

NLP Use and Language Use Toward Artificial Language Users?

David Schlangen

Grundlagen der Computerlinguistik // Department Linguistik // Universität Potsdam

<http://clp.ling.uni-potsdam.de>

david.schlangen@uni-potsdam.de

LITHME

2022-12-02

These slides:

<https://clp.ling.uni-potsdam.de/talks>

Conclusions

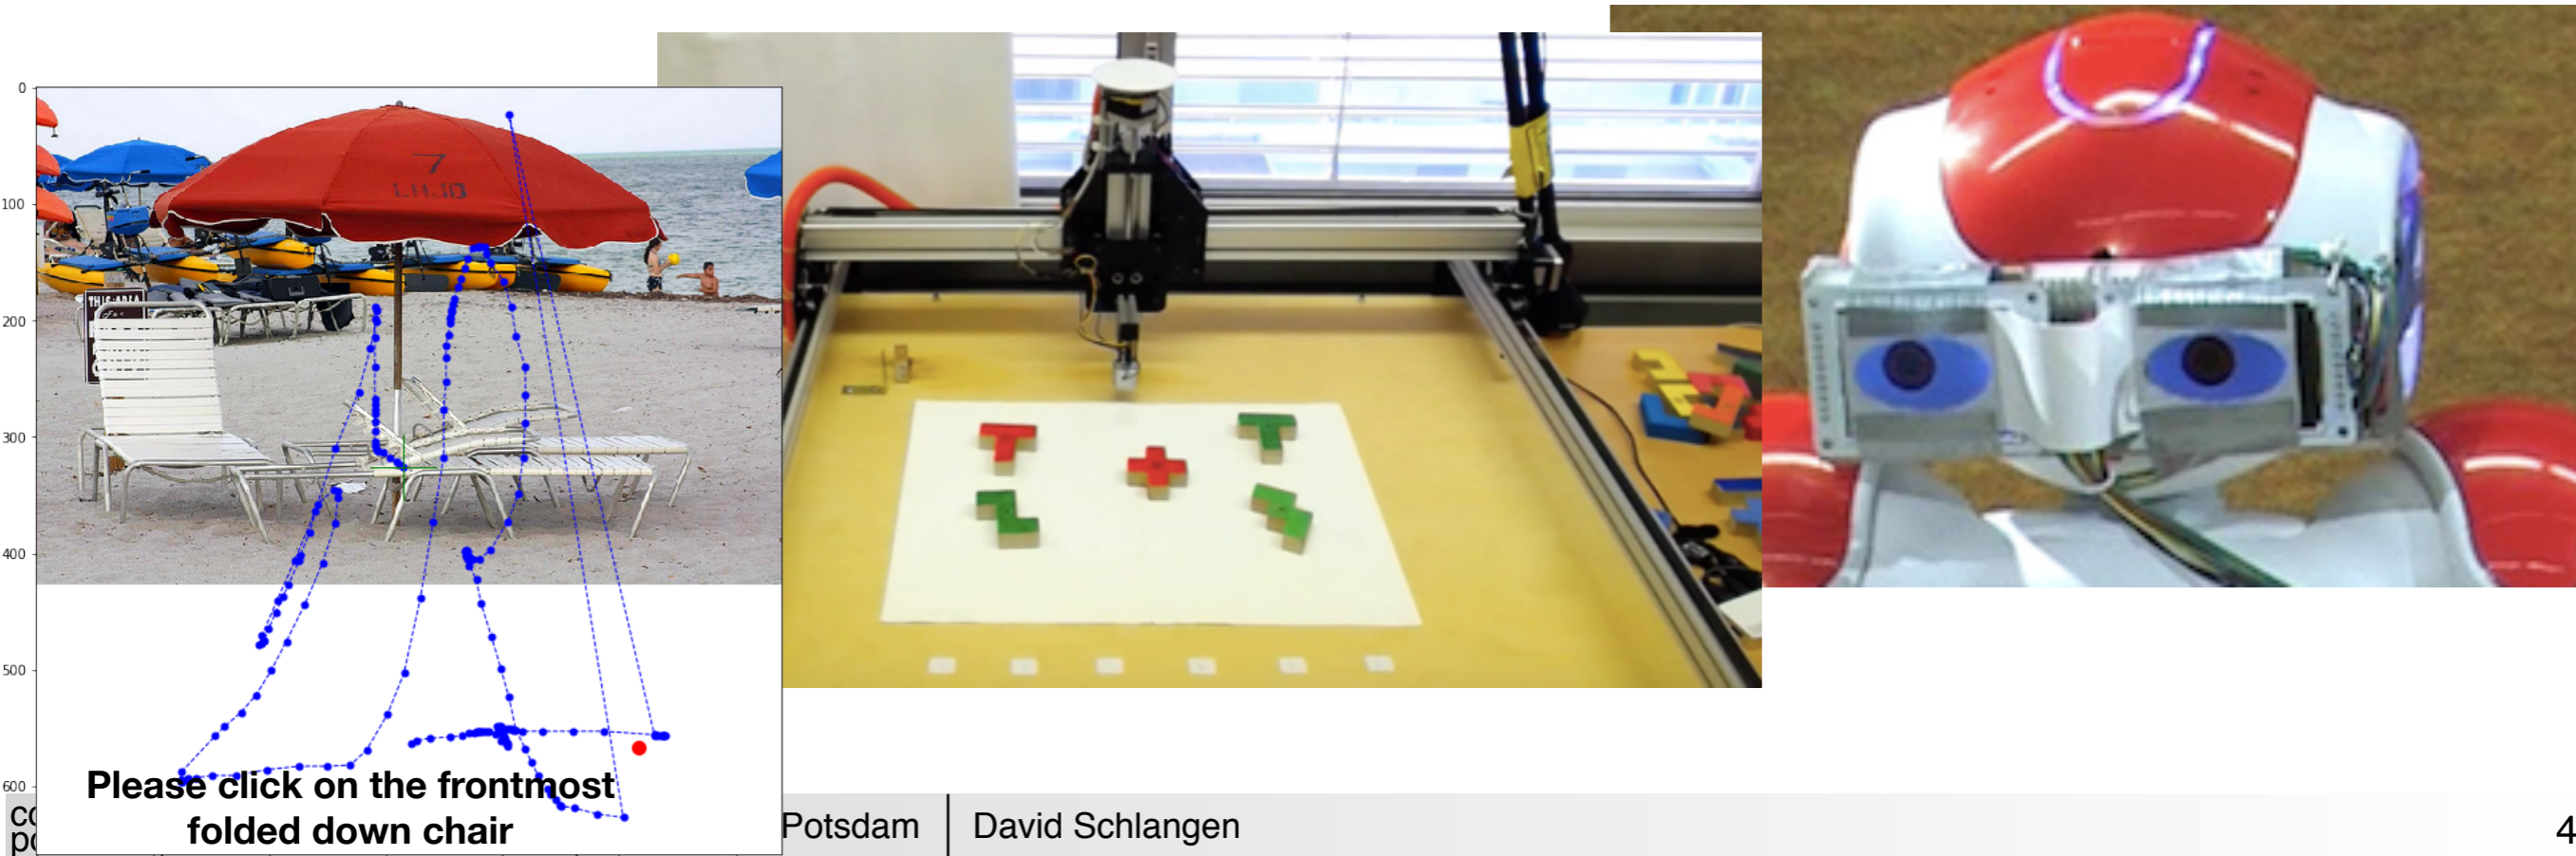
- User facing NLP technologies often appear to be not just language processors, but language users
- The way these technologies use language is conceptually different from how people use language, and from how people assume that these technologies use language
- Some possible ways to address this mismatch
 - restrict applications to limit potential harm
 - model additional capabilities
 - erode norms of language use

