

Zero-Shot Generalization using Intrinsically Motivated Compositional Emergent Protocols

Rishi Hazra*, Sonu Dixit*, Sayambhu Sen

{rishihazra, sonudixit, sayambhusen}@iisc.ac.in

Indian Institute of Science, Bangalore

Abstract

Human language has been described as a system that makes *use of finite means to express an unlimited array of thoughts*. Of particular interest is the aspect of compositionality, whereby, the meaning of a compound language expression can be deduced from the meaning of its constituent parts. If artificial agents can develop compositional communication protocols akin to human language, they can be made to seamlessly generalize to unseen combinations. Studies have recognized the role of curiosity in enabling linguistic development in children. In this paper, we seek to use this intrinsic feedback in inducing a systematic and unambiguous protolanguage. We demonstrate how compositionality can enable agents to not only interact with unseen objects but also transfer skills from one task to another in a zero-shot setting: *Can an agent, trained to ‘pull’ and ‘push twice’, ‘pull twice’?*

1 Introduction

In the recent past, there has been a great deal of research in the field of emergent language in artificial agents interacting in simulated environments (Kirby, 2001; Havrylov and Titov, 2017; Tieleman et al., 2019; Gupta et al., 2020). However, the real question here is, to what extent do these evolved protocols resemble natural language? Recent studies have revealed the following about emergent languages: (i) they do not conform to Zipf’s Law of Abbreviation (Chaabouni et al., 2019); (ii) communication protocols either do not follow compositionality patterns of natural language (Kottur et al., 2017) or are not always interpretable (Lowe et al., 2019); (iii) emerged protocols are sensitive to experimental conditions (Lazaridou et al., 2018).

Although compositionality is not crucial to achieving generalization, more compositional protocols have been shown to display higher zero-shot

performance (Ren et al., 2020). While work on incorporating compositionality into emergent languages is still in its early stages, certain works (Mordatch and Abbeel, 2018; Chaabouni et al., 2020) have proposed to use limited channel capacity as a means to achieve composition. However, we argue that agents may fail to develop meaningful communication protocols in such a restricted setting. Motivated by human behavior, we formulate intrinsic rewards (Gopnik et al., 2001; Baldassarre and Mirolli, 2013) to provide incentives to the agents for paying attention to communication despite having a limited channel capacity. Forced to deal with it to earn more intrinsic rewards, the agents must learn to use a more systematic and unambiguous protolanguage.

As proof of concept, we push the boundaries of compositionality to a more challenging multi-task settings, arguing that it can also support the acquisition of a more complex repertoire of skills (performing a *pull twice* task when it has been trained to *pull, push and push twice*), in addition to generalizing over novel composition of object properties (pushing *red square* when it has been trained to push a *red circle* and a *blue square*)¹.

2 Problem Setup

We analyze a typical signalling game (Lewis, 1969), comprising a stationary Speaker-Bot (*speaker*) and a mobile Listener-Bot (*listener*), by modelling it in form of a Markov Decision Process specified by the tuple $(\mathcal{S}, \mathcal{O}, \mathcal{A}, r, \mathcal{T}, \gamma)$. At the beginning of each round, the speaker receives a natural language instruction (*push a red circle*) and communicates the same using discrete messages $m_{i=1}^{n_m}$, sampled from a message space \mathcal{M} , to the listener over a communication channel. Here, d_m is the dimension of the message m_i , and n_m is the number of messages (these constitute the channel capacity).

