



CIS-700
Spring 2020


Reasoning for Natural Language Understanding

Dan Roth

Computer and Information Science
University of Pennsylvania

Introduction Part II

This class



- Understand early and current work on Reasoning
 - (Learn to) read critically, present, and discuss papers
- Understand some of the difficulties in NLU from the perspective of reasoning
 - Conceptual and technical
- Try some new ideas


- How:
 - Presenting/discussing papers
 - Probably: 2 presentations each; 4 discussants
 - Writing a few critical reviews
 - “Small” individual project (reproducing);
 - Large project (pairs)
 - Tentative details are on the web site.

- Machine Learning
 - 519/419
 - 520
 - Other?
- NLP
 - Yoav Goldberg’s book
 - Jurafsky and Martin
 - Jacob Eisenstein

- Attendance is mandatory
- Participation is mandatory

- Time of class?
- **Expectations?**


What's Important in order to make progress in NLU



- How to make progress towards natural language understanding
 - Learning and Reasoning; knowledge
- Dispel with [some] of the currently hot trends
 - If we want to reach the moon...
- What is Reasoning?
- Reasoning for Natural Language Understanding
- Today: Examples & Discussion




Questions




- What is Reasoning?
 - Do we do reasoning? Yes, we do.

June 21, 240 B.C.: The Earth Is Round, and It's This Big



- Eratosthenes knew that at noon on the day of the summer solstice, the sun was observed to be directly overhead at Syene (modern-day Aswan):
 - You could see it from the bottom of a deep well, and a sundial cast no shadow.
 - Yet, to the north at Alexandria, a sundial cast a shadow even at the solstice midday, because the sun was not directly overhead there.
 - Therefore, the Earth must be round
- Assuming the sun to be sufficiently far away to be casting parallel rays at Syene and Alexandria, it would be possible to figure out the Earth's circumference
- $\theta = 7.2^\circ$; distance = 800km (5,000 stades) → **circumference = $360/7.2 * 800 = 40,000\text{km}$**
- (He made a few mistakes: Alexandria is not exactly north of Aswan; Aswan is not exactly in the tropic of cancer; the distance is not exactly 800km; earth is not a perfect sphere, but current calculations put it at 40,030km.)




- <https://www.youtube.com/watch?v=Mw30CgaXiQw>
- See also “The Enigma of Reason”

- What is Reasoning?
 - Do we do reasoning? Yes, we do.
 - How can we formulate it?

- Reasoning requires knowledge
 - How do we represent it?
 - What types of knowledge
 - What types of representations?

- We want to think about these in the context of natural language understanding
 - In what ways does it change the game?
 - Is Reasoning for/in NLU different than “Reasoning”?



AN EXAMPLE FOR NATURAL LANGUAGE UNDERSTANDING AND THE AI PROBLEMS IT RAISES

John McCarthy

Computer Science Department

Stanford University


Stanford, CA 94305

`jmc@cs.stanford.edu`

`http://www-formal.stanford.edu/jmc/`

1976


A New York Times Story



“A 61-year old furniture salesman was pushed down the shaft of a freight elevator yesterday in his downtown Brooklyn store by two robbers while a third attempted to crush him with the elevator car because they were dissatisfied with the \$1,200 they had forced him to give them.

The buffer springs at the bottom of the shaft prevented the car from crushing the salesman, John J. Hug, after he was pushed from the first floor to the basement. The car stopped about 12 inches above him as he flattened himself at the bottom of the pit.

A New York Times Story (Cont.)




Mr. Hug was pinned in the shaft for about half an hour until his cries attracted the attention of a porter. The store at 340 Livingston Street is part of the Seaman's Quality Furniture chain.

Mr. Hug was removed by members of the Police Emergency Squad and taken to Long Island College Hospital. He was badly shaken, but after being treated for scrapes of his left arm and for a spinal injury was released and went home. He lives at 62-01 69th Lane, Maspeth, Queens.


