

A Signal-Centric Perspective on the Evolution of Symbolic Communication

Quintino Francesco Lotito
University of Trento
Trento, Italy
quintino.lotito@studenti.unitn.it

Leonardo Lucio Custode
University of Trento
Trento, Italy
leonardo.custode@unitn.it

Giovanni Iacca
University of Trento
Trento, Italy
giovanni.iacca@unitn.it

ABSTRACT

The evolution of symbolic communication is a longstanding open research question in biology. While some theories suggest that it originated from sub-symbolic communication (i.e., iconic or indexical), little experimental evidence exists on how organisms can actually evolve to define a shared set of symbols with unique interpretable meaning, thus being capable of encoding and decoding discrete information. Here, we use a simple synthetic model composed of sender and receiver agents controlled by Continuous-Time Recurrent Neural Networks, which are optimized by means of neuro-evolution. We characterize signal decoding as either regression or classification, with limited and unlimited signal amplitude. First, we show how this choice affects the complexity of the evolutionary search, and leads to different levels of generalization. We then assess the effect of noise, and test the evolved signaling system in a referential game. In various settings, we observe agents evolving to share a dictionary of symbols, with each symbol spontaneously associated to a 1-D unique signal. Finally, we analyze the constellation of signals associated to the evolved signaling systems and note that in most cases these resemble a Pulse Amplitude Modulation system.

CCS CONCEPTS

• **Computing methodologies** → **Artificial life; Multi-agent systems; Cooperation and coordination; Neural networks**; • **Theory of computation** → *Multi-agent learning*.

KEYWORDS

Symbolic Communication, Neuro-evolution, NEAT, Continuous-Time Recurrent Neural Networks, Agent-Based Model.

ACM Reference Format:

Quintino Francesco Lotito, Leonardo Lucio Custode, and Giovanni Iacca. 2021. A Signal-Centric Perspective on the Evolution of Symbolic Communication. In *2021 Genetic and Evolutionary Computation Conference (GECCO '21)*, July 10–14, 2021, Lille, France. ACM, New York, NY, USA, 9 pages. <https://doi.org/10.1145/3449639.3459273>

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

GECCO '21, July 10–14, 2021, Lille, France

© 2021 Association for Computing Machinery.

ACM ISBN 978-1-4503-8350-9/21/07...\$15.00

<https://doi.org/10.1145/3449639.3459273>

1 INTRODUCTION

The origin of language is considered as the last known major transition in evolution [36]. Language is in turn perceived as the culmination of an evolutionary history that led from the iconic/indexical forms of sub-symbolic communication to a symbolic one [13, 34]. Unravelling this history by gathering biological evidence is, however, extremely hard. Yet, studying the evolution of symbolic communication is a fundamental step toward understanding how language originated. One possible way to gain insights into this evolutionary process consists in analyzing how primordial forms of symbolic communication can evolve *in silico*, in agent-based models. Previous literature has shown that, under given conditions, it is possible to evolve various forms of symbolic communication in Artificial Life simulations. On the other hand, only few works have analyzed the specific evolved communication systems, especially w.r.t. the emerging signals and their association to symbols.

In this work, we present a synthetic model composed of sender and receiver agents controlled by Continuous-Time Recurrent Neural Networks (CTRNNs) [5], which are able to generate and decode continuous-time signals. The CTRNNs are optimized by means of the Neuro-Evolution of Augmenting Topologies (NEAT) algorithm [39]. First, we show that the agents are able to evolve a signaling system in several settings, even in the presence of noise. Then, we analyze the evolved signaling systems by looking at the properties of the signals evolved under the different conditions. This analysis shows, quite interestingly, that the evolved systems are in most cases similar to well-known signal modulation techniques such as Pulse Amplitude Modulation (PAM) [4]. Compared to the previous literature, our work is distinct in these two main aspects:

- We characterize signal decoding as either regression or classification, with limited and unlimited signal amplitude, and show how this choice affects the complexity of the evolutionary search, and leads to different levels of generalization.
- We perform a thorough analysis of the evolved signaling system in terms of *constellations of signals* [17] and clustering [1], and show how different classes of signals emerge from evolution.

The rest of the paper is organized as follows. In the next section, we provide an overview of the related work. Then, in Section 3 we introduce the background on CTRNNs and NEAT. In Section 4, we describe the experimental setup. Then, we present the numerical results in Section 5, and we analyze the evolved signaling systems in Section 6. Finally, we draw the conclusions of this study in Section 7.

2 RELATED WORK

Over the years, the evolution of communication (both sub-symbolic and symbolic) has attracted a great research interest from theoretical biologists, mathematicians, linguists, and computer scientists.

Several works analyze the change in communication performance due to the available *implicit* vs. *explicit* communication channels and their properties [8, 12, 18, 40]. One special case of communication based on an implicit channel is *referential communication*, which occurs when agents exchange information by observing the *behavior* of other agents. This form of communication has been largely studied e.g. in [1, 2, 11, 14, 30, 35, 43].

Social factors, including sexual reproduction [19], have been shown to play an important role in the evolution of communication. In [37], the author studies the emergence of agreement between symbols and meanings in a “culture” (i.e., a population of agents that somehow “agree” on a given representation of the reality). In [23], the authors study the effects of social influence and transmission bottlenecks on communication.

Other works study various specific properties of emerging communication. For instance, Tuci [42] investigate on the “sound-to-silence” ratio of an evolved communication system. Lacroix [26] studies the information transfer in an “impoverished signaling game”. Froese et al. [16], study how the behavior of the communicating agents changes depending on the scenario. In [25, 27], Lacroix analyzes the agents’ salience, i.e., their expectations w.r.t. the others’ behavior. Kadish et al. [22] study the effect of noise on the communication between agents evolved in noise-free contexts. Nowak et al. [32] study the effect of the number of symbols on the error rate of the language. Talamini et al. [41] study how the effectiveness of the communication varies in different settings.

An active research direction in this field focuses on the evolution of *compositional* languages. A language is defined to be compositional when multiple symbols can be combined to define a meaning. For instance in [31] the authors review the literature related to the emergence of communication and mention the emergence of compositional communication as one of the main future challenges.

