

Lecture 23

Context in Language Processing

Nathan Schneider

ENLP | 6 May 2021


Be aware she can move in any direction

Be aware she can move in any direction





AI Ambitions




Semantic Analysis

- We've seen tasks that analyze the meanings or topics of documents, words, and sentences
 - ▶ document classification
 - ▶ topic models
 - ▶ word representations & similarity
 - ▶ word sense disambiguation
 - ▶ semantic role labeling
- These are challenging tasks. But even if we could automate them perfectly, we'd still be a long way from human-like automatic language processing.

Understanding: Beyond Semantics

What is required to understand this conversation?



- **Semantics of the expressions themselves**
 - ▶ coffee refers to the drink, not the tree or bean (WSD)
 - ▶ 4:00 and 3:00 are times (NER)
 - ▶ “at 4:00”: semantic role marking the time of an event
 - ▶ “?” indicates question
- **But there’s a lot more to understanding than just the explicit language....**



- **Encyclopedic knowledge about the world**
 - ▶ Nobody would think this means “Does coffee exist at 4:00?” We know about social activities associated with coffee.
 - ▶ Likely 4:00pm, because people are normally asleep at 4:00am. (And people generally don’t go for coffee in their sleep.) Unless.....



- **Knowledge of the situation/
conversational context/
common ground**
 - ▶ Perhaps it's 2:30am and we're working to finish something for a deadline. Or we both are back from a conference and are severely jetlagged.
 - ▶ Perhaps we have a habit of going to a certain place for coffee, so it can be left implicit.



- **Discourse coherence**

- ▶ We normally assume that interlocutors are “cooperative” (H.P. Grice): They respond with relevant information, say what they believe to be true, don’t change topics without suitable pause or warning, etc.
- ▶ Here, we interpret the second question as *proposing an alternative* time, and requesting confirmation.

