Rational choices about complex linguistic representations

Christopher Potts

Stanford Linguistics and the Stanford NLP Group

ESSLLI 2021: Approaches to implicature

Fundamental claims

- There is no tension between exhaustification-based theories and rational choice theories.
- Stronger: the two theories can explain more together.
- Perspectives:
  - How do people reason about complex linguistic representations?
  - Rational choice theories facilitate experimental evaluations of representational theories.

Theoretical positioning

- Grice (1975): inspiring, but insufficiently precise.
- Rational choice models are inspired by Grice but unencumbered:
  - Full characterization of ‘cooperativity’
  - No maxims
  - More aligned with Bayesian cognitive science
- No devotion to rational choice theories per se; just want to explain what people actually do with language!

Beyond representations

- Pragmatic free variables – how are they valued?
- Scope-taking ambiguities – what factors shape actual inferences?
- A range of implicatures is attested – but what actually happens between speakers and listeners?

Coordinating on a logical form in context

"Some of the players made some of their shots."

1. *Exh* [*Exh* [some of the players] made *Exh* [some of their shots]]
2. *Exh* [*Exh* [some of the players] made [some of their shots]]
3. *Exh* [[some of the players] made *Exh* [some of their shots]]
4. *Exh* [[some of the players] made [some of their shots]]
5. [*Exh* [some of the players] made *Exh* [some of their shots]]
6. [*Exh* [some of the players] made [some of their shots]]
7. [[some of the players] made *Exh* [some of their shots]]
8. [[some of the players] made [some of their shots]]

Embedded scalars

Quick overview

Central observation

In the scope of a non-monotone quantifier, strengthening *some* to *only some* does not strengthen the semantic content of the sentence:

Exactly one student passed some of their exams.

$\not\Leftarrow$

Exactly one student passed only some of their exams.

Predictions

- Grice (?): Doesn't sound like quantity! (Maybe manner?)
- Standard RSA: *only some* construals should be unavailable.
- Representational theories: there is a representation in which *some* is interpreted as *only some*.

A brief history

Central observation

In the scope of a non-monotone quantifier, strengthening *some* to *only some* does not strengthen the semantic content of the sentence.

1. Geurts and Pouscoulous (2009): made the central observation and found low rates of *only some* construals.
2. Chemla and Spector (2011): modified experimental design and higher rates of *only some* construals.
3. Potts et al. (2016): further simplified experimental design; *only some* construals consistently observed.
4. Franke et al. (2016): prosodic effects and relative preferences of different construals.
5. van Tiel et al. (2018): factors affecting *only some* construals in these and many other cases.

Scalar under universal quantifiers

Every player hit some of his shots.

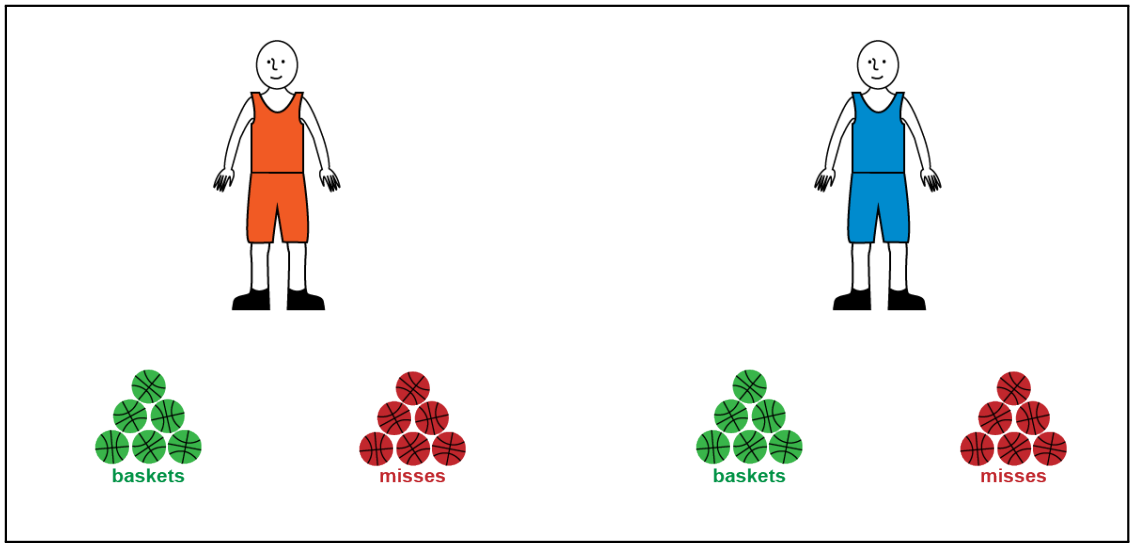

a.	Worlds:	NN NS NA SN SS SA AN AS AA	
b.	Literal:	SS SA AS AA	‘all hit at least some’
c.	Implicature:	NN NS NA SN SS SA AN AS	‘not all hit all’
d.	Result:	SS SA AS	‘all hit some; not all hit all’
e.	Aux. premise:	NN SS AA	‘uniform outcomes’
f.	Communicated:	SS	‘all hit only some’

Scalar under non-monotone quantifier

Exactly one player hit some of his shots.

