Emergent Compositionality in Signaling Games

Nicholas Tomlin [1] Ellie Pavlick [2]

[1]University of California, Berkeley [2]Brown University

Abstract

Understanding the origins of linguistic compositionality is a fundamental challenge in evolutionary linguistics. Prior work has explored this topic through dynamical computational modeling and experiments in iterated learning. We explore these questions using RL agents tasked with developing cooperative communication strategies in a signaling game. We analyze how various mechanisms (such as Bayesian pragmatic reasoning) and constraints (such as limited memory) may affect compositionality and generalizability in the invented communication protocols. In particular, our preliminary results suggest that incremental pragmatic reasoning induces a bias towards lexical compositionality. To evaluate the extensibility of our model, we compare the behavior of the RL agents to the behavior of humans on the same task. That is, we ask humans to coordinate in a reference game task by repeatedly composing non-linguistic symbols. We discuss ways in which the resulting protocol mirrors and differs from that produced by the RL agents.

Compositional Referring Behavior

We seek to measure the compositionality of simple referring expressions. In this domain, *compositionality* and *groundedness* are interchangeable.

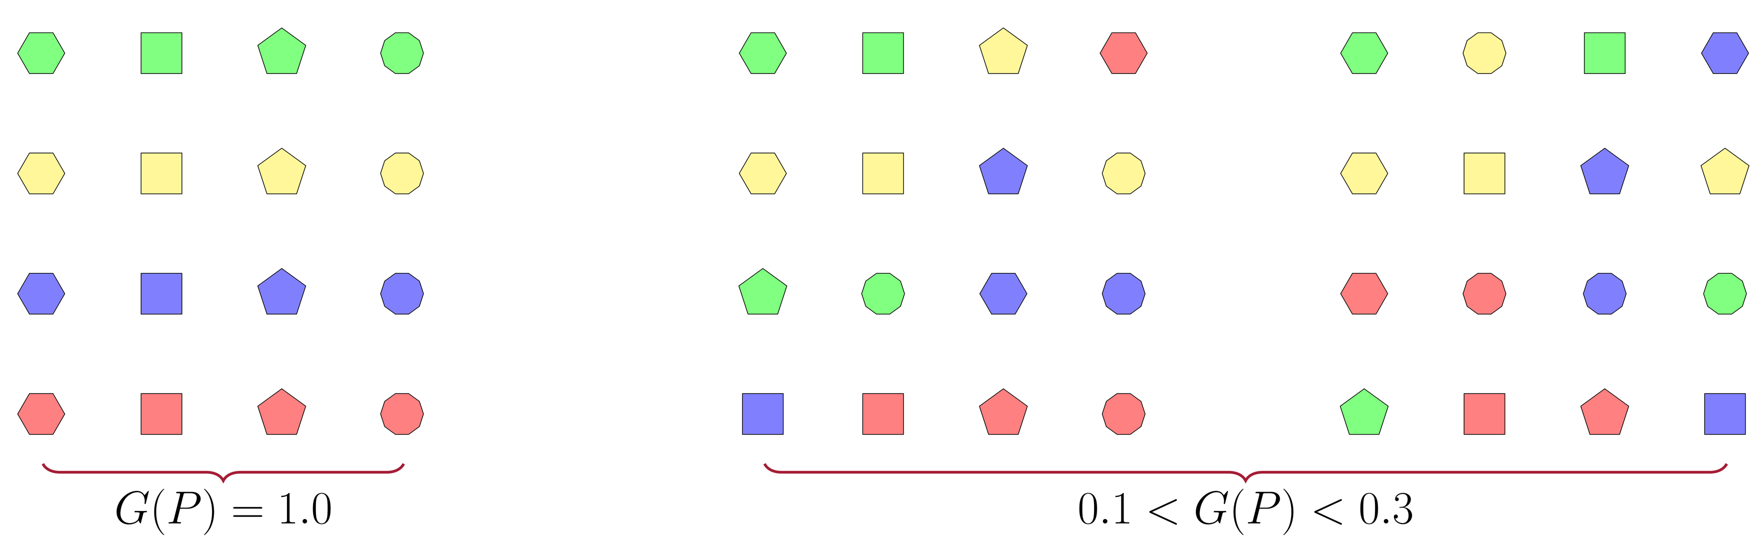

Figure 1. Three possible partitioning strategies

Given a corpus of objects and their descriptions, we can measure groundedness as a continuous attribute of a single token u or of the entire communication policy P, as follows:

$$G(u) = \frac{\max_a (C(u,a))}{C(u)} \qquad (1) \qquad\qquad G(P) = \sum_{u \in V} \frac{G(u)}{|V|} \qquad (2)$$

Importantly, we analyze models without directly optimizing for groundedness scores.