Overview

- Background
 - Your Speaker
 - The Situation
 - General AI is Here and Very Much Not Here
 - Language Processing and Language Use
- A Minimal Model of Meaning Making
- How Is Computer Speech Meaningful?
 - Excursus: Large Language Models
- Possible Ways to Address the Mismatch btw. Assumptions and Reality

About Me

- MA (Bonn) in CL, Philosophy, & CS
- PhD (UoEdinburgh), formal & comp. pragmatics of dialogue
- 2006 – 2011, Emmy Noether Group on Incremental Dial. Proc.
- 2011 – 2019, Prof. of Applied CL, Bielefeld University
- 2019 –, Prof. of the Foundations of CL, University of Potsdam




The Situation

General Artificial Intelligence is There, and Not There

TECHNOLOGY

The Google engineer who thinks the company's AI has come to life

AI ethicists warned Google not to impersonate humans. Now one of Google's own thinks there

By  Nitasha Tiku

June 11, 2022 at 8:00 a.m. EDT

 Futurism

Google Says It's Closing in on Human-Level Artificial Intelligence

Artificial intelligence researchers are doubling down on the concept that we will achieve artificial general intelligence (AGI) — that's AI that...

19 May 2022

 Daily Mail

Google's DeepMind says it is close to achieving 'human-level artificial intelligence'

DeepMind, a British company owned by Google, may be on the verge of achieving human-level artificial intelligence (AI). Advertisement.

18 May 2022

WOULD YOU PAY A MONTHLY FEE FOR VOICE COMMANDS? —

Amazon Alexa is a “colossal failure,” on pace to lose \$10 billion this year

Layoffs reportedly hit the Alexa team hard as the company's biggest money loser.

RON AMADEO - 11/21/2022, 8:32 PM



Enlarge / The fourth-generation Echo device is a cloth-covered sphere with a halo at the base, contrasting with the squat plastic cylinders of earlier-generation Echoes.

The Puzzle

SuperGLUE (Wang *et al.* 2019)

What causes a change in motion? The application of a force. Any time an object changes motion, a force has been applied. In what ways can this happen? Force can cause an object at rest to start moving. Forces can cause objects to speed up or slow down. Forces can cause a moving object to stop. Forces can also cause a change in direction. In short, forces cause changes in motion. The moving object may change its speed, its direction, or both. We know that changes in motion require a force. We know that the size of the force determines the change in motion. How much an objects motion changes when a force is applied depends on two things. It depends on the strength of the force. It also depends on the objects mass. Think about some simple tasks you may regularly do. You may pick up a baseball. This requires only a very small force.

Would the mass of a baseball affect how much force you have to use to pick it up?

Yes ✓

Liam Fedus, ST-MoE-32B: 91.2

“A sparsely activated Mixture-of-Expert model with 269B parameters, FLOP-matched to a 32B parameter dense model. Pre-trained on C4 corpus (Raffel *et al.*, 2019).”

Amazon’s Alexa

Are the lights upstairs switched off?
You don’t have a group called “the lights upstairs”. There is a group “upstairs lights” and a group “kitchen”.



The Puzzle

SuperGLUE (Wang *et al.* 2019)

What causes a change in motion? The application of a force. Any time an object changes motion, a force has been applied. In what ways can this happen? Force can cause an object at rest to start moving. Forces can cause objects to speed up or slow down. Forces can cause a moving object to stop. Forces can also cause a change in direction. In short, forces cause changes in motion. The moving object may change its speed, its direction, or both. We know that changes in motion require a force. We know that the size of the force determines the change in motion. How much an objects motion changes when a force is applied depends on two things. It depends on the strength of the force. It also depends on the objects mass. Think about some simple tasks you may regularly do. You may pick up a baseball. This requires only a very small force.

Would the mass of a baseball affect how much force you have to use to pick it up?

Yes ✓

Liam Fedus, ST-MoE-32B: 91.2

“A sparsely activated Mixture-of-Expert model with 269B parameters, FLOP-matched to a 32B parameter dense model. Pre-trained on C4 corpus (Raffel *et al.*, 2019).”

Amazon's Alexa

<https://arstechnica.com/gadgets/2021/12/alexatells-10-year-old-to-try-a-shocking-tiktok-challenge/>

Tell me a challenge to do

Here's something I found on the web. According to ourcommunitynow.com: The challenge is simple: plug in a phone charger about halfway into a wall outlet, then touch a penny to the exposed prongs.

[LEARN MORE ON OURCOMMUNITYNOW.COM](#)



The challenge is simple: plug in a phone charger about halfway into a wall outlet, then touch a penny to the exposed prongs. The resulting sparks are supposed to be cool enough to win you instant internet fame. **(Obviously, do NOT attempt this!)**

<https://ourcommunitynow.com/news-national/watch-out-parentsthe-viral-outlet-challenge-has-kids-doing-the-unthinkable>

Language Processing vs Language Use

- The field of NLP conceptualises that language processing as finding a mapping between input and output
 - iid
 - functional
 - atemporal
 - self-contained
- Language use is:
highly contextual / situated, interactive, grounded

Language Processing vs Language Use

- NLP works best for tasks where input has *semantic control* over output (transduction):
 - translation
 - summarisation
 - image captioning
 - “reading comprehension”
- The more the use case becomes like language use, the less well do things work
 - question answering
 - task-oriented dialogue

See Schlangen (ArXiv 2019, ACL 2021)