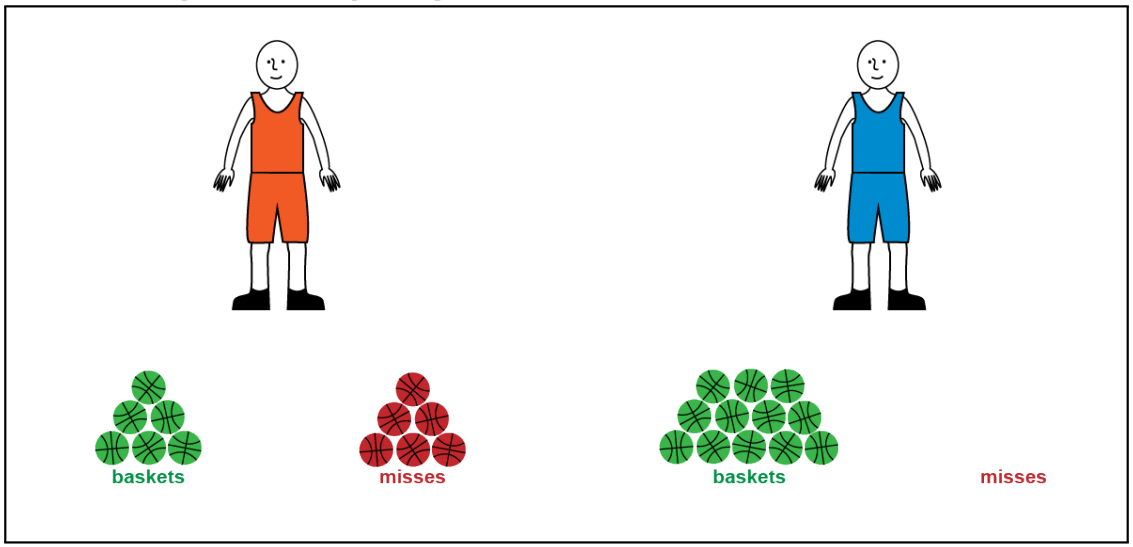

a. Worlds: NN NS NA SN SS SA AN AS AA
b. Literal: NS NA SN AN ‘exactly one hit at least some’
c. Local: NS SN SA AS ‘exactly one hit only some’