Mechanisms Leading to Compositionality

We identify four potential general cognitive mechanisms which may cause or otherwise be prerequisite to compositional referring behavior:

1) Iterated transmission effects (Smith et al., 2003; Kirby et al., 2014)
2) Noisy-channel model
3) Compression effects (Kirby et al., 2015; Yang, 2016)
4) Pragmatic reasoning

These mechanisms are neither exclusive nor exhaustive.

Task & Talk Reference Game

Kottur et al. (2017): *Natural Language Does Not Emerge 'Naturally' in Multi-Agent Dialog.*

- Sender observes referent I
- Discrete vocab $V = \{X, Y, Z, 1, 2, \ldots\}$
- Receiver observes dialogue and task G
- Shared reward if guess is correct

Task iterates until convergence to perfect accuracy.

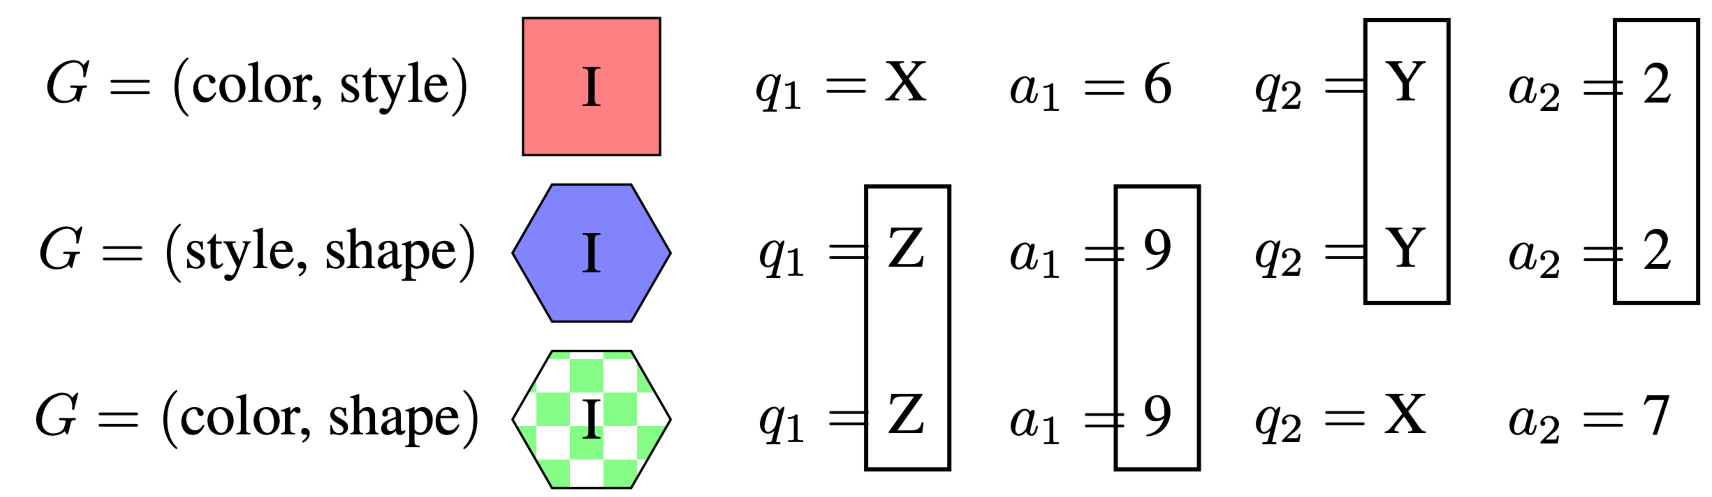

Figure 2. Example dialogues from the *Task & Talk* reference game. Note that since the first two referents share the same style (solid) the dialogues use consistent tokens $q_2 = \text{Y}$ and $a_2 = 2$ to refer to this attribute. "Y" may be interpreted as asking about the style of G, while "2" may be interpreted as the answer "solid." This idealized dialogue is an example of grounded communication, since there is a one-to-one correspondence between referent attributes and dialogue tokens.