¹demos:<https://sites.google.com/view/compositional-comm>

* Equal Contribution

ity, $|C| = c_{d_m}^{n_m}$). At each step t , the listener receives an observation $\mathbf{o}^{(t)} \in \mathcal{O}$, comprising the 2D grid-view and the received messages $m_{i=1}^{n_m}$, and takes an action $\mathbf{a}^{(t)} \in \mathcal{A}$. The goal of the listener is to choose optimal actions according to a policy $\pi : (\mathcal{O}, m_{i=1}^{n_m}) \mapsto \Delta(\mathcal{A})$, to maximize its long-term reward $\mathcal{R} = \sum_t \gamma^t r^{(t)}$. Here, γ is the discount factor and \mathcal{T} is the transition function $\mathcal{T} : \mathcal{S} \times \mathcal{A} \mapsto \mathcal{S}$. The environment generates a 0-1 (sparse) reward, i.e., the listener gets a reward of $r = 1$ if it achieves the specified task, otherwise $r = 0$. However, the listener has no information about either the task or the target, and relies on the speaker for the same. Given a language $\mathcal{L}(\cdot) : \mathcal{C} \mapsto \mathcal{M}$, we use topographic similarity (*topsim*) (Brighton and Kirby, 2006) between \mathcal{C} (set of concepts) and \mathcal{M} (set of messages) as a measure of compositionality. Our work is contrasted with that of gSCAN (Ruis et al., 2020) which focuses on rule-based generalization using a supervised learning framework.

3 Approach

3.1 Environment Description

In our experiments, we use a 4×4 grid. Cells in the grid contain objects characterized by certain attributes like shape, size, color and weight. These objects can either be the *target* object or the *distractor* objects. Distractors have either the same color or the same shape (or both) as that of the target. We keep the number of distractors fixed ($= 2$). The listener and the objects may spawn at any random location on the grid. Given an instruction, it is first processed using a parser to $\langle \text{VERB}, \{\text{ADJ}_i\}_{i=1}^3, \text{NOUN} \rangle^2$. The speaker transmits the same using a set of messages to the listener which, then, processes the grid representation and the received messages to achieve the given task. In our experiments, we use a $\{0, 1\}^{d_{grid} \times 4 \times 4}$ vector array for the grid representation, where each cell has a d_{grid} -dimensional encoding.

3.2 Model Description

(1,2) The speaker receives the parsed input instruction parsed. (3) The speaker uses an encoder to map the concept input to a hidden representation $\in \mathbb{R}^{n_m \times d_h}$. From this representation, a set of one-

²VERB: task ('walk', 'push', 'pull'); ADJ: object attributes like color ('red', 'blue', 'yellow', 'green'), size ('small', 'big') and weight ('light', 'heavy'); NOUN: object shape ('square', 'circle', 'cylinder', 'diamond')




Figure 1: Model Description

hot encoded messages $m_{i=1}^{n_m} \in \{0, 1\}^{d_m}$ are sampled (during training) using Categorical sampling, which are then transmitted over the communication channel. The number of messages n_m is set to $|\langle \text{VERB}, \{\text{ADJ}_i\}_{i=1}^3, \text{NOUN} \rangle|$. During evaluation, sampling is replaced with an $\arg \max(\cdot)$. We use the Straight Through trick (Jang et al., 2017) to retain differentiability. (4) At each step, the grid input is mapped in the Grid Encoder to an output $\mathcal{G}_t \in \mathbb{R}^{d_g \times 4 \times 4}$. (5) Next, we compute the attention weights $\alpha_{i=1}^{16}$ for each grid cell by taking a normalized dot product between z and $\mathcal{G}_t^{i \times d_g}$. A weighted combination is fed to the sub-policy networks.

The listener learns to (i) identify the target object in the context of distractors, (ii) interact with the target object by executing a task specified by the speaker. We use a hierarchical-RL framework (Sutton et al., 1999) for training. There are two sub-policies corresponding to the PUSH and the PULL tasks. (6, 7) In each round, the master policy selects either sub-policies using the received messages³. The sub-policies have a shared input, which includes the grid encoder and the attention network. The whole framework is trained end-to-end using REINFORCE (Williams, 1992).

In order to induce a more efficient training, we keep a measure of the Learning Progress (LP) of the listener for all tasks on a *held-out* set, where LP for task i is given as $\text{LP}_i = |r_i - \mu_i|$. Here, μ_i denotes the running mean of rewards for task i . The tasks are sampled from a Categorical distribution with probabilities $p(i) = \frac{\text{LP}_i}{\sum_j \text{LP}_j}$ and, consequently, episodes corresponding to the sampled tasks are generated. This way, the listener can keep track of goals that are already learned, or can insist on goals that are currently too hard.

