Introduction

Why do languages partition mental concepts into words the ways they do?

Introduction

Recent work suggests that language is shaped by pressure for efficient communication[1]. This involves an information-theoretic tradeoff between

- Cognitive load, or Complexity
- Informativeness

[1] See for example: Gibson et al. (2017), Kemp et al. (2018) and Zaslavsky et al. (2019)

Numeral systems across languages reflect a need for efficient communication[2].

- Approximate systems.
- Exact systems.
- Recursive systems.

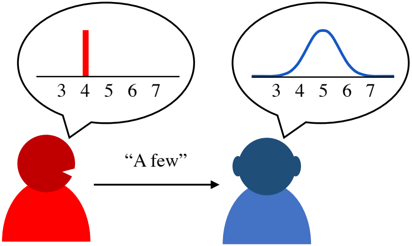

Figure 1: Communication setup studied in Xu et al. 2020

[2] Xu et al. (2020)

Introduction

Is there a computational learning mechanism that leads to efficient communication?

Introduction

Is there a computational learning mechanism that leads to efficient communication?

In this work we show how efficient approximate and exact numeral systems emerge via Reinforcement Learning.

Efficient Communication

What do we mean by efficient communication?

Efficient Communication

We measure communication cost as expected surprisal[3]

$$C = -\sum_{n,w} p(n)p(w|n)\log p(n|w).$$

We measure complexity, or cognitive load, as the number of terms used in the numeral system.

[3] Gibson et al. (2017).

Efficient Communication

We measure communication cost as expected surprisal[3]

$$C = -\sum_{n,w} p(n)p(w|n)\log p(n|w).$$

We measure complexity, or cognitive load, as the number of terms used in the numeral system.
$\rightarrow$ *An efficient numeral system should, given a certain number of terms, minimize the expected surprisal.*

[3] Gibson et al. (2017).

Signaling Game and Reinforcement Learning

We consider a Multi-Agent Reinforcement Learning setup consisting of two agents playing a signaling game with

- finite set of numbers $\mathcal{N}$
- finite set of words $\mathcal{W}$.

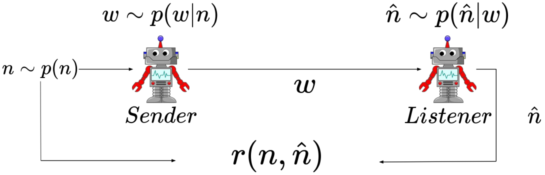

Signaling Game and Reinforcement Learning

- 3 different need probabilities estimated from human data.
- Uniformed need probability
- Power-law estimated as in Xu et al. (2020).
- We consider the range $[1, 20]$.

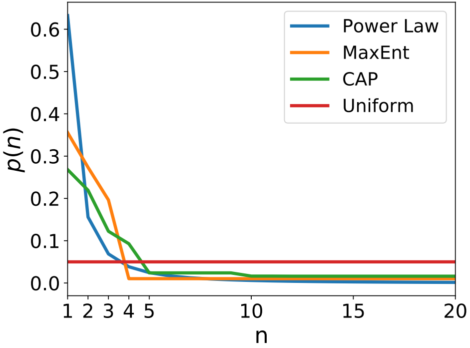

Signaling Game and Reinforcement Learning

We consider the following three reward functions

- $r_{\text{linear}}(n, \hat{n}) = 1 - \frac{|n-\hat{n}|}{20}$.
- $r_{\text{inverse}}(n, \hat{n}) = \frac{1}{1+|n-\hat{n}|}$.
- $r_{\exp}(n, \hat{n}) = e^{-|n-\hat{n}|}$.

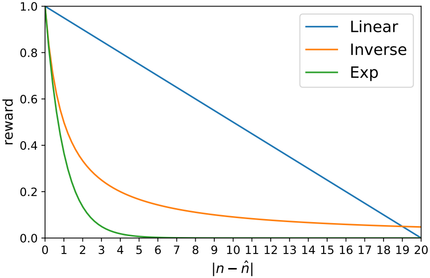

Signaling Game and Reinforcement Learning

In Q-learning an agent estimates the expected reward for each state-action pair

$$F_S : \mathcal{N} \times \mathcal{W} \longrightarrow [0,1]$$
$$F_L : \mathcal{W} \times \mathcal{N} \longrightarrow [0,1].$$

Here F_S and F_L are parameterized by a neural network with one hidden layer of 50 neurons and with ReLU activation.

