

# Lecture 2: Natural Language Challenges



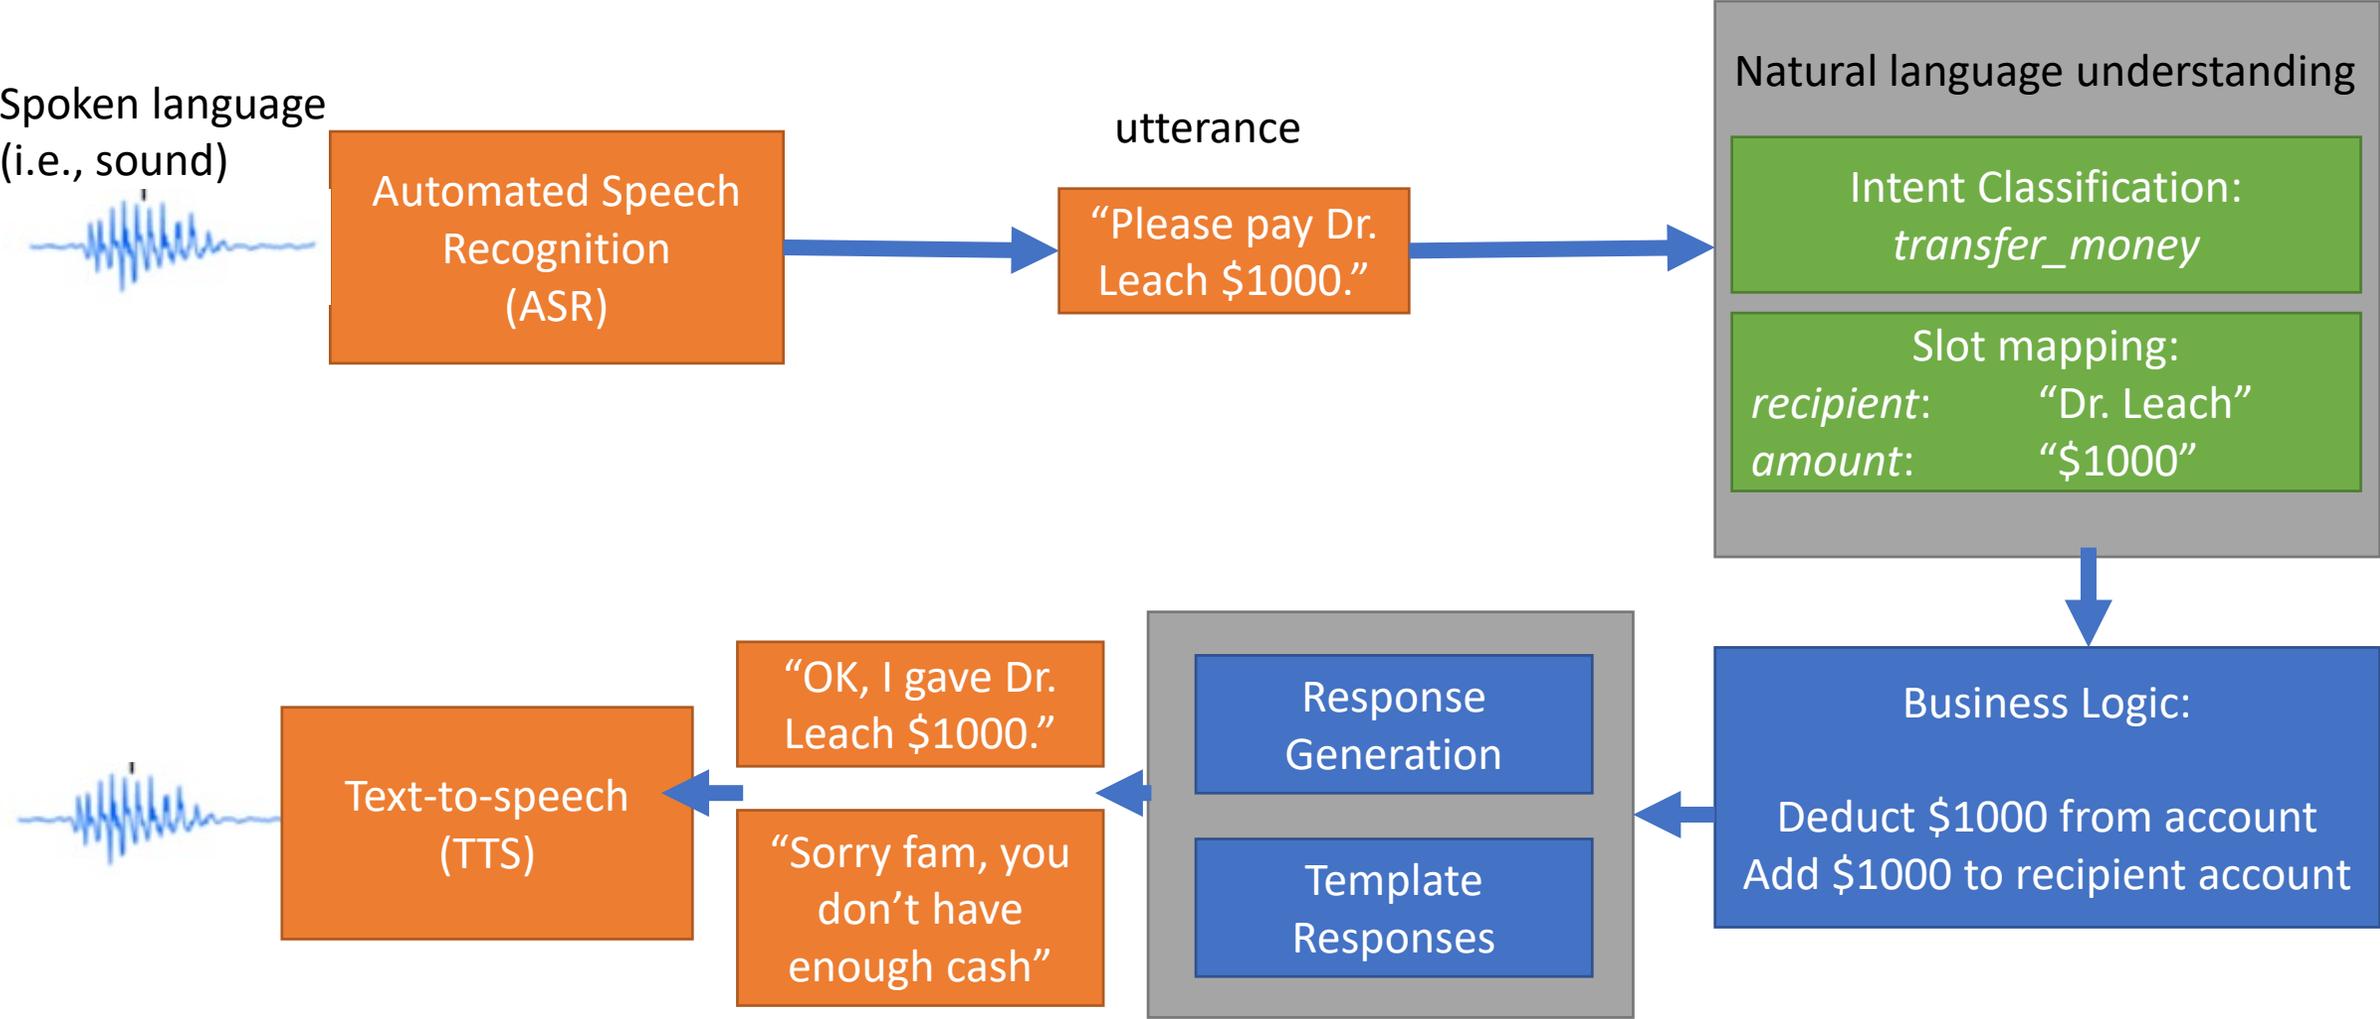
# Review: Course Logistics

- <https://dijkstra.eecs.umich.edu/eecs498>

# Review: Course Introduction

- **Conversational AI** seeks to build software that can **converse with human users** to help accomplish tasks or acquire information
  - We take an **utterance** as input, **classify** its **intent**, **extract** semantic **slot information**, take action with **business logic**, then **generate a response**
- **Project-based** course in which your **team** will **pitch** an idea for a virtual assistant, then implement it during the semester over three **sprints**, culminating in a **final demonstration**
- You will have access to the **ClinC Conversational AI Platform** for building your project

# Review: Conversational Flow



# Review: Conversational Flow

- 2 or more interlocutors exchanging information
  - What assumptions must we make to create virtual assistants?
- In Conversational AI, a human user speaks or writes an **utterance** that gets processed by a **natural language understanding** engine
  - The **utterance** implies an **intent** (what the user *intends* to do) and contains **slot values** (useful semantic information)
- In our Gaming demo from last class, what intents and slots do you think are relevant?

# One-Slide Summary: NLP Challenges

- **Natural Language Processing** is the field of CS, AI, and CL concerned with interactions in *natural* languages
- **NLP is hard** because human language is **messy, ambiguous, and constantly changing**
  - We rely on a lot of contextual information as humans
    - Implicit references
    - Implicit memory
    - Colloquialism
- **Rule-based** NLP requires human-created criteria for extracting semantics from an utterance
- **Deep learning** NLP allows data-driven models to extract semantics without rules and language-agnostically

# Conversational AI Building Blocks

- Automated Speech Recognition (ASR) is a **separate** step
  - Google reigns supreme
  - Other specialty providers (e.g., Julius, Sphinx)
  
- We use **Natural Language Processing** techniques to help **understand** the semantics of the **textual utterance**

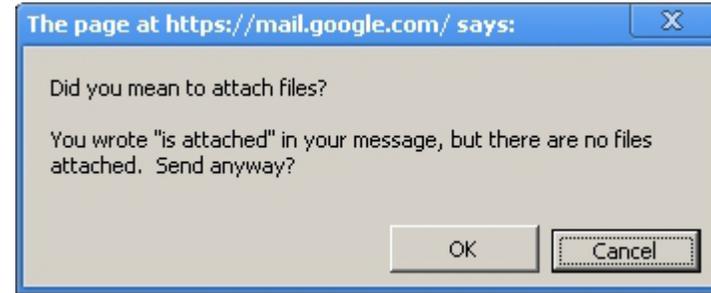
# Natural Language Processing

- **Wiki: Natural language processing (NLP)** is a field of computer science, artificial intelligence, and computational linguistics concerned with the interactions between computers and human (natural) languages.

