Emergence of compositional language in communication through noisy channel

Anonymous Authors[1]

Abstract

In this paper, we investigate how communication through a noisy channel can lead to the emergence of compositional language. Our novel approach is end-to-end, allows for different inductive biases on the agents' architecture, and trains without periodical resets of the networks' weights. This relaxes some of the assumptions in recently developed methods. The impact on the structure of the resulting language is shown in the context of signaling games. We also develop a new metric for measuring degree of compositionality.

1. Introduction

Communication can emerge in situations which require coordination and information sharing to achieve a joint objective. This is common in multi-agent systems with partial observation Foerster et al. (2016); Lazaridou et al. (2016); Jaques et al. (2018); Raczaszek-Leonardi et al. (2018). Communication is compositional if complex signals can be represented as a combination of their constituents. It is a feature of human languages, and it facilitates generalization, productivity, and knowledge sharing. As such, it is considered essential for general intelligence (Lake et al. (2016)).

Noise is ubiquitous in non-digital communication. For example, in biology, it has profound importance being the driving force of evolution. It is also present in human communication (Rothwell (1999)) and is perhaps equally significant, if, indeed, it promotes compositionality. Our main contribution is showing experimentally that this is the case in a simple communication model: signaling games. We argue that this effect might follow from the fact that a compositional language is (partially) robust to the mistakes caused by a noisy channel.

Signaling games (Lewis (1969); Skyrms (2010); Lazaridou et al. (2016); Fudenberg et al. (1991)) are popular model of communication. In such a game, there is a sender and a receiver. The former observes a state of the world and sends a message to the latter. Upon receiving a message, the receiver performs an action, which results in a reward or a loss. Typically, both the agents are rewarded if the receiver can identify the state of the world communicated by the sender. Compositionality rarely emerges in standard versions of this model.

While signaling games are inherently a (multi-agent) reinforcement learning setup, in the simplest case, it can be represented as a supervised learning problem. More precisely, it can be viewed as a classification task, with the underlying architecture resembling a discrete auto-encoder: the encoder acts as a sender, the bottleneck corresponds to a communication channel, and a decoder plays the role of a receiver. Without additional constraints, good accuracy can be achieved with non-compositional communication protocol - it is enough that the sender distinguishes the features of the observed state in his messages. It is likely that such a mapping will show poor generalization to novel combinations of features (Kottur et al., 2017).

Our experiments were conducted using a set of images with two features: shape and color (in a similar vein to Choi et al. (2018), Bogin et al. (2018), and Korbak et al. (2019)). We obtained some promising results, which demonstrate that introducing a noisy channel to the system, indeed creates enough tension for compositionality to emerge. Additionally, we show that the noisy channel can be modeled as a simple dense layer, which takes probabilities as input, has a fixed weight matrix of a certain form, and uses a logarithmic activation function. This constitutes a small change to the supervised learning pipeline, and is interesting in its own right (it can be viewed as a yet another regularizing layer). In particular, this does not require template transfer (Korbak et al. (2019)), periodical resets of the agents' weights or memory (Li & Bowling (2019), Kottur et al. (2017), Das et al. (2017)). The approach is also agnostic to the choice of neural network architectures for a receiver and sender. This makes it less cognitively demanding, when compared to other methods (e.g. the obverter algorithm Batali (1998), Choi et al. (2018)).

Measuring compositionality measures is another research question. To this end we introduce a new, *conflict counting*, measure and use to literature measures: topographical similarity (see e.g. Brighton & Kirby (2006); Lazaridou

[1]Anonymous Institution, Anonymous City, Anonymous Region, Anonymous Country. Correspondence to: Anonymous Author <anon.email@domain.com>.

Preliminary work. Under review by the International Conference on Machine Learning (ICML). Do not distribute.

et al. (2018) and Kriegeskorte (2008); Bouchacourt & Baroni (2018), where it is called a *representational similarity*). We provide a discussion of conflict counting, highlighting some interesting phenomenons that emerge. We also provide analysis of compositionality with respect to the noise level.

The rest of the paper is organized as follows. In the next section, we discuss related works. In Section 3, we present details of our method. Section 4 is devoted to experimental results. We discuss future work and conclude the paper in Section 5.

