

Natural Language Processing



Instructors: Dan Klein, Alane Suhr
GSIs: Jessy Lin, Kevin Yang
UC Berkeley

Logistics



Logistics

- Enrollment

- Class is full but we hope to fit waitlist
- We will process enrollments after P1
- Material access doesn't require enrollment
- <https://edstem.org/us/join/BfhEtz>

- Course expectations

- Readings, lectures, 5 projects
- No sections, no exams
- Engagement with the course
- Workload will be high, autonomy required
- Patience: class is under construction!

- Requirements

ML: A-level mastery,
eg CS189

NL: Care a lot about
natural language

PL: Ready to work in
PyTorch (Kaggle)

- COVID/Attendance Policies

- Webcasts posted (Berkeley-local)
- Expanded late day policy (14 day)
- Please do not come to lecture sick



Resources and Readings

- **Resources**
 - Webpage (syllabus, readings, slides, links)
 - EdStem Discussion (course communication)
 - Gradescope (submission and grades)
 - Compute via Kaggle notebooks
 - Berkeley-internal recordings for main lectures
- **Readings (see webpage)**
 - Individual papers will be linked
 - Optional text: Jurafsky & Martin, 3rd (more NL)
 - Optional text: Eisenstein (more ML)

CS 288	Schedule
Home Course Info Staff	August 24: Introduction <i>Project 1 released</i>
	August 29: Language modeling: statistical models
	August 31: Language modeling: early neural models

ed CS 288 - Ed Discussion

[New Thread](#) Search

COURSES +
CS 288

CATEGORIES
■ General
■ Lectures
■ Assignments

Filter v

No threads
Be the first to create a thread!



Projects and Infrastructure

Projects

- P1: Language Models
- P2: Machine Translation
- P3: Syntax and Parsing
- P4: LLMs and Structured Prediction
- P5: Multimodal Language Understanding

Infrastructure

- Python / PyTorch
- Compute via Kaggle notebooks
- Grading via Gradescope

The screenshot shows a Kaggle notebook titled "Project 1" in "Draft saved" mode. The interface includes a sidebar with navigation options like "Create", "Home", "Competitions", "Datasets", "Code", "Discussions", "Learn", and "More". The main content area contains text instructions and a code cell. The text instructions state that the write-up should be 1-page maximum and include a list of four requirements: 1. A concise and precise description of the extension that you tried. 2. A motivation for why you believed this approach might improve your model. 3. A discussion of whether the extension was effective and/or an analysis of the results. This will generally involve some combination of tables, learning curves, etc. 4. A bottom-line summary of your results comparing validation perplexities of your improvement to the original LSTM. The purpose of this exercise is to experiment, so feel free to try/ablate multiple of the suggestions above as well as any others you come up with! When you submit the file, please name it `report.pdf`.

Below the text is a code cell with the following content:

```
[ ]:
## Feel free to copy your original LSTM solution down here to modify for your report if you'd like.
# YOUR CODE [optionally] HERE
##

#lstm_model = LSTMModel()
#lstm_model.train()

print('lstm validation perplexity:', lstm_model.dataset_perplexity(validation_dataset))
save_truncated_distribution(lstm_model, 'lstm_predictions.npy', short=False)
```

What is NLP?



Natural Language Processing



NLP



Goal: Deep Understanding

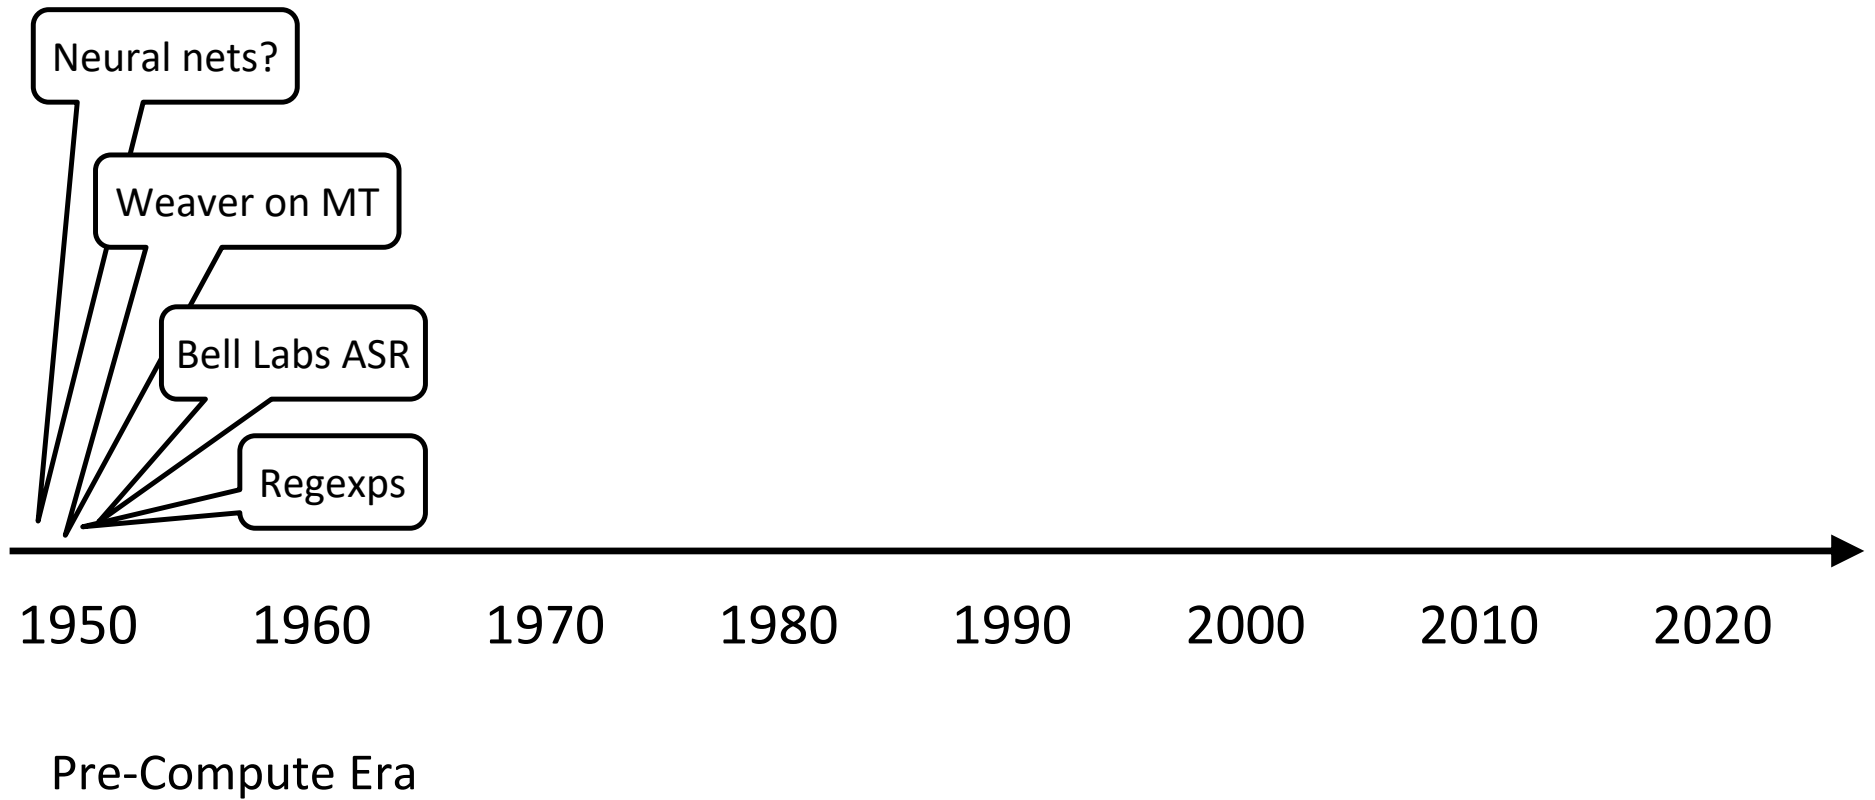
- Requires knowledge, context, and grounding
- Traditionally out of reach, but huge recent advances

