COMP 790.139 (Fall 2016)
Natural Language Processing
(with some vision, robotics, and deep learning)

Mohit Bansal

Class Info/Logistics

- COMP 790.139 ‘Natural Language Processing’
- **3 UNITS**
- Instructor: Mohit Bansal (SN258, http://www.cs.unc.edu/~mbansal/)
- Time: Wed 10.10am-12.40pm
- Room: FB008
- Office Hours: Wed 12.40-1.40pm (by appointment), SN258
- Course Webpage: http://www.cs.unc.edu/~mbansal/teaching/nlp-seminar-fall16.html
- **Course Email**: nlpcomp790unc@gmail.com

About Me

- Asst. Professor, CS, UNC (joined Fall 2016)
- Research Asst. Professor, TTI-Chicago, 2013-2016
- PhD, UC Berkeley, 2008-2013

- Research Interests:
  - Past: Syntactic parsing, coreference resolution, taxonomy induction, world knowledge and commonsense induction
  - Current: Multimodal and embodied semantics (i.e., language with vision and speech, for robotics); human-like language generation and Q&A/dialogue; interpretable and structured deep learning

- Office SN258

- Webpage: http://www.cs.unc.edu/~mbansal/, Email: mbansal@cs.unc.edu

Your Introductions

- Please say your:
  - Name
  - Department/degree/major
  - Research interests (ML/AI/NLP/CV experience? Coding?)
  - Fun fact ☺

(send me a few lines of intro at nlpcomp790unc@gmail.com)

About the Course (and its Goals)

- Research-oriented seminar course! We will read lots of interesting papers, brainstorm, and do fun novel projects!
- We'll start with some basics of NLP
- Then cover some specific, latest research topics via several paper readings
- E.g., we will discuss connections of NLP with vision and robotics, and several deep learning for NLP models
- Brainstorm regularly and code + write up fun/novel projects!
- Some lecture(s) on academic/research quality paper writing
- No NLP background needed but some ML and coding experience highly recommended!

Expectations/Grading

- Paper presentation (20%)
- Paper written summaries (15%)
- Class participation, discussion and brainstorming (25%)
- Project reports and presentations (40%)

Paper Presentation

- Lead discussion for 1-2 papers on a topic some week (may be done in pairs/groups depending on class size)
- Read related papers and present background to audience
- Present task and model details of given paper
- Present demo's of related code, etc.
- Ask interesting questions to initiate brainstorming
- Mention some next steps, future work, extension ideas!

Paper Written Summaries

- 0.5-1 page (per paper) write-up for every week's paper(s)
- Describe the task
- Summarize the method
- Explain the novelty
- Discuss the next steps or potential improvements

Class Participation and Brainstorming

- Audience students expected to take part in lively discussion every paper reading!
- After every topic gets completed (i.e., several papers in 2-3 weeks), we will have a brainstorming and 'idea-generation' session!
- Exact details to be announced soon but students expected to submit and discuss novel idea(s) on the whole general topic, e.g., new related task or dataset, new approach to existing task, combinations of tasks/approaches, etc.
- Don't hesitate to propose fancy ideas ☺, but try to keep them grounded/feasible and think of how to approach them realistically (in terms of datasets, models, speed, memory, etc.)

Project

- This is a ‘Reading, Coding, and Writing’ class!
- Students will pick (early) their favorite topic among latest cutting-edge research topics covered in class
- And will try a novel idea (implementing+extending or original) -- I am happy to discuss details!
- Midterm and final report + presentation (and possibly some updates)
- Might be in pairs/groups depending on final class size
- Use ACL conference style files and aim for conference-quality papers
- Will have some lecture(s) on research-quality paper writing

Reference Books

- **SLP2**: D. Jurafsky & James H. Martin. “Speech and Language Processing: An Introduction to Natural Language Processing, Computational Linguistics and Speech Recognition”. Prentice Hall, Second Edition, 2009.
- **SLP3**: Some draft chapters of the third edition are available online at https://web.stanford.edu/~jurafsky/slp3/
- **FSNLP**: Chris Manning and Hinrich Schütze, Foundations of Statistical Natural Language Processing, MIT Press. Cambridge, MA: May 1999. http://nlp.stanford.edu/fsnlp/
- **ML Background**: Andrew Ng’s Coursera Machine Learning course https://www.coursera.org/learn/machine-learning

Course Syllabus/Topics (tentative)

- NLP Basics, Foundations, and Core Tasks (in brief):
  - Language Modeling
  - Part-of-speech tagging
  - Syntactic parsing: Constituent, Dependency, CCG, others
  - Coreference Resolution
  - Distributional Semantics: PMI, NNs, CCA
  - Compositional Semantics: Logical-form, Semantic Parsing, Vector-form
  - Question Answering: Factoid-based, Passage-based
  - Sentiment Analysis
  - Generation: Summarization, Dialogue Models
  - Machine Translation
- Human-like Language Understanding and Generation:
  - Ambiguities: Attachment ambiguities, Coreference ambiguities
  - Non-literal: Metaphors/Idioms, Politeness, Sarcasm, Humor
  - Generating Non-literal/Ambiguous Language, Coherent and Intelligent Dialogue
- Language and Vision: Image-Text Alignment, Language Disambiguation via Images, Image/Video Captioning, Image/Video Question Answering, Text to Image Generation, Visual Story Entailment
- Language for Robotics: Instructions for Navigation, Articulation, Manipulation, Skill learning
- Several interesting machine/deep learning models along the way, e.g., deep+structured models, interpretable models, adversarial models, reward-based models (reinforcement learning)

What is NLP?

