


About Me

- Graduated from Illinois 2012
- Postdoc at Maryland 2013-14
- Started this year at Purdue University!
- Research Interests:
 - Applying Machine Learning (**ML**) methods to Natural Language Processing (**NLP**) problems
 - *Very Fruitful Interaction between ML and NLP!*
 - **Maybe it's better to start with where we were supposed to be 13 years ago!**

2001




Maybe it's better to start with where we were supposed to be 13 years ago, **according to 1960's sci-fi**

Dave : *Open the pod bay doors, HAL.*



HAL: *I'm sorry Dave. I'm afraid I can't do that.*



Clearly unrealistic! Hal understood:

- (1) the meaning of the words
- (2) grounded their meaning in a physical environment
- (3) understood the situation and intent of the speaker

Now consider the “perfect” personal assistant on your smartphone stopping you from sending that ANGRY email to your boss.

Realistic?

NLP research is about pushing the limits of realistic applications

NLP in Practice

- Personal assistants that interact in natural language
 - *Simple voice activation turned into complex language analysis*
 - Several recent high profile applications
 - *Interaction is often humorous and enjoyable*




NLP in Practice

- **Information Extraction**

- Parse unstructured text into structured information
- Now a standard part of most email services



Hi Dan,
I just wanted to let you know that we scheduled the meeting for Monday 9:30, at the Lawson Commons. It will take two hours.



Event: Meeting
Date: Monday, Sep 15
Start: 9:30am
End: 11:30
Location: Lawson Commons

NLP in Practice

- **Sentiment Analysis**
 - Meaningful interpretation of product reviews
 - Identify the product aspects users care about
 - *Deception detection*

I just bought company-A newest laptop. The display is awesome, the speakers are not that great. I'm happy with the performance, but I think they charge too much for it!



Display: Positive
Speakers: Negative
Performance: Positive
Price: Negative

NLP can be Very Challenging!

Kiel first encountered Moore's James Bond in 1977's "The Spy Who Loved Me," where his silent hitman Jaws repeatedly menaced Bond with his sharp metal teeth. Although repeatedly thwarted by the British spy, Jaws proved resilient and even sort of likable: Near the movie's end he survived a brush with a killer shark by biting the creature.

Jaws was such a popular character that the producers of the Bond series brought him back two years later for "Moonraker," which was set partly in space. He and Bond battled each other in an opening skydiving sequence and in a memorable scene atop an aerial tram in Rio de Janeiro. Later, Jaws switched allegiances to Bond upon learning that his employer, the villain Drax, planned to exterminate him.

Opinion: Why Jaws was best 'Bond' villain ever

In later years, Kiel turned his hand to writing and producing as well as acting, including in the 1991 movie "The Giant of Thunder Mountain," according to the IMDB.

He also had a small role in "Happy Gilmore," the 1996 Adam Sandler golf comedy.

"Richard Kiel was one of the nicest, funniest guys I've ever met. I'll never forget hanging out with him & how good he was to everyone," Sandler tweeted Thursday.

Kiel's family posted this message on Facebook:

"It is with very heavy hearts that we announce that Richard has passed away, just three days shy of his 75th birthday. Richard had an amazing joy for life and managed to live every single day to the fullest. Though most people knew of him through his screen persona, those who were close to him knew what a kind and generous soul he was.

"His family was the most important thing in his life and we are happy that his last days were spent surrounded by family and close friends. Though his passing was somewhat unexpected, his health had been declining in recent years. It is nice to think that he can, once again, stand tall over us all."