Another important aspect of the evolution of communication is the possibility to “ground” a language, i.e., create a correspondence between an evolved language and human languages. In [28], the authors train a multi-agent system, composed of a sender and a receiver, to play a referential game. Moreover, they provide a method to ground the language used by NNs in human languages. Havrylov et al. [21], inspired by [28], train a communication system based on an encoder-decoder structure, to play a referential game. In [10], the authors show that when training a speaker-listener couple of agents, they do not naturally adopt a form of efficient encoding. In [20], the author describes the symbol grounding problem and proposes a scheme in which he hypothesizes that symbol grounding happens in a bottom-up manner, i.e., starting from perceptions and ending in symbols which describe those perceptions.

While evolutionary approaches are one way to evolve communication systems, they are not the only one. In [3], the authors study the emergence of communication in a multi-agent system by using simple reinforcement learning in three different versions of a composition game (a kind of referential game). Foerster et al. [15] propose two approaches based on Deep Q Networks to make the communication emerge in sender-receiver settings: Reinforced Inter-Agent Learning (RIAL) and Differentiable Inter-Agent Learning (DIAL). In RIAL, sender and receiver learn to communicate by learning independently; in DIAL, the agents learn to communicate by propagating the gradients from receiver to sender.

Finally, the evolution of communication has been object of works tailored to its applications. In [7], the authors present their perspective on the combination of swarm robotics and emergence of language. In [24], the authors implement a framework to ease the training of multi-agent systems that must learn to communicate.

3 BACKGROUND

In the following, we briefly introduce the two main computational tools used in this work, namely CTRNNs and NEAT.

3.1 CTRNNs

CTRNNs [5] are computational NNs in which the internal state of each neuron i is time-dependent and it is determined by the following differential equation:

$$\tau_i \frac{\partial y_i}{\partial t} = -y_i + f_i \left(\beta_i + \sum_{j \in A_i} w_{ij} y_j \right) \quad (1)$$

where τ_i , y_i , f_i and β_i indicate respectively the time constant, the potential, the activation function and the bias of neuron i , A_i is the set of indices of neurons that provide an input to neuron i , and w_{ij} is the weight of the connection from neuron j to neuron i .

The time evolution of a CTRNN can be computed using the forward Euler method:

$$y_i(t + \Delta t) = y_i(t) + \Delta t \frac{\partial y_i}{\partial t} \quad (2)$$

One interesting property of CTRNNs is that they inherently allow recurrence and self-loops. Because of this, complex internal patterns can emerge. The resulting temporal dynamics makes CTRNNs particularly suited to model systems that are capable to produce an output signal [6] or to receive an input signal.

3.2 NEAT

Neuro-evolution, that is the application of EAs to optimize NNs, is a growing field in computational intelligence [38]. While seminal works in this direction were oriented toward the optimization of NNs with fixed topology, in the last two decades methods have been proposed to evolve the topology along with the weights.

One of the main algorithms falling into this category is NEAT [39]. NEAT combines tuning the network weights with *complexification*, meaning that the algorithm also evolves the network topology to fit the complexity of a given problem. Usually, NEAT starts with a population made of the simplest possible NNs, and changes in the network topology are introduced by *structural mutations* that can add or remove edges and nodes. This process allows NEAT to select, through fitness evaluations, the right topology with an appropriate level of complexity with respect to the problem to solve. On a very high level, to allow structural mutations, each gene in each genome refers to an edge in the NN, and has information about the two nodes being connected and the weight of that connection.

NEAT speciates the population based on genomes’ similarity and historical ancestors (*historical markings*). In this way, individuals are constrained to compete primarily within their own niches (*explicit fitness sharing*), and then with the whole population. This way of proceeding ensures topological innovations to be protected enough and have time to optimize their structure before having to

compete with other species in the population. For a more in-depth description of NEAT, we refer to the reader to [39].

4 EXPERIMENTAL SETUP

Before diving into the description of our experimental setup, we provide some definitions that we will use in the rest of the text.

Definition 4.1 (Concept). A concept is broadly defined as an idea or a concrete object that a population of agents is interested in being able to communicate.

Definition 4.2 (Symbol). A symbol is defined as a social agreement on the representation of a concept. In our case, the representation is a 1-D signal, and we assume a 1-to-1 concept-symbol mapping.

Definition 4.3 (Vocabulary). A vocabulary, denoted as V , is defined as the collection of concepts shared (or intended to be shared) by a population of agents, and their associated symbols.

The general setting of the problem of our interest is the following. We consider two agents, one called the *sender*, and the other one called the *receiver*. Both agents share a common set of concepts, and our goal is to evolve an agreement on a signal representation for all concepts (i.e., agents need to create a common vocabulary). Once a common vocabulary is established, the sender is able to propagate information to the receiver (this might be used, e.g., to communicate to the receiver what to pick in a collection of objects, or to indicate in which direction to move in a navigation task).

In the first part of our experiments, concepts are represented by one feature (Section 5.1-5.3), i.e., the “id” of the concept (represented by a number). Later, we also investigate the case in which concepts are represented by m features, introducing a referential game [28] (Section 5.4). In all cases, we focus on the effect of two specific aspects on the evolved signal representation:

- The **decoding** paradigm, i.e., whether modeling signal decoding as a *regression* or a *classification* problem has any influence on the evolutionary process. In the regression setting, the receiver’s NN has only one output node, i.e., the concept it decoded. In the classification setting, the receiver’s NN has $|V|$ output nodes (see Table 2), representing the probability distribution over the vocabulary V . While this latter approach is more robust (as it allows probabilistic decoding, thus being potentially less affected by noise), regression is interesting in that it is associable to numerical meaning of counting, one of the main concepts that humans acquire in the first 3-4 years of childhood. This, in turn, may facilitate the acquisition of higher-level concepts [9].
- The available **amplitude** of the signals generated by senders, which can be *limited* or *unlimited* depending on the sender activation function (see Table 2). Intuitively, evolving a symbolic representation for a set of concepts in a limited amplitude setting is much more difficult with respect to an unlimited setting, since signal representations need to compete for a limited bandwidth and avoid overlapping, which may lead to ambiguity.

Thus, we consider four settings in total, that we refer to *regression limited/unlimited* and *classification limited/unlimited* respectively. For each experimental setting under investigation, we are interested in analyzing the evolution of the agreement of the signal representation in the population of senders and receivers.

