

On the Spontaneous Emergence of Discrete and Compositional Signals

Nur Geffen Lan¹ Emmanuel Chemla² Shane Steinert-Threlkeld³

¹ Computational Linguistics Lab, Tel Aviv University

² EHESS, PSL University, CNRS, Ecole Normale Supérieure

³ Department of Linguistics, University of Washington

nurlan@mail.tau.ac.il chemla@ens.fr shanest@uw.edu

Abstract

We propose a general framework to study language emergence through signaling games with neural agents. Using a continuous latent space, we are able to (i) train using backpropagation, (ii) show that discrete messages nonetheless naturally emerge. We explore whether categorical perception effects follow and show that the messages are not compositional.

1 Introduction

In a signaling game, artificial agents learn to communicate to achieve a common goal: a sender sees some piece of information and produces a message, which is then sent to a receiver that must take some action (Lewis, 1969; Skyrms, 2010). If the action is coherent with the sender’s initial piece of information, the choice of the message and its interpretation is reinforced. For instance, in a referential game, sender and receiver see a set of objects, and the sender knows which of these the receiver must pick; the sender then sends a message to the receiver, who must interpret it to pick up the right object (Lazaridou et al., 2017, 2018; Havrylov and Titov, 2017; Chaabouni et al., 2019).

This setting has been used to study the factors influencing the emergence of various fundamental properties of natural language, such as *compositionality* (Kirby et al., 2015; Franke, 2016; Steinert-Threlkeld, 2016; Mordatch and Abbeel, 2018; Lazaridou et al., 2018; Choi et al., 2018). In this paper, we add focus on two other so-called ‘design features’ of natural language (Hockett, 1960): *discreteness* (i.e. words form clusters in acoustic space), and *displacement* (i.e. efficient communication can occur about objects and facts beyond the immediate context of the conversation).

From an implementation point of view, we follow the recent literature which has shown that a sig-

naling game is essentially an autoencoder setting, with the encoder playing the role of the sender, and the decoder the role of the receiver (see Fig. 1). In this literature, however, the discreteness of the communication protocol is assumed, since the networks then traditionally use a (normally sequential and) discrete latent space (Havrylov and Titov, 2017; Chaabouni et al., 2019; Kharitonov et al., 2019).

Our main contribution is a generalization of the current implementation of signaling games as autoencoders. Our implementation covers a broader variety of signaling games, and it crucially incorporates the possibility of displacement and makes no *a priori* assumption of discreteness. Our main result is that under appropriate conditions, discreteness emerges spontaneously: if the latent space is thought about as a continuous acoustic space, then trained messages form coherent clusters, just like regular words do. We also show that the messages are not compositional.

In addition to contributing to our understanding of the emergence of communication protocols with features like natural language, our results have technical significance: by using a continuous communication protocol, with discreteness spontaneously emerging, we can train end-to-end using standard backpropagation, instead of reinforcement learning algorithms like REINFORCE and its refinements (Williams, 1992; Schulman et al., 2015; Mnih et al., 2016), which are difficult to use in practice.

2 Related Work

A related line of work attempts to avoid the difficulties of reinforcement learning—used when there are stochastic nodes in a computation graph—by reparameterization and/or non-stochastic estimators (Bengio et al., 2013; Schulman et al., 2015). In the emergent communication case, where the stochastic nodes are discrete (e.g. sampling a

message from a sender distribution), the Gumbel-Softmax estimator has become increasingly popular (Jang et al., 2017; Maddison et al., 2017).

That work enables standard backpropagation to be used for training by optimizing approximations to the true reinforcement learning signal. By contrast, we do not approximate the discrete RL learning signal, but rather ask under what conditions discreteness will emerge.

Several earlier papers explore similar topics in the emergence of discrete symbols. Nowak et al. (1999) show that the division of the acoustic space is an emergent property of language use under noise. It assumes that speakers have a fixed language and asks which such ones are stable. In our setting, the language itself is changing as the result of reinforcement from communication and transmission itself is not noisy.

