some objections)

What is an intelligent system?

I mentioned “intelligent systems”. What is it?

the Turing test

Is this a good test? (Hint: there are some objections)

Striking argument: the system must be lying at some point in some configurations.

Particularities and other objections

(Levesque, KR'12)

The Turing test is heavily based on language (“understanding” and generation). More specifically on conversations, which:

- ▶ facilitate deception and trickery (ELIZA 66, PARRY 72);
- ▶ are hard to evaluate.

Do you know any other system that aims at distinguishing between human and computers?

Some tools to distinguish humans from computers

- ▶ CAPTCHA
- ▶ Recognizing Textual Entailment
- ▶ Winograd schemes

CAPTCHA

- ▶ recognizing letters and or numbers
- ▶ efficient (to date) at distinguishing human from computers
- ▶ would we really describe a system solving (consistently) CAPTCHAs as intelligent?
- ▶ but arguably the most useful test to date!

Recognizing Textual Entailment

- A Time Warner is the world's largest media and internet company.
- B Time Warner is the world's largest company.

Recognizing Textual Entailment

- A Time Warner is the world's largest media and internet company.
- B Time Warner is the world's largest company.

- A Norway's most famous painting, The Scream by Edvard Munch, was recovered Saturday.
- B Edvard Munch painted The Scream.

Recognizing Textual Entailment

- A Time Warner is the world's largest media and internet company.
- B Time Warner is the world's largest company.

- A Norway's most famous painting, The Scream by Edvard Munch, was recovered Saturday.
- B Edvard Munch painted The Scream.
- B' The recovered painting was worth more than \$1000.

Winograd schemes

The trophy would not fit in the brown suitcase because it was too big. What was too big?

- ▶ Answer 0: the trophy
- ▶ Answer 1: the suitcase

Winograd schemes

The trophy would not fit in the brown suitcase because it was too big. What was too big?

- ▶ Answer 0: the trophy
- ▶ Answer 1: the suitcase

Joan made sure to thank Susan for all the help she had given. Who had given the help?

- ▶ Answer 0: Joan
- ▶ Answer 1: Susan

Features of Winograd schemes

1. two parties are mentioned in a sentence by noun phrases.
2. A pronoun or possessive adjective is used in the sentence in reference to one of the parties, but is also of the right sort for the second party (he/him,his/ for males,she/her/her for females, ...)
3. the question involves determining the referent of the pronoun or possessive adjective. Answer 0 is the first party mentioned, Answer 1 the second
4. there is a word (the *special* word) that appear in the sentence and possibly in the question. When it is replaced by another word, everything still makes perfect sense, but the answer changes

Winograd schemes

The city councilmen refused the demonstrators a permit because they feared violence. Who feared violence?

Winograd schemes

The city councilmen refused the demonstrators a permit because they feared violence. Who feared violence?

Answer given by a human: **the councilmen.**

Winograd schemes

The city councilmen refused the demonstrators a permit because they feared violence. Who feared violence?

Answer given by a human: **the councilmen.**

The city councilmen refused the demonstrators a permit because they advocated violence. Who advocated violence?

Winograd schemes

The city councilmen refused the demonstrators a permit because they feared violence. Who feared violence?

Answer given by a human: **the councilmen.**

The city councilmen refused the demonstrators a permit because they advocated violence. Who advocated violence?

Answer given by a human: **the demonstrators.**

Finding *alternate* words

The trophy would not fit in the brown suitcase because it was too big. What was too big?

- ▶ Answer 0: the trophy
- ▶ Answer 1: the suitcase

Finding *alternate* words

The trophy would not fit in the brown suitcase because it was too big. What was too big?

- ▶ Answer 0: the trophy
- ▶ Answer 1: the suitcase

Joan made sure to thank Susan for all the help she had given. Who had given the help?

- ▶ Answer 0: Joan
- ▶ Answer 1: Susan

The qualities of a good Winograd scheme

- ▶ easily disambiguated by a human (native English-speaking) reader;
- ▶ not solvable by simple techniques;
- ▶ Google-proof;

Examples of Google-proof questions

“Can an alligator run the hundred-meter hurdles?”

Examples of Google-proof questions

“Can an alligator run the hundred-meter hurdles?” → is it still google proof?

Examples of Google-proof questions

“Can an alligator run the hundred-meter hurdles?” → is it still google proof?

“Can an alligator dance salsa?”

Examples of Google-proof questions

“Can an alligator run the hundred-meter hurdles?” → is it still google proof?

“Can an alligator dance salsa?”

“Can an elephant skydive?”

Why do we care?

Do we really want to care about these questions?

Why do we care?

Do we really want to care about these questions?

