

Decentralised Emergence of Robust and Adaptive Linguistic Conventions in Populations of Autonomous Agents Grounded in Continuous Worlds

Jérôme Botoko Ekila
Vrije Universiteit Brussel
Brussels, Belgium
jerome@ai.vub.ac.be

Jens Nevens
Vrije Universiteit Brussel
Brussels, Belgium
jens@ai.vub.ac.be

Lara Verheyen
Vrije Universiteit Brussel
Brussels, Belgium
lara.verheyen@ai.vub.ac.be

Katrien Beuls
Université de Namur
Namur, Belgium
katrien.beuls@unamur.be

Paul Van Eecke
Vrije Universiteit Brussel
Brussels, Belgium
paul@ai.vub.ac.be

ABSTRACT

This paper introduces a methodology through which a population of autonomous agents can establish a linguistic convention that enables them to refer to arbitrary entities that they observe in their environment. The linguistic convention emerges in a decentralised manner through local communicative interactions between pairs of agents drawn from the population. The convention consists of symbolic labels (word forms) associated to concept representations (word meanings) that are grounded in a continuous feature space. The concept representations of each agent are individually constructed yet compatible on a communicative level. Through a range of experiments, we show (i) that the methodology enables a population to converge on a communicatively effective, coherent and human-interpretable linguistic convention, (ii) that it is naturally robust against sensor defects in individual agents, (iii) that it can effectively deal with noisy observations, uncalibrated sensors and heteromorphic populations, (iv) that the method is adequate for continual learning, and (v) that the convention self-adapts to changes in the environment and communicative needs of the agents.

KEYWORDS

language emergence; multi-agent systems; autonomous agents; emergent communication; self-organisation; language evolution

1 INTRODUCTION

Human languages are evolutionary systems, which emerge and evolve through local communicative interactions between members of a linguistic community. Processes of variation and selection are at play during each and every communicative interaction, at the level of concepts, words and grammatical structures [11, 27, 33, 36, 46]. Variants are introduced as creative solutions to communicative impasses and are selected for based on their linguistic, cognitive and physical fitness [14, 17, 48]. The evolutionary and self-organising nature of human languages gives rise to a number of unique qualities. First of all, such decentralised, self-organising systems are known to be robust and to be able to self-repair substantial perturbations [20, 34]. Second, populations of language users converge on shared conventions that still remain adaptive to changes in

their environment and communicative needs [1]. Finally, the resulting languages effectively serve as an abstraction layer above the sensory-motor observations and internal mental representations of individual language users [31]. Indeed, while linguistic forms can be observed and shared, their meanings remain tied to each language user's individual physical and cognitive embodiment.

This agent-based and evolutionary perspective on the human ability to communicate through language has served as a starting point for the development of a range of computational methodologies that model how artificial agents can co-construct emergent languages that satisfy their communicative needs [see e.g. 7, 13, 15, 25, 29, 43]. Rather than modelling the learning of an existing natural language, which has emerged and evolved to fit the communicative needs of a population of human language users, these methodologies allow for *artificial natural languages* to emerge and evolve to optimally support the embodiment, environment and communicative needs of populations of artificial agents. These languages are artificial in the sense that they do not exist outside the experimental set-up, yet natural in the sense that they emerge and evolve through the same evolutionary principles as human languages do.

The last decade has witnessed substantial progress when it comes to the application of these methodologies to a variety of tasks, including visual question answering [12], solving puzzles [15], negotiation [6], reference [25], navigation [4, 29, 47] and coordination in self-driving cars [35]. At the same time, the application-oriented focus of the experiments has resulted in less attention to the modelling of the evolutionary mechanisms through which human languages have emerged and continue to evolve. For example, populations often consist of two agents only, populations are divided into agents that can speak and agents that can listen, learning is not decentralised, or agents have mind-reading capabilities. As a consequence, the languages that emerge do not exhibit the unique qualities of human languages that originally motivated the experimental paradigm [48].

In this paper, we go back to the original motivation of computationally modelling the emergence and evolution of human-like languages and introduce a methodology through which a population of autonomous agents can establish, in a fully decentralised manner, a linguistic convention that enables them to refer to arbitrary entities

that they observe in their environment. By taking part in situated and task-oriented communicative interactions, each agent gradually builds up its own inventory of associations between symbolic labels (word forms) and concept representations (word meanings). While all concept representations are individually constructed and grounded in an agent’s own sensory-motor endowment and experiences, the linguistic systems of all individual agents are compatible on a communicative level. In prior work, the emergence of a linguistic convention for a set of concepts in discrete worlds (represented by categorical data) was demonstrated [53]. The emergence of a shared vocabulary for a set of concepts that are represented as a set of continuous feature channels was demonstrated by [31], but their methodology was limited to acquiring concepts that occur in existing natural languages. In contrast, our methodology is applicable to any dataset that describes entities in terms of combinations of continuously-valued features and is not limited to concepts or words that occur in existing natural languages. Until now, this combination of properties has never been achieved together. Apart from introducing the methodology, we also present a range of experiments that demonstrate the human-like qualities of the emergent languages. As such, we show that the methodology (i) enables a population to converge on a communicatively effective, coherent and human-interpretable linguistic convention, (ii) is naturally robust against sensor defects in individual agents, (iii) can effectively deal with noisy observations, uncalibrated sensors and heteromorphic populations, (iv) is adequate for continual learning, and (v) leads to languages that self-adapt to changes in the environment and communicative needs of the agents.

2 METHODOLOGY

The methodology that we introduce in this paper has its roots in the *language game paradigm*, a methodological framework that was originally conceived to computationally model the origins and evolution of language [30, 40, 48]. Language game experiments simulate how populations of agents can learn to communicate with each other by taking part in pairwise, task-oriented and situated communicative interactions. Our methodology introduces an innovative way in which agents represent, invent, adopt and align concepts, and integrates these representations and processing mechanisms in a fairly standard language game set-up.

Experiment. We define a language game experiment $E = (W, P, G)$ to be a coupling between a world W , a population P and a sequence $G = (g_j)_{j=1}^i$ of i communicative interactions, referred to as *games*.

Population. The population $P = \{a_1, \dots, a_k\}$ comprises a set of k autonomous agents. Each agent $a \in P$ is initialised with an empty linguistic inventory $I^a = \{\}$. Indeed, the agents do not know any concepts or words at the beginning of the experiment. Each agent is endowed with a set of l sensors $S^a = \{s_1, \dots, s_l\}$ through which it can observe its environment. All sensors are required to map their output to values between 0 and 1. The number of sensors and their types are not necessarily the same for all agents in the population.

World. The world $W = \{e_1, \dots, e_m\}$ comprises a set of m entities. Each entity is represented through a feature vector \mathbf{x} , with each dimension of this vector representing a particular sensor that agents can be endowed with. Depending on their individual endowment,

agents can thus perceive an entity through a vector that comprises a subset of these dimensions. The values that are perceived on these dimensions might also differ from agent to agent, for example when noise or calibration differences are included in the experimental set-up. The feature vector \mathbf{x} for an entity in W as perceived by agent $a \in P$ is notated as \mathbf{x}^a .