4.1 Agents

A sender-receiver pair can be seen as an encoder-decoder structure. More specifically, the *sender* is controlled by a CTRNN that takes as input the representation of an object (that may be either its id or a set of properties, as shown in Section 5.4) and sends to a channel a continuous-time signal (which is sampled into a fixed number of equally-spaced samples over a given time window) that encodes a symbol associated to that object. On the other hand, the *receiver* is controlled by a CTRNN that takes as input the sampled signal sent by the sender and has to *decode* it into the corresponding object. With reference to Eq. (2), we assume $\Delta t = 1$. We should note that this configuration is purposely supposed to “force” the communication between the agents: in fact, by driving the evolutionary process towards the communication between the agents, we expect to observe the emergence of a symbolic communication system.

4.2 Co-evolutionary algorithm

In order to evolve the CTRNNs, we used the implementation of NEAT provided by the NEAT-Python library¹, and built on top of it a co-evolutionary algorithm to allow the evolution of sender/receiver pairs². Tables 1-2 show the main parameters used in the experiments, which are based on the default values set in NEAT-Python.

The co-evolutionary algorithm (whose full pseudo-code is omitted for brevity) consists of two parallel independent executions of NEAT, one handling a population of senders and one handling a population of receivers, which are synchronized at each generation on the evaluation process. The latter, whose pseudo-code is shown in Algorithm 1, consists in the following. Sender/receiver pairs are formed by the Cartesian product between all sender genomes and receiver genomes. Each pair p uses its time window to propagate the signal (from sender to receiver), after which the fitness of p (shared between sender and receiver), to be maximized, is computed as:

$$p.\text{fitness} = - \sum_{\text{concepts}} \mathbb{I}[\text{concept}_{\text{sent}} \neq \text{concept}_{\text{received}}] \quad (3)$$

where concepts indicates the set of concepts considered for the evaluation of the pair, and $\mathbb{I}[\text{concept}_{\text{sent}} \neq \text{concept}_{\text{received}}]$ is an indicator function returning 1 if the concept sent is different from the concept received and 0 otherwise.

Since each single sender/receiver genome is part of multiple pairs, eventually we consider as its final fitness the best one (maximum) among the fitnesses of all the pairs to which it participated.

The co-evolutionary process stops either when a pair is found with fitness equal to zero (indicating no communication error), or after a total budget of 10000 generations for each of the two populations. Each population is randomly reset in case of 50 generations without improvement.

5 RESULTS

We present now the results of our experiments: first we consider simple symbolic communication, then we test zero-shot communication, and communication in the presence of noise. Finally, we apply our evolutionary model to a multi-feature referential game.

¹ Available at: <https://github.com/CodeReclaimers/neat-python>.

² Our code is publicly available at: <https://github.com/FraLotito/evol-signal-comm>.

Algorithm 1: Co-evolutionary evaluation process.

```

Let  $S, R$  be the sets of sender and receiver genomes;
Let  $P = S \times R$  be the set of sender/receiver pairs to evaluate;
Let  $x.fitnesses$  be the set of fitnesses computed  $\forall x \in S, R$ ;
Let  $computeFitness()$  be the function reported in Eq. (3);
foreach  $p \in P$  do
     $p.fitness = computeFitness(p)$ ;
     $p.sender.fitnesses = p.sender.fitnesses \cup p.fitness$ ;
     $p.receiver.fitnesses = p.receiver.fitnesses \cup p.fitness$ ;
foreach  $x \in S, R$  do
     $x.fitness = \max(x.fitnesses)$ ;

```

Table 1: Common parameter settings.

Parameter	Value	Parameter	Value
Sender population	20	Receiver population	20
Time window	10	Sender input nodes	1
Sender hidden nodes	0	Sender output nodes	1
Receiver input nodes	1	Receiver hidden nodes	0
Connection add prob.	0.5	Connection delete prob.	0.5
Node add prob.	0.2	Node delete prob.	0.2
Elitism (species)	1	Elitism (individual)	1
Reset on extinction	1	Max generations (stagnation)	50
Feed forward	False	Max generations (total)	10000

Table 2: Specific parameter settings.

Name	Regression		Classification	
	Unlimited	Limited	Unlimited	Limited
No. receiver outputs	1	1	$ V $	$ V $
Sender activation func.	Identity	Sigmoid	Identity	Sigmoid
Receiver activation func.	Identity	Identity	Sigmoid + Softmax	Sigmoid + Softmax

5.1 Symbolic communication

The first question we try to address is whether the evolution of symbolic communication is facilitated or not by: 1) the use of a limited or unlimited signal amplitude and 2) formulating the decoding step as either a regression or classification problem.

In order to answer this question, we consider the four aforementioned settings, and evolve agents to share a collection of $|V|$ concepts (i.e., a vocabulary of size $|V|$), each one described by one feature. We assume that there exists an order (shared in the population) between the concepts, such that each concept is correctly identified by its index. The concept i is described by the feature $\frac{i}{|V|}$, therefore the feature representation is limited in the interval $[0, 1]$.

Figure 1 shows the evolutionary trend an example run on the regression setting with limited amplitude, with the corresponding signals evolved at different steps of the process. In the first generations, the subdivision of the amplitude between the different signals is very coarse, leading to ambiguity in the symbolic representation, and therefore to low fitness value. The evolution eventually leads to a change in the shape of the signals (from triangular to constant) and to an unambiguous subdivision of the amplitude.

Figure 2 shows the number of generations required for the evolution to converge to a zero-fitness solution, as a function of the

vocabulary size. For each vocabulary size, the corresponding value is averaged over 20 evolutionary runs. We observe that there is a large gap in the convergence time (which can be seen as a proxy to complexity of the problem) when comparing the regression and classification settings, the first one being easier to solve than the latter. We further observe that limiting the available amplitude for the signals increases the hardness of the problem in both settings.

Figure 3 shows the fitness trend (averaged over 20 runs) in each of the four settings, with a vocabulary size $|V| = 5$ (4 for the classification setting with limited amplitude, for which it was not possible to converge to a solution within the 10000 generations, see Figure 2b). Also in this case we note that the regression settings are quickly solved in a few tens of generations, while in the classification settings the evolution shows a much slower convergence.

Table 3: Zero-shot experimental results (no. of correctly communicated concepts, including $|V_T|$).

