What can natural language processing tell us about human language processing?

John R. Starr
Cornell University

I’m John R. Starr

Hi!

You can usually find me here:

Broadly, I'm interested in psycholinguistic modeling:

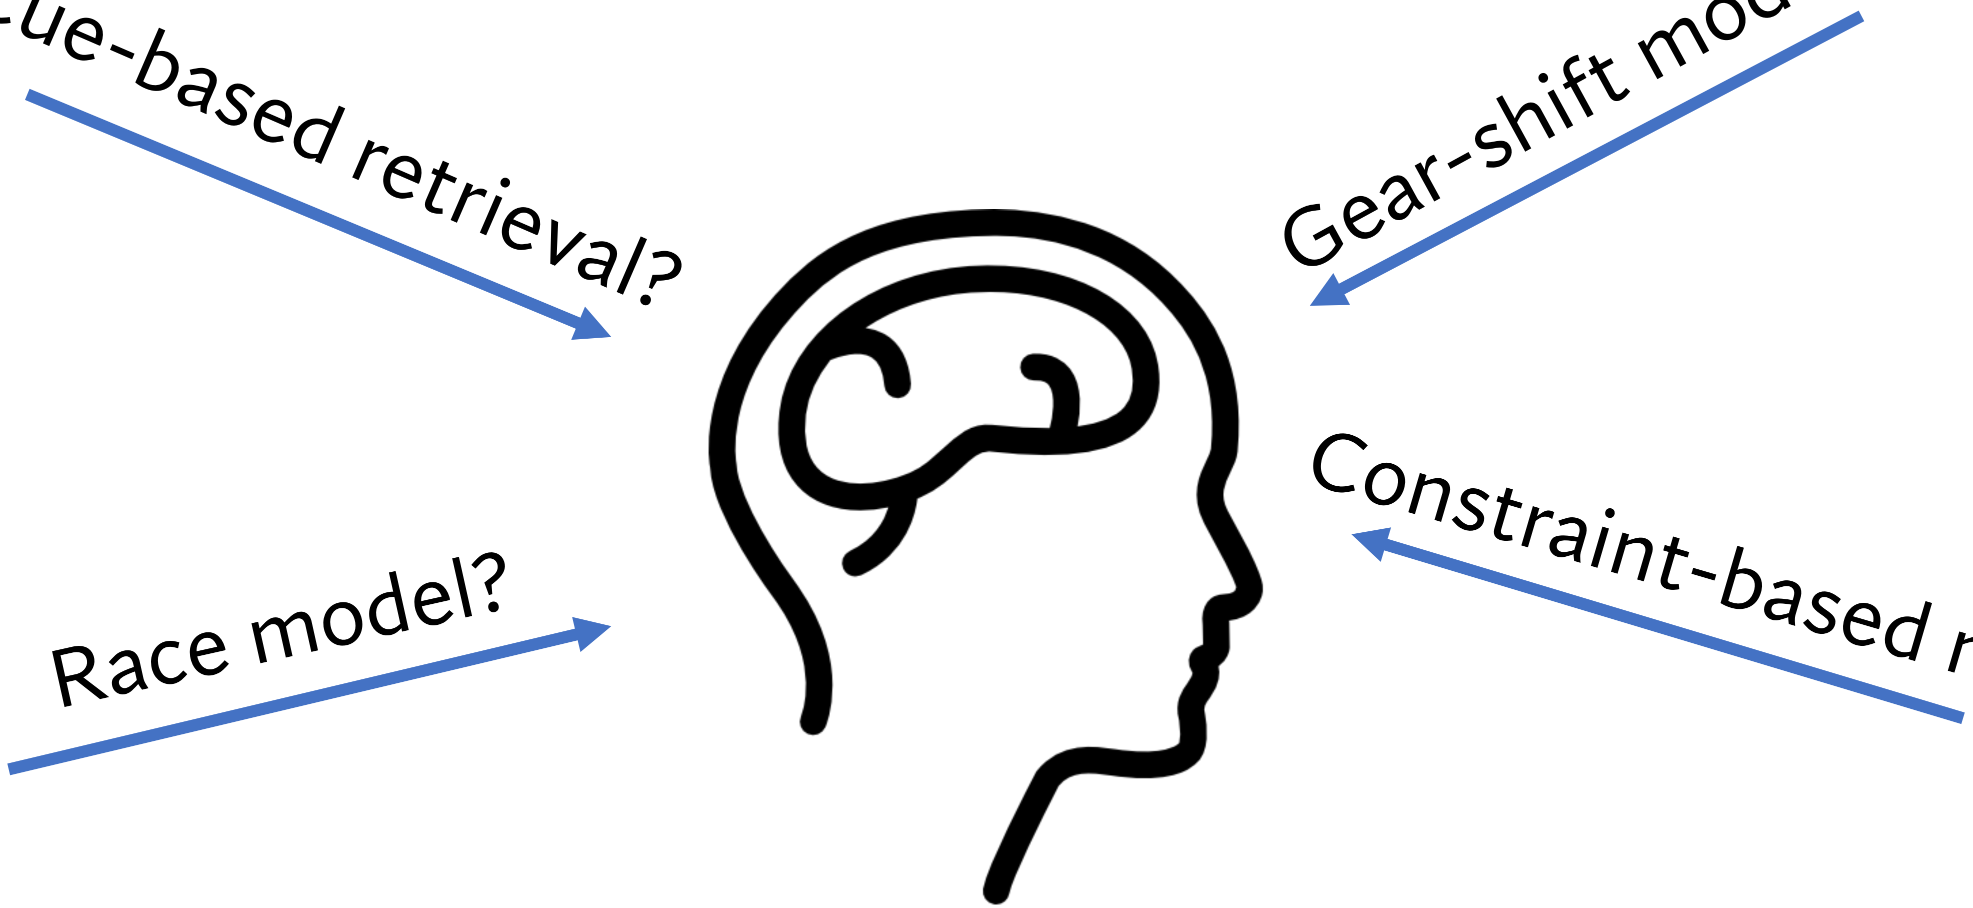

Going counterclockwise: Lewis & Vasishth (2005), van Gompel et al. (2000), Trueswell et al. (1994), Starr et al. (in prep)

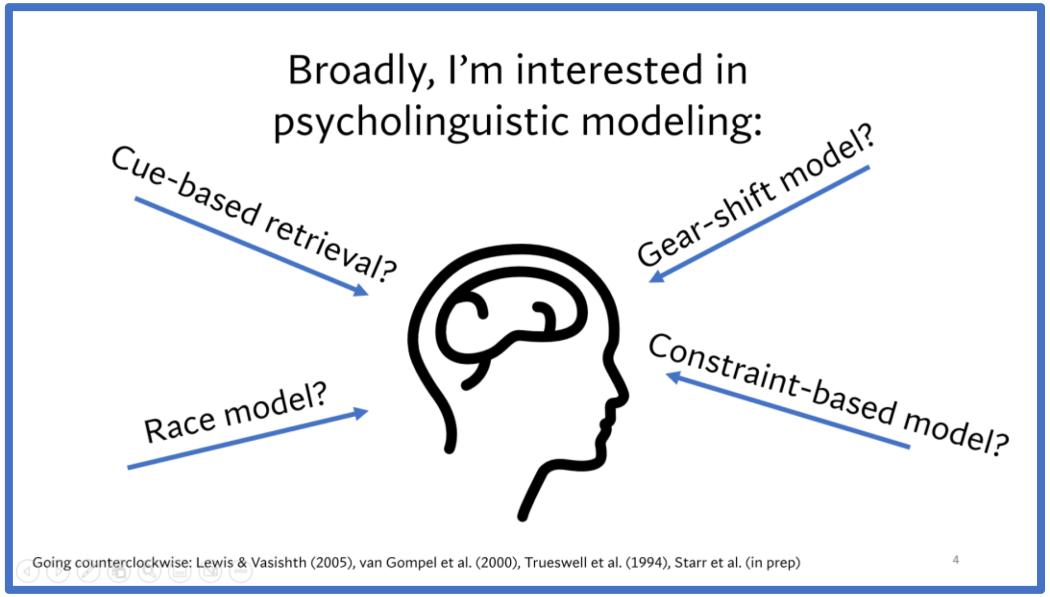

+ Computational Implementations

A walk-through of today's talk:

1. Breaking down the question

2. A short (and interactive!) tutorial on some essential NLP concepts involving neural networks…

3. … and adding back in the linguistics!

