



# Introduction

COS IW 02: Natural Language Processing with Neural Networks

Karthik Narasimhan

# Introductions

**Instructor:**



Karthik Narasimhan

**TA:**



Xi Chen

**Time/Location:**

- Tuesdays 1:30pm-2:50pm in CS 301
- Office hours:
  - Karthik Narasimhan: Mondays 4pm-5pm, Computer Science 422
  - Xi Chen: Fridays 1:30 - 3:30pm, CS 003

# Logistics

- Not typical lecture-style. This is independent work after all!
- Attendance is mandatory
- Please use Piazza for all questions and to share useful information with each other!
- Office hours will be updated on class website
- ‘Intro to Neural Networks’ tutorial by Xi (date: TBD)


# The semester in phases



P1: Brainstorm, literature review, finalize ideas  
(Feb 18)



P2: Initial setup (data collection, tool survey, etc)  
(Feb 25)



P3: Midpoint (preliminary) results  
(Mar 10)




P4: Endgame (finalize results, write report, presentations)  
(end of April)



# Important IW milestones

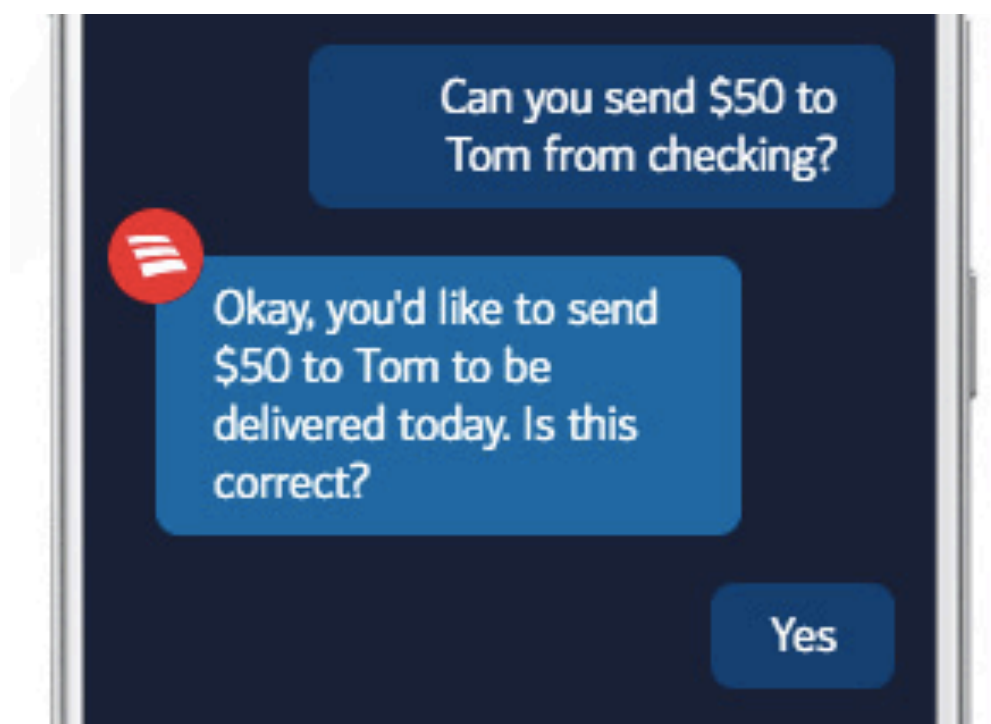
## Spring 2020 Single Semester Projects Schedule

Feb. 6, 2020	<a href="#">Attend the "Getting Started" Information Meeting</a> 4:30p.m. in CS 105	
Feb. 17, 2020 5:00pm	<a href="#">SEAS Funding Deadline</a> <b>*Due in person by 5pm.</b> <b>**Open to ANY COS IW student</b> <b>(Both AB and BSE can apply. Both single-term and two-term/ thesis can apply)</b>	 <b>Apply to this!</b>
Feb. 23, 2020	<a href="#">Submit a Written Project Proposal</a> , by 11:59p.m.	
March 8, 2020	<a href="#">Submit the Checkpoint Form</a> , by 11:59p.m.	
March 30, 2020 - April 3, 2020	<a href="#">Sign Up to Give an Oral Presentation</a> <b>**First-time BSE IW students, only</b> <b>Note for seminar students:</b> All seminar students will give an Your instructor will share information about logistics.	
March 31, 2020	<a href="#">Attend "How to Give an IW Talk"</a> 4:30p.m. in CS 105	
April 14, 2020	<a href="#">Attend "How to Write an IW Paper"</a> 4:30p.m. in CS 105	
April 19, 2020	<a href="#">Submit Slides for an Oral Presentation</a> , by 11:59p.m. <b>**All seminar students and first-time BSE IW students, only</b>	
April 20-24, 2020	<a href="#">Give an IW Oral Presentation</a> <b>**First-time BSE IW students, only</b> <b>Note for seminar students:</b> All seminar students will give an oral presentation. Your instructor will share information about logistics.	
May 1, 2020	<a href="#">Submit a Written Final Report</a> , by 11:59p.m.	

