

# Grounded language understanding: Neural RSA

Christopher Potts

Stanford Linguistics

CS224u: Natural language understanding



# Papers employing these techniques

- Andreas and Klein 2016
- Fried et al. 2018
- Monroe et al. 2017
- Monroe et al. 2018

# Motivation

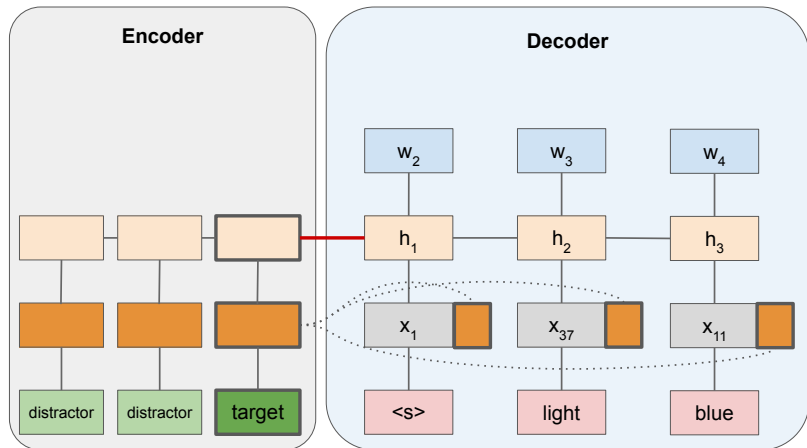
- Discriminative image labeling
- Image captioning
- Machine translation
- Collaborative problem solving
- Interpreting complex descriptions
- Optical Character Recognition
- Scalability
- Sensitivity to variation
- Bounded rationality
- New kinds of model assessment
- Impact

# Colors in context

	Context			Utterance
			blue	
			The darker blue one	
			teal not the two that are more green	
			dull pink not the super bright one	
			not any of the regular greens	
			Purple	
			blue	

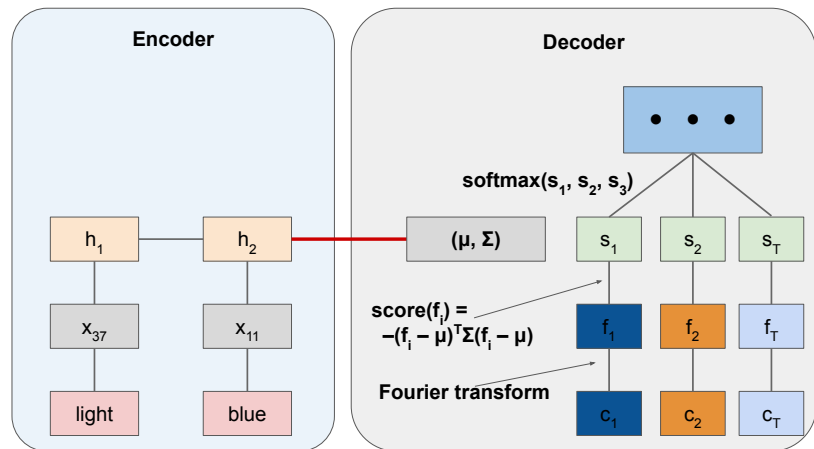
Stanford Colors in Context corpus  
(Monroe et al. 2017)

# Literal neural speaker $S_{lit}^{\theta}$



Monroe et al. 2017

# Neural literal listener



Monroe et al. 2017

# Neural pragmatic agents

Neural pragmatic speaker (Andreas and Klein 2016)

$$\mathbf{s}_{\text{prag}}^{\theta}(msg | state) = \frac{\mathbf{L}_0^{\theta}(state | msg)}{\sum_{msg' \in X} \mathbf{L}_0^{\theta}(state | msg')}$$

with  $X$  a sample from  $\mathbf{s}_{\text{lit}}^{\theta}(msg | state)$  such that  $msg \in X$ .

Neural pragmatic listener

$$\mathbf{L}_1^{\theta}(state | msg) \propto \mathbf{s}_{\text{prag}}^{\theta}(msg | state)$$

Blended neural pragmatic listener

Weighted combination of  $\mathbf{L}_0^{\theta}$  and  $\mathbf{L}_1^{\theta}$ .

## Other related work

- Golland et al. (2010): Recursive speaker/listener reasoning as part of interpreting complex utterances compositionally, with grounding in a simple visual world.
- Wang et al. (2016): Pragmatic reasoning helps in online learning of semantic parsers.
- Tellex et al.'s (2014) Inverse Semantics: Robot utterances are scored by models similar to RSA's pragmatic speakers.
- Khani et al. (2018): Collaborative games with pragmatic reasoning.
- Cohn-Gordon and Goodman (2019): RSA for translation
- Cohn-Gordon et al. (2018, 2019): Word- and character-level RSA
- Monroe and Potts (2015): "RSA as a hidden activation function"
- Mao et al. 2016: pragmatic learning objectives



# 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.
- 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.
- Daniel Fried, Jacob Andreas, and Dan Klein. 2018. [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.
- Dave Golland, Percy Liang, and Dan Klein. 2010. [A game-theoretic approach to generating spatial descriptions](#). In *Proceedings of the 2010 Conference on Empirical Methods in Natural Language Processing*, pages 410–419, Stroudsburg, PA. ACL.
- Fereshte Khani, Noah D. Goodman, and Percy Liang. 2018. [Planning, inference and pragmatics in sequential language games](#). *Transactions of the Association for Computational Linguistics*, 6:543–555.
- 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.
- Will Monroe and Christopher Potts. 2015. Learning in the Rational Speech Acts model. In *Proceedings of 20th Amsterdam Colloquium*, Amsterdam. ILLC.

# References II

- 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*.
- Sida I. Wang, Percy Liang, and Christopher D. Manning. 2016. [Learning language games through interaction](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2368–2378. Association for Computational Linguistics.