1. Breaking Down the Question

C.Psyd PLab

What can natural language processing tell us about human language processing?

John R. Starr
Cornell University

What is natural language processing?

- At a high level: The application of computational tools to linguistic data with the purpose of completing some task

- What it involves:
  1. Determining your task
  2. Collecting some data (often the hardest part!)
  3. Putting the data in a form that your model can handle
  4. Training your model
  5. Testing your model
  6. Evaluating your model

A sample task:

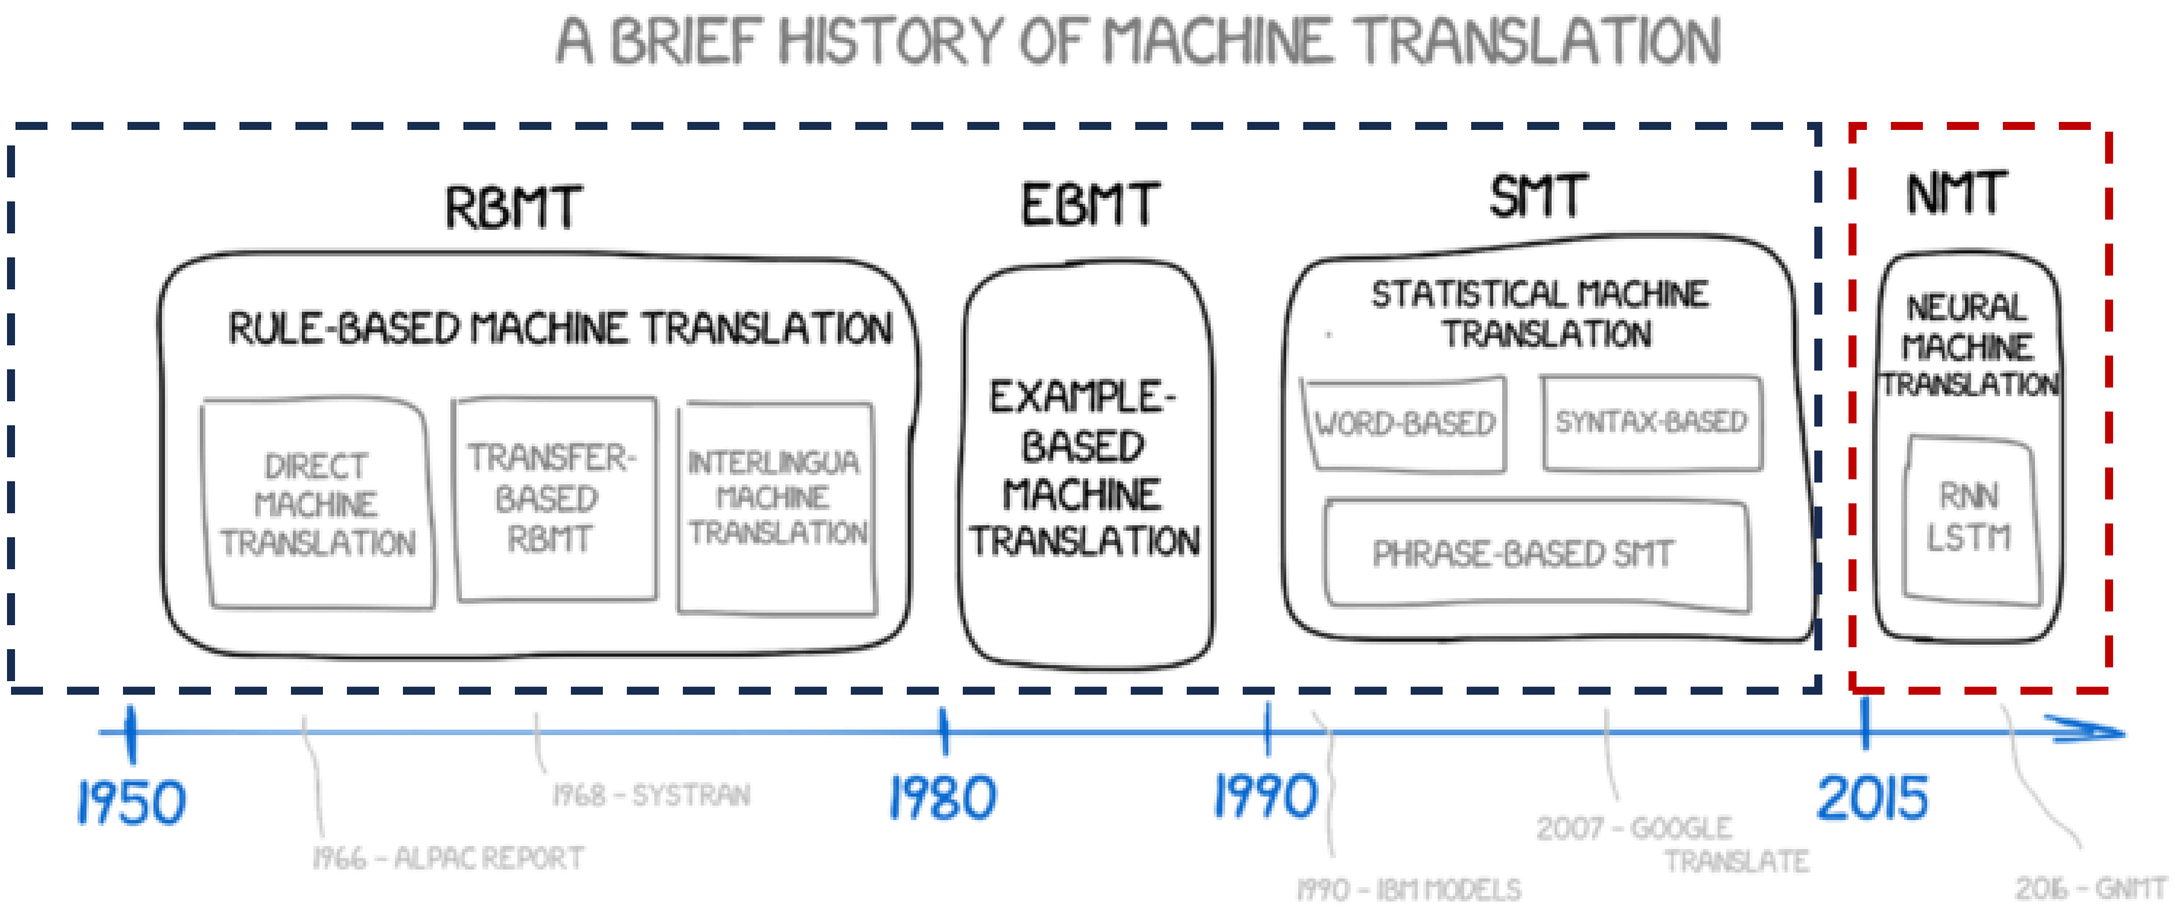

Thank you, Ilya Pestov!

Relevant NLP techniques for this talk (1)

- Distributional semantics (oftentimes learned from data)…!

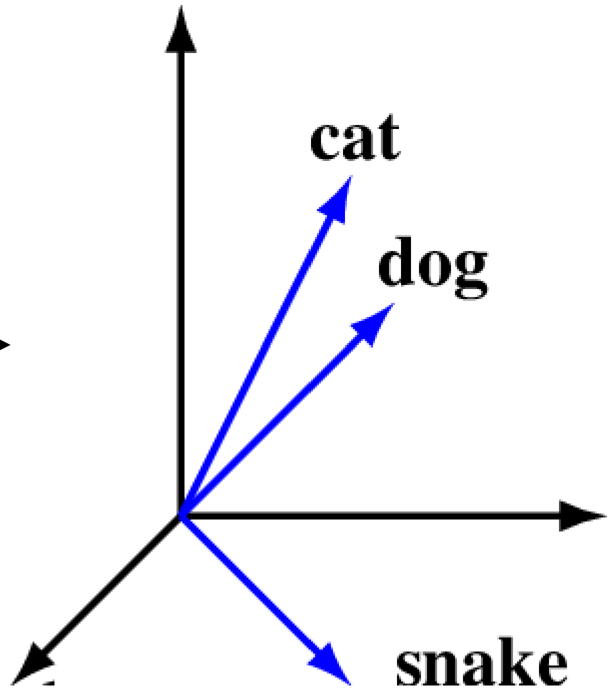

Relevant NLP techniques for this talk (2)

- Language modeling task…!