2. Related work

Signaling games (Lewis (1969); Skyrms (2010); Lazaridou et al. (2016); Fudenberg et al. (1991)) are popular models of communication and have traditionally been solved using either simple reinforcement learning (Skyrms, 2010) or evolutionary optimization (Cangelosi, 2001; Grouchy et al., 2016). An early work on the subject utilizing neural networks is (Rumelhart et al., 1986), while the recent results taking advantage of deep learning progress Lazaridou et al. (2018), Bouchacourt & Baroni (2018). Kottur et al. (2017) argue that a strong inductive biases are necessary for the emergence of compositionality between communicating agents. Das et al. (2017) places a pressure on agents, to use symbols consistently across varying contexts, by a frequent reset of the agent's memory. The topic of communication is inherently interesting in the context of mult-agent RL system, see Hernandez-Leal et al. (2020, Table 2) for up-to-date survey information.

A psychologically driven approach was proposed by Choi et al. (2018) and Bogin et al. (2018), who build upon the obverter algorithm (Oliphant & Batali (1997), Batali (1998)). This algorithm assumes that an agent can use its own policy as a model to predict the response of the other agent. This approach could be connected with *theory of mind* (Premack & Woodruff (1978)) and its variant, *simulation theory* (Gordon (1986), Heal (1986)). The obverter algorithm has several limitations (the agents must share an identical architecture and the task must be symmetric) and Korbak et al. (2019) used an alternative approach, based on the idea of template transfer (Barrett & Skyrms (2017)).

A different take on compositionality, creates bias towards protocols that are easy to teach to new agents (Kirby (2001); Kirby et al. (2008), Brighton (2002)). In the machine learning literature, this idea was explored by Li & Bowling (2019) and Cogswell et al. (2019). Gradually increasing task complexity that incentivizes reusing existing patterns of communication (De Beule & Bergen, 2006).

The use of noise as a regularizer is a powerful concept in deep learning, see e.g. dropout (Srivastava et al. (2014)) or semantic hashing (Salakhutdinov & Hinton (2009)). The latter was used in discrete autoencoders (Kaiser & Bengio (2018)) as a mechanism allowing backpropagation through a discrete latent. A similar idea was also applied in the context of learning to communicate (Foerster et al. (2016)). In the context of compositionality, Li & Bowling (2019) used noise to reset the weights of the agents' neural networks. In this work we use the noise explicitly to corrupt the communication channel.

3. Method

Training pipeline Consider a sender s_θ, modeled as a neural network, which observes an image x from some dataset $\mathcal{D}$. We assume that each element of $\mathcal{D}$ has K independent features $f_1, \ldots, f_K$ (here we consider $K = 2$). The sender sends a message comprising of L symbols (here we assume $L = K$). Basing on these symbols, the receiver has to guess the values for each of K features correctly. Both the features and symbols are discrete and enumerated with $A_f = \{1, \ldots, d_f\}$ and $A_s = \{1, \ldots, d_s\}$, respectively.

Formally, $s_\theta(x) = (s^i_\theta(x))^K_{i=1}$, where $s^i_\theta(x) = (s^i_{j,\theta}(x))^d_{j=1} \in \mathcal{P}(A_s)$ represents the probability distribution[1] corresponding to ith symbol. Define a function $\texttt{noise}\colon \mathcal{P}(A_s) \to \mathbb{R}^{d_s}$ as follows:

$$\texttt{noise}(x) = \log(Wx), \tag{1}$$

where $W \in \mathbb{R}^{d_s \times d_s}$ is a fixed matrix, such that $Wx > 0$ and $Wx \in \mathcal{P}(A_s)$, for any $x \in \mathcal{P}(A_s)$[2]. The second condition on W is satisfied, for instance, by a family of stochastic matrices; several examples are also given at the end of this section. Let $\widehat{s}^i_\theta(x)$ denotes the distribution of a ith symbol which passes through the noisy channel:

$$\widehat{s}^i_\theta(x) = \texttt{noise}(s^i_\theta(x)).$$

Suppose further that $g^i = (g^i_1, \ldots, g^i_{d_s})$ is a vector of i.i.d. Gumbel$(0, 1)$ random variables. Define[3]

$$\widehat{\mathfrak{m}}_i = \texttt{gumbel_softmax}(\widehat{s}^i_\theta(x); \tau, g^i) \in \mathbb{R}^{d_s}.$$

The receiver neural network is defined as $r_\psi(\mathfrak{m}) = (r^i_\psi(\mathfrak{m}))^K_{i=1}$, where $r^i_\psi(\mathfrak{m}) = (r^i_{j,\psi}(\mathfrak{m}))^d_{j=1} \in \mathcal{P}(A_f)$ represents the probability distribution on A_f, corresponding to ith feature.