De Boer (2000) simulates the emergence of vowel systems in artificial agents modeled after phonetic production and perception in humans, resulting in a self-discretizing acoustic space and a vowel system that resembles human ones. This makes the agents much closer to what we know about humans, but also limits its scope. Results about emergent communication can tell us both about the emergence of human language, but also about communication protocols in general, that may be used by very different agents, e.g. autonomous ones, or animals (Steinert-Threlkeld et al., 2020).

3 Function Games

We here introduce a general communication game setting, which we call Function Games. Our games contain three basic components: (i) a set of contexts C , (ii) a set of actions A , (iii) a family of functions F , from contexts to actions. One play of a Function Game game runs as follows:

1. Nature chooses $f \in F$ and a context $c \in C$.
2. Sender sees the context c and f .
3. Sender sends a message m to Receiver.
4. Receiver sees a possibly different context c' and the message m and chooses an action a' .
5. Both are 'rewarded' iff $a' = f(c')$.

Abstractly, the function f represents some piece of knowledge available primarily for Sender, and which determines what action is appropriate in any given context. Two concrete interpretations will help illustrate the variety of communication protocols and goals that this framework encompasses.

Generalized referential games. A reference game is one in which Sender tries to get Receiver to pick the correct object out of a given set (Skyrms, 2010; Lazaridou et al., 2017, 2018; Havrylov and Titov, 2017; Chaabouni et al., 2019). Here, contexts are sets of objects (i.e. an $m \times n$ matrix, with m objects represented by n features). Normally (though we will drop this assumption later), $c' = \text{shuffled}(c)$: Sender and Receiver see the same objects, but in a different arrangement. Actions are the objects, and the functions $f \in F$ are *choice functions*: $f(c) \in c$ for every context c .

Belief update games. We will mostly focus on the previous interpretation, but illustrate the generality of the setting with another interpretation here. Contexts can represent the (possibly different) belief states of the agents. 'Actions' can represent updated belief states ($A = C$), the different functions in F then representing how to update an agent's beliefs in the light of learning a particular piece of information (passed directly to Sender, and only through the message to Receiver).

4 Experiment

Because we are interested in the simultaneous emergence both of discrete and of compositional signals, we use a Function Game called the Extremity Game designed to incentivize and test rich compositionality (Steinert-Threlkeld, 2018, 2020). In this game, one may think of the n dimensions of the objects as gradable properties, e.g. size and darkness, so that a 2D object is determined by a given size and shade of gray. For the functions, we set $F = \{\arg \min_i, \arg \max_i : 0 \leq i < n\}$. An emerging language may contain compositional messages like 'MOST + BIG', 'LEAST + DARK'.

4.1 Model

Our model (Figure 1) resembles an encoder-decoder architecture, with Sender encoding the context/target pair into a message, and Receiver decoding the message (together with its context c') into an action. Both the encoder and decoder are multi-layer perceptrons with two hidden layers of 64 ReLU units (Nair and Hinton, 2010; Glorot et al., 2011). A smaller, intermediate layer without an activation function bridges the encoder and decoder and represents the transformation of the input information to messages.

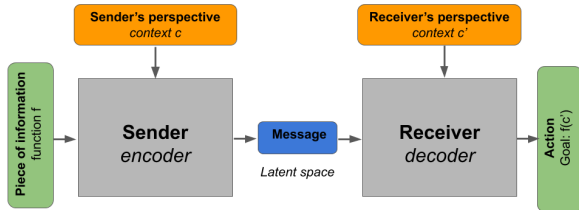


Figure 1: Our model architecture, mixing terminology from the autoencoder and signaling game traditions.

4.2 Game Parameters

We manipulate the following parameters:

- **Context identity.** In the *shared* setting, Receiver sees a shuffled version of Sender’s context ($c' = \text{shuffled}(c)$). In the *non-shared* setting, Receiver’s context c' is entirely distinct from Sender’s. This forces displacement and may incentivize compositional messages, since Sender cannot rely on the raw properties of the target object in communication.
- **Context strictness.** In *strict* contexts, there is a one-to-one (and onto) correspondence between F and A (as in the original Extremity Game from Steinert-Threlkeld, 2018, 2020). In *non-strict* contexts, an object may be the arg max or arg min of several dimensions, or of no dimension.