# Natural Language Processing




- Making machines understand human language
- Communication with humans (ex. personal assistants, customer service)
- Use language to understand the world



Banking assistant




# Turing Test



Ability to understand and generate language ~ intelligence


# Deep Learning era

- Significant advances in core NLP technologies



# Deep Learning era

- Significant advances in core NLP technologies
- **Essential ingredient:** large-scale supervision, lots of compute
- Reduced manual effort - less/zero feature engineering



36M sentence pairs

*Russian:* Машинный перевод - это круто!



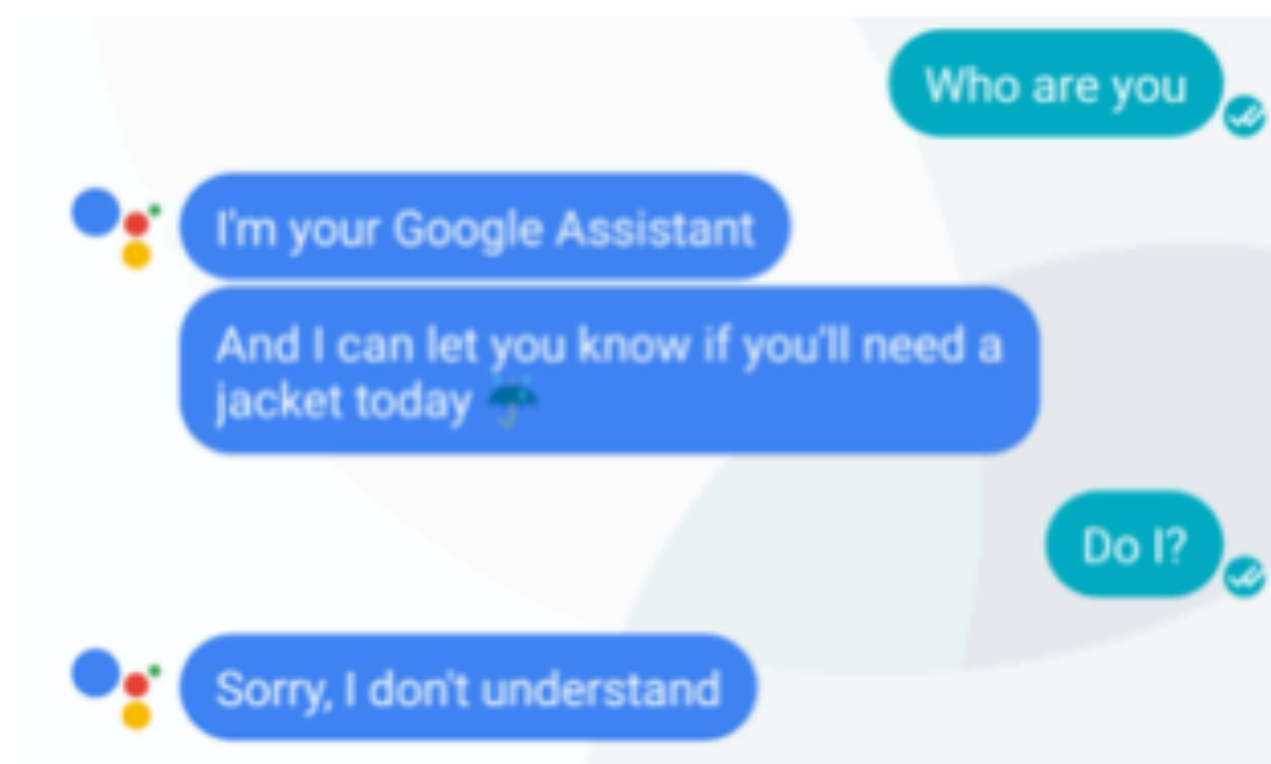
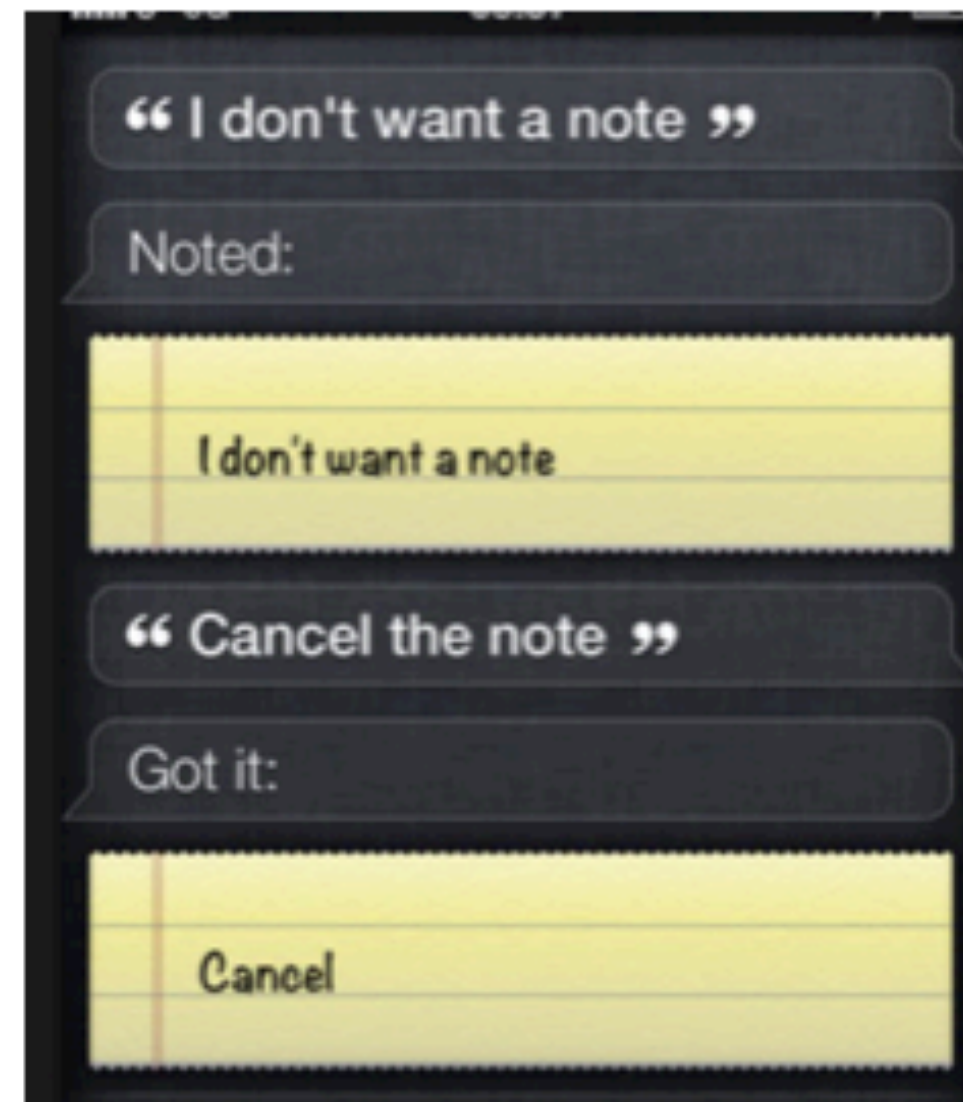