Setting	$ V_T = 3$	$ V_T = 5$	$ V_T = 7$
Regression, unlimited	4.75 ± 2.4	7.90 ± 1.7	9.20 ± 1.2
Regression, limited	3.15 ± 0.5	5.10 ± 0.3	7.30 ± 0.5
Classification, unlimited	3.15 ± 0.4	5.00 ± 0.0	NA
Classification, limited	3.00 ± 0.0	NA	NA

5.2 Zero-shot communication

The second question we address is whether the evolved pairs of agents under the four settings are capable of performing zero-shot communication, i.e., correctly communicate concepts not seen at fitness evaluation time. In other words, to assign an unseen concept to the sender and see if the sender/receiver communication is correctly carried out, even if there was no previous agreement on the signal to use. More formally: let V be the vocabulary and $V_T \subset V$ be the set of concepts considered for fitness evaluation, we are interested in observing if the agents in the best evolved pair p are able to communicate, at the end of the evolutionary process, also the concepts in $V - V_T$.

In these experiments, we set $|V| = 10$, and we measure the number of concepts in V (including those in $|V_T|$) that are correctly communicated by the best pair at the end of the evolution, for three different values of $|V_T|$. The results are reported in Table 3, averaged over 20 evolutionary runs. In the table, “NA” means that the setting did not converge to a solution. We observe that zero-shot communication occurs effectively only in the regression cases with unlimited amplitude, where 2 to 3 unseen concepts are communicated correctly. On the contrary, the classification settings seem to be unable to generalize to unseen concepts.

5.3 Communication in the presence of noise

So far we have assumed the channel of communication to be reliable, a strong assumption considering that many real-world channels are actually subject to noise. In order to investigate the effect of noise, we replicate the experiments of Section 5.1, and apply Gaussian noise $\mathcal{N} \sim (0, \sigma^2)$ over the signals generated by senders. Also in this case we set $|V| = 5$ for all cases except the classification setting with limited amplitude, for which we set $|V| = 4$.




Figure 1: Evolution of signals across generations in the regression setting with limited amplitude. The agents first evolve a coarse subdivision of the amplitude, and then progressively refine it, until all concepts are correctly communicated.

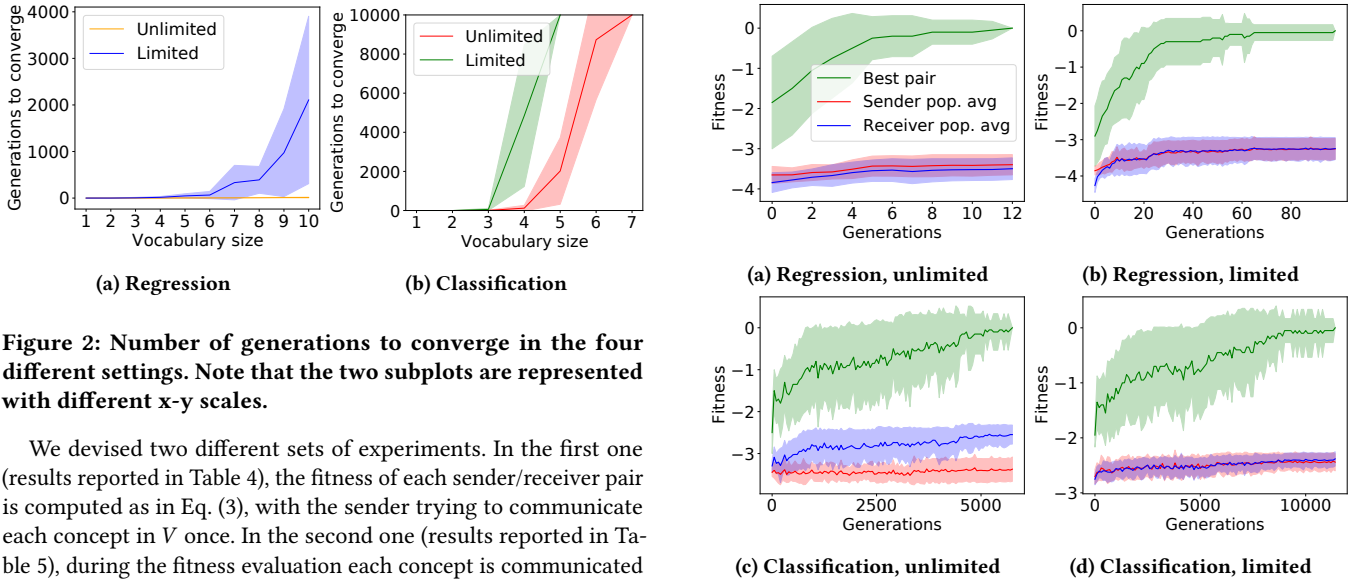


Figure 2: Number of generations to converge in the four different settings. Note that the two subplots are represented with different x-y scales.

We devised two different sets of experiments. In the first one (results reported in Table 4), the fitness of each sender/receiver pair is computed as in Eq. (3), with the sender trying to communicate each concept in V once. In the second one (results reported in Table 5), during the fitness evaluation each concept is communicated 3 times (i.e., there are $3 \times |V|$ trials of communication), and Eq. (3) is modified accordingly to account for the 3 trials per concept. In this second experiment, the convergence is expected to be more difficult than in the first one, since a sender/receiver pair needs to be able to succeed in more communication trials subject to noise, thus requiring to evolve a much more robust communication system.

After evolution, we then consider a test phase during which the best evolved pair performs, for each concept in V , 20 communication trials (25 times for $|V| = 4$), thus for a total of 100 test cases for each pair. To control the amount of noise in the system, we consider four different values for σ , applied both at evolution and test time.

In Tables 4 and 5 we report, for each setting, the results in terms of successfully communicated concepts over 100 test cases, averaged over 20 runs. In the tables, “NA” means that the setting did not converge to a zero-fitness solution. Overall, the percentage of successful communication tends to decrease as the noise level

Figure 3: Fitness trend (averaged over 20 evolutionary runs) in the four different settings with a vocabulary of size 5 (4 for the classification setting with limited amplitude).

increases. A notable exception to this trend is the regression unlimited amplitude case with $\sigma = 1$ and 3 trials per concept (Table 5), where the percentage of success increases w.r.t. $\sigma = 0.5$. This effect seems to suggest that there is a certain noise threshold after which large noise levels force signals to be better separated (especially in the unlimited amplitude case where there is no constraint on the inter-signal distance), thus leading to lower ambiguity/higher robustness and a higher percentage of successful communication. Although, the same effect is not observed in the corresponding classification case. This intuition is however confirmed by observing the equally spaced constellations of signals corresponding to

this case, reported in the Supplementary Material (available on our GitHub repository).