3.3 Inducing Compositionality

We would ideally want the concept to message mapping to be *injective* (one-to-one), i.e. $\forall c, \tilde{c} \in \mathcal{C}, \mathcal{M}(c) = \mathcal{M}(\tilde{c}) \implies c = \tilde{c}$. Furthermore, the

³actions spaces: master policy: {A, B, Null}; subpolicy A/B: {left, right, forward, backward, push/pull}

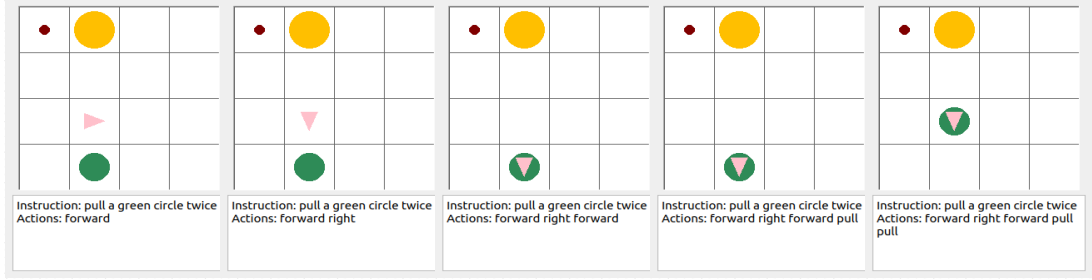


Figure 2: [Best viewed in color] Demonstration of Intrinsic Speaker on the numeral split for task PULL TWICE. Here, the green circle is heavy, hence the listener has to apply two units of force (TWICE) to pull it.

messages in \mathcal{M} must exhibit a systematic structure (in holistic languages, one can satisfy the injective property without being compositional). Studies on language evolution have proposed limiting the channel capacity of communication as a constraint for achieving compositionality (Nowak and Krakauer, 1999). Yet, in the course of our experiments, on increasing $|\mathcal{C}|$, we observed rather predictably that, with a limited channel capacity, it becomes increasingly difficult for the speaker to converge upon a consistent and unambiguous mapping from \mathcal{C} to \mathcal{M} . Consequently, the listener would either ignore the information from the speaker (*speaker abandoning*), or may exploit the inadequate information (*undercoverage*⁴) to converge on a local optimum (learning a fixed sequence of actions). To that end, we propose two types of intrinsic rewards to address these issues.

Undercoverage: The limited channel capacity acts as an information bottleneck, impeding the speaker’s ability to transmit unambiguously. Therefore, it becomes difficult for the listener to infer the decoded messages at its end. To address this issue, we formulate a notion of compositionality from recent works in disentanglement (Higgins et al., 2017). We propose to use the Mutual Information (MI) between the concepts and the messages $I(\mathcal{C}, \mathcal{M})$ as an intrinsic reward:

$$\begin{aligned} I(\mathcal{C}, \mathcal{M}) &= H(\mathcal{C}) - H(\mathcal{C}|\mathcal{M}) \\ &= H(\mathcal{C}) + \mathbb{E}_{c \sim \mathcal{C}, m \sim \mathcal{M}(c)} \log p(c|m) \end{aligned}$$

Given that the training episodes are generated independent of the object specifications, $H(\mathcal{C})$ can be assumed to be constant. We approximate the last term using Jensen’s inequality ($\mathbb{E}_{c \sim \mathcal{C}, m \sim \mathcal{M}(c)} [\log p(c|m)] \geq$

⁴Inspired by machine translation works (Tu et al., 2016), we define coverage as a mapping from a particular concept element to its appropriate message element. Full coverage refers to a distinct mapping of the whole concept input to corresponding symbols in \mathcal{M} .

