CSE 5525: Speech and Language Processing

Instructor: Alan Ritter

Administrative Details

- Course Webpage
  - http://aritter.github.io/courses/5525_fall16.html

CSE 5525: Speech and Language Processing

Fundamentals of natural language processing, automatic speech recognition and speech synthesis; lab projects concentrating on building systems to process written and/or spoken language.

Details
Instructor: Alan Ritter (ritter.1492@osu.edu)
Wednesday: Friday, 12:45 - 2:05
Place: Bolz Hall 318
Office Hours: Wednesdays 4:00-5:00pm, Dreese 595

Administrative Details

Grading

Grading will be based on:

Participation and in-class Exercises (10%)

You will receive credit for asking and answering questions related to the homework on Piazza, and also for engaging in class discussion.

Homeworks (50%)

The homeworks will include both written and programming assignments. Homework should be submitted to the Dropbox folder in Carmen by 11:59pm on the day it is due (unless otherwise instructed). Late homework will be accepted up to 48 hours later for 50% credit. After 48 hours, late homework will not be accepted. Please email your homework to the instructor if there are any technical issues with submission.

Midterm (20%)

There will be an in-class midterm on March 2.

Final Projects (20%)

The final project is an open-ended assignment, with the goal of gaining experience applying the techniques presented in class to real-world datasets. Students should work in groups of 3-4. It is a good idea to discuss your planned project with the instructor to get feedback. The final project report should be 4 pages and is due on April 30. The report should describe the problem you are solving, what data is being used, the proposed technique you are applying in addition to what baseline is used to compare against.

Administrative Details

- Piazza (discussion and resources)
  - https://piazza.com/osu/spring2016/5525/home
- Carmen (homework submission and grades)
  - https://carmen.osu.edu/d2l/home/11684583

Course Details

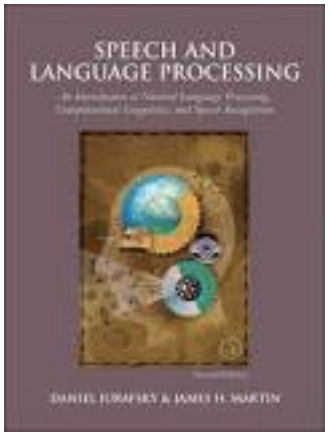

- Book
  - Jurfasky and Martin (2nd edition)
  - Will also use selections from the 3rd edition (unpublished)
  - Some other readings as well
- Prerequisits / Assumptions
  - Basic Probability
  - Basic Linear Algebra
  - Python or ability to learn Python quickly
  - Linux/Unix (For windows users: https://www.cygwin.com/)
  - Numpy/scipy

What to Expect

- Lots of math and programming
- A bit of Linguistics
- Computing Resources:
  - Experiments could take hours to run depending the efficiency of your code. We recommend you start early.

- Questions?

The Role of NLP in Artificial Intelligence

- Goal: give computers the ability to think (and speak)
- Long history in computer science
- Turing Test (Alan Turing 1950)
- Loebner Prize
  - Little interest from NLP community…
  - Very simple programs can fool some judges some of the time