In all experiments, the latent space (message) dimension is always 2, and objects have 5 dimensions. Strict contexts therefore contain 10 objects, while non-strict contexts contain 5, 10, or 15 objects.

4.3 Training Details

We use the Adam optimizer (Kingma and Ba, 2015) with learning rate 0.001, $\beta_1 = 0.9$, and $\beta_2 = 0.999$. The model is trained for 5,000 steps by feeding the network mini-batches of 64 contexts concatenated with one-hot function selectors. The network’s loss is taken as the MSE between the target object $f(c')$ and the object generated by the Receiver. For each setting of the above parameters, we run 20 trials with different random seeds.¹

5 Results

5.1 Communicative success

We measure the communicative success of the network by calculating the accuracy of recovering the correct object from c' . Receiver’s prediction is considered correct if its output is closer to $f(c')$ than

¹The project’s code for extension and reproduction is available at <https://github.com/0xnurl/signaling-auto-encoder>.

	Shared	Non-shared
Strict		
10 objects	63.78% \pm 1.63	60.22% \pm 1.56
Non-strict		
5 objects	49.37% \pm 1.67	43.55% \pm 1.69
10 objects	33.06% \pm 1.47	31.89% \pm 1.63
15 objects	27.58% \pm 1.30	27.95% \pm 1.24

Table 1: Communicative success, as measured by object recovery accuracy.




Figure 2: Sampled messages for contexts of 10 objects of size 5 for (a) an untrained and (b) a trained network. Colors represent the $f_i \in F$ input part of the Sender.

to all other objects in c' . Accuracy of the different settings is reported in Table 1. While the network handles displacement well (*non-shared contexts*), the model struggles with non-strict contexts. Note that although accuracy is not 100%, it is still well above chance, since e.g. for a context of 10 objects random guessing yields an expected accuracy of 10% (which we observe in our model before training).

5.2 Discrete signals

Figure 2 depicts message vectors sampled from the latent space layer, before and after training. It is apparent that discrete messages emerge from the imposed learning regime. We measure cluster tendency more quantitatively through two measures, one considering Sender’s *production*, and the other Receiver’s *perception*.

First, we sample 100 contexts, and collect the output of the trained encoder for each of these contexts combined with each possible function f . We apply an unsupervised clustering algorithm to this set of produced messages (DBSCAN, Ester et al., 1996, with $\epsilon = 0.5$). A label is assigned to each cluster using the ground truth: the label of a cluster is the function f that was most often at the source of a point in this cluster. This allows us to compute F1-scores, which are reported in Table 2. The model reached near-optimal clusteriza-

	Shared	Non-shared
Strict		
10 objects	1.00 ± 0.00	0.90 ± 0.09
Non-strict		
5 objects	0.99 ± 0.02	0.54 ± 0.15
10 objects	1.00 ± 0.00	0.99 ± 0.01
15 objects	1.00 ± 0.00	1.00 ± 0.00

Table 2: Discreteness in production, as measured by F1 scores for automatically clusterized messages.

	Shared	Non-shared
Strict		
10 objects	63.39% ± 1.45	55.37% ± 3.43
Non-strict		
5 objects	46.94% ± 1.70	29.40% ± 5.59
10 objects	32.63% ± 1.43	31.51% ± 1.62
15 objects	28.24% ± 1.11	27.94% ± 1.20

Table 3: Discreteness in perception, as measured by object recovery accuracy from artificial messages.

tion measures in 7 out of 8 parameter settings, with the Non-strict, Non-shared context with 5 objects being the exception.

The second approach is akin to studying perception. Given the clusterization of the message space, we sample new messages from each cluster, and test Receiver’s perception of these ‘artificial’ messages, which have never been produced by Sender. To sample artificial messages, we take the average of 10 messages from a (now labelled) cluster. These artificial messages are fed to Receiver for 100 different contexts. The output object accuracy for these artificial messages is shown in Table 3. The model achieves recovery accuracy similar to when interpreting actual messages.