$\mathbb{E}_{c \sim \mathcal{C}, m \sim \mathcal{M}(c)} [\log q_\phi(c|m)]$) to obtain a lower bound for $I(\mathcal{C}, \mathcal{M})$. Here, $q_\phi(c|m)$ is a learned discriminator module which takes the (concatenated) messages and tries to predict the concept labels (i.e. elements of $\langle \text{VERB}, \{\text{ADJ}_i\}_{i=1}^3, \text{NOUN} \rangle$) and $\mathbb{E}_{c \sim \mathcal{C}, m \sim \mathcal{M}(c)} \log q_\phi(c|m)$ is its negative cross-entropy loss. The final intrinsic reward is:

$$I(\mathcal{C}, \mathcal{M}) \geq H(\mathcal{C}) + \mathbb{E}_{c \sim \mathcal{C}, m \sim \mathcal{M}(c)} \log q_\phi(c|m) \quad (1)$$

Intuitively, it suggests that it should be easy to infer the concepts from the messages. Conversely, the confusion (high error) arising from the speaker’s inability to express concepts will lead to lower rewards. Note, that the reward will be highest when the conditions of full coverage and one-to-one mapping are satisfied (the discriminator will then be able to predict all the concept elements with high probability). We add the $I(\mathcal{C}, \mathcal{M})$ reward at the last step of the episode, given as: $r[-1] + \lambda_1 I(\mathcal{C}, \mathcal{M})$, where λ_1 is a tunable hyperparameter. The discriminator q_ϕ is periodically trained using batches sampled from a memory buffer, where we store the pair $\langle c_i, m_i \rangle$. Note, that we block the discriminator gradients to the speaker and use it merely as an auxiliary means to provide intrinsic feedback.

Speaker Abandoning Existing works (Lowe et al., 2019) have shown that while training RL-agents augmented with a communication channel, it is likely that the speaker fails to influence the listener’s actions. To address this, we propose to add another intrinsic reward to maximize the mutual information between the speaker’s messages and the listener’s actions, given the grid information.

At each step, we simulate k intermediate steps to sample pseudo messages \tilde{m} from the message distribution \mathcal{M} . Together with the original message m , we compute two sets of probability values corresponding to actions of the listener: (i) $\pi(a_t|m, \mathcal{G}_t)$ which is the listener’s policy conditioned on both the messages and the output of the grid encoder

\mathcal{G}_t ; (ii) $p(a_t|\mathcal{G}_t)$ or the probability distribution over the listener’s actions conditioned on just the output of the grid encoder. We then calculate the mutual information for each step as follows:

$$\begin{aligned} I(a_t, m|\mathcal{G}_t) &= \sum_{a_t, m} p(a_t, m|\mathcal{G}_t) \log \frac{p(a_t, m|\mathcal{G}_t)}{p(a_t|\mathcal{G}_t)p(m|\mathcal{G}_t)} \\ &= \sum_{a_t, m} p(m|\mathcal{G}_t)p(a_t|m, \mathcal{G}_t) \log \frac{p(a_t|m, \mathcal{G}_t)}{p(a_t|\mathcal{G}_t)} \\ &= \mathbb{E}_{m \sim \mathcal{M}} [D_{KL}(p(a_t|m, \mathcal{G}_t)||p(a_t|\mathcal{G}_t))] \end{aligned}$$

Note that $p(m|\mathcal{G}_t) = p(m)$ since messages and grid-view are independently processed. Here $p(a_t|\mathcal{G}_t)$ is obtained by marginalizing over the joint probability distribution, given as, $\sum_{\tilde{m}} p(a_t, \tilde{m}|\mathcal{G}_t) = \sum_{\tilde{m}} p(a_t|\tilde{m}, \mathcal{G}_t)p(m)$. We use Monte Carlo approximation to replace the Expectation by sampling messages from \mathcal{M} . The final reward equation for k pseudo-steps is given as:

$$\begin{aligned} I(a_t, m|\mathcal{G}_t) &= \frac{1}{k} \sum_m D_{KL} [\boldsymbol{\pi}(a_t|m, \mathcal{G}_t) || \sum_{\tilde{m}} \boldsymbol{\pi}(a_t|\tilde{m}, \mathcal{G}_t)p(m)] \quad (2) \end{aligned}$$

Maximizing Equation 2 leads to a higher speaker influence on the listener’s actions. The net reward at each step is given as: $r_t + \lambda_3 I(a_t, m|\mathcal{G}_t)$, where λ_3 is a tunable hyperparameter.

