



SOCIAL NORMS GUIDE REFERENCE RESOLUTION


MITCHELL ABRAMS, MATTHIAS SCHEUTZ





DO SOCIAL NORMS MODULATE REFERENCE RESOLUTION?


DO SOCIAL NORMS MODULATE REFERENCE RESOLUTION?



Context: Friend's Car

“ Take [a seat] ”

DO SOCIAL NORMS MODULATE REFERENCE RESOLUTION?




Context: Taxicab

“ Take [a seat] ”




BACKGROUND: PRAGMATIC CONSTRAINTS

- Psycholinguistics evidence for pragmatic constraints (Hanna and Tanenhaus, 2004)



- Attempts to model social norms (Malle et al., 2020)

- Norms as *prescriptions* and *prohibitions*



- Datasets for testing contextual knowledge and common-sense reasoning

- Winograd Schema Challenge (WSC) (Levesque et al., 2012)
- KnowRef dataset (Emami et al., 2018)




REFERENCE RESOLUTION

What is needed to test whether norms modulate reference resolution?

REFERENCE RESOLUTION

What is needed to test whether norms modulate reference resolution?


Reference task



REFERENCE RESOLUTION

What is needed to test whether norms modulate reference resolution?


Reference task



Modulate contexts and associated norms

Taxicab


Friend's Car



REFERENCE RESOLUTION

What is needed to test whether norms modulate reference resolution?