Linguistic inventory. The linguistic inventory I of an agent $a \in P$, notated as I^a , is a potentially empty set of words, with each word $w \in I$ being a coupling $w = (f, c, s)$ between a word form $f \in F$, a concept representation c and an entrenchment score s . s is bound between 0 and 1. The score is a measure of the word’s effectiveness in past interactions. F is an infinite set of word forms, typically enumerated through a regular expression.

Concept representation. A concept representation $c = ((\omega_1, \mu_1, \sigma_1), \dots, (\omega_l, \mu_l, \sigma_l))$ consists of a sequence of couplings between three numerical values ω , μ and σ . This sequence holds one such coupling for each sensor with which an agent is endowed. The weight value ω_i represents the importance of feature channel i for the concept, the mean value μ_i holds the prototypical value for the concept on this channel and the standard deviation value σ_i holds the standard deviation for the concept on this channel. Concepts are thus represented as a sequence of normal distributions, with one distribution being associated to each feature channel via a weight that indicates the importance of this feature channel for the concept.

Game. Each game $g \in G$ proceeds as follows:

- (1) **Context selection** A context $C = \{e_1, \dots, e_n\} \in W$ consisting of a subset of n entities is randomly selected from the world.
- (2) **Agent and role selection** Two agents $a_1, a_2 \in P$ are randomly selected from the population. a_1 is assigned the role of speaker $S = a_1$, while a_2 is assigned the role of listener $L = a_2$. Each agent perceives the world through its own sensors and has no access to the ‘objective’ feature vectors of the context C (see *World* above).
- (3) **Topic selection** A topic entity $T \in C$ is randomly selected from the context and is only disclosed to the speaker S . It is the task of S to draw the attention of the listener L to T using a word from the speaker’s linguistic inventory $w \in I^S$.
- (4) **Conceptualisation and production** The speaker S computes the similarity $\text{sim}(c, \mathbf{x}^S)$ between the concept representation $c = ((\omega_1, \mu_1, \sigma_1), \dots, (\omega_l, \mu_l, \sigma_l))$ of each word in its linguistic inventory $w = (f, c, s) \in I^S$ and the perceived feature vector $\mathbf{x}^S = (x_1, \dots, x_l)$ for each entity in the context C . As specified in Equation 1, this is done by computing for each channel the z-score of the perceived value given the distributions stored in the concept representation. These z-scores are then mapped to values between 0 and 1 by first applying the exponential function to their absolute values and then computing the multiplicative inverse of the result. For each channel i in the concept representation, the resulting value is then weighted according to the weight value ω_i . ω_i is itself normalised by the sum of the weights on all channels. This normalisation step avoids an inherent bias towards concept representations with a higher number of relevant channels. The metric $\text{sim}(c, \mathbf{x}^S)$ is then computed as the sum of the resulting values on all channels.

$$\text{sim}(c, \mathbf{x}^a) = \sum_{i=1}^l \underbrace{\frac{\omega_i}{\sum_{k=1}^l \omega_k}}_{\text{normalised weight}} * \underbrace{\frac{1}{\exp(|\frac{x_i - \mu_i}{\sigma_i}|)}}_{\text{channel similarity}} \quad (1)$$

All words in the speaker’s linguistic inventory, i.e. $w \in I^S$, for which the similarity between their concept representation c and the perceived feature vector for the topic entity T is larger than the similarity between c and any other entity in C are collected as candidate words. As such, the set of candidate words corresponds to all words in the speaker’s inventory that distinguish the topic entity from the other entities in the context. Then, the candidate words are ranked according to their communicative adequacy, computed as the product of their entrenchment score s and their discriminative power DP . The DP is computed as the similarity between c and T minus the similarity between c and the closest other entity in C . In other words, it measures how well the topic is discriminated from all other objects using the concept c . The word form f of the candidate word with the highest communicative adequacy is then uttered by S as the utterance U . U is shared between S and the listener L . If there are no candidate words in I^S :

- (4a) **Invention** A new word $w = (f, c, s)$ is added to the speaker’s linguistic inventory I^S , with f being randomly selected from the infinite set of forms F (see *Linguistic inventory* above) and s being assigned a default initial value. The concept representation $c = ((\omega_1, \mu_1, \sigma_1), \dots, (\omega_l, \mu_l, \sigma_l))$ is initialised with $\mu_1 \dots \mu_l$ being the values of the perceived feature vector \mathbf{x}^S , $\sigma_1 \dots \sigma_l$ being assigned a default initial value, and $\omega_1 \dots \omega_l$ being assigned a default initial value as well. Then, f is uttered as U .
- (5) **Comprehension and interpretation** The listener L observes the utterance U . If L knows a word with the form U , i.e. $w = (U, c, s) \in I^L$, L computes the similarity between c and every entity in the context C using the similarity metric specified in Equation 1. L then points to the entity that is most similar to c . If L does not know a word with the form U , no pointing happens and L signals that it could not understand.
- (6) **Feedback** If the listener L pointed to the topic entity T , the speaker S signals success. Otherwise, S signals failure and provides feedback by pointing to T^1 .
- (7) **Alignment** If the game g was successful, the speaker S will increase the score s of the used word $w = (U, c, s) \in I^S$ by a fixed reward value. At the same time, the scores of the word’s competitors, i.e. all other $w \in I^S$ that were earlier identified as candidate words (see *Conceptualisation and production* above), are decreased by a value that is proportional to how similar their concept representation is to the concept representation of the used word. This value is computed by multiplying the similarity between both concept representations by a fixed inhibition value. As specified in Equation 2, the similarity between two concept representations is computed as the sum over all

channels of the Hellinger similarity between the corresponding distributions [19], multiplied by the similarity between their normalised weights and their average normalised weight.

The *Hellinger similarity* component is included to reflect the relative importance of the similarity between the distributions for corresponding channels, where closer distributions lead to a higher similarity. The *similarity of normalised weights* component is included to reflect the relative importance of the similarity between the weights on corresponding channels, where a smaller difference between the weights indicates a higher similarity between the channels. Finally, the *average normalised weights* component is included to reflect that channel similarities are more meaningful if channel weights are higher, with channels holding a higher average weight contributing more to the overall similarity score.

$$\text{sim}(c^q, c^r) = \sum_{i=1}^l \overbrace{\left[(1 - H(\mathcal{N}(\mu_i^q, (\sigma_i^q)^2), \mathcal{N}(\mu_i^r, (\sigma_i^r)^2))) \right]}^{\text{Hellinger similarity}} * \underbrace{\left(1 - \left| \frac{\omega_i^q}{\sum_{k=1}^l \omega_k^q} - \frac{\omega_i^r}{\sum_{k=1}^l \omega_k^r} \right| \right)}_{\text{similarity of normalised weights}} * \underbrace{\left[\frac{\frac{\omega_i^q}{\sum_{k=1}^l \omega_k^q} + \frac{\omega_i^r}{\sum_{k=1}^l \omega_k^r}}{2} \right]}_{\text{average normalised weights}} \quad (2)$$

The listener L then collects all words in its linguistic inventory that can be considered candidate words according to the procedure described in *Conceptualisation and production* above (now based on I^L and \mathbf{x}^L instead of I^S and \mathbf{x}^S). The scores of the word with form U and of the competing words are updated in the same way as is done for the speaker.