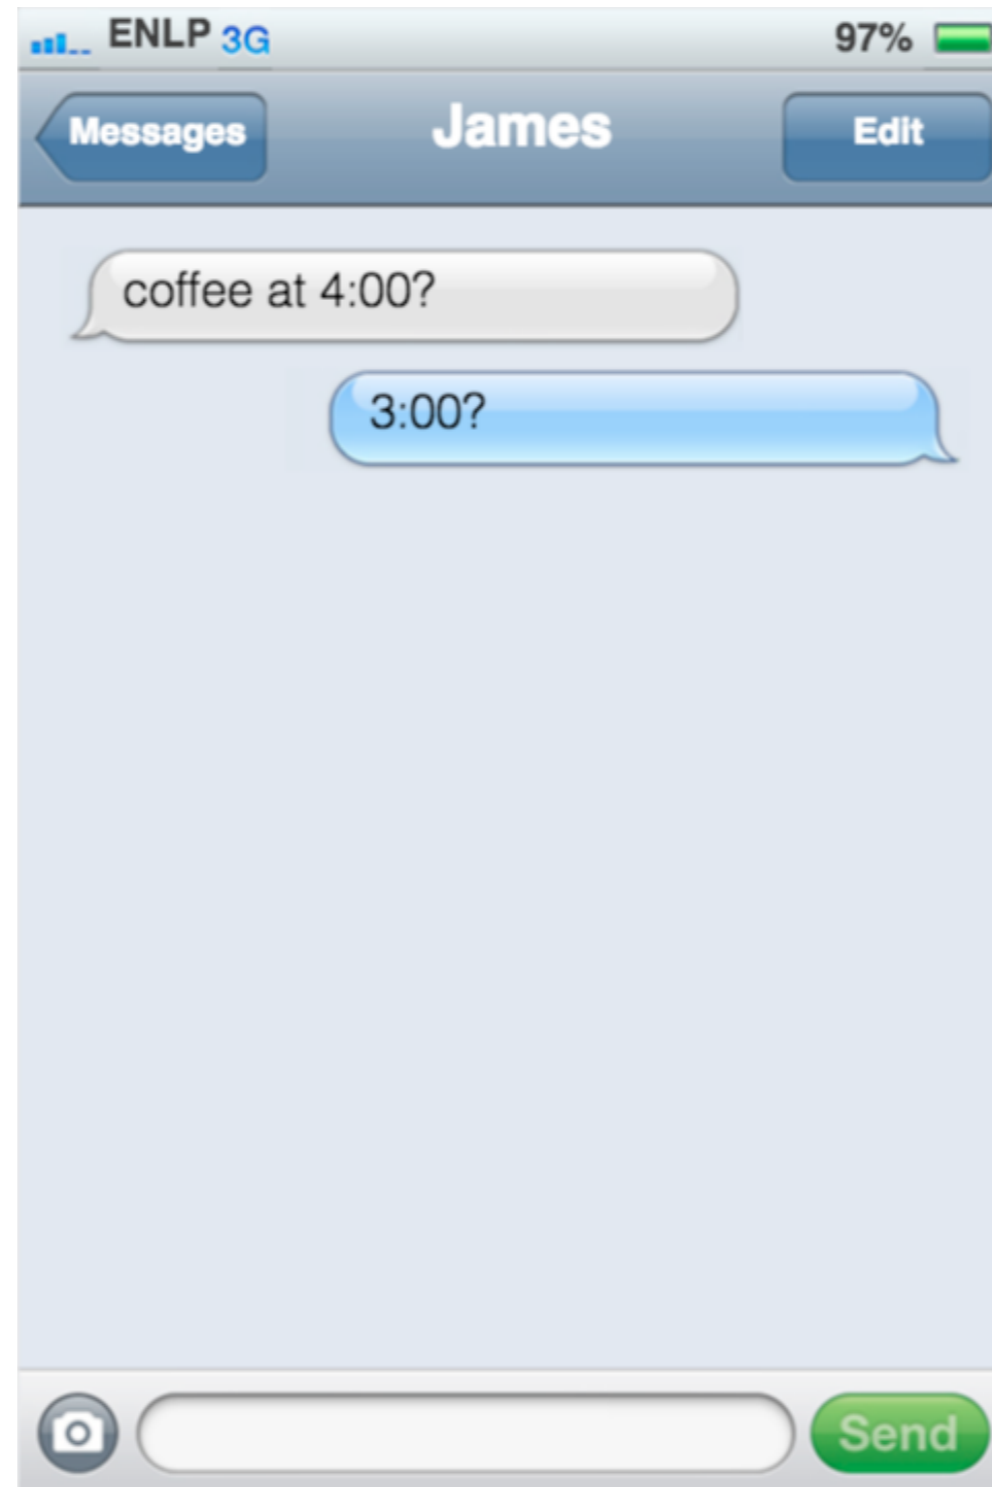


- **Relationship to action**

- ▶ A truly intelligent app would offer information that would help my decision (e.g., when the café closes)
- ▶ and put the event on my calendar at the agreed-upon time
- ▶ and remind me to leave in time to arrive at the agreed-upon meeting place at that time.
- ▶ If it is unsure of details, it should confirm with me rather than do the wrong thing.
- ▶ Industry is already moving in this direction with personal assistants.



What else can be inferred from this conversation?



- **Relationship between interlocutors**

- ▶ This conversation is informal. We might infer that the speakers are friends.



Types of Context

- **Conversational** (what has been said already, whose turn it is to talk)
- **Situational** (what is going on at the moment of conversation)
- **Social**
 - ▶ relationship between interlocutors—e.g. status/formality
 - ▶ their sociolinguistic identities—e.g. accent, expression of gender
- **Shared knowledge** (e.g. that we are in the same class and there is an exam coming up)
- **General knowledge** (e.g. that in order to stay awake late at night it might be helpful to drink coffee)


Understanding: It's Not Just the Words

- Actually understanding such conversations requires a lot of inferences based on world knowledge and context (**pragmatics**).