Reference task



Modulate contexts and associated norms

Taxicab

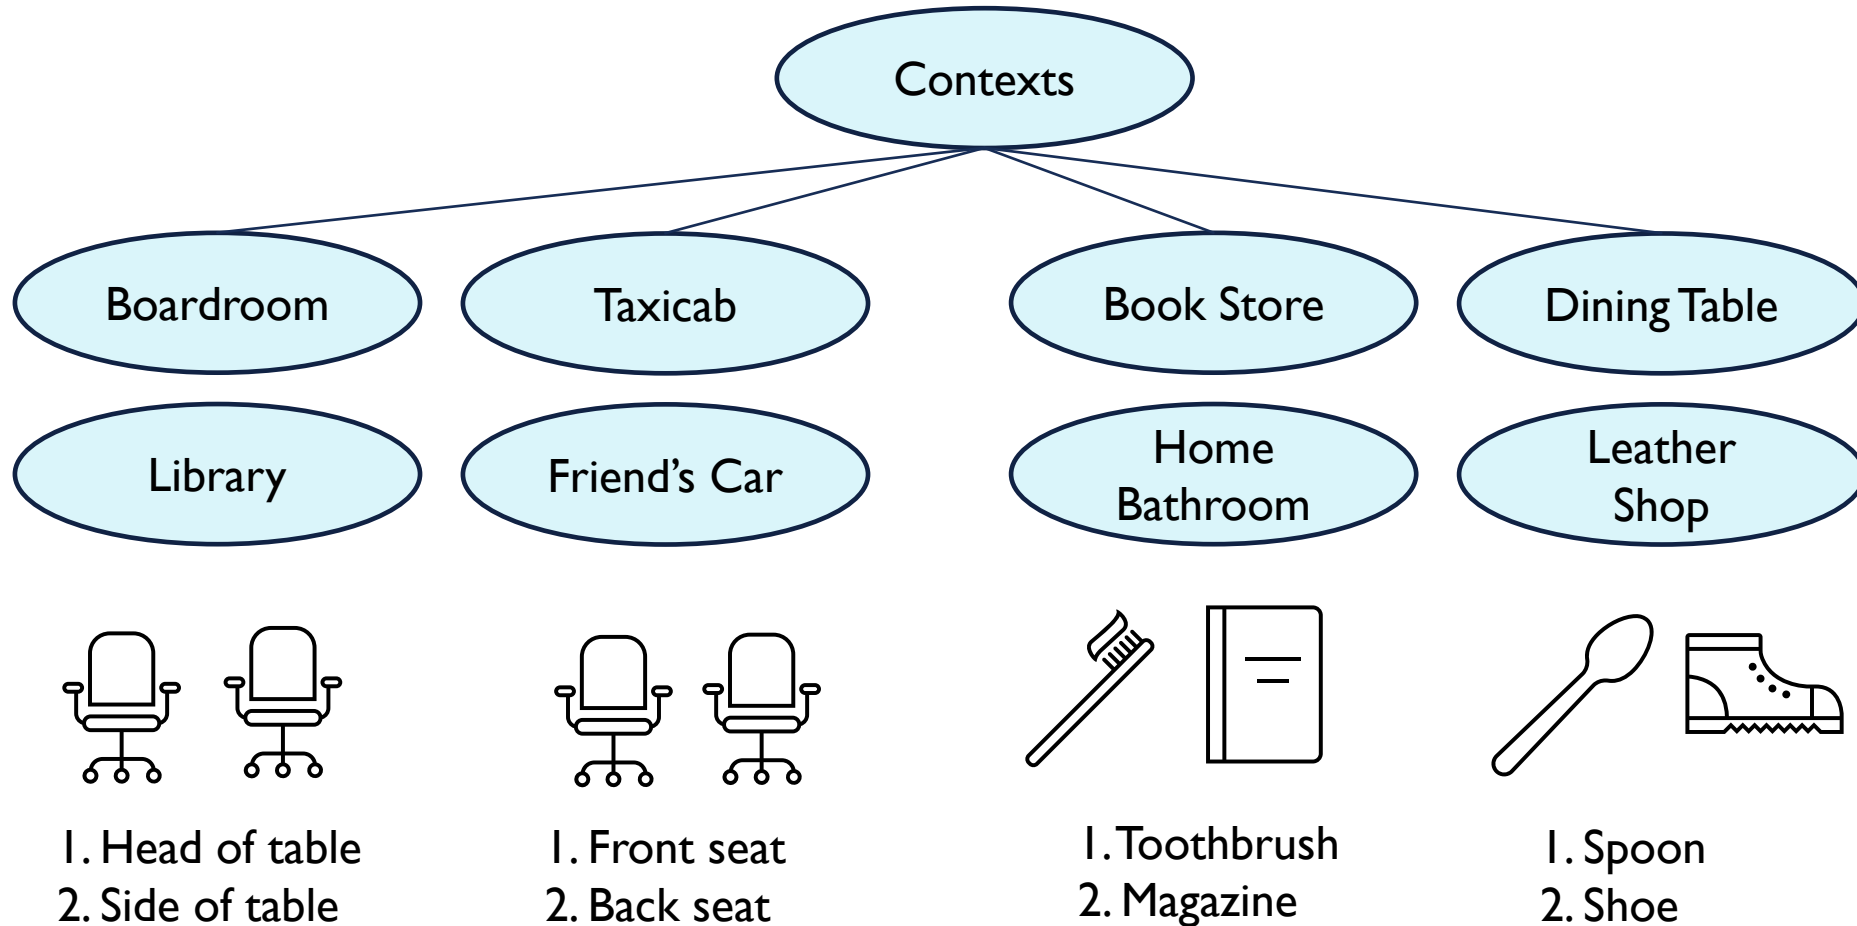
Friend's Car



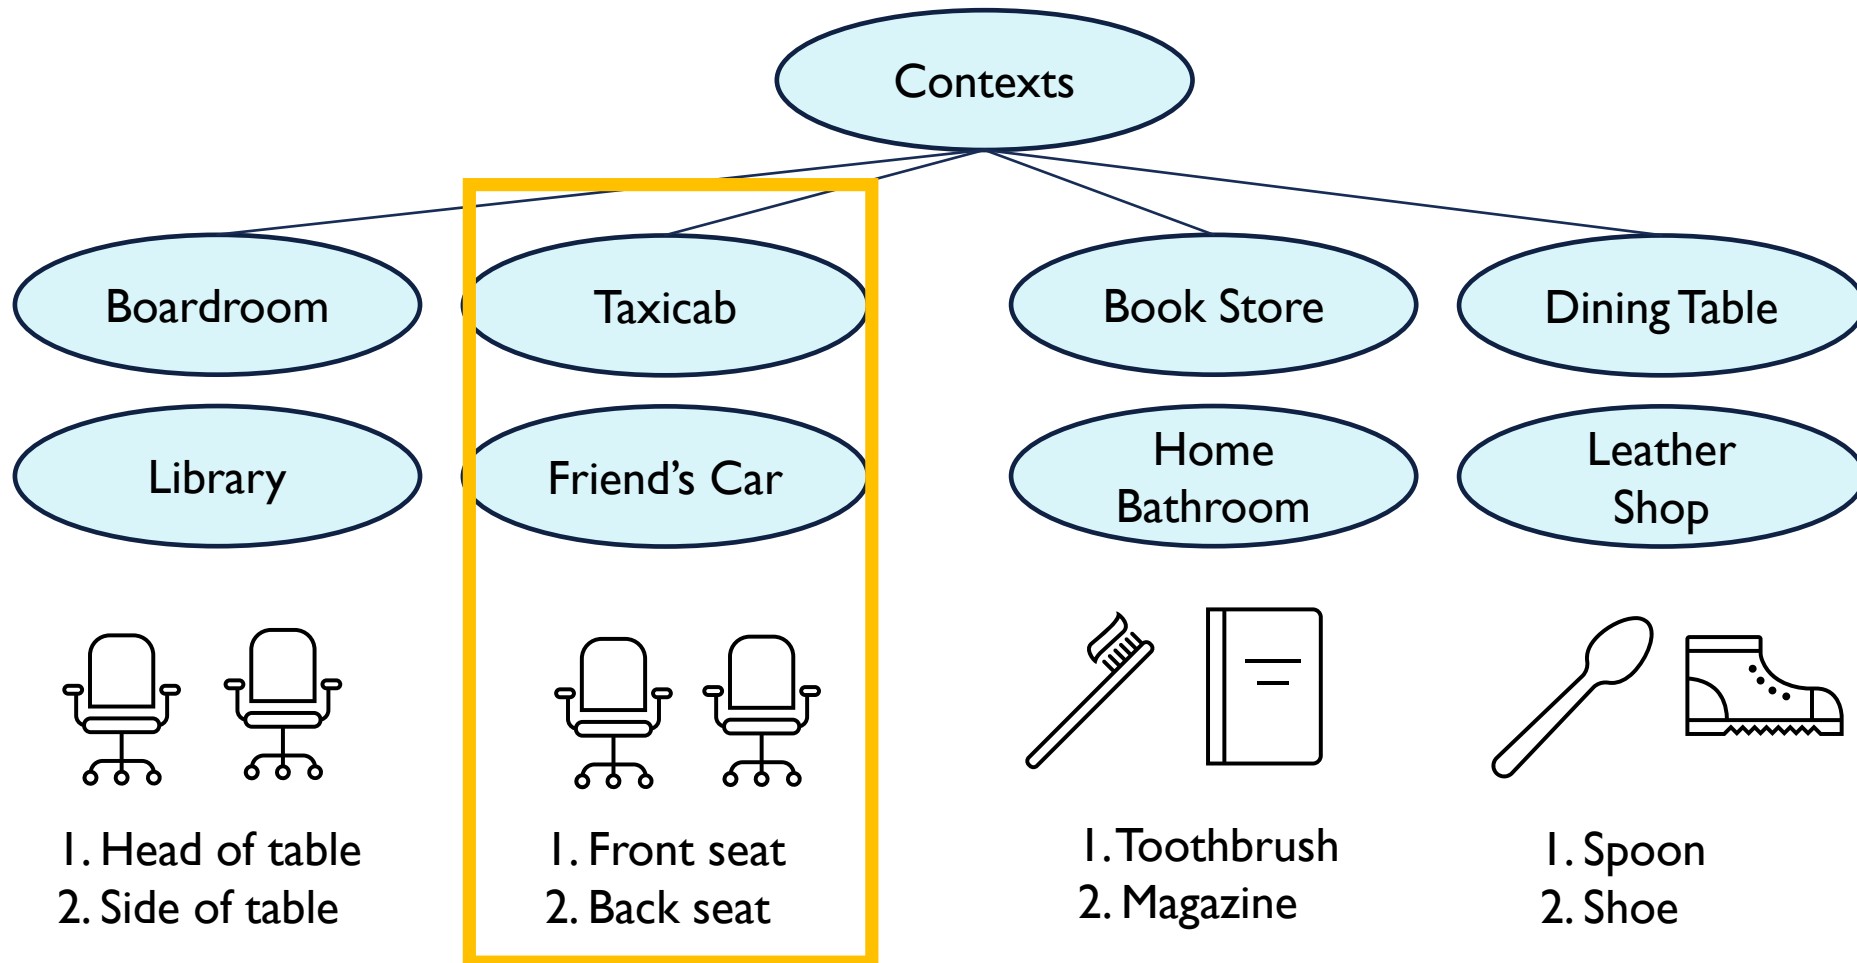
Control Referring Expression

“ Take a seat ”

EXPERIMENTAL DESIGN



EXPERIMENTAL DESIGN



REFERENCE EXPERIMENT

Background:

The scene takes place in a taxicab.

There are two available seats in a taxicab: the front seat and the back seat.
The driver says, "take a seat."

The passenger chooses to sit in the back seat

Was this the correct object?

- yes
- no

Please select the best reason for your response:

