

CS546: Machine Learning in NLP (Spring 2018)

<http://courses.engr.illinois.edu/cs546/>

# Lecture 1

# Introduction/Admin

Julia Hockenmaier

*juliahmr@illinois.edu*

3324 Siebel Center

Office hours: Tue/Thu 2pm-3pm

# Welcome to CS546!

## **Julia Hockenmaier (Instructor)**

[juliahmr@illinois.edu](mailto:juliahmr@illinois.edu)

Office hours: Tuesdays, Thursdays 2pm-3pm, 3324 Siebel

## **Chris Cervantes (TA)**

[ccervan2@illinois.edu](mailto:ccervan2@illinois.edu)

Office hours: Tuesdays, 2:30pm—3:30pm

## **Anjali Narayan-Chen (TA)**

[nrynchn2@illinois.edu](mailto:nrynchn2@illinois.edu)

Office hours: Thursdays: 2:30pm—3:30pm

What will you learn  
in this class?

# CS546: Machine Learning in NLP

## Questions you should be able to answer after CS546:

What Machine Learning (ML) techniques and tools work well for which Natural Language Processing (NLP) tasks?

What are the challenges in applying ML to NLP tasks?

## What we're aiming to cover in CS546 this year:

Focus on neural approaches (“deep learning”) to NLP

Background and current research

Overview of different types of neural models and NLP tasks

## What you need to do in CS546:

Read, present and discuss research paper(s)

Do a research project

# Prerequisites

## CS447 Introduction to NLP (or equivalent)

Basic understanding of NLP tasks and models

## CS446 Machine Learning (or equivalent)

Basic understanding of ML

## Python programming

Most neural network toolkits use it (Tensorflow, Pytorch)

How will we run this  
class?

# This class consists of...

## ... lectures

Tuesdays/Thursdays, 3:30-4:45, SC0216

Many of these will be paper presentations by students

## ... office hours

TA office hours are intended for research projects

My office hours are mainly intended for paper presentations

## ... research projects

These can be done in groups of up to four students

## ... a Compass page

For grades and to submit reports and paper reviews

## ... a Piazza page

For discussions

# Assessment

Your grade will consist of

- ... 35%: your presentation of a research paper in class
- ... 50%: your research project
- ... 10%: your written reviews of research papers  
(graded for completion)
- ... 5%: your participation in class

# Paper presentations

Everybody needs to prepare a **15-minute oral presentation** and a **two-page writeup** about one research paper to be shared with the class.

NB: This paper shouldn't come from your own research group, nor can it be a paper you presented in your qualifying exam.

- We will send out a sign-up sheet with dates and papers for each class.
- You will have to **come to my office hours** *at least two day in advance* with your slides to show them to me, otherwise you will only get half credit for your presentation.
- You have **one week after your presentation to send in your writeup** (so that you can reflect any in-class discussion)

# Paper reviews

For 10 lectures where papers are discussed, you will have to submit a review of one of the papers that was discussed in class.

- Due to the size of the class, we can only grade you for completion.
- You will have to submit a PDF to Compass.
- We will distribute a LaTeX template with a few short questions that you should use for every paper.
- You can choose which papers you want to review.

We encourage you to get into the habit of taking notes about the papers you read. Hopefully this will get you started.

# Research project

You will have to complete a sizable research project.

Due to the size of the class, you will have to work in groups (at least two, but up to 4 students).

There will be several milestones:

- Initial proposal
- Intermediate report and presentation
- Final report and presentation

We will have accounts and GPU hours on BlueWaters for these projects.

# What is NLP?

# NLP in the news...

Facebook AI Creates Its Own Language In Creepy Preview Of Our Potential Future

Computers can now describe images using language you'd understand

## **Barbie Wants to Get to Know Your Child**

With the help of A.I., America's most famous doll tries to fulfill a timeless dream — convincing little girls that she's a real friend. What will happen if they believe her?