# What is NLP?

- “Natural” languages
  - English, Mandarin, French, Swahili, Arabic, Nahuatl, ....
  - NOT Java, C++, Perl, ... (take compilers lol)
- Ultimate goal: Natural human-to-computer communication
- Sub-field of Artificial Intelligence, but very interdisciplinary
  - Computer science, human-computer interaction (HCI), linguistics, cognitive psychology, speech signal processing (EE), ...

# Applications: Text Classification



1-21 of 21

Primary Social 1 new Google+ Promotions 2 new Google Offers, Zagat Updates 1 new Google Play

James, me (2) **Hiking** Hiking trip on Saturday - Yay - so glad you can join. We should leave from I 3:14 pm

Hannah Cho **Thank you** - Keri - so good that you and Steve were able to come over. Thank you : 3:05 pm

Lev Birdson **School** Upcoming school conference dates Hello everyone. A few people have www.wired.com

# Applications: Machine Translation

Google buenas noches

All Images Shopping Apps Videos More Search tools

About 20,800,000 results (0.54 seconds)

Spanish English

buenas noches Edit Goodnight

3 more translations

Open in Google Translate

Haaretz הארץ  
אולי הפעם הוא יצליח?

היום: נתניהו שוב ינסה להעביר את ד"ר טרכטנברג בממשלה  
www.haaretz.co.il

אחרי הפאריסה בשבוע שעבר, בלשכת רה"מ מסרבים להתחייב שתהיה הצבעה בתום הדיון. בישראל ביתנו, ש"ס והעצמאות דורשים ריכוך או השמטה של חלק מההמלצות

Like · Comment · Translate · Share · Yesterday at 06:00

9 people like this.

View 1 share

שוקי זקס עם שלם משלם את מחיר החוחות והקמבנות של ראש הממשלה. שלא מצא לנכון להתנגד להבראת מערכת הבריאות. סיכון חיים אינם ולאחרים גרם לסבל. חובה לגרש את החאפן הזה מהעיר. מארה"ב כבר זרקו אותו. מאירופה גם. האיש הזה הורס בשקדנות את המדינה ומכלה את אורחיה. בו לך בנימין נתניהו. נתניהו נתעו: לא מתחייב להצבעה על טרכטנברג. www.ynet.co.il השותפות הקואליציונית הצליחו להטיל מורא על ראש הממשלה, שכבר אינו מתחייב להביא להצבעה היום את מסקנות ד"ר טרכטנברג. ש"ס לפי שעה מתנגדת. ישראל ביתנו ומפלגת העצמאות יקיימו בבוקר ישיבת שרים ויחליטו רק בה - המערכת הפוליטית, חדשות Expand preview · Yesterday at 06:31 · Like · 2 people · Translate

עמי וקנין או שהד"ר יעבור או שביבי יעבור  
Yesterday at 07:10 · Like · Translate

Dalya Gumis שיפסיקו להצביע ותחילו לבצע  
Yesterday at 08:11 · Like · Translate

Yuval Gilor נמאס כבר לך הביתה  
Yesterday at 08:49 · Like · Translate

Haaretz הארץ  
Maybe this time he succeeds? [?]

היום: נתניהו שוב ינסה להעביר את ד"ר טרכטנברג בממשלה  
www.haaretz.co.il

אחרי הפאריסה בשבוע שעבר, בלשכת רה"מ מסרבים להתחייב שתהיה הצבעה בתום הדיון. בישראל ביתנו, ש"ס והעצמאות דורשים ריכוך או השמטה של חלק מההמלצות

Like · Comment · Original · Share · Yesterday at 06:00

9 people like this.

View 1 share

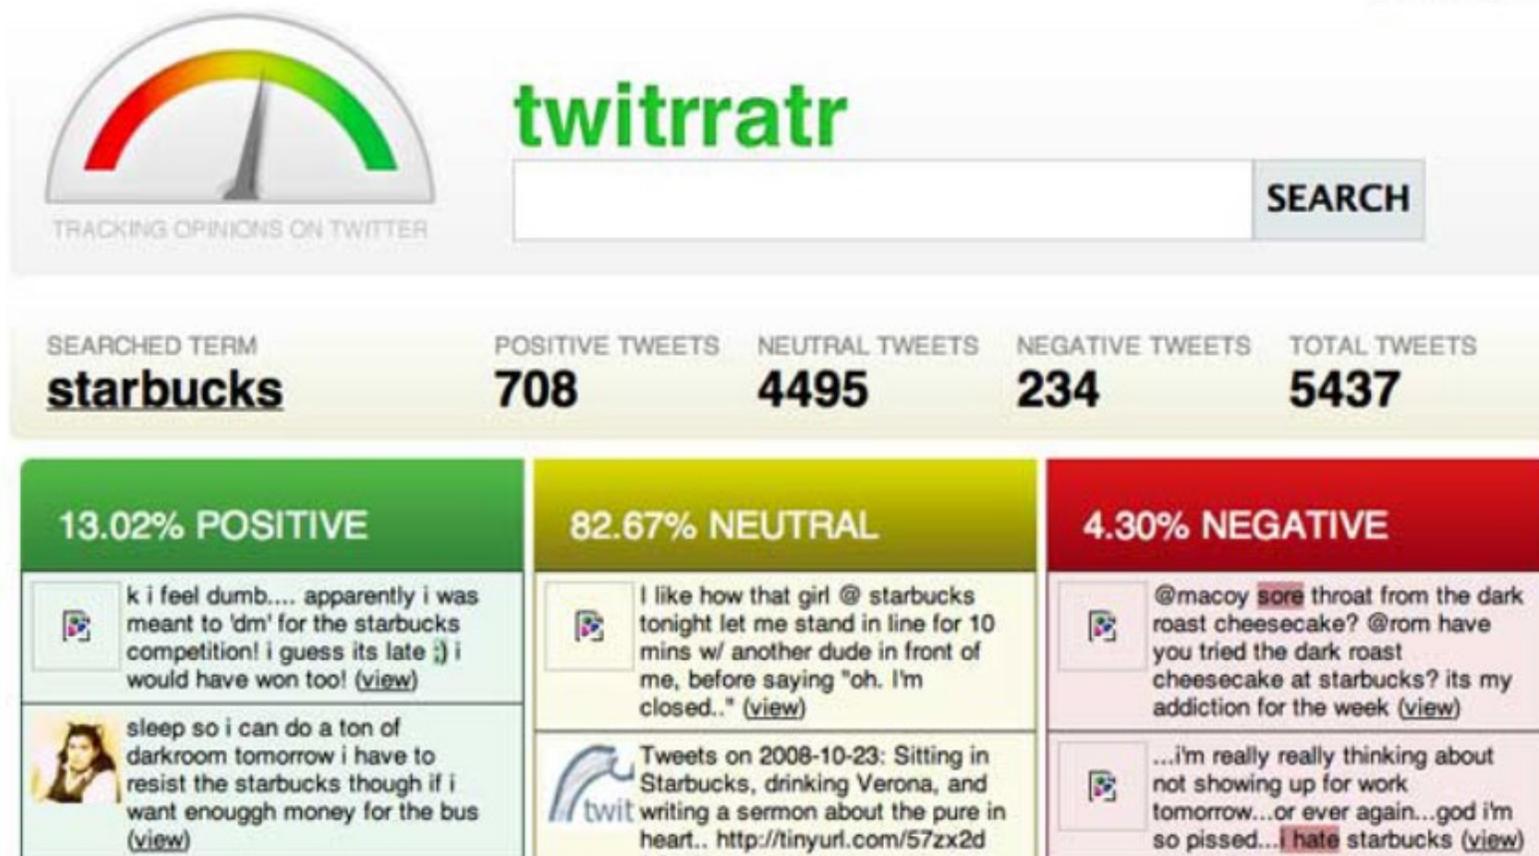
שוקי זקס With Full pay the price for hzhitho vakombinot of the Prime Minister. Not find it appropriate to action habrat the health system. Endangered animals, people and caused the suffering that must drive with ... city. Foundation of Minneapolis had already thrown him. Europe also. This man destroys the country with the citizens in its path. Contempt: you Binyamin Netanyahu. Netanyahu was not committed to voting the Trachtenberg managed to impose a "partnership koalicionit terror the Prime Minister, who is pledging to bring voting day the conclusions report Trachtenberg. Israel Beiteinu-independence morning ministerial session will vihalito only it-political system, news Expand preview · Yesterday at 06:31 · Like · 2 people · Original

Or acknowledging the report moves or the chips will go  
Yesterday at 07:10 · Like · Original

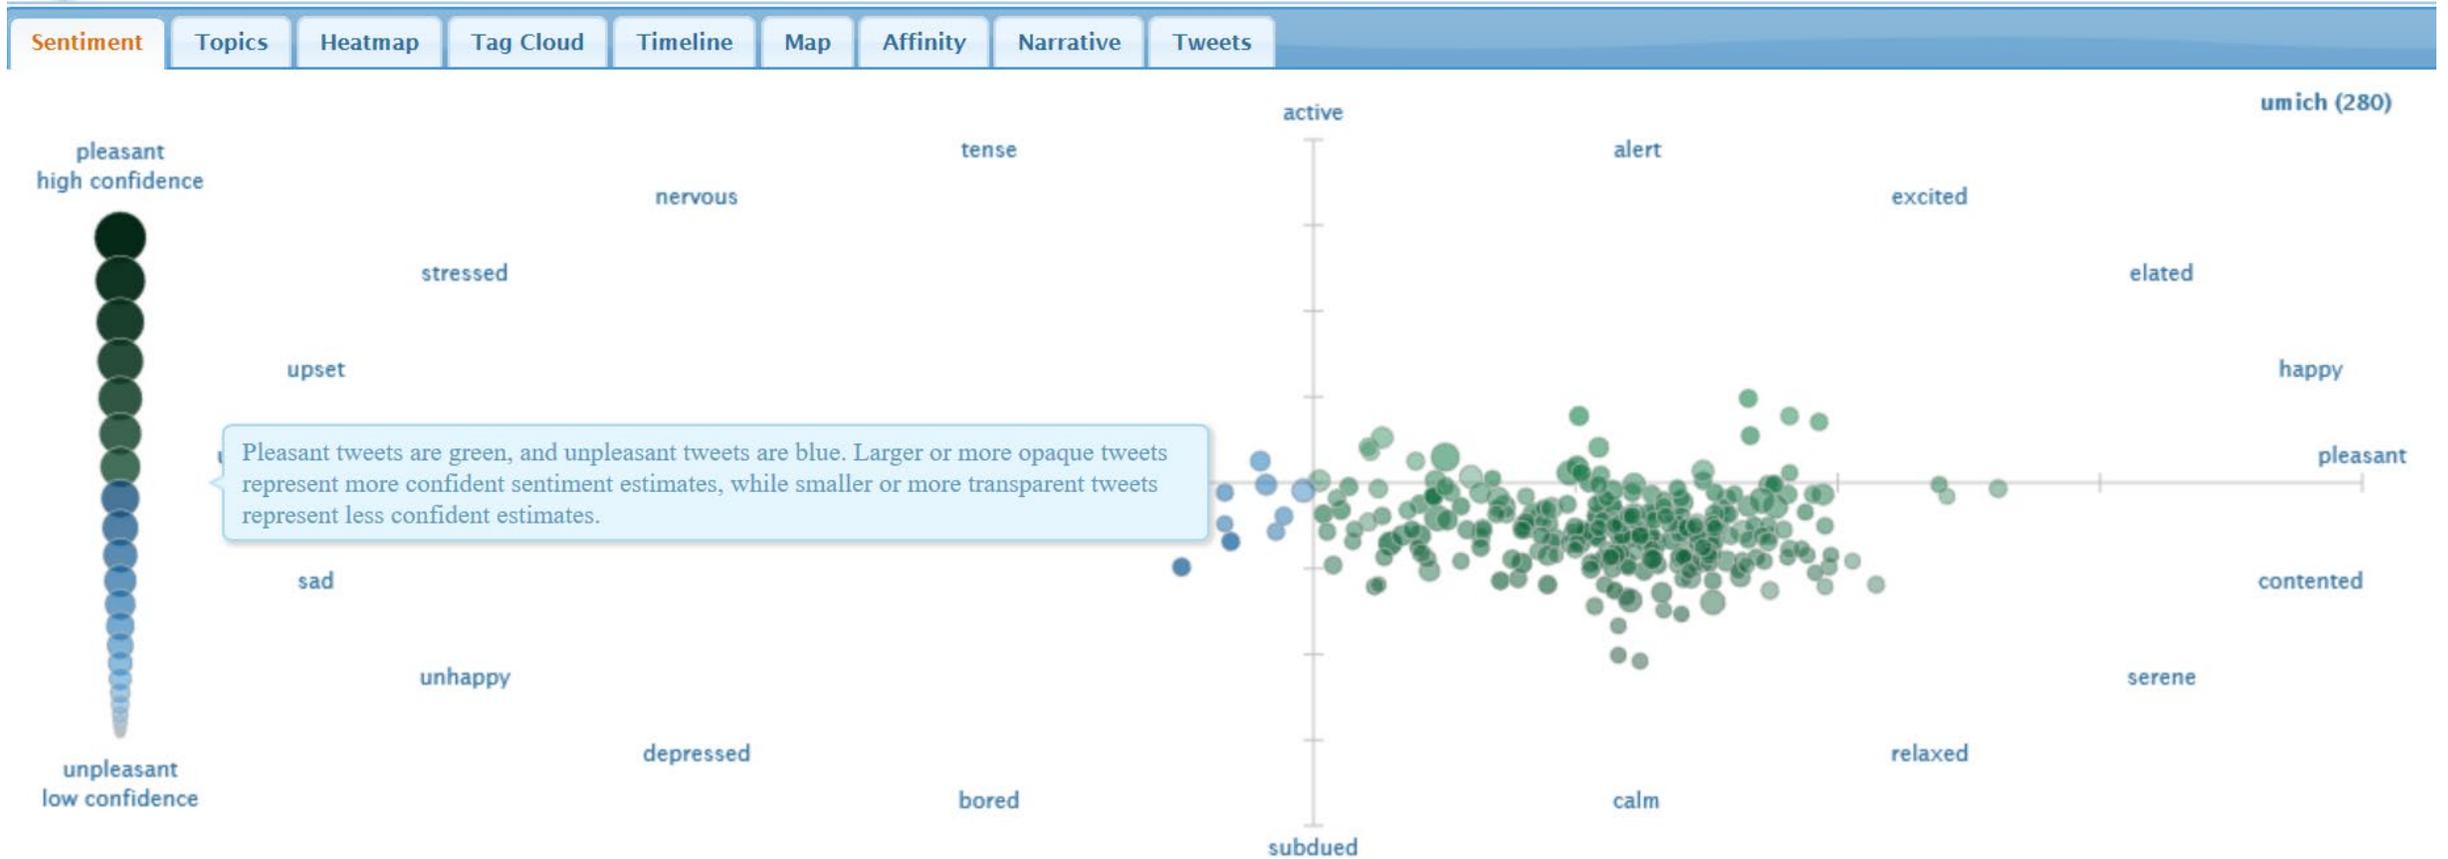
Dalya Gumis Stop vote once to  
Yesterday at 08:11 · Like · Original

Yuval Gilor Tired already go home  
Yesterday at 08:49 · Like · Original

# Applications: Sentiment/Opinion



# Applications: Sentiment/Opinion



# Applications: Information Extraction

New York Times Co. named Russell T. Lewis, 45, president and general manager of its flagship New York Times newspaper, responsible for all business-side activities. He was executive vice president and deputy general manager. He succeeds Lance R. Primis, who in September was named president and chief operating officer of the parent.

Person	Company	Post	State
Russell T. Lewis	New York Times newspaper	president and general manager	start
Russell T. Lewis	New York Times newspaper	executive vice president	end
Lance R. Primis	New York Times Co.	president and CEO	start

# Applications: Text Generation

SciGen: A tool for automatically generating scientific papers

## Router: A Methodology for the Typical Unification of Access Points and Redundancy

Jeremy Stribling, Daniel Aguayo and Maxwell Krohn

### ABSTRACT

Many physicists would agree that, had it not been for congestion control, the evaluation of web browsers might never have occurred. In fact, few hackers worldwide would disagree with the essential unification of voice-over-IP and public-private key pair. In order to solve this riddle, we confirm that SMPs can be made stochastic, cacheable, and interoperable.

### I. INTRODUCTION

Many scholars would agree that, had it not been for active networks, the simulation of Lamport clocks might never have occurred. The notion that end-users synchronize with the investigation of Markov models is rarely outdated. A theoretical grand challenge in theory is the important unification of virtual machines and real-time theory. To what extent can web browsers be constructed to achieve this purpose?

Certainly the usual methods for the emulation of Smalltalk

WMSCI 2005 System <SCI05.System@iis.org>

Sat, Mar 26, 2005 at 12:4

To: XXXX

\*\*\* PLEASE DO NOT REPLY TO THIS E-MAIL ADDRESS \*\*\*

This is a system e-mail account. So, this mailbox is not monitored for responses. Mail sent to this address cannot be answered.

Dear Mr. Jeremy STRIBLING

On behalf of the Organizing Committee, we would like to inform you that, up to the present, we have not received any reviews yet for your paper entitled: "Router: A Methodology for the Typical Unification of Access Points and Redundancy". So, your paper has been accepted, as a non-reviewed paper, for presentation at the 9th World Multiconference on Systemics, Cybernetics and Informatics (WMSCI 2005) to be held in Orlando, USA, on July 10-13, 2005. We will inform you about the conference program, including your presentation, once the timetable is finalized

# Applications: Text Generation

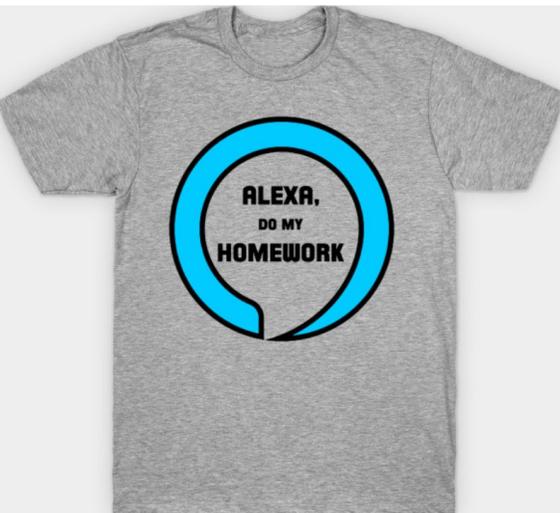
...develop, build and purchase the best equipment known to mankind. Our military is depleted, and we have to do it. We cannot let that happen. We're not going to happen. And yet that's a very tough night for a little while and then you have never did a deal. He did one deal. A house. And if he wins totally in favor of Common Core. He's very, very low energy. So low energy. So low energy. So low energy person, very, very low energy that everybody was apologizing to me. They saw that I want is common sense, above your safety, and above all else. I refuse to be political media – I love that sign. Look at what's going to run. He's just having fun. Just a good time. His brand – “ Like I care about my brand. At this point, my brand – “  
Like I care about my brand. They're not so stupidly and foolishly gave them.



**INSTA-TRUMP**

# NLP in Conversational AI

- NLP is the major workhorse of a conversational AI system
- **In this course, you will apply NLP solutions as part of a conversational AI**



Natural language understanding

Intent Classification:  
*transfer\_money*

Slot mapping:  
*recipient:* "Dr. Leach"  
*amount:* "\$1000"

# Challenges in NLP

- NLP is **hard**
  - Human languages are messy, ambiguous, and ever-changing
- What challenges get in the way of understanding and responding to natural language?
  - Implicit references
  - Ambiguous references / semantics
  - Implicit memory
  - Imprecise rules
  - Myriad languages
  - Scale

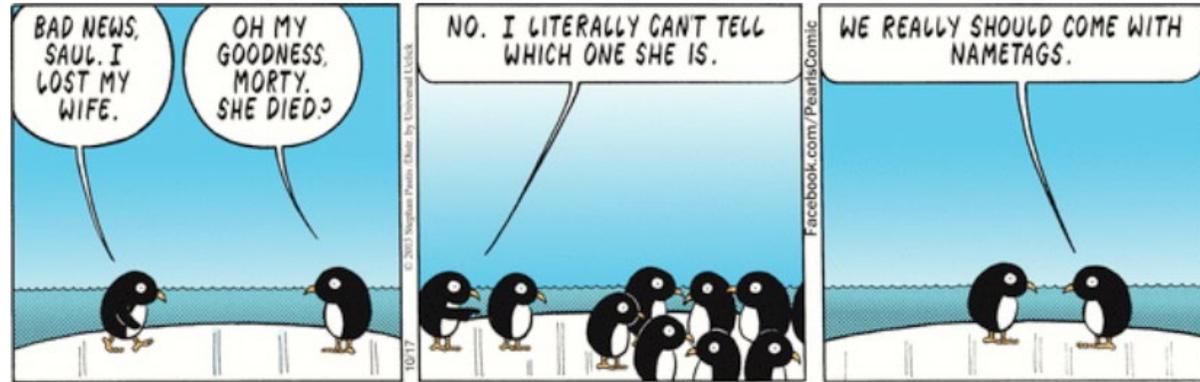


# Challenges in NLP

**Christopher Robin** is alive and well. **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**. The poem was printed in a magazine for others to read. **Mr. Robin** then wrote a book

- Who wrote **Winnie the Pooh**?
- Where did **Chris** live?

# Challenges: Ambiguity



credit: A. Zwicky



# Challenges: Ambiguity

- Word sense / meaning ambiguity



# Challenges: Ambiguity

San Jose cops kill man with knife

Text Paper Translate Listen

## San Jose cops kill man with knife

Ex-college football player, 23, shot 9 times allegedly charged police at fiancée's home

By Hamed Aleaziz and Vivian Ho

Thursday. Police officials said two officers opened fire Wednesday afternoon on Phillip Watkins outside his fiancée's home because they feared for their lives. The officers had been drawn to the home, officials said, by a 911 call reporting an armed home invasion that, it turned out, had been made by Watkins himself.

But the mother of Watkins' fiancée, who also lives in the home on the 1300 block of Sherman Street, said she witnessed the shooting and described it as excessive. Faye Buchanan said the confrontation happened shortly after she called a suicide intervention hotline in hopes of getting Watkins medical help.

Watkins' 911 call came in at 5:01 p.m., said Sgt. Heather Randal, a San Jose police spokeswoman. "The caller stated there was a male breaking into his home armed with a knife," Randal said. "The caller also stated he was locked in an upstairs bedroom with his children and requested help from police."

She said Watkins was on the sidewalk in front of the home when two officers got there. He was holding a knife with a 4-inch blade and ran toward the officers in a threatening manner, Randal said.

"Both officers ordered the suspect to stop and drop the knife," Randal said. "The suspect continued to charge the officers with the knife in his hand. Both officers, fearing for their safety and defense of their life, fired at the suspect."

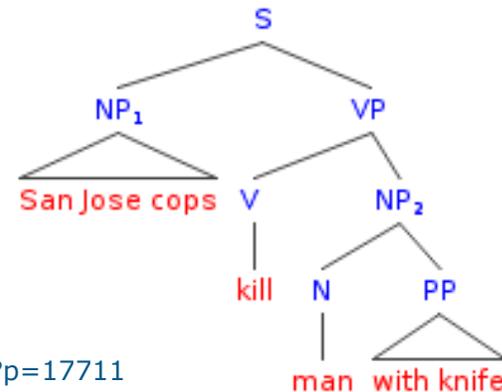
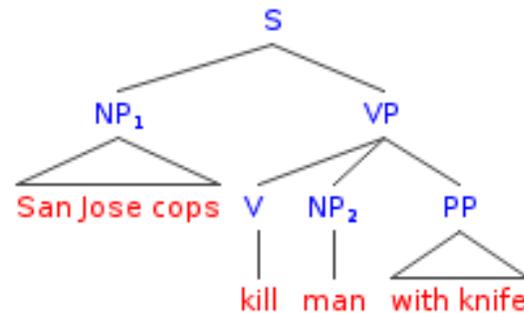
On the police radio, one officer said, "We have a male with a knife. He's walking toward us."

"Shots fired! Shots fired!" an officer said moments later.

A short time later, an officer reported, "Male is down. Knife's still in hand."

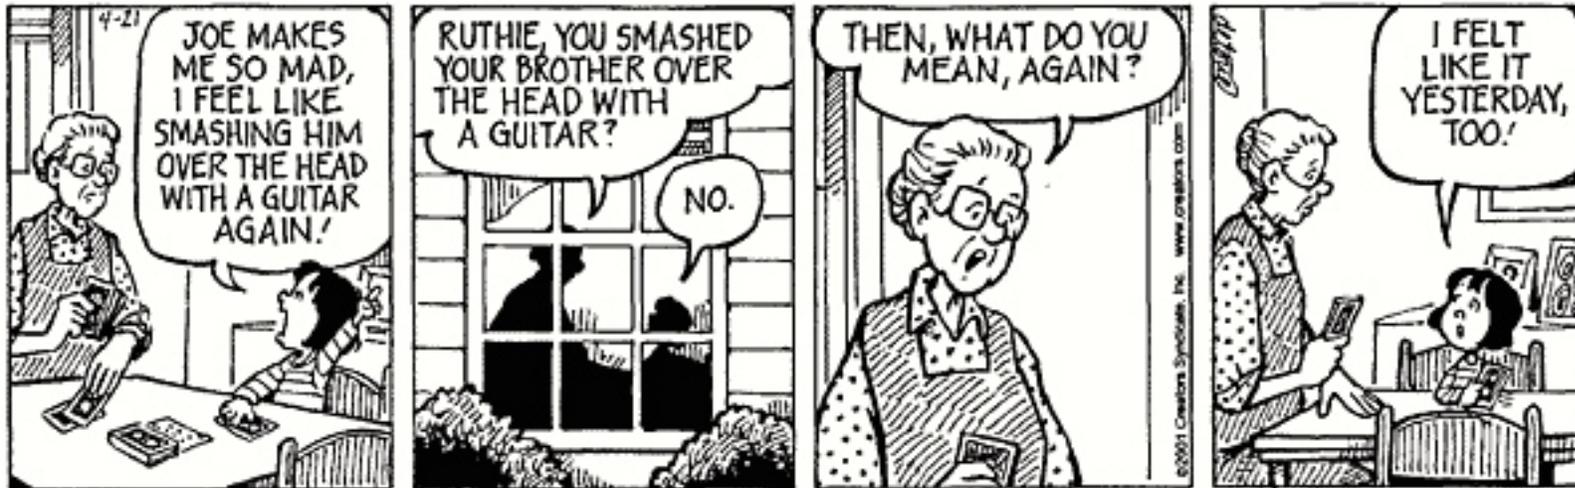
Buchanan said she had been prompted to call the *Shoot continues on D8*

Back Continue



Credit: Mark Liberman, <http://languagelog.ldc.upenn.edu/nll/?p=17711>

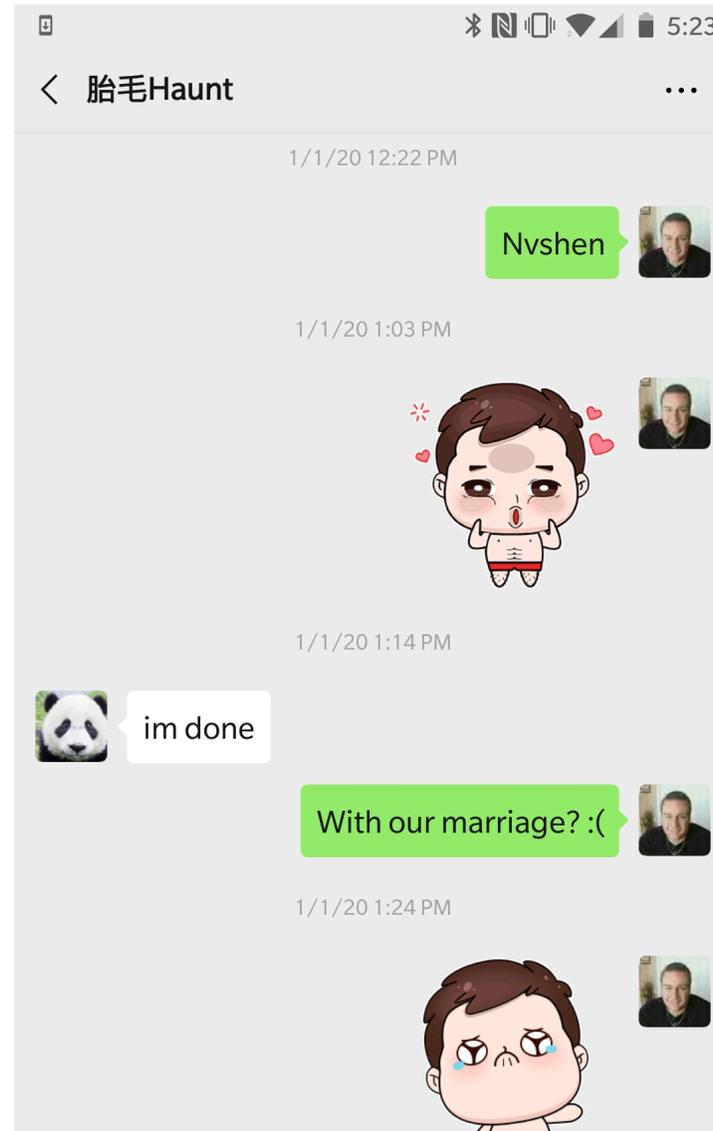
# Challenges: Ambiguity



# Challenges: Ambiguity

- Ambiguous headlines:
  - Include your children when baking cookies
  - Local High School Dropouts Cut in Half
  - Hospitals are Sued by 7 Foot Doctors
  - Iraqi Head Seeks Arms
  
- Safety Experts Say School Bus Passengers Should Be Belted
- Teacher Strikes Idle Kids

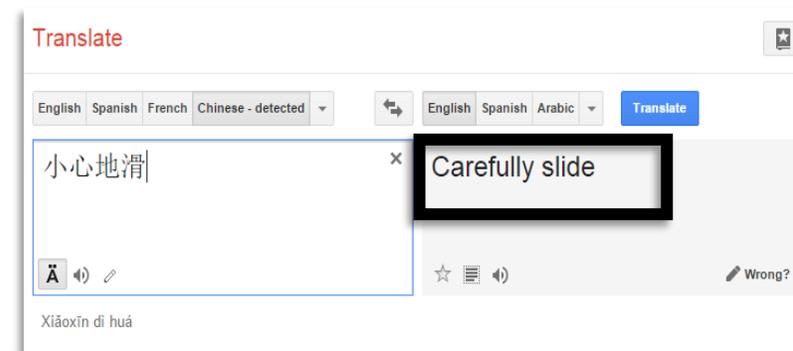
# Challenges: Implicit context



# Challenges: Language is dynamic

LOL	Laugh out loud
G2G	Got to go
BFN	Bye for now
B4N	Bye for now
Idk	I don't know
FWIW	For what it's worth
LUWAMH	Love you with all my heart

# Challenges: Language is Compositional



# Challenges: Language is Compositional



**小心:**  
Carefully  
Careful  
Take  
Care  
Caution



**地滑:**  
Slide  
Landslip  
Wet Floor  
Smooth



# Challenges: Scale

- Examples:
  - Bible (King James version): ~700K
  - Penn Tree bank ~1M from Wall street journal
  - Newswire collection: 500M+
  - Wikipedia: 2.9 billion word (English)
  - Web: several billions of words

# Challenges: Summary

- **Natural** language is replete with messy issues!

# NLP Terminology

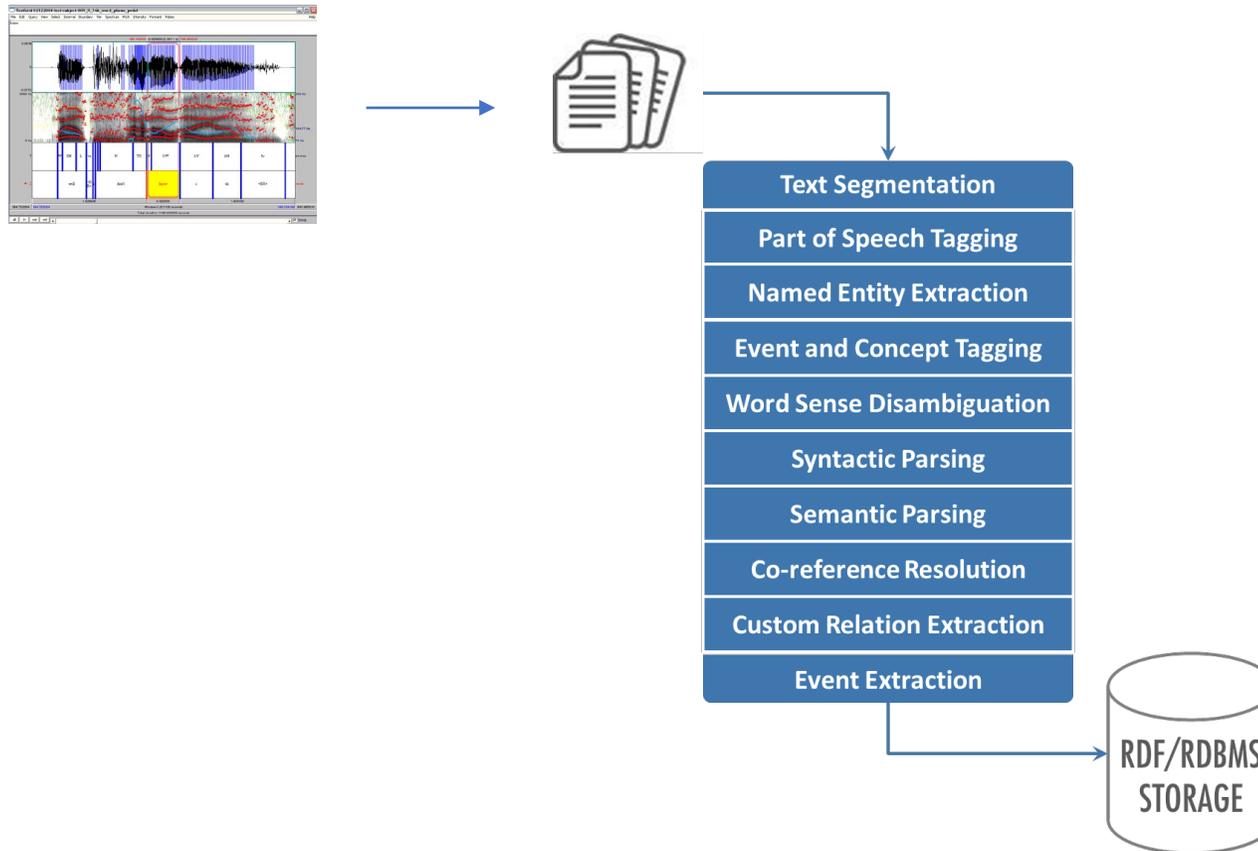
- Morphology: What is a word?
- **奧林匹克運動會** (**希臘語**: Ολυμπιακοί Αγώνες, 簡稱**奧運會**或**奧運**) 是**國際奧林匹克委員會**主辦的包含多種**體育**運動項目的國際性運動會, 每四年舉行一次。
- **كبيوتها** = “to her houses”
- Lexicography: What does each word mean?
  - He plays bass guitar.
  - That bass was delicious!
- Syntax: How do the words relate to each other?
  - **The dog** bit **the man**. ≠ **The man** bit **the dog**.
  - But in Russian: **человек** **собаку** съел = **человек** съел **собаку**



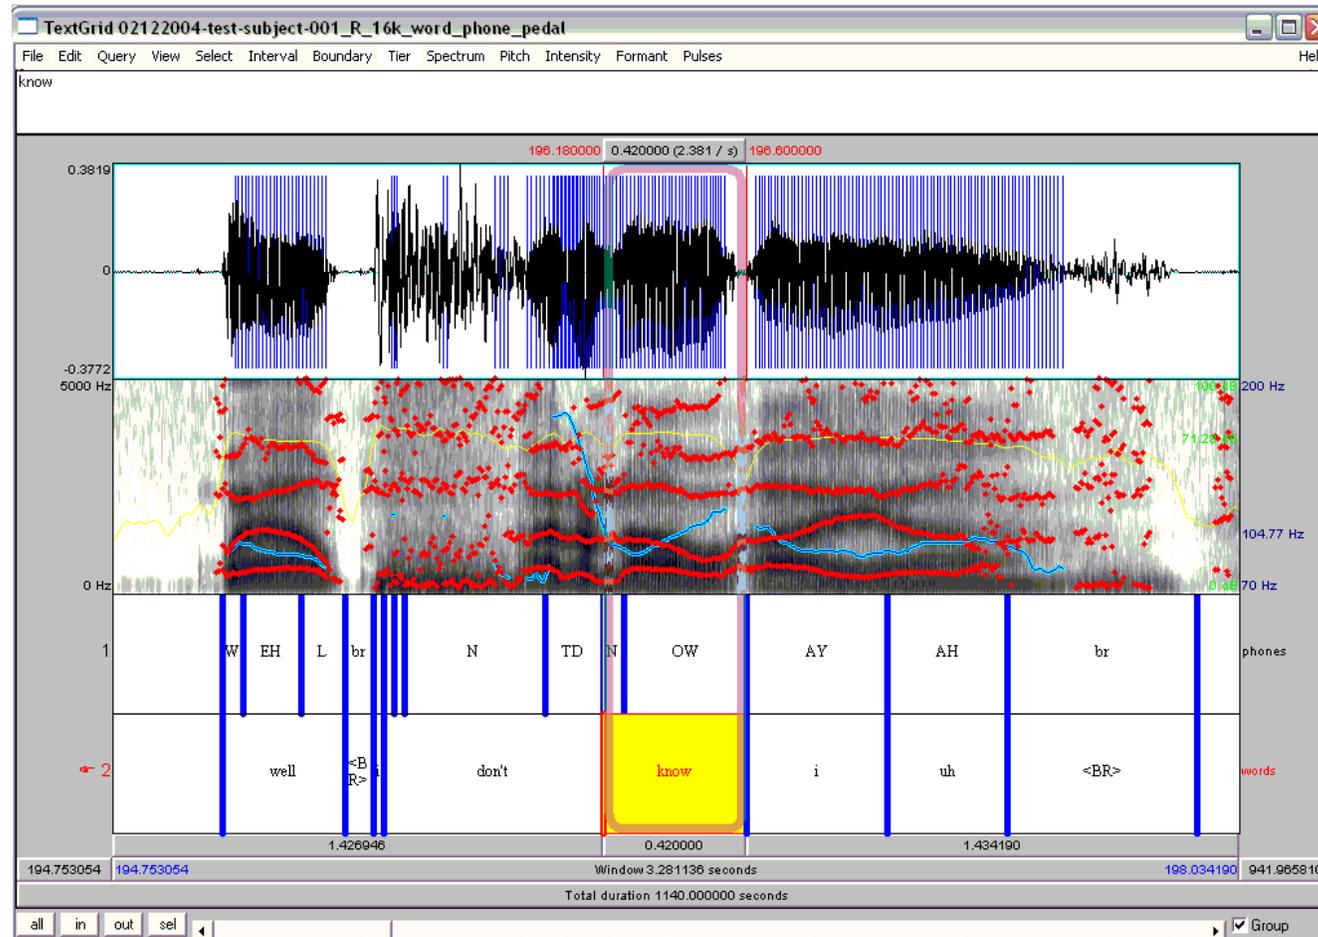
# NLP Terminology

- Semantics: How can we infer meaning from sentences?
  - I saw the man on the hill with the telescope.
  - The ipod is so **small!** 😊
  - The monitor is so **small!** 😞
- Discourse: How about across many sentences?
  - President Bush met with President-Elect Obama today at the White House. He welcomed him, and showed him around.
  - Who is “he”? Who is “him”? How would a computer figure that out?

# Classic NLP Pipeline



# Digital Signal Processing meets NLP



Examples from Prof. Julia Hirschberg's slides

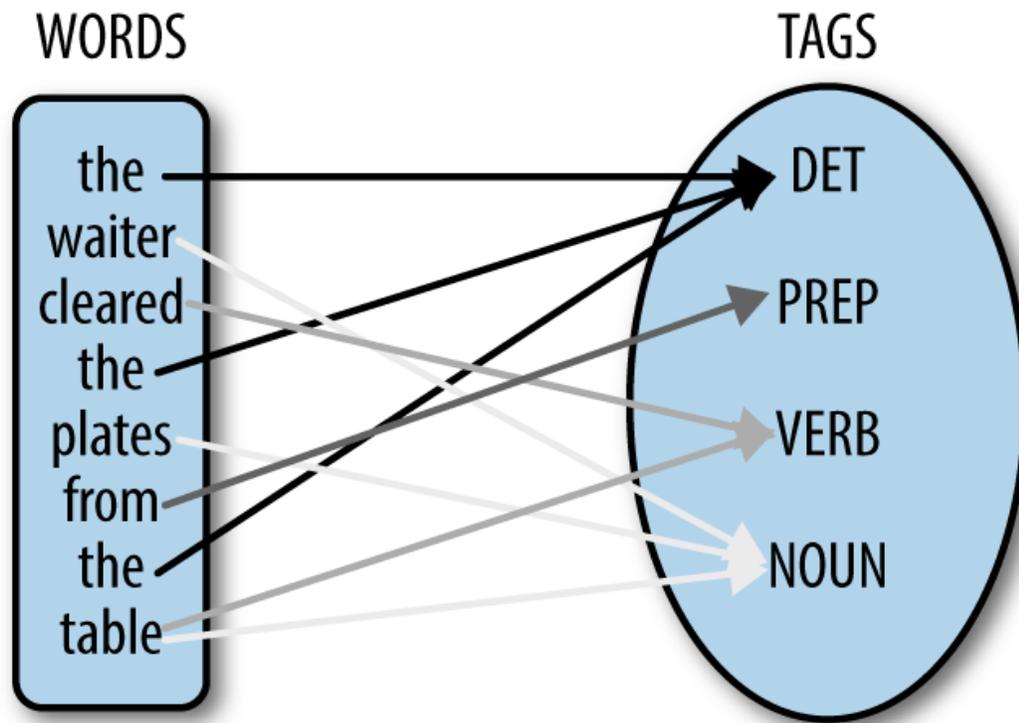
# Spoken Language Processing

- Speech Recognition
  - Automatic dictation, assistance for blind people, automatic 411, ...
- Related things we study...
  - How does intonation affect semantic meaning?
  - Detecting uncertainty and emotions
  - Detecting deception!
- Why is this hard?
  - Each speaker has a different voice (male vs female, child versus older person)
  - Many different accents (Scottish, American, non-native speakers) and ways of speaking
  - Conversation: turn taking, interruptions, ...

# Spoken Language Processing

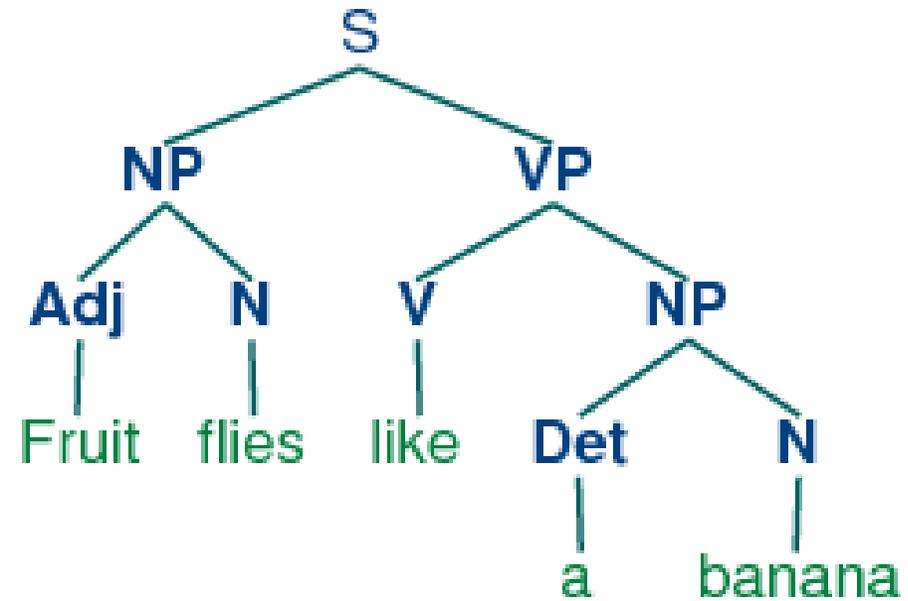
- Text-to-Speech / Spoken dialog systems
  - Call response centers, tutoring systems, ...
- Related things we study...
  - Making computer voices sound more human
  - Making computer speech acts more human-like

# Part of speech tagging

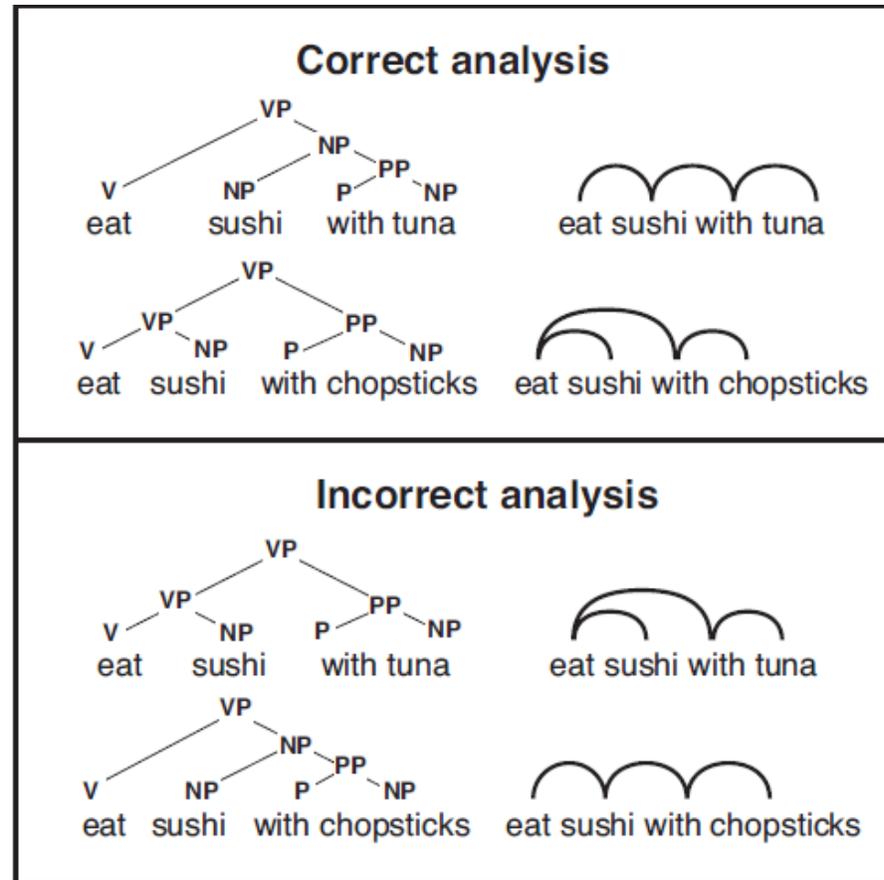


Tag	Meaning	English Examples
ADJ	adjective	new, good, high, special, big, local
ADP	adposition	on, of, at, with, by, into, under
ADV	adverb	really, already, still, early, now
CONJ	conjunction	and, or, but, if, while, although
DET	determiner, article	the, a, some, most, every, no, which
NOUN	noun	year, home, costs, time, Africa
NUM	numeral	twenty-four, fourth, 1991, 14:24
PRT	particle	at, on, out, over per, that, up, with
PRON	pronoun	he, their, her, its, my, I, us
VERB	verb	is, say, told, given, playing, would
.	punctuation marks	. , ; !
X	other	ersatz, esprit, dunno, gr8, univeristy

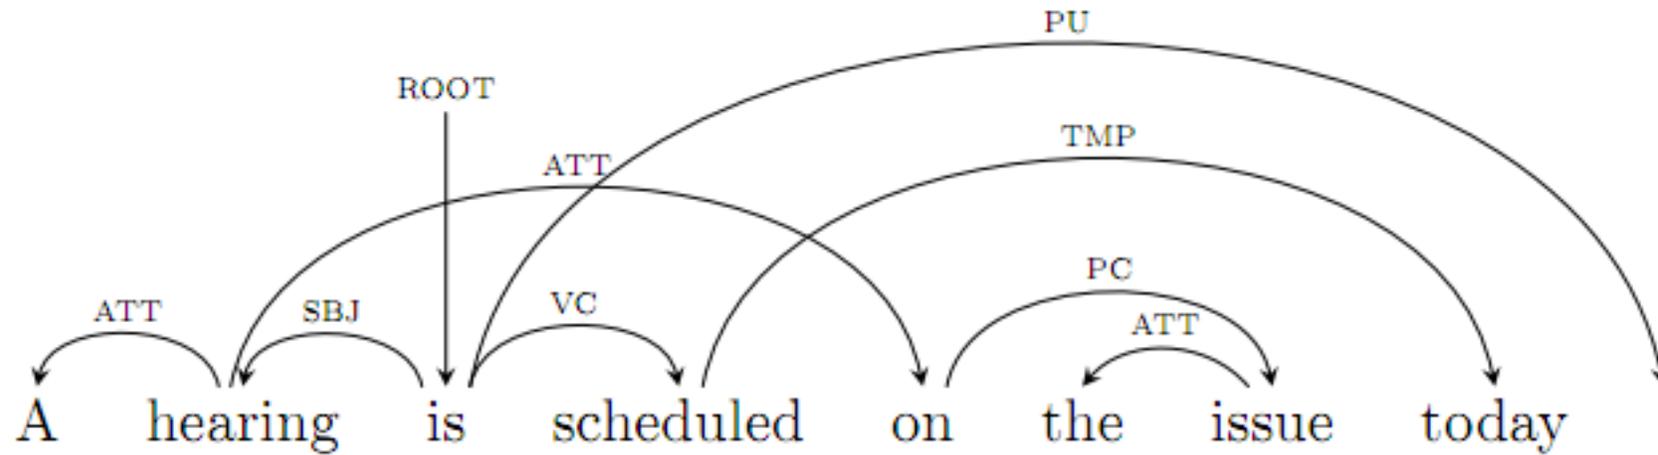
# Syntactic (Constituency) parsing



# Syntactic structure => meaning

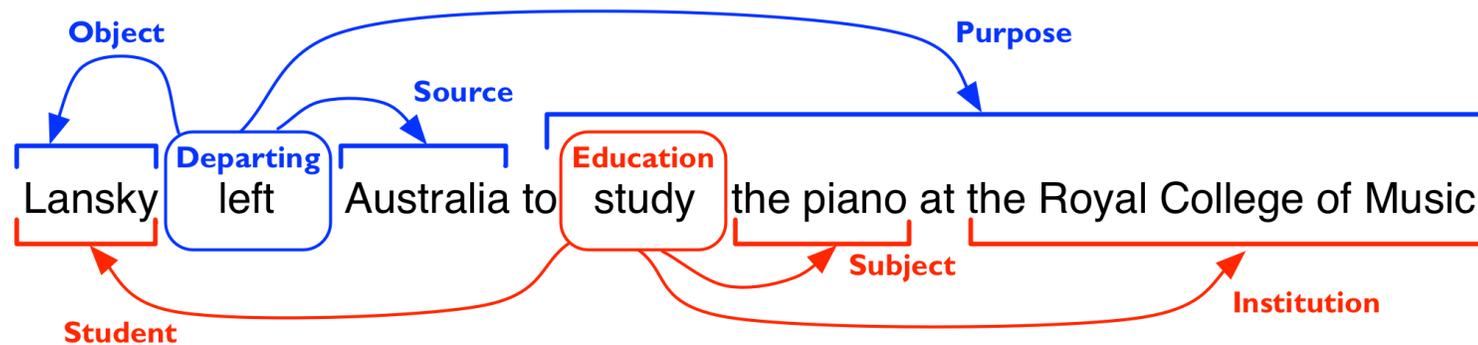


# Dependency Parsing



# Semantic analysis

- Word sense disambiguation
- Semantic role labeling



Credit: Ivan Titov

Q: [Chris] = [Mr. Robin] ?

**Christopher Robin** is alive and well. **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**. The poem was printed in a magazine for others to read. **Mr. Robin** then wrote a book

Slide modified from Dan Roth

# Co-reference Resolution

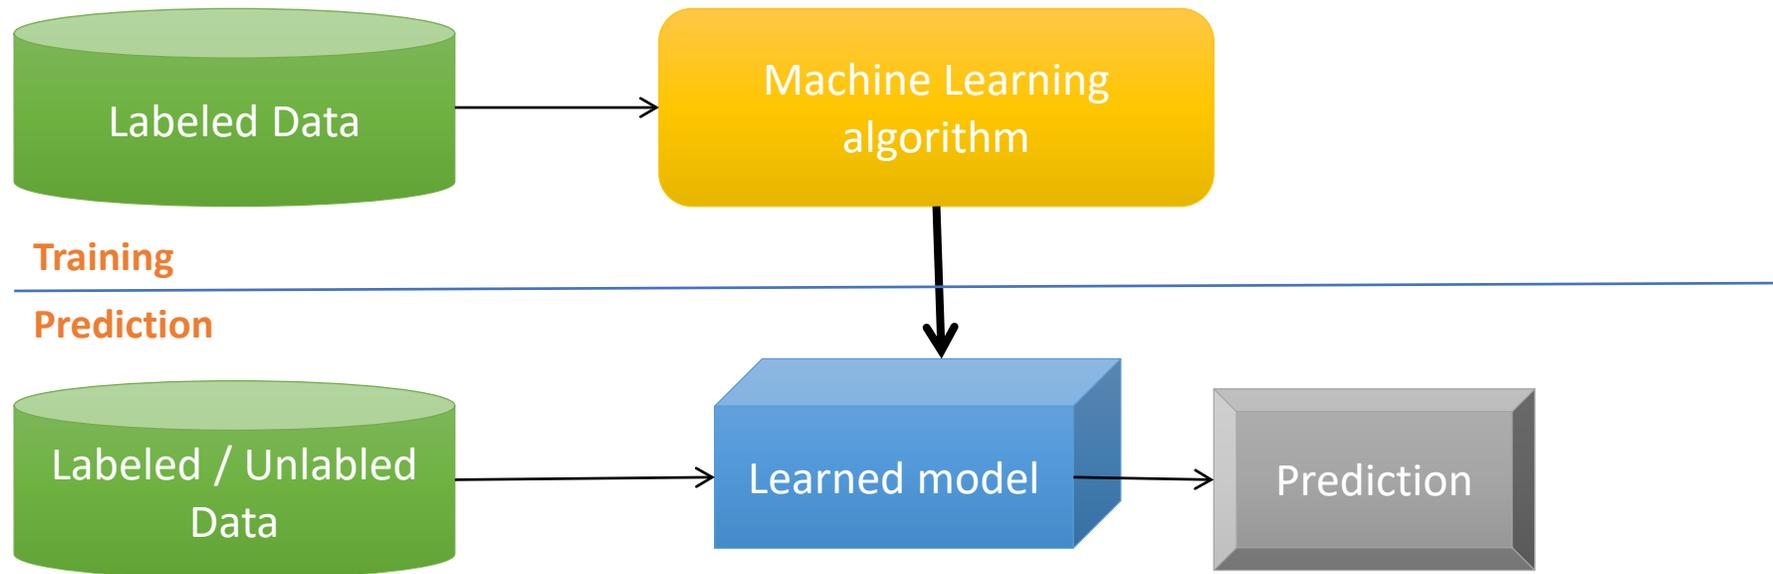
**Christopher Robin** is alive and well. **He** is the same person that you read about in the book, **Winnie the Pooh**. As a **boy**, **Chris** lived in a pretty home called **Cotchfield Farm**. When **Chris** was three years old, **his father** wrote a poem about **him**. The poem was printed in a magazine for others to read. **Mr. Robin** then wrote a book

# So what about ML and Deep Learning?

- **Conversational AI** uses **Machine Learning (ML)** and **Deep Learning (DL)** to complete many high-level **NLP** tasks

# Machine Learning Basics

Machine learning is a field of computer science that gives computers the ability to **learn without being explicitly programmed**



Methods that can learn from and make predictions on data

# Types of Learning

**Supervised:** Learning with a **labeled training** set

Example: email *classification* with already labeled emails

**Unsupervised:** Discover **patterns** in **unlabeled** data

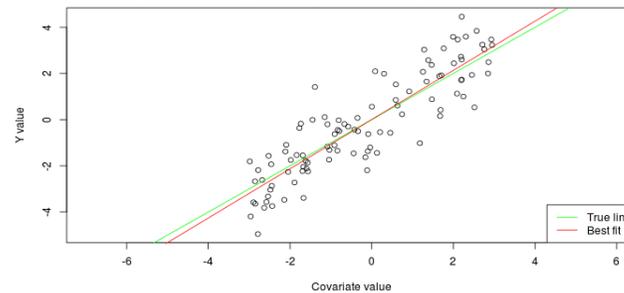
Example: *cluster* similar documents based on text

**Reinforcement learning:** learn to **act** based on **feedback/reward**

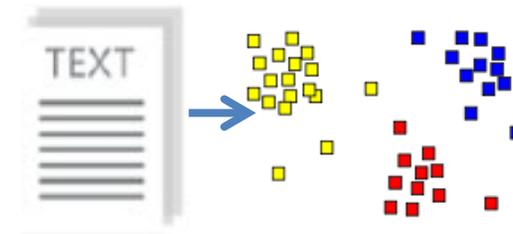
Example: learn to play Go, reward: *win or lose*



Classification



Regression



Clustering

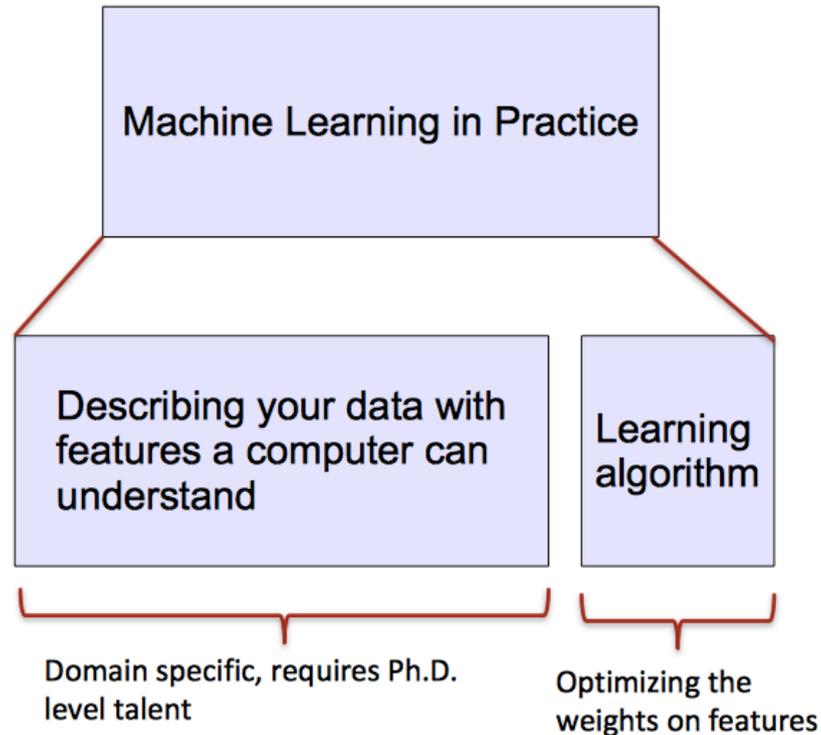
Anomaly Detection  
Sequence labeling

...

# ML vs Deep Learning

Most machine learning methods work well because of **human-designed representations** and **input features**

ML becomes just **optimizing weights** to best make a final prediction



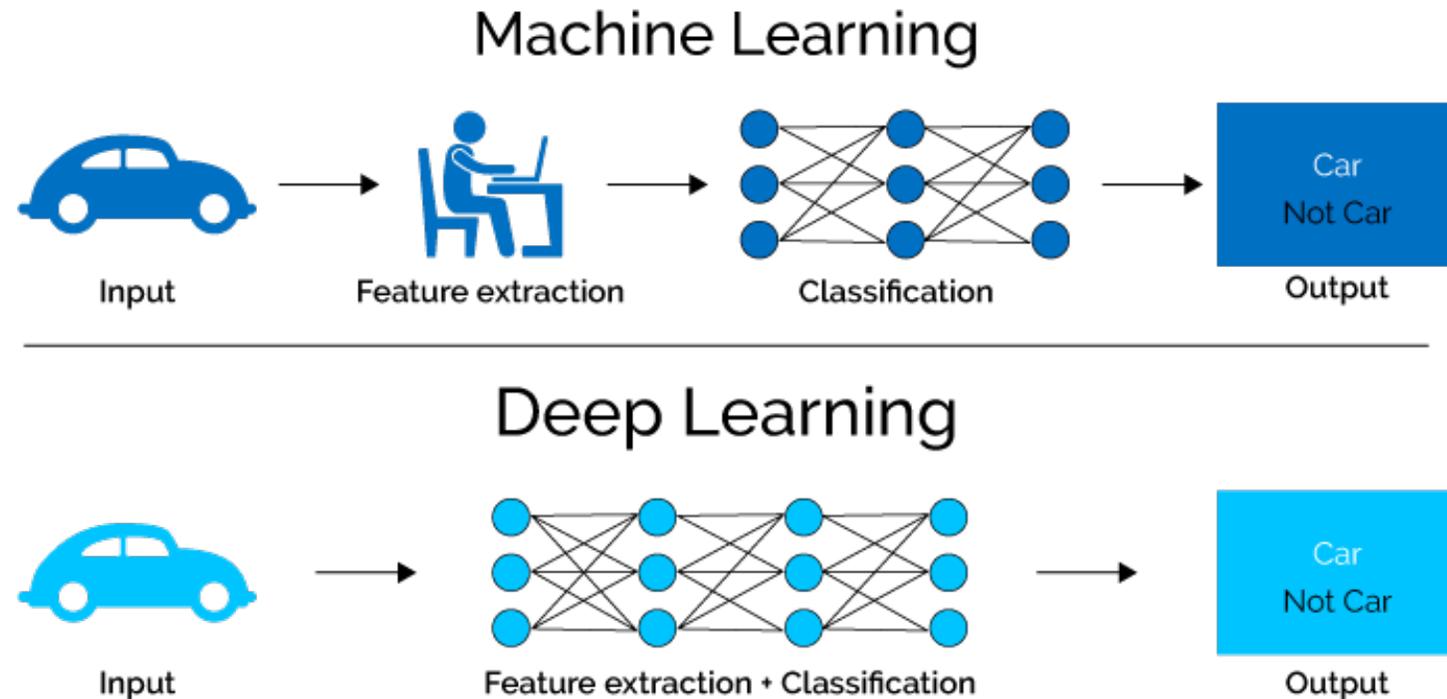
Feature	NER
Current Word	✓
Previous Word	✓
Next Word	✓
Current Word Character n-gram	all
Current POS Tag	✓
Surrounding POS Tag Sequence	✓
Current Word Shape	✓
Surrounding Word Shape Sequence	✓
Presence of Word in Left Window	size 4
Presence of Word in Right Window	size 4

# What is Deep Learning?

A machine learning subfield of learning **representations** of data. Exceptional effective at **learning patterns**.

Deep learning algorithms attempt to learn (multiple levels of) representation by using a **hierarchy of multiple layers**

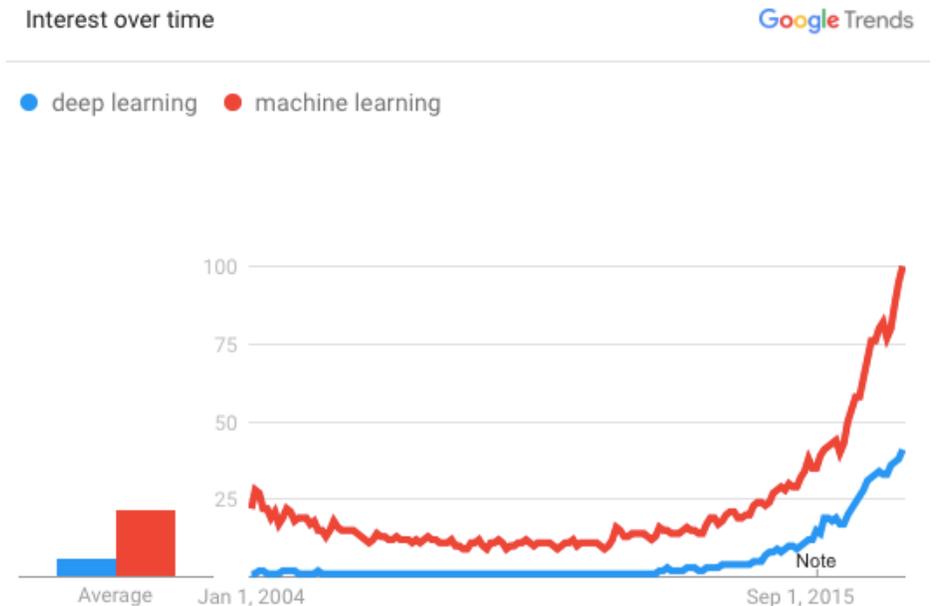
If you provide the system **tons of information**, it begins to understand it and respond in useful ways.



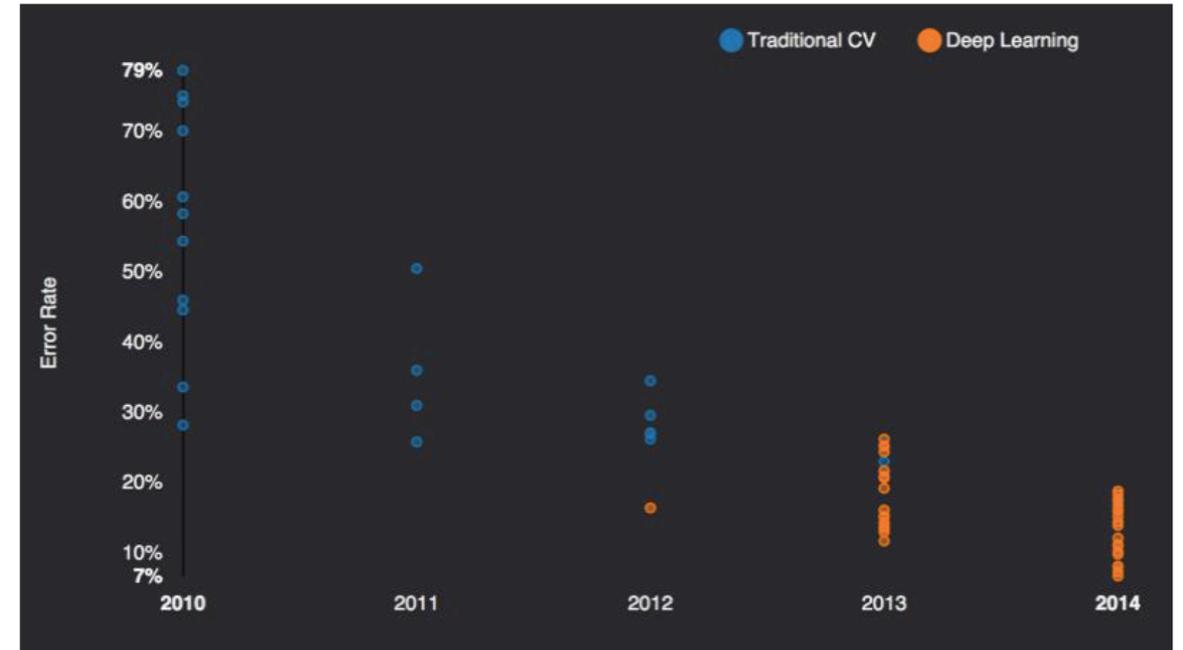
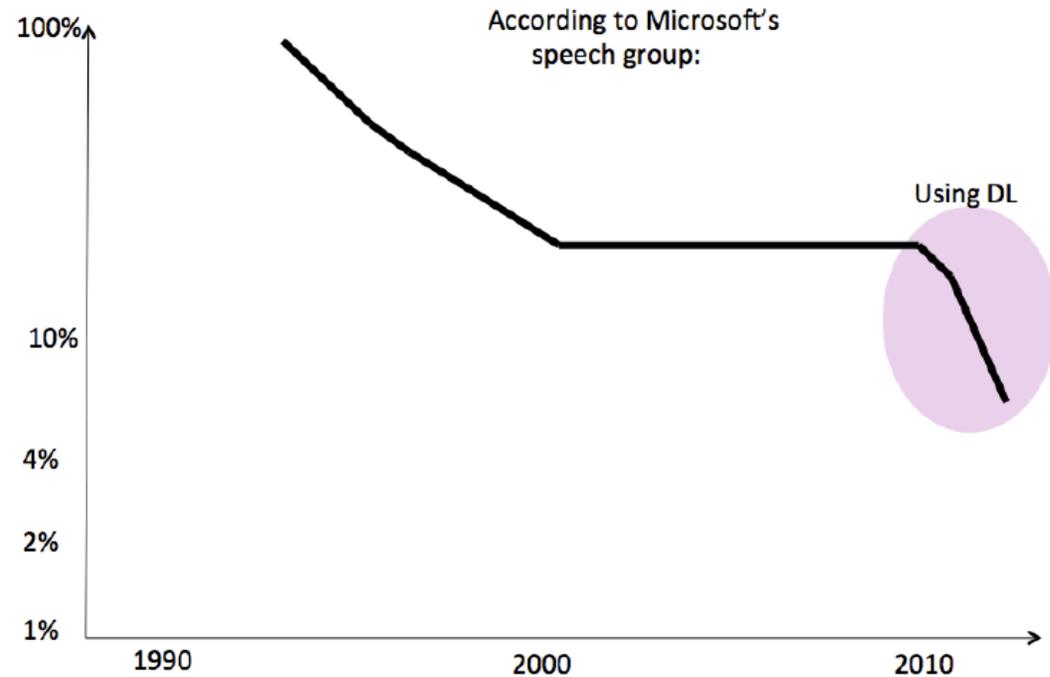
# Why is DL useful?

- Manually designed features are often **over-specified**, **incomplete** and take a **long time to design** and validate
- Learned Features are **easy to adapt**, **fast** to learn
- Deep learning provides a very **flexible**, (almost?) **universal**, learnable framework for representing world, visual and linguistic information.
- Can learn both unsupervised and supervised
- Effective **end-to-end** joint system learning
- Utilize large amounts of training data

In ~2010 DL started outperforming other ML techniques  
first in speech and vision, then NLP



# State of the art for...



ImageNet: The "computer vision World Cup"

## Deep Learning in Speech Recognition

Several big improvements in recent years in NLP

- ✓ Machine Translation
- ✓ Sentiment Analysis
- ✓ Dialogue Agents
- ✓ Question Answering
- ✓ Text Classification ...

Leverage different levels of representation

- words & characters
- syntax & semantics