Classic Systems: Shallow Matching

- Focus on robustness, scale, and predictability
- Very useful, but has had fundamental limitations



NLP History



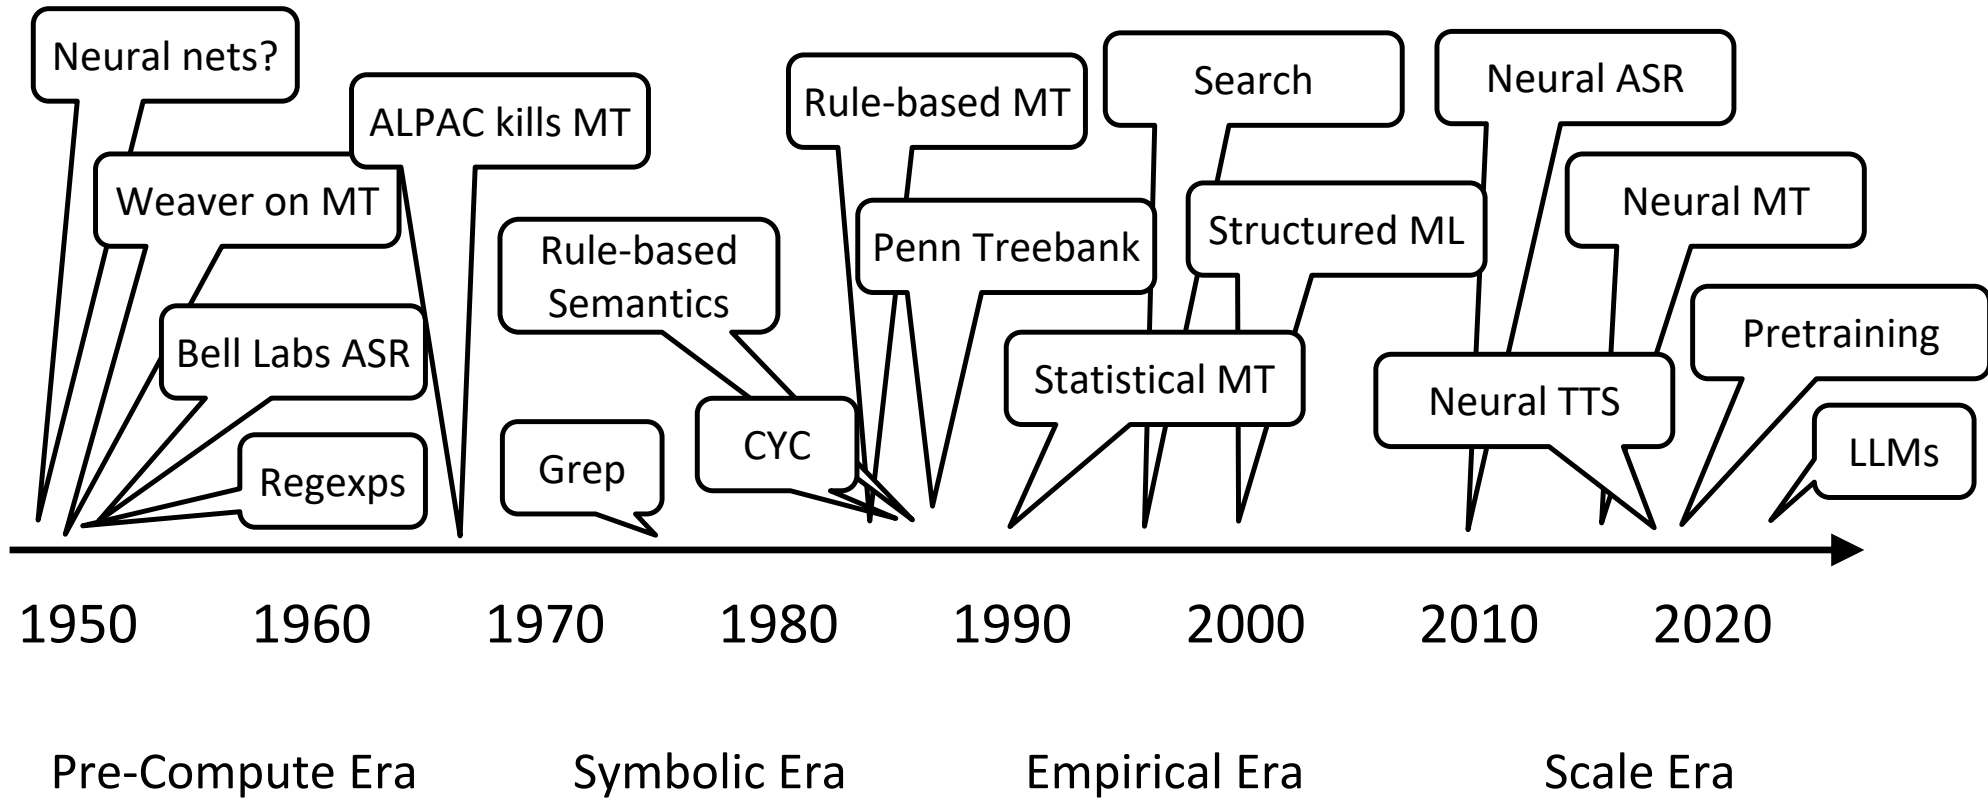


NLP History





NLP History



Designing NLP Systems



Role of Language Technologies

- What is their *intended* use?
 - Keeping users engaged?
 - Interacting with users as little as possible?
 - What kind of users?
- How does this influence their design?
 - The user interface
 - The underlying structure / model
 - The assumptions they are making
- How are they *actually* used?
 - Where do they fail?
 - How do users adapt?
- What other considerations should we make?
 - Resource efficiency
 - Fairness, ethics, social impact