Experiment display

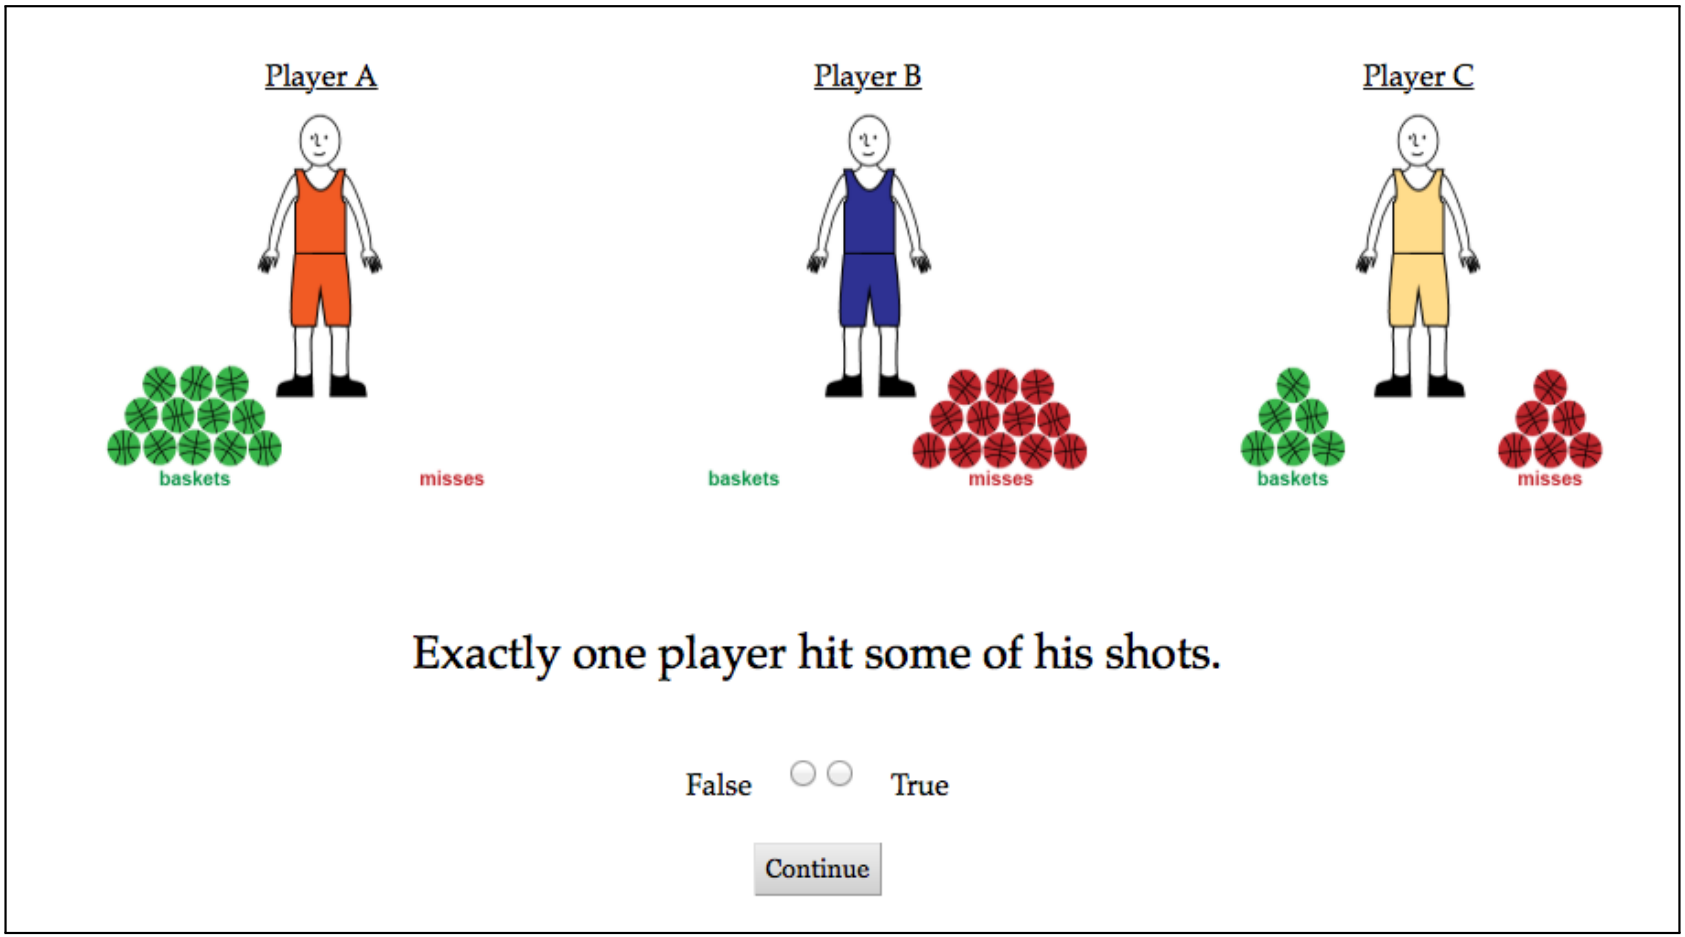

Potts et al. 2016

Experiment display

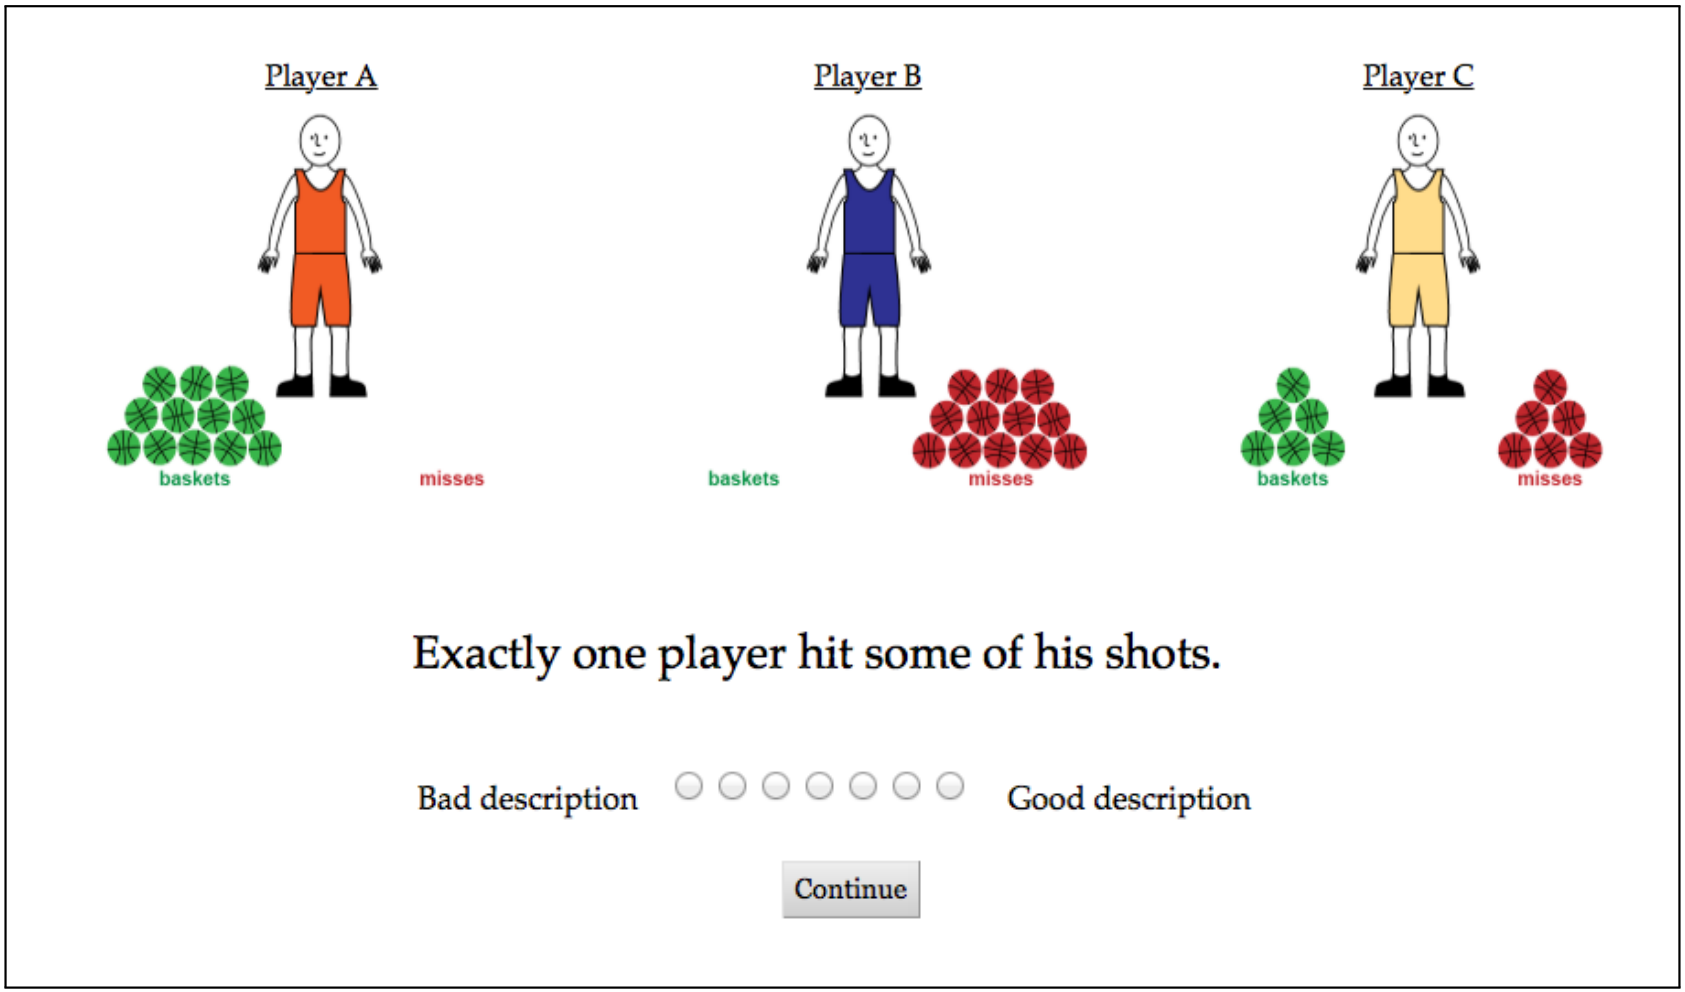

Potts et al. 2016

Other experiment details

- 800 participants recruited via Mechanical Turk (no participants or responses excluded)
- Between-subjects design
- 3 training items; 23 fillers; 9 target sentences:

$$\left\{\begin{array}{l}\text{Every}\\ \text{Exactly one}\\ \text{No}\end{array}\right\} \text{ player hit } \left\{\begin{array}{l}\text{all}\\ \text{none}\\ \text{some}\end{array}\right\} \text{ of his shots.}$$

- Worlds: {NNN, NNS, NNA, NSS, NSA, NAA, SSS, SSA, SAA, AAA}