- Question answering

What is NLP?

- Question answering

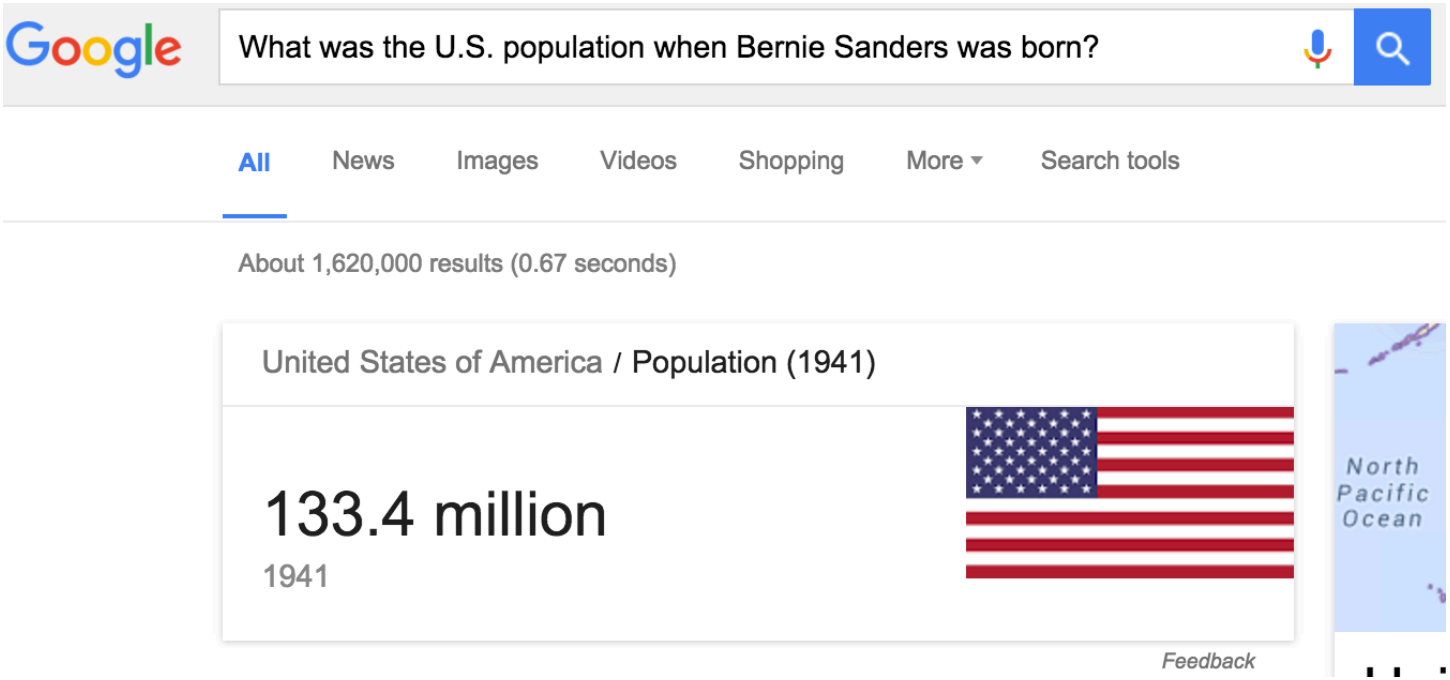

What is NLP?

- Question answering

What is NLP?