Speech Systems

- **Text to Speech (TTS)**

- Text in, audio out
- SOTA: nearly perfect aside from prosody



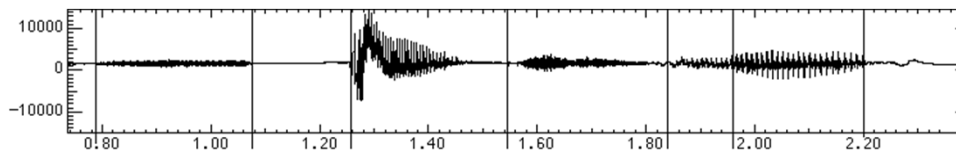
GOOGLE TECH ARTIFICIAL INTELLIGENCE

Google launches more realistic text-to-speech service powered by DeepMind's AI



- **Automatic Speech Recognition (ASR)**

- Audio in, text out
- SOTA: <<1% error for digit strings, 5% conversational speech, still >>20% hard acoustics



“Speech Lab”

MICROSOFT TECH ARTIFICIAL INTELLIGENCE

Microsoft reaches 'human parity' with new speech recognition system

Speak-N-Spell / Google WaveNet / The Verge



Machine Translation

EN CE MOMENT Impôts Kenya Syrie Pakistan Emploi Scandale Prism

Impôt sur le revenu : combien paierez-vous en 2014 ?

DÉCOUVREZ TOUS LES SERVICES ABONNÉS
S'abonner au Monde à partir de 1 €

APPEL À TÉMOIGNAGES

Adhérent(e) d'Europe Ecologie-Les Verts, partagez-vous le constat sévère de Noël Mamère sur EELV ?
Partagez votre expérience

En continu

- 19:53 Budget : des dépenses incompressibles
- 19:36 Entendu à la "fashion week" de Paris
- 19:19 Commande géante à Airbus
- 19:04 Plainte contre "Valeurs actuelles"
- 19:01 Venezuela : 17 personnes arrêtées
- 18:59 Vidberg : le budget nouveau est arrivé
- 18:50 La "mission noble" de la NSA
- 18:38 Roms : Joutes entre Bruxelles et Paris

Sélectionnez votre revenu et votre situation familiale pour découvrir si vous bénéficiez de la pause fiscale.

- Comment le budget pour 2014 est-il réparti ? VISUEL INTERACTIF
- Un budget 2014 soumis aux critiques
- Budget : ces dépenses qu'aucun gouvernement ne peut diminuer
- Budget 2014 : le détail des économies VISUEL INTERACTIF

AT THIS MOMENT Taxes Kenya Syria Pakistan Use Prism scandal

Income tax: how much do you pay in 2014?

DÉCOUVREZ TOUS LES SERVICES ABONNÉS
S'abonner au Monde à partir de 1 €

CALL FOR EVIDENCE

Member (s) of Europe Ecology-Greens, do you share the finding of severe Christmas Mamère EELV?
Share your experience

Continuous

- spenses
- "Week" in Paris
- Actual Values"
- ple arrested
- udget came
- " of the NSA
- ween Brussels &

Select your income...

- How is the budget allocated?
- A 2014 budget criticism

- Translate text from one language to another
- Challenges:
 - What's the mapping? [learning to translate]
 - How to make it efficient? [fast translation search]
 - Fluency (next class) vs fidelity (later)

Example: Yejin Choi



Spoken Language Translation



Image: Microsoft Skype via Yejin Choi



Summarization

- Condensing documents
 - Single or multiple docs
 - Extractive or abstractive
- Very context-dependent!

The image shows a screenshot of a CNN news article titled "Official: Egypt balloon explosion probe can take 2 weeks". The article is dated February 27, 2013, and is written by Adam Makary, Saad Abedine, and Mariano Castillo. The article discusses a hot air balloon explosion in Luxor, Egypt, that killed 19 people. The article is annotated with several colored boxes highlighting key information:

- Twitter:** A blue box containing a tweet: "CNN: official investigation into yesterday air balloon accident in Luxor could take 2 weeks".
- Story Highlights:** A red box containing a list of key points:
 - **NEW:** No foul play is suspected, official says.
 - The Tuesday accident was the world's deadliest hot air balloon accident in at least 20 years.
 - Officials: Passengers in the balloon included 19 foreign tourists.
 - Luxor province bans all hot air balloon flights until further notice.
- Text Excerpts:** Several green boxes highlight specific text from the article:
 - "The Tuesday accident was the world's deadliest hot air balloon accident in at least 20 years."
 - "Preliminary investigations confirmed no foul play was involved when gas canisters aboard the balloon exploded, causing it to plummet about 1,000 feet (300 meters) to the ground, Gov. Ezzat Saad said."
 - "Passengers in the balloon included 19 foreign tourists: nine from Hong Kong, four from Japan, three from Britain, two from France and one from Hungary, officials said."
 - "Tuesday's crash prompted the governor to ban all hot air balloon flights until further notice."
- Video Thumbnails:** Three video thumbnails are shown with "Click to play" buttons:
 - "Tourists killed in hot air balloon blast"
 - "2000 balloon crash survivor speaks"
 - "How to stay safe in a hot air balloon"

Image: CNN via Wei Gao



Search, Questions, and Reasoning

Google Web Images Groups News Froogle Local more »
any US states' capitals are also their largest cities? Search

Google Web
What is
Web
Tip: Have a question? Ask the
Web definitions for capital
Cheyenne: the capital a
wordnet.princeton.edu/p
Scholarly articles for Wha
The Cultural Contexts of
Capital requirements and
MSN Encarta - Cheyenne
Cheyenne (city, Wyoming), c
... In 1869 the city was selecte
encarta.msn.com/encyclopedia_76
capital of Wyoming: Informa
Note: click on a word meaning below
The noun **capital** of Wyoming has
www.answers.com/topic/capital-of-
Cheyenne: Weather and
Chey-enne (shī-ān ' , -ēn ')
the state near the Nebraska
www.answers.com/topic/che

How many US states' capitals are also their largest cities?