He has worked for seven years at the store, on the corner of Nevins Street, and this was the fourth time he had been held up in the store. The last time was about one year ago, when his right arm was slashed by a knife-wielding robber.”

New York Times Story: Questions



- An intelligent person or program should be able to answer the following questions based on the information in the story:
 - The article proceeds with 22 questions:
 1. Who was in the store when the events began?
 - Probably Mr. Hug alone, although the robbers might have been waiting for him, but if so, this would have been stated.
 2. What did the porter say to the robbers?
 - Nothing, because the robbers left before he came.
 20. Why did Mr. Hug yell from the bottom of the elevator shaft?
 - So as to attract the attention of someone who would rescue him.
- “The above list of questions is rather random. I doubt it covers all facets of understanding the story.”


McCarthy's Challenges



The QA module is not being trained
Once the program knows English, and has the relevant background knowledge, it should answer the questions

- A formalism capable of expressing the assertion of the sentences free from dependence on the grammar of the English language. (“Artificial Natural Language”, ANL)
 - Semantic Parser
- An “understander” that constructs the “facts” from the text.
 - Information Extraction: Entities, Relations, Temporal, Quantities,...
- Expression of the “general information” about the world that could allow getting the answers to the questions from the “facts” and the “general information”
 - Background Knowledge
- A “problem solver” that could answer the above questions on the basis of the “facts”.
 - Question Answering Engine


- What can we learn from this example?
 - Difficulties of NLU
 - Importance of reasoning
 - Part of Reasoning here seems to be “providing the reasons”, not only the “answers”
 - Decoupling learning from reasoning
 - McCarthy thinks that there is a need for some level of abstraction – an abstract representation of the text and the relevant knowledge so that a generic module can work on it and “do the reasoning”.
- Is this important/Essential?




- You will spend the next 10 minutes on:
 - Suggest a reasoning problem.
 - Describe it
 - Suggest a way to formulate it so that you can write a program that solves it

 - Think about knowledge needed
 - Describe the type of knowledge you think is needed and why/when
 - Suggests ways to formulate it: represent it and use it


More Examples




(1) I Have Some Questions About....



- Scenario:
 - You are reading the book, but left it for a couple of weeks.
 - You need a refresher: some of the events, entities, the current relationships between David and James Steerforth.
- Conversing about it is challenging:
 - Many chapters, multiple voices, long periods of time,...
 - The novel features the character [David Copperfield](#), his journey of change and growth from infancy to maturity, as many people enter and leave his life and he passes through the stages of his development. (**Fiction, and you know it**)
 - London and England in the 19-th century; socio-economic state, child exploitation; schools, prisons, emigration to Australia (**true historical facts**)
- [What computational tasks should we think about?](#)




(2) I Want to Talk about this News Story




Who was in the store when the events began?

- The story doesn't say.
- Probably Mr. Hug alone, although the robbers might have been waiting for him, but if so, this would have been stated.



■ Why was he crying?

- Maybe he was scared.
- Maybe he was injured.
- Maybe he called for help.



■ Understanding the **story** and **conversing** about it require **Situated Reasoning**: Model-based Reasoning

“A 61-year old furniture salesman was pushed down the shaft of a freight elevator yesterday in his downtown Brooklyn store by two robbers while a third attempted to crush him with the elevator car because they were dissatisfied with the \$1,200 they had forced him to give them.


The buffer springs at the bottom of the shaft prevented the car from crushing the salesman, John J. Hug, after he was pushed from the first floor to the basement. The car stopped about 12 inches above him as he flattened himself at the bottom of the pit.

Mr. Hug was pinned in the shaft for about half an hour until his cries attracted the attention of a porter. The store at 340 Livingston Street is part of the Seaman's Quality Furniture chain.