- Machine Translation

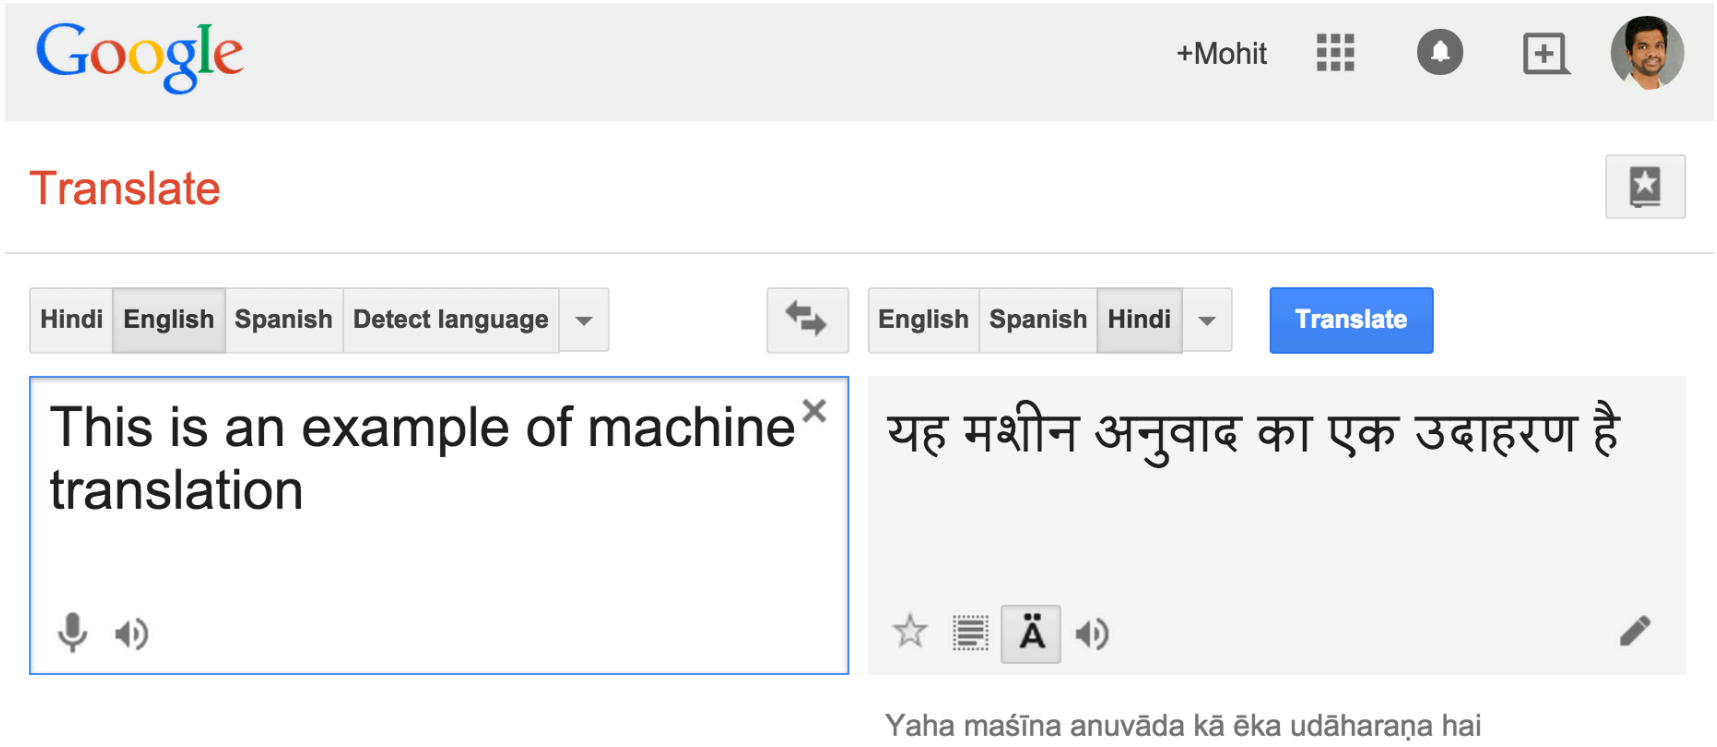

What is NLP?

- Sentiment Analysis

Sentiment Analysis with Python NLTK Text Classification

This is a demonstration of **sentiment analysis** using a `NLTK 2.0.4` powered **text classification** process. It can tell you whether it thinks the text you enter below expresses **positive sentiment**, **negative sentiment**, or if it's **neutral**. Using **hierarchical classification**, *neutrality* is determined first, and *sentiment polarity* is determined second, but only if the text is not neutral.

Analyze Sentiment

Language

english

Enter text

```
It always amazes me how Universal never
cares to create anything remotely clever
when it comes to their animations, and so
once again they come up with a harmless
little story that wants to be cute and
funny (which it is sometimes) but is only
bound to be quickly forgotten.
```

Enter up to 50000 characters

Analyze

Sentiment Analysis Results

The text is **neg**.

The final sentiment is determined by looking at the classification probabilities below.

Subjectivity

- neutral: 0.3
- **polar: 0.7**

Polarity

- pos: 0.2
- **neg: 0.8**

What is NLP?

- Natural Language Generation: Summarization

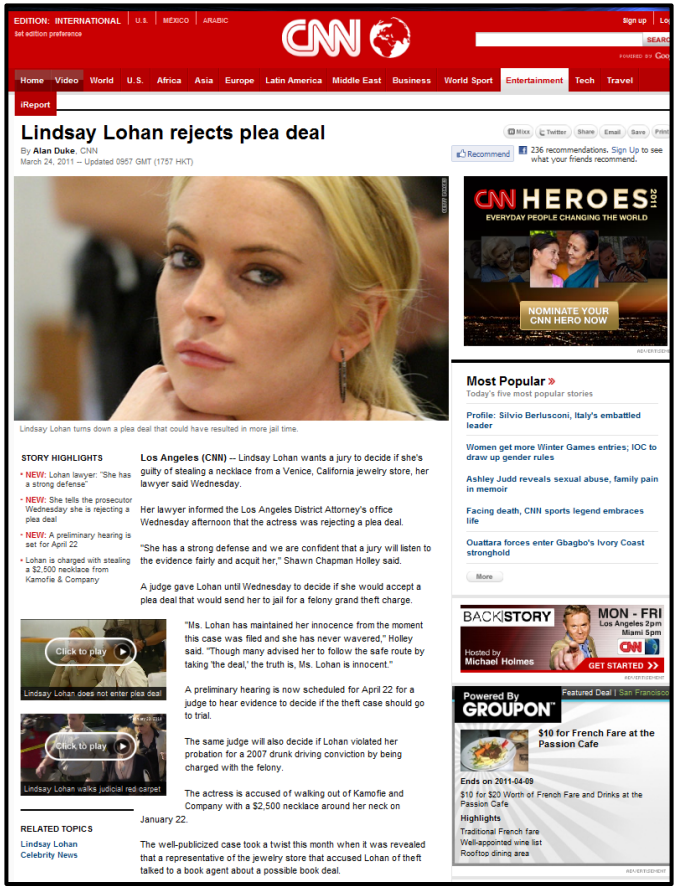

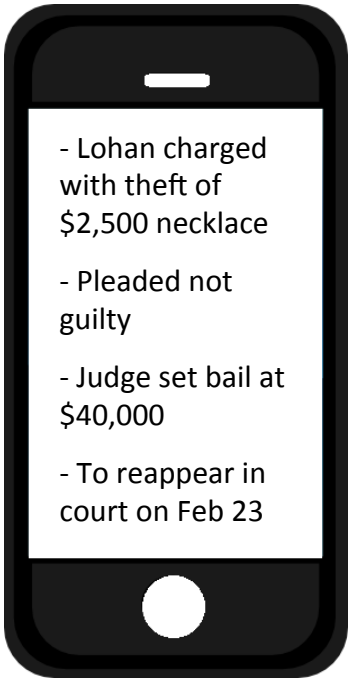

What is NLP?