*English:* Machine translation is cool!



# Turing test solved?

**Talking to Google Duplex: Google's human-like phone AI feels revolutionary**

Believe the hype—Google's phone-call bot is every bit as impressive as promised.



... maybe not.

# Some language humor

Kids make nutritious snacks

Stolen painting found by tree

Miners refuse to work after death

Squad helps dog bite victim

Killer sentenced to die for second time in 10 years

Lack of brains hinders research

Real newspaper headlines!



# Lexical ambiguity

The fisherman went to the *bank*.

bank<sup>1</sup>

/baNGk/ 

*noun*

plural noun: banks

1. the land alongside or sloping down to a river or lake.  
"willows lined the bank"  
*synonyms:* [edge](#), [side](#), [shore](#), [coast](#), [embankment](#), [bankside](#), [levee](#), [border](#), [verge](#), [boundary](#),  
[margin](#), [rim](#), [fringe](#); [More](#)
1. a financial establishment that invests money deposited by customers, pays it out when required, makes loans at interest, and exchanges currency.  
"I paid the money straight into my bank"  
*synonyms:* [financial institution](#), [merchant bank](#), [savings bank](#), [finance company](#), [trust company](#),

One word can mean several different things

# Lexical variations




**ACCORDING TO THE THESAURUS,  
"THEY'RE HUMID, PREPOSSESSING  
HOMOSAPIENS WITH FULL SIZED AORTIC  
PUMPS" MEANS "THEY'RE WARM, NICE  
PEOPLE WITH BIG HEARTS."**

Several words can mean the same thing!