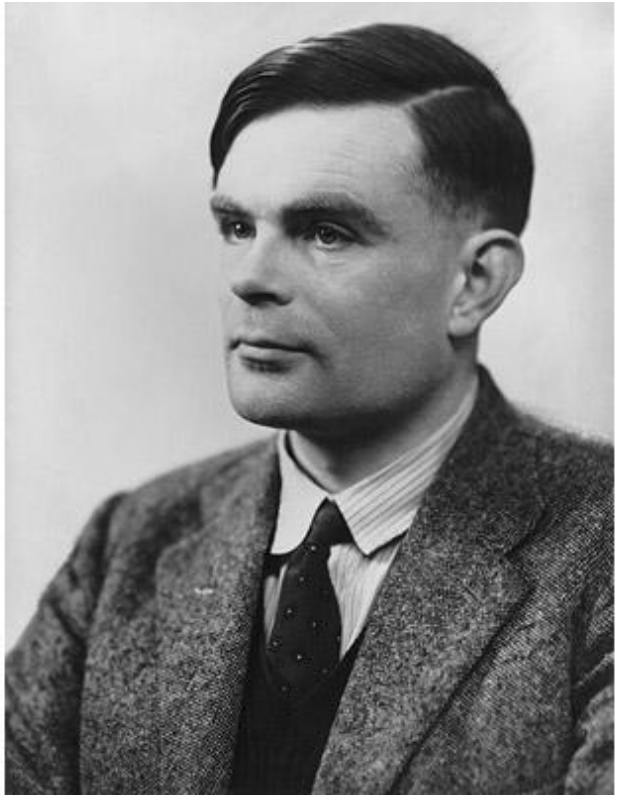

The Chinese Room (Searle 1980)

NLP in Science Fiction

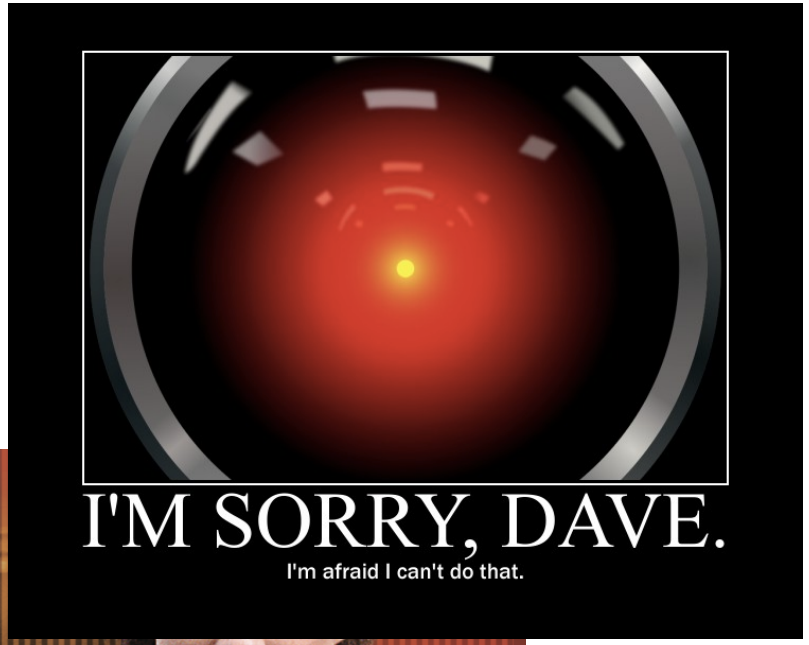

Today's NLP Applications

- Speech Interfaces
- Machine Translation
- Search Engines
- Information Extraction
- Summarization
- …