- Natural Language Generation: Conversation/Dialogue

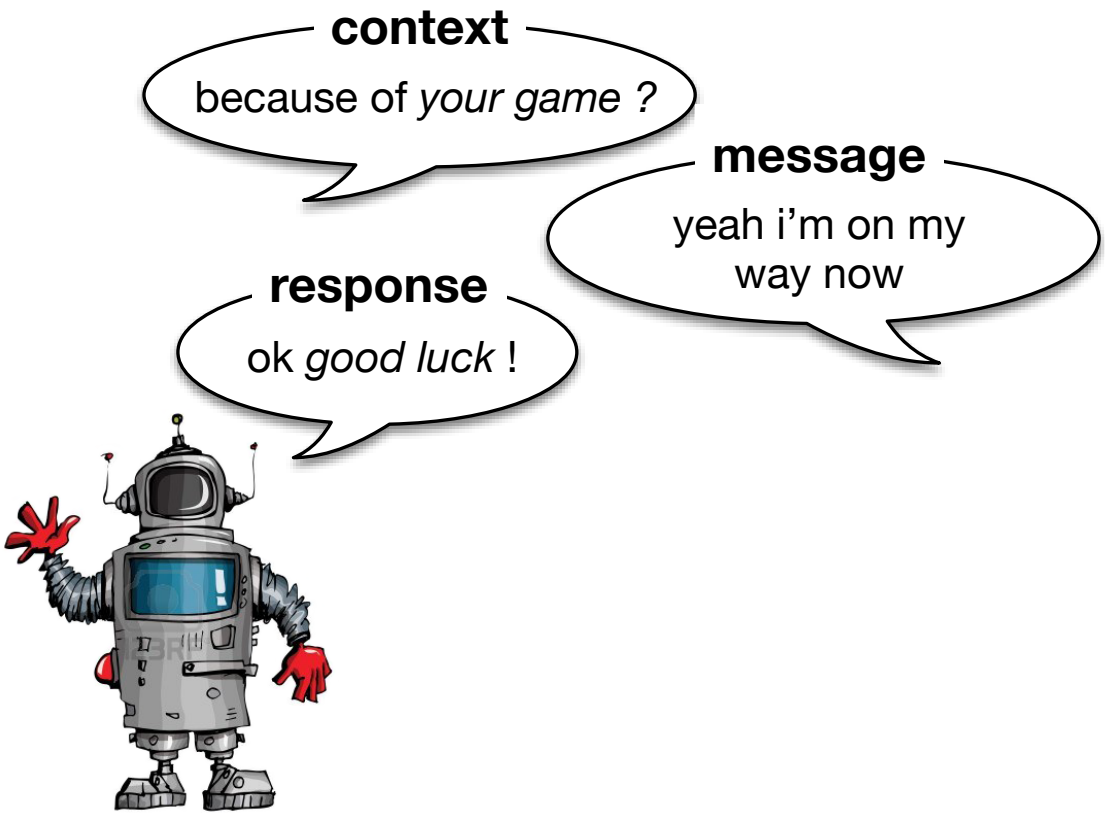

[Sordoni et al., 2015]

What is NLP?

- Natural Language Generation: Image Captioning

www.cs.toronto.edu/~nitish/nips2014demo/results/8832804.html

Results

Tags

- authors
- scones
- luncheon
- breakfast
- seder

Nearest Caption in the Training Dataset

a man cuts a cake while children sit around at the table , looking on .

Generated Captions

- two people at a table with a cake .
- the two people are having a meal that is in a party .
- a man and two children in a blue table with a cake .
- a man sitting at a table with a bunch of cake on it .
- a man and woman sitting at a table with cake at a party .

back

[UToronto]

What is NLP?

- Natural Language Generation: 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?

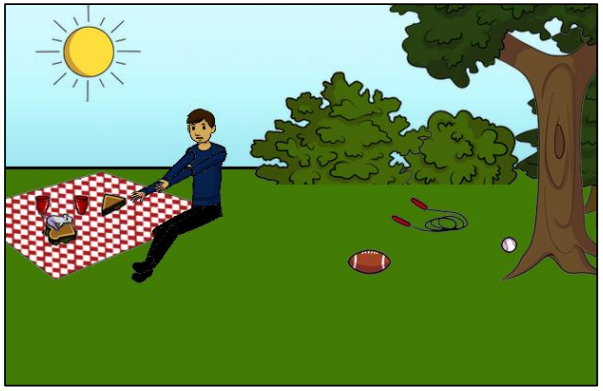

Is this person expecting company?
What is just under the tree?

Does it appear to be rainy?
Does this person have 20/20 vision?

[Antol et al., 2015]

What is NLP?