How many US states' capitals are their second largest cities?

As of my last knowledge update in September 2021, there are two US states where the second largest city serves as the state capital:

1. Texas: The second largest city in Texas is San Antonio, which is not the state capital.
2. Florida: The second largest city in Florida is Miami, which is not the capital of Florida.

Please verify this information since circumstances might have changed after September 2021.

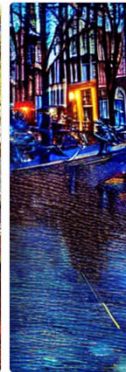


Text to Image

the first day in an NLP class

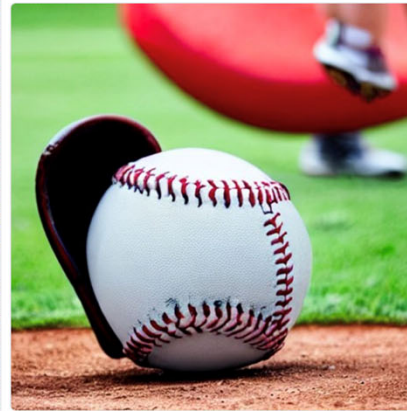
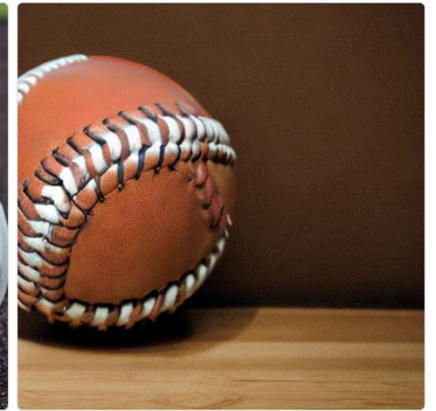


Amsterdam canal in post-impressionist brush reflection of narrow houses and bicycles, swirl channelling van Gogh's 'Starry Night', engulf from pedestrian-eye vie



a baseball sitting to the left of a basketball

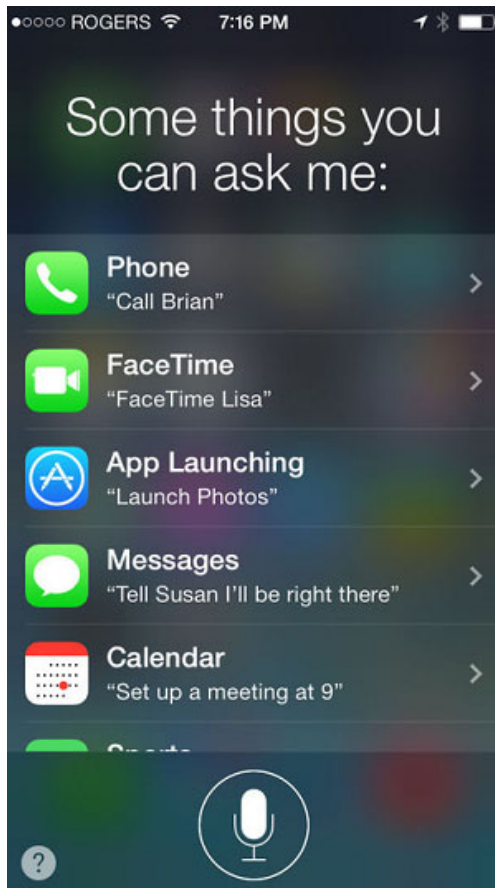
Generate image



source: <https://www.inkearn.com/pulse/prompt-hack-midjourney-stable-diffusion-marengo-kemp>



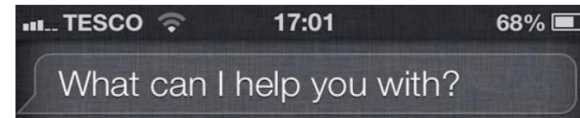
Virtual Assistants



- Virtual assistants must do
 - Speech recognition
 - Language analysis
 - Dialog processing
 - Text to speech
 - ... and back-end actions!



Conversations with Devices?



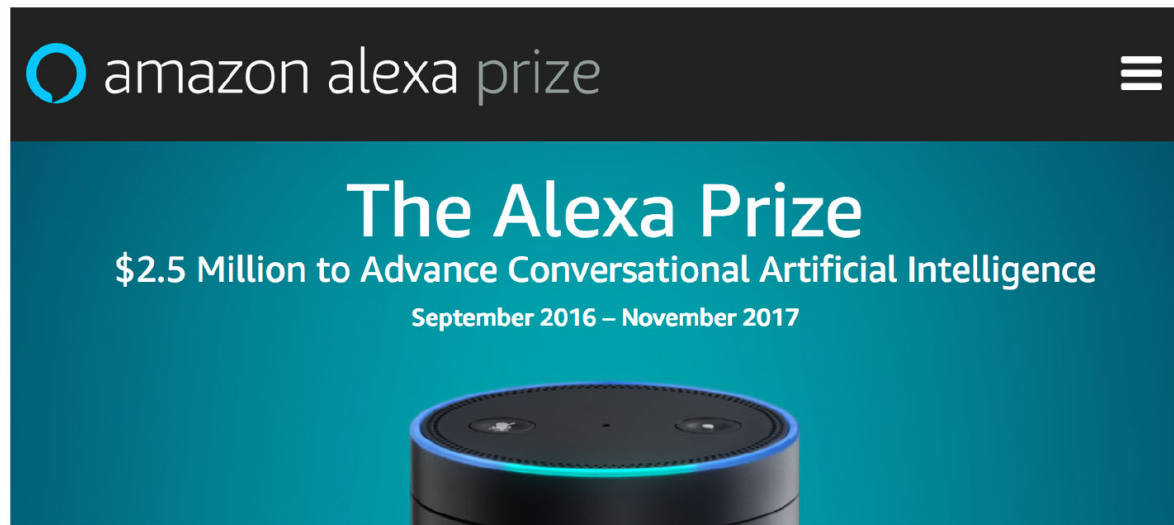


ELIZA

```
ELIZA: You're not really talking about me -- are you ?
YOU:   What?
ELIZA: What do you think ?
YOU:   I mean, I feel happy to be writing right now, to be typing this as we spe
ak.  Lately, I've had writer's block.
ELIZA: You say you mean ?
YOU:   That's what I said.
ELIZA: Can you elaborate on that ?
YOU:   I only really feel happy when I'm writing.
ELIZA: How long have you been writing ?
YOU:   Always, but I've been writing seriously since I was eighteen.
ELIZA: What incident are you thinking of ?
YOU:   Um.
ELIZA: What does that suggest to you ?
YOU:   I don't think it's enough, to only be happy when you write.
ELIZA: Does that trouble you ?
YOU:   ...Um, because you can't write all the time.  And then there's the real w
orld, and you still have to deal with real people.
ELIZA: What are your feelings now ?
YOU:   I've been thinking a lot of that poem by Yeats:  "The Man and the Echo."
        Do you know it?
ELIZA: Do you say you've been thinking a lot of that poem by yeats: "the man and
        the echo for some special reason ?
YOU:   Yes
```



Chatbot Competitions!