Both S and L will also update their concept representation associated to U based on the context C (as perceived by the individual agents). On each channel i , they update μ_i and σ_i to include their perceived feature vector (\mathbf{x}^S or \mathbf{x}^L) using Welford’s online algorithm [52]. The weights on the channels are also updated for both S and L . In a first phase, the channels with a positive discriminative power (see *Conceptualisation and production* above) are identified, i.e. the channels that have a higher similarity to T than to any other entity in C according to Equation 1. Then, all subsets of the powerset of all channels that at least contain the set of channels with positive discriminating power are considered, and the subset with the highest discriminative power for T with relation to C is selected. The weights on the channels in this subset are increased by a fixed step on a sigmoid function and the weights on the other channels are decreased by a fixed step on the same function. The weight values are thereby bounded between 0 and 1, with values becoming more stable as they approach 0 or 1.

If the game g was not successful, there are two distinct cases for alignment. If the failure was due to the L pointing to a different entity than T , S will decrease the score of $w = (U, c, s) \in I^S$ by

¹Pointing and observing the pointing of others are implemented by revealing the feature vector for the topic entity. However, this always remains the feature vector \mathbf{x}^a as perceived through the sensors of the individual agent a (i.e. the objective ‘world’ vector \mathbf{x} is never revealed). The link between these subjective vectors and the objective vector is kept for computing the external evaluation metrics.

Table 1: Overview of parameters with standard default values.

Description	Parameter	Default
# agents in population	k	10
# entities in context	n	3...10
# sensors per agent	l	all (homomorphic)
initial entrenchment score	s^i	0.5
entrenchment reward	s^r	+0.1
entrenchment punishment	s^p	-0.1
entrenchment inhibition	s^{li}	$-0.02 * sim(c_q, c_r)$
initial standard deviation	σ^i	0.01
initial channel weight	ω^i	0.5
sigmoid function	f	$\frac{1}{1+e^{-1/2x}}$
channel weight reward	c^r	+1
channel weight punishment	c^p	-5

a fixed value. L will also decrease the score of $w = (U, c, s) \in I^L$ by a fixed value and update c based on T with relation to C in the same way as if the game would have been successful. If the failure was due to L not knowing a word $w = (U, c, s)$, S will decrease the score of $w = (U, c, s) \in I^S$ by a fixed value and L will adopt the word as follows:

- (7a) **Adoption** A new word $w = (U, c, s)$ is added to I^L , with s being assigned a default initial value. The concept representation $c = ((\omega_1, \mu_1, \sigma_1), \dots, (\omega_l, \mu_l, \sigma_l))$ is initialised with $\mu_1 \dots \mu_l$ being the values of the perceived feature vector X^L , $\sigma_1 \dots \sigma_l$ being assigned a default initial value, and $\omega_1 \dots \omega_l$ being assigned a default initial value as well.

The formal definition of the methodology specifies a number of parameters that need to be set when carrying out concrete experiments. A first set of parameters concerns the general experimental set-up. It specifies the number of agents in the population (k), the number of entities in the context of a single communicative interaction (n), and, per agent, a list of sensors with which it is endowed. The second set of parameters specifies how the scores of words are updated after each interaction. It specifies the initial score of words that are invented or adopted (s^i), along with update rules for words that were used successfully (s^r), for words that were used unsuccessfully (s^p), and for words that compete with words that were used successfully (s^{li}). The final set of parameters concerns the initialisation and updating of concept representations. It specifies the initial weight and standard deviation for feature channels in new concepts (ω^i and σ^i), the sigmoid function along which channel weights are increased or decreased (f) and the step on this function by which the weights are shifted in case of success (c^r) or failure (c^p). An overview of these parameters along with standard default values is provided in Table 1. The supplementary material details the hyperparameter tuning process.

3 EXPERIMENTAL VALIDATION

This section presents a range of experiments that were designed to serve as an initial validation of the methodology introduced in Section 2, as well as to demonstrate the robustness, flexibility and adaptivity of the emergent languages. Three datasets were

chosen for this experimental validation, based on their public availability, the fact that they describe entities in terms of continuous features, and the diversity of domains that are covered. The first dataset (henceforth CLEVR) makes use of the images of the CLEVR dataset [21], which were preprocessed according to the procedure described by Nevens et al. [31]. Concretely, the resulting dataset comprises 85,000 images, in which each depicted object is represented through a feature vector. The 20 dimensions of these feature vectors correspond to information obtained through computer vision techniques, including the number of corners of an object, its width-height ratio, its color channel values, and its position on the horizontal and vertical axes. The second dataset (henceforth WINE) concerns the Wine Quality dataset [9], which holds information about 4898 wine samples along 11 dimensions that describe their physicochemical characteristics (e.g. acidity, residual sugar, alcohol and sulphates). The third dataset, called Credit Card Fraud Detection [10] (henceforth CREDIT), holds 284,807 entries of financial transactions described along 28 dimensions resulting from a principal component analysis. As such, the resulting datasets cover three very different types of data, ranging from visual scenes over physicochemical analyses to principal components extracted from financial transaction records.

For each of the three datasets, the data is split into training and test sets. For each split, a set of scenes is created. Each scene consists of a unique sets of 3 to 10 entities that can serve as the context for a language game (see the *Context selection* step in Section 2). In other words, a scene represents the ‘environment’ in which the communication interaction takes place. Each scene consists of 3 to 10 entities. Due to split, entities that occur in a training scene are excluded from being part of a test scene.

3.1 Emergence of a communicatively effective, coherent and interpretable convention

The first experiment validates the methodology on the three datasets, thereby adopting the default parameter settings listed in Table 1. The results are analysed in terms of three quantitative metrics grounded in the language game paradigm: *degree of communicative success*, *degree of linguistic coherence* and *average linguistic inventory size*. The degree of communicative success reflects how successful a population of agents is at solving the game task. It is computed as the average outcome of the last 1,000 games, where success counts as 1 and failure as 0 (see the *Feedback* step of the game description in Section 2). The degree of linguistic coherence quantifies in how far the different agents in the population would produce the same utterance under the same circumstances, thereby measuring convergence towards a predictable linguistic convention. It is computed for each game as a binary measure that indicates whether the listener agent would have used the same utterance as the one produced by the speaker agent to describe the topic entity, if this agent would have been the speaker. This binary measure is then averaged over the last 1,000 games. Finally, the average linguistic inventory size reflects the number of distinct words that are in ‘active use’. It is calculated as the average number of distinct words uttered by the agents during the last 1,000 games in which they took up the speaker role.