Challenges with Human Experiments

- **Planning effects** – because senders roughly know the reference space, they may preemptively plan compositional strategies without input from the receiver
- **Number of referents** – participants should encounter memory constraints, but the task should be achievable
- **Unnaturalness** – as with other types of artificial language learning experiments, humans are influenced by their existing vocabulary and language behavior

Human Experiments

Experiment I: Dense State Space

We first replicate the *Task & Talk* signaling game with a small referent space. We observe that humans typically plan directly compositional strategies, resulting in high groundedness scores. Consequently, humans obtain near-perfect accuracy on heldout referents.

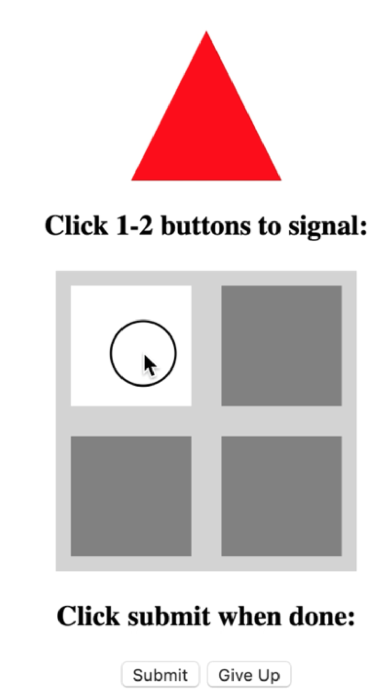

Experiment II: Sparse State Space

To prevent planning strategies, we consider a setting with the same number of referents, where it is impossible to achieve a maximum groundedness score $G(P) = 1.0$. In this setting, humans achieve high accuracy only after modifying their initial signaling strategy.

Computational Agent Results

We test four different reinforcement learning models. For each model, we run 100 iterations of *Task & Talk* until convergence and calculate mean groundedness scores.

		4x4 Baseline	4x4 Multitask
	Tabular Q-Learning	0.153	0.181
	Tabular Q-Learning (MC)	0.151	0.182
RNN	REINFORCE	0.203	0.254
	Pragmatic REINFORCE	0.210	**0.890**
LSTM	REINFORCE	0.150	0.188
	Pragmatic REINFORCE	0.153	**0.874**

Table 1. Mean policy groundedness scores (Equation 2)

While many reinforcement learning algorithms do not result in compositional referring behavior, inducing the correct biases may lead to groundedness. Further, highly compressed networks achieve slightly higher groundedness scores.