Another interesting finding is that in the unlimited amplitude cases the classification setting consistently outperforms the regression setting in terms of robustness to noise. Finally, we observe that the setting with 3 trials per concept is more challenging since there are cases in which the evolutionary process does not converge to a solution; however, when convergence occurs this setting leads to higher percentages of success w.r.t. the corresponding cases with 1 trial per concept.

Table 4: Noise experimental results (1 trial per concept).

Setting	$\sigma = 0.1$	$\sigma = 0.2$	$\sigma = 0.5$	$\sigma = 1$
Regression, unlimited	93.00 \pm 8.60	72.40 \pm 19.5	52.35 \pm 30.2	37.15 \pm 27.2
Regression, limited	62.70 \pm 8.30	36.90 \pm 16.1	18.45 \pm 3.30	17.50 \pm 3.90
Classification, unlimited	96.55 \pm 9.70	94.55 \pm 7.20	88.25 \pm 9.30	72.75 \pm 29.2
Classification, limited	47.45 \pm 18.2	28.25 \pm 9.90	24.50 \pm 6.30	26.40 \pm 4.60

Table 5: Noise experimental results (3 trials per concept).

Setting	$\sigma = 0.1$	$\sigma = 0.2$	$\sigma = 0.5$	$\sigma = 1$
Regression, unlimited	95.90 \pm 4.6	88.60 \pm 8.00	82.60 \pm 9.3	94.45 \pm 8.9
Regression, limited	67.45 \pm 5.1	65.65 \pm 14.0	NA	NA
Classification, unlimited	97.00 \pm 5.5	96.50 \pm 4.00	93.35 \pm 4.6	90.95 \pm 5.3
Classification, limited	74.70 \pm 7.6	NA	NA	NA

5.4 Referential game

To further analyze the capabilities of the evolved communication system, we introduce a simple referential game [28]. Both sender and receiver are provided with a collection of N objects, each made up of m binary features. The order of the objects is preserved in both collections, such that each object is identifiable by its position without ambiguity. The sender chooses an object, and, after a time window allotted for the communication, the receiver is requested to pick the same object chosen by the sender, see Figure 4.

In this game, the sender/receiver pair is required to evolve an effective communication protocol based on 1-D signals which describe m -dimensional feature representations of objects.




Figure 4: Conceptual scheme of the referential game.

We rely again on the co-evolutionary algorithm described in Section 4.2, computing the fitness of each sender/receiver pair p (to be maximized) as follows:

$$p.\text{fitness} = - \sum_{\text{objects}} \mathbb{I}[\text{object.id}_{\text{sender}} \neq \text{object.id}_{\text{receiver}}] \quad (4)$$

where:

- objects represents the set of objects available in the experiment;
- $\mathbb{I}[\text{object.id}_{\text{sender}} \neq \text{object.id}_{\text{receiver}}]$ is an indicator function returning 1 if the object picked by the sender is different from the object picked by the receiver and 0 otherwise.

Therefore, to compute the fitness of each pair, the referential game is repeated once for every object in the set, and the fitness considers the number of objects correctly recognized. Also in this case we take as final fitness of each agent the maximum among the fitnesses of all the pairs to which it participated. The stop criteria and all the other parameters are set as Section 4.2. In this case the optimal solution is to have a communication protocol that allows the receiver to recognize every object picked by the sender.

Each object is described by 3 binary features, yielding a maximum of 8 different objects. We perform experiments with 3 different sizes of the set of the objects to be shared (3, 5 and 8). In the case of 3 and 5 objects, these are chosen randomly (once for all the runs of each setting) among the 8 possible triples of bits. For each size and setting, we execute 20 evolutionary runs, and average the results.

Considering the unlimited amplitude case, both under regression and classification settings, the evolution achieves 100% success rate in the case of 3 and 5 objects, while it fails to converge to a solution in the case of 8 objects. For illustration purposes, we report the fitness trend regarding the regression and classification settings with 5 objects in Figure 5, where once again we can notice the slower convergence of the classification setting.

In the experiments with limited amplitude, both the regression and classification settings fail to reach convergence for any number of objects from 3 to 8, confirming our previous observations on the increased difficulty due to the use of limited signal amplitudes.

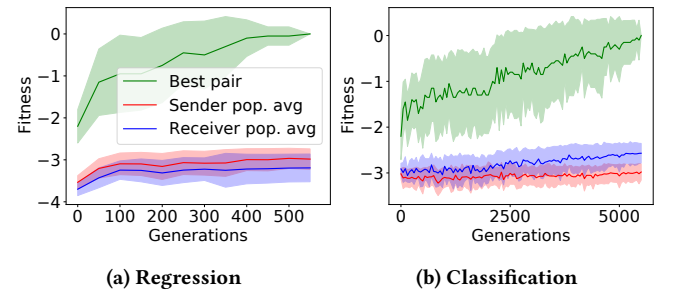


Figure 5: Fitness trend of the referential game experiments in the regression and classification settings with unlimited amplitude and a vocabulary made up of 5 objects.

6 ANALYSIS OF THE EVOLVED SIGNALING SYSTEMS

As seen earlier, the evolved signaling systems can be seen as an agreement between two parties (i.e., the sender and the receiver) that must optimize their behavior in order to find an optimal encoding-decoding strategy: this, in turn, can be seen as a trade-off between simplicity (for the receiver to understand the meaning of signals) and expressiveness (for the sender to convey the meaning of signals). In this section, we will analyze the communication systems evolved, and we will show how the signals evolved tend to form clusters characterized by different specific waveforms.