- Automatic Speech Recognition

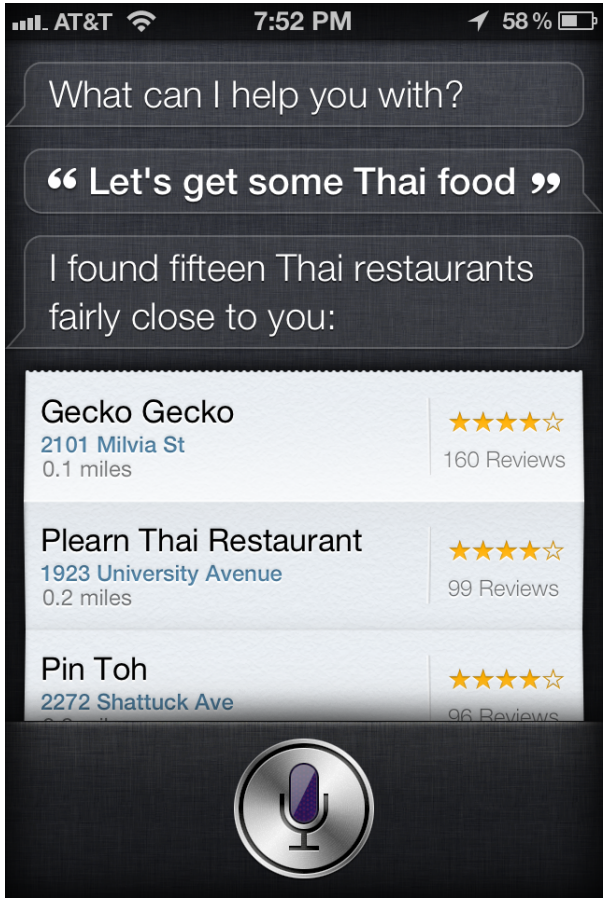

10-min break?

Some Exciting NLP Challenges

1) Human-like Language Understanding: metaphors/idioms, humor, sarcasm, politeness/rudeness

2) Language Generation and Dialogue: document summarization, database to language summary, coherent and intelligent conversation models

3) Grounded Language with Vision and Speech: image-text alignment, language disambiguation via images, image/video captioning, image/video question answering, text to image generation, visual story entailment

4) Embodied Language for Robotic Tasks: instructions for navigation, articulation, manipulation, skill learning

5) Machine Learning Models: deep+structured models, interpretable models, adversarial models, reward-based models (reinforcement learning)

Human-like Ambiguous Language

- Non-literal: Idioms, Metaphors

You: *I am under the weather today.*
Siri: *Here is the weather today…50 F*

Human-like Ambiguous Language

- Non-literal: Idioms, Metaphors

Human-like Ambiguous Language

- Humor, Sarcasm, Politeness/Rudeness

Human-like Ambiguous Language

- Prepositional Attachment, Coreference Ambiguities

Human-like Ambiguous Language

- Prepositional Attachment, Coreference Ambiguities

Visually Grounded Language

- Text-Image Alignment: Most of our daily communication language points to several objects in the visual world

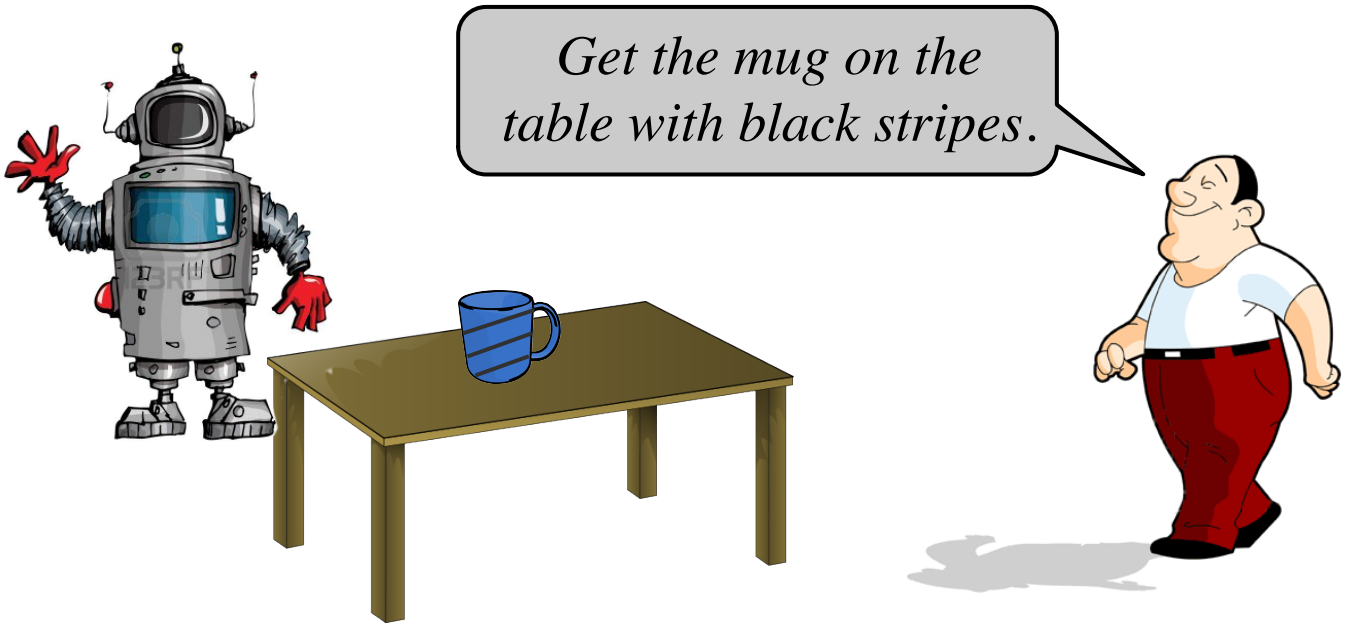

Visually Grounded Language

- Visual Question Answering: Humans asking machines about pictures/videos, e.g., for visually impaired, in remote/ dangerous scenarios, in household service settings

