# Communication and Language Emergence Among Populations and Clusters of Agents

Graham Todd
Stanford University, CS224N

## Problem

Recent research has investigated the emergence of linguistic communication through the use of **communication games**.

However, most efforts focus on only a **single pair** of agents (one **speaker** and one **listener**).

We combine approaches and methods from multiple researchers, and extend a classic communication game to a **population of agents**.

## Data / Task

We study an extension of the **Lewis Signaling Game**..

At each step, a speaker receives a **target** object (which has many attributes). It then produces a **message** of discrete symbols. The listener receives the message, the target, and a set of **distractor** objects. It interprets the message and then attempts to select the target correctly from the distractor set. If it correctly identifies the target, we call it a **communication success**.

We extend the game to a population of agents with **clusters.** Each agent can both speak and listen and is more likely to communicate **within its cluster**.

We use a variant of the **ViSA dataset**, with 250 concepts and 300 possible attributes.

## Approach

Both speaker and listener systems are primarily driven by a **single-layer LSTM**. As demonstrated by the diagram below, from [1].

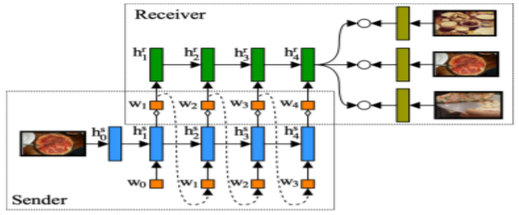


Normally, discrete choices (like symbol selection for the messages) cannot be **backpropagated through.** However, with the **Gumbel Softmax function** (described in [1]), we are able to approximate the categorical distribution with a differentiable function. The Straight-Through Gumbel Softmax trick uses discrete symbols in the forward pass and a differentiable approximation in the backwards pass, allowing for standard neural network optimization techniques.

For our extension of the game, we modify the **"hardness"** of the clustering—the probability that an agent selects a partner from within its cluster. We're interested in whether hardness has a **non-linear** effect on **overall population performance**.

Since agents in our version can both speak and listen, we allow agents to **listen to themselves**. This encourages a single agent's speaker and listener systems to converge to a single language.

## Results

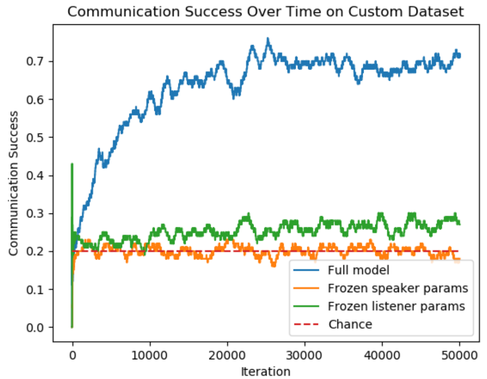


Above is a plot of the **communication success** of a **single pair** of agents performing the standard signaling task. We see that they are able to **significantly outperform** chance when the **full model** is used.

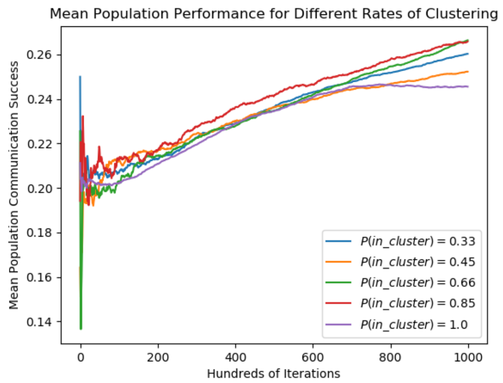


Above is a plot of the performance (mean communication success) of **multiple populations** for different values of clustering hardness. We see that optimal results are achieved with intermediate hardness values, implying **nonlinearity.**

## Analysis / Conclusion

Our experiment provides some **tentative evidence** for a nonlinear relationship between clustering hardness and population performance.

However, the differences in performance were **quite small** and the overall performance in the population signaling game was **quite low**.

This is partly due to the **high difficulty** of the population signaling game. Adding more agents to the population produces an **interference effect,** making converging to a single language much harder. Clustering is designed to help with this, but likely does not fully alleviate the challenge.

It may also be due to the fact that agents did not make use of their **entire vocabulary**.

Ultimately, we present some **early indication** that population dynamics among artificial agents are rich and nonlinear and may warrant **addition research**.

## References

[1] Serhii Havrylov and Ivan Titov. Emergence of language with multi-agent games: Learning to communicate with sequences of symbols. *CoRR*, abs/1705.11192, 2017

[2] Angeliki Lazaridou, Karl Moritz Hermann, Karl Tuyls, and Stephen Clark. Emergence of linguistic communication from referential games with symbolic and pixel input. *CoRR*, abs/1804.03984, 2018.