6.1 Constellations

In order to understand how the signaling systems evolved, we apply the procedure described in [17] to analyze the constellations of signals produced by the system in the regression setting with limited amplitude, $|V| = 5$ and without noise. In our model, the signals s_k , $k = 1 \dots |V|$, represent time series of 10 samples each (size of the time window). The procedure from [17] works as follows. In the first step, we take the first signal from the evolved signaling system, s_1 , and compute φ_1 , i.e., the first base of a vector space, as:

$$\varphi_1 = \frac{s_1}{\|s_1\|}$$

where $\|\cdot\|$ indicates the 2-norm. Then, we iterate on each signal s_k , $k = 2 \dots |V|$, and compute $(s_k)_{\perp(k-1)}$, that is its orthogonal component with respect to all the previous bases:

$$(s_k)_{\perp(k-1)} = s_k - \sum_{i=0}^{k-1} \langle s_k, \varphi_i \rangle \varphi_i$$

where $\langle \cdot, \cdot \rangle$ indicates the scalar product. With this, we calculate a new base φ_k as follows:

$$\varphi_k = \frac{(s_k)_{\perp(k-1)}}{\|(s_k)_{\perp(k-1)}\|}.$$

In our analysis, we set φ_k to a zero-vector if $\|(s_k)_{\perp(k-1)}\| < 10^{-4}$. Furthermore, we observe that in all our settings the evolved signaling systems generate constellation of signals that can be projected on a vector space of up to 2 dimensions. This allows us to represent the signals in a 2-D plane by simply computing the projection of each signal on φ_1 and φ_2 .

A typical evolutionary process (in the considered regression setting) proceeds as follows:

- In the first generation (Figure 6a), the signals lie on a plane. As we can see, the margin between signals is small, meaning that some pairs of signals are less easy to distinguish from others.
- Then, for a certain number of generations (Figures 6b-6e) the signals lie on a line (i.e., there are no orthogonal components, somewhat similar to a PAM system [4]) but, as in the first generation, there are still pairs of signals separated by a small margin (with respect to other pairs).
- Eventually (Figure 6f), the spacing between signals increases, leading to a more robust signaling system (optimal zero-fitness).

It is important to note that in this specific setting the best signaling system does not necessarily correspond to a system with equally spaced signals. This is due to the absence of noise, which entails that an equal margin between the signals does not constitute an evolutionary advantage.

In comparison, the same analysis on one of the runs of the regression case with limited amplitude in the presence of noise ($\sigma = 0.1$) with 3 trials per concept (results shown in Table 5) reveals that, while in the first part of the evolutionary process the constellations are similar to the ones shown in Figure 6a-6e, eventually the signaling system at the end of the evolutionary process (Figure 7) consists in almost-equally spaced signals that allow a better adaptation to the noise present on the channel. As mentioned earlier, the analysis reported in the Supplementary Material for the other cases with noise show similar results.

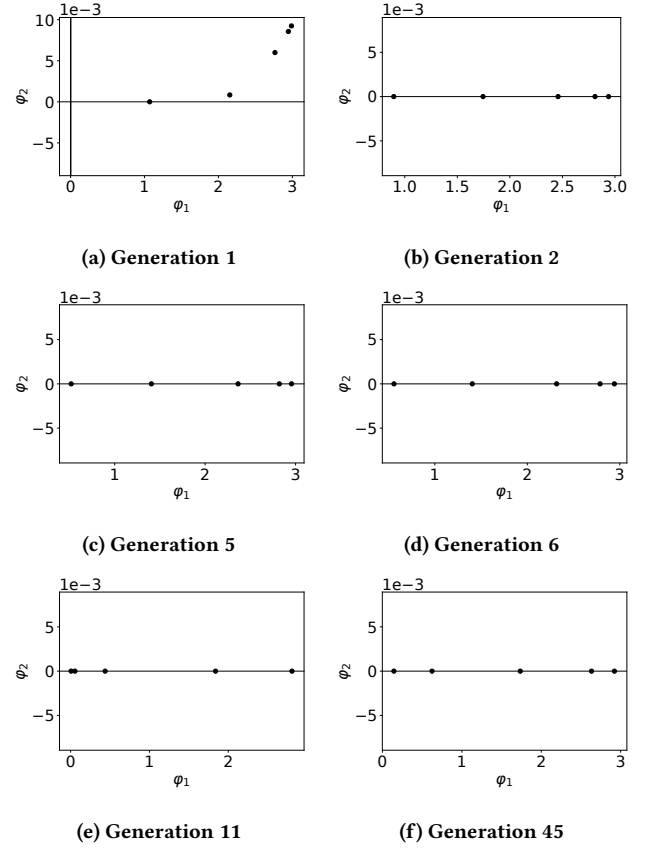


Figure 6: Constellation of signals of the best signaling system evolved during a typical evolutionary run on the regression setting with limited amplitude, vocabulary of size 5, and without noise.




Figure 7: Constellation of signals of the best signaling system evolved on the regression setting with limited amplitude, vocabulary of size 5, and noise ($\sigma = 0.1$, 3 trials).

6.2 Clustering of the evolved signaling systems

To understand whether the evolved signaling systems show common features, we perform a clustering of the best signaling systems evolved in the 20 runs of the settings tested in Section 5.1. We exclude from this analysis the classification with limited amplitude setting, due to the different vocabulary size w.r.t. the other settings. Since our goal is to discover if and how many clusters of signaling systems exist, we apply a clustering algorithm that, as opposed to non-density-based approaches (e.g. k-means), does not require to fix the number of clusters. Specifically, we use OPTICS [1], a density-based clustering approach available in `scikit-learn` [33]. The parameters we used for the algorithm are: `min_samples = 4`; `metric = Chebyshev distance`; `xi = 0.1`. All the other parameters are set to their default values.

The result of the clustering is shown in Figure 8, where we present a 2-D representation of the data by using the t-SNE algorithm [29]. We can see that most of the clusters contain signaling systems resulting from the same setting. Furthermore, for each cluster we compute the central signal by selecting the point that minimizes the mean distance from the other points in the cluster. The central signals are shown in Figure 9. We observe four clusters of signals, of which three appear more clearly identifiable: the first one (Figure 9a) contains signals that become constant at the first timestep; the third one (Figure 9c) contains signals with an exponential-like shape; the fourth one (Figure 9d) contains signals that look like triangular waves. On the other hand, the second cluster (Figure 9b) is less clearly defined. In fact, it mostly contains signals that become constant after the second timestep, but also signals that look like the ones present in the first cluster.