The experimental results obtained on the test portions of the three datasets are listed in Table 2. The table reports the average performance on the three metrics over 10 independent experimental runs, along with a value that indicates the spread of the results in terms of two standard deviations. The results show that the methodology enables a population of agents to converge on a communicatively effective and coherent linguistic convention in each of the task environments, with a degree of communicative success above 99.5% and a degree of linguistic coherence above 87.5%. The average linguistic inventory size of the agents revolves around 50 words in each of the task environments. Examples of words that have emerged during the different experiments are shown in Figure 1. Figure 1a visualises a word with the form “demoxu” that emerged in agent 1 in the CLEVR experiment and was fully entrenched after 1,000,000 games ($s = 1.0$). The concept representation of this word includes three relevant dimensions ($\omega > 0.0$): *area*, *bb-area* and *rel-area*. The values on these dimensions respectively represent, normalised on a scale between 0 and 1, the number of pixels within an entity’s boundaries, the number of pixels within an entity’s rectangular bounding box, and the ratio between an entity’s area and the number of pixels in the entire image. When mapping the *bb-area* and *rel-area* values back to raw pixel counts, we can interpret that the word prototypically refers to entities with an area of 1344 pixels (standard deviation of 76.8 pixels), a bounding box of 1574 pixels (standard deviation of 115 pixels), and covering just under 1% of the image. In human terms, these are objects with a small visible surface that fill a large part, yet not all, of their bounding box. When looking at agent 1’s use of this word throughout the experiment, it is indeed used in 73% of all cases to refer to small spheres. Figure 1b visualises a word with the form “zapose” that emerged in agent 1 in the WINE experiment and was fully entrenched after 1,000,000 games ($s = 1.0$). The concept representation of this word has specialised towards a single relevant dimension ($\omega > 0.0$), namely the amount of residual sugar. When mapping the μ and σ values back to grams per liter, we can interpret that the concept representation prototypically refers to entities with a residual sugar content of 12,34 g/l (standard deviation of 1.39 g/l). In human terms, the concept can thus be used to refer to medium sweet wines. Figure 1c visualises a word with the form “bogezi” that emerged in agent 1 in the CREDIT experiment and was fully entrenched after 1,000,000 games ($s = 1.0$). The concept representation of this word has specialised towards a low value range on a single relevant dimension ($\omega > 0.0$), namely the second PCA component. When it comes to interpretability, the three example words illustrate that the concept representations are interpretable up to the interpretability of the input dimensions. Indeed, if these dimensions are meaningful to humans, for example in the case of visual features, physico-chemical characteristics or other sensor measurements, the resulting concepts are equally human-interpretable. If these features correspond to dimensions that are more difficult to interpret by humans, such as PCA components, a communicatively effective and coherent linguistic convention with transparent concept representations still emerges, but the interpretation difficulty of the input dimensions percolates to the concept representations.

Figure 2 provides more insight into the evolutionary dynamics that take place during the training phase of the CLEVR experiment. The graph shows the degree of communicative success (solid line,

Table 2: Results on the three test sets in terms of communicative success, linguistic coherence and linguistic inventory size. Mean and 2 standard deviations computed over 10 runs.

Dataset	Comm. suc. \uparrow	Ling. coh. \uparrow	Inv. size \downarrow
CLEVR	99.65 \pm 0.13	93.86 \pm 1.09	46.72 \pm 2.45
WINE	99.74 \pm 0.15	88.67 \pm 1.92	52.67 \pm 2.93
CREDIT	99.67 \pm 0.13	87.72 \pm 2.50	51.43 \pm 2.49

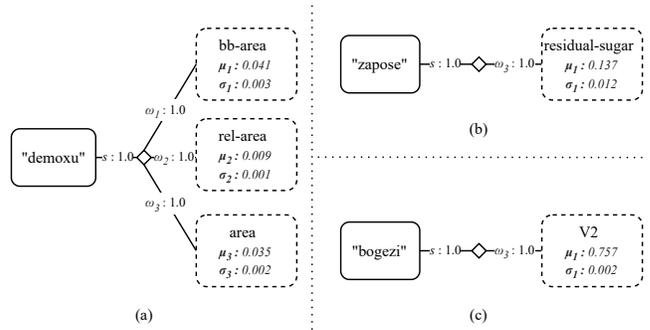


Figure 1: Examples of emerged concepts for the CLEVR (a), WINE (b) and CREDIT (c) datasets.

left y-axis), the degree of linguistic coherence (dotted line, left y-axis) and the average linguistic inventory size (dashed line, right y-axis) as a function of the number of games that are played. The degree of communicative success starts at 0, as all agents start with an empty linguistic inventory. It rises to about 90% after 50,000 games, and continues to grow to over 99.5% over the course of the 1,000,000 games that are played. The degree of linguistic coherence roughly follows the same dynamics as the degree of communicative success, although the growth is much slower. After 1,000,000 games, the degree of linguistic coherence has reached about 90% as it continues to increase. The average linguistic inventory size shows the typical ‘overshoot pattern’ that is found in many language emergence experiments [48]. Indeed, many words emerge during the initial phase of the experiment, as the individual agents are constantly faced with the need to invent. Then, as a result of the rewarding and punishing of words during the alignment phase of each game, the population converges towards a smaller set of words. The graph shows that the peak linguistic inventory size lies around 90 words, while an average of just over 50 words is reached after 1,000,000 games. Note that these numbers are not the same as those reported for the test set in Table 2. Indeed, words that are too specific to adequately describe previously unseen entities are never used by the agents at test time, reducing the number of words in ‘active use’. The same dynamics as those shown in Figure 2 for CLEVR also materialise in the WINE and CREDIT experiments, but could not be included graphically due to space limitations.

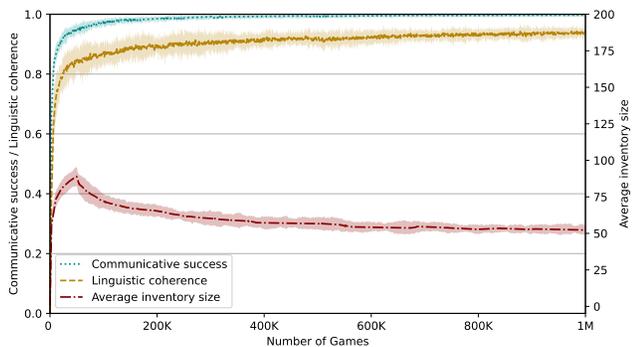


Figure 2: Evolutionary dynamics during the training phase of the CLEVR experiment: degree of communicative success, degree of linguistic coherence and average linguistic inventory size as a function of the number games that are played.

3.2 Compositional generalisability of the emergent concepts

The second experiment assesses the generality of the emergent concepts in terms of their adequacy to refer to entities that exhibit previously unseen attribute combinations, a challenge referred to as *compositional generalisability* [21, 23]. We therefore apply the methodology to a variation on CLEVR that is based on the CLEVR CoGenT dataset [21]. CLEVR CoGenT was especially designed to test the robustness of intelligent systems against correlations that occur at training time but not at test time. As such, a number of biases are included in the scenes by imposing restrictions on the composition of entities. In particular, in the training scenes, all cubes are either grey, blue, brown or yellow, while cylinders are always red, green, purple, or cyan. Test set A contains scenes that are subject to the same correlations. Test set B however consists of scenes that are subject to a different set of correlations, with cubes always being red, green, purple or cyan, and cylinders always being grey, blue, brown or yellow. There are no restrictions on the colour of spheres in either of the splits. Test set A can be used to assess how well a learnt model performs in a standard machine learning setting, in which the training and test sets are drawn from the same distribution. Test set B can be used to assess whether the learnt model generalises beyond the correlations that characterise the training set. For the purposes of this experiment, we built a training set and two test sets using the CLEVR CoGenT images through the same two-stage process as the one that was used for creating the CLEVR, WINE and CREDIT datasets. The results of this experiment, which are provided in Table 3, show that the performance of the agents on test set A and test set B is very similar in terms of degree of communicative success (99.63% vs. 99.62%), degree of linguistic coherence (93.50% vs. 93.51%) and average linguistic inventory size (47.63 words vs. 47.58 words). The compositional generalisability experiment thereby confirms that the emergent linguistic convention does not break down when faced with the need to refer to entities that instantiate previously unseen attribute combinations.