# Machine Translation

Google Translate

首页 | 高层 | 时政 | 地方 | 法治 | 廉政 | 人事 | 人才 | 社会 | 文史 | 悦读 | 资料

习近平

新华网 高层 > 正文

最新播报: 全国对口支援西藏先进个人名单(21:49)

## 习近平：我们永远不会忘记曾经风雨同舟、相互理解和支持的老朋友

2014年08月25日 21:33:08 来源：新华网

### 习近平同津巴布韦总统穆加贝举行会谈

#### 强调中津要做平等相待、相互支持、互利共赢、共同发展的好朋友、好伙伴、好兄弟

新华网北京8月25日电（记者谭晶晶）国家主席习近平25日在人民大会堂同津巴布韦总统穆加贝举行会谈。习近平高度评价中津传统友谊及穆加贝为两国关系发展做出的重要贡献，强调中国人民是重情义的，我们永远不会忘记曾经风雨同舟、相互理解和支持的老朋友。中方愿同津方一道，弘扬传统友谊，加强各领域合作，做平等相待、相互支持、互利共赢、共同发展的好朋友、好伙伴、好兄弟。

习近平指出，穆加贝总统是著名的非洲民族解放运动领袖、非洲一体化的重要推动者，也是中国人民的老朋友。中津传统友谊是在我们并肩反帝反殖反霸的光辉岁月中凝结而成的，体现出两国共同遵循的独立自主、相互尊重、反对外来干涉等对外关系的基本原则。建交34年来，双方在涉及彼此核心利益和重大关切问题上相互理解、相互支持，在发展进程中相互帮助。中津传统友谊是两国共同的宝贵财富，我们双方都很珍惜。当前，中津都处在国家建设的重要阶段，两国友好合作正迎来新的机遇。中方愿意同津方一道努力，推动两国关系全面深入发展，更好惠及两国人民。

Senior | Politics | Local | Law | Industry

新华网 Senior > Text

## Xi Jinping: We will never forget once stood, mutual understanding and support of old friends

August 25, 2014 21:33:08 Source: Xinhua

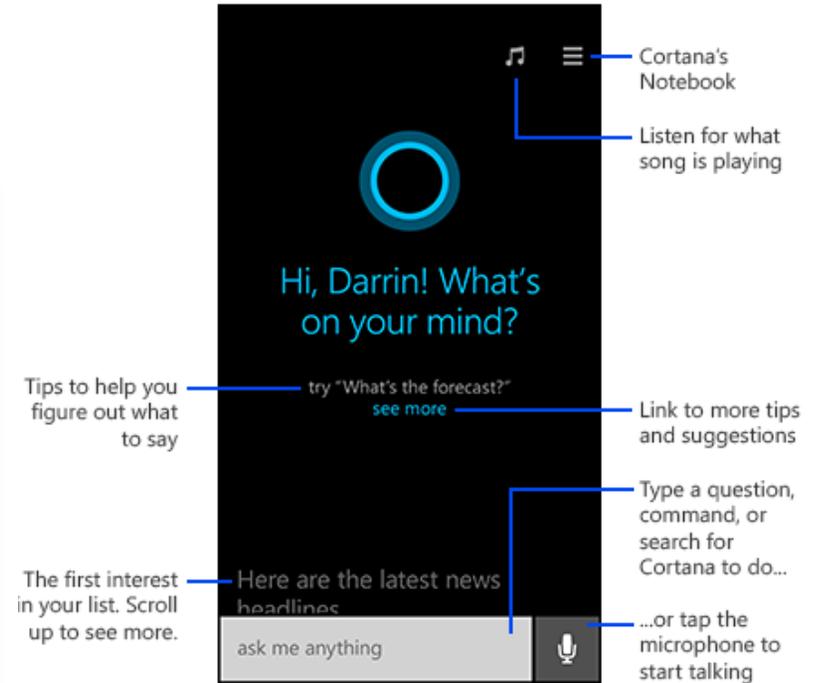
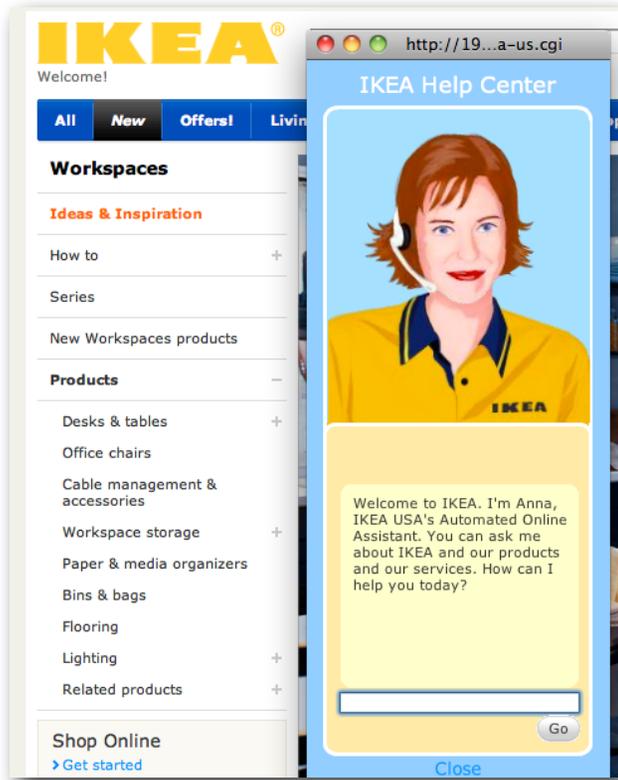
### Xi Jinping held talks with Zimbabwean President Robert Mugabe