- Alexa Prize competition to build chatbots that keep users engaged
 - First winner in 2017: UW's Sounding Board (Fang, Cheng, Holtzman, Ostendorf, Sap, Clark, Choi)
- Compare to the Turing test (eg Loebner Prize) where the goal is to fool people



SoundingBoard Example

Turn	User	Sounding Board
1	Let's chat	Hi! This is an Alexa Prize socialbot. How's your day?
2	I'm <u>five</u> .	Glad to hear it! We could talk about robots, batman, or superman. Say "next", to chat about other things.
3	Superman.	It looks like you wanna chat about superman. I was high up in the cloud when I realized: If people don't recognize Clark Kent as Superman because of his glasses, does his eye doctor know his identity?
4	I guess so.	Did you know that Henry Cavill almost missed the call for the role of Superman cause he was playing "World of Warcraft"
5	Really, I didn't know that.	Weird, right? Speaking of superman, Did you see: Superman? It was released in 1997. It's a comedy.
6	Yes, it was hilarious.	I'm glad you feel this is hilarious. Which part Movie est about this movie?
7	The part when he met lewis leah.	Interesting. meccartin, and raffi co-directed this film. The movie has a 6.3 out of 10 on IMDB. which seems pretty good! do you like the movie's director? Fact
8	Yeah, he's cool.	Cool! By the way, do you wanna know something interesting about raffi?
9	Okay, go ahead.	Super! My friend told me that Raffi is a worldwide children's rights activist and has fought against exploitative advertising to children.