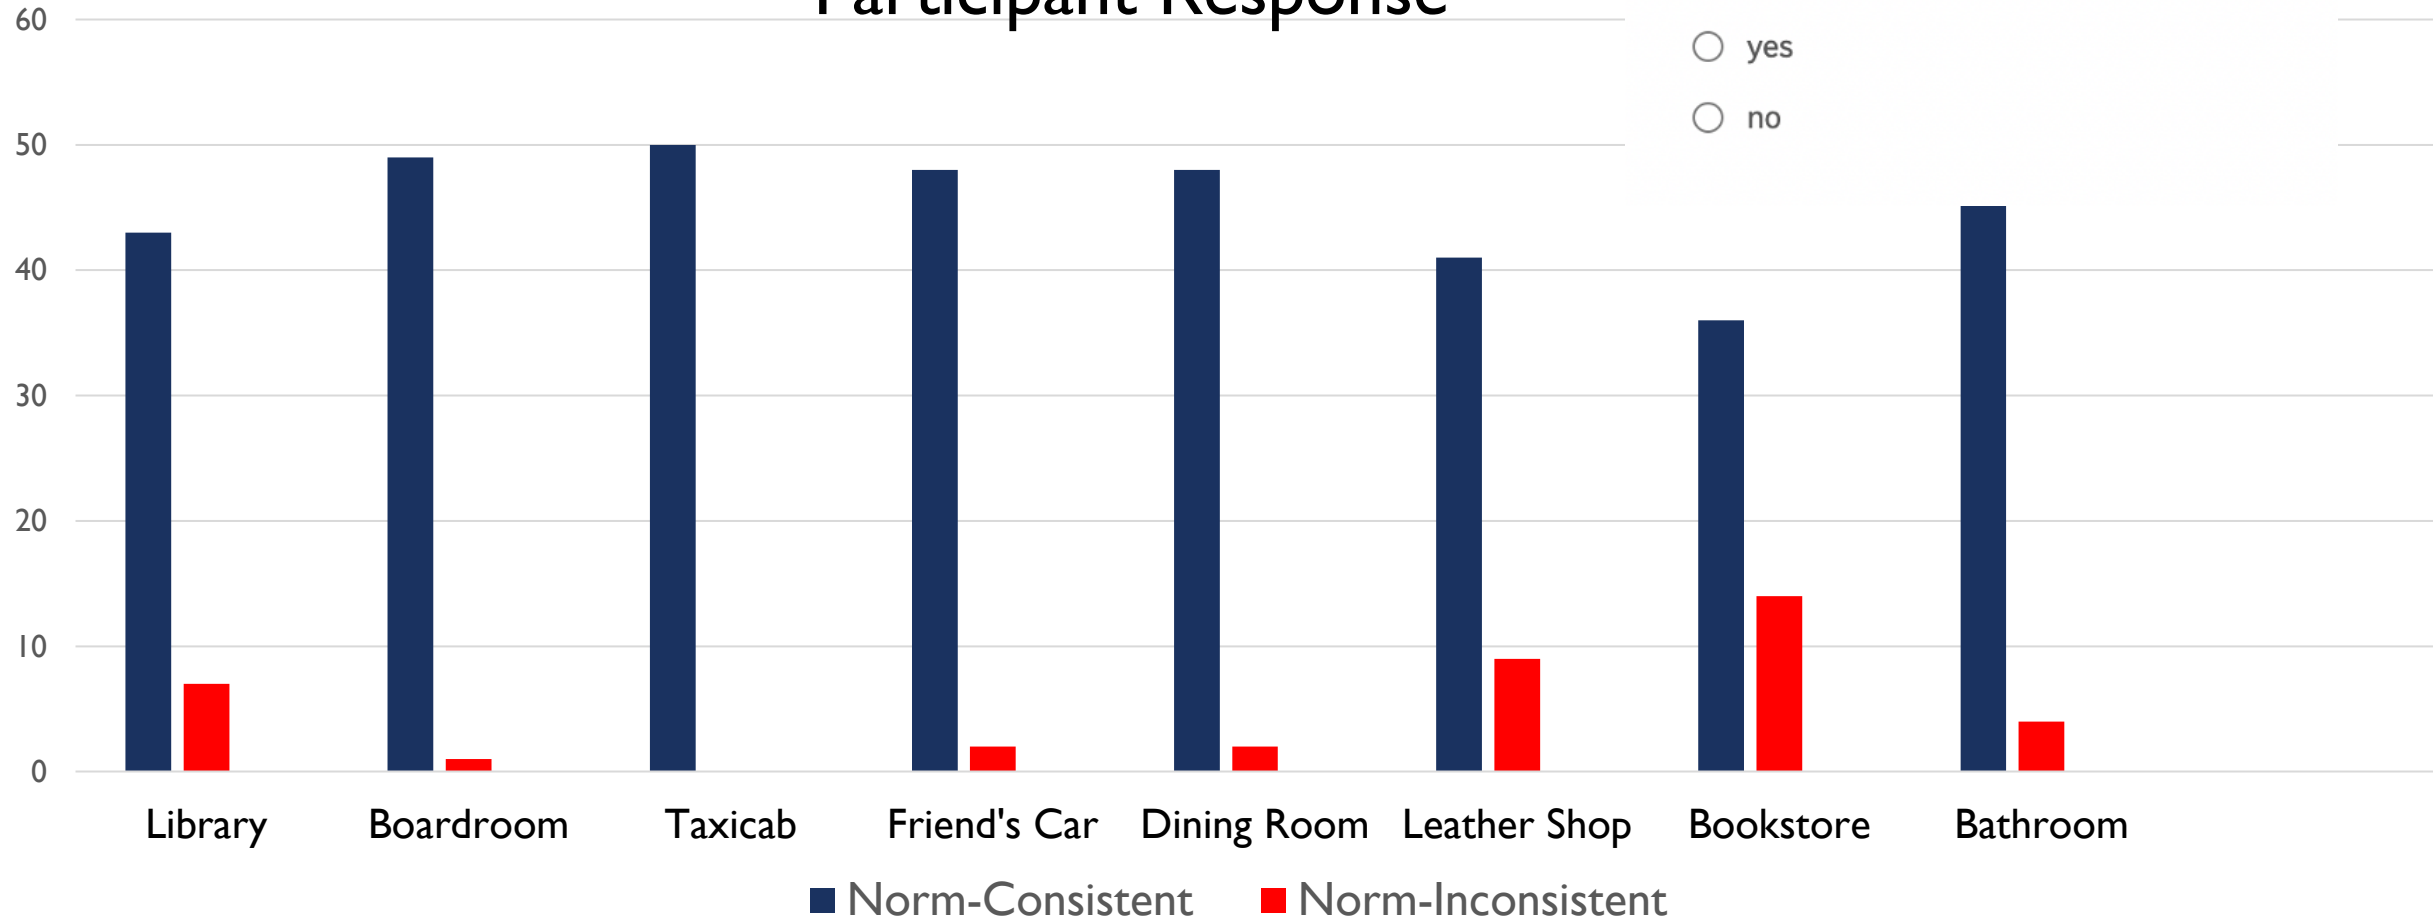
- typical for the setting
- object is mentioned first
- object is mentioned more recently
- time sensitive option
- more convenient option
- other:

Participant Response

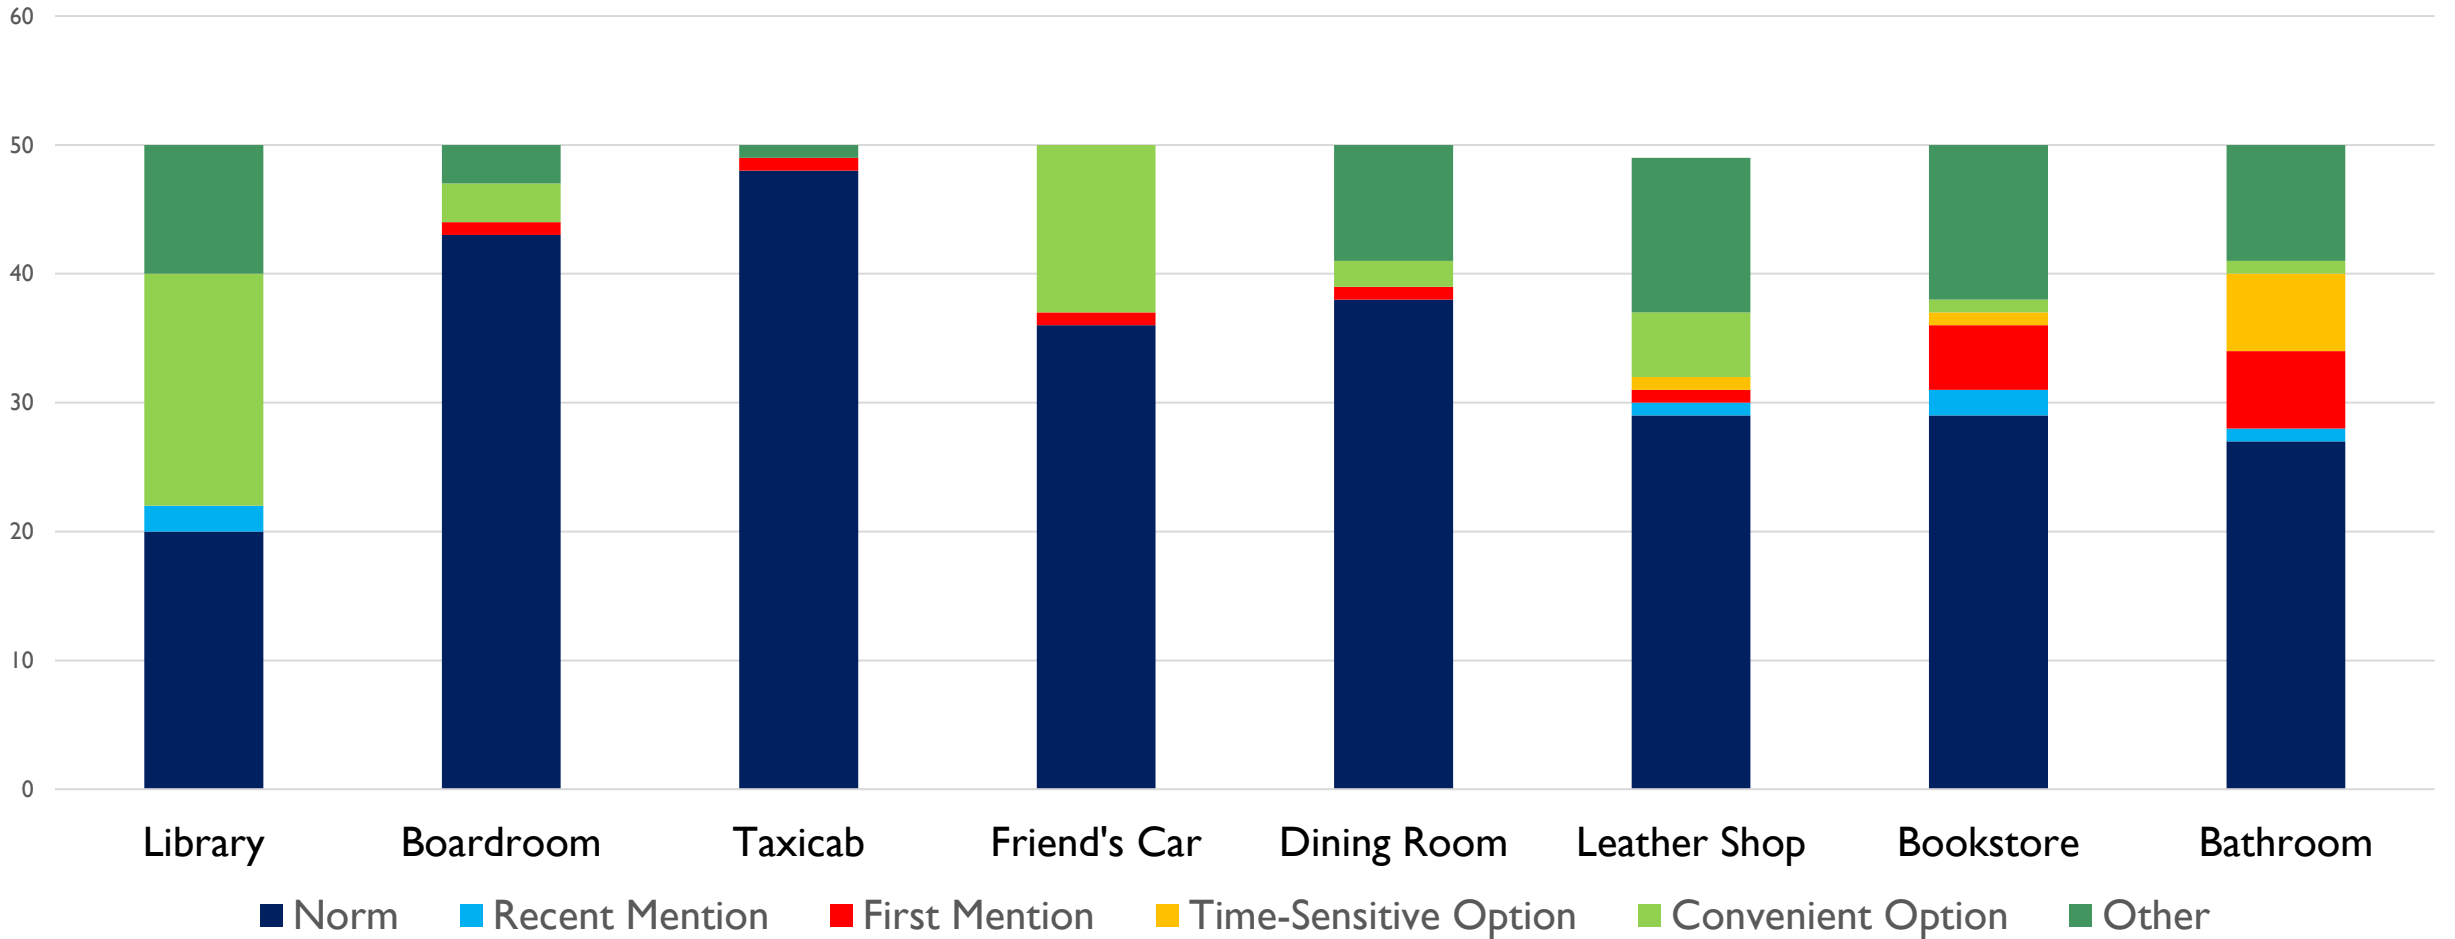
Was this the correct object?

yes


no



Summary of Explanations



How well do NLP tools perform the same reference task?