Human Subject Results

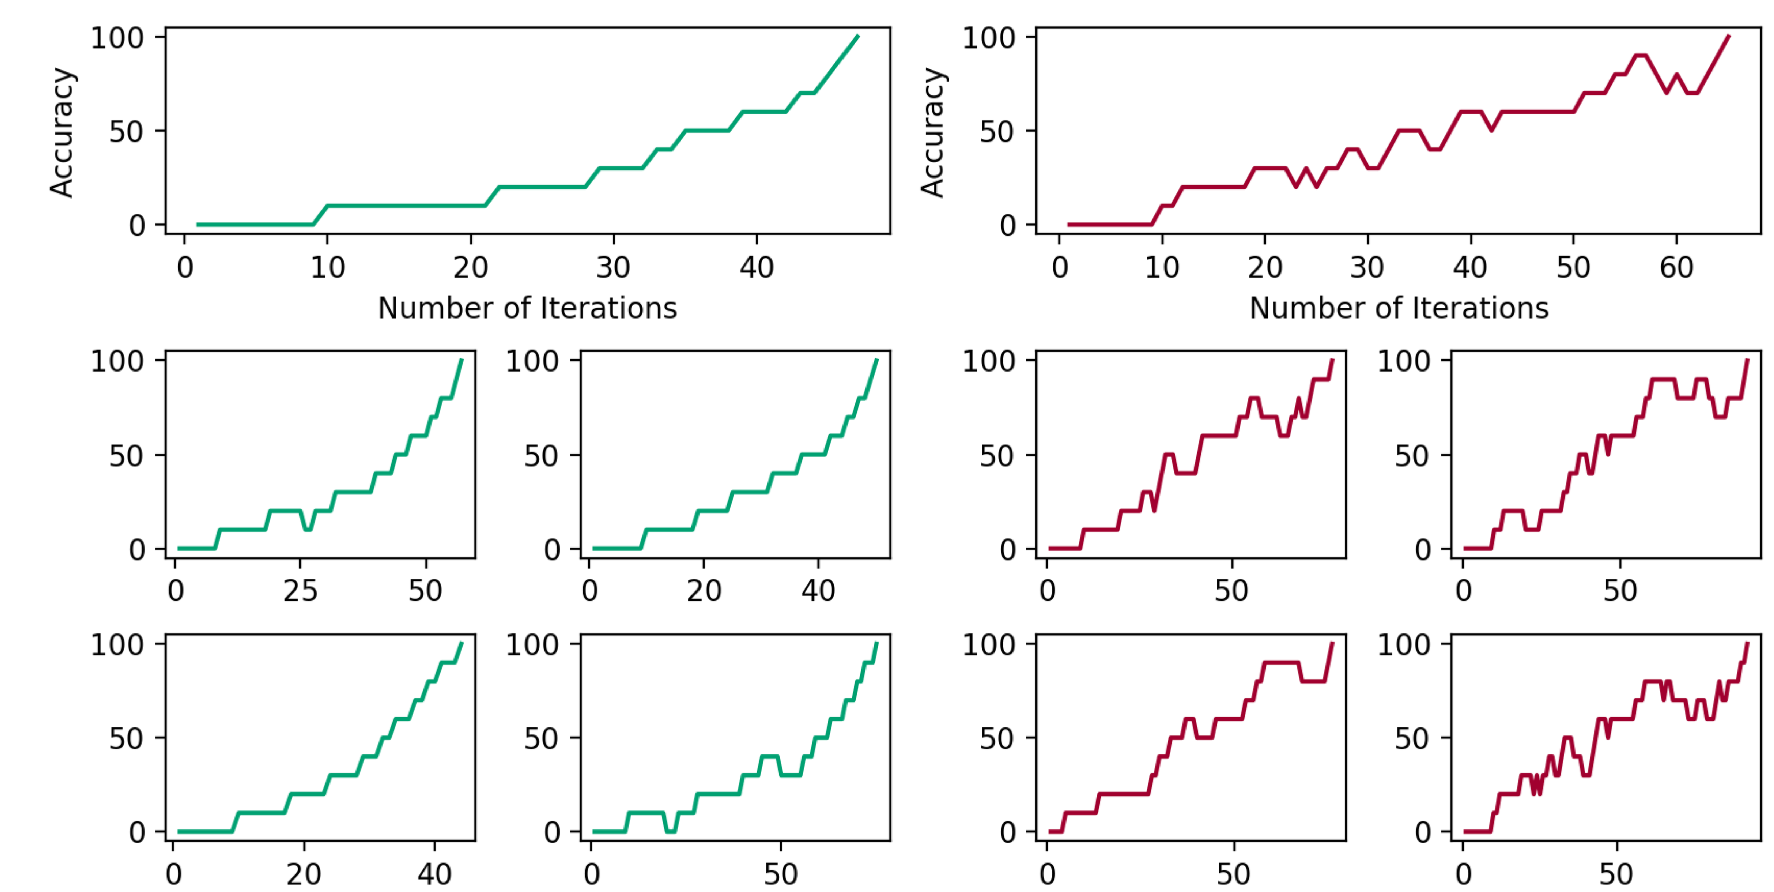

Figure 3. Learning curves from 20 human participants. Data from Experiment I is presented on the left (green), while data from Experiment II is on the right (red). Observe a U-shaped learning curve.

References

Kirby, S., Griffiths, T., and Smith, K. (2014). Iterated learning and the evolution of language. *Current opinion in neurobiology*, 28:108–114.

Kirby, S., Tamariz, M., Cornish, H., and Smith, K. (2015). Compression and communication in the cultural evolution of linguistic structure. *Cognition*, 141:87–102.

Kottur, S., Moura, J. M., Lee, S., and Batra, D. (2017). Natural language does not emerge 'naturally' in multi-agent dialog. *arXiv preprint arXiv:1706.08502*.

Smith, K., Brighton, H., and Kirby, S. (2003). Complex systems in language evolution: the cultural emergence of compositional structure. *Advances in Complex Systems*, 6(04):537–558.

Yang, C. (2016). The price of linguistic productivity.