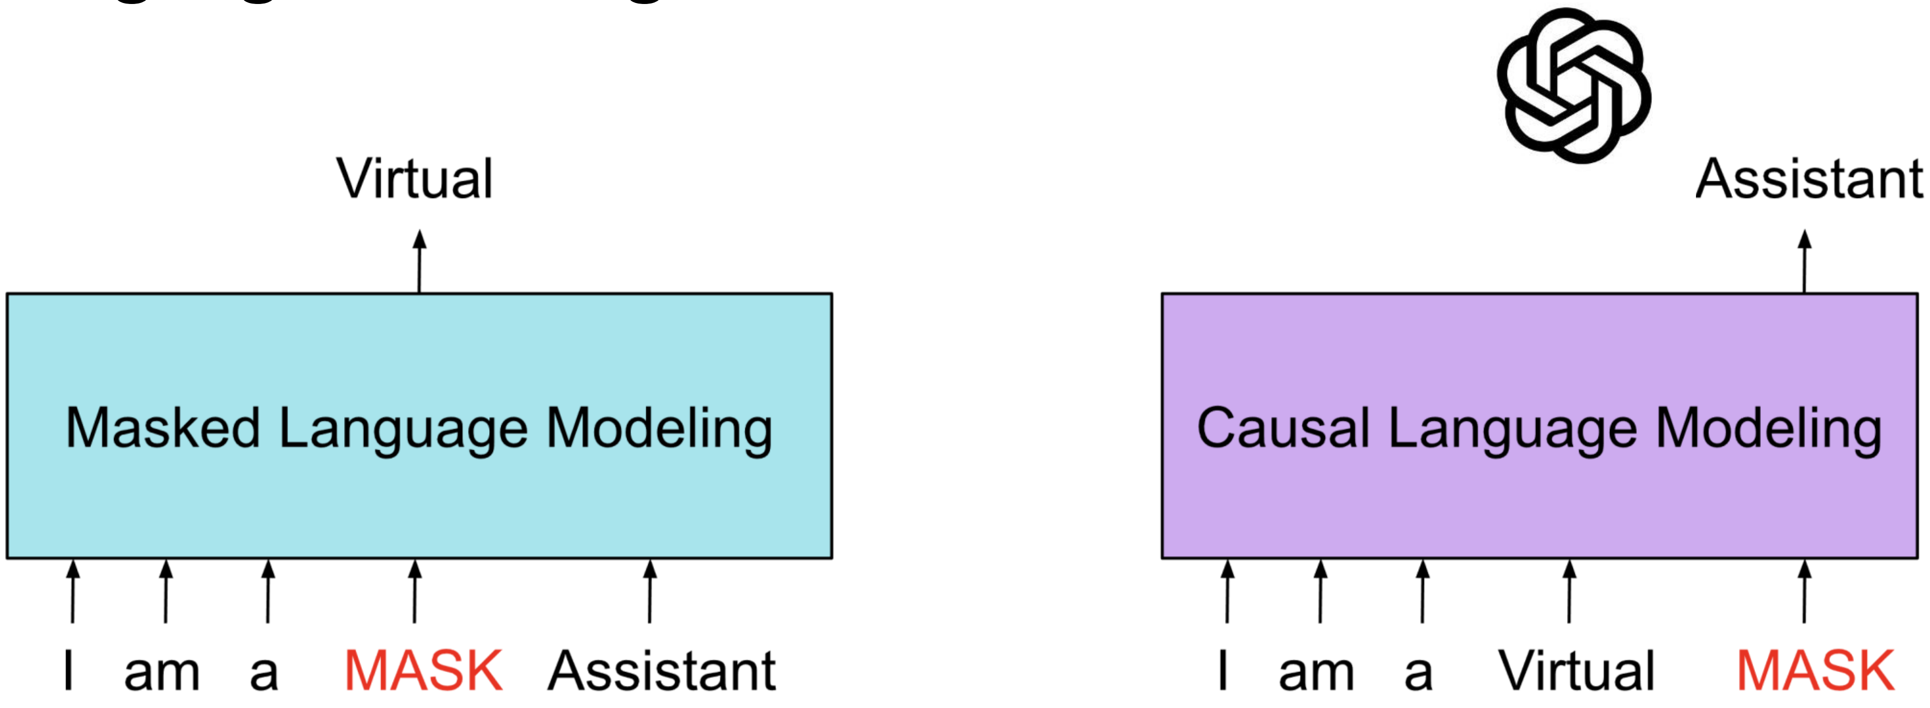

(a) (b)

~~What can natural language processing~~ tell us about human language processing?

John R. Starr
Cornell University

Who is us?

- Researchers!
- Non-researchers!
- … and other AI models??

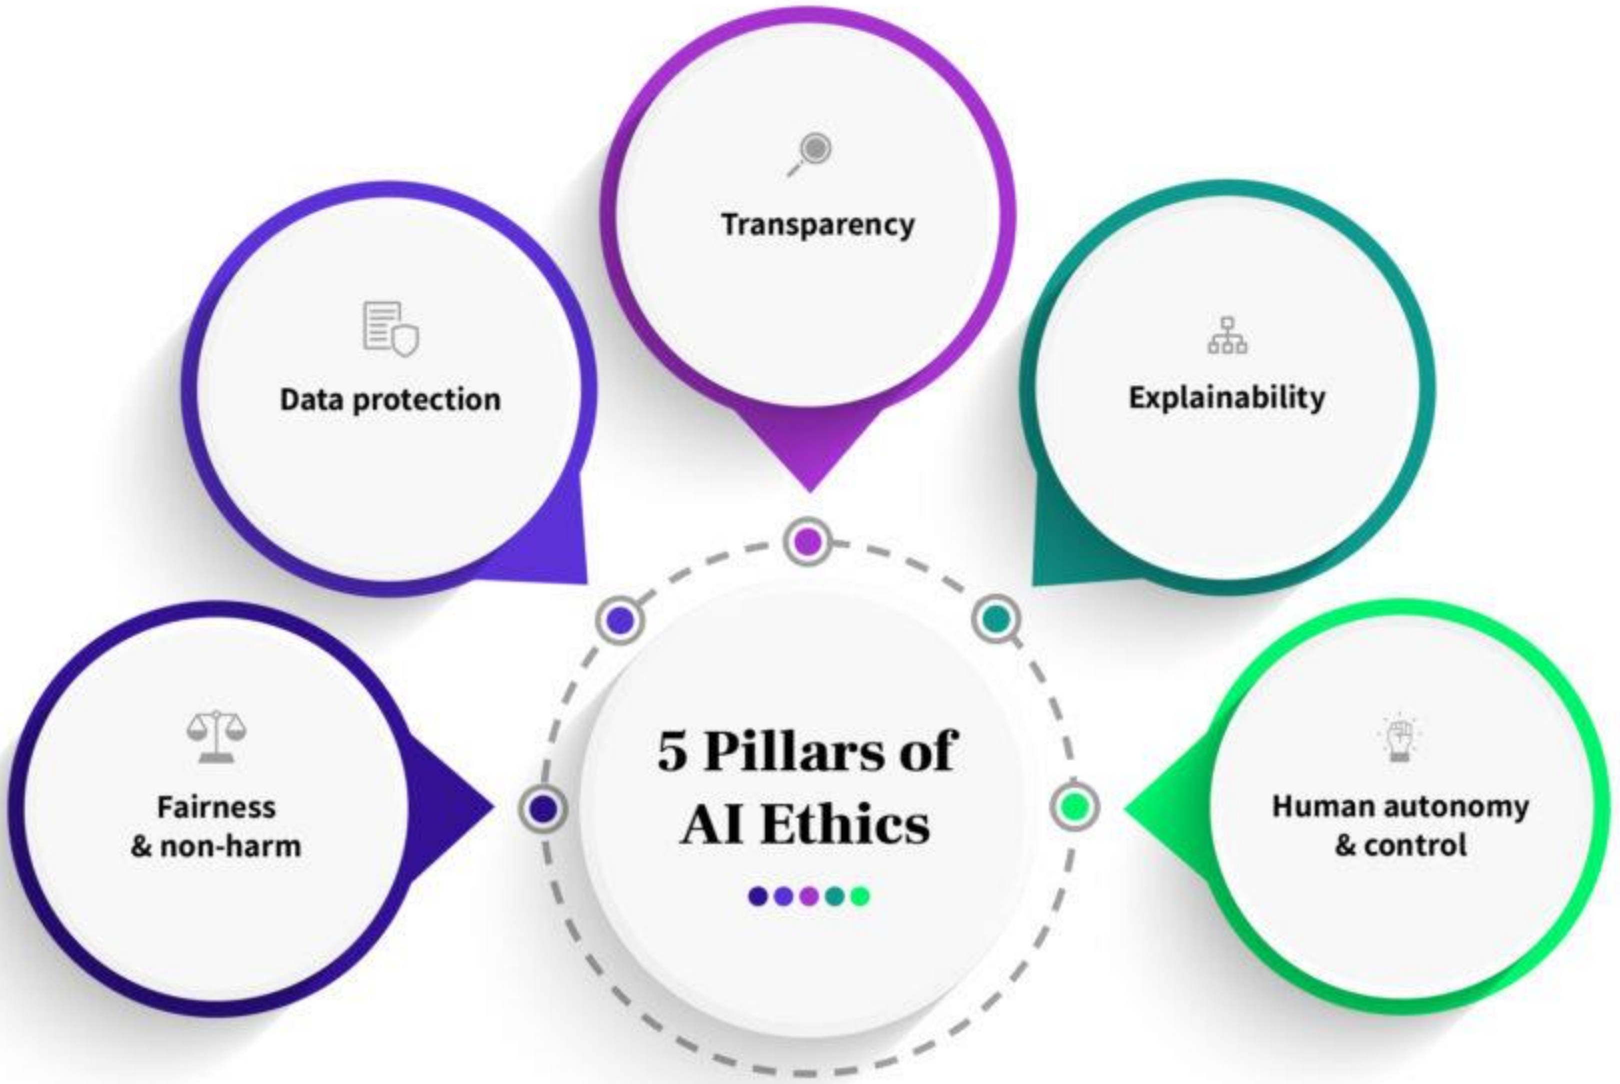
Transparency
Data protection
Explainability
5 Pillars of AI Ethics
Fairness & non-harm
Human autonomy & control