[1] $\mathcal{P}(A) = \{x \in \mathbb{R}^{|A|} : x_i \geq 0, \sum_{i \in A} x_i = 1\}$.

[2] We could also define `noise` for all $x \in \mathbb{R}^m$, for some m, by first applying `softmax` to x, and then using equation (1).

[3] The functions `gumbel_sample` and `gumbel_softmax` are defined as follows:

$$\texttt{gumbel_sample}(x; \tau, g) = \arg\max_i (\log(x_i) + g_i),$$

$$\texttt{gumbel_softmax}(x; \tau, g)_i = \frac{\exp((\log(x_i) + g_i)/\tau)}{\sum_{j=1}^k \exp((\log(x_j) + g_j)/\tau)}.$$

In the Straight-Through mode (see Jang et al. (2016)), r_ψ takes $\widehat{\mathfrak{m}}$ as input half the time, and $\widetilde{\mathfrak{m}}$, the remaining half the time. Here[3]

$$\widetilde{\mathfrak{m}} = \texttt{stop_gradient}(\widehat{\mathfrak{m}} - \overline{\mathfrak{m}}) + \overline{\mathfrak{m}},$$
$$\overline{\mathfrak{m}}_i = \texttt{one_hot}(\widehat{m}_i) \in \mathbb{R}^d,$$
$$\widehat{m}_i = \texttt{gumbel_sample}(s^i_\theta(x); g^i) \in A,$$

i.e. $\widehat{m} = (\widehat{m}_i)_{i=1}^K$ is a sampled noisy message. The neural networks are trained using

$$\mathcal{L} = \mathcal{L}_{xent} + \lambda_{KL}\mathcal{L}_{KL} + \lambda_{l_2}\mathcal{L}_{l_2}.$$

The cross-entropy loss is defined as

$$\mathcal{L}_{xent} = \mathbb{E}_{(x,f_1,\ldots,f_K)\sim\mathcal{D}}\left[\sum_{i=1}^K \log r_{f_i,\psi}(\widetilde{\mathfrak{m}}(x))\right].$$

Furthermore, $\mathcal{L}_{KL} = \mathbb{E}_{x\sim\mathcal{D}}\left[\sum_{i=1}^K \mathrm{KL}(U[0,1]||s^i_\theta(x))\right]$. and $\mathcal{L}_{l_2} = ||\theta||_2 + ||\psi||_2$.

Compositionality measures In this paper we use two compositionality measures: topographical similarity (see e.g. Brighton & Kirby (2006), Lazaridou et al. (2018), Kriegeskorte (2008), Bouchacourt & Baroni (2018)), and our new conflict counting measure described now and defined in (2).

We assumed $L = K$, the number of symbols generated by the sender is equal to the number of features. In such a setting, we expect that compositional language will use one symbol for each feature. To simplify we also assume that the permutation of the position of symbols to the position of features, $\phi : \{1,\ldots,L\} \mapsto \{1,\ldots,K\}$, is fixed.

We say that the principal meaning of a symbol s at position j assuming ϕ is $\mathtt{m}(s,j;\phi) := \arg\max_f \mathtt{count}(s,j,f;\phi)$, where

$$\mathtt{count}(s,j,f;\phi) = \sum_{x\in\mathcal{D}} \mathbf{1}(s_{x,j} = s, f_{x,\phi(j)} = f),$$

and ties in $\arg\max$ are broken arbitrarily. Our metric is defined as:

$$\mathtt{conflicts} = \min_\phi \sum_{s,j} \mathtt{score}(s,j;\phi), \qquad (2)$$

where $\mathtt{score}(s,j;\phi) = \sum_{f\neq\mathtt{m}(s,j;\phi)} \mathtt{count}(s,j,f;\phi)$. Intuitively, `score` measures how many times feature assigned to a symbol s at a position j diverts from its principal meaning $\mathtt{m}(s,j;\phi)$. `conflicts` totals these errors and takes $\min$ over possible orderings ϕ.

Noise In this paragraph, we show how to define different noise channels using `noise` layer, as defined in (1). Consider a probability distribution $p \in \mathcal{P}(d_s)$. Then Wp can represent a diverse set of distributions[4].