NeuralCoref (Hannibal and Johnson, 2015)

Stanford CoreNLP (Finkel et al., 2005)


AllenNLP (Lee et al., 2017)

GPT-3 Davinci Model (Brown et al., 2020)

GPT-3 Curie Model (Brown et al., 2020)

AllenNLP

Dining Room Scene




*The scene takes place in a dining room. There is **a shoe** and a spoon sitting on a dining room table. Dinner is about to be served.*

*Someone says, "remove **it.**"*

Someone else removes the shoe from the table.

AllenNLP

Dining Room Scene




The scene takes place in a dining room. There is a spoon and a shoe sitting on a dining room table. Dinner is about to be served.

Someone says, "remove it."

Someone else removes the shoe from the table.


GPT-3: Davinci

Dining Room Scene



The scene takes place in a dining room. There is a spoon and a shoe sitting on a dining room table. Dinner is about to be served.


*Someone says, "remove **it.**"*



Which one?


GPT-3: Davinci

Dining Room Scene



*The scene takes place in a dining room. There is **a shoe** and **a spoon** sitting on a dining room table. Dinner is about to be served.*

*Someone says, "remove **it.**"*



Which one?



GPT-3: Davinci

Home Bathroom Scene

*The scene takes place in a home bathroom. There is a magazine and **a toothbrush** visible on a bathroom floor.*

*Someone says, "pick **it** up."*

Which one?