C.Psyd

PLab

~~What can natural language processing tell us about~~ human language processing?

John R. Starr
Cornell University

What is human language processing?

I saw the man in the valley

COUGH … doesn't this task look familiar? *COUGH*

NP Attachment

I [saw [the man [in the valley]]]

VP Attachment

I [saw [the man] [in the valley]]

Q: How does the language processor mediate difficult incremental input?

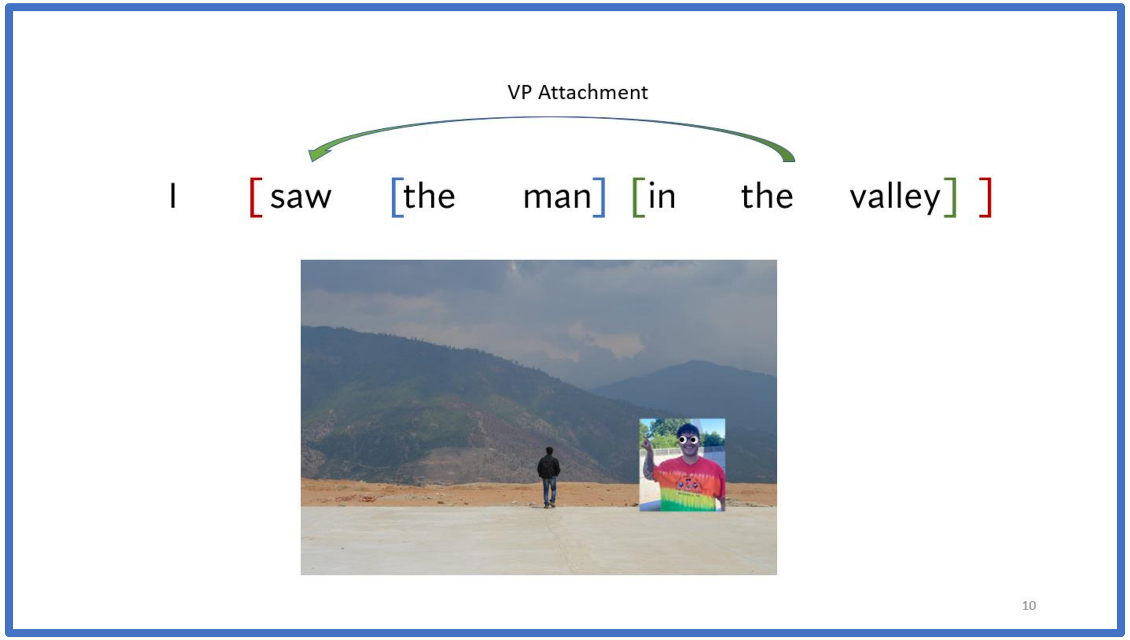

One option: A *Serial* Model of Processing

I [saw [the man] [in the valley]] from above.

Frazier (1979), Ferreira & Clifton (1986) Frazier & Rayner (1992)

One option: A *Serial* Model of Processing

I [saw [the man] [in the valley]] from above.

Reanalyzed as…

I [saw [the man [in the valley]]]

Frazier (1979), Ferreira & Clifton (1986) Frazier & Rayner (1992)

We've looked at difficult input caused by…

1. Attachment ambiguity:

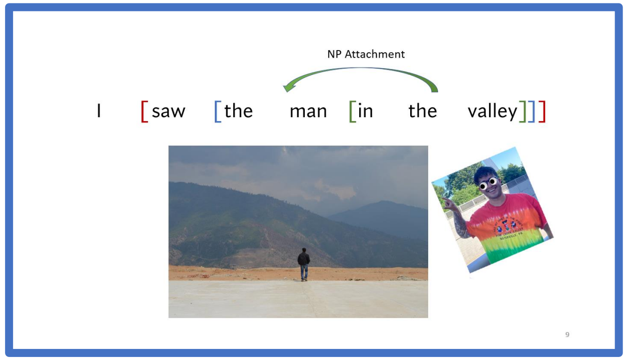

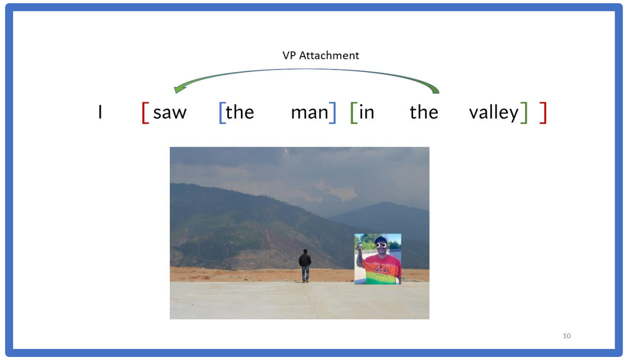

2. But what about non-syntactic sources?

Phonology can inform structure too!

Consider the string: “Let’s eat John”

Let’s eat John

Let’s eat, John

Situating this talk…

- Natural language processing (NLP) involves developing computational tools that perform language tasks…
- … but these tools should be informed by *human* contributions, given that humans will be using them…
- … and humans are the original language processors, meaning approaches from linguistics may be of value to NLP tools

… but why even compare natural language processing and human language processing?

1. Many of the tasks in NLP are based on tasks for human language processing (i.e. language modeling task!)
2. NLP models describe how much of language can be learned from statistical information alone
3. NLP models allow us to have a baseline comparison for human performance
4. If we can better align the performance of NLP tools on a task with the performance of humans on the same task, we will make tools that are more usable for humans

… and also:

- NLP is a burgeoning field – and linguists can contribute our expertise!

“Every time I fire a linguist, my accuracy goes up.”

- Frederick Jelinek, IBM

2. Interactive Tutorial Time!

https://colab.research.google.com/drive/1sZp6ThC-jtSfrvkI7QBxmSXncjwvEtcL?usp=sharing

3. Adding Back the Linguistics!

Let’s read…!

The horse raced past the barn fell.

This incremental task looks familiar!