4 Experiments

Zero-Shot Generalization Splits: (i) **Visual split:** All episodes not containing the ‘red square’ as a target object, were used for training the model. During evaluation, we examine whether the trained model can generalize to the following instructions: *walk to a red square*; *push/pull a red square*. (ii) **Numeral split:** The training set contains instructions with *Push*, *Push Twice* and *Pull*, whereas, test set contains *Pull Twice* task. Here the modifier *Twice* is used to denote a heavier object i.e., listener should execute two consecutive ‘pull’ actions to move the object. The listener must infer from its training that a symbol corresponding to *heavy* requires twice as many actions.

Baselines: We compare our Intrinsic Speaker model with the following baselines. (i) **Oracle Listener:** For each cell, we zero-pad the grid encoding with an extra bit, and set it ($= 1$) for the cell containing the target object. This way, the listener has complete information about the target in context of the distractors. We use this baseline as our upper limit of performance. (ii) **Perfect Speaker:**




Figure 3: [Best viewed in color] **Left:** Comparison of Intrinsic Speaker with other baselines on a single policy module for WALK task. It can be observed that Intrinsic Speaker performs as well as the Perfect Speaker baseline; **Right:** Comparison of topsim metric of Intrinsic Speaker (with and without feedback) and Simple Speaker. All plots have been obtained by averaging the validation rewards obtained over 5 independent runs. [X-axis: 1 unit = 50 episodes]

Task	Model	Zero-Shot Accuracy
<i>walk to a red square</i> (visual split)	Simple Speaker	73.43%
	Intrinsic Speaker	80.24%
<i>push a red square</i> (visual split)	Simple Speaker	67.17%
	Intrinsic Speaker	72.45%
<i>pull a red square</i> (visual split)	Simple Speaker	66.80%
	Intrinsic Speaker	73.29%
<i>pull a red square twice</i> (numeral split)	Simple Speaker	65.25%
	Intrinsic Speaker	69.77%

Table 1: Comparison of simple speaker and intrinsic speaker zero-shot performance on different splits.

The speaker uses an Identity matrix that channels the input directly to the listener. Thus, it is perfectly compositional. (iii) **Simple Speaker:** Here the speaker-listener is trained end-to-end without using the intrinsic rewards.

5 Results

(i) The proposed Intrinsic Speaker outperforms the Simple Speaker in terms of both, convergence rewards and topsim score (Figure 3). In fact, the Intrinsic Speaker matches the performance of the Perfect Speaker, thus, showing that the emergent communication is highly compositional (≈ 0.9). (ii) The zero-shot generalization accuracy in Table 1 shows that the Intrinsic Speaker consistently outperforms the Simple Speaker on both splits. (iii) In order to test the effectiveness of intrinsic rewards in inducing compositionality, we trained the Intrinsic Speaker with no external reward from the environment. As shown in Fig 3 (right), the intrinsic rewards were alone capable of generating a topsim score ≈ 0.6 . For further details, the readers are requested to refer the paper (Hazra et al., 2020).