NLP Today: Speech Interfaces

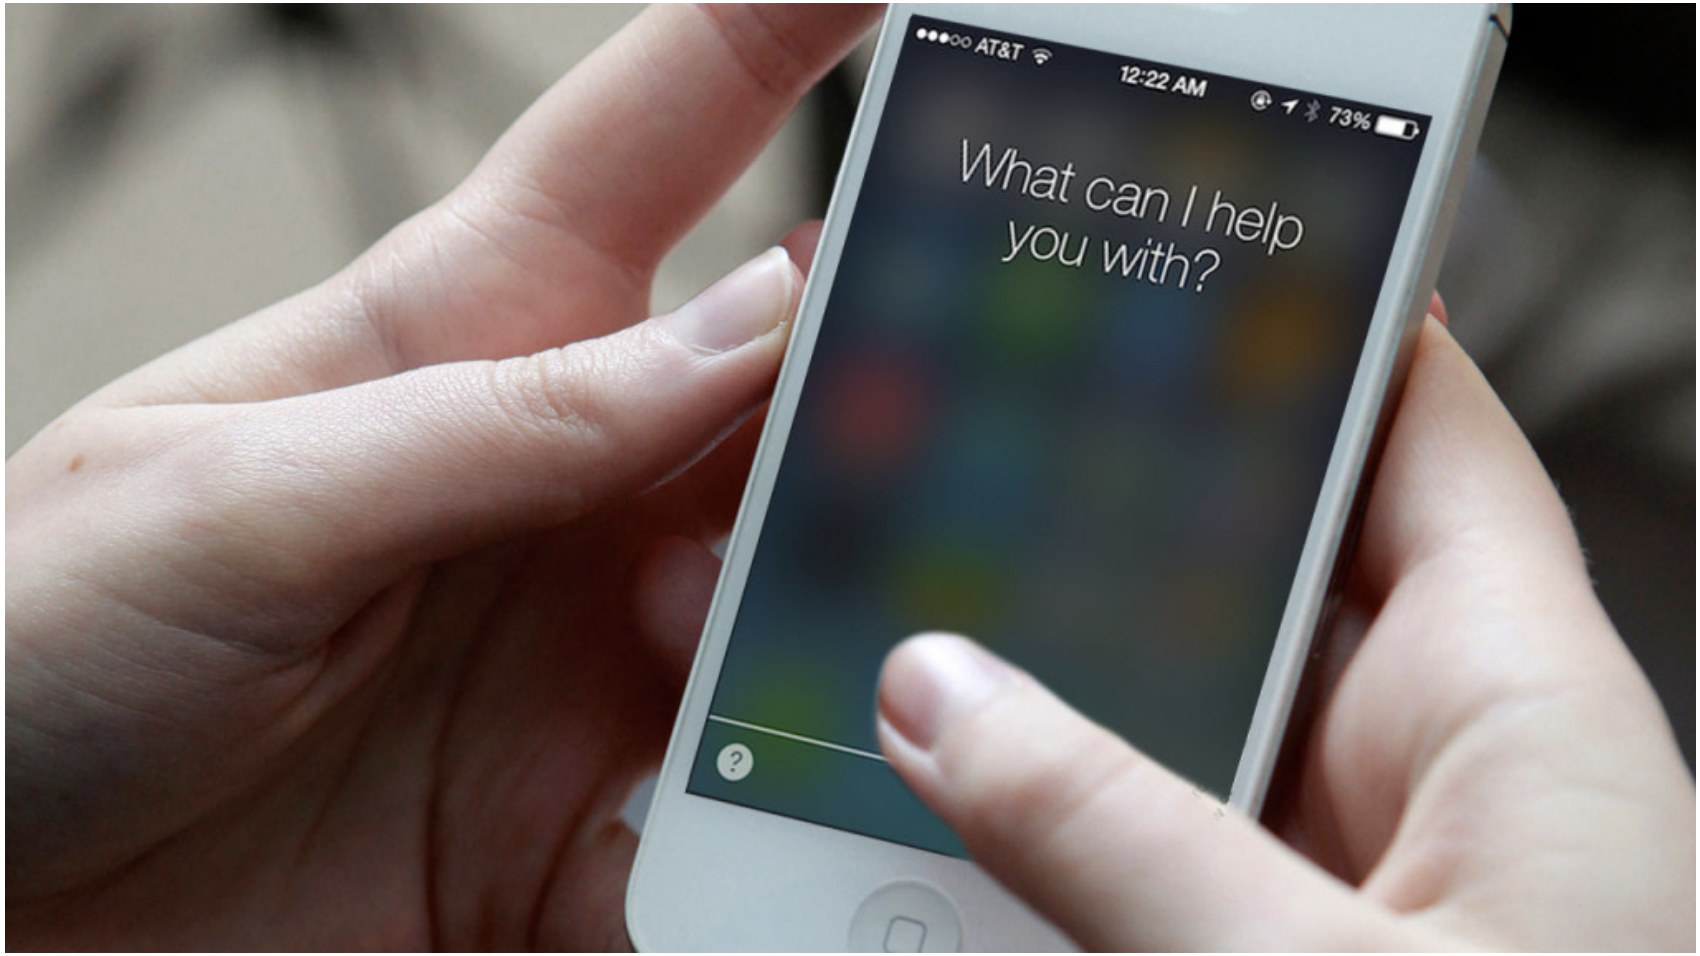

Speech Interfaces

- Automatic Speech Recognition
  - Audio In, text out
  - SOTA: 0.3% error for digit strings, 5% dictation 50%+ for TV