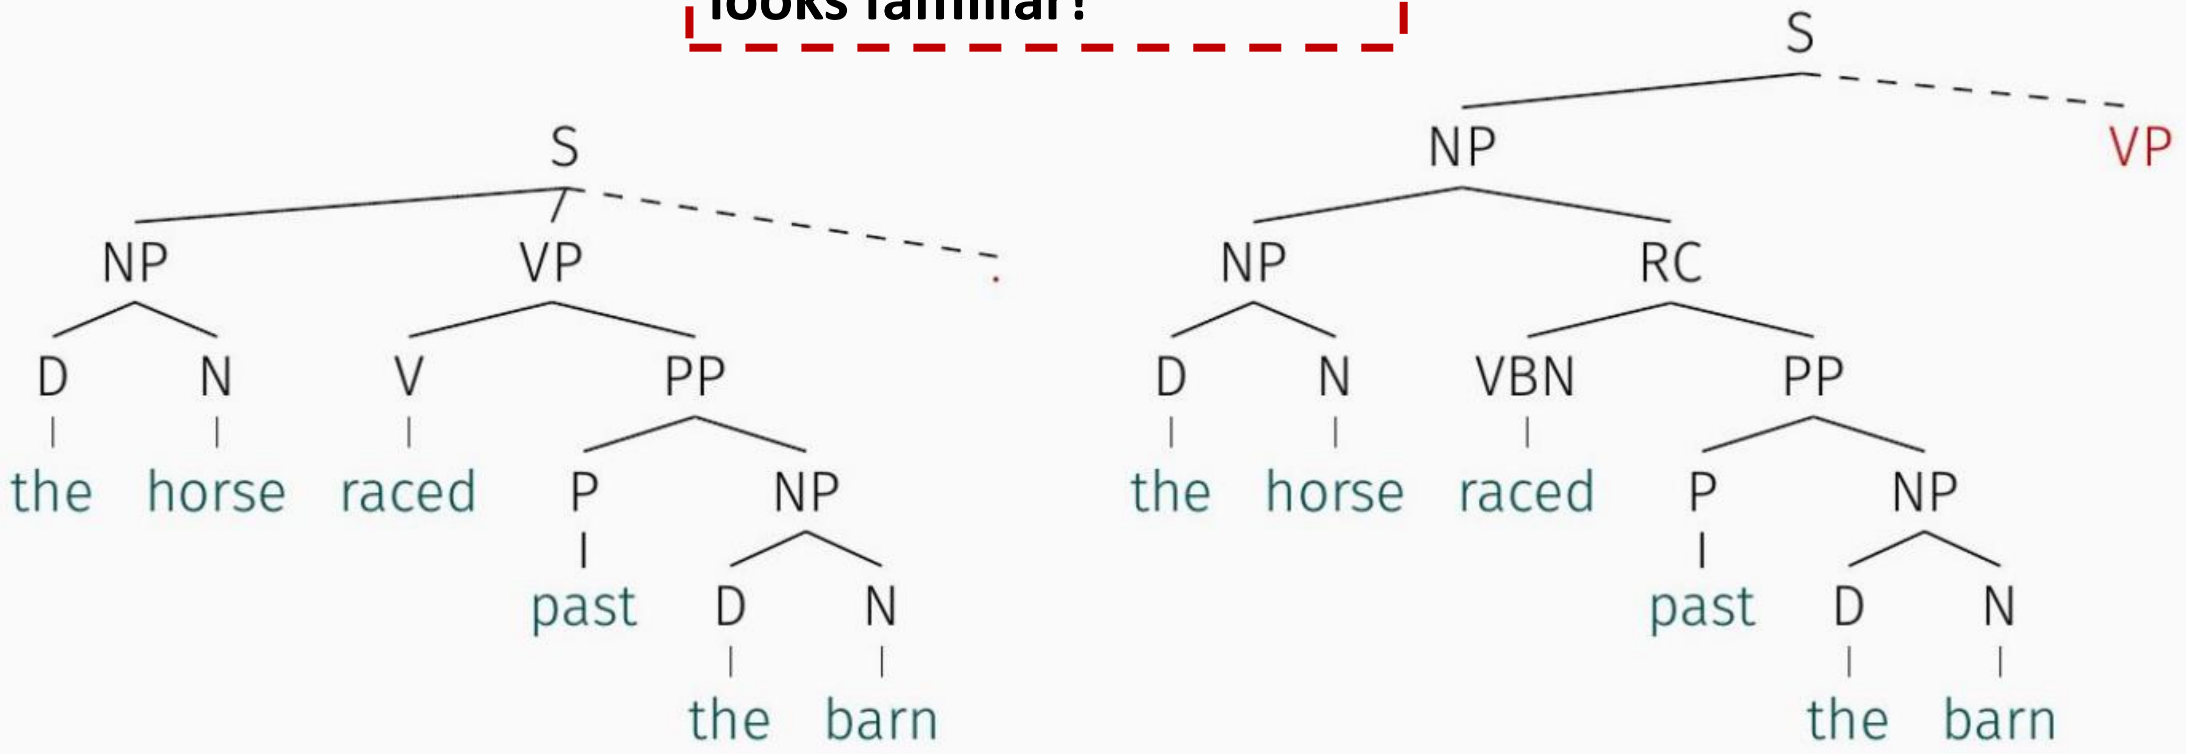

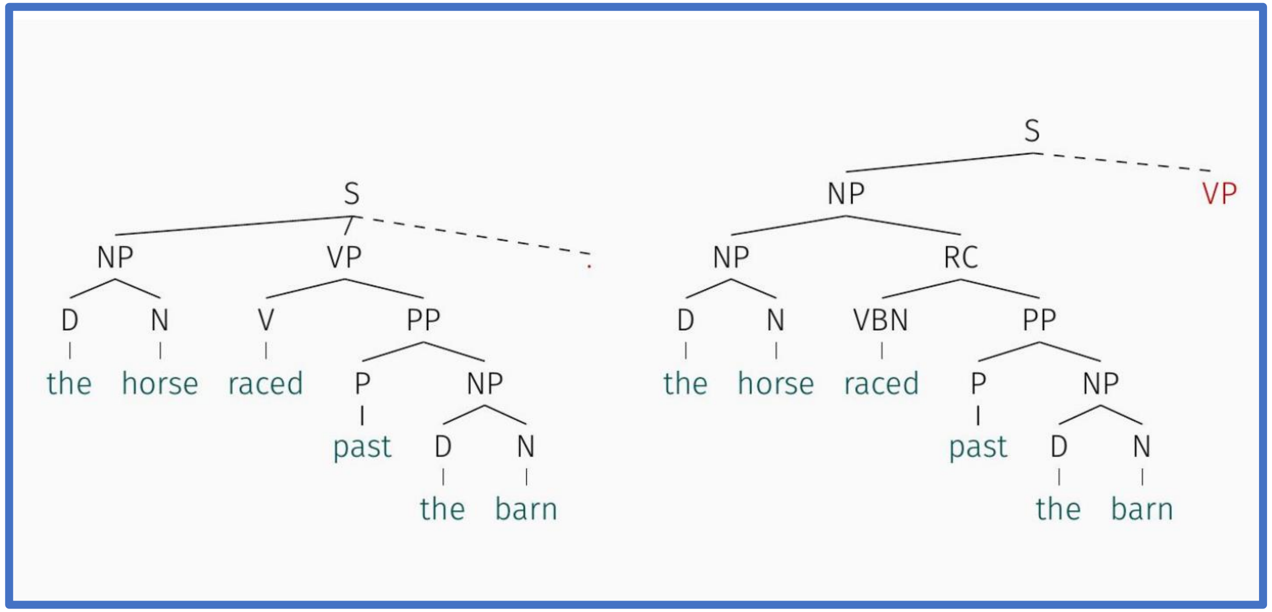

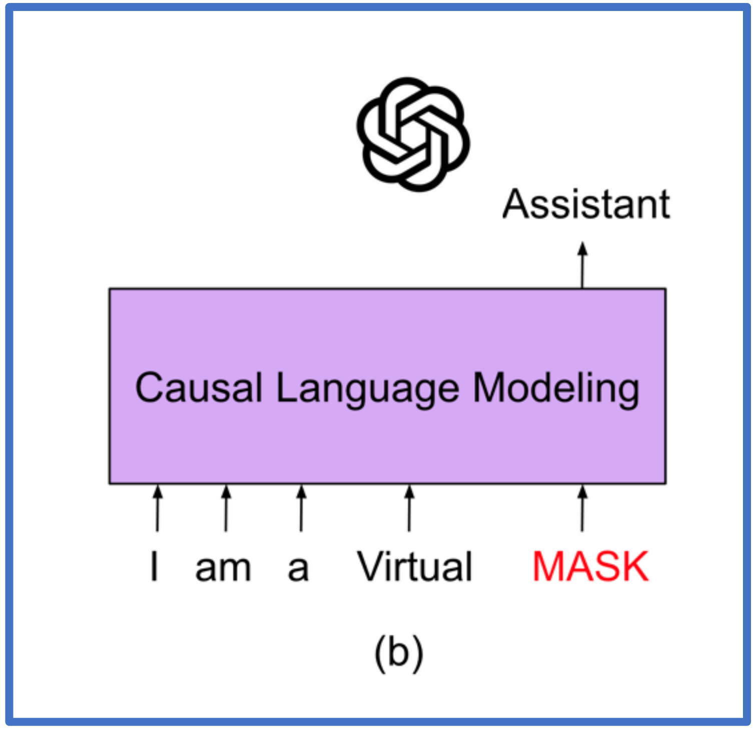

Human reading times ⟷ Word probability

What's the relationship between human reading times and language model probabilities?

