Efficient learning for emergent communication

Emil Carlsson, Devdatt Dubhashi, and Fredrik Johansson

Department of Computer Science and engineering, Chalmers University of Technology, Gothenburg, Sweden.
{caremil, dubhashi, fredrik.johansson}@chalmers.se

1 Introduction

This project investigates central questions at the intersection of cognitive science and machine learning: how is language used as a means of efficiently communicating mental concepts and how does it emerge? An influential approach in cognitive science treats language as a means to communicate efficiently in an information theoretic sense. This is illustrated schematically in figure 1 as an interaction between a sender who has a concept in mind that she wishes to convey to a receiver by transmitting a discrete message over a communication channel.

Given a concept c, the sender utters a word from a vocabulary according to a distribution $w \sim p(w|c)$. When given a word, the receiver guesses a concept according a distribution $\hat{c} \sim p(\hat{c}|w)$. The efficiency of the communication scheme is measured by how well the receiver recreates the original concept as measured by the KL divergence between the receiver and sender distributions. Cognitive scientists have argued that human languages have evolved to optimize communication efficiency in this sense. Here we take a computational approach and address two fundamental questions:

- Can general computational learning mechanisms lead to the emergence of efficient communication schemes?
- What is the best achievable computational complexity of such learning schemes as measured by the efficiency of learning algorithms?

To address these questions, we place them in the setting of multi-agent games with reward signals. Once a receiver has reconstructed the concept as above, a shared reward $r(c, \hat{c})$ is given to both agents which measures how good the reconstruction was. The agents then update their belief distributions according to this reward. Central frameworks for learning in such settings include multiarm bandit algorithms and reinforcement learning. In contrast to the cognitive science setting, we start both agents from *tabula rasa*: The words in the vocabulary will be initialized without any meaning and it is up to the two agents to agree on their meaning during several rounds of the game. Hence, the agents are inventing their own language while playing the game. Thus studying this system also gives insights in the emergence of language and its universal principles.

In the case where the number of possible concepts is larger than the vocabulary size, the agents will have to assign several concepts to the same words and perform lossy compression. In the extreme, the concepts exists in some continuous space and the words uttered by the sender will have so symbolise a certain

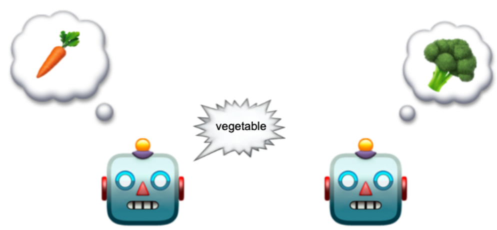

Fig. 1: The sender thinks about the concept carrot and utters the word vegetable. The receiver then guesses the concept it thinks the sender was referring to.

part of the space. In both these cases it is intresting to study how fast the agents converges, i.e agree on a language, and how efficient the resulting language is w.r.t the environment $p(c)$ and the reward function $r(c, \hat{c})$.

2 Research direction

A recent study considers the emergent language when two agents plays a referential game based on colors [1]. The agents was represented as neural networks and was updated using the policy gradient method REINFORCE. A color was represented as a vector $c \in \mathbf{R}^3$ and sampled from the Munsell chart used in the World Color Survey. It was found that artificial agents produced partitions of the color space close to partitions that are found in real human languages.

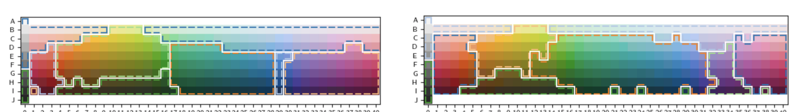

(a) Artificial language of 10 terms. (b) Human language of 10 terms.

Fig. 2: Comparison between a language emerged from our multi-agent setting and a human language. The encapsulated areas symbolises words.

In future work, we would like to investigate the theoretical properties when it comes to the convergence of language in our setting. We would also like to investigate if there exists theoretical guaranties on the efficiency of the resulting language.

At last, we would like to extend the study done in [1] to other domains than colors. For example numeral systems or semantic relations similar to WordNet.

References

1. Kågebäck, M., Dubhashi, D.P., Sayeed, A.B.: Deepcolor: Reinforcement learning optimizes information efficiency and well-formedness in color name partitioning. Cognitive Science (2018)