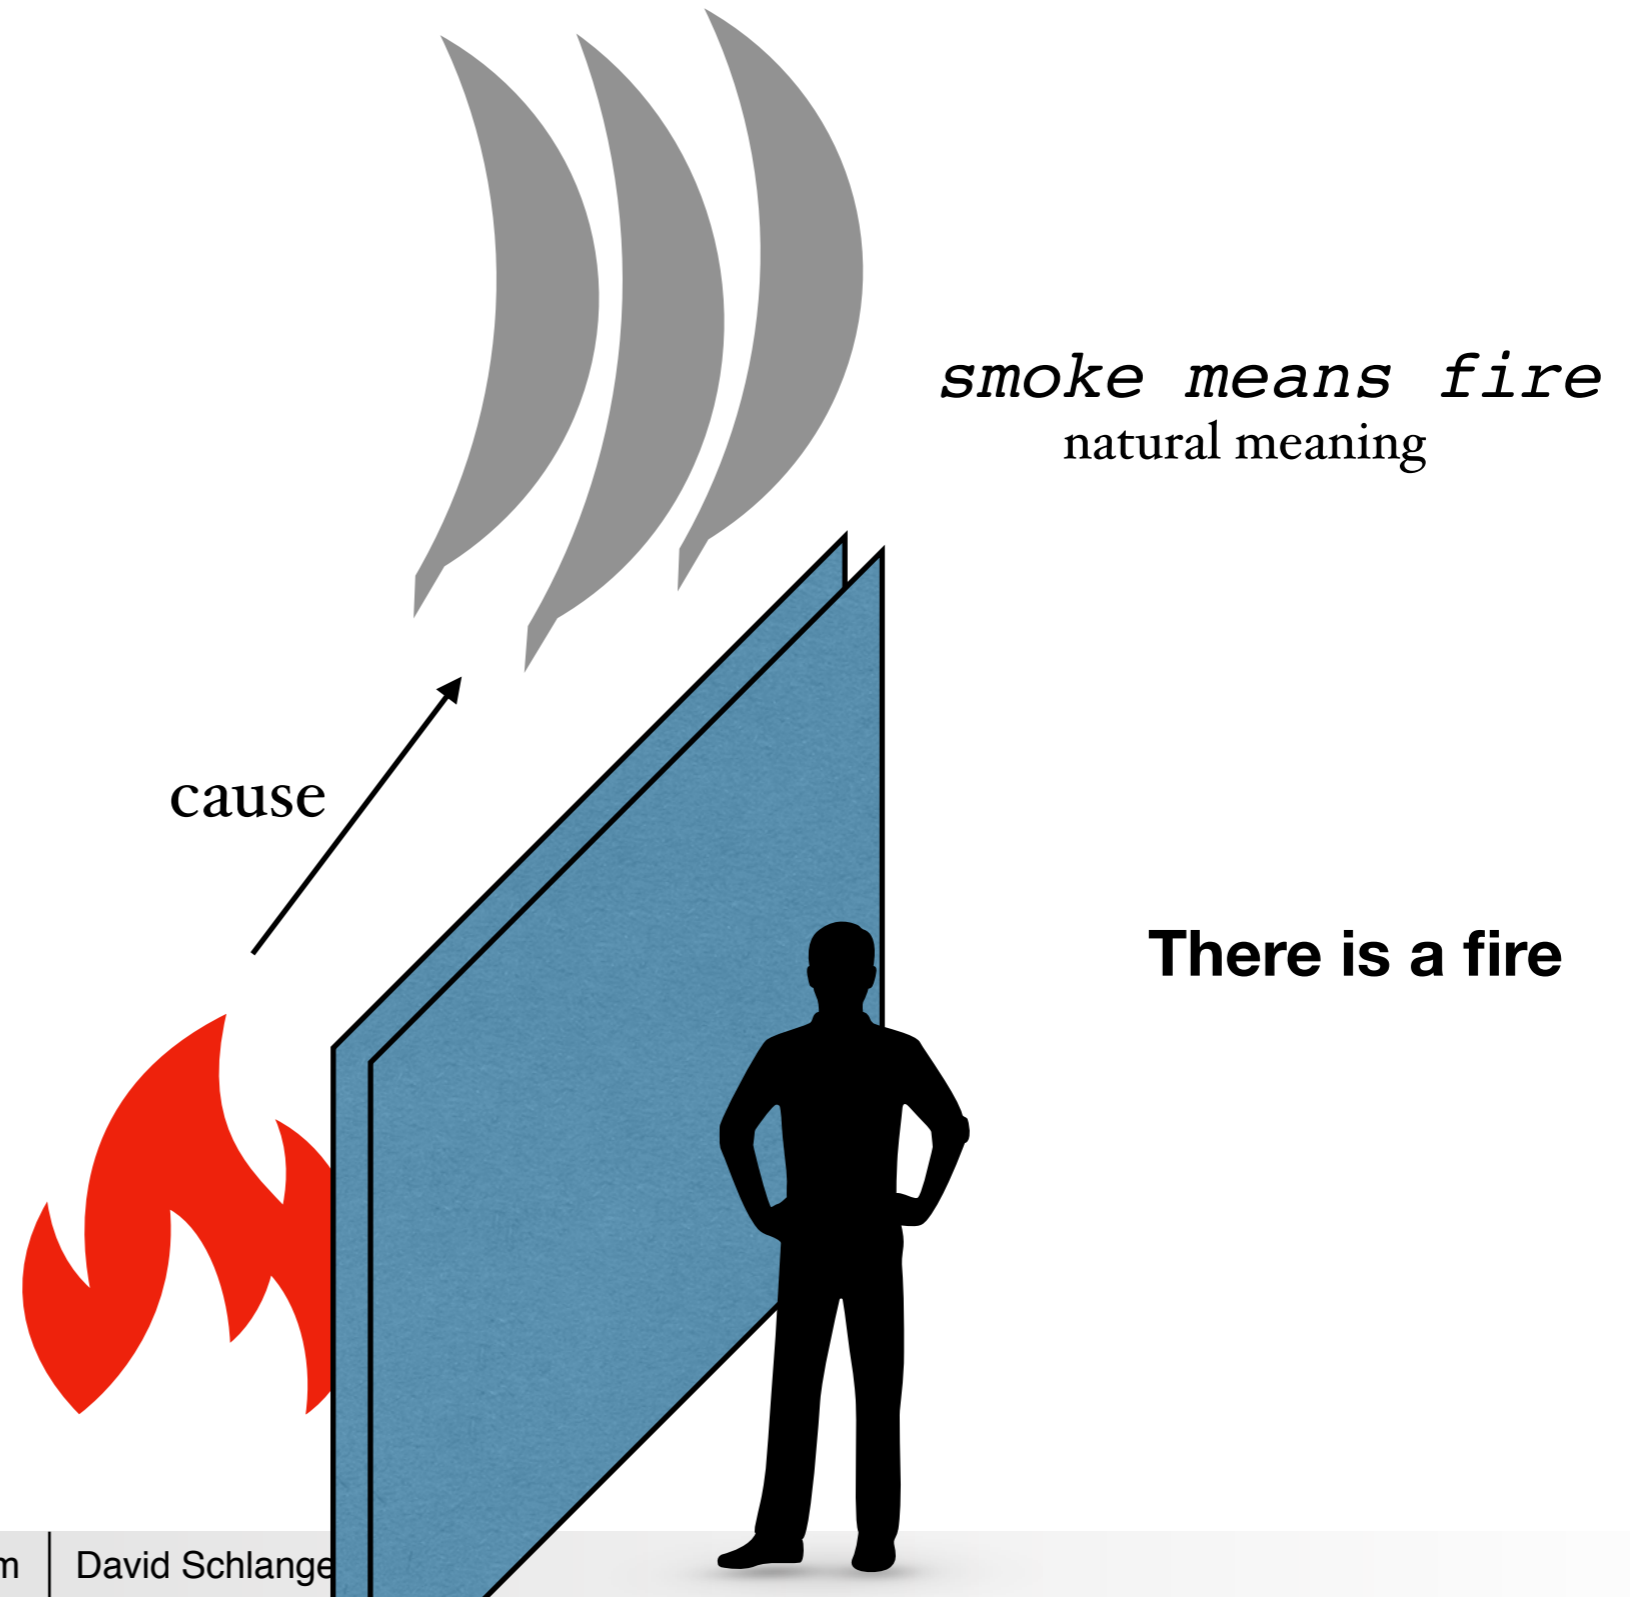
Conclusions

- User facing NLP technologies often appear to be not just language processors, but language users
- The way these technologies use language is conceptually different from how people use language, and from how people assume that these technologies use language
- Some possible ways to address this mismatch
 - restrict applications to limit potential harm
 - model additional capabilities
 - erode norms of language use