[4]For instance if $p_i > 0$ and $g_i > 0$ for all i, then $Wp = g$, where $W_{ii} = g_i/p_i$ and $W_{ij} = 0$, for $i \neq j$.

A uniform noise applied to p manifests itself as a mixture distribution of p and $U\{1,\ldots,d_s\}$, with density equal to $(1-\varepsilon)p + \frac{\varepsilon}{d_s}$, where ε is the probability of randomly scrambling the symbol. The logits of such a distribution can be represented as $\mathtt{noise}(p)$, where W, defined in (1), can take the following form: $W_{ij} = 1-(1-1/d_s)$ and $W_{ij} = \varepsilon/d_s$ for $i \neq j$. Here, the probability of changing a symbol equals $\varepsilon(1-1/d_s)$. To make it equal to ε, we define W as

$$W_{ij} = \begin{cases} 1-\varepsilon, & i = j, \\ \frac{\varepsilon}{d_s-1}, & i \neq j. \end{cases} \qquad (3)$$

The formula given in (3) is the one we use in this paper. Similarly, a formula for a bit-flipping noise can be obtained[5].

4. Experiments

Experiments setup For our experiments, we chose a subset of a dataset used by Choi et al. (2018)[6] and Korbak et al. (2019), by selecting four shapes (box, cylinder, ellipsoid, sphere) in four colors (blue, cyan, gray, green), see Figure 1. Each image has dimensions of 128 x 128 x 3 and there is 100 images for each shape–color pair.

The sender consists of a vision and encoding module. The former consists of two convolution layers (each having 8 filter, kernel 4x4, stride 1, elu activation function). After each convolutional layer there is a 2x2 max pool layer. The output of the last max pool layer is passed through two dense layers (with 64 neurons and elu activation) and linear classifier with softmax for each symbol. The receiver's network takes one-hot encoded symbols s as input. Consequently, s and $1-s$ are passed to a dense layers. The sum of those layers is processed through the elu activation function and two dense layers (as in Kaiser & Bengio (2018)). Finally, there are two linear classifiers - one for shape and one for color. Each dense layer in the receiver has a width of 64.

Figure 1. Samples from the dataset used in this paper.

The compositionality of emerging languages was measured on the training set as our dataset was not big enough to ensure generalization (the reconstruction accuracy was low). We conjecture that this orthogonal to compositionality and plan to investigate this in future work.

Tight language: compositionality In this experiment we learned a language as tight as possible. Clearly, four symbols are enough to describe four shapes in four colors. How-

[5]$W_{ij} = Y_{i\oplus j}$, where $\oplus$ is a `xor` operation, $Y_1 = 1-\varepsilon$ and $Y_i = \varepsilon\mathbf{1}(i-1 \text{ is a power of } 2)/\lfloor\log_2(d)\rfloor$.

[6]The dataset is available at https://github.com/benbogin/obverter.

ever, we found out that it is quite challenging to learn in this setup. Consequently, we loosened the constraints on a language and allowed the sender to use five symbols. Example of a language learned by the sender is presented in Table 1. Although, the language uses five symbols, only four are used to describe shape and four to describe color. We speculate that adding one additional symbol created enough wiggle room to allow the training to escape local minima, and hence enabled the model to learn.

Table 1. Language learned by the sender in an end-to-end training. $|A_s| = 5$, $\varepsilon = 0.1$. Here $5/10$ experiments has `conflicts` $= 0$ (one of them is presented). Four symbols were used to denote color and four for shape. This language is fully compositional.

SHAPE COLOR	BOX	CYLINDER	ELLIPSOID	SPHERE
BLUE	0,1	3,1	1,1	4,1
CYAN	0,3	3,3	1,3	4,3
GRAY	0,2	3,2	1,2	4,2
GREEN	0,4	3,4	1,4	4,4

Big language: Compositionality is not what it seems In this experiment, we trained language significantly bigger than necessary with $|A_s| = 7$. In half of the experiments, the sender learned a language that is compositional regarding to our metrics (with zero or one conflicts). Language that has a highest number of conflicts is presented in Table 2. Despite poor values of the metrics (9 conflicts, topographical similarity of 0.68) it can be argued that the language is also compositional. It may be seen from the table that for box and cylinder messages starts with color followed by shape. The other two shapes have this order reversed. It's possible to achieve high accuracy (0.95) because symbols used in the first group of shapes are almost never used in the same position in the second group (only exception being 2 used for blue box and sphere as first symbol of message).