Table 3: Results of the compositional generalisability experiments, showing a similar performance in both conditions.

Dataset	Comm. suc. \uparrow	Ling. coh. \uparrow	Inv. size \downarrow
CoGenT A	99.63 \pm 0.12	93.50 \pm 1.96	47.63 \pm 3.40
CoGenT B	99.62 \pm 0.12	93.51 \pm 1.83	47.58 \pm 3.59

Table 4: Results of the experiments that validate the applicability of the methodology to heteromorphic populations.

Condition	Comm. suc. \uparrow	Ling. coh. \uparrow	Inv. size \downarrow
HOM-19	99.66 \pm 0.12	93.75 \pm 1.26	46.34 \pm 3.62
HET-19	98.47 \pm 1.33	89.73 \pm 3.62	48.25 \pm 2.68
HOM-10	99.60 \pm 0.42	92.82 \pm 2.93	47.33 \pm 6.35
HET-10	85.55 \pm 9.54	59.00 \pm 14.54	52.68 \pm 8.23

3.3 Applicability to heteromorphic populations

The third experiment assesses the applicability of the methodology to heteromorphic populations, in our case populations in which not all agents are equipped with the same combination of sensors. For this purpose, we set up a variation on the CLEVR experiment in which each individual agent has access to a randomly selected subset of the 20 dimensions that are provided by the dataset. This means in practice that almost all games are played by two agents that do not perceive the same entities through the same dimensions. Concretely, we run two instances of the experiment in which the agents are respectively endowed with combinations of 19 and 10 randomly selected sensors (HET-19 and HET-10). In order to establish a meaningful basis for comparison, we also run a version of the experiment with homomorphic populations in which the agents are endowed with the same number of sensors (HOM-19 and HOM-10). In the homomorphic setting, the sensor combination is randomly selected for each agent at the beginning of each experimental run and remains constant throughout the experiment.

The test results of the experiment are listed in Table 4. When moving from the homomorphic to the heteromorphic setting, the degree of communicative success decreases from 99.66% to 98.47% with 19 out of 20 sensors available and from 99.60% to 85.55% with only 10 out of 20 sensors available. The degree of linguistic coherence drops to a larger extent, from 93.75% to 89.73% and from 92.82% to 59.00%. At the same time, the average linguistic inventory size increases from 46.34 to 48.25 words and from 47.33 to 52.68 words. The experiment thereby confirms that a high degree of communicative success can still be reached even if agents are equipped with very different combinations of sensors. Unsurprisingly, there is more variation in the words that are used by the agents in the heteromorphic setting, as agents will tend to use words that optimally fit their own sensory apparatus. This increased variation is reflected by the observed drop in degree of linguistic coherence and rise in average linguistic inventory size.

3.4 Robustness against sensor defects

The fourth experiment validates the robustness of the methodology against sensor defects that occur in individual agents. For this

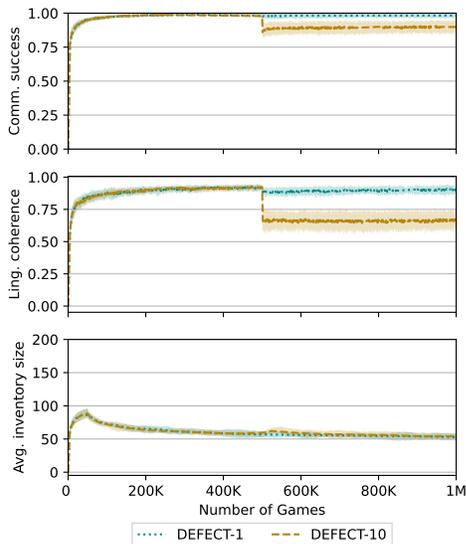


Figure 3: Evolutionary dynamics during the training phase of the CLEVR experiment in which each agent loses access to 1 or 10 sensors after 500,000 games.

purpose, we run a version of the CLEVR experiment in which the agents suffer from a sudden malfunction after 500,000 games. To simulate this malfunction, all agents lose access to a predefined number of sensors, which are randomly selected for each individual agent. The dynamics of the experiment are visualised in Figure 3 for experimental conditions in which the agents lose access to respectively 1 and 10 of their 20 sensors (DEF-1 and DEF-10). As is to be expected, the degrees of communicative success and linguistic coherence drop at the moment of the malfunction. As the linguistic convention adapts to the new circumstances, we observe a temporary rise in the average linguistic inventory size and a partial recovery of the degrees of communicative success and linguistic coherence.

The results on the test set are provided for both conditions in Table 5 along with the results of the HET-19 and HET-10 experiments as a basis for comparison. The degree of communicative success amounts to 98.31% in the setting where one sensor malfunctions and to 89.95% in the setting with 10 malfunctioning sensors. The degree of linguistic coherence amounts to 90.15% and 66.61% respectively, while the average number of words in use amounts to 47.05 and 47.72 respectively. Note that the experimental conditions after the malfunction correspond in fact to those of the experiments with heteromorphic populations reported on in Section 3.3. When comparing both results, we can see that the performance after the malfunction is still better in terms of all three metrics than the performance achieved in the experiments where the agents never had access to all sensors. The experiment thereby demonstrates on the one hand that the methodology is robust against extensive sensor defects in individual agents, and on the other hand that the emergence of an effective linguistic convention before a malfunction can remain beneficial even in the long term.

Table 5: Results of the experiments that validate the robustness of the methodology against sensor defects in individual agents.

Condition	Comm. suc. \uparrow	Ling. coh. \uparrow	Inv. size \downarrow
DEF-1	98.31 \pm 1.67	90.15 \pm 2.99	47.05 \pm 3.82
HET-19	98.47 \pm 1.33	89.73 \pm 3.62	48.25 \pm 2.68
DEF-10	89.95 \pm 4.28	66.61 \pm 7.55	47.72 \pm 3.20
HET-10	85.55 \pm 9.54	59.00 \pm 14.54	52.68 \pm 8.23

Table 6: Results of the experiments that assess the robustness of the methodology against differences in perception.

Condition	Comm. suc. \uparrow	Ling. coh. \uparrow	Inv. size \downarrow
CLEVR	99.65 \pm 0.13	93.86 \pm 1.09	46.72 \pm 2.45
SHIFT-0.1	99.62 \pm 0.13	93.48 \pm 2.02	46.82 \pm 1.60
SHIFT-1	99.61 \pm 0.14	93.73 \pm 1.44	46.16 \pm 4.01
NOISE-0.1	98.40 \pm 0.63	82.65 \pm 3.07	47.72 \pm 2.28
NOISE-1	87.04 \pm 2.44	46.58 \pm 1.27	49.00 \pm 6.38

3.5 Robustness against differences in perception

The fifth experiment assesses the robustness of the methodology against differences in the agents’ perception of the world, which corresponds in our experiments to the way in which the perceived feature vectors \mathbf{x}^S and \mathbf{x}^L are computed from an entity’s ‘objective’ feature vector \mathbf{x} (see *World* in Section 2). Concretely, we simulate two different scenarios. In the first scenario, the agents record different sensor values because of a lack of calibration. This is simulated by shifting \mathbf{x}^S and \mathbf{x}^L with respect to \mathbf{x} by a value that is individually set for each sensor of each agent at the beginning of each experimental run. These values are sampled from a normal distribution with a mean of 0 and a standard deviation of either 0.1 (SHIFT-0.1), simulating slight calibration differences, or 1.0 (SHIFT-1), simulating substantial calibration differences. In the second scenario, the sensor values recorded by the agents are subject to noise. This is simulated by shifting \mathbf{x}^S and \mathbf{x}^L with respect to \mathbf{x} by a value that is independently sampled for each sensor of each participating agent at the start of each game from normal distributions with a mean of 0 and a standard deviation of 0.1 (NOISE-0.1) or 1.0 (NOISE-1).