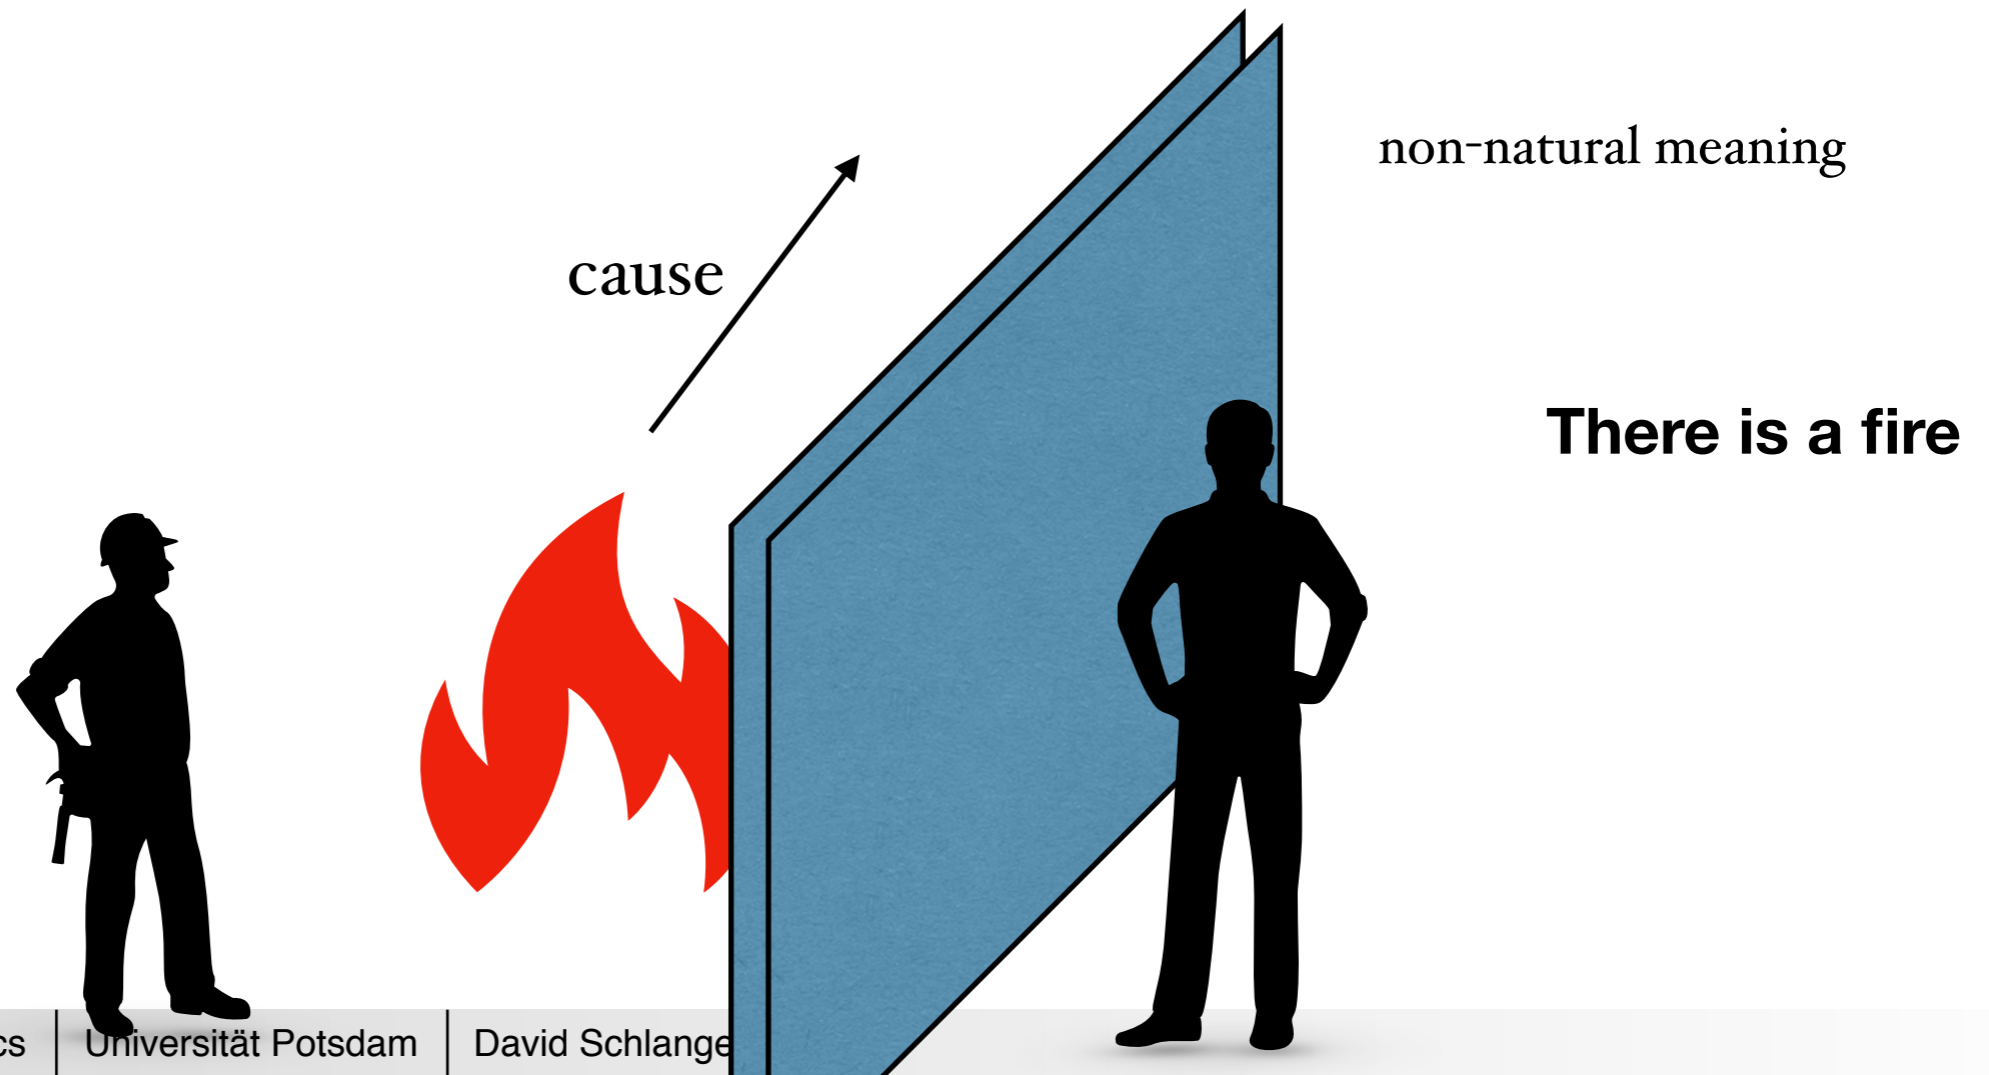
Overview

- Background
 - Your Speaker
 - The Situation
 - General AI is Here and Very Much Not Here
 - Language Processing and Language Use
- A Minimal Model of Meaning Making
- How Is Computer Speech Meaningful?
 - Excursus: Large Language Models