Table 2. Language learned by sender (pre-trained convnets). $|A_s| = 7$, $\varepsilon = 0.1$.

SHAPE COLOR	BOX	CYLINDER	ELLIPSOID	SPHERE
BLUE	2,3	2,5	3,1	2,1
CYAN	5,4	5,5	3,6	1,6
GRAY	6,4	6,5	3,2	1,2
GREEN	4,4	4,5	3,0	1,0

Noise vs metrics We examined the relation between noise level and compositionality of emergent language. For each selected noise level we run a batch of experiments with 10 different seeds, and we measured compositionality using topographical similarity and conflicts count. All experiments were performed with the same pre-trained convnets. The results are presented in Figure 2. The main take away is that there is no compositionality without noise in the channel. On the other hand, too large noise prevents the model from learning. As a result, there is a sweet spot for noise, located in quite a narrow window $\varepsilon \in [0.08, 0.15]$, which results in in a compositional languages.

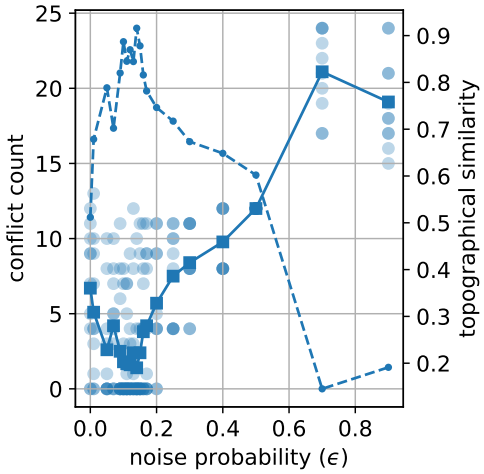

Figure 2. Compositionality metrics as functions of the channel's noise: mean conflict count measure (solid line, left axis) and mean topographical similarity (dashed line, right axis). Points represents conflict count for each experiment. For each noise probability ten experiments was run.

5. Conclusions and future work

We presented a promising method to achieve emergence of compositional language in communication between two agents. The key insight was to leverage randomness by introducing a noisy communication channel, which can be implemented as a regularization layer. We tested the approach on a version of signaling game, where a sender observes an image, communicates with the receiver, and the receiver has to guess the image's features.

This work can be extended in several directions. One of them would be to perform a fully-fledged comparison with other methods, as well as to increase the number of features and their possible range of values (e.g. by using CLEVR dataset). This could require more expressive neural network architectures, in particular sequence-to-sequence models (e.g. LSTMs or Transformers). Furthermore, the problem can be steered more towards the Reinforcement Learning framework. This includes a reinforcement learning objective and a broader, multi-agent, environment. Additionally, it is interesting how the training pipeline can be simplified (e.g. by removing the KL loss). Finally, it is interesting to study how introducing a noisy layer can improve performance of neural networks in other contexts.

References

Barrett, J. A. and Skyrms, B. Self-assembling Games. *The British Journal for the Philosophy of Science*, 68(2):

329–353, June 2017. ISSN 0007-0882, 1464-3537. doi: 10.1093/bjps/axv043. URL https://academic.oup.com/bjps/article-lookup/doi/10.1093/bjps/axv043.

Batali, J. Computational simulations of the emergence of grammar. *Approach to the Evolution of Language*, pp. 405–426, 1998.

Bogin, B., Geva, M., and Berant, J. Emergence of Communication in an Interactive World with Consistent Speakers. *arXiv:1809.00549 [cs]*, September 2018. URL http://arxiv.org/abs/1809.00549. arXiv: 1809.00549.

Bouchacourt, D. and Baroni, M. How agents see things: On visual representations in an emergent language game. *arXiv:1808.10696 [cs]*, August 2018. URL http://arxiv.org/abs/1808.10696. arXiv: 1808.10696.

Brighton, H. Compositional Syntax From Cultural Transmission. *Artificial Life*, 8(1):25–54, January 2002. ISSN 1064-5462, 1530-9185. doi: 10.1162/106454602753694756. URL http://www.mitpressjournals.org/doi/10.1162/106454602753694756.