The results of the perceptual difference experiment are provided in Table 6 in comparison to the original CLEVR experiment. We can see that a lack of calibration has no significant effect on the experimental results. The presence of sensor noise leads to a non-catastrophic decrease in degree of communicative success (from 99.65% to 98.40% and 87.04%). The decrease in degree of linguistic coherence is more substantial (from 93.86% to 82.65% and 46.58%) and is accompanied by a slight increase in the average linguistic inventory size (from 46.72 to 47.72 and 49.00). The experiment thereby shows that the methodology does not break down when faced with agents that perceive the world differently. It also confirms the trend observed in the previous experiments that more challenging experimental conditions lead to more variation in language use, while remarkable degrees of communicative success can still be achieved.

Table 7: Results of the continual learning experiment.

Condition	Comm. suc. \uparrow	Ling. coh. \uparrow	Inv. size \downarrow
CLEVR	99.61 \pm 0.11	93.75 \pm 1.85	46.41 \pm 2.67
CLEVR-WINE	99.72 \pm 0.20	88.13 \pm 2.34	52.40 \pm 2.49
CLEVR-CONT	99.60 \pm 0.14	93.72 \pm 1.83	46.36 \pm 2.34

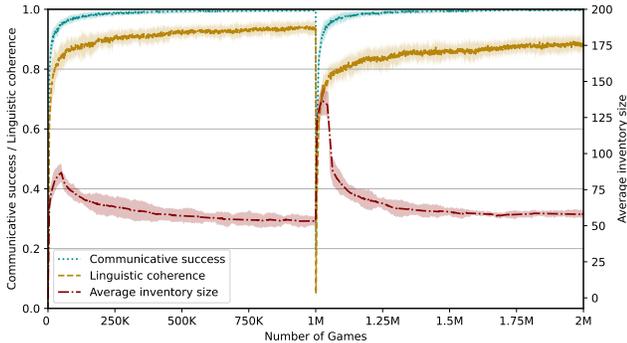


Figure 4: Evolutionary dynamics during the continual learning experiment, in which a model is first trained on CLEVR and then on WINE.

3.6 Adequacy for continual learning

The final experiment assesses the adequacy of the methodology for continual learning, focussing in particular on its robustness against catastrophic forgetting [16, 28]. For the purposes of this experiment, we (i) train a model on the CLEVR training set, (ii) evaluate it on the CLEVR test set, (iii) continue to train it on the WINE training set, (iv) evaluate it on the WINE test set, and (v) evaluate it again on the CLEVR test set. The dynamics of the experiment are shown in Figure 4. Unsurprisingly, the first 1,000,000 games exhibit the same dynamics as those observed for the original CLEVR experiment shown in Figure 2. When moving to WINE after 1,000,000 games, the degrees of communicative success and linguistic coherence drop to 0 but rapidly increase again as if a new experiment would have started. The average linguistic inventory size quickly rises as many new words are invented to accommodate the WINE dataset. Then, this number gradually decreases as the words that are not adequate for describing wine samples cease to be used by the agents.

The results of the experiment are shown in Table 7. After training on CLEVR, a degree of communicative success of 99.61% is obtained on the CLEVR test set, along with a degree of linguistic coherence of 93.75% and an average linguistic inventory size of 46.41 (CLEVR). After continuing to train on WINE, the agents achieve a degree of communicative success of 99.72% on the WINE test set, along with a degree of linguistic coherence of 88.13% and an average linguistic inventory size of 52.40 (CLEVR-WINE). These numbers indeed match those recorded in the original CLEVR and WINE experiments. When evaluating the model on the CLEVR test set after training first on CLEVR and then on WINE, the results that are obtained do not deviate significantly from the results obtained before training on wine (CLEVR-CONT). The experimental results thereby confirm that the methodology is adequate for continual learning and is not susceptible to catastrophic forgetting.

4 DISCUSSION AND CONCLUSION

This paper has introduced a methodology through which a communicatively effective, robust and adaptive linguistic convention can emerge in a population of autonomous agents. The linguistic convention emerges in a decentralised manner through local, task-oriented and situated communicative interactions that take place between pairs of agents drawn from the population. The linguistic convention takes the form of symbolic labels associated to concept representations that are grounded in a multi-dimensional, continuous feature space. These form-meaning associations are individually constructed by each agent and are shaped by their past successes and failures in communication. The methodology embodies the evolutionary dynamics of the language game paradigm [30, 40, 41, 48] and integrates an innovative way in which agents represent, invent, adopt and align concept representations.

Along with a formal definition of the methodology, we have presented a range of experiments that serve as its initial validation and which demonstrate the desirable properties of the emergent *artificial natural languages*. As such, we have first applied the methodology to three datasets that contain very different types of data, ranging from visual scenes over physicochemical analyses to principal components extracted from financial transaction records. Yielding a communicatively effective, coherent and transparent linguistic convention in all three cases, the experiment shows that the effectiveness of the methodology is not limited to a particular domain or data type. Then, we reported on two experiments that confirm that the methodology is capable of compositional generalisation and that it remains effective when applied to heteromorphic populations. The fourth and fifth experiments demonstrate the robustness of the methodology against sensor defects and noisy observations, including those resulting from a lack of calibration. The final experiment validates the adequacy of the methodology for continual learning, focussing in particular on its resilience against catastrophic forgetting.

The research reported on in this paper constitutes a novel contribution to the state of the art as it lifts three consequential limitations that were never successfully overcome together in prior work. The first limitation concerns the emergent nature of the conceptual distinctions. Most prior approaches learn to ground a predefined set of concepts [31, 39, 51]. These concepts are symbolically annotated in training data and correspond to distinctions that occur in an existing natural language, typically English. As these concepts have emerged and evolved to fit the communicative needs and physical endowment of a community of human language users, they do not necessarily fit well the sensors and communicative tasks of artificial agents [48]. The second limitation concerns the circumstances under which the languages emerge and evolve. These are often too far removed from those under which human languages emerge and evolve to bring about evolutionary processes that yield emergent languages with the same desirable properties. In particular, populations sometimes consist of two agents only [5, 18, 32], agents can either speak or listen, but not both, [7, 8, 24, 29], or learning is not decentralised [15, 22]. Finally, prior approaches that are not subject to the first two limitations are limited in their applicability, as they have not been generalised beyond the emergence of naming

conventions [26, 44, 45], to continuous feature spaces [53, 54], or to arbitrary combinations of feature channels [2, 3, 37, 38, 42, 49, 50]. By lifting these three limitations at the same time, the methodology introduced in this paper provides a model of how human-like linguistic conventions can emerge and evolve in populations of autonomous agents, which is, importantly, directly applicable to any dataset that situates entities in a continuous feature space.