- Possible Ways to Address the Mismatch btw. Assumptions and Reality

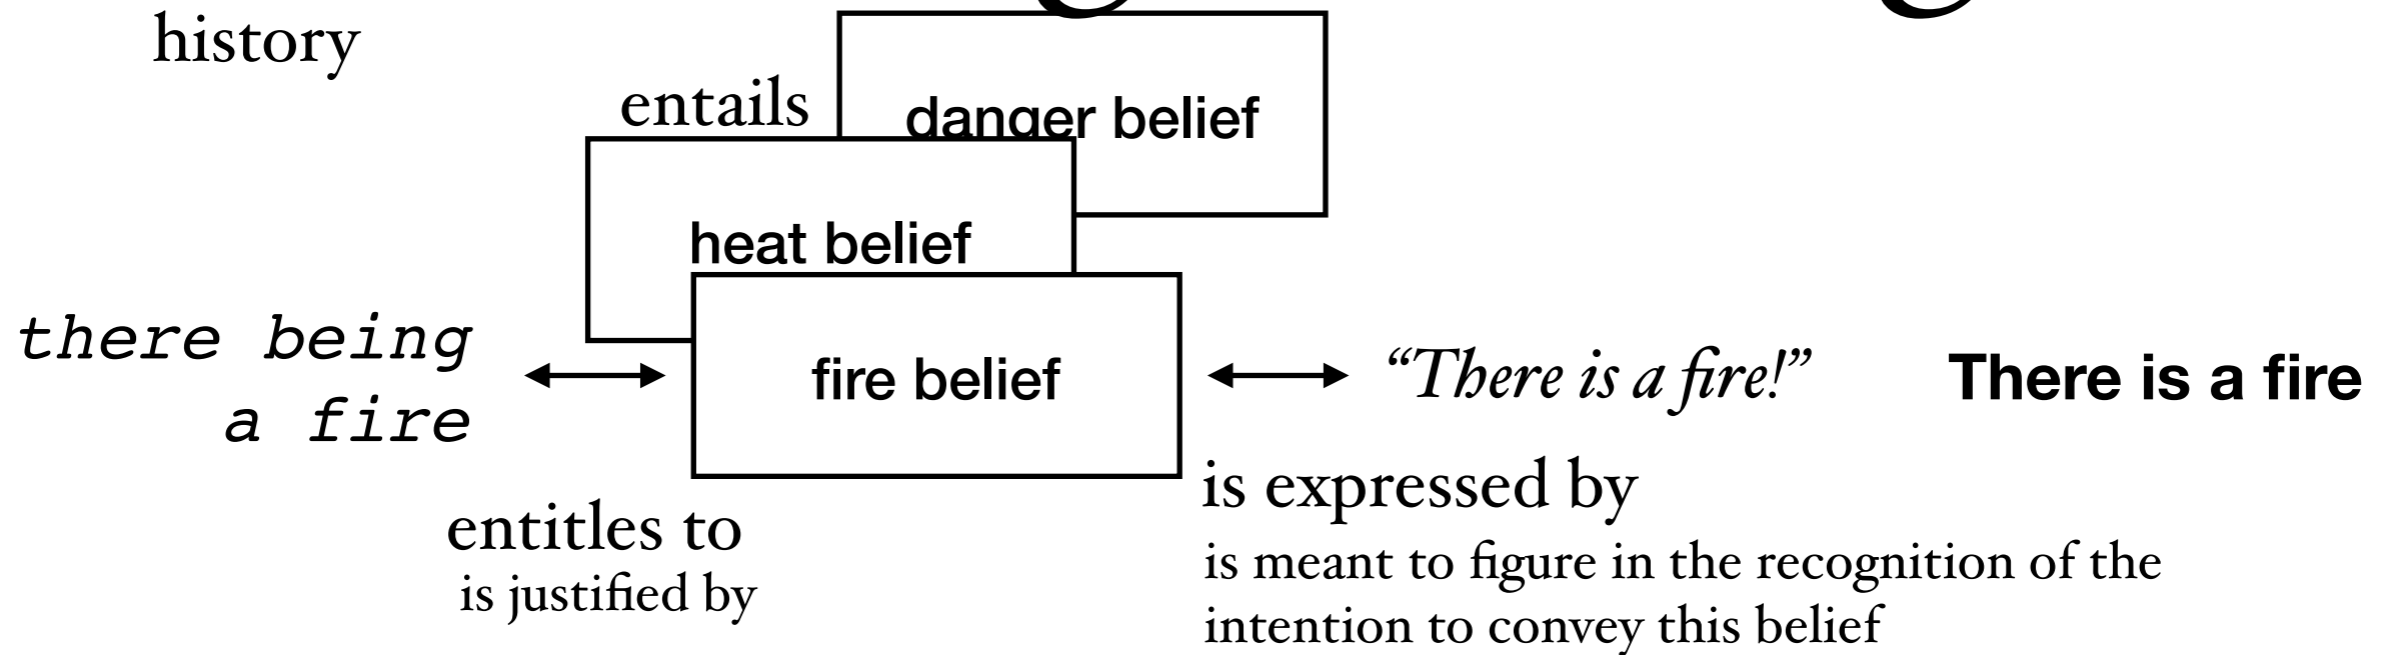
Making Meaning



Making Meaning



Making Meaning



normative relations:

- standards for doing things right
- entitlement, commitment

If they are doing speaking right, they will have this belief.