- Average 80 responses per target–world pair
- Visual display of worlds and jersey colors randomized

Selected results

Exactly one player hit { all / none / some } of his shots.

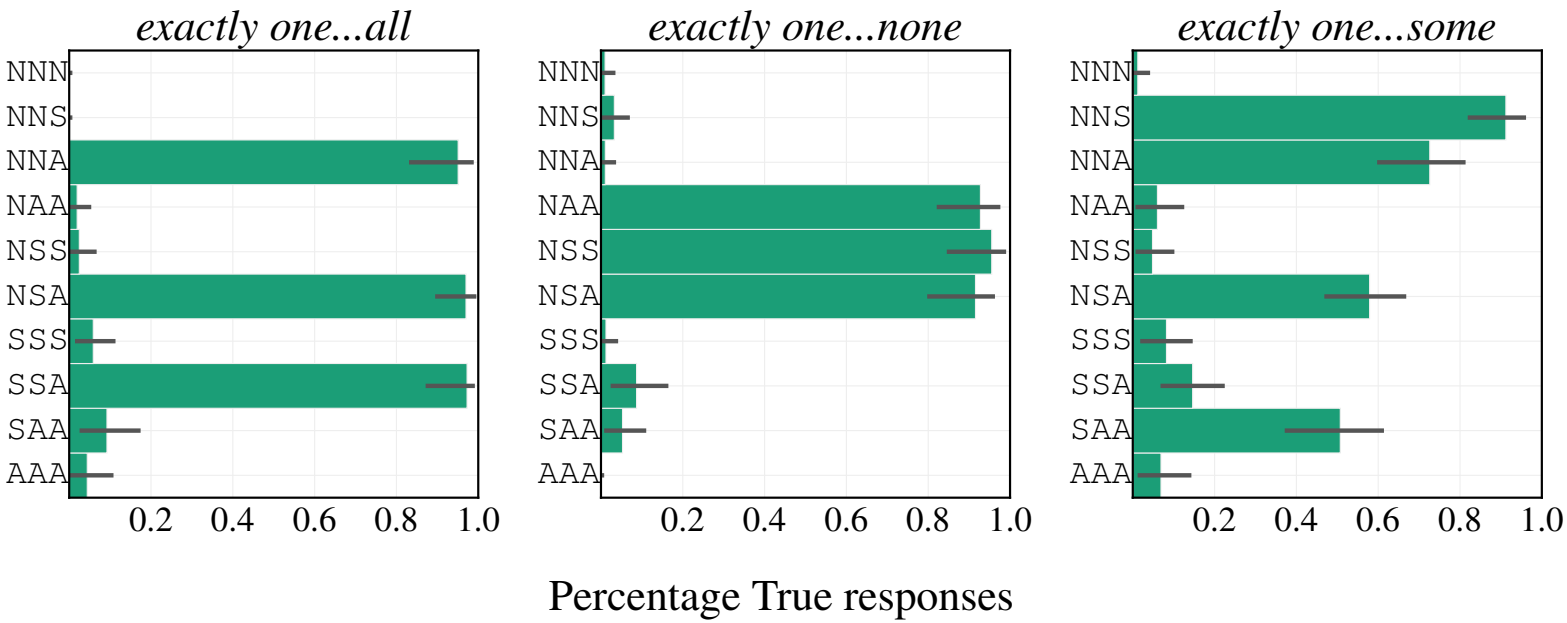

Models evaluated

1. **Literal semantics**

2. **Standard RSA**

3. **Unconstrained refinement**: L_{prag} reasons about a large, unstructured space of possible construals.

4. **Neo-Gricean refinement**: as in 'Unconstrained refinement', but with just neo-Gricean constraints on possible construals.

   Are exhaustification-based theories included here?