Understanding: It's Not Just the Words

WHAT IS SAID

WHAT IS UNDERSTOOD



Understanding: It's Not Just the Words

- Actually understanding such conversations requires a lot of inferences based on world knowledge and context (**pragmatics**).
- But is that only true of conversations? What about unidirectional language use (books, articles)?

“Sherwood Park had its third fire in less than a month on Tuesday. However, there were no injuries” (<http://www.sherwoodparknews.com/2016/01/14/no-injuries-in-park-fire>)

- **Semantics:** Sherwood Park is a neighborhood (not a literal park); this fire is no longer active
- **Discourse + world knowledge:**
 - ▶ *no humans injured in this fire* (unknown whether any ants were harmed)
 - ▶ “However” signals a contrast with an expectation raised by the first sentence: injuries might have been expected from an unintentional fire
 - ▶ Harm to humans is highly newsworthy, so it’s important for the story to inform us of an event that DIDN’T occur
 - ▶ Likely inference: there is a pattern of fires in Sherwood Park (why?)
 - ▶ Were there injuries in previous fires? Unspecified.
 - ▶ What would have to change for the information to be presented in the opposite order?

Understanding: It's Not Just the Words

- Different aspects of meaning are required to be explicit in different languages. E.g., **lexicalization** patterns in Hebrew vs. English:

eish אש

Google



אש 🔍



All

Images

Videos

News

Maps

More

Settings



Tools

View saved

SafeSearch




שריפה *srefa*

Google  שריפה 

All News Images Videos Shopping More Settings Tools View saved SafeSearch

purple white black blue red green water space land air ice lightning rain house building barn wood coal charcoal oil fireplace



The image grid contains 30 individual photographs of fire-related incidents. The top row shows a grass fire, a large fire in a field, a multi-story building on fire with firefighters, a close-up of intense flames, a person standing near a large fire, a house on fire, and a fire at night. The second row features a fire at night with a person in the foreground, a fire at night with a person, a large fire at night, a fire in a field, firefighters at a large fire, and a fire with a flag in the background. The third row includes a fire in a field, a fire in a field, a fire on a building, a house on fire, a firefighter at a fire, and a fire in a building. The bottom row shows a fire in a field, a fire at night with a play button icon, a fire truck at night, a firefighter at a fire, a fire with a chair in the foreground, and a fire in a field.

Understanding: It's Not Just the Words

- Different aspects of meaning are required to be explicit in different languages. E.g., **lexicalization** patterns in Hebrew vs. English:
 - ▶ EN “fire” ↔ HE {*eish* ‘purposeful fire’, *srefa* ‘destructive fire’}
 - ▶ EN {“color”, “paint”} ↔ HE *tseva*
- **formality/social status:** Which 2nd person pronoun to use in German or French?
- **evidentiality:** How does the speaker know the information? (directly observed, secondhand, etc.)
- **spatial systems:** absolute (compass directions) or relative

Understanding: It's Not Just the Words

- Some information can be made “minimally explicit”, requiring discourse-level inference.
- **anaphora** (pronouns): *He* sells the greatest soup *you*'ve ever eaten.
 - ▶ Need to decide which pronouns are **referential**, and resolve their antecedents.
 - ▶ Special case of **coreference resolution** (grouping referring expressions that indicate the same entity).
- **pro-drop**: In many languages, pronominal subjects can be dropped (verb agreement helps disambiguate): *Quiero un taco.*

Perspective in language

- The choice of language can put a “spin” on the information being conveyed, emphasizing certain nuances or dimensions of meaning. Sometimes called **construal**.
- May indicate a social perspective (**framing**)
 - ▶ *Mistakes were made.*
 - ▶ “thrifty” vs. “stingy”
 - ▶ “terrorists” vs. “freedom fighters”
- May be mundane and subtle: **on** *the bus* vs. **in** *the bus*

Understanding: It's Complicated

- Lots of implicit information, even in expository text.
- How to even evaluate whether a system is comprehending the story?
 - ▶ Give the system an **exam**—multiple choice or fill-in-the-blank. Challenge datasets based on actual exam questions (reading comprehension, mathematical reasoning, biology).
 - ▶ Test the system's **decision-making** skills, such as controlling a robot or making moves in a game based on language. Requires link between comprehension and action/grounding.
 - ▶ **Multimodal**: Link text to image, video, or action.