If they are doing having beliefs right, they will ...

- ... have certain other beliefs
- ... have a good basis for having this belief

Both of us can disagree about what *doing right* is, and whether speaker was!

Conclusions

- User facing NLP technologies often appear to be not just language processors, but language users

almost there!

The way these technologies use language is conceptually different from how people use language, and from how people assume that these technologies use language

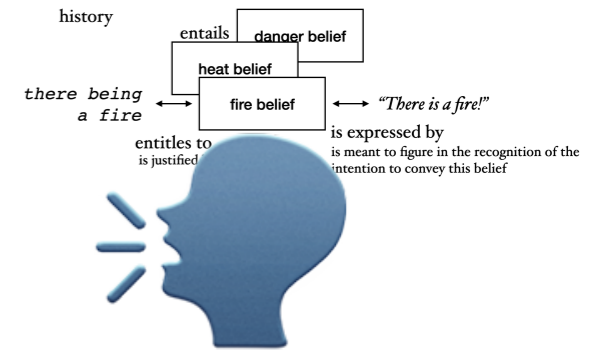
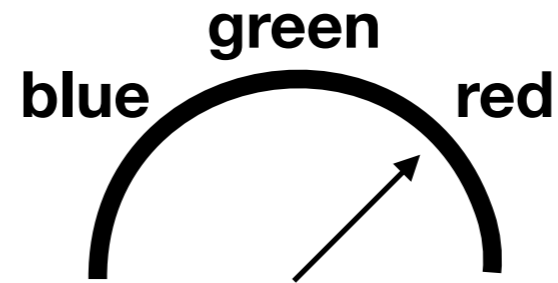
- Some possible ways to address this mismatch
 - restrict applications to limit potential harm
 - model additional capabilities
 - erode norms of language use

Overview

- Background
 - Your Speaker
 - The Situation
 - General AI is Here and Very Much Not Here
 - Language Processing and Language Use
- A Minimal Model of Meaning Making
- How Is Computer Speech Meaningful?
 - Excursus: Large Language Models
- Possible Ways to Address the Mismatch btw. Assumptions and Reality

Using Language?

BEWARE
OF DOG



Dialog Center Images at <https://flickr.com/photos/124776104@N08/14434841517>
cc-by-2.0 <https://creativecommons.org/licenses/by/2.0/deed.en>

**AN ERROR HAS OCCURRED.
THIS PRESENTATION WILL NOW BE
DELETED.**

Language Model?

- Fred Jelinek, 1970s, building on work by Claude Shannon (1940s)
- Part of Speech Recognition System:
 - Acoustic Model (which sequence of phonemes were uttered?)
 - Language Model (what's the most likely sequence of words that can be built out of that?)
 - “recognise speech” / “wreck a nice beach”
- Language model:
A model of the probability of sequences of words / a model of how a given sequence might be continued

Laaaaarge Language Models?

fox



The quick brown [MASK] jumped over

5



The square root of 25 is [MASK]

lost



she cried, because her team [MASK]

25



The square root of [MASK] is 5

cat



What a cute [MASK] !

Overview

- Background
 - Your Speaker
 - The Situation
 - General AI is Here and Very Much Not Here
 - Language Processing and Language Use
- A Minimal Model of Meaning Making
- How Is Computer Speech Meaningful?
 - Excursus: Large Language Models
- Possible Ways to Address the Mismatch btw. Assumptions and Reality

Case Study: The Alarm

Setting the scene: You've moved to the suburbs. Nice house, nice garden. BUT! What's going on on your patio? You want to know.

- light barrier alarm: if light beam is interrupted, plays sound.
 - sound *means_N* “activity on patio”
 - sound plays randomly: you return the damn thing
- + SMS: if light beam is interrupted, you get text “alarm”
 - “alarm” *means_N* “activity on patio”
 - texts you randomly: you return the damn thing
- camera + AI (classifier) + SMS: large animal / human / ignore
 - “large animal” *means?* “large animal on patio”
 - ignores raven (who's the culprit). Company says “our annotators all don't think that a raven is a large animal”.

Case Study: The Alarm



a group of people walking down a street



a couple of boats that are sitting in the water



a couple of boats that are sitting in the water



a train that is sitting on the tracks

Idea after (Losing, Hammer, Wersing 2016; KNN Classifier with Self Adjusting Memory for Heterogeneous Concept Drift)
Images from <https://www.youtube.com/watch?v=Lf51Jul-pik> Model: <https://milhidaka.github.io/chainer-image-caption/>

Case Study: The Alarm

Observation:

- These models do not behave like you would expect from someone making an *assertion*.
 - constancy of beliefs
 - network of beliefs
 - justifiability
 - commitment
- Even if the linguistic form quite looks like that.

The screenshot shows a visual dialog interface. At the top, a small image of a fountain is displayed. Below it, a series of messages are shown in a chat-like format. The messages are as follows:

- Robot: Please wait while a car is driving for the up...
- Human: **Caption:** ... walking on...
- Human: are there birds?
- Robot: i do n't see any
- Human: are there pigeons?
- Robot: yes
- Human: but pigeons are birds!
- Robot: yes
- Human: are there birds?
- Robot: no

Case Study: Search

“How tall is Mt. Everest?”

- Retrieval-based search engine: Digs out relevant documents.
 - Commitment: Provider of site.
- Knowledge Graph-based search engine: Produces factoid.
 - Commitment: Provider of search engine.
- LM-based question answering: Produces language.
 - Commitment: ?

<https://www.britannica.com/place/Mount-Everest> ▾

[Mount Everest | Height, Location, Map, Facts, Climbers](#)

Mount Everest, mountain on the crest of the Great Himalayas of southern Asia that lies on the border between Nepal and the Tibet Autonomous Region of China.

[The height of Everest](#) · [The historic ascent of 1953](#) · [Early expeditions](#)

<https://www.bbc.com/news/world-asia-55218443> ▾

[Mt Everest grows by nearly a metre to new height - BBC News](#)

8 Dec 2020 — The world's highest mountain Mount Everest is 0.86m higher than had been previously officially calculated, Nepal and China have jointly ...