Mr. Hug was removed by members of the Police Emergency Squad and taken to Long Island College Hospital. He was badly shaken, but after being treated for scrapes of his left arm and for a spinal injury was released and went home. He lives at 62-01 69th Lane, Maspeth, Queens.

He has worked for seven years at the store, on the corner of Nevins Street, and this was the fourth time he had been held up in the store. The last time was about one year ago, when his right arm was slashed by a knife-wielding robber.”

(3) Some questions to my Sports' Assistant



Modified version of a question for AI2's DROP dataset

Coming off a road win over the Cowboys, the Redskins traveled to Lincoln Financial Field for a Week 5 NFC East duel with the Philadelphia Eagles. In the first quarter, the Redskins trailed early as **RB Brian Westbrook scored on a 9-yard TD run** and **the Eagles DeSean Jackson returned a punt 68 yards for a touchdown.**

Washington still trailed at half time **14:9, with field goals from Shaun Suisham.** In the third quarter, the Redskins took the lead on a trick play as WR Antwaan Randle El threw an 18-yard **TD pass to TE Chris Cooley.** In the fourth quarter, the Redskins increased their lead when **Clinton Portis scored on a 4-yard TD run.** The Eagles managed one more score in the final quarter for a **final score of 17:23.**



Football Scoring Rule Book:


- TD could be 6, 7, or 8 points.
 - Kick....
- Field goal is worth 3 points
-

General rules that are to be inst

aned from the recap.
ograms on the text
g team))
n is necessary.

What computational tasks
should we think about?

(4) Let's Talk about Dinner



But, there is another big difference between these two scenarios.

- The first applies only today.
- The second is a general rule that I'd like the Assistant to remember.

- → Let's talk about dinner.

- A:** Where do you want to go?

- → I had a big lunch

- [This is not an answer; can the Assistant figure it out?]
- It's probably just a hint that we should go for a light dinner