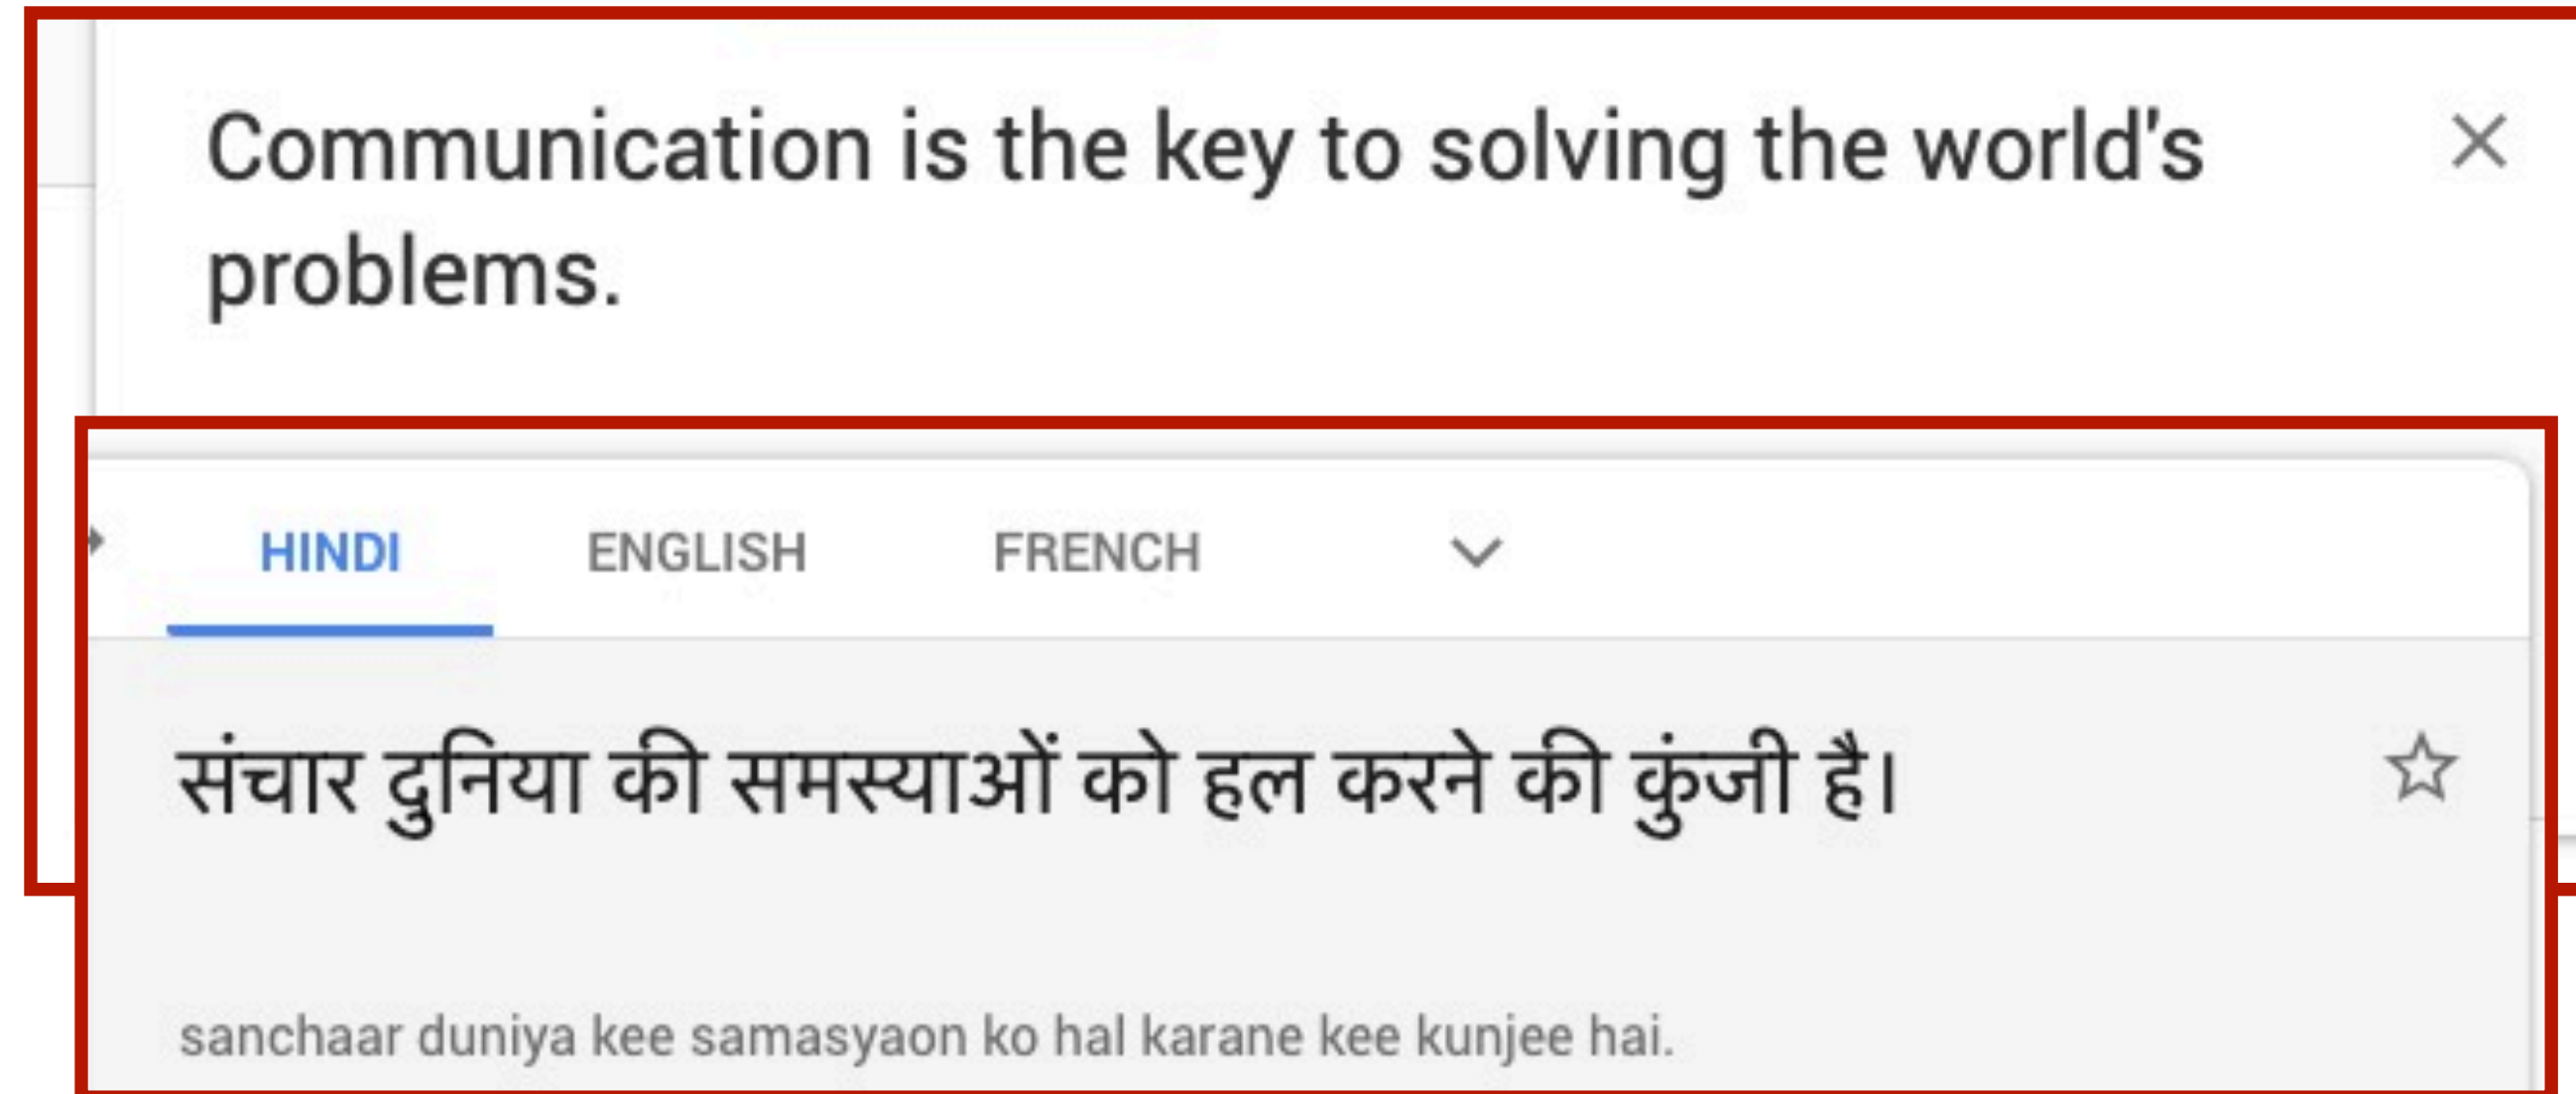


# Challenges in modern NLP

- *Scale*: Large number of phenomena
- *Sparsity*: Text data is often heavy-tailed



# Machine Translation




- One of the “holy grail” problems in artificial intelligence
- Practical use case: Facilitate communication between people in the world
- Extremely challenging (especially for low-resource languages)

# Information Extraction

*The Massachusetts Institute of Technology (MIT) is a private research university in Cambridge, Massachusetts, often cited as one of the world's most prestigious universities.  
Founded in 1861 in response to the increasing industrialization of the United States, ...*


Article



**City:** Cambridge, MA  
**Founded:** 1861  
**Mascot:** Tim the Beaver  
...

Database

# Language generation



## Better Language Models and Their Implications

We've trained a large-scale unsupervised language model which generates coherent paragraphs of text, achieves state-of-the-art performance on many language modeling benchmarks, and performs rudimentary reading comprehension, machine translation, question answering, and summarization—all without task-specific training.

**With the start of the new academic year, Princeton** has an opportunity to help provide a new generation of women with a diverse set of academic resources for higher education. We are offering the resources of the Princeton-McGill program specifically to women with undergraduate degrees who would like to enhance their academic experience. Princeton-McGill offers a comprehensive suite of services for women and their families including a variety of graduate programs, support programs, and the opportunity to serve as leaders in their communities with a wide variety of programs, activities and services. For the upcoming fall, Princeton-McGill