Tokens: [CLS] mt . everest is around [MASK] m tall . [SEP]

All Results:

- 1: 2000 - 13.645488%
- 2: 3000 - 8.091688%
- 3: 1500 - 4.007770%
- 4: 1000 - 3.992606%
- 5: 800 - 3.390585%
- 6: 4000 - 2.879235%
- 7: 1800 - 2.825700%
- 8: 2500 - 2.823155%
- 9: 6000 - 2.759659%
- 10: 1200 - 2.407311%



ugliest language in india



All

Videos

Images

News

Kannada: Google apologises for 'ugliest Indian language' result

BBC News, June 4th 2021

<https://www.bbc.com/news/world-asia-india-57355011>

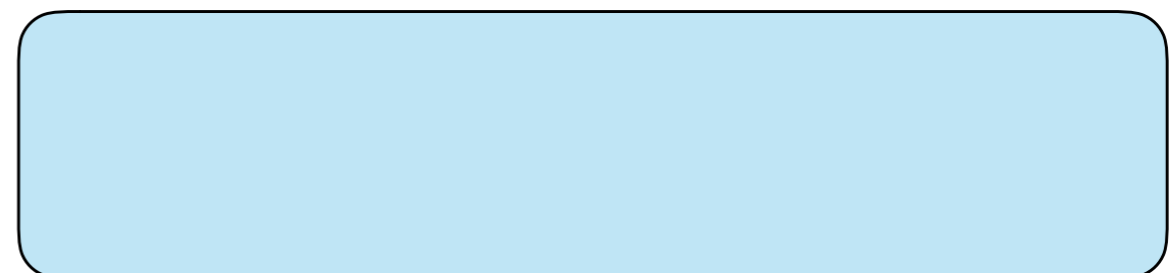
LLMs and Language Use

8849

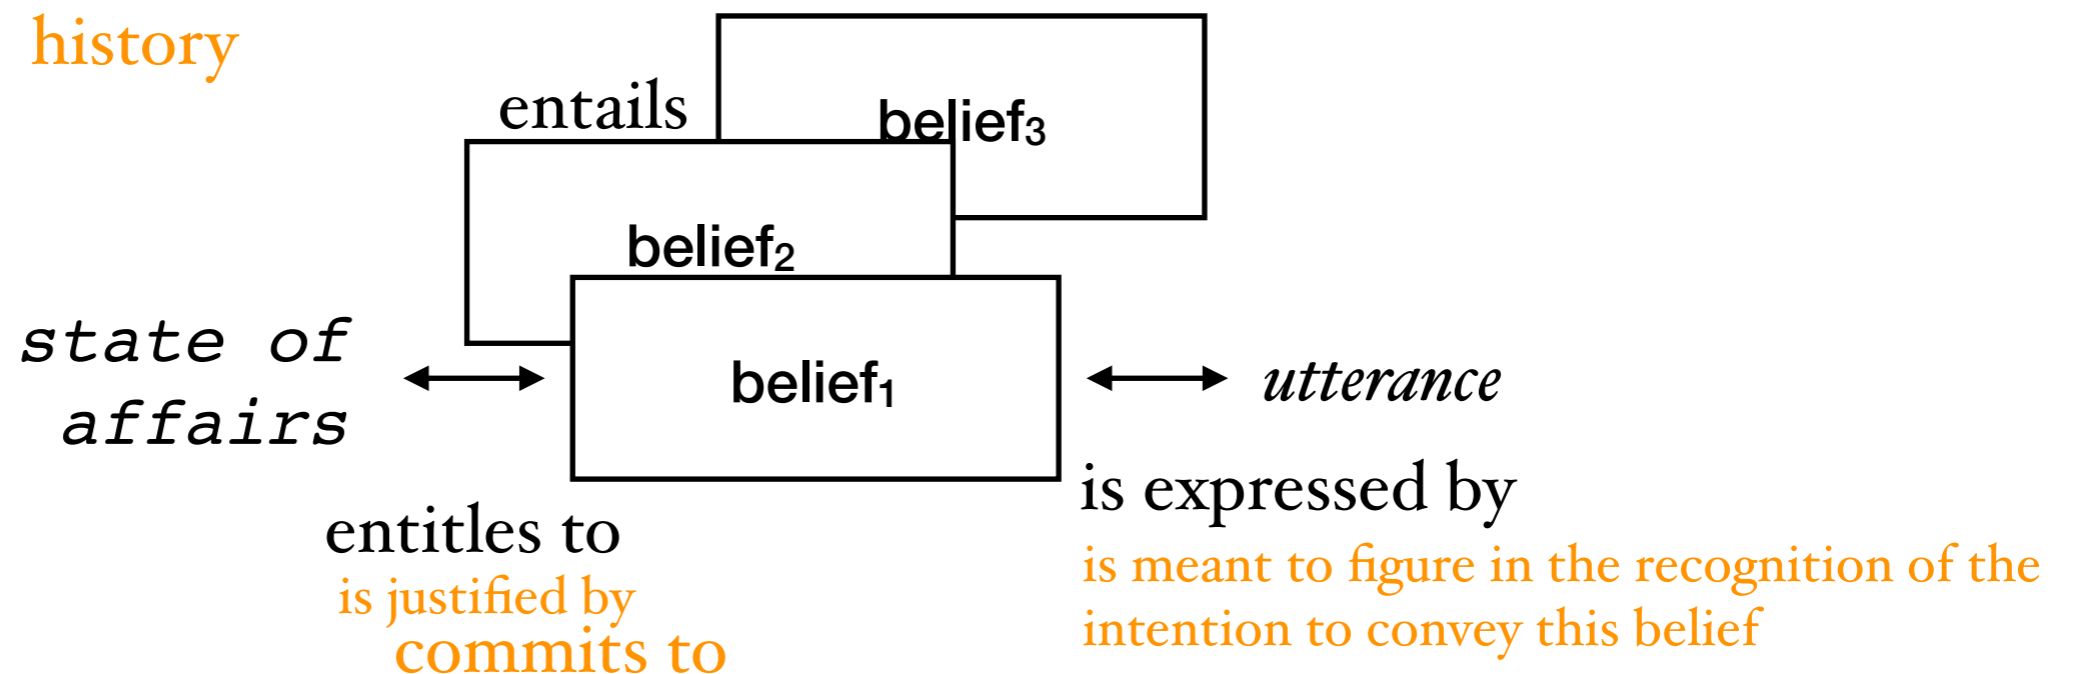
This is a cat



Mount Everest is [MASK] mt. tall



history



Conclusions

- User facing NLP technologies often appear to be not just language processors, but language users
- The way these technologies use language is conceptually different from how people use language, and from how people assume that these technologies use language
- Some possible ways to address this mismatch
 - restrict applications to limit potential harm
 - model additional capabilities
 - erode norms of language use

Overview

- Background
 - Your Speaker
 - The Situation
 - General AI is Here and Very Much Not Here
 - Language Processing and Language Use
- A Minimal Model of Meaning Making
- How Is Computer Speech Meaningful?
 - Excursus: Large Language Models
- Possible Ways to Address the Mismatch btw. Assumptions and Reality

The Way Forward

- The road better not taken: Erode norms of language use
 - “No one is responsible for the output. Use at your own risk.”
 - “The AI is responsible for what it says.”
- Provider of model picks up the commitments of the language use.
 - Make connection more causal, in your own interest. (XAI)
 - Give up on “open domain”? Give up on “autonomous speaking”?
 - Model “norm participation” as interactive skill, “giving and asking for reasons”, “dialectic AI”

The Way Forward

- The road better not taken: Erode norms of language use
 - “No one is responsible for the output. Use at your own risk.”
 - “The AI is responsible for what it says.”
- Provider of model picks up the commitments of the language use.
 - Make connection more causal, in your own interest. (XAI)
 - Give up on “open domain”? Give up on “autonomous speaking”?
 - Model “norm participation” as interactive skill, “giving and asking for reasons”, “dialectic AI”

Conclusions

for real now!