Finally, by performing nearest neighbour classification (using the Chebyshev distance) of the signals we obtain that: 53.3% of the signaling systems are in the first cluster; 13.3% in the second one; 23.3% in the third one; 10% in the fourth one. Thus the first cluster appears to be a strong attractor for the evolutionary process.




Figure 8: The evolved signaling systems in a 2-D vector space obtained with t-SNE. Each color represents a cluster (points classified as “noise” are in gray). The central signal for each cluster are represented in Figure 9: green (Figure 9a); yellow (Figure 9b); blue (Figure 9c); orange (Figure 9d).

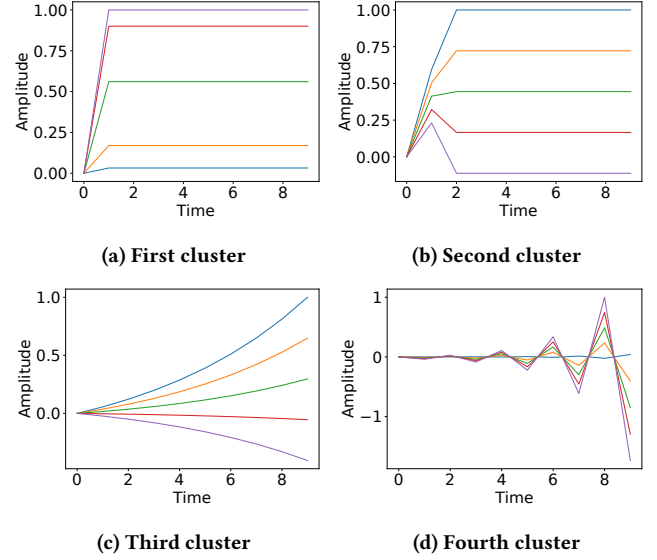


Figure 9: Central signals for each cluster shown in Figure 8.

7 CONCLUSIONS

In this paper, we studied the evolution of symbolic communication in two contexts: a simple symbolic communication task and a referential game. The agents communicate through waveforms sent on a dedicate channel. We analyzed the communication systems evolved in several settings under different points of view: convergence time, generalization capabilities, robustness to noise and properties of the evolved signals.

We found that the definition of the problem (i.e., classification or regression) has a significant impact on both the number of generations required to solve the task and the generalization capabilities. Furthermore, we observed that adding noise to the channel leads to communication systems that try to maximize the “distance” between the signals. Finally, we observed that the evolutionary process usually led to similar types of signaling systems, which can be grouped in four main categories, of which one appears dominant.

This study has, of course, some limitations. The first one concerns the agent controller. In fact, while the results obtained show that forms of symbolic communication can be evolved by using continuous-time recurrent neural networks, other forms of controllers/function approximators based e.g. on vanilla RNNs, Genetic Programming or Markov Brains may not necessarily lead to similar results. Another limitation regards the simplicity of our approach w.r.t. the evolution of communication in real-world scenarios. For instance, in our work we define and constrain one channel of communication, while in the real world communication may occur on multiple channels, and its effectiveness may depend on the context. Nevertheless, we believe that some of our findings (e.g. on the complexity of the problem and the effect of modelling it as either classification or regression) may hold true also in real-world scenarios and can be extended also to other controller architectures. Future work will address these two issues, and possibly focus on more complex (e.g. two-way) communication tasks as well as adversarial perturbations on the signals.