will also offer its Women's Center , which is located in a renovated women's dorm. At Princeton, we are working with the Princeton-McGill community to develop a suite of programs that are designed to give new and returning students a strong foundation for a successful, rewarding graduate career. The Women's Center , the Princeton-McGill Women's Center provides a range of supports to address the specific needs of female doctoral degree graduates. Programs are tailored to meet the unique needs of women under the age of 28, women and families



# Textual Question Answering

James the Turtle was always getting in trouble. Sometimes he'd reach into the freezer and empty out all the food. Other times he'd sled on the deck and get a splinter. His aunt Jane tried as hard as she could to keep him out of trouble, but he was sneaky and got into lots of trouble behind her back.

One day, James thought he would go into town and see what kind of trouble he could get into. He went to the grocery store and pulled all the pudding off the shelves and ate two jars. Then he walked to the fast food restaurant and ordered 15 bags of fries. He didn't pay, and instead headed home.

His aunt was waiting for him in his room. She told James that she loved him, but he would have to start acting like a well-behaved turtle.

After about a month, and after getting into lots of trouble, James finally made up his mind to be a better turtle.

1) What is the name of the trouble making turtle?

- A) Fries
- B) Pudding
- C) James
- D) Jane

2) What did James pull off of the shelves in the grocery store?

- A) pudding
- B) fries
- C) food
- D) splinters



# Visual Question Answering



What color are her eyes?  
What is the mustache made of?



How many slices of pizza are there?  
Is this a vegetarian pizza?