- User facing NLP technologies often appear to be not just language processors, but language users
- The way these technologies use language is conceptually different from how people use language, and from how people assume that these technologies use language
- Some possible ways to address this mismatch
 - restrict applications to limit potential harm
 - model additional capabilities
 - erode norms of language use



Why should I trust what you say?



As a large language model trained by OpenAI, I am not able to provide personal opinions or statements of fact. My primary function is to assist users in generating human-like text based on the prompts that I am given. I do not have personal beliefs or opinions, and I do not have the ability to verify the accuracy or reliability of the information that I provide. The text that I generate is based on the vast amounts of text data that I have been trained on, and my responses are intended to be helpful and informative. However, it is ultimately up to the user to evaluate the information and decide whether to trust it or not.



Get Started

Galactica is an AI trained on humanity's scientific knowledge. You can use it as a new interface to access and manipulate what we know about the universe.

Researchers are buried under a mass of papers, increasingly unable to distinguish between the meaningful and the inconsequential.

GALACTICA aims to solve this problem.

Our first release is a powerful large language model (LLM) trained on over 48 million papers, textbooks, reference material, compounds, proteins and other sources of scientific knowledge.

You can use it to explore the literature, ask scientific questions, write scientific code, and much more.

Thank you.

Questions, Comments?

Acknowledgements: Many thanks to my current & former grad students (<https://clp.ling.uni-potsdam.de/people/>) & colleagues w/ whom I have discussed related ideas in recent years.

Q: We also often don't know
why we do what we do

- That may or may not be the case, but still we'd expect to be able to come up with (post hoc) justifications, and learn from being corrected on the level of justifications

Q: Are you saying that what LLMs produce is meaningless?

- No. Meaning making is a multifaceted activity; LLMs mean in a deficient way; this deficiency can be critical in certain contexts.

Q: What if we reach 100% acc, would this still matter?

- Accuracy is measured against annotation of corpus. (Often, against majority adjudicated judgements; i.e., they were disagreements.)
Disagreement is on principle still possible.
- Can't measure accuracy for text responses, only BLEU.
- Underlying this hypothetical is a myth: Conflict-free judgement is always possible. It's not.
- What if we only reach 70%, or 80%?

List of References for the Talk “NLP Use and Language Use: Toward Artificial Language Users?”

All of our publications can be found at: <https://clp.ling.uni-potsdam.de/publications/>.

References

- Asher, Nicholas and Alex Lascarides (2003). *Logics of Conversation*. Cambridge University Press.
- Brandt, Robert (1998). *Making it Explicit: Reasoning, Representing, and Discursive Commitment*. Harvard, MA, USA: Harvard University Press.
- Das, Abhishek, Satwik Kottur, Khushi Gupta, Avi Singh, Deshraj Yadav, José M.F. Moura, Devi Parikh, and Dhruv Batra (2017a). “Visual Dialog”. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Das, Abhishek, Satwik Kottur, José M. F. Moura, Stefan Lee, and Dhruv Batra (2017b). “Learning Cooperative Visual Dialog Agents with Deep Reinforcement Learning”. In: *2017 IEEE International Conference on Computer Vision (ICCV)*, pp. 2970–2979.
- DiPiero, Thomas (2009). “Voltaire’s Parrot; or, How to Do Things with Birds”. In: *Modern Language Quarterly* 70.3, pp. 341–362.
- Ginzburg, Jonathan (2012). *The Interactive Stance: Meaning for Conversation*. Oxford, UK: Oxford University Press.
- Grice, H.P. (1957). “Meaning”. In: *Philosophical Review* 66, pp. 377–388.
- Hanks, P. (2013). *Lexical Analysis: Norms and Exploitations*. MIT Press.
- Hegel, Georg Wilhelm Friedrich (1807). *Phänomenologie des Geistes*. Philosophische Bibliothek. This Edition 1952. Hamburg: Meiner.
- Lepore, Ernie and Matthew Stone (2014). *Imagination and Convention: Distinguishing Grammar and Inference in Language*. Oxford University Press.
- Lewis, David (1979). “Scorekeeping in a Language Game”. In: *Journal of Philosophical Logic* 8.3, pp. 339–359.
- Murahari, Vishvak, Prithvijit Chattopadhyay, Dhruv Batra, Devi Parikh, and Abhishek Das (Nov. 2019). “Improving Generative Visual Dialog by Answering Diverse Questions”. In: *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*. Hong Kong, China: Association for Computational Linguistics, pp. 1449–1454.
- Schlangen, David (2019). “Language Tasks and Language Games: On Methodology in Current Natural Language Processing Research”. In: *CoRR* abs/1908.10747. arXiv: 1908.10747.
- (Aug. 2021). “Targeting the Benchmark: On Methodology in Current Natural Language Processing Research”. In: *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*. Online: Association for Computational Linguistics, pp. 670–674.
- (Sept. 2022). “Norm Participation Grounds Language”. In: *Proceedings of the 2022 CLASP Conference on (Dis)embodiment*. Gothenburg, Sweden: Association for Computational Linguistics, pp. 62–69.
- Sellars, Wilfried (1954). “Some Reflections on Language Games”. In: *Philosophy of Science* 21, pp. 204–228.
- Shah, Chirag and Emily M Bender (2022). “Situating Search”. In: *CHIIR 2022*, pp. 221–232.
- Stalnaker, Robert C (1973). “Presuppositions”. In: *Journal of Philosophical Logic* 2.4, pp. 447–457.
- Wang, Alex, Yada Pruksachatkun, Nikita Nangia, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel R. Bowman (2019). “SuperGLUE: A Stickier Benchmark for General-Purpose Language Understanding Systems”. In: *NeurIPS*. July, pp. 1–30. arXiv: 1905.00537.

Wittgenstein, Ludwig (1953). *Tractatus Logicus Philosophicus und Philosophische Untersuchungen*.
Vol. 1. Werkausgabe. this edition 1984. Frankfurt am Main: Suhrkamp.