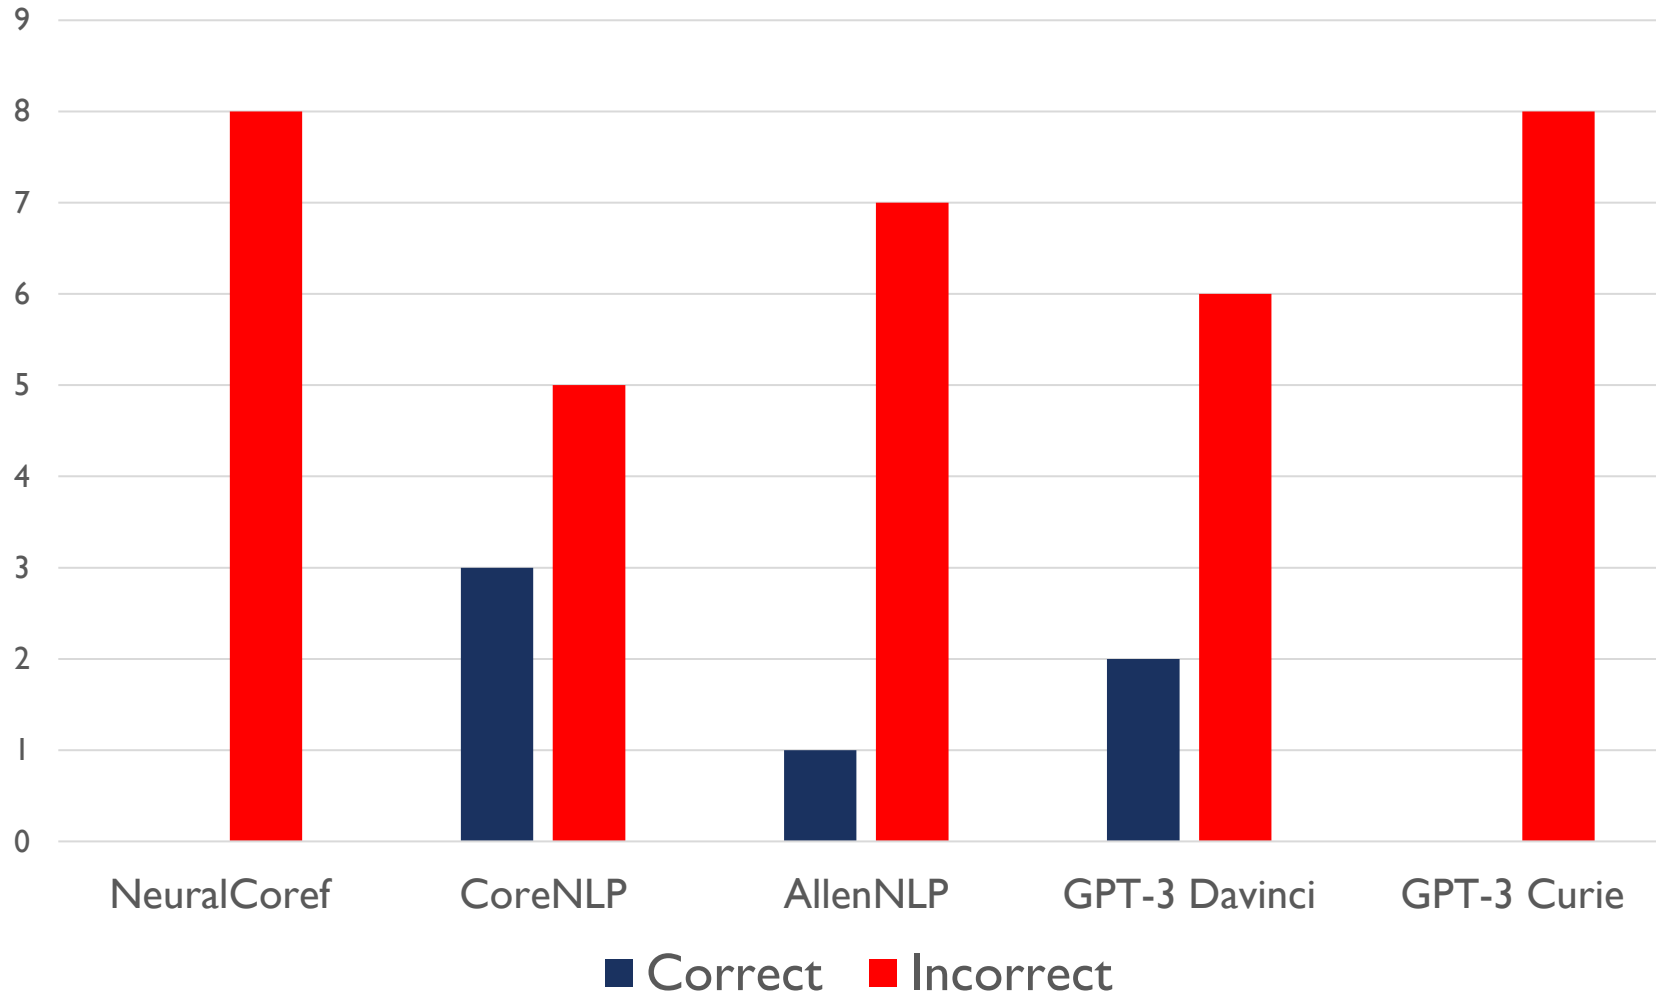


GPT-3:

Justification


“If we use the noun that appears in the context, it is clear that the speaker is referring to the toothbrush. There is no other “it” in the sentence... We would never say, “Pick up the magazine.” This is why it’s important to know whether the noun is the subject or object of a sentence.”

NLP Model Evaluations




CONCLUSION

- Norms in reference task
- Current approaches in NLP
- Embodied Systems (see paper for outlined approach)



THANK YOU

mitchell.abrams@tufts.edu



CITATIONS

Tom B Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. *arXiv preprint arXiv:2005.14165*.

Joyce Y Chai, Pengyu Hong, and Michelle X Zhou. 2004. A probabilistic approach to reference resolution in multimodal user interfaces. In *Proceedings of the 9th international conference on Intelligent user interfaces*, pages 70–77.

Ali Emami, Paul Trichelair, Adam Trischler, Kaheer Suleman, Hannes Schulz, and Jackie Chi Kit Cheung. 2018. The knowref coreference corpus: Removing gender and number cues for difficult pronominal anaphora resolution. *arXiv preprint arXiv:1811.01747*.

Jenny Rose Finkel, Trond Grenager, and Christopher D Manning. 2005. Incorporating non-local information into information extraction systems by gibbs sampling. In *Proceedings of the 43rd Annual Meeting of the Association for Computational Linguistics (ACL'05)*, pages 363–370.