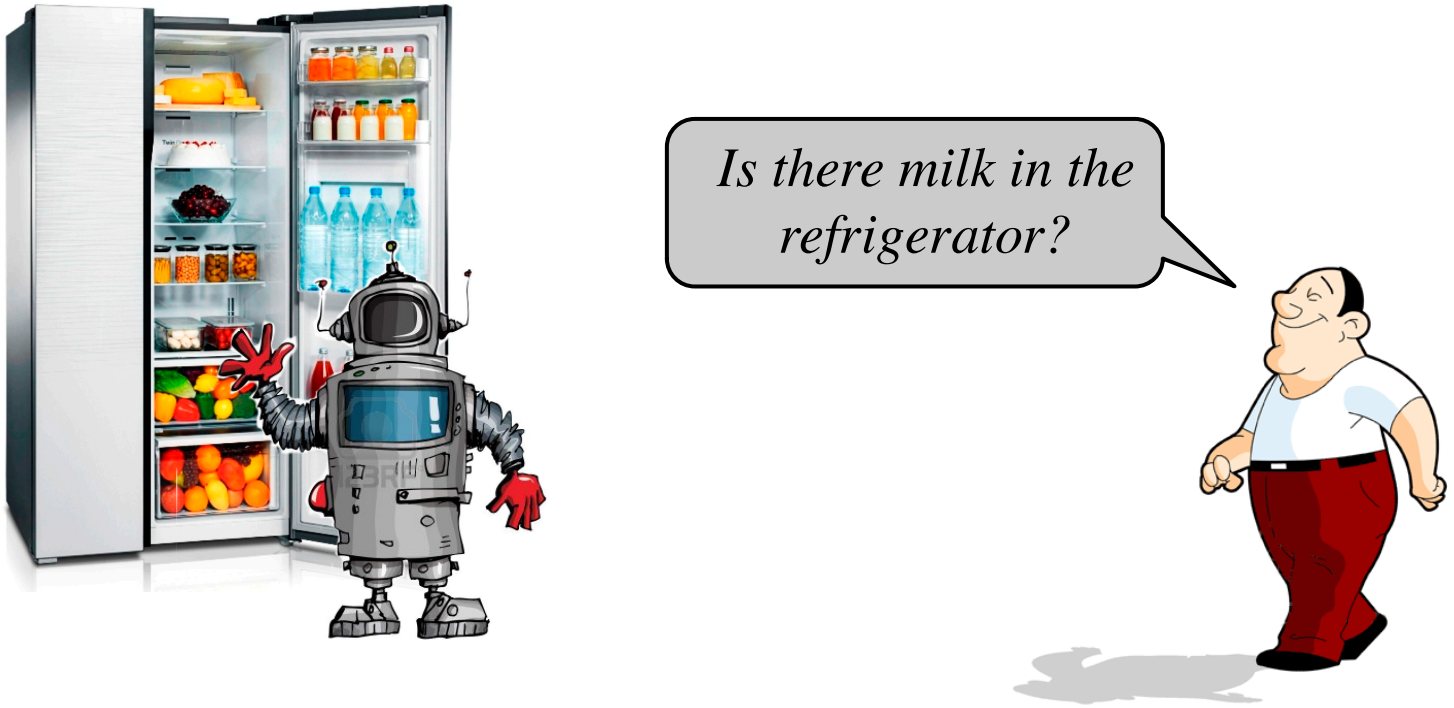

Embodied Language (Robot Instructions)

- Task-based instructions, e.g., navigation, grasping, manipulation, skill learning