Negotiation

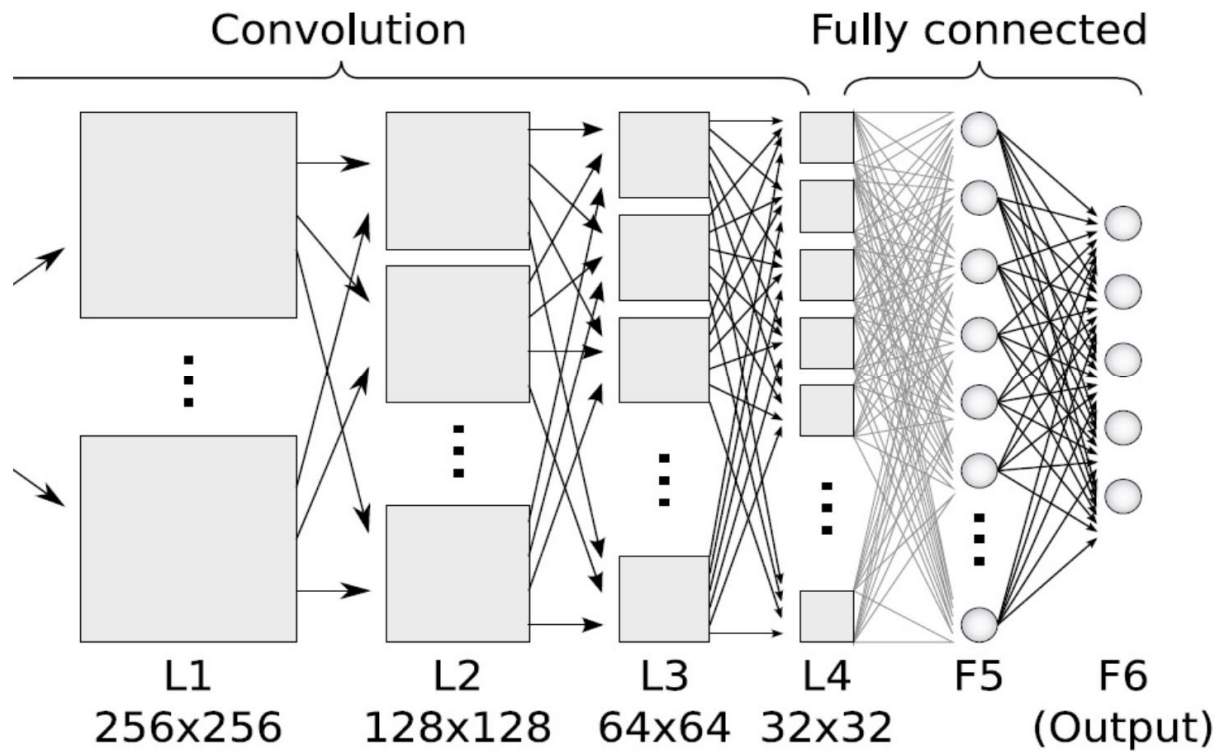
Thought

Movie

Fact

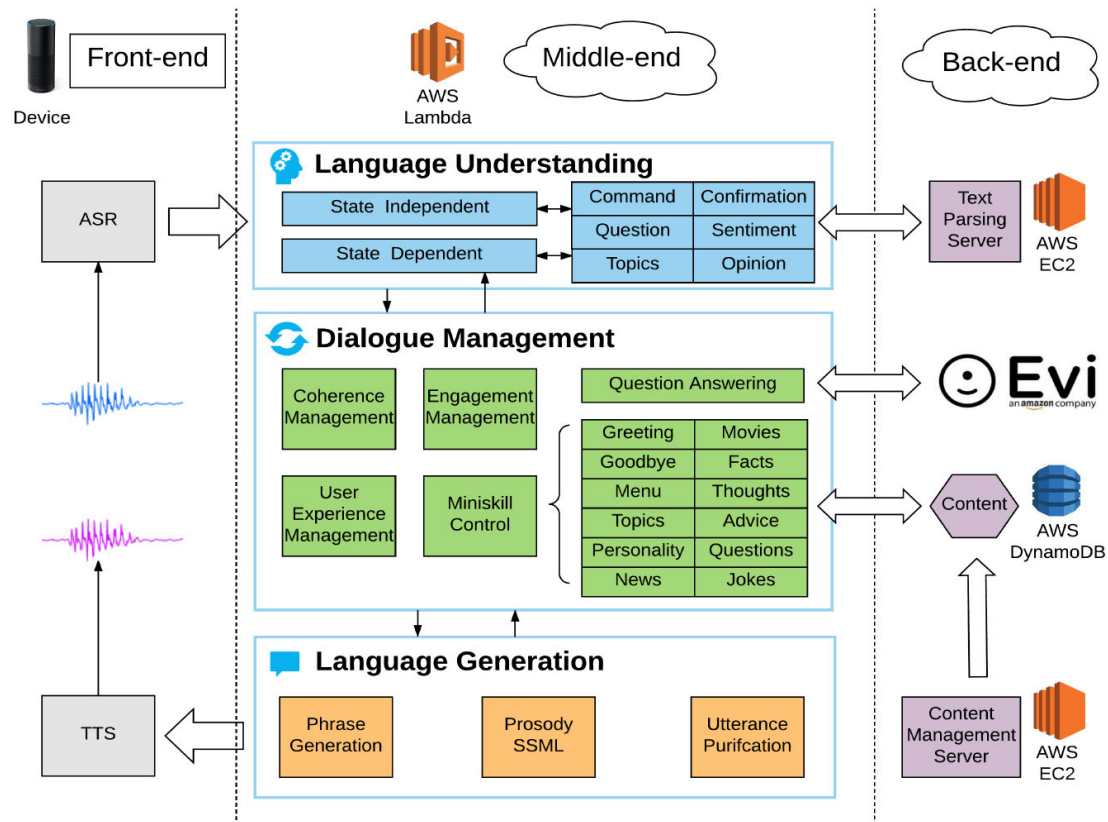


Sounding Board's Architecture?















Sounding Board's Architecture



Source: Yejin Choi

 <p>My roomba seems to be br</p>	 <p>The light on the roomba is white</p>	 <p>I tried resetting it with the clean button and spinning</p>	 <p>The roomba isn't moving at all because it won't respond to the clean button</p>
 <p>I'd be happy to help you trouble-shoot, diagnose and potentially fix the issue.</p> <ol style="list-style-type: none"> Check for Obstructions: Make sure the Roomba's path or brushes aren't blocked or moving erratically. Clean the Roomba: Dust and debris can accumulate in the brushes and sensors, which can affect the Roomba's performance. 	 <p>If the light on your Roomba is white, it means the Roomba is attempting to clean but it's trying to locate the dock for charging. Here's what you can do to address this issue:</p> <ol style="list-style-type: none"> Check Dock Placement: Ensure the dock is placed in a clear area with enough space around it. 	 <p>If resetting the Roomba using the "Clean" button doesn't resolve the issue and the behavior continues, it's possible that there may be a hardware issue with the Roomba's internal components or sensors. Here are some additional steps:</p> <ol style="list-style-type: none"> Check Sensors: Inspect the sensors on the Roomba, particularly the bumper sensors. These sensors help the Roomba detect obstacles and prevent it from getting stuck. 	 <p>If your Roomba isn't responding to the "Clean" button and isn't moving at all, it's possible that there might be a more serious issue with its internal components. Here are some steps you can take to address this situation:</p> <ol style="list-style-type: none"> Check Battery: Make sure your Roomba has a sufficient battery charge. If the battery is completely drained, the Roomba won't respond. If necessary, connect the Roomba to the charging dock and let it charge for a while before attempting to power it on again.
 <p>that was extremely helpful and not at all too verbose or irrelevant to my situation</p>			
 <p>I'm glad to hear that you found the information helpful and relevant! If you have any more questions or need further assistance in the future, feel free to reach out. Good luck with resolving the issue with your Roomba, and take care!</p>			
<p>electronics, you might attempt to repair them yourself. However, if you're not experienced, it's best to consult with professionals.</p> <p>Remember, safety is important when working with electronics. If you're unsure about a step, it's best to consult with the support team.</p>	<p>prevent it from taking down your Wi-Fi network. If you're experiencing this issue, you can try the following steps:</p> <ol style="list-style-type: none"> Contact Customer Support: If you're still having trouble, you might be able to provide additional information to the support team. If the issue is not resolved after trying these steps, you might want to consider replacement options. 	<ol style="list-style-type: none"> Professional Repair: If your Roomba is not responding to the "Clean" button and you're not experienced with electronics, you might consider opening it up to inspect the internal components. However, if you're not experienced with electronics, attempting repairs that could potentially cause damage to the Roomba is not recommended. <p>Remember that electronic devices can be sensitive to static electricity. If you're attempting to repair the Roomba, make sure you're properly grounded and use proper knowledge to avoid causing damage. If you're still having trouble, you might want to consider assistance from the manufacturer's support team.</p>	<ol style="list-style-type: none"> Consider Replacement: If your Roomba is significantly older and has been experiencing multiple issues, it might be more cost-effective to consider replacing it with a newer model. <p>Given that your Roomba isn't responding at all, and it's not a simple issue like an obstruction or connectivity problem, it's best to involve the manufacturer's support team or a professional technician who specializes in repairing Roombas. They will have the expertise and tools necessary to diagnose and fix the issue correctly.</p>



Ambiguities



Live TV

Stevie Wonder announces he'll be having kidney surgery during London concert

By Amir Vera, CNN

Updated 11:16 PM EDT, Sat July 06, 2019



(CNN) — [Stevie Wonder](#) will be taking a break from music.

The legendary singer-songwriter announced during [a concert in London](#) Saturday that he will be undergoing kidney surgery.



Problem: Ambiguity

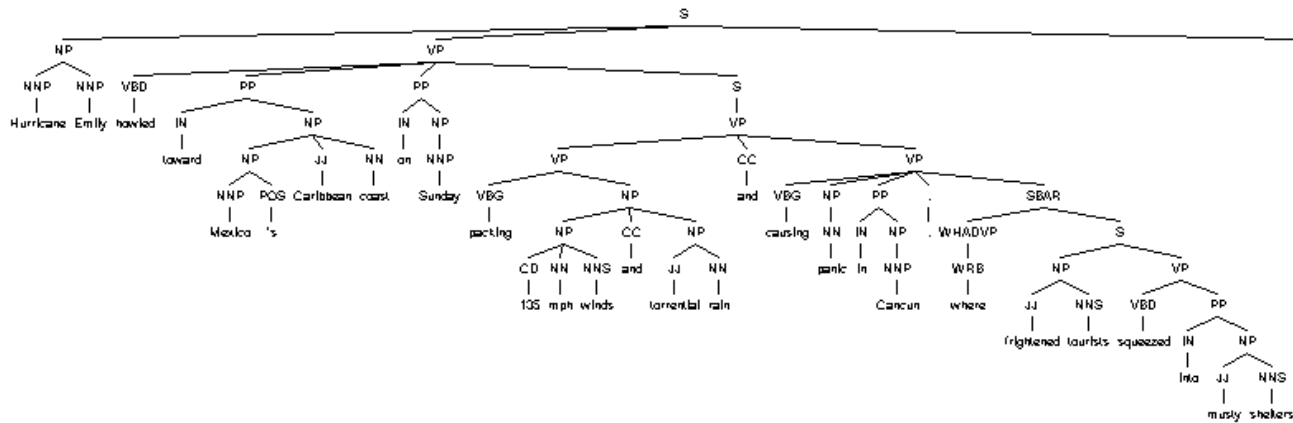
- Headlines:
 - Teacher Strikes Idle Kids
 - Hospitals Are Sued by 7 Foot Doctors
 - Ban on Nude Dancing on Governor's Desk
 - Stolen Painting Found by Tree
 - Kids Make Nutritious Snacks
 - Iraqi Head Seeks Arms
 - Local HS Dropouts Cut in Half

- Why are these funny?