Overall results

	Pearson		Spearman		MSE	
Literal semantics	.938	(.926 – .947)	.762	(.754 – .770)	.0065	(.0057 – .0075)
Standard RSA	.924	(.911 – .932)	.757	(.749 – .766)	.0079	(.0072 – .0090)
Unconstrained uncertainty	.945	(.936 – .950)	.794	(.767 – .820)	.0038	(.0035 – .0044)
Neo-Gricean uncertainty	.959	(.950 – .962)	.809	(.808 – .820)	.0034	(.0031 – .0040)

Selected results

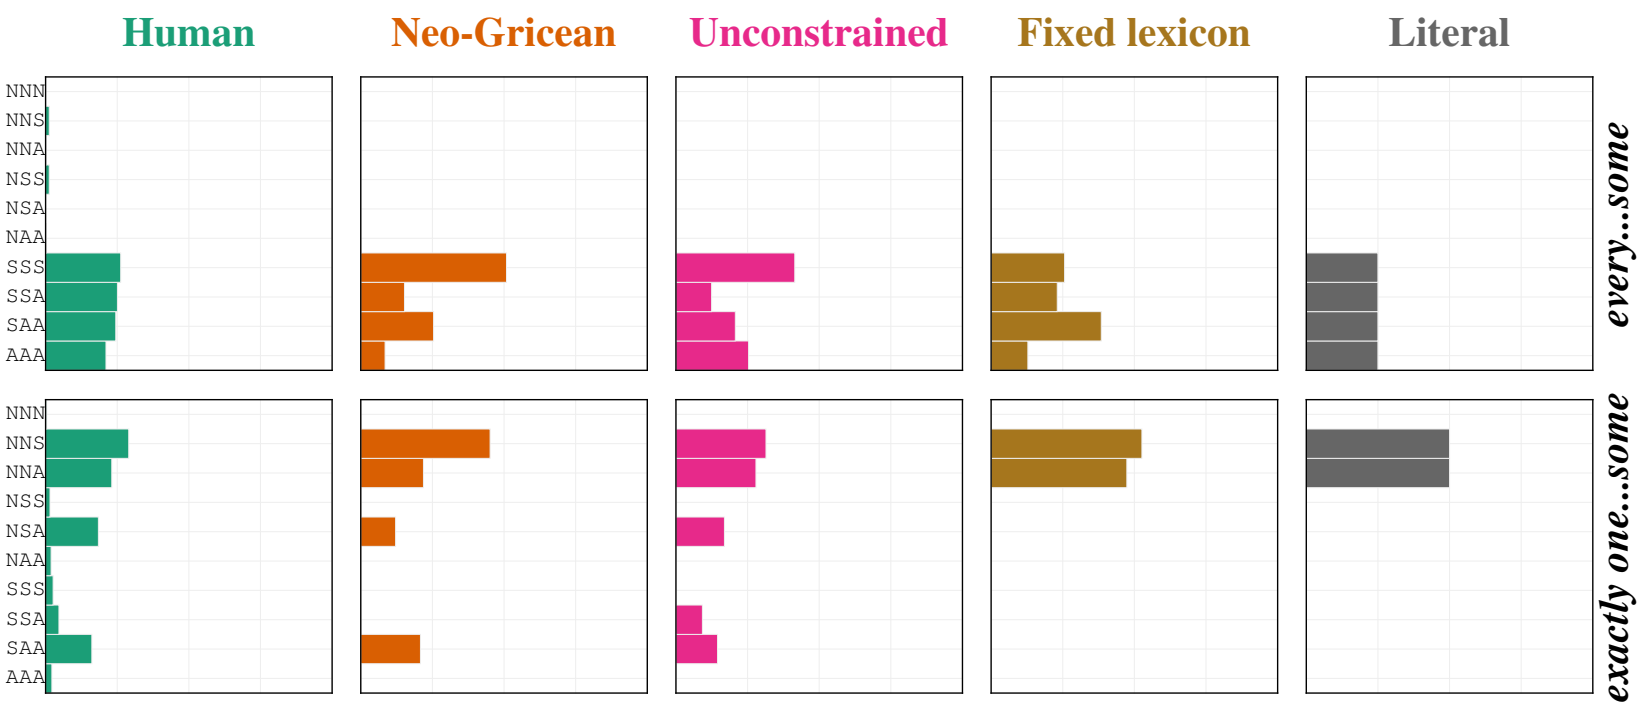

Selected results

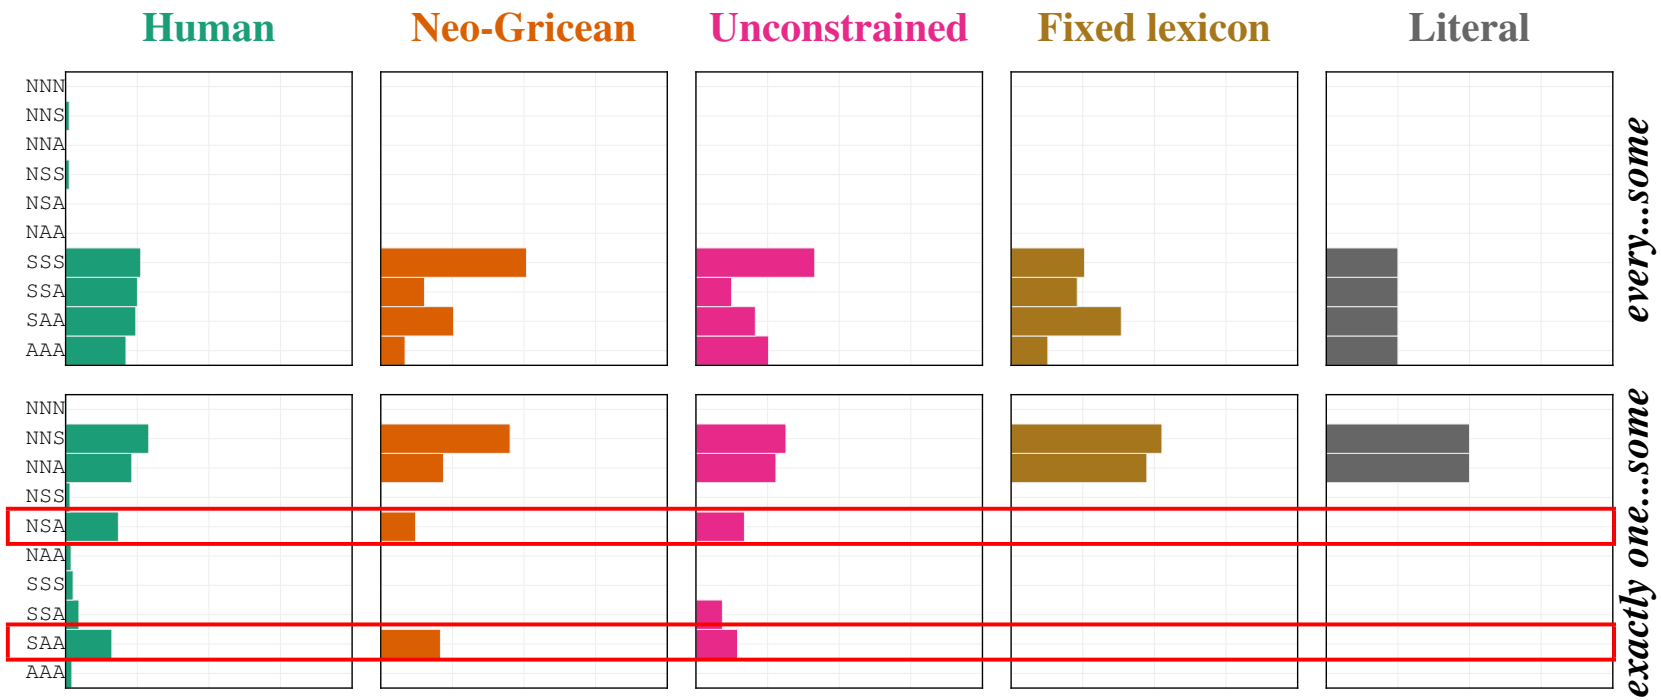

Only enriched

Selected results

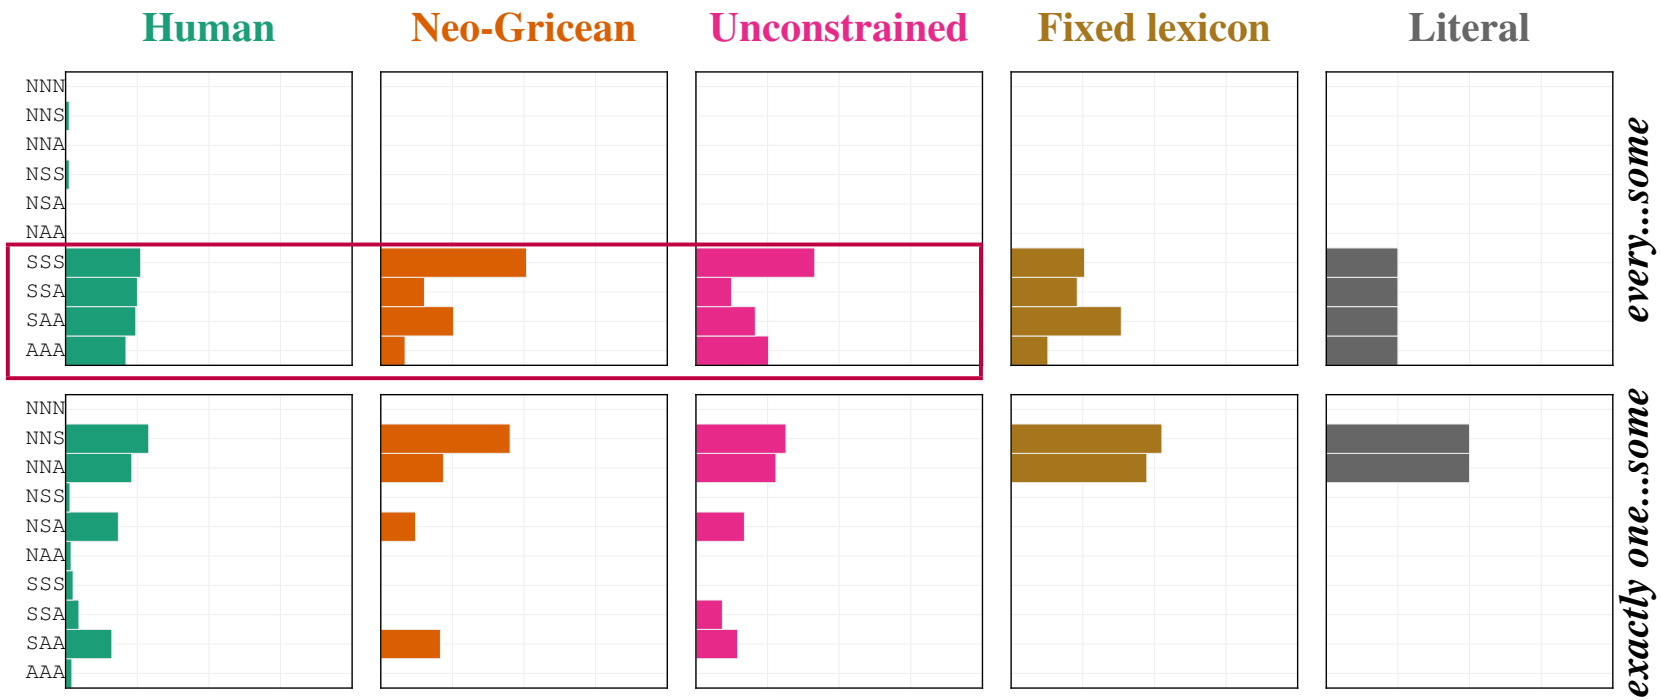

Implicature less likely than the best models predict

Franke and Bergen (2020)

