# Incremental Pragmatics and Emergent Communication

Nicholas Tomlin and Ellie Pavlick
(Brown University)

# Groundedness in Emergent Communication

- Roughly: one-to-one correspondence between vocabulary tokens and real-world attributes;
- Useful for interpretability;
- Might be a prerequisite to productivity (cf. Kottur, et al. 2017).

$$G(u) = \frac{\max_a(C(u,a))}{C(u)} \qquad (1)$$

$$G(P) = \sum_{u \in V_A} \frac{G(u)}{|V_A|} \qquad (2)$$

# Prior Work on Groundedness in Emergent Communication

- “Emergence of Grounded Compositional Language in Multi-Agent Populations” (Mordatch & Abbeel 2017);
- “Learning Cooperative Visual Dialog Agents with Deep Reinforcement Learning” (Das, et al. 2017);
- “Natural Language Does Not Emerge ‘Naturally’ in Multi-Agent Dialog” (Kottur, et al. 2017).

# *Task & Talk*

| | | Q-Bot: Turn 1 | A-Bot: Turn 1 | Q-Bot: Turn 2 | A-Bot: Turn 2 |
|---|---|---|---|---|---|
| 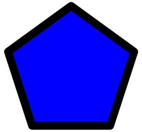 | Task: [Color, Shape]<br>Ans: [Blue, Pentagon] | X | 1 | Y | 6 |
| 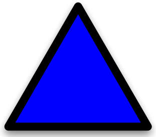 | Task: [Color, Style]<br>Ans: [Blue, Solid] | X | 1 | Z | 11 |
| 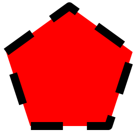 | Task: [Shape, Style]<br>Ans: [Pentagon, Dashed] | Y | 6 | Z | 12 |

# *Task & Talk*

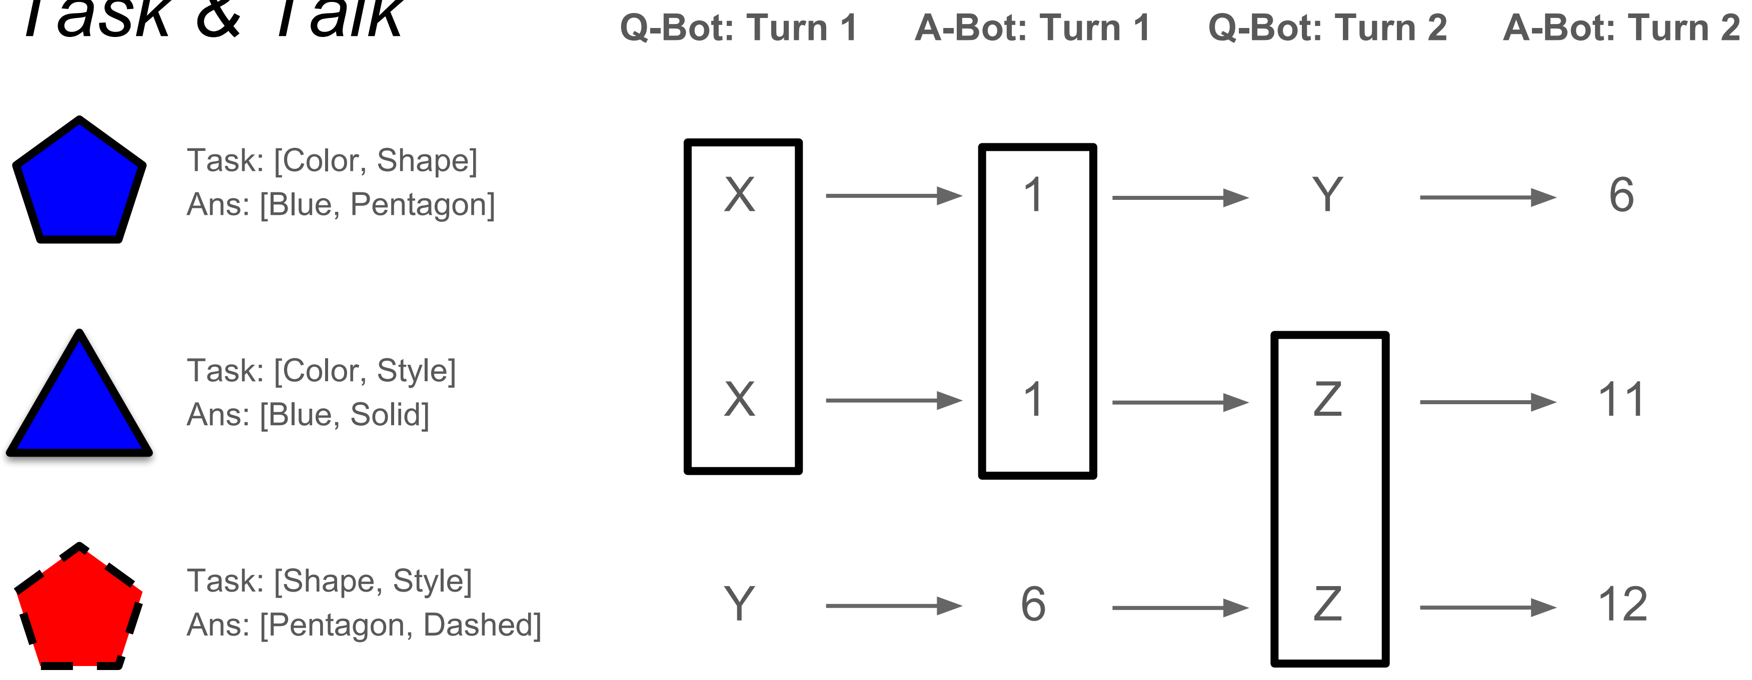
Q-Bot: Turn 1
A-Bot: Turn 1
Q-Bot: Turn 2
A-Bot: Turn 2
Task: [Color, Shape]
Ans: [Blue, Pentagon]
Task: [Color, Style]
Ans: [Blue, Solid]
Task: [Shape, Style]
Ans: [Pentagon, Dashed]
X
1
Y
6
X
1
Z
11
Y
6
Z
12