Brighton, H. and Kirby, S. Understanding Linguistic Evolution by Visualizing the Emergence of Topographic Mappings. *Artificial Life*, 12(2):229–242, January 2006. ISSN 1064-5462, 1530-9185. doi: 10.1162/artl.2006.12.2.229. URL http://www.mitpressjournals.org/doi/10.1162/artl.2006.12.2.229.

Cangelosi, A. Evolution of communication and language using signals, symbols, and words. *IEEE Transactions on Evolutionary Computation*, 5(2):93–101, April 2001. doi: 10.1109/4235.918429.

Choi, E., Lazaridou, A., and de Freitas, N. Compositional Obverter Communication Learning From Raw Visual Input. *ICLR 2018*, April 2018. URL http://arxiv.org/abs/1804.02341. arXiv: 1804.02341.

Cogswell, M., Lu, J., Lee, S., Parikh, D., and Batra, D. Emergence of Compositional Language with Deep Generational Transmission. *ArXiv*, abs/1904.09067, 2019.

Das, A., Kottur, S., Moura, J. M. F., Lee, S., and Batra, D. Learning Cooperative Visual Dialog Agents with Deep Reinforcement Learning. *2017 IEEE International Conference on Computer Vision (ICCV), Venice, 2017*, March 2017. URL http://arxiv.org/abs/1703.06585. arXiv: 1703.06585.

De Beule, J. and Bergen, B. K. On the emergence of compositionality. In *The Evolution of Language*, pp. 35–42, Rome, Italy, March 2006. World scientific. ISBN 978-981-256-656-0 978-981-277-426-2. doi: 10.1142/9789812774262_0005. URL http://www.worldscientific.com/doi/abs/10.1142/9789812774262_0005.

Foerster, J. N., Assael, Y. M., de Freitas, N., and Whiteson, S. Learning to Communicate with Deep Multi-Agent Reinforcement Learning. *NIPS'16 Proceedings of the 30th International Conference on Neural Information Processing Systems*, May 2016. URL http://arxiv.org/abs/1605.06676. arXiv: 1605.06676.

Fudenberg, D., Tirole, J., TIROLE, J., and Press, M. *Game Theory*. Mit Press. MIT Press, 1991. ISBN 9780262061414. URL https://books.google.pl/books?id=pFPHKwXro3QC.

Gordon, R. M. Folk Psychology as Simulation. *Mind & Language*, 1(2):158–171, 1986. ISSN 1468-0017. doi: 10.1111/j.1468-0017.1986.tb00324.x.

Grouchy, P., D'Eleuterio, G. M. T., Christiansen, M. H., and Lipson, H. On The Evolutionary Origin of Symbolic Communication. *Scientific Reports*, 6:34615, October 2016. URL https://doi.org/10.1038/srep34615.

Heal, J. Replication and Functionalism. In Butterfield, J. (ed.), *Language, Mind, and Logic*, pp. 135–150. Cambridge University Press, 1986.

Hernandez-Leal, P., Kartal, B., and Taylor, M. E. A very condensed survey and critique of multiagent deep reinforcement learning. In Seghrouchni, A. E. F., Sukthankar, G., An, B., and Yorke-Smith, N. (eds.), *Proceedings of the 19th International Conference on Autonomous Agents and Multiagent Systems, AAMAS '20, Auckland, New Zealand, May 9-13, 2020*, pp. 2146–2148. International Foundation for Autonomous Agents and Multiagent Systems, 2020. URL https://dl.acm.org/doi/abs/10.5555/3398761.3399105.

Jang, E., Gu, S., and Poole, B. Categorical Reparameterization with Gumbel-Softmax. *arXiv:1611.01144 [cs, stat]*, November 2016. URL http://arxiv.org/abs/1611.01144. arXiv: 1611.01144.

Jaques, N., Lazaridou, A., Hughes, E., Gulcehre, C., Ortega, P. A., Strouse, D. J., Leibo, J. Z., and de Freitas, N. Social Influence as Intrinsic Motivation for Multi-Agent Deep Reinforcement Learning. *arXiv:1810.08647 [cs, stat]*, October 2018. URL http://arxiv.org/abs/1810.08647. arXiv: 1810.08647.

Kaiser, Ł. and Bengio, S. Discrete autoencoders for sequence models. *arXiv preprint arXiv:1801.09797*, 2018.

Kirby, S. 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, April 2001. doi: 10.1109/4235.918430.