In sum, we can identify discrete, abstract regions of the latent space corresponding to different functions in the input, just like words form clusters in acoustic space.

5.3 Compositionality

Our agents are capable of communicating in abstract situations, namely some in which their contexts are different in the first place. This generalizability suggests that the messages may be ‘compositional’. We here probe for a candidate compositional structure to the latent space, by asking how the messages relate to the structure of the family of

functions F .

First, the pioneering Mikolov et al., 2013 looks for compositionality at the level of word embeddings (WE) through addition, most classically asking whether $WE(\text{queen})=WE(\text{king})-WE(\text{man})+WE(\text{woman})$. In the current Game, we can ask whether the messages are related as follows, for any dimensions i and j : $M(c, \arg \max_i)=M(c, \arg \max_j)-M(c, \arg \min_j)+M(c, \arg \min_i)$. For each such pair of object dimensions we calculate the right-hand side of the equation above for 100 contexts, feed it to Receiver, compare Receiver’s output to the output that would have been obtained if $M(c, \arg \max_i)$ (the left-hand side) had been sent in the first place. This leads to important degradation of average communicative success: a drop of at least 24 percentage points across parameter combinations, to around chance level. Full results are in the left column of Table 4.

Second, we note as others that the composition-as-addition assumption is disputable, both in general and in the original application case (Linzen, 2016; Chen et al., 2017). To abstract away from this issue, we train a ‘composition network’ (an MLP with 2 hidden layers of 64 ReLU units) on the task of predicting $M(c, \arg \max_i)$ from $M(c, \arg \max_j)$, $M(c, \arg \min_j)$ and $M(c, \arg \min_i)$, therefore letting it discover any function for mixing values, and not involving addition *a priori*. We leave out one dimension i_0 from training, and feed Receiver with the message predicted by the ‘composition network’ from $M(c, \arg \max_j)$, $M(c, \arg \min_j)$ and $M(c, \arg \min_{i_0})$. If the language was compositional, this predicted message should behave like $M(c, \arg \max_{i_0})$, but we found that, as in the case of addition, the average communication accuracy for all taken-out parameters dropped dramatically (again, at least 24 percentage points drop). Full results are in the right column of Table 4.

5.4 Categorical perception

Above we essentially propose an analysis of discreteness both in production and perception. This can lead to more psycholinguistic-like queries about these emergent languages. For instance, one may ask whether classical ‘Categorical Perception’ (CP) effects obtain, whereby two messages at a short distance in the latent space may be discriminated easily if (and only if) they are on two sides of a categorical boundary for interpretation purposes

	Compositionality by Addition		Composition Network	
	Shared	Non-shared	Shared	Non-shared
Strict				
10 objects	7.82% ± 2.40	11.94% ± 2.13	13.70% ± 6.85	10.18% ± 6.15
Non-strict				
5 objects	16.86% ± 3.23	17.14% ± 3.54	15.10% ± 2.05	14.35% ± 2.74
10 objects	5.82% ± 2.37	6.46% ± 1.79	5.00% ± 2.62	5.92% ± 2.12
15 objects	3.72% ± 1.42	4.00% ± 1.54	1.59% ± 1.31	2.48% ± 1.05

Table 4: Communicative success using messages ‘inferred’ by assuming a systemic relation within $\arg \min_i / \arg \max_i$ message pairs. The ‘compositionality by addition’ method assumes that $M(c, \arg \max_i) = M(c, \arg \max_j) - M(c, \arg \min_j) + M(c, \arg \min_i)$. The ‘compositional network’ is an MLP trained to predict $M(c, \arg \max_i)$ from the other three messages. Table values are object recovery accuracies averaged for all i .

(see Liberman et al., 1957, and Damper and Harnad, 2000 for early discussions in the context of neural architectures).