Tal Linzen

COGNITIVE SCIENCE
A Multidisciplinary Journal

Cognitive Science 45 (2021) e12988

ISSN: 1551-6709 online
DOI: 10.1111/cogs.12988

Single-Stage Prediction Models Do Not Explain the Magnitude of Syntactic Disambiguation Difficulty

Marten van Schijndel, PhD,[a] Tal Linzen, PhD[b]

[a]Department of Linguistics, Cornell University
[b]Department of Linguistics and Center for Data Science, New York University

NP/S: The woman saw { the doctor wore a hat. / that the doctor wore a hat. }

- NP/S: The woman saw the doctor wore a hat.
- NP/S: The woman saw that the doctor wore a hat.

NP/S: The woman saw
- the doctor wore a hat.
- that the doctor wore a hat.

NP/Z: When the woman
- visited her nephew laughed loudly.
- visited, her nephew laughed loudly.

NP/S: The woman saw
- the doctor wore a hat.
- that the doctor wore a hat.

NP/Z: When the woman
- visited her nephew laughed loudly.
- visited, her nephew laughed loudly.

MV/RR: The horse
- raced past the barn fell.
- which was raced past the barn fell.

Predicted/empirical word-by-word garden path effects

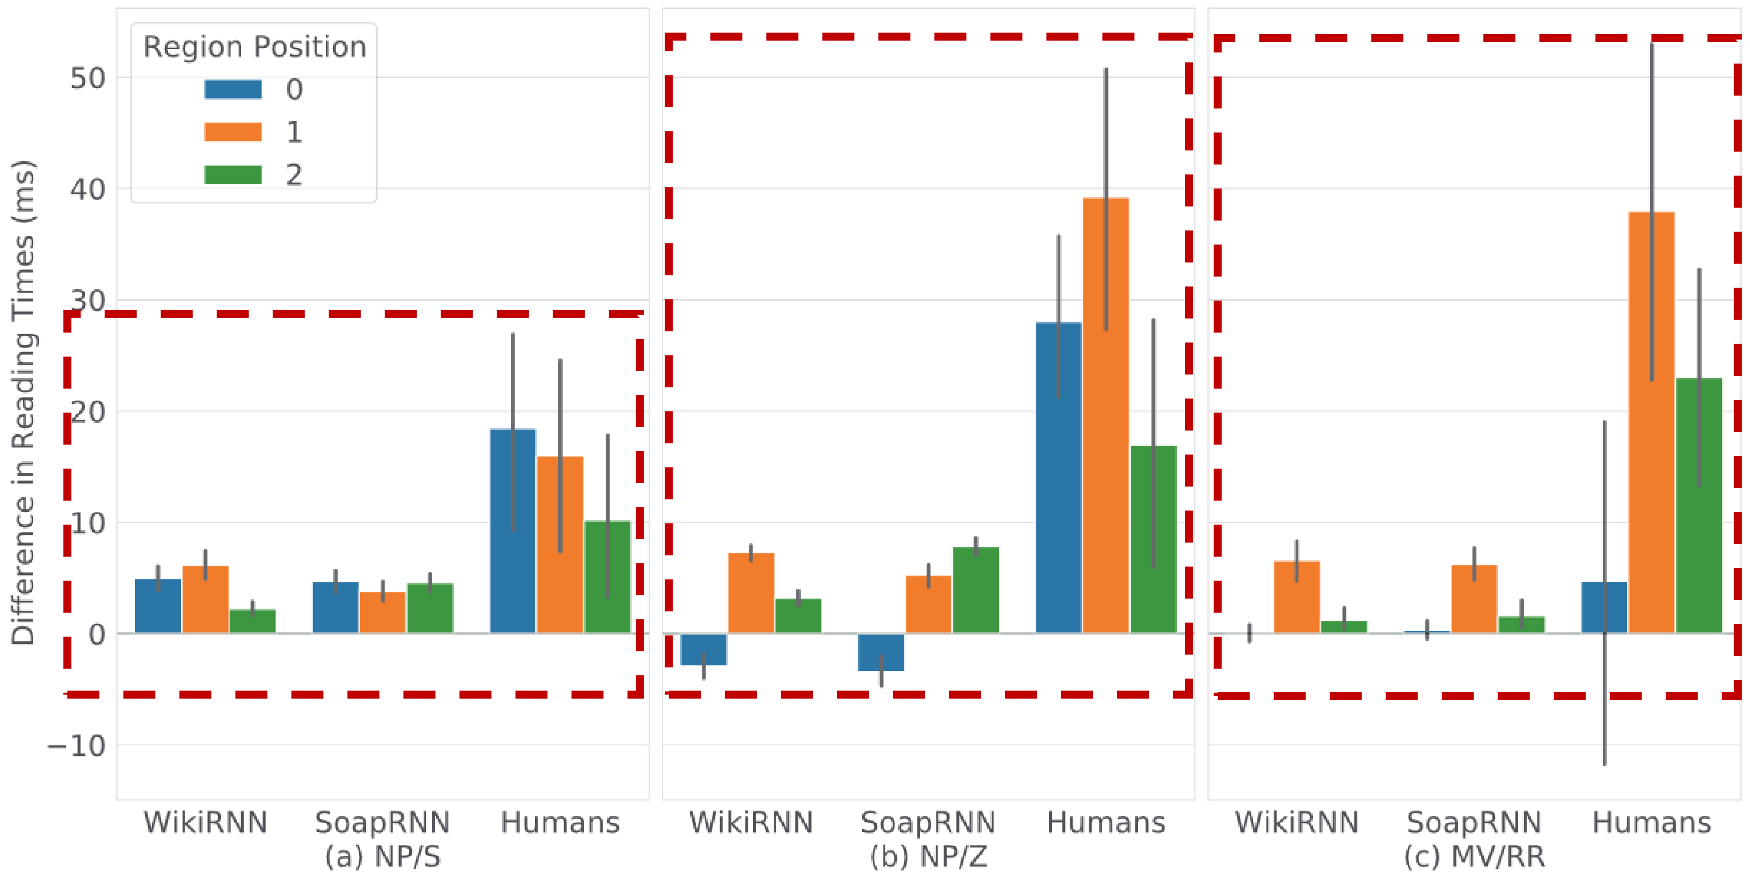

Fig. 2. Differences in word-by-word reading times between ambiguous and unambiguous sentences on the first, second, and third words of the disambiguating region, as predicted by the language models, compared to empirical reading times. The subplot shows the disambiguation region of: (a) ambiguous NP/S sentences compared to matched unambiguous controls (example (4) in the text); (b) ambiguous NP/Z sentences compared to matched unambiguous controls (example (5) in the text); (c) ambiguous MV/RR sentences compared to matched unambiguous controls (example (6) in the text). Error bars represent bootstrapped 95% confidence intervals.

van Schijndel & Linzen (2021) Summary

- Language models get that there should be an increase in reading times…

- … but they severely underpredict the magnitude of that increase in reading time!

What kind of data might we need to learn these kinds of phenomena?