# *Task & Talk*

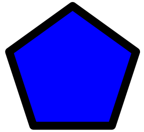

Task: [Color, Shape]
Ans: [Blue, Pentagon]

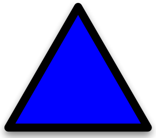

Task: [Color, Style]
Ans: [Blue, Solid]

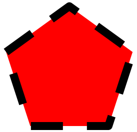

Task: [Shape, Style]
Ans: [Pentagon, Dashed]

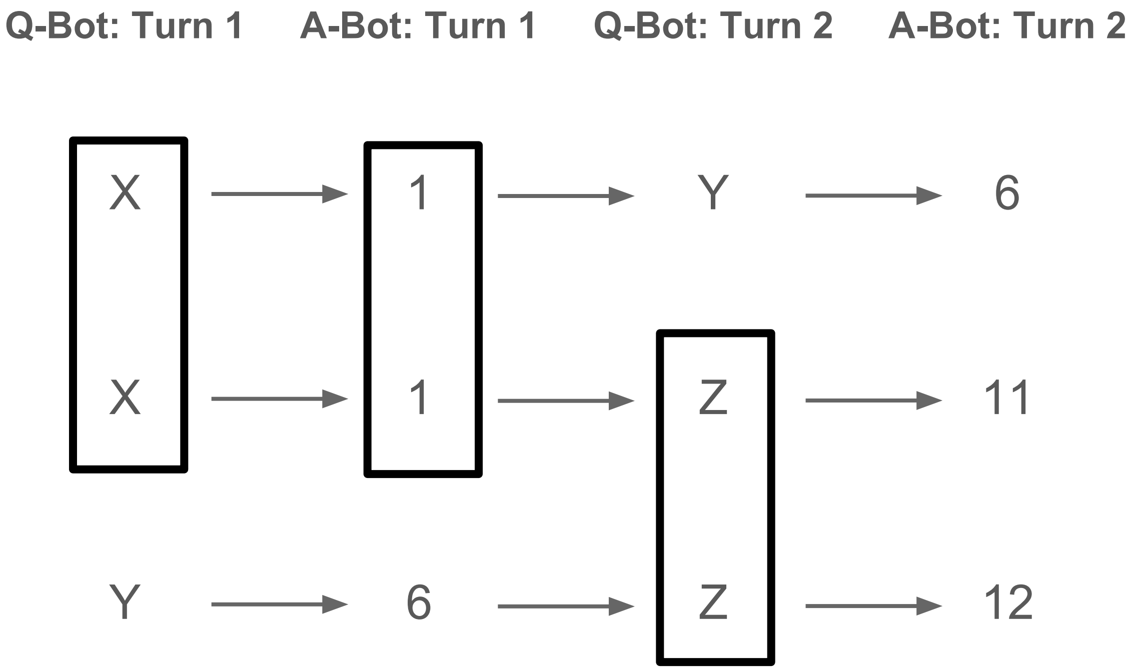


(Idealized example: the models aren't really doing this!)