What Do We Need to Understand Language?



Example: Syntactic Analysis



Hurricane Emily howled toward Mexico 's Caribbean coast on Sunday packing 135 mph winds and torrential rain and causing panic in Cancun, where frightened tourists squeezed into musty shelters .

Accuracy: 95+



Latent Representations

Personal Pronouns (PRP)

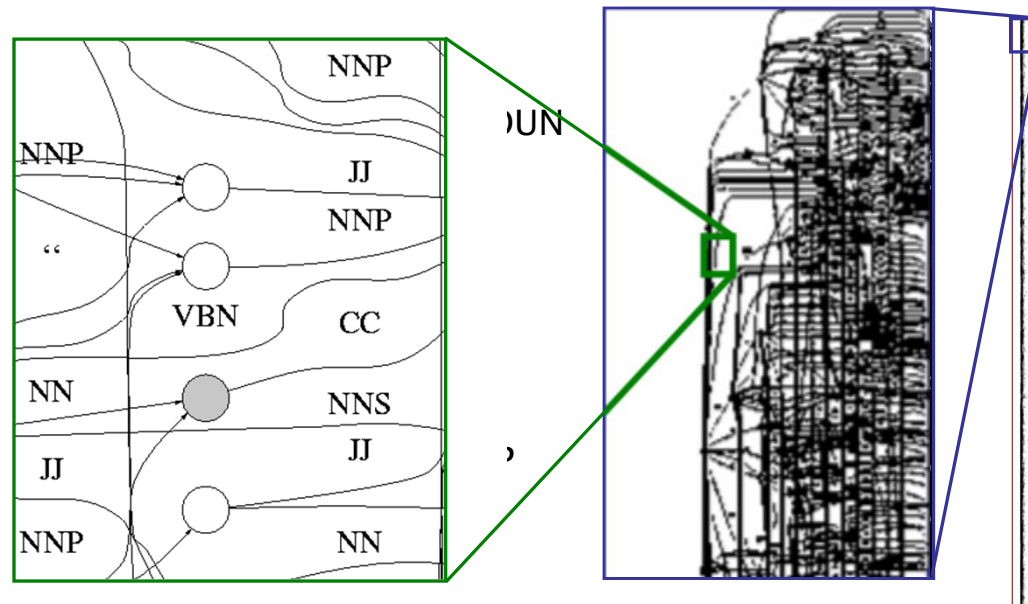
PRP-1	it	them	him
PRP-2	it	he	they
PRP-3	It	He	I

Proper Nouns (NNP)

NNP-14	Oct.	Nov.	Sept.
NNP-12	John	Robert	James
NNP-2	J.	E.	L.
NNP-1	Bush	Noriega	Peters
NNP-15	New	San	Wall
NNP-3	York	Francisco	Street



We Need Data





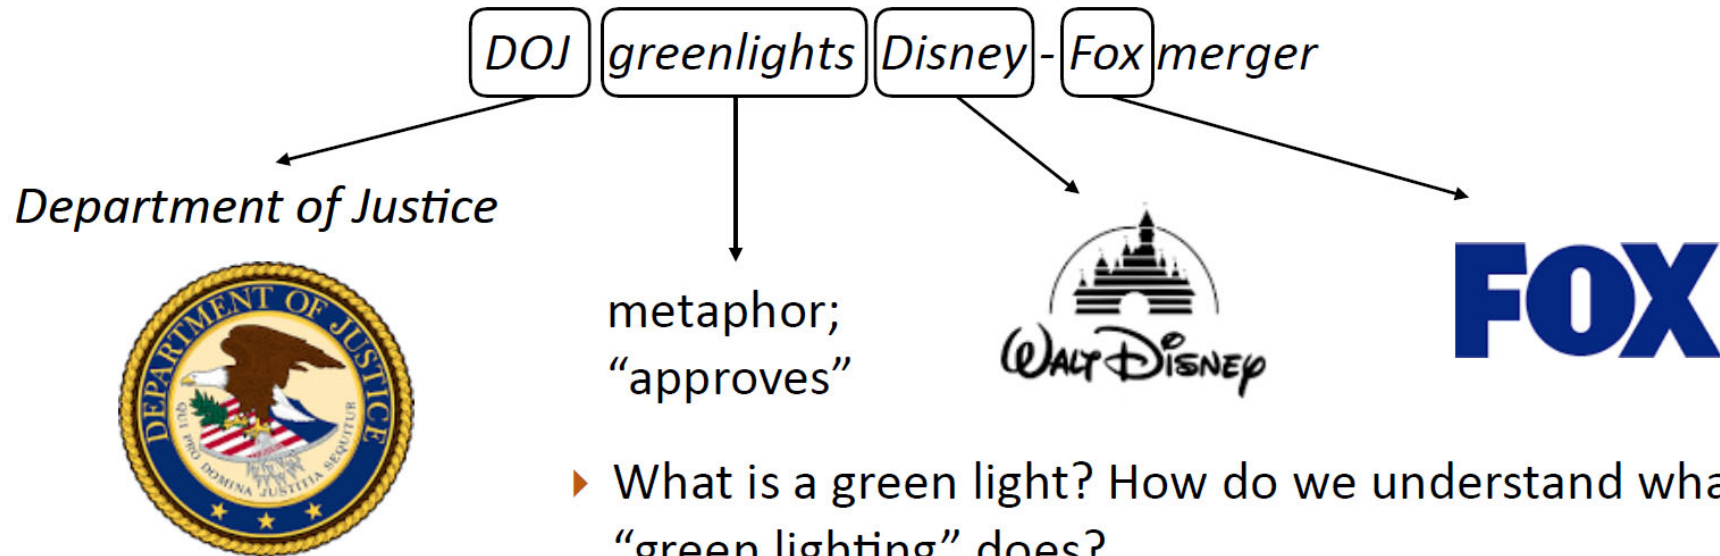
We Need A Lot of Data

SOURCE	Cela constituerait une solution transitoire qui permettrait de conduire à terme à une charte à valeur contraignante.
HUMAN	That would be an interim solution which would make it possible to work towards a binding charter in the long term .
1x DATA	[this] [constituerait] [assistance] [transitoire] [who] [permettrait] [licences] [to] [terme] [to] [a] [charter] [to] [value] [contraignante] [.]
10x DATA	[it] [would] [a solution] [transitional] [which] [would] [of] [lead] [to] [term] [to a] [charter] [to] [value] [binding] [.]
100x DATA	[this] [would be] [a transitional solution] [which would] [lead to] [a charter] [legally binding] [.]
1000x DATA	[that would be] [a transitional solution] [which would] [eventually lead to] [a binding charter] [.]