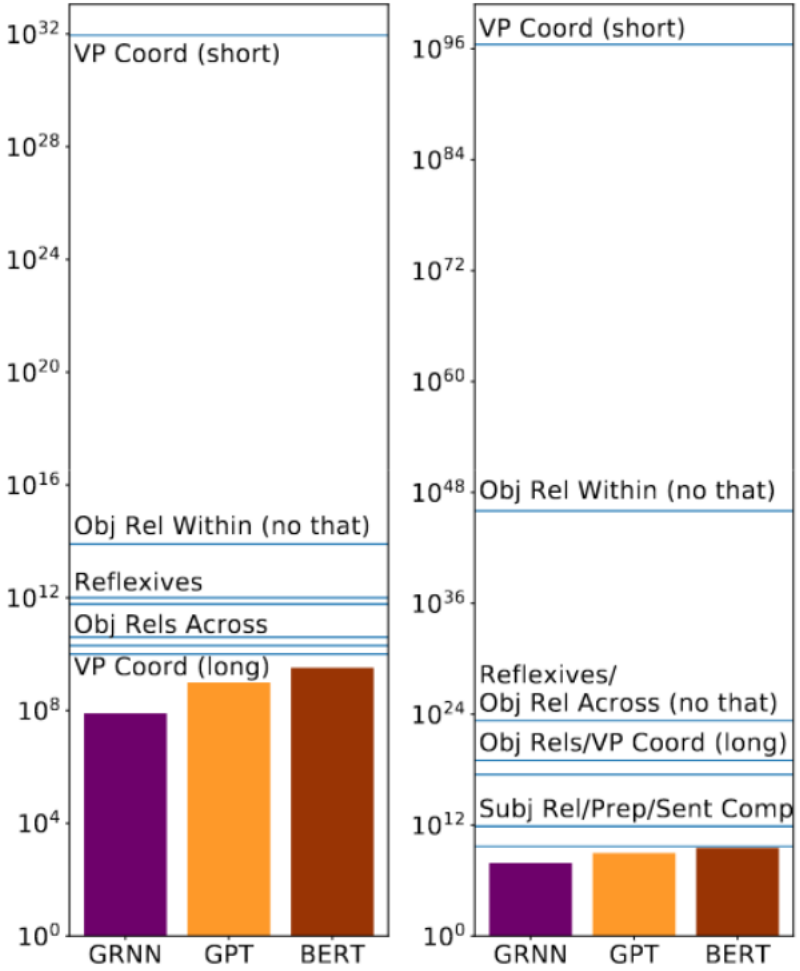

(a) Human-like (b) 99.99%

NNs fail at simple patterns like:

The authors laugh and reads books.

BERT was trained on 3e9 words

T5 was trained on 1e11 words

Would require 1e32 words to learn this structure as well as a human.

Do language model representations mimic human representations of language?

- Consider the phrase “gave in” in the following two sentences:

a. The teacher *gave in* to the student’s demands.
b. The exam that the teacher *gave in* class was difficult.

- A. is a “light verb construction”, where *gave in* acts as a unit.
- B. is a non-light verb construction,, where *gave in* do not act as a unit

Do language model representations mimic human representations of language?

- Theoretically, humans generate some differences in representations of these things… (Wittenberg et al. 2014)

- But do language model representations reflect these distinctions?

Some patterns emerge initially...