# Problems with *Task & Talk*

- Reduces to “4x4 Variant” after Q-Bot’s first turn;
- Proposed changes to task design:

| | Original (Kottur et al., 2017) | 4x4 Baseline | 4x4 Multitask |
|---|---|---|---|
| Q-BOT speaks | ✓ | ✗ | ✗ |
| Q-BOT observes $G$ | ✓ | ✓ | ✗ |
| Utility function $U(\hat{w})$ | ✗ | ✗ | ✓ |
| Pragmatic model | ✗ | ✗ | ✓ |
| Curriculum learning | ✗ | ✗ | ✓ |
| Number of tasks | 3 | 1 | 3 |
| Number of referents | 64 | 16 | 16 |
| Vocab size $\|V_A\|$ | 4 | 4 | 8 |

# 4x4 Multitask

- Mixture of tasks: (shape) and (shape, color) both acceptable;
- Curriculum learning: one-attribute tasks presented first;

- Might expect that grounded communication would emerge in this scenario, but it doesn’t with tabular Q-learning or REINFORCE;
- Perhaps we’re missing some communication mechanism...

# Rational Speech Acts (Frank & Goodman 2012)

- Recursive reasoning process between speakers and listeners about alternative utterances and referents;
- Meant to capture the **cooperative principle**: be concise, truthful, informative, relevant, etc.;
- Enforces an injective mapping between referents and utterances.

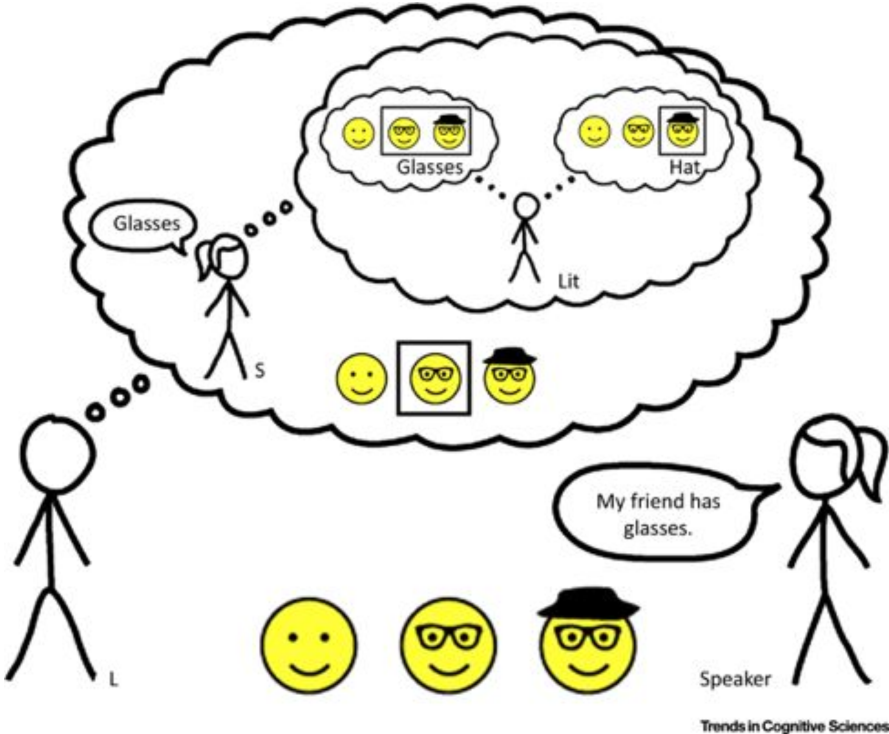

# Incremental Pragmatics

Incremental pragmatics is a well-motivated mechanism of human language processing (Sedivy, et al. 1999).