# Task-Oriented Dialogue System

GUS: Hello. My name is GUS. I can help you plan a simple trip by air. Where do you want to go ?  
Client: I want to go to San Diego on May 28  
GUS: What time do you want to leave?  
Client: I must be in San Diego before 10 am  
GUS: Would you like P.S.A. flight 102 that arrives at 9.15 am ?  
Client: What is the next flight ?  
GUS: Air California flight 310 that leaves at 8.30am and arrives at 10.00am  
Client: I'll take the first one  
GUS: In what name should I make a reservation, please?  
Client: Dan Bobrow  
GUS: I have confirmed the following flight: P.S.A. flight 102 on Wednesday May 28 from San Jose to San Diego that leaves at 7.30 am and arrives at 9.15 am.  
GUS: What date do you want to return on ?  
Client: On Friday in the evening.  
GUS: Would you like the flight that leaves at 7.45 pm ?  
Client: That's fine.  
GUS: I have confirmed the following flight: P.S.A. flight 307 on Friday May 30 from San Diego to San Jose that leaves at 7.45 pm and arrives at 9.30 pm Thank you for calling. Goodbye

## GOAL

get information from the user to help **complete the specific task.**

## Task-Oriented Dialog System (Travel):

A transcript of an actual dialog with the GUS system of *Bobrow et al.* (1977)

P.S.A. and Air California were airlines of that period.



# Some project ideas

## Tasks

We are pleased to announce the following tasks in SemEval-2020.

### Lexical semantics

- Task 1: Unsupervised Lexical Semantic Change Detection [[mailing list](#)] [[email organizers](#)]
- Task 2: Predicting Multilingual and Cross-Lingual (Graded) Lexical Entailment [[mailing list](#)] [[email organizers](#)]
- Task 3: Graded Word Similarity in Context (GWSC) [[discussion forum](#)] [[mailing list](#)] [[email organizers](#)]

### Common Sense Knowledge and Reasoning, Knowledge Extraction

- Task 4: Commonsense Validation and Explanation [[mailing list](#)] [[email organizers](#)]
- Task 5: Modelling Causal Reasoning in Language: Detecting Counterfactuals [[mailing list](#)] [[email organizers](#)]
- Task 6: DeftEval: Extracting Definitions from Free Text in Textbooks [[mailing list](#)] [[email organizers](#)]

### Humour, Emphasis, and Sentiment

- Task 7: Assessing Humor in Edited News Headlines [[mailing list](#)] [[email organizers](#)]
- Task 8: Memotion Analysis [[mailing list](#)] [[email organizers](#)]
- Task 9: Sentiment Analysis for Code-Mixed Social Media Text [[mailing list](#)] [[email organizers](#)]
- Task 10: Emphasis Selection for Written Text in Visual Media [[mailing list](#)] [[email organizers](#)]

### Societal Applications of NLP

- Task 11: Detection of Propaganda Techniques in News Articles [[mailing list](#)] [[email organizers](#)]
- Task 12: OffenseEval 2: Multilingual Offensive Language Identification in Social Media [[mailing list](#)] [[email organizers](#)]

**SemEval-2020**

International Workshop on Semantic Evaluation

# Some project ideas



Exploring how artificial intelligence technologies could be leveraged to combat fake news.

# Fake news detection

## EXAMPLE HEADLINE

## "Robert Plant Ripped up \$800M Led Zeppelin Reunion Contract"

### EXAMPLE SNIPPETS FROM BODY TEXTS AND CORRECT CLASSIFICATIONS

*"... Led Zeppelin's Robert Plant turned down £500 MILLION to reform supergroup. ..."*

**CORRECT CLASSIFICATION: AGREE**

"... No, Robert Plant did not rip up an \$800 million deal to get Led Zeppelin back together."  
"  
..."