Summarization: Capture the essence of a long text



Richard Kiel, actor who played Bond villain "Jaws" dies at 74

Summarization is still an open problem

Representative of difficult problems in NLP:

- **Word-level decisions VS high level tasks that require deeper understanding**

More than Words

- Natural language is inherently ambiguous

Teacher strikes idle kids

Hospital sued by 7 foot doctors

Local High School **Dropouts** Cut in Half

- **Machine learning** an effective tool for resolving ambiguities
 - Build statistical prediction models based on annotated data
- **Not a magical solution**: can be difficult to learn some problems
 - For example, twitter posts do not look like WSJ or NYT articles
- **Higher level NL understanding tasks rely on complex analysis that is difficult to annotate!**

The Age of Digital Interaction

We have access to unprecedented amount of NL resources!



Human-Computer Interaction

Natural Language Interfaces

Personal Assistants



Computer-Mediated Interaction

*Large-scale human communication
mediated by computers*

Reshapes the way humans interact!

How can this massive amount of information be used?

Key observation:

NL communication is followed by interactions!

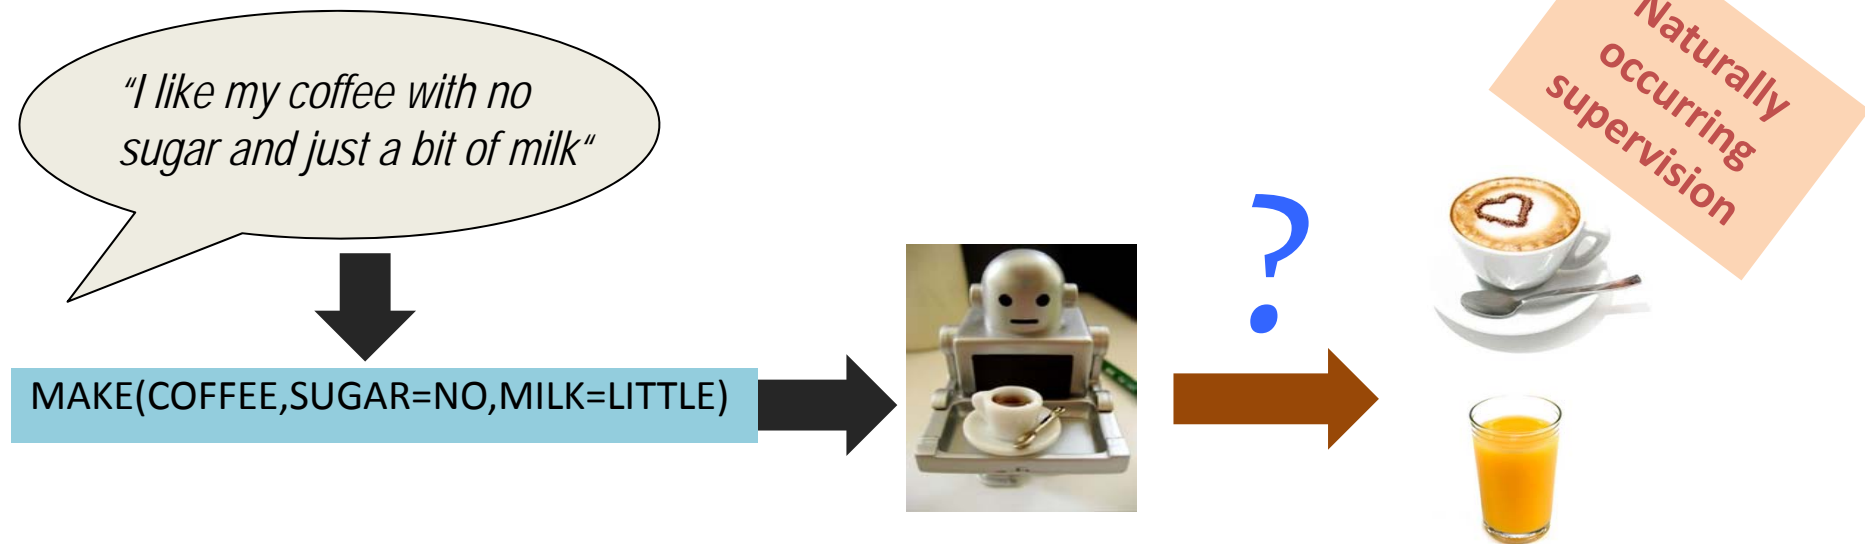
Learning from Interactions

Semantic Interpretation: moving from NL to a formal language

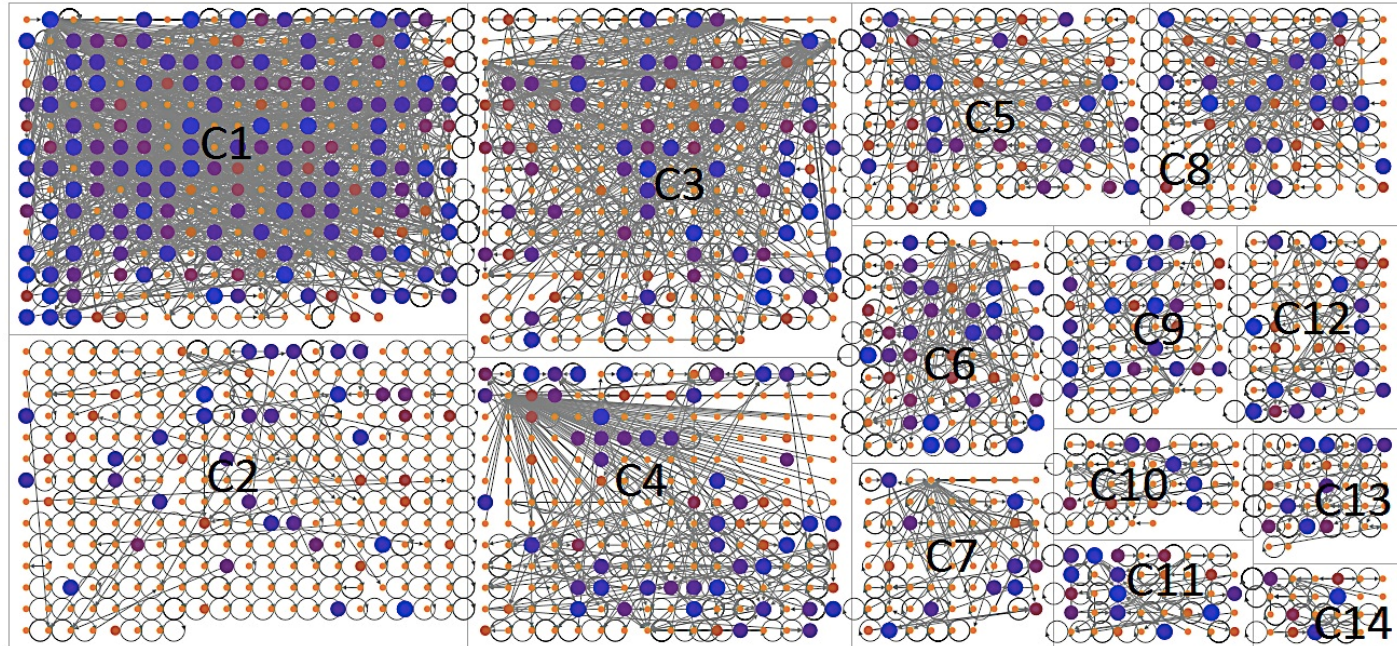
→ This process is commonly known as **Semantic Parsing**

Typically this mapping is learned from **annotated data**
pairs of (**Sentence, Logical Form**)

→ Requires **a lot** of training data : “hot java” = “coffee”



Naturally Occurring Reactions



Can we rely on **social interactions** to guide language understanding?

Summary: Research Challenges

- **Pushing the limits of NLP**
 - **Understanding human interaction:** *data is plentiful but not carefully annotated*
- **ML perspective: learning with limited supervision**
 - How to learn effectively from reactions?
 - How to use world knowledge? *Social knowledge?*

Dan Goldwasser
dgoldwas@purdue.edu