As an initial foray, we can investigate the sharpness of the boundaries of our discrete messages (i.e. distribution in latent space). For representation purposes, we sample pairs of messages, call them M_{-1} and M_{+1} generated by Sender for two choice functions F_{-1} and F_{+1} . We explore a continuous spectrum of messages in the dimension connecting these two messages ($M_t = \frac{(1-t)M_{-1} + (1+t)M_{+1}}{2}$, continuously shifting from M_{-1} to M_{+1} as the continuous variable t moves from -1 to $+1$). The messages M_t are fed to Receiver together with contexts C' , and for each function F_{-1} and F_{+1} in turn, we calculate object recovery accuracy. This is plotted in Figure 3 for an Extremity Game model trained in a strict, non-shared context setting with object size 5. The model shows that clusters have relatively sharp boundaries, especially in the direction of a message belonging to another cluster (the area where x is between -1 and $+1$ in Fig. 3).




Figure 3: Categorical perception effect, demonstrated by accuracy of object recovery using messages shifted between two ‘meanings’.

We can thus identify a boundary around a cluster, and its width, providing the necessary setup to investigate CP effects: whether pairs of messages crossing such a boundary behave differently (e.g., are easier to discriminate) than a pair of equally distant messages both on one side of this boundary.

6 Conclusion

We propose a general signaling game framework in which fewer *a priori* assumptions are imposed on the conversational situations. We use both production and perception analyses, and find that under appropriate conditions, which are met by most studies involving neural signaling games, messages become discrete without the analyst having to force this property into the language (and having to deal with non-differentiability issues). We find no evidence of compositional structure using vector analogies and a generalization thereof but do find sharp boundaries between the discrete message clusters. Future work will explore other measures and alternative game settings for the emergence of compositionality, as well as more subtle psychological effects (Categorical Perception) of continuous biological systems exhibiting discrete structure, like the auditory system.

Acknowledgments

We acknowledge the funding support from ANR-17-EURE-0017, and greatly thank Marco Baroni, Diane Bouchacourt, Rahma Chaabouni, Emmanuel Dupoux, Roni Katzir, Philippe Schlenker, Benjamin Spector, Jakub Szymanik, and three ACL reviewers.