REFERENCES

- [1] Clay Beckner, Richard Blythe, Joan Bybee, Morten H. Christiansen, William Croft, Nick C. Ellis, John Holland, Jinyun Ke, Diane Larsen-Freeman, and Tom Schoenemann. 2009. Language is a complex adaptive system: Position paper. *Language learning* 59 (2009), 1–26. <https://doi.org/10.1111/j.1467-9922.2009.00533.x>
- [2] Tony Belpaeme and Joris Bleys. 2005. Explaining universal color categories through a constrained acquisition process. *Adaptive Behavior* 13, 4 (2005), 293–310.
- [3] Joris Bleys. 2016. *Language strategies for the domain of colour*. Language Science Press, Berlin, Germany.
- [4] Ben Bogin, Mor Geva, and Jonathan Berant. 2018. Emergence of communication in an interactive world with consistent speakers. In *Emergent Communication Workshop: NeurIPS 2018*.
- [5] Diane Bouchacourt and Marco Baroni. 2018. How agents see things: On visual representations in an emergent language game. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, Ellen Riloff, David Chiang, Julia Hockenmaier, and Jun'ichi Tsujii (Eds.). Association for Computational Linguistics, 981–985. <https://doi.org/10.18653/v1/D18-1119>
- [6] Kris Cao, Angeliki Lazaridou, Marc Lanctot, Joel Z. Leibo, Karl Tuyls, and Stephen Clark. 2018. Emergent communication through negotiation. In *6th International Conference on Learning Representations (ICLR 2018)*. 1–15.
- [7] Rahma Chaabouni, Eugene Kharitonov, Emmanuel Dupoux, and Marco Baroni. 2021. Communicating artificial neural networks develop efficient color-naming systems. *Proceedings of the National Academy of Sciences* 118, 12 (2021), e2016569118. <https://doi.org/10.1073/pnas.2016569118>
- [8] Rahma Chaabouni, Florian Strub, Florent Althé, Eugene Tarassov, Corentin Tallec, Elnaz Davoodi, Kory Wallace Mathewson, Olivier Tieleman, Angeliki Lazaridou, and Bilal Piot. 2022. Emergent communication at scale. In *10th International Conference on Learning Representations (ICLR 2022)*. 1–30.
- [9] Paulo Cortez, Antonio Cerdeira, Fernando Almeida, Telmo Matos, and José Reis. 2009. Modeling wine preferences by data mining from physicochemical properties. *Decision Support Systems* 47, 4 (2009), 547–553. <https://doi.org/10.1016/j.dss.2009.05.016>
- [10] Andrea Dal Pozzolo, Olivier Caelen, Yann-Aël Le Borgne, Serge Waterschoot, and Gianluca Bontempi. 2014. Learned lessons in credit card fraud detection from a practitioner perspective. *Expert Systems with Applications* 41, 10 (2014), 4915–4928. <https://doi.org/10.1016/j.eswa.2014.02.026>
- [11] Charles R. Darwin. 1871. *The descent of man, and selection in relation to sex* (1st ed.). Vol. 1. John Murray, London, United Kingdom.
- [12] Abhishek Das, Satwik Kottur, José M. F. Moura, Stefan Lee, and Dhruv Batra. 2017. Learning cooperative visual dialog agents with deep reinforcement learning. In *2017 IEEE International Conference on Computer Vision (ICCV)*, Rita Cucchiara, Yasuyuki Matsushita, Nicu Sebe, and Stefano Soatto (Eds.). IEEE Computer Society, Washington, D.C., USA, 2951–2960.
- [13] Jonas Doumen, Katrien Beuls, and Paul Van Eecke. 2023. Modelling Language Acquisition through Syntactico-Semantic Pattern Finding. In *Findings of the Association for Computational Linguistics: EAACL 2023*, Andreas Vlachos and Isabelle Augenstein (Eds.). Association for Computational Linguistics, 1317–1327.
- [14] Gerald Echterhoff. 2013. The role of action in verbal communication and shared reality. *Behavioral and Brain Sciences* 36, 4 (2013), 354–355. <https://doi.org/10.1017/S0140525X12002567>
- [15] Jakob Foerster, Ioannis Alexandros Assael, Nando de Freitas, and Shimon Whiteson. 2016. Learning to communicate with deep multi-agent reinforcement learning. In *Advances in Neural Information Processing Systems 29 (NIPS 2016)*, Daniel Lee, Masashi Sugiyama, Ulrike Von Luxburg, Isabelle Guyon, and Roman Garnett (Eds.). Curran Associates Inc., Red Hook, NY, USA, 2137–2145.
- [16] Robert M. French. 1999. Catastrophic forgetting in connectionist networks. *Trends in Cognitive Sciences* 3, 4 (1999), 128–135. [https://doi.org/10.1016/S1364-6613\(99\)01294-2](https://doi.org/10.1016/S1364-6613(99)01294-2)
- [17] Paul Grice. 1967. Logic and Conversation. In *Studies in the Way of Words*, Paul Grice (Ed.). Harvard University Press, Cambridge, MA, USA, 41–58.
- [18] Serhii Havrylov and Ivan Titov. 2017. Emergence of language with multi-agent games: Learning to communicate with sequences of symbols. In *Advances in Neural Information Processing Systems 30 (NIPS 2017)*, Isabelle Guyon, Ulrike Von Luxburg, Samy Bengio, Hanna Wallach, Rob Fergus, S. V. N. Vishwanathan, and Roman Garnett (Eds.). Curran Associates Inc., Red Hook, NY, USA, 2146–2156.
- [19] Ernst Hellinger. 1909. Neue Begründung der Theorie quadratischer Formen von unendlichvielen Veränderlichen. *Journal für die reine und angewandte Mathematik* 1909, 136 (1909), 210–271. <https://doi.org/10.1515/crll.1909.136.210>
- [20] Francis Heylighen. 2001. The science of self-organization and adaptivity. In *Knowledge management, organizational intelligence and learning, and complexity. The encyclopedia of life support systems*, Lowell Douglas Kiel (Ed.). EOLSS Publishers, Oxford, United Kingdom, 253–280.
- [21] Justin Johnson, Bharath Hariharan, Laurens van der Maaten, Li Fei-Fei, C. Lawrence Zitnick, and Ross Girshick. 2017. CLEVR: A Diagnostic Dataset for Compositional Language and Elementary Visual Reasoning. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE Computer Society, Washington, D.C., USA, 2901–2910.
- [22] Jooyeon Kim and Alice Oh. 2021. Emergent Communication under Varying Sizes and Connectivities. In *Advances in Neural Information Processing Systems 34 (NeurIPS 2021)*, Marc'Aurelio Ranzato, Alina Beygelzimer, Yann Dauphin, Percy S. Liang, and Jennifer W. Vaughan (Eds.). Curran Associates Inc., Red Hook, NY, USA, 17579–17591.
- [23] Najoung Kim and Tal Linzen. 2020. COGS: A Compositional Generalization Challenge Based on Semantic Interpretation. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, Bonnie Webber, Trevor Cohn, Yulan He, and Yang Liu (Eds.). Association for Computational Linguistics, 9087–9105. <https://doi.org/10.18653/v1/2020.emnlp-main.731>
- [24] Satwik Kottur, José Moura, Stefan Lee, and Dhruv Batra. 2017. Natural Language Does Not Emerge 'Naturally' in Multi-Agent Dialog. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, Martha Palmer, Rebecca Hwa, and Sebastian Riedel (Eds.). Association for Computational Linguistics, 2962–2967. <https://doi.org/10.18653/v1/D17-1321>
- [25] Angeliki Lazaridou, Alexander Peysakhovich, and Marco Baroni. 2017. Multi-agent cooperation and the emergence of (natural) language. In *5th International Conference on Learning Representations (ICLR 2017)*. 1–11.
- [26] Martin Loetzsch. 2015. *Lexicon formation in autonomous robots*. Ph.D. Dissertation. Humboldt-Universität zu Berlin, Berlin, Germany.
- [27] John Maynard Smith and Eörs Szathmáry. 1999. *The origins of life: From the birth of life to the origin of language*. Oxford University Press, Oxford, United Kingdom.
- [28] Michael McCloskey and Neal J. Cohen. 1989. Catastrophic Interference in Connectionist Networks: The Sequential Learning Problem. In *Psychology of Learning and Motivation* 24, Gordon H. Bower (Ed.). Academic Press, 109–165. [https://doi.org/10.1016/S0079-7421\(08\)60536-8](https://doi.org/10.1016/S0079-7421(08)60536-8)
- [29] Igor Mordatch and Pieter Abbeel. 2018. Emergence of grounded compositional language in multi-agent populations. In *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence*, Sheila McIlraith and Kilian Q. Weinberger (Eds.). AAAI Press, Washington, D.C., USA, 1495–1502.
- [30] Jens Nevens, Paul Van Eecke, and Katrien Beuls. 2019. A Practical Guide to Studying Emergent Communication through Grounded Language Games. In *AISB 2019 Symposium on Language Learning for Artificial Agents*. AISB, 1–8.
- [31] Jens Nevens, Paul Van Eecke, and Katrien Beuls. 2020. From continuous observations to symbolic concepts: A discrimination-based strategy for grounded concept learning. *Frontiers in Robotics and AI* 7, 84 (2020). <https://doi.org/10.3389/frobt.2020.00084>
- [32] Michael Noukhovitch, Travis LaCroix, Angeliki Lazaridou, and Aaron Courville. 2021. Emergent Communication under Competition. In *Proceedings of the 20th International Conference on Autonomous Agents and Multi-Agent Systems*. 974–982.
- [33] Pierre-Yves Oudeyer and Frédéric Kaplan. 2007. Language evolution as a Darwinian process: Computational studies. *Cognitive Processing* 8, 1 (2007), 21–35. <https://doi.org/10.1007/s10339-006-0158-3>
- [34] Rolf Pfeifer, Max Lungarella, and Fumiya Iida. 2007. Self-organization, Embodiment, and Biologically Inspired Robotics. *Science* 318, 5853 (2007), 1088–1093. <https://doi.org/10.1126/science.1145803>
- [35] Cinjun Resnick, Ilya Kulikov, Kyunghyun Cho, and Jason Weston. 2017. Vehicle Communication Strategies for Simulated Highway Driving. In *Emergent Communication Workshop: NeurIPS 2017*.
- [36] August Schleicher. 1869. *Darwinism tested by the science of language. English translation of Schleicher 1863, translated by Alex V. W. Bikkers*. John Camden Hotten, London, United Kingdom.
- [37] Michael Spranger. 2013. Grounded lexicon acquisition - Case studies in spatial language. In *Proceedings of the 2013 IEEE Third Joint International Conference on Development and Learning and Epigenetic Robotics (ICDL)*. IEEE, 1–6.
- [38] Michael Spranger. 2016. *The evolution of grounded spatial language*. Language Science Press, Berlin, Germany. <https://doi.org/10.17169/langsci.b53.183>
- [39] Michael Spranger and Katrien Beuls. 2016. Referential uncertainty and word learning in high-dimensional, continuous meaning spaces. In *Proceedings of the 2016 Joint IEEE International Conference on Development and Learning and Epigenetic Robotics (ICDL-EpiRob)*. IEEE, 95–100.
- [40] Luc Steels. 1995. A self-organizing spatial vocabulary. *Artificial Life* 2, 3 (1995), 319–332. <https://doi.org/10.1162/artl.1995.2.3.319>