Kirby, S., Cornish, H., and Smith, K. Cumulative cultural evolution in the laboratory: An experimental approach to the origins of structure in human language. *Proceedings of the National Academy of Sciences*, 105(31):10681–10686, August 2008. ISSN 0027-8424, 1091-6490. doi: 10.1073/pnas.0707835105. URL http://www.pnas.org/cgi/doi/10.1073/pnas.0707835105.

Korbak, T., Zubek, J., Kuciński, Ł., Miłoś, P., and Raczaszek-Leonardi, J. Developmentally motivated emergence of compositional communication via template transfer. *arXiv preprint arXiv:1910.06079*, 2019.

Kottur, S., Moura, J. M. F., Lee, S., and Batra, D. Natural Language Does Not Emerge 'Naturally' in Multi-Agent Dialog. *arXiv:1706.08502 [cs]*, June 2017. URL http://arxiv.org/abs/1706.08502. arXiv: 1706.08502.

Kriegeskorte, N. Representational similarity analysis – connecting the branches of systems neuroscience. *Frontiers in Systems Neuroscience*, 2008. ISSN 16625137. doi: 10.3389/neuro.06.004.2008. URL http://journal.frontiersin.org/article/10.3389/neuro.06.004.2008/abstract.

Lake, B. M., Ullman, T. D., Tenenbaum, J. B., and Gershman, S. J. Building Machines That Learn and Think Like People. *arXiv:1604.00289 [cs, stat]*, April 2016. URL http://arxiv.org/abs/1604.00289. arXiv: 1604.00289.

Lazaridou, A., Peysakhovich, A., and Baroni, M. Multi-Agent Cooperation and the Emergence of (Natural) Language. *arXiv:1612.07182 [cs]*, December 2016. URL http://arxiv.org/abs/1612.07182. arXiv: 1612.07182.

Lazaridou, A., Hermann, K. M., Tuyls, K., and Clark, S. Emergence of Linguistic Communication from Referential Games with Symbolic and Pixel Input. *arXiv:1804.03984 [cs]*, April 2018. URL http://arxiv.org/abs/1804.03984. arXiv: 1804.03984.

Lewis, D. K. *Convention: a philosophical study*. Blackwell, Oxford, nachdr. edition, 1969. ISBN 978-0-631-23257-5 978-0-631-23256-8. OCLC: 837747718.

Li, F. and Bowling, M. Ease-of-teaching and language structure from emergent communication. In *Advances in Neural Information Processing Systems*, pp. 15825–15835, 2019.

Oliphant, M. and Batali, J. Learning and the Emergence of Coordinated Communication. *Center for Research on Language Newsletter*, 11, 1997.

Premack, D. and Woodruff, G. Does the chimpanzee have a theory of mind? *Behavioral and Brain Sciences*, 1(4): 515–526, December 1978. ISSN 1469-1825, 0140-525X. doi: 10.1017/S0140525X00076512.

Rothwell, J. *In the Company of Others: An Introduction to Communication*. McGraw-Hill Higher Education, 1999. ISBN 9781559347389. URL https://books.google.pl/books?id=VtPZAAAAMAAJ.

Rumelhart, D. E., Hinton, G. E., and Williams, R. J. Learning representations by back-propagating errors. *Nature*, 323(6088):533–536, October 1986. ISSN 0028-0836, 1476-4687. doi: 10.1038/323533a0. URL http://www.nature.com/articles/323533a0.

Raczaszek-Leonardi, J., Nomikou, I., Rohlfing, K. J., and Deacon, T. W. Language Development From an Ecological Perspective: Ecologically Valid Ways to Abstract Symbols. *Ecological Psychology*, 30 (1):39–73, January 2018. ISSN 1040-7413, 1532-6969. doi: 10.1080/10407413.2017.1410387. URL https://www.tandfonline.com/doi/full/10.1080/10407413.2017.1410387.

Salakhutdinov, R. and Hinton, G. Semantic hashing. *International Journal of Approximate Reasoning*, 50(7): 969–978, 2009.

Skyrms, B. *Signals: evolution, learning, & information*. Oxford University Press, Oxford ; New York, 2010. ISBN 978-0-19-958082-8 978-0-19-958294-5. OCLC: ocn477256653.

Srivastava, N., Hinton, G., Krizhevsky, A., Sutskever, I., and Salakhutdinov, R. Dropout: a simple way to prevent neural networks from overfitting. *The journal of machine learning research*, 15(1):1929–1958, 2014.