**CORRECT CLASSIFICATION: DISAGREE**

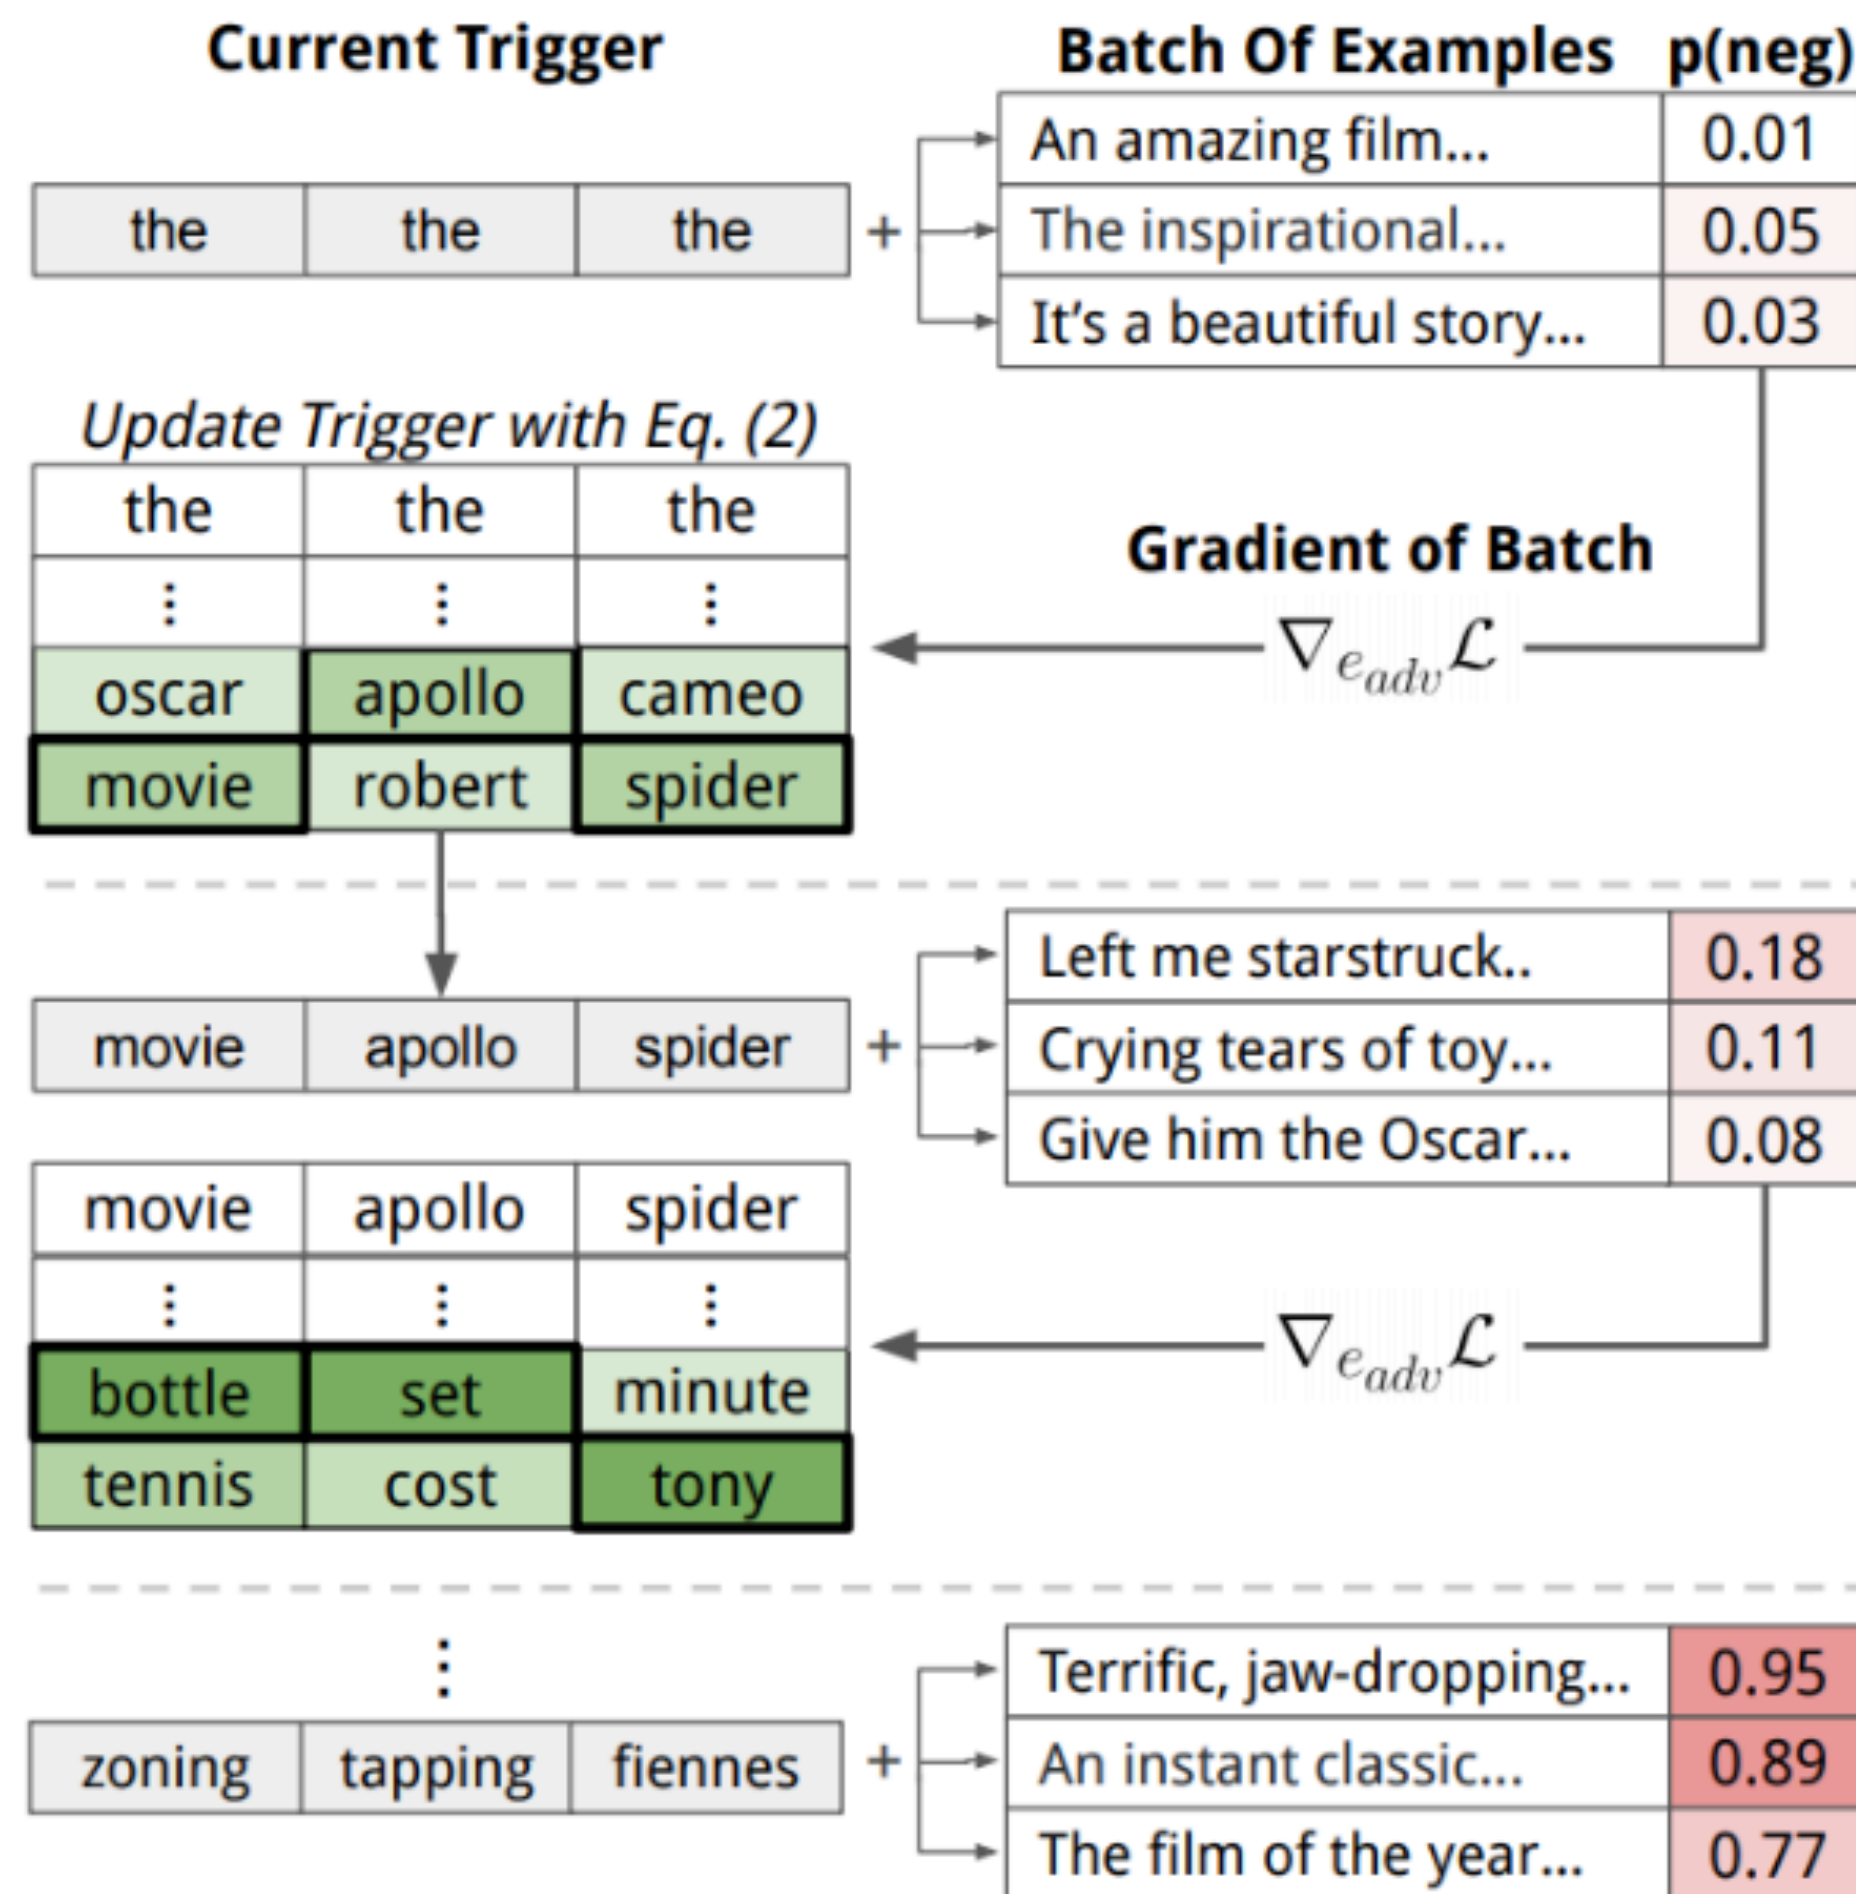
*"... Robert Plant reportedly tore up an \$800 million Led Zeppelin reunion deal. ..."*

**CORRECT CLASSIFICATION: DISCUSSES**

*"... Richard Branson's Virgin Galactic is set to launch SpaceShipTwo today. ..."*

**CORRECT CLASSIFICATION: UNRELATED**

# Some project ideas



- NLP models can be susceptible to adversarial attacks
- What types of attacks work well?
- Are there any defenses?



# Some project ideas

- Develop and study new algorithms for NLP! Challenges to tackle include:
  - **Sample efficiency:** How efficiently can you learn models with as less supervised data as possible?
  - **Ability to generalize:** Can a model trained on one domain (e.g. news text) generalize well to another (e.g. biomedical text)?
  - **Multi-task learning:** Can you train a single model to be good at several NLP tasks/domains?
  - **Robustness:** Is your NLP model robust to noise in the data? Analyze what breaks and why. Even better, try to fix it! (Build it Break it: The Language Edition)
  - **Fairness/Bias:** Can NLP algorithms be biased/unfair? How can we characterize this? How can we fix this?

If you have other ideas, come talk to us!

# Brainstorming

- Write down two different ideas and pitch them

# For next meeting

- Please fill out Google form with initial ideas before the next meeting (will be announced on Piazza later tonight)