France had in 2008 10500 laws and 127 000 decree (not counting European laws as far as I got).

Why do we care?

Do we really want to care about these questions?

France had in 2008 10500 laws and 127 000 decree (not counting European laws as far as I got).

“Ignorance of the law excuses no one”

Importance of *explaining why something holds or not*.

Other criteria

- ▶ Easily disambiguated by a native English-speaking human: this notion may also vary, though more slowly
- ▶ not solvable by simple techniques: this notion may also vary (is chess “solvable” by simple techniques? Go? Arimaa?)

An example of easily solvable Winograd scheme

“The women stopped taking the pills because they were pregnant/carcinogenic. Which individuals were pregnant/carcinogenic?”

- ▶ the women
- ▶ the pills

Use of “selectional restriction”.

- ▶ I love to drink coffee.
- ▶ I love to drink cars.

An example of scheme that is “too hard”

Frank was jealous/happy when Bill said that he was the winner of the competition. Who was the winner?

- ▶ Answer 0: Frank
- ▶ Answer 1: Bill

What can be said if a system passes the test?

Levesque claims:

“with a very high probability, anything that answers correctly is engaging in behavior that we would say shows thinking in people.”

A link with KRR

What is obvious to a human depends a lot on what he knows.
The man could not lift his son because he was so weak/heavy.
Who was weak/heavy?

- ▶ Answer 0: the man
- ▶ Answer 1: his son

A link with KRR

What is obvious to a human depends a lot on what he knows.
The man could not lift his son because he was so weak/heavy.
Who was weak/heavy?

- ▶ Answer 0: the man
- ▶ Answer 1: his son

The large ball crashed right through the table because it was made of steel/styrofoam. What was made of steel/styrofoam?"

- ▶ the ball
- ▶ the table

A link with KRR

What is obvious to a human depends a lot on what he knows.
The man could not lift his son because he was so weak/heavy.
Who was weak/heavy?

- ▶ Answer 0: the man
- ▶ Answer 1: his son

The large ball crashed right through the table because it was made of steel/styrofoam. What was made of steel/styrofoam?"

- ▶ the ball
- ▶ the table

.
Necessity to have some kind of knowledge about what styrofoam (and steel) are → knowledge bases and reasoning required.

Does the following fulfill these requirements?

The racecar zoomed by the school bus because it was going so fast/slow. What was going so fast/slow?

- ▶ Answer 0: the racecar
- ▶ Answer 1: the school bus

Does the following fulfill these requirements?

The racecar zoomed by the school bus because it was going so fast/slow. What was going so fast/slow?

- ▶ Answer 0: the racecar
- ▶ Answer 1: the school bus

Levesque says that no, because can be solved by “simple techniques”

Does the following fulfill these requirements?

The racecar zoomed by the school bus because it was going so fast/slow. What was going so fast/slow?

- ▶ Answer 0: the racecar
- ▶ Answer 1: the school bus

Levesque says that no, because can be solved by “simple techniques” Can you improve this scheme?

“Anna did a lot better/worse than her good friend Lucy on the test because she had studied so hard. Who studied hard?”

Your turn...

Design some good Winograd schemes.

What we have seen

- ▶ popular “specifications” of “intelligent systems” implies the use of human language
- ▶ several proposals have been made
- ▶ they are designed such that a purely statistical approach can not work
- ▶ KRR may bring something for this

On the other hand NLP may allow KRR researches to have a wider range of applications.

Starting point: where to start from?

What are Description Logics?

→ The Handbook of Description Logics

Starting point: where to start from?

What are Description Logics?

→ The Handbook of Description Logics

What is NLP?

→ The Handbook of Computational Linguistics and Natural Language Processing (HCLNLP)

A terminological remark

What is the difference between Computational Linguistics and Natural Language Processing?

A terminological remark

What is the difference between Computational Linguistics and Natural Language Processing?

First sentence of the introduction of [HCLNLP]:

“The field of computational linguistics (CL), together with its engineering domain of natural language processing (NLP), has exploded in recent years.”

It seems that both expression are often used interchangeably.

Logic...

...does not seem to have a prominent part in current approaches.

- ▶ few occurrences to first-order logic (mainly in the Complexity chapter);
- ▶ none to fuzzy logic;
- ▶ none to probabilist logic;
- ▶ none to modal logic.

Topic of the course

Explore what have been the proposed approaches that use both NLP and logic-based approaches.

A bit more detail

- ▶ POST: a classical NLP problem, a classical approach
- ▶ some reasons why “logic is not well suited to NLP”
- ▶ Markov logic and some applications to NLP problems
- ▶ Natural Logic
- ▶ from natural language to structured data
- ▶ possibly some other topics