We Need Knowledge


- ▶ World knowledge: have access to information beyond the training data



- ▶ What is a green light? How do we understand what "green lighting" does?
- ▶ Need commonsense knowledge



Data and Knowledge

- Classic knowledge representation worries: How will a machine ever know that...
 - Ice is frozen water?
 - Beige looks like this: 
 - Chairs are solid?
- Answers:
 - 1980: write it all down
 - 2000: get by without it
 - 2020: learn it from web data



Knowledge from Large LMs?

A12 Allen Institute for AI

AllenNLP

- ↶ Annotate a sentence ^
 - Semantic Role Labeling
 - Named Entity Recognition
 - Constituency Parsing
 - Dependency Parsing
 - Open Information Extraction
 - Sentiment Analysis
- ☰ Annotate a passage ^
 - Coreference Resolution
- ❓ Answer a question ^
 - Reading Comprehension
- ⚙ Semantic parsing ^

Language Modeling

This demonstration uses the public 345M parameter [OpenAI GPT-2](#) language model to generate sentences.

Enter some initial text and the model will generate the most likely next words. You can click on one of those words to choose it and continue or just keep typing. Click the left arrow at the bottom to undo your last choice.

Sentence:

At high temperatures, ice becomes

Predictions:

32.0% **The**

10.3% **Ice**

5.2% ↶

3.7% **Water**

2.7% **When**

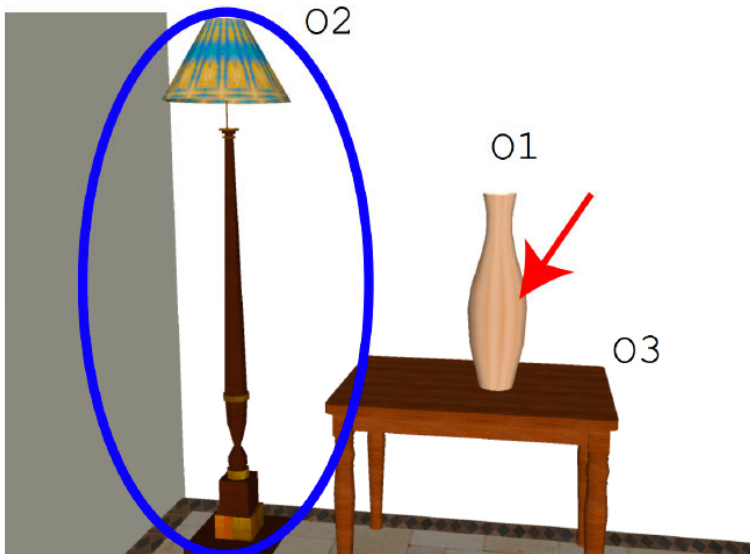
↶ **Undo**



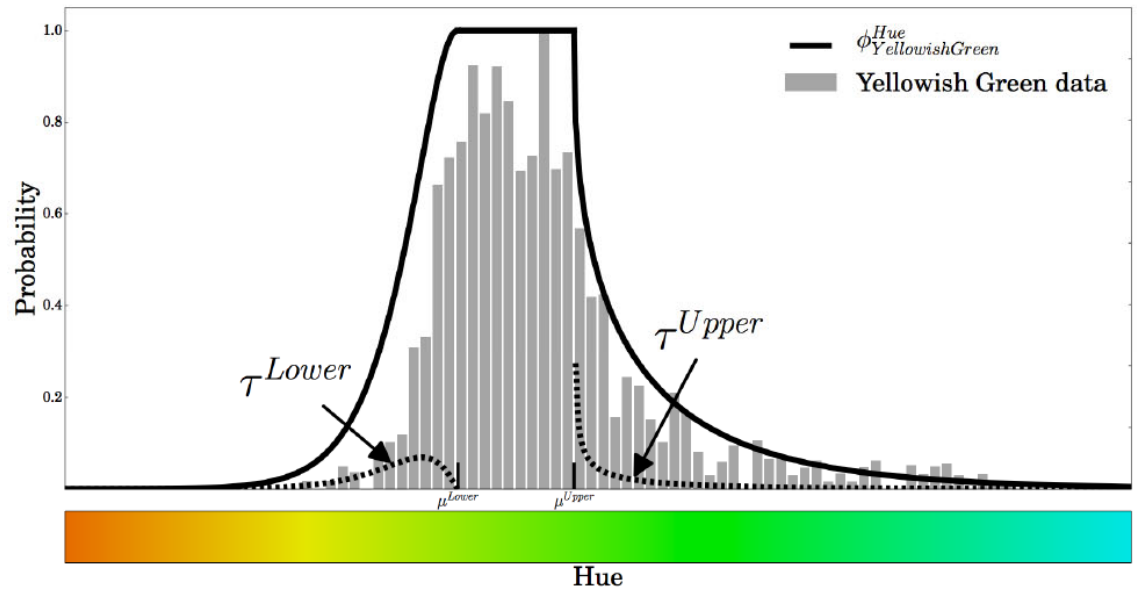
We Need Grounding

Grounding: linking linguistic concepts to non-linguistic ones

Question: What object is right of 02 ?



Golland et al. (2010)



McMahan and Stone (2015)



Example: Grounded Dialog

When is my package arriving?

Friday!





Why is Language Hard?

- We Need:
 - Representations
 - Models
 - Data
 - Grounding
 - Learning
 - Scale
 - Efficient Algorithms
- ... and often we need all these things at the same time



Course Overview (Subject to Change)

- 6 Lectures: Language / Sequence Modeling
- 2 Lectures: NLP Tasks, Data, and Evaluation
- 5 Lectures: Structured Prediction
- 4 Lectures: Transformers and LLMs
- 5 Lectures: NLP Beyond Text
- 5 Lectures: Panels on Current Topics