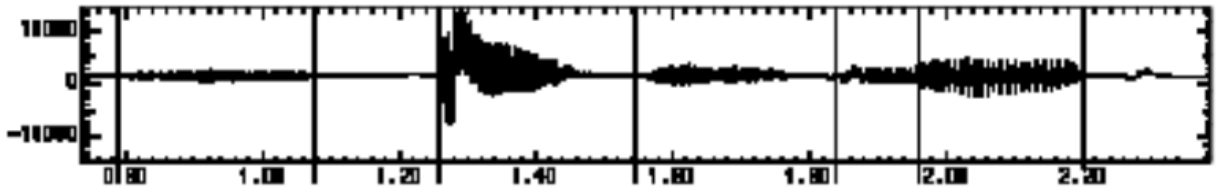

"Speech Lab"

Natural Language Understanding

- Convert words into semantic representation that can be used to query a database

cuisine location

Where can I find **Indian** food in **Clintonville**?

1. Mughal Darbar

42 reviews

$$ · Indian

University District

2321 N High St
Columbus, OH 43202
(614) 477-6065

This restaurant takes reservations — Find a Table

This restaurant accepts pickup orders — Start Order

Speech Interfaces

- Text to Speech
  - Text in Audio Out

Question Answering

Summarization

User reviews:

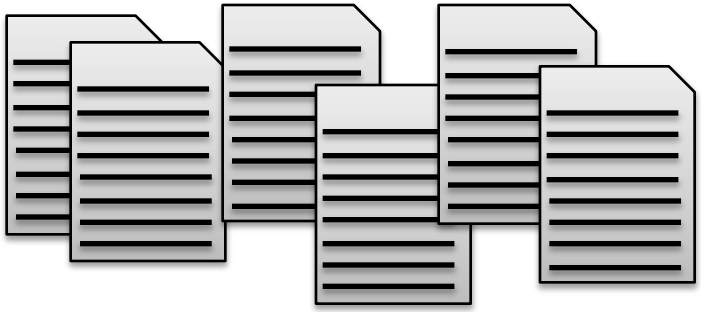

"**Ash** always makes visitors feel like her personal guests and never disappoints as a hostess." in 3 reviews

"Great customer service and **delicious food**." in 2 reviews

"I did opt to eat on their **patio** during the nice summer weather and I would have to say that it was pleasant." in 7 reviews

Outdoor Seating: Yes

Information Extraction

“Yess! Yess! Its official Nintendo announced today that they Will release the Nintendo 3DS in north America march 27 for $250”

Information Extraction

"Yess! Yess! Its official ***Nintendo*** *announced today that they Will release the* ***Nintendo 3DS*** *in* ***north America march 27*** *for* ***$250***"

Information Extraction

"Yess! Yess! Its official ***Nintendo*** *announced today that they Will release the* ***Nintendo 3DS*** *in* ***north America march 27*** *for* ***$250****"*

COMPANY	PRODUCT	DATE	PRICE	REGION

PRODUCT RELEASE

Information Extraction

"Yess! Yess! Its official ***Nintendo*** *announced today that they Will release the* ***Nintendo 3DS*** *in* ***north America*** ***march 27*** *for* ***$250****"*