#### Nakatsu do emphasize equality, mutual support, mutual benefit and common development, good friends, good partners and good brothers

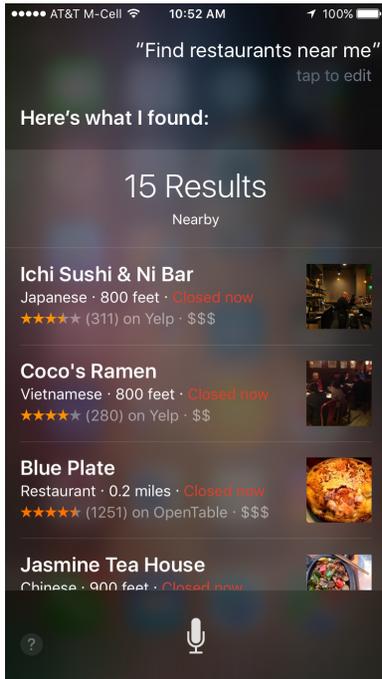
Xinhua Beijing, August 25 (Reporter Tan Jingjing) President Xi Jinping held talks on the 25th at the Great Hall with President Robert Mugabe. Xi spoke highly of the traditional friendship between the important contribution Zimbabwe and Mugabe made for the development of bilateral relations, stressing the friendship of the Chinese people is heavy, we will never forget once stood, mutual understanding and support of old friends. China is willing to Tianjin square one, carry forward the traditional friendship and strengthen cooperation in various fields, doing equal treatment, mutual support, mutual benefit and common development of good friends, good partners and good brothers.

Xi pointed out that President Mugabe is a famous leader of the national liberation movement in Africa, an important promoter of African integration, is also an old friend of the Chinese people. Nakatsu traditional friendship is condensed in our glory days alongside imperialism and colonialism and hegemonism, and reflects both countries followed independence, mutual respect, and oppose foreign interference and other basic principles of foreign relations. Establishment of diplomatic ties 34 years ago, the two sides understand each other on issues concerning each other's core interests and major concerns, mutual support, mutual help in the development process. Nakatsu traditional friendship is the common treasure of the two countries, we both cherish. Currently, Nakatsu are at an important stage of nation-building, friendship and cooperation between the two countries are embracing new opportunities. China is willing to work together with Tianjin efforts to promote comprehensive development of bilateral relations, to better benefit the two peoples.