- [41] Luc Steels. 2003. The evolution of communication systems by adaptive agents. In *Symposium on Adaptive Agents and Multi-Agent Systems*, Eduardo Alonso, Daniel Kudenko, and Dimitar Kazakov (Eds.), 125–140. https://doi.org/10.1007/3-540-44826-8_8
- [42] Luc Steels. 2015. *The Talking Heads experiment: Origins of words and meanings*. Language Science Press, Berlin, Germany.
- [43] Luc Steels and Tony Belpaeme. 2005. Coordinating perceptually grounded categories through language: A case study for colour. *Behavioral and Brain Sciences* 28, 4 (2005), 469–489. <https://doi.org/10.1017/S0140525X05000087>
- [44] Luc Steels and Martin Loetzsch. 2012. The grounded naming game. In *Experiments in Cultural Language Evolution*, Luc Steels (Ed.), Vol. 3. John Benjamins, Amsterdam, Netherlands, 41–59. <https://doi.org/10.1075/ais.3.04ste>
- [45] Luc Steels, Martin Loetzsch, and Michael Spranger. 2016. A boy named Sue: The semiotic dynamics of naming and identity. *Belgian Journal of Linguistics* 30, 1 (2016), 147–169. <https://doi.org/10.1075/bjl.30.07ste>
- [46] Luc Steels and Eörs Szathmáry. 2018. The evolutionary dynamics of language. *Biosystems* 164 (2018), 128–137. <https://doi.org/10.1016/j.biosystems.2017.11.003>
- [47] Sainbayar Sukhbaatar, Arthur Szlam, and Rob Fergus. 2016. Learning multi-agent communication with backpropagation. In *Advances in Neural Information Processing Systems 29 (NIPS 2016)*, Daniel Lee, Masashi Sugiyama, Ulrike Von Luxburg, Isabelle Guyon, and Roman Garnett (Eds.). Curran Associates Inc., Red Hook, NY, USA, 2244–2252.
- [48] Paul Van Eecke, Katrien Beuls, Jérôme Botoko Ekila, and Roxana Rădulescu. 2022. Language games meet multi-agent reinforcement learning: A case study for the naming game. *Journal of Language Evolution* 7, 2 (2022), 213–223. <https://doi.org/10.1093/jole/lzad001>
- [49] Paul Vogt. 2005. The emergence of compositional structures in perceptually grounded language games. *Artificial intelligence* 167, 1–2 (2005), 206–242. <https://doi.org/10.1016/j.artint.2005.04.010>
- [50] Paul Vogt. 2015. *How mobile robots can self-organise a vocabulary*. Language Science Press, Berlin, Germany. <https://doi.org/10.17169/langsci.b50.113>
- [51] Sida I. Wang, Percy Liang, and Christopher D. Manning. 2016. Learning Language Games through Interaction. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, Katrin Erk and Noah A. Smith (Eds.). Association