1. $\left\{ \begin{array}{l} \text{None} \\ \text{Some} \\ \text{All} \end{array} \right\}$ of the aliens drank $\left\{ \begin{array}{l} \text{none} \\ \text{some} \\ \text{all} \end{array} \right\}$ of their water.
2. Assume exhaustification theories.
3. Experiments for both production and comprehension.
4. Models:
   a. **Literal semantics**
   b. **Standard RSA**
   c. **Neo-Gricean refinement**
   d. **Lexical intentions**: Speakers communicate both about the world and about their preferred lexicon.
   e. **Global intentions**: Speakers communicate about sentence–LF pairs.
5. Global intentions is the best model of speakers . . .
6. and it is the one that most tightly integrates representational and rational choice theories.

Main takeaways from the literature

- Representational theories cannot be assessed against experimental data without explicit linking hypotheses.
- Rational choice theories provide a rich class of such hypotheses and have independent interest.
- Combined evaluations of these theories reveals sensitivity to alternatives.
- No single rational choice model is clearly superior. Needed:
  - New pragmatic models
  - New representational theories

Scaling up and out

Central question

- Rational choice theories have already achieved a very large impact in cognitive psychology and various parts of artificial intelligence.
- How can we increase this impact by leveraging insights from representational theories?

Scalar implicatures in the wild

Scalar implicatures just aren't that consistent or prominent in real data:

- Degen (2013): in conversational data (Switchboard), the rate of *only some* implicatures is not higher than 45%.
- Li (2021): similar study focused on disjunction, finding high variability in scalar implicature.
- Jasbi (2018): in CHILDES, exclusive readings of *or* are prominent but just one of at least six different prominent construals.
- Hendriks et al. (2009): in a large-scale experiment in Dutch, rates of *only some* implicatures were generally above 80% but sensitive to many surface cues.
- van Tiel et al. (2016): 'scalar diversity' even where alternatives are made salient.

Pragmatics in the wild

- Vagueness “resolution” (Lassiter and Goodman 2013)
- Implicit comparison (Monroe et al. 2017)
- Answerhood conditions (Hawkins and Goodman 2019)
- Metaphor (Kao et al. 2014a)
- Hyperbole (Kao et al. 2014b)
- Politeness (Yoon et al. 2016, 2017)
- Irony (Cohn-Gordon and Bergen 2019)
- Social meaning (Burnett 2019; Qing and Cohn-Gordon 2019)
- Incremental pragmatic inference (Cohn-Gordon et al. 2019)
- Adaptation (Schuster and Degen 2019; Hawkins et al. 2019)

Pragmatics in Artificial Intelligence

- Referential descriptions (Andreas and Klein 2016; Monroe et al. 2017, 2018)
- Navigational instructions (Fried et al. 2018a)
- Question answering (Hawkins and Goodman 2020)
- Machine translation (Cohn-Gordon and Goodman 2019)
- Image description (Mao et al. 2016; Vedantam et al. 2017; Cohn-Gordon et al. 2018)
- QUDs for controlled text generation (Nie et al. 2020)
- Collaborative problem solving (Tellex et al. 2014)
- Summarization (Fried et al. 2018b)
- Adaptive agents (Hawkins et al. 2019)

Pragmatics in Artificial Intelligence

- Informativity in context

Pragmatics in Artificial Intelligence

- Informativity in context
- Scalability
- Sensitivity to variation
- General mechanisms for pragmatic inference

Conclusions

- An exciting moment for pragmatics
- Rational choices about complex linguistic representations
- What's behind all these results?
  - Interdisciplinarity
  - Insights about representation and reasoning
  - A lot of hard-won insights about language data
- Representations in NLP, and a call for help

Thanks!

References I

Jacob Andreas and Dan Klein. 2016. Reasoning about pragmatics with neural listeners and speakers. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 1173–1182. Association for Computational Linguistics.

Heather Burnett. 2019. Signalling games, sociolinguistic variation and the construction of style. *Linguistics and Philosophy*, 42(5):419–450.

Emmanuel Chemla and Benjamin Spector. 2011. Experimental evidence for embedded scalar implicatures. *Journal of Semantics*, 28(3):359–400.

Reuben Cohn-Gordon and Leon Bergen. 2019. Verbal irony, pretense, and the common ground. Ms. Stanford University and UC San Diego.

Reuben Cohn-Gordon and Noah Goodman. 2019. Lost in machine translation: A method to reduce meaning loss. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 437–441, Minneapolis, Minnesota. Association for Computational Linguistics.

Reuben Cohn-Gordon, Noah D. Goodman, and Christopher Potts. 2018. Pragmatically informative image captioning with character-level inference. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 439–443, Stroudsburg, PA. Association for Computational Linguistics.

Reuben Cohn-Gordon, Noah D. Goodman, and Christopher Potts. 2019. An incremental iterated response model of pragmatics. In *Proceedings of the Society for Computation in Linguistics*, pages 81–90, Washington, D.C. Linguistic Society of America.

Judith Degen. 2013. A corpus-based study of *Some* (but not *All*) implicatures. Ms., University of Rochester.

Michael Franke and Leon Bergen. 2020. Theory-driven statistical modeling for semantics and pragmatics: A case study on grammatically generated implicature readings. *Language*, 96(2):e77–e96.

Michael Franke, Fabian Schlotterbeck, and Petra Augurzky. 2016. Embedded scalars, preferred readings and prosody: An experimental revisit. *Journal of Semantics*, 34(1):153–199.