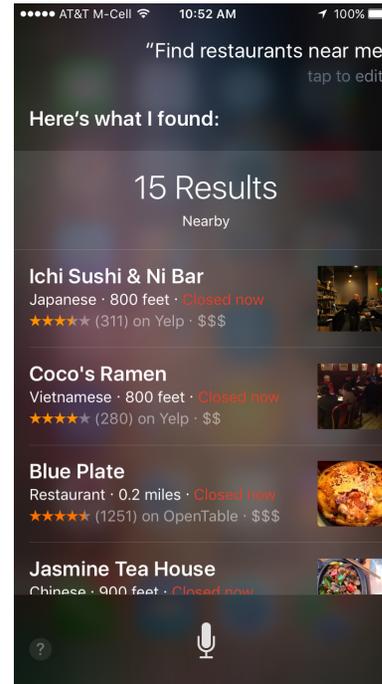
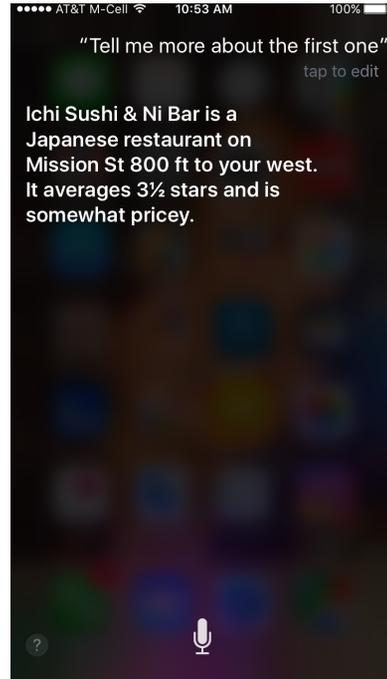
# Dialog systems



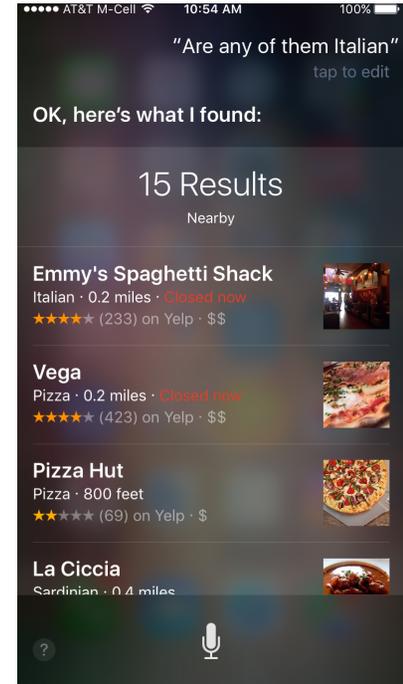
# Siri [Jurafsky & Martin, 3rd ed]



(a)



(b)



# Sentiment/opinion analysis

**Customer Reviews**  
**Haier HLT71 7-Inch Handheld LCD TV** by Haier

**Average Customer Rating**  
★★★★☆ (688 customer reviews)

5 star:	(214)
4 star:	(197)
3 star:	(69)
2 star:	(55)
1 star:	(153)

[Image quality](#) ..... ★★★★★ (359)  
[Portability](#) ..... ★★★★★ (359)  
[Ease of use](#) ..... ★★★★★ (356)  
[Features](#) ..... ★★★★★ (353)

> [See and rate all 15 attributes.](#)

[Create your own review](#)

**The Most Helpful Reviews**

**The most helpful favorable review**

1,085 of 1,127 people found the following review helpful

★★★★★ **FANTASTIC 7 inch portable LCD TV !**

I am really thrilled with my purchase of this Haier HLT71 7-inch, ATSC 2009-compatible portable LCD TV! I have never heard of the name brand Haier, and I had heard of the nightmare stories about some of the name brands having poor pictures on little screens, so I thought I'd give the Haier a try.