References

- Yoshua Bengio, Nicholas Léonard, and Aaron Courville. 2013. [Estimating or Propagating Gradients Through Stochastic Neurons for Conditional Computation](#).
- Rahma Chaabouni, Eugene Kharitonov, Emmanuel Dupoux, and Marco Baroni. 2019. [Anti-efficient encoding in emergent communication](#). In *Proceedings of the 33rd Conference on Neural Information Processing Systems (NeurIPS 2019)*.
- Dawn Chen, Joshua C. Peterson, and Thomas L. Griffiths. 2017. [Evaluating vector-space models of analogy](#). In *Proceedings of the 39th Annual Conference of the Cognitive Science Society*.
- Edward Choi, Angeliki Lazaridou, and Nando de Freitas. 2018. [Compositional Obverter Communication Learning from Raw Visual Input](#). In *International Conference of Learning Representations (ICLR 2018)*, pages 1–18.
- R.I. Dampier and S.R. Harnad. 2000. [Neural network models of categorical perception](#).
- Bart De Boer. 2000. [Self-organization in vowel systems](#). *Journal of Phonetics*, 28(4):441–465.
- Martin Ester, Hans-Peter Kriegel, Jörg Sander, and Xiaowei Xu. 1996. A density-based algorithm for discovering clusters in large spatial databases with noise. In *KDD*, volume 96, pages 226–231.
- Michael Franke. 2016. [The Evolution of Compositionality in Signaling Games](#). *Journal of Logic, Language and Information*.
- Xavier Glorot, Antoine Bordes, and Yoshua Bengio. 2011. Deep Sparse Rectifier Neural Networks. In *14th International Conference on Artificial Intelligence and Statistics (AISTATS)*, pages 315–323.
- 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 Conference on Neural Information Processing Systems (NeurIPS 2017)*.
- Charles F Hockett. 1960. The Origin of Speech. *Scientific American*, 203:88–111.
- Eric Jang, Shixiang Gu, and Ben Poole. 2017. [Categorical Reparameterization with Gumbel-Softmax](#). In *International Conference of Learning Representations (ICLR)*.
- Eugene Kharitonov, Rahma Chaabouni, Diane Bouchacourt, and Marco Baroni. 2019. [EGG: a toolkit for research on Emergence of lanGUage in Games](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP): System Demonstrations*, pages 55–60, Stroudsburg, PA, USA. Association for Computational Linguistics.
- Diederik P. Kingma and Jimmy Ba. 2015. [Adam: A Method for Stochastic Optimization](#). In *International Conference of Learning Representations (ICLR)*.
- Simon Kirby, Monica Tamariz, Hannah Cornish, and Kenny Smith. 2015. [Compression and communication in the cultural evolution of linguistic structure](#). *Cognition*, 141:87–102.
- 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 of Learning Representations (ICLR 2018)*.
- Angeliki Lazaridou, Alexander Peysakhovich, and Marco Baroni. 2017. [Multi-Agent Cooperation and the Emergence of \(Natural\) Language](#). In *International Conference of Learning Representations (ICLR2017)*.
- David Lewis. 1969. *Convention*. Blackwell.
- Alvin M Liberman, Katherine Safford Harris, Howard S Hoffman, and Belver C Griffith. 1957. The discrimination of speech sounds within and across phoneme boundaries. *Journal of Experimental Psychology*, 54(5):358.
- Tal Linzen. 2016. [Issues in evaluating semantic spaces using word analogies](#). In *Proceedings of the 1st Workshop on Evaluating Vector-Space Representations for NLP*, pages 13–18, Berlin, Germany. Association for Computational Linguistics.
- Chris J Maddison, Andriy Mnih, Yee Whye Teh, United Kingdom, and United Kingdom. 2017. [The Concrete Distribution: A Continuous Relaxation of Discrete Random Variables](#). In *International Conference of Learning Representations (ICLR)*.
- Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013. [Efficient Estimation of Word Representations in Vector Space](#). *arXiv:1301.3781 [cs]*. ArXiv: 1301.3781.
- Volodymyr Mnih, Adrià Puigdomènech Badia, Mehdi Mirza, Tim Harley, Timothy P Lillicrap, David Silver, and Koray Kavukcuoglu. 2016. [Asynchronous Methods for Deep Reinforcement Learning](#). In *International Conference on Machine Learning (ICML)*.
- Igor Mordatch and Pieter Abbeel. 2018. [Emergence of Grounded Compositional Language in Multi-Agent Populations](#). In *The Thirty-Second AAAI Conference on Artificial Intelligence (AAAI 2018)*.
- Vinod Nair and Geoffrey E Hinton. 2010. [Rectified Linear Units Improve Restricted Boltzmann Machines](#). In *Proceedings of the 27th International Conference on Machine Learning (ICML)*.

- Martin A. Nowak, David C. Krakauer, and Andreas Dress. 1999. [An error limit for the evolution of language](#). *Proceedings of the Royal Society B: Biological Sciences*, 266(1433):2131–2136.
- John Schulman, Nicolas Heess, Theophane Weber, and Pieter Abbeel. 2015. [Gradient Estimation Using Stochastic Computation Graphs](#). In *Advances in Neural Information Processing Systems 28 (NIPS 2015)*.
- Brian Skyrms. 2010. *Signals: Evolution, Learning, and Information*. Oxford University Press.
- Shane Steinert-Threlkeld. 2016. [Compositional Signaling in a Complex World](#). *Journal of Logic, Language and Information*, 25(3):379–397.
- Shane Steinert-Threlkeld. 2018. [Paying Attention to Function Words](#). In *Emergent Communication Workshop @ NeurIPS 2018*.
- Shane Steinert-Threlkeld. 2020. [Towards the Emergence of Non-trivial Compositionality](#). *Philosophy of Science*.
- Shane Steinert-Threlkeld, Philippe Schlenker, and Emmanuel Chemla. 2020. [Referential and General Calls in Primate Semantics](#).
- Ronald J Williams. 1992. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine Learning*, 8(3-4):229–256.