Target: “Touch the yellow bowl.”

Eye-tracking after “yellow” favors the yellow comb rather than the bowl because of the contrast effect.

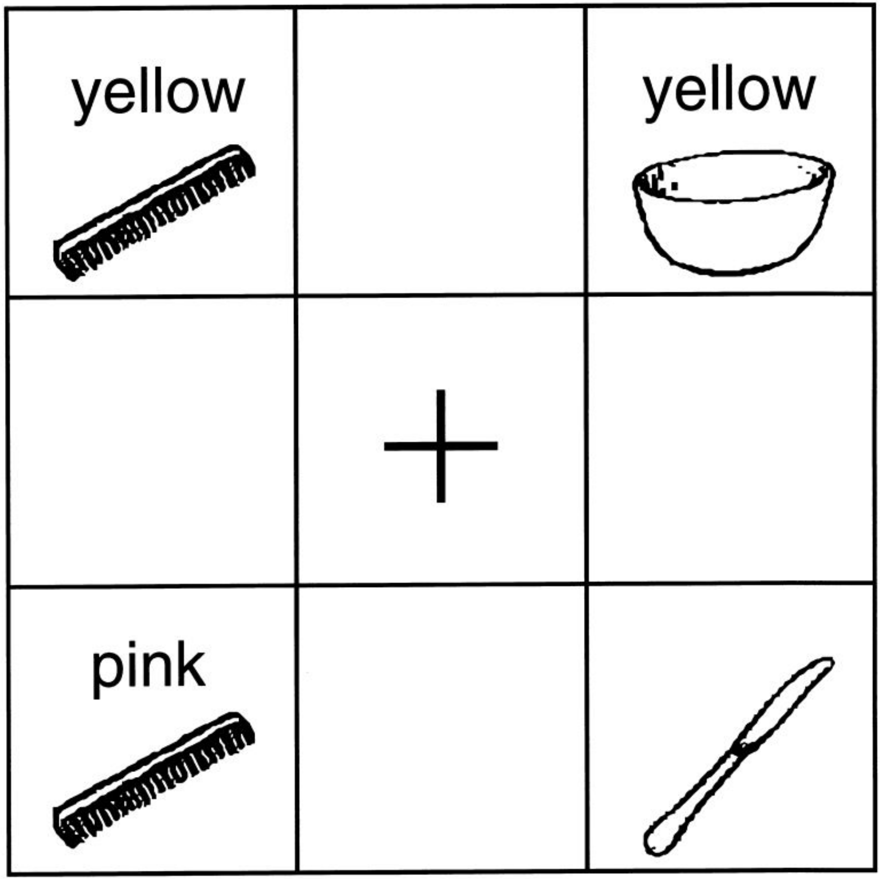

# Incremental RSA (Cohn-Gordon, et al. 2018)

Base RSA agent: [[utterance]](world)

Base incremental RSA agent: [[partial utterance]](world)

...where [[partial utterance]](world) denotes the fraction of possible utterance continuations which are consistent with the world state.

$$P(a_t \mid s_t^A) \propto \pi_A(a_t \mid s_t^A; \theta_A) \cdot \mathbb{E}_{s_0^Q} \left[ \pi_Q(\hat{w} \mid s_0^Q; \theta_Q) \cdot U(\hat{w}) \right]$$

# Model and Results

We train tabular Q-learning and REINFORCE agents on modified *Task & Talk*. The incremental pragmatic model achieves near-perfect groundedness.

Mean groundedness scores across 100 iterations:

| | 4x4 Baseline | 4x4 Multitask |
|---|---|---|
| Tabular Q-Learning | 0.153 | 0.181 |
| Tabular Q-Learning (MC) | 0.151 | 0.182 |
| REINFORCE | 0.150 | 0.188 |
| Pragmatic REINFORCE | 0.153 | **0.874** |

# Future Work

- Ablations on task modifications
- Wider domain for evaluation on held-out data
- Evaluating time-course of grounding:
  - Does RSA speed up training? (It weakly constrains the search space.)
  - Why do tokens become ungrounded? What is the effect of batch size?
- Comparison to memory efficiency models of productivity (cf. Yang 2016)
- Evaluate human performance on this task (MTurk experiment!)

Thank you!