Daniel Fried, Jacob Andreas, and Dan Klein. 2018a. Unified pragmatic models for generating and following instructions. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 1951–1963, New Orleans, Louisiana. Association for Computational Linguistics.

Daniel Fried, Ronghang Hu, Volkan Cirik, Anna Rohrbach, Jacob Andreas, Louis-Philippe Morency, Taylor Berg-Kirkpatrick, Kate Saenko, Dan Klein, and Trevor Darrell. 2018b. Speaker-follower models for vision-and-language navigation. *arXiv preprint arXiv:1806.02724*.

Bart Geurts and Nausicaa Pouscoulous. 2009. Embedded implicatures?!? *Semantics and Pragmatics*, 2(4):1–34.

H. Paul Grice. 1975. Logic and conversation. In Peter Cole and Jerry Morgan, editors, *Syntax and Semantics*, volume 3: Speech Acts, pages 43–58. Academic Press, New York.

References II

Robert X. D. Hawkins and Noah D. Goodman. 2019. Why do you ask? The informational dynamics of questions and answers. PsyArXiv, 6.

Robert X. D. Hawkins and Noah D. Goodman. 2020. Why do you ask? the informational dynamics of questions and answers. Ms., Princeton University and Stanford University.

Robert X. D. Hawkins, Minae Kwon, Dorsa Sadigh, and Noah D. Goodman. 2019. Continual adaptation for efficient machine communication. In *ICML Workshop on Adaptive & Multitask Learning: Algorithms & Systems*.

Petra Hendriks, John Hoeks, Helen de Hoop, Irene Krammer, Erik-Jan Smits, Jennifer Spenader, and Henriette de Swart. 2009. A large-scale investigation of scalar implicature. In *Semantics and Pragmatics: From Experiment to Theory*, Palgrave Studies in Pragmatics, Language and Cognition, pages 30–50, Houndmills, Basingstoke, Hampshire. Palgrave Macmillan.

Masoud Jasbi. 2018. *Learning Disjunction*. Ph.D. thesis, Stanford University, Stanford, CA.

Justine T. Kao, Leon Bergen, and Noah D. Goodman. 2014a. Formalizing the pragmatics of metaphor understanding. In *Proceedings of the 36th Annual Meeting of the Cognitive Science Society*, pages 719–724, Wheat Ridge, CO. Cognitive Science Society.

Justine T. Kao, Jean Y. Wu, Leon Bergen, and Noah D. Goodman. 2014b. Nonliteral understanding of number words. *Proceedings of the National Academy of Sciences*, 111(33):12002–12007.

Daniel Lassiter and Noah D. Goodman. 2013. Context, scale structure, and statistics in the interpretation of positive-form adjectives. In *Proceedings of Semantics and Linguistic Theory 23*, pages 587–610, Ithaca, NY. CLC Publications.

Elissa Li. 2021. Predicting scalar inference and alternativehood with neural sentence encoders. Master's thesis, Stanford University.

Junhua Mao, Jonathan Huang, Alexander Toshev, Oana Camburu, Alan L. Yuille, and Kevin Murphy. 2016. Generation and comprehension of unambiguous object descriptions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 11–20. IEEE.

Will Monroe, Robert X. D. Hawkins, Noah D. Goodman, and Christopher Potts. 2017. Colors in context: A pragmatic neural model for grounded language understanding. *Transactions of the Association for Computational Linguistics*, 5:325–338.

Will Monroe, Jennifer Hu, Andrew Jong, and Christopher Potts. 2018. Generating bilingual pragmatic color references. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 2155–2165, Stroudsburg, PA. Association for Computational Linguistics.

Allen Nie, Reuben Cohn-Gordon, and Christopher Potts. 2020. Pragmatic issue-sensitive image captioning. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 1924–1938, Online. Association for Computational Linguistics.

Christopher Potts, Daniel Lassiter, Roger Levy, and Michael C. Frank. 2016. Embedded implicatures as pragmatic inferences under compositional lexical uncertainty. *Journal of Semantics*, 33(4):755–802.

References III

Ciyang Qing and Reuben Cohn-Gordon. 2019. Use-conditional meaning in Rational Speech-Act models. In *Proceedings of Sinn und Bedeutung 23*, volume 2, pages 253–265.

Sebastian Schuster and Judith Degen. 2019. Speaker-specific adaptation to variable use of uncertainty expressions. In *Proceedings of the 41st Annual Meeting of the Cognitive Science Society*, pages 2769–2775, Montreal. Cognitive Science Society.

Stefanie Tellex, Ross A. Knepper, Adrian Li, Thomas M. Howard, Daniela Rus, and Nicholas Roy. 2014. Asking for help using inverse semantics. In *Proceedings of Robotics: Science and Systems*.

Bob van Tiel, Emiel van Miltenburg, Natalia Zevakhina, and Bart Geurts. 2016. Scalar diversity. *Journal of Semantics*, 33(1):137–175.

Bob van Tiel, Ira Noveck, and Mikhail Kissine. 2018. Reasoning with 'some'. *Journal of Semantics*, page ffy012.

Ramakrishna Vedantam, Samy Bengio, Kevin Murphy, Devi Parikh, and Gal Chechik. 2017. Context-aware captions from context-agnostic supervision. *arXiv:1701.02870*.

Erica J Yoon, Michael Henry Tessler, Noah D Goodman, and Michael C Frank. 2016. Talking with tact: Polite language as a balance between kindness and informativity. In *Proceedings of the 38th Annual Conference of the Cognitive Science Society*.

Erica J Yoon, Michael Henry Tessler, Noah D Goodman, and Michael C Frank. 2017. "i won't lie, it wasn't amazing": Modeling polite indirect speech. In *Proceedings of the 39th Annual Conference of the Cognitive Science Society*.