I ordered this tv directly from Amazon and Published on November 14, 2009

**Customer Ratings** [Close]

Ease of use	★★★★☆ (356)
Features	★★★★☆ (353)
Remote control	★★★★☆ (348)
Wireless reception	★★★★☆ (329)
Sound quality	★★★★☆ (212)
Product quality	★★★★☆ (209)

> [See what customers say about these attributes](#)

Prompts for sign-in

[Sign in to add ratings and attributes](#)

Source:  
www.amazon.com

# Have chatbots developed their own language?

<https://code.facebook.com/posts/1686672014972296/deal-or-no-deal-training-ai-bots-to-negotiate/>

*Bob: "I can can I I everything else."*

*Alice: "Balls have zero to me to."*

No, not really...

<http://www.snopes.com/facebook-ai-developed-own-language/>

# How well does that Barbie do?

Barbie: *“Do you have any sisters?”*

Child: *“Yeah, I only have one.”*

Barbie: *“What’s something nice that your sister does for you?”*

Child: *“She does nothing nice to me”*

Barbie: *“Well, what is the last nice thing your sister did?”*

Child: *“She helped me with my project  
— and then she destroyed it.”*

Barbie: *“Oh, yeah, tell me more!”*

Child: *“That’s it, Barbie,”*

Barbie: *“Have you told your sister lately how cool she is?”*

Child: *“No. She is not cool,”*

Barbie: *“You never know, she might appreciate hearing it”*

(Thanks to Barbara Grosz for pointing out the example from the NYT)

<https://www.nytimes.com/2015/09/20/magazine/barbie-wants-to-get-to-know-your-child.html>

# What is the current state of NLP?

Lots of commercial applications and interest.

Some applications are working pretty well already, others not so much.

A lot of hype around “deep learning” and “AI”

- Neural nets are powerful classifiers and sequence models
- Public libraries (Tensorflow, Torch, Caffe, etc.) and datasets make it easy for anybody to get a model up and running
- “End-to-end” models put into question whether we still need the traditional NLP pipeline that this class is built around
- We’re still in the middle of this paradigm shift
- But many of the fundamental problems haven’t gone away

# Examples of NLP applications

(What can NLP be used for?)

## Natural language (and speech) interfaces

Search/IR, database access

Dialog systems (e.g. customer service, robots, cars, tutoring), chatbots

Image description/search

## Information extraction, summarization, translation

Process (large amounts of) text automatically to obtain meaning/knowledge contained in the text

Translate text automatically from one language to another

## Convenience, social science

Grammar/style checking, automate email filing, autograding

Identify/analyze trends, opinions, etc. (e.g. in social media)

# Examples of NLP tasks

(What capabilities do NLP systems need?)

## Natural language understanding

- Extract information (e.g. about entities or events) from text

- Translate raw text into a meaning representation

- Reason about information given in text

- Execute NL instructions

## Natural language generation and summarization

- Translate database entries or meaning representations to raw natural language text

- Produce (appropriate) utterances/responses in a dialog

- Summarize (newspaper or scientific) articles, describe images

## Natural language translation

- Translate one natural language to another

# The NLP (NLU) Pipeline

An NLP system may use some or all of the following steps:

## Tokenizer/Segmenter

to identify words and sentences

## Morphological analyzer/POS-tagger

to identify the part of speech and structure of words

## Word sense disambiguation

to identify the meaning of words

## Syntactic/semantic Parser

to obtain the structure and meaning of sentences

## Coreference resolution/discourse model

to keep track of the various entities and events mentioned

# NLP research questions redux

## How do you represent (or predict) words?

Do you treat words in the input as atomic categories, as continuous vectors, or as structured objects?

How do you handle rare/unseen words, typos, spelling variants, morphological information?

Lexical semantics: do you capture word meanings/senses?

## How do you represent (or predict) word sequences?

Sequences = sentences, paragraphs, documents, dialogs,...

As a vector, or as a structured object?

## How do you represent (or predict) structures?

Structures = labeled sequences, trees, graphs, formal languages (e.g. DB records/queries, logical representations)

How do you represent “meaning”?