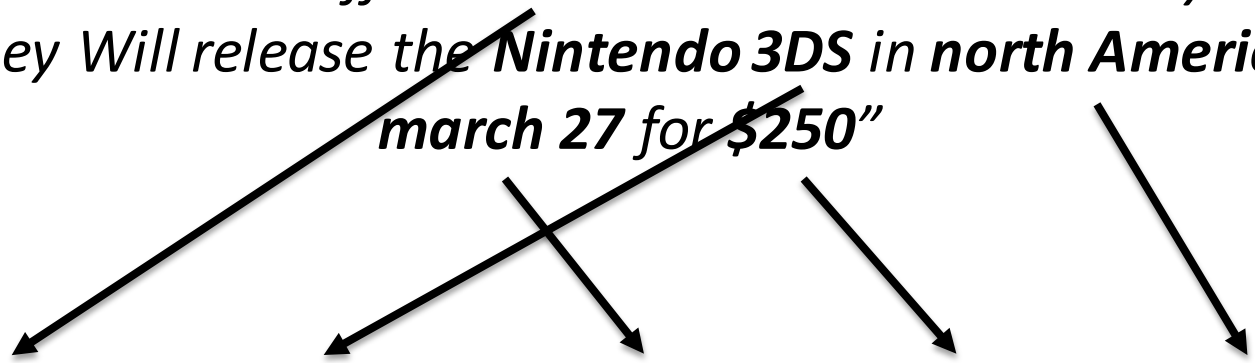

COMPANY	PRODUCT	DATE	PRICE	REGION
Nintendo	3DS	March 27	$250	North America

PRODUCT RELEASE

Information Extraction

Samsung Galaxy S5 *Coming to All Major* ***U.S.*** *Carriers*

- State of the art is maybe 80%, for single easy fields: 90%+
- Redundancy helps a lot!
- Much of human knowledge is waiting to be harvested from the Web!

COMPANY	PRODUCT	DATE	PRICE	REGION
Samsung	Galaxy S5	April 11	?	U.S.
Nintendo	3DS	March 27	$250	North America

PRODUCT RELEASE

NLP as a Field (and subfields)

- Tasks (each have their own datasets and evaluations)
  - Syntactic Parsing
  - Summarization
  - Machine Translation
  - Information Extraction
  - …
- Methods
  - Heuristic / Rule-based Approaches (dominated until the 1990s statistical revolution)
  - Probabilistic Graphical Models
  - Neural Networks
  - …

Q: Why is NLP hard?

- Why is it easy for computers to parse Python, but not English?

A: Ambiguity!

Ambiguity

Example: Some Funny News Headlines

- Milk Drinkers Turn to Powder
- Shouting Match Ends Teachers' Hearing
- Town OKs Animal Rule
- Aging Expert Joins University Faculty
- British Left Waffles on Falkland Islands
- Local High School Dropouts Cut in Half