References

- Gianluca Baldassarre and Marco Mirolli. 2013. *Intrinsically Motivated Learning Systems: An Overview*, pages 1–14. Springer Berlin.
- Henry Brighton and Simon Kirby. 2006. Understanding linguistic evolution by visualizing the emergence of topographic mappings. *Artif. Life*.
- Rahma Chaabouni, Eugene Kharitonov, Diane Bouchacourt, Emmanuel Dupoux, and Marco Baroni. 2020. Compositionality and generalization in emergent languages. *CoRR*, abs/2004.09124.
- Rahma Chaabouni, Eugene Kharitonov, Emmanuel Dupoux, and Marco Baroni. 2019. Anti-efficient encoding in emergent communication. *CoRR*, abs/1905.12561.
- Alison Gopnik, Andrew Meltzoff, and Patricia Kuhl. 2001. *The Scientist in the Crib: Minds, Brains and How Children Learn*, volume 189. HarperCollins.
- Shubham Gupta, Rishi Hazra, and Ambedkar Dukkipati. 2020. Networked multi-agent reinforcement learning with emergent communication. In *Proceedings of the 19th International Conference on Autonomous Agents and MultiAgent Systems, AAMAS '20*, page 1858–1860. International Foundation for Autonomous Agents and Multiagent Systems.
- Serhii Havrylov and Ivan Titov. 2017. Emergence of language with multi-agent games: Learning to communicate with sequences of symbols. In *Proceedings of the 31st International Conference on Neural Information Processing Systems, NIPS'17*, page 2146–2156.
- Rishi Hazra, Sonu Dixit, and Sayambhu Sen. 2020. Infinite use of finite means: Zero-shot generalization using compositional emergent protocols. *CoRR*, abs/2012.05011.
- I. Higgins, Loïc Matthey, A. Pal, C. Burgess, Xavier Glorot, M. Botvinick, S. Mohamed, and Alexander Lerchner. 2017. beta-vae: Learning basic visual concepts with a constrained variational framework. In *International Conference on Learning Representations (ICLR) 2017*.
- Eric Jang, Shixiang Gu, and Ben Poole. 2017. Categorical reparameterization with gumbel-softmax. In *International Conference on Learning Representations (ICLR) 2017*.
- S. Kirby. 2001. Spontaneous evolution of linguistic structure—an iterated learning model of the emergence of regularity and irregularity. *IEEE Transactions on Evolutionary Computation*, 5(2):102–110.
- Satwik Kottur, José Moura, Stefan Lee, and Dhruv Batra. 2017. Natural language does not emerge ‘naturally’ in multi-agent dialog. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 2962–2967. Association for Computational Linguistics.
- Angeliki Lazaridou, Karl Moritz Hermann, Karl Tuyls, and Stephen Clark. 2018. Emergence of linguistic communication from referential games with symbolic and pixel input. In *International Conference on Learning Representations (ICLR) 2018*.
- David K. Lewis. 1969. *Convention: A Philosophical Study*. Wiley-Blackwell.
- Ryan Lowe, Jakob Foerster, Y-Lan Boureau, Joelle Pineau, and Yann Dauphin. 2019. On the pitfalls of measuring emergent communication. In *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems, AAMAS '19*, page 693–701. International Foundation for Autonomous Agents and Multiagent Systems.
- Igor Mordatch and Pieter Abbeel. 2018. Emergence of grounded compositional language in multi-agent populations. In *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence*, pages 1495–1502.
- Martin A. Nowak and David C. Krakauer. 1999. The evolution of language. *Proceedings of the National Academy of Sciences*.
- Yi Ren, Shangmin Guo, Matthieu Labeau, Shay B. Cohen, and Simon Kirby. 2020. Compositional languages emerge in a neural iterated learning model. In *International Conference on Learning Representations (ICLR) 2020*.
- Laura Ruis, Jacob Andreas, Marco Baroni, Diane Bouchacourt, and Brenden M Lake. 2020. A benchmark for systematic generalization in grounded language understanding. In *Advances in Neural Information Processing Systems*, volume 33, pages 19861–19872. Curran Associates, Inc.
- Richard S. Sutton, Doina Precup, and Satinder Singh. 1999. Between mdps and semi-mdps: A framework for temporal abstraction in reinforcement learning. *Artificial Intelligence*.
- Olivier Tieleman, Angeliki Lazaridou, Shibli Mourad, Charles Blundell, and Doina Precup. 2019. Shaping representations through communication: community size effect in artificial learning systems. *CoRR*.
- Zhaopeng Tu, Zhengdong Lu, Yang Liu, Xiaohua Liu, and Hang Li. 2016. Modeling coverage for neural machine translation. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics*, pages 76–85. Association for Computational Linguistics.
- Ronald J Williams. 1992. Simple statistical gradient following algorithms for connectionist reinforcement learning. *Machine Learning*.