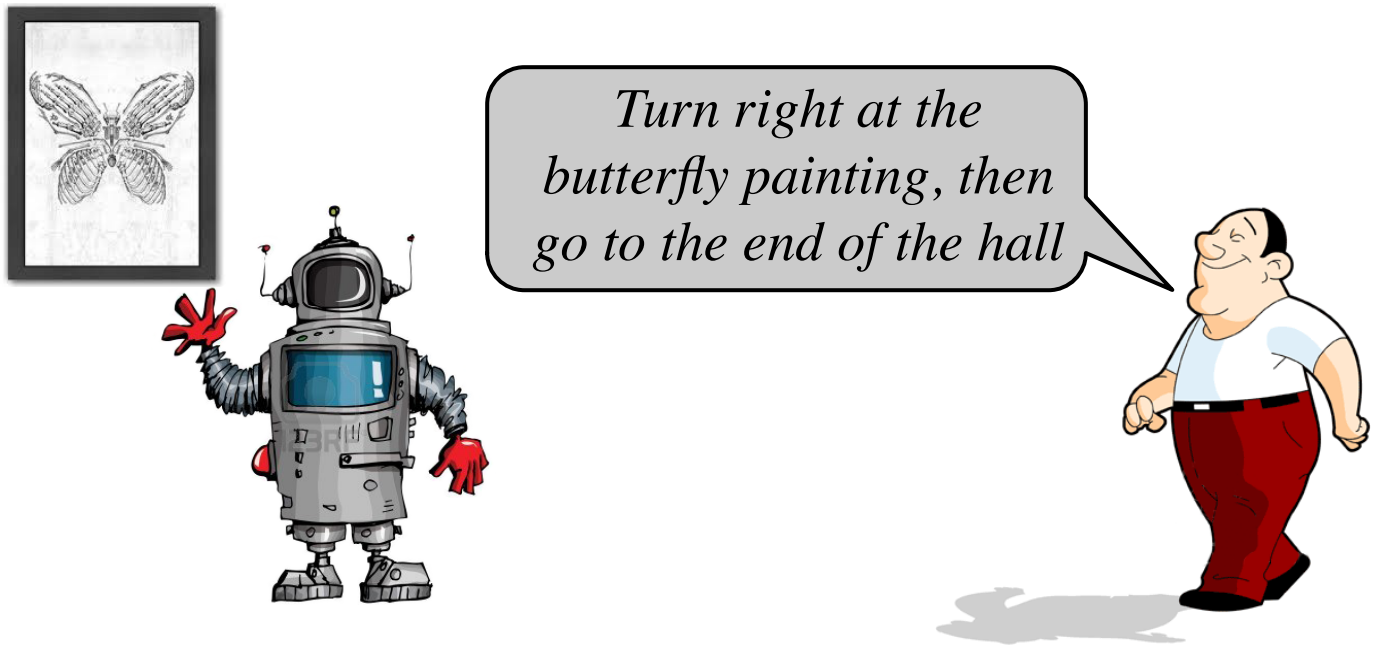

Embodied Language (Robot Instructions)

- Task-based instructions, e.g., navigation, grasping, manipulation, skill learning

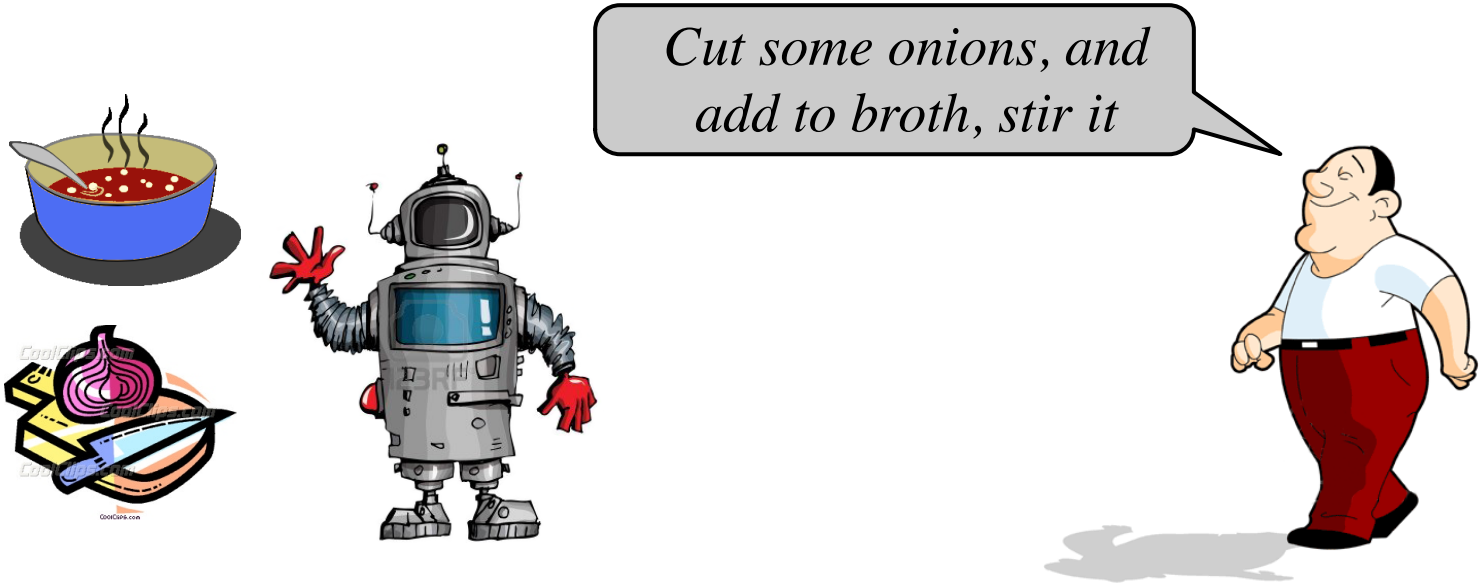

Grounded Language Generation/Dialogue

- Both for answering human questions, and to ask questions back, and for casual chit-chat