Layers of Linguistic Annotation

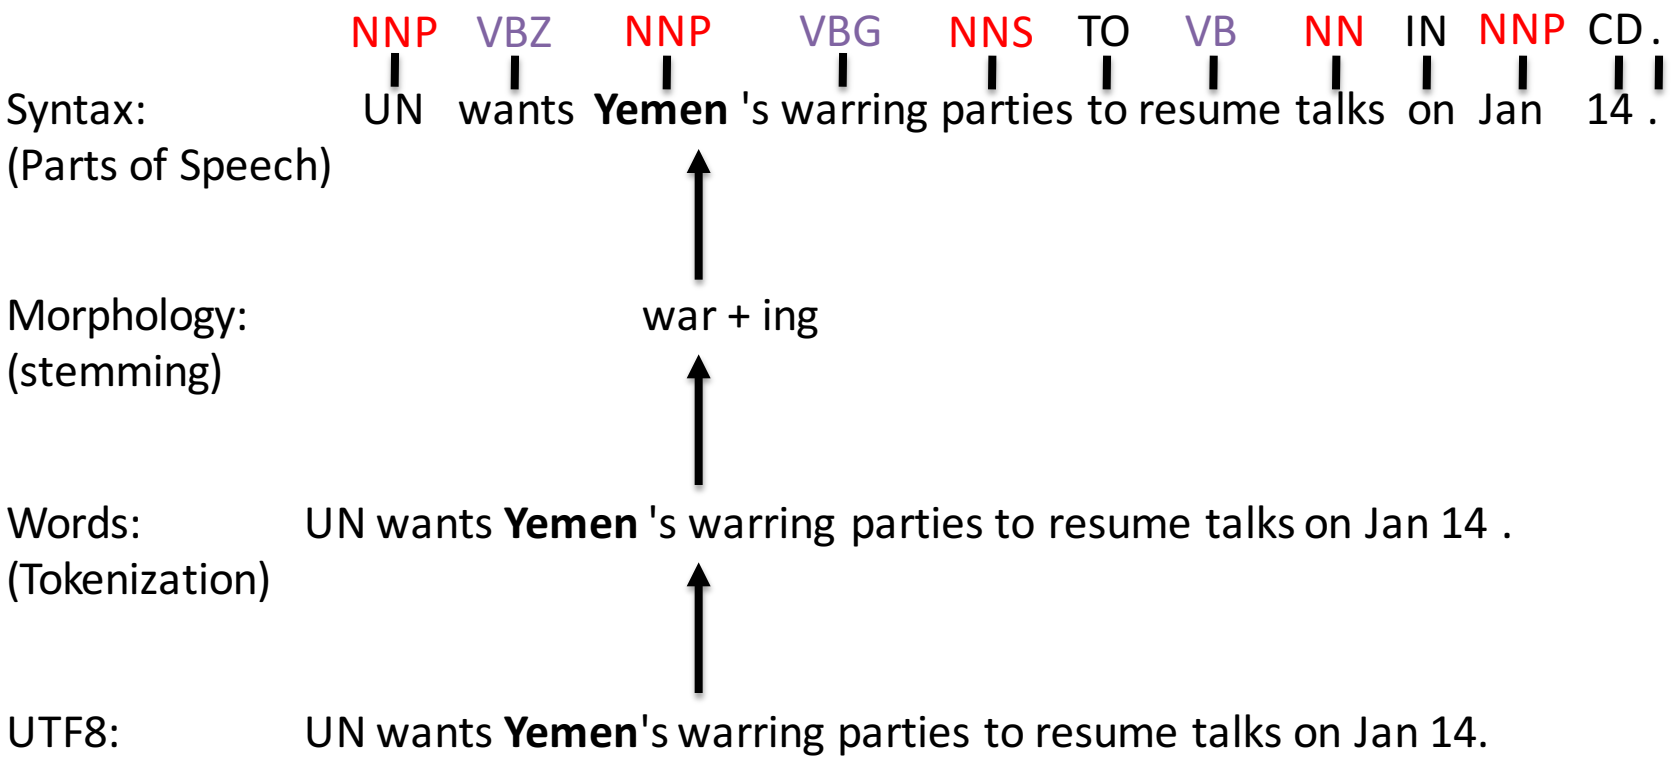

Layers of Linguistic Annotation

Syntax:
(Constituents)

```
(ROOT
  (S
    (NP (NNP UN))
    (VP (VBZ wants)
      (S
        (NP
          (NP (NNP Yemen) (POS 's))
          (VBG warring) (NNS parties))
        (VP (TO to)
          (VP (VB resume)
            (NP (NNS talks))
            (PP (IN on)
              (NP (NNP Jan) (CD 14)))))))
    (. .)))
```

Syntax:
(Parts of Speech)

Layers of Linguistic Annotation

Semantics/
Discourse:

Talks-Event
Between: Yemen's warring parties
Mediator: UN
Date: 1/14/2016

Syntax:
(Constituents)

```
(ROOT
  (S
    (NP (NNP UN))
    (VP (VBZ wants)
      (S
        (NP
          (NP (NNP Yemen) (POS 's))
          (VBG warring) (NNS parties))
        (VP (TO to)
          (VP (VB resume)
            (NP (NNS talks))
            (PP (IN on)
              (NP (NNP Jan) (CD 14)))))))
    (. .)))
```

Goals for This Class

- By the end of the class you should:
  - Be able to build basic NLP tools (translate algorithms into code)
  - Be able to read (and re-implement) current research papers in NLP
  - (hopefully) Be able to notice gaps and propose novel solutions.

Homework 1

- Probability Refresher
- Should be fairly trivial
- Due at the beginning of class on Friday
  - Hand in paper copy