REFERENCES

- [1] Mihael Ankerst, Markus M Breunig, Hans-Peter Kriegel, and Jörg Sander. 1999. OPTICS: ordering points to identify the clustering structure. *ACM Sigmod record* 28, 2 (1999), 49–60.
- [2] Jordi Arranz, Jason Noble, and Eric Silverman. 2011. The origins of communication revisited. In *ECAL*. MIT Press, Cambridge, MA, USA, 8.
- [3] Jeffrey A Barrett, Brian Skyrms, and Calvin Cochran. 2018. Hierarchical Models for the Evolution of Compositional Language. Preprint.
- [4] John R. Barry, Edward A. Lee, and David G. Messerschmitt. 2004. *Pulse-Amplitude Modulation*. Springer US, Boston, MA, 131–201.
- [5] Randall D. Beer. 1995. On the Dynamics of Small Continuous-Time Recurrent Neural Networks. *Adaptive Behavior* 3, 4 (Jan. 1995), 469–509.
- [6] Oliver Bown and Sebastian Lexer. 2006. Continuous-Time Recurrent Neural Networks for Generative and Interactive Musical Performance. In *EuroGP*. Springer-Verlag, Berlin, Heidelberg, 652–663.
- [7] Nicolas Cambier, Roman Miletitch, Vincent Frémont, Marco Dorigo, Eliseo Ferrante, and Vito Trianni. 2020. Language Evolution in Swarm Robotics: A Perspective. *Frontiers in Robotics and AI* 7 (2020), 12.
- [8] A. Cangelosi. 2001. Evolution of communication and language using signals, symbols, and words. *IEEE Transactions on Evolutionary Computation* 5, 2 (April 2001), 93–101.
- [9] Susan Carey. 2000. The origin of concepts. *Journal of Cognition and Development* 1, 1 (2000), 37–41.
- [10] Rahma Chaabouni, Eugene Kharitonov, Emmanuel Dupoux, and Marco Baroni. 2019. Anti-efficient encoding in emergent communication. arXiv: 1905.12561.
- [11] Gautier Dagan, Dieuwke Hupkes, and Elia Bruni. 2020. Co-evolution of language and agents in referential games. arXiv:2001.03361 [cs.CL]
- [12] Joachim de Greeff and Stefano Nolfi. 2010. Evolution of Implicit and Explicit Communication in Mobile Robots. In *Evolution of Communication and Language in Embodied Agents*, Stefano Nolfi and Marco Mirolli (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 179–214.
- [13] Terrence William Deacon. 1998. *The symbolic species: The co-evolution of language and the brain*. W. W. Norton & Company, New York, NY, USA.
- [14] Katrina Evtimova, Andrew Drozdov, Douwe Kiela, and Kyunghyun Cho. 2018. Emergent Communication in a Multi-Modal, Multi-Step Referential Game. arXiv:1705.10369 [cs.LG]
- [15] Jakob N. Foerster, Yannis M. Assael, Nando de Freitas, and Shimon Whiteson. 2016. Learning to Communicate with Deep Multi-Agent Reinforcement Learning. arXiv: 1605.06676.
- [16] Tom Froese and Ezequiel A. Di Paolo. 2010. Modelling social interaction as perceptual crossing: an investigation into the dynamics of the interaction process. *Connection Science* 22, 1 (March 2010), 43–68.
- [17] Robert Gallager. 2012. *Principles of digital communication*. Cambridge University Press, Cambridge, United Kingdom.
- [18] C. Lee Giles and Kam-Chuen Jim. 2003. Learning Communication for Multi-agent Systems. In *Innovative Concepts for Agent-Based Systems*, Gerhard Goos, Juris Hartmanis, Jan van Leeuwen, Walt Truszkowski, Mike Hinchey, and Chris Rouff (Eds.). Vol. 2564. Springer Berlin Heidelberg, Berlin, Heidelberg, 377–390.
- [19] Paul Grouchy, Gabriele M. T. D'Eleuterio, Morten H. Christiansen, and Hod Lipson. 2016. On The Evolutionary Origin of Symbolic Communication. *Scientific Reports* 6, 1 (Dec. 2016), 34615.
- [20] Stevan Harnad. 1990. The symbol grounding problem. *Physica D: Nonlinear Phenomena* 42, 1–3 (1990), 335–346.
- [21] Serhii Havrylov and Ivan Titov. 2017. Emergence of Language with Multi-agent Games: Learning to Communicate with Sequences of Symbols. arXiv: 1705.11192.
- [22] David Kadish and Sebastian Risi. 2020. Adapting to a changing environment: Simulating the effects of noise on animal sonification. In *ALIFE*. MIT Press, Cambridge, MA, USA, 687–695.
- [23] Dariusz Kalociński, Marcin Mostowski, and Nina Gierasimczuk. 2018. Interactive Semantic Alignment Model: Social Influence and Local Transmission Bottleneck. *Journal of Logic, Language and Information* 27, 3 (Sept. 2018), 225–253.
- [24] Eugene Kharitonov, Rahma Chaabouni, Diane Bouchacourt, and Marco Baroni. 2019. EGG: a toolkit for research on Emergence of lanGuage in Games. In *EMNLP-IJCNLP*. Association for Computational Linguistics, Hong Kong, China, 55–60.
- [25] Travis Lacroix. 2018. The correction game or, how pre-evolved communicative dispositions might affect communicative dispositions. Preprint.
- [26] Travis LaCroix. 2020. Communicative bottlenecks lead to maximal information transfer. Preprint.
- [27] Travis LaCroix. 2020. On salience and signaling in sender–receiver games: partial pooling, learning, and focal points. *Synthese* 197, 4 (April 2020), 1725–1747.
- [28] Angeliki Lazaridou, Alexander Peysakhovich, and Marco Baroni. 2017. Multi-Agent Cooperation and the Emergence of (Natural) Language. arXiv: 1612.07182.
- [29] Laurens van der Maaten and Geoffrey Hinton. 2008. Visualizing data using t-SNE. *Journal of machine learning research* 9, Nov (2008), 2579–2605.
- [30] Santosh Manicka. 2012. Analysis of Evolved Agents Performing Referential Communication. In *ALIFE*. MIT Press, Cambridge, MA, USA, 393–400.
- [31] Stefano Nolfi and Marco Mirolli. 2010. Evolving Communication in Embodied Agents: Assessment and Open Challenges. In *Evolution of Communication and Language in Embodied Agents*, Stefano Nolfi and Marco Mirolli (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 215–220.
- [32] M. A. Nowak, D. C. Krakauer, and A. Dress. 1999. An error limit for the evolution of language. *Proceedings of the Royal Society of London. Series B: Biological Sciences* 266, 1433 (Oct. 1999), 2131–2136.
- [33] Fabian Pedregosa, Gaël Varoquaux, Alexandre Gramfort, Vincent Michel, Bertrand Thirion, Olivier Grisel, Mathieu Blondel, Peter Prettenhofer, Ron Weiss, Vincent Dubourg, et al. 2011. Scikit-learn: Machine learning in Python. *Journal of Machine Learning research* 12 (2011), 2825–2830.
- [34] Joshua B Plotkin and Martin A Nowak. 2001. Major transitions in language evolution. *Entropy* 3, 4 (2001), 227–246.
- [35] Naohiro Shibuya, Hiroyuki Iizuka, and Masahito Yamamoto. 2018. Evolution of communication through differentiation of communicative and goal-directed behaviors. *Artificial Life and Robotics* 23, 2 (June 2018), 225–234.
- [36] John Maynard Smith and Eors Szathmáry. 1997. *The major transitions in evolution*. Oxford University Press, Oxford, UK.
- [37] Kenny Smith. 2002. The cultural evolution of communication in a population of neural networks. *Connection Science* 14, 1 (March 2002), 65–84.
- [38] Kenneth O Stanley, Jeff Clune, Joel Lehman, and Risto Miikkulainen. 2019. Designing neural networks through neuroevolution. *Nature Machine Intelligence* 1, 1 (2019), 24–35.
- [39] Kenneth O. Stanley and Risto Miikkulainen. 2002. Evolving Neural Networks through Augmenting Topologies. *Evolutionary Computation* 10, 2 (June 2002), 99–127.
- [40] Jacopo Talamini, Eric Medvet, and Alberto Bartoli. 2019. Communication-Based Cooperative Tasks: How the Language Expressiveness Affects Reinforcement Learning. In *SAC*. ACM, New York, NY, USA, 898–905.
- [41] Jacopo Talamini, Eric Medvet, Alberto Bartoli, and Andrea De Lorenzo. 2020. Communication in Decision Making: Competition favors Inequality. In *ALIFE*. MIT Press, Cambridge, MA, USA, 570–577.
- [42] Elio Tuci. 2009. An investigation of the evolutionary origin of reciprocal communication using simulated autonomous agents. *Biological Cybernetics* 101, 3 (Sept. 2009), 183–199.
- [43] Paul L Williams, Randall D Beer, and Michael Gasser. 2008. Evolving Referential Communication in Embodied Dynamical Agents. In *ALIFE*. MIT Press, Cambridge, MA, USA, 702–709.