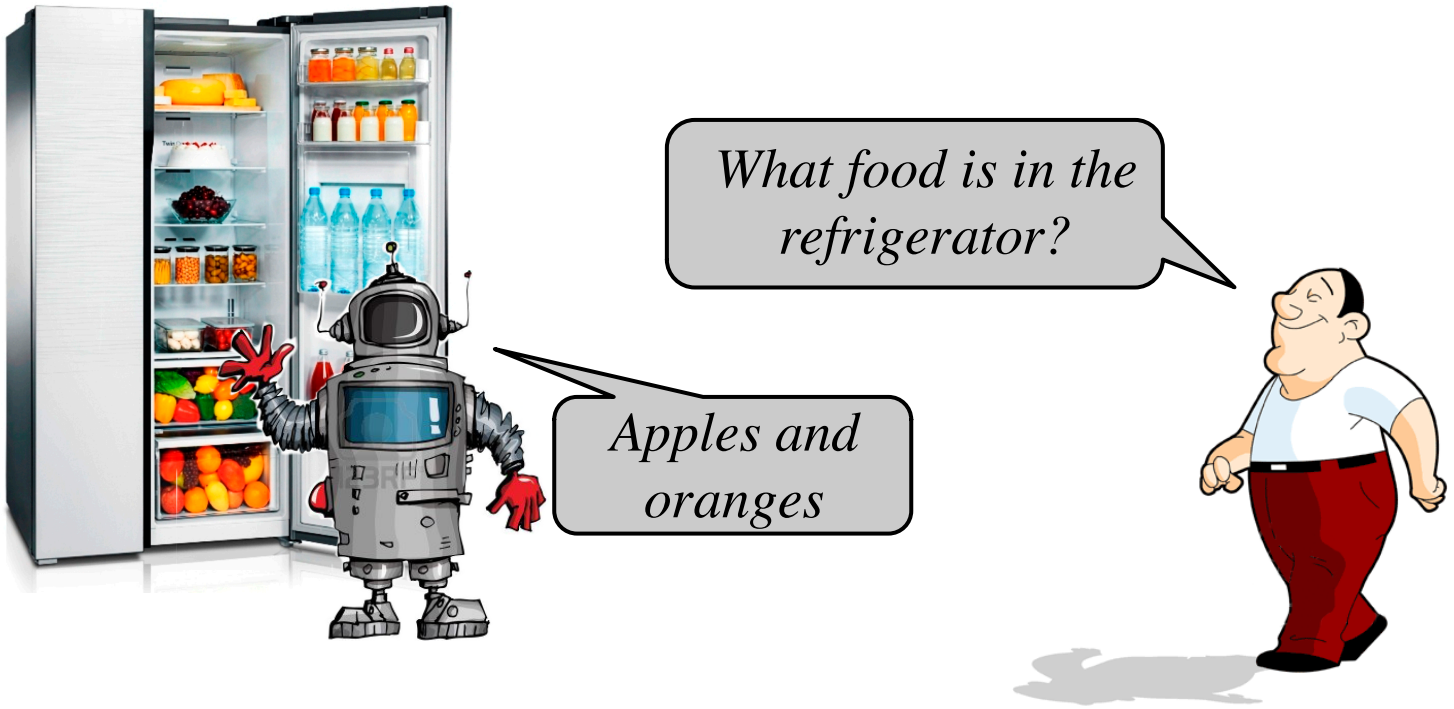

Grounded Language Generation/Dialogue

- Both for answering human questions, and to ask questions back, and for casual chit-chat

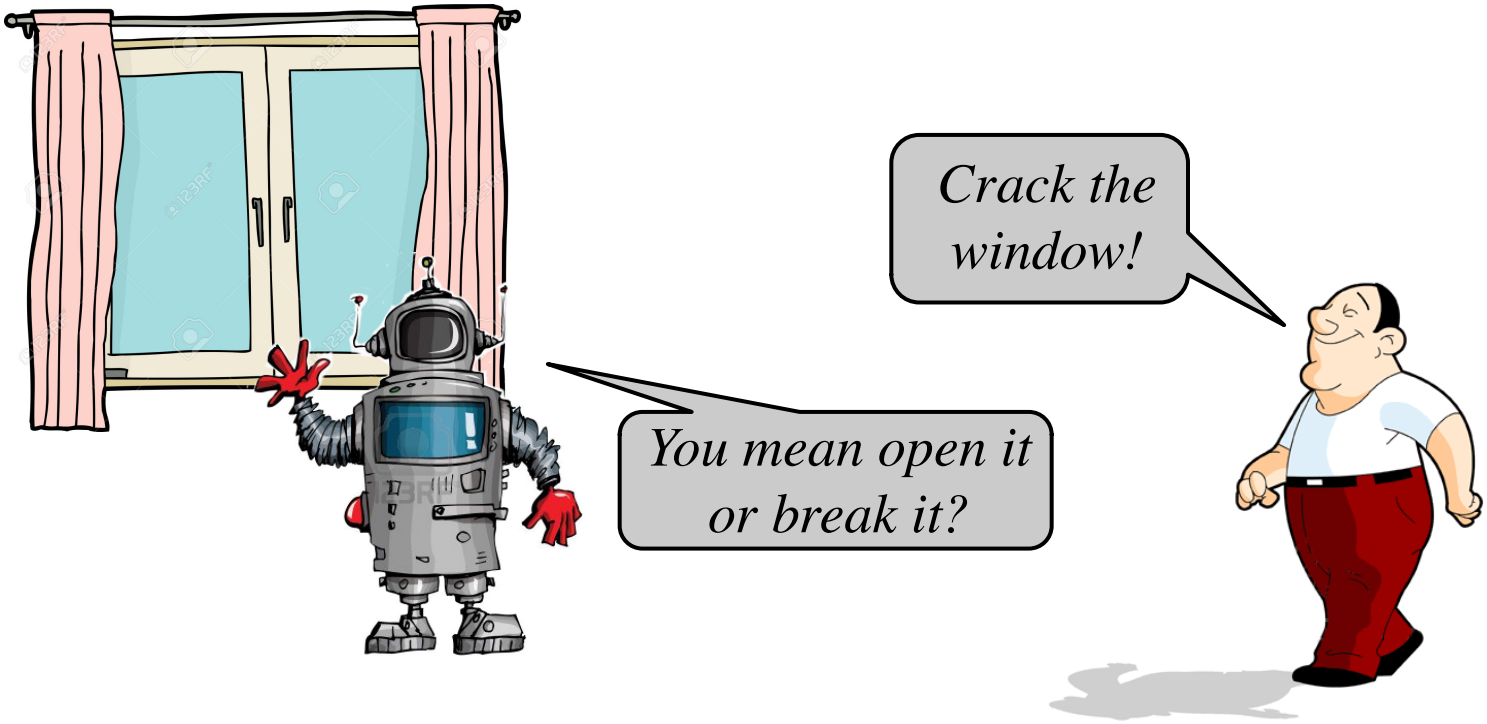

Next Week

- We will start with some basic of NLP, e.g., tagging, parsing, coreference, etc.
- Website will be updated with initial paper list very soon
- Once paper list is up, start volunteering for paper presentation!
- We might do pairs/groups of students depending on class size
- Project details and dates will be announced soon