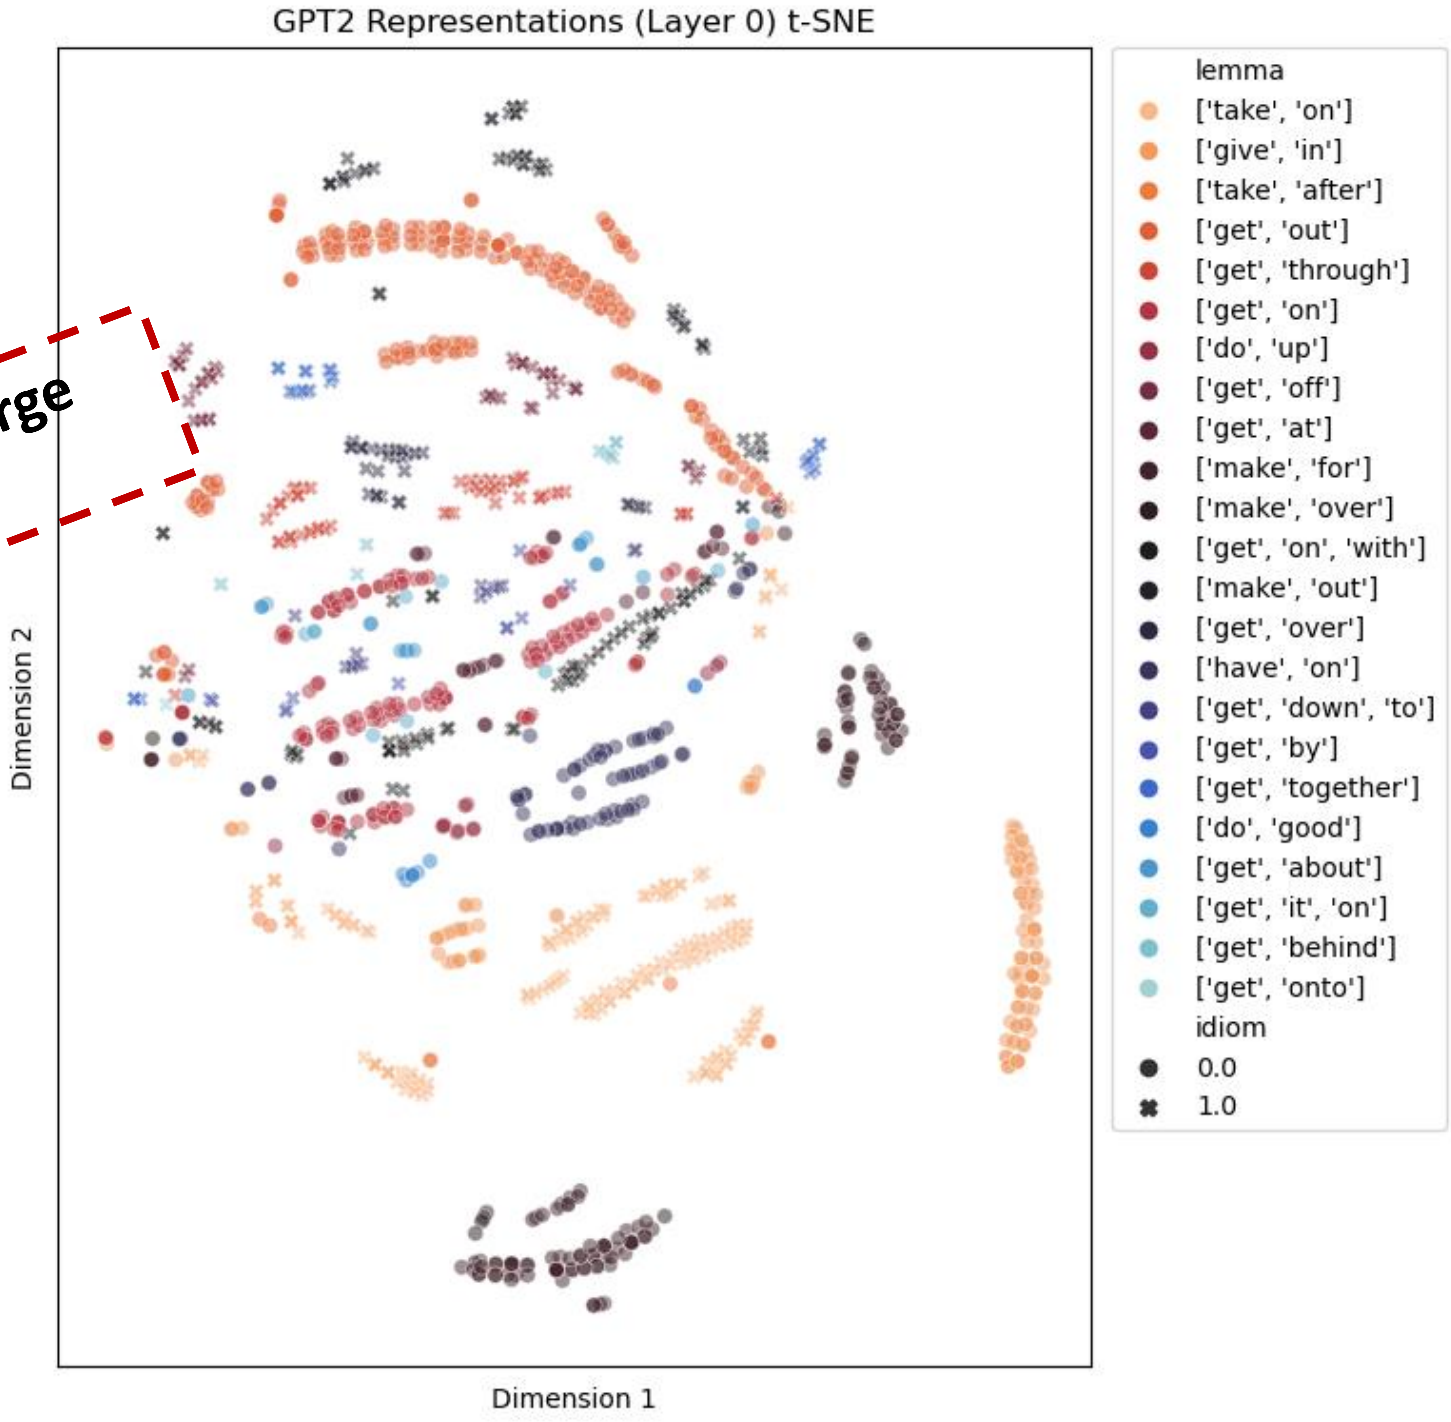

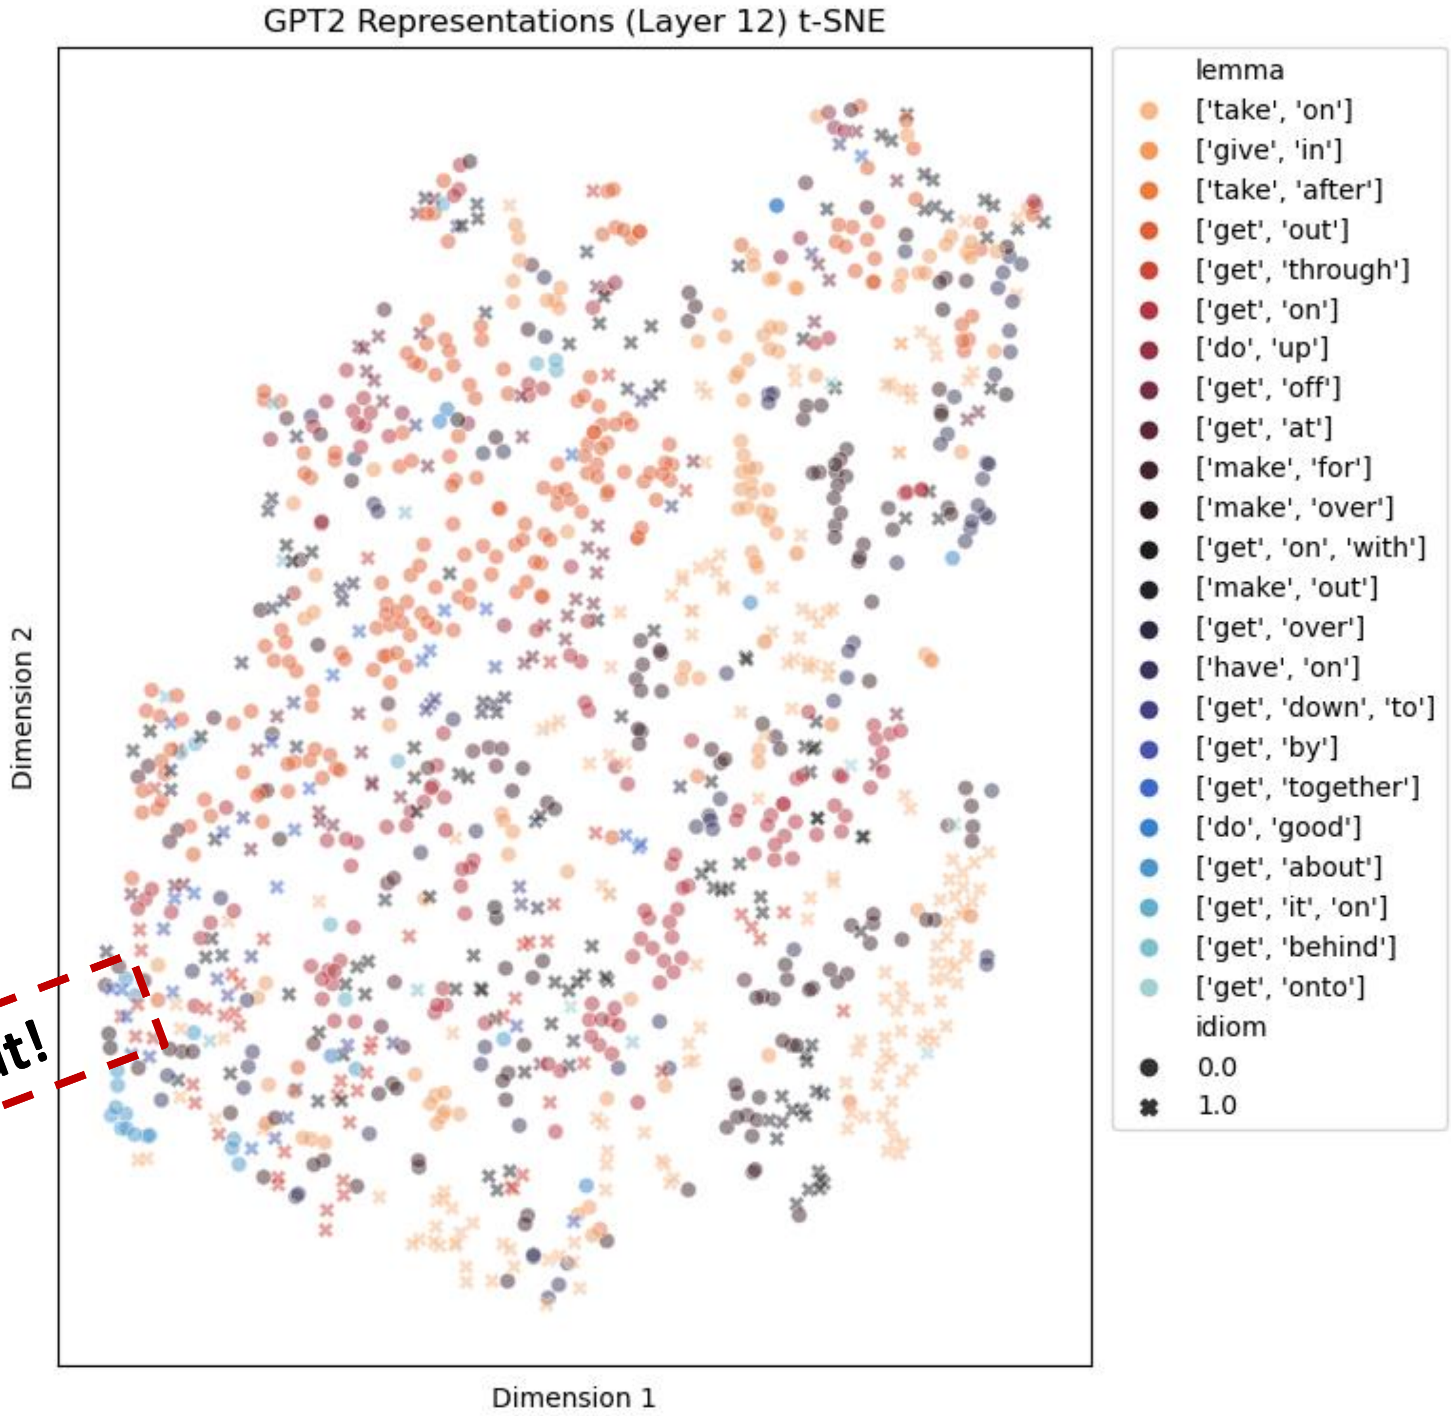

But not at the output!

Starr et al. (submitted)

- Language models do not appear to generate consistent representations of light verb constructions…

- … but they are still able to disambiguate whether a construction is a light verb or not…?

Conclusion

- We still have a long way to go!

- The relationship between natural language processing and human language processing can be a two-way street.

- … get involved!

Acknowledgements

Marten
van Schijndel

Draga
Zec

Helena
Aparicio

Acknowledgements