Automatic Caption Generation

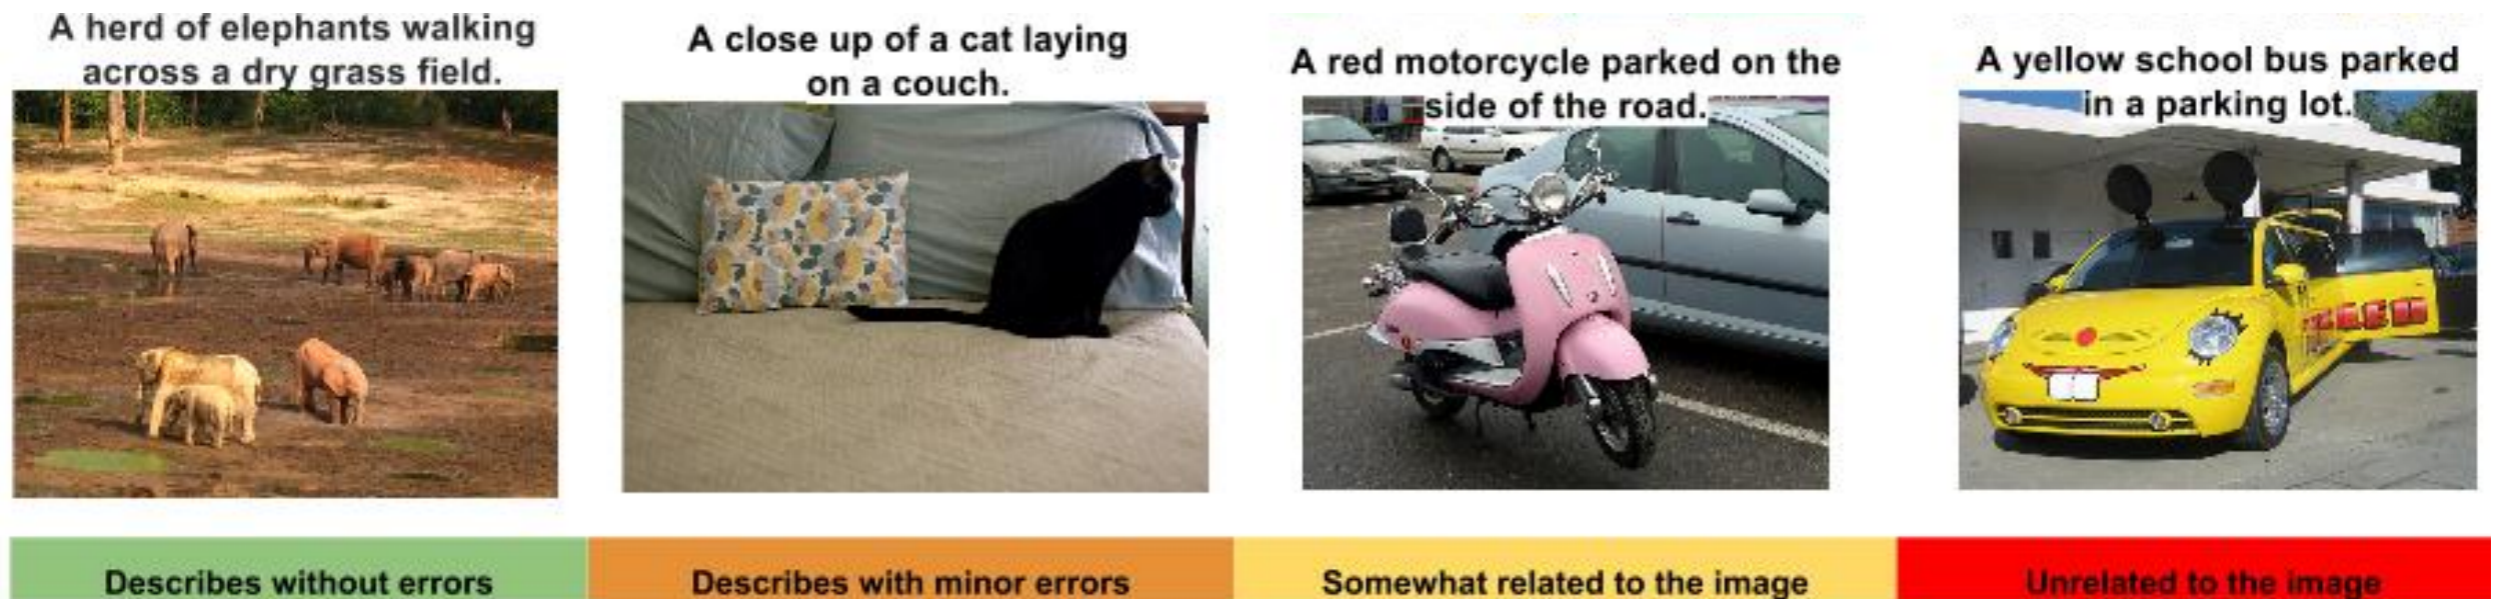


Figure 5. A selection of evaluation results, grouped by human rating.

Vinyals et al., CVPR 2015

http://www.cv-foundation.org/openaccess/content_cvpr_2015/papers/Vinyals_Show_and_Tell_2015_CVPR_paper.pdf

Language-Directed Robot Navigation



https://www.youtube.com/watch?v=7nUq28utuGM&list=PL6SYoj2z5jWfBFhZQdxF_luQ-sgpDXAO4&index=1

NLP in Academia



Muppets Models



NLP in the Real World



Social Context

- NLP technologies have users.
- Sometimes characteristics of the users matter to the task.
- It's one thing to optimize F-score in a research paper. But this might not capture the true cost of certain errors to a user.
- Do users trust systems? Do they trust them *too much*?
- How can systems be designed to be more trustworthy and transparent?

Societal Context

How is the NLP/AI technology going to be used?
What positive or negative effects might it have on society?


- **ACTIVITY:** Give examples of how technology can be used in ways that are
 - ▶ Advantageous to society
 - ▶ Unintentionally harmful
 - ▶ Intentionally harmful

Societal Context

How is the NLP/AI technology going to be used?
What positive or negative effects might it have on society?

- ▶ Promote societal goals—health care, crime prevention?
- ▶ Unfairness—implicit bias against certain groups of people (gender, race, dialect, ...)?
- ▶ Unethical or malicious uses—violating privacy, deception?

Recommendations

- Timnit Gebru et al.: [Datasheets for Datasets](#)
- Margaret Mitchell et al.: [Model Cards for Model Reporting](#)
- Emily Bender et al.: [On the Dangers of Stochastic Parrots: Can Language Models Be Too Big?](#) 
- + Social Factors in CL & AI course!

Summary

- The techniques discussed in this course were aimed at classifying documents, or analyzing words and sentences.
- But much of human language exploits our awareness of discourse, pragmatics, perspectives, other modalities, and the world. This is really, really hard (AI complete).
- Different languages have different requirements for what needs to be explicit.
- Tasks like **coreference resolution**, **automatic caption generation**, and **language-directed robot tasks** are important for stimulating research along dimensions beyond local and explicit linguistic communication.
- NLP technologies have an impact on users and society. Important to consider potential malicious or unintentionally problematic implications.