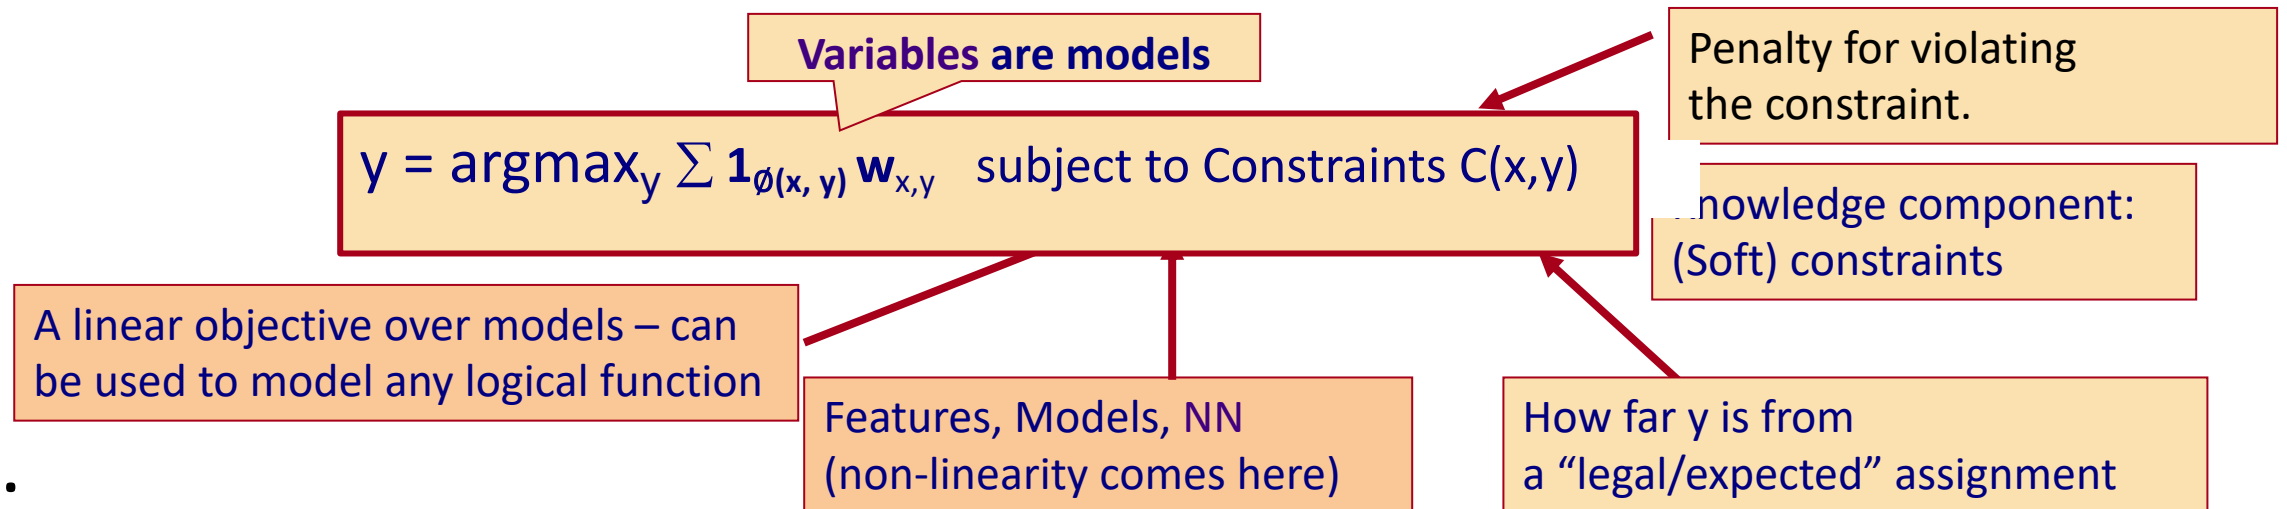
- → I don't like crowds

- [This is not an answer; can the Assistant figure it out?]
- Perhaps a preference for small venues?

- → I had a lot of pizza the last few weeks

- [Again; not a direct; how do we understand it? How do we represent it?]

- These **reasoning problems** are best thought of as **abductive reasoning** problems:
 - Reasoning for the best explanation, given the observations and the knowledge available.
- Formulating it as a constraint optimization problems provides a way to incorporate learned models (as variables) along with knowledge (Boolean Fns; Logical; Constraints)



- **Learning** formulations such as **Posterior Regularization** [Ganchev et al. 10] and **CoDL** [Chang et al. 07, 12] provide ways to learn such objectives [but are not sufficient]
 - Users have **expectation** from the output behavior, and they communicate to clarify it. It is imposed on the posterior

- What is the role of **formal theories** of reasoning and representation?
 - They assume that we can map text and world knowledge to a “symbolic” representation; given that, the problem is solved (so people think).
 - Note that this is true even when people use neural networks for all/part of the computation
 - If this is wrong – where is it wrong?
 - Is it the infeasibility of the mapping?
 - Is it that our formal theories of reasoning are missing something?
- Think also about the statements I expressed last time
 - Reasoning is about giving reasons
- What are the implications of this (whether you agree with it or not) on the need to have “symbols”?

- Eventually, we may want to think about a neural implementation of Reasoning
 - Is it necessary? Is this where the challenge is?
 - Note that neuro-symbolic AI goes back many years
 - Is it ok to first think about formal theories, and then encode them in neural architectures?
 - Or, there is an advantage to directly thinking about neural representations.
 - This latter view means that there is **no other representation** of what neural architectures are doing.
 - E.g., is an embedding of a sentence different than other representations of it in some principled way?
 - Or is it just a more compact representation?
- Next, we will describe things from the perspective of Learning to Reason.
 - The presentation will mostly focus on the logical approach, but similar ideas can be extended to other formalisms.
 - This will hopefully serve as introduction to formalisms,