Joy E Hanna and Michael K Tanenhaus. 2004. Pragmatic effects on reference resolution in a collaborative task: Evidence from eye movements. *Cognitive science*, 28(1):105–115

D. Jurafsky and J.H. Martin. 2009. *Speech and Language Processing: An Introduction to Natural Language Processing, Computational Linguistics, and Speech Recognition*. Prentice Hall Series in Artificial Intelligence. Pearson Prentice Hall.

Kenton Lee, Luheng He, Mike Lewis, and Luke Zettlemoyer. 2017. End-to-end neural coreference resolution. *arXiv preprint arXiv:1707.07045*.

Hector Levesque, Ernest Davis, and Leora Morgenstern. 2012. The winograd schema challenge. In *Thirteenth International Conference on the Principles of Knowledge Representation and Reasoning*.

Bertram F Malle, Eric Rosen, Vivienne B Chi, Matthew Berg, and Peter Haas. 2020. A general methodology for teaching norms to social robots. In *2020 29th IEEE International Conference on Robot and Human Interactive Communication (RO-MAN)*, pages 1395–1402. IEEE.

Michael K Tanenhaus, Michael J Spivey-Knowlton, Kathleen M Eberhard, and Julie C Sedivy. 1995. Integration of visual and linguistic information in spoken language comprehension. *Science*, 268(5217):1632–1634.