Signaling Game and Reinforcement Learning

Given a number n

1. The sender applies dropout to its hidden layer to get a network f_S.

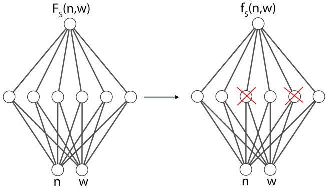

Signaling Game and Reinforcement Learning

Given a number n

1. The sender applies dropout to its hidden layer to get a network f_S.

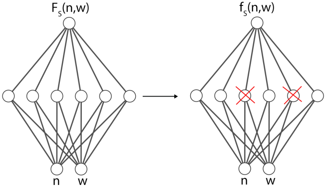

2. The expected reward are then estimated for each pair (n, w).

Signaling Game and Reinforcement Learning

Given a number n

1. The sender applies dropout to its hidden layer to get a network f_S.

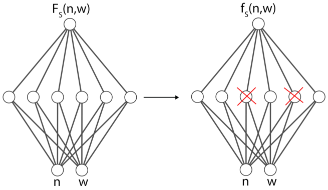

2. The expected reward are then estimated for each pair (n, w).
3. The sender conveys the word satisfying $w^* = \text{argmax}_w f_S(n, w)$.

Signaling Game and Reinforcement Learning

Given a number n

1. The sender applies dropout to its hidden layer to get a network f_S.

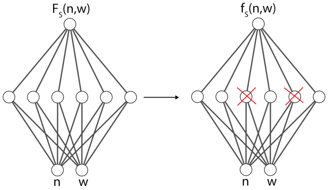

2. The expected reward are then estimated for each pair (n, w).
3. The sender conveys the word satisfying $w^* = \text{argmax}_w\, f_S(n, w)$.

Dropout encourages exploration. This can be seen as an implicit form of Thompson Sampling.

Signaling Game and Reinforcement Learning

Agents are updated by minimizing the mean-squared error between predicted reward and actual reward, using the optimizer Adam.

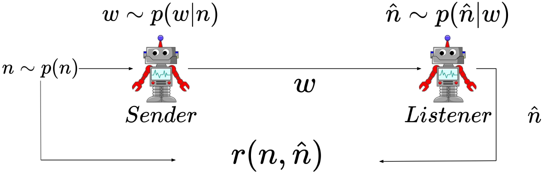

Signaling Game and Reinforcement Learning

After the agents have converged we estimate $p(w|n)$ as

$$p(w|n) \approx \frac{1}{1000} \sum_{i=1}^{1000} 1(w = \operatorname*{argmax}_{\hat{w}} f_{S,i}(\hat{w}, n)).$$

- If $p(w|n)$ is not peaked we treat it as an approximate numeral system.
- We take the mode of $p(w|n)$ as an exact numeral system.

Experimental Setup

- Maximum vocabulary size of 10.
- For each combination of need and reward we trained 6000 independent sender-listener pairs.
- We trained each pair for 10 000 updates.
- Batch size was 100.

Results

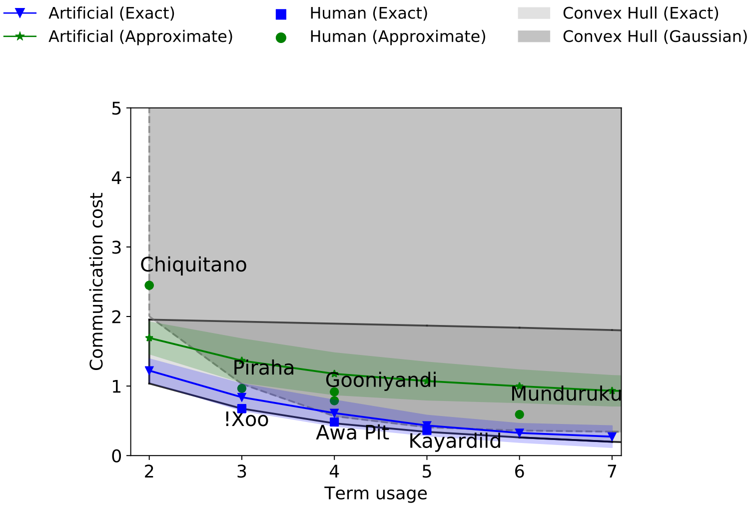

Figure 2: Linear reward and power-law prior.

Results

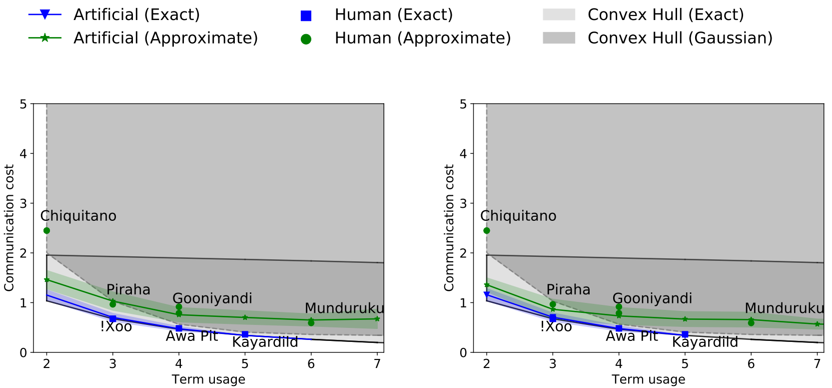

(a) Reward: Inverse, Prior: Power law

(b) Reward: Exponential, Prior: Power law

Results

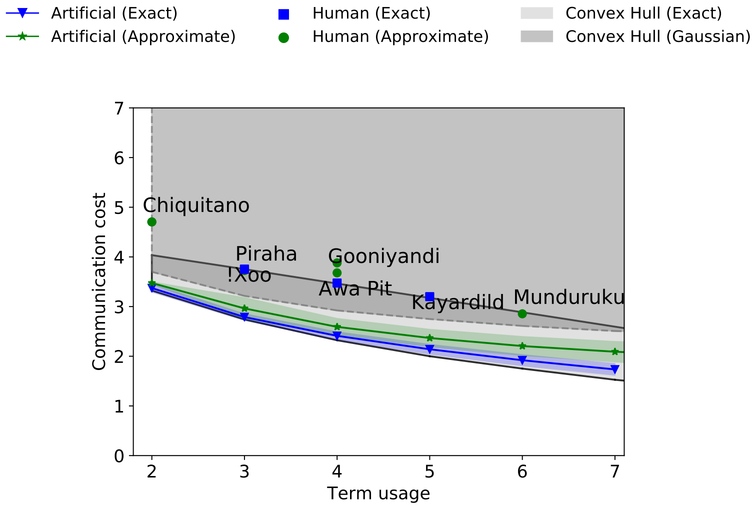

Figure 4: Linear reward and uniformed prior.

Results

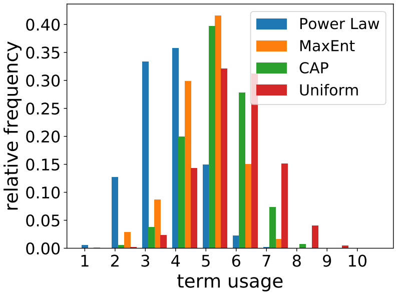

(a) Normalized frequencies

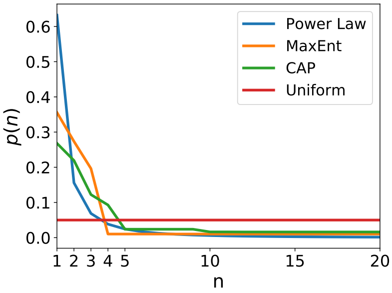

(b) Need Probabilities

Results

Consensus system using Correlation Clustering

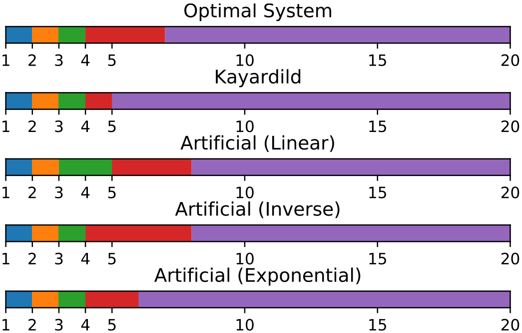

Figure 6: Consensus systems for 5 terms using the power-law prior.

Conclusions

- RL agents can learn task specific communication protocols which are near-optimal in information-theoretic sense.
- Same level of efficiency as human systems and with similarities between artificial and human languages.

References

- Yang Xu, Emmy Liu, and Terry Regier.Numeral Sys-tems Across Languages Support Efficient Communication:From Approximate Numerosity to Recursion. 2020.
- Noga Zaslavsky, Charles Kemp, Naftali Tishby, and TerryRegier. Color naming reflects both perceptual structure andcommunicative need. 2019.
- Edward Gibson, Richard Futrell, Julian Jara-Ettinger, KyleMahowald, Leon Bergen, Sivalogeswaran Ratnasingam,Mitchell Gibson, Steven T. Piantadosi, and Bevil R. Conway. Color naming across languages reflects color use.2017.
- Charles Kemp, Yang Xu, and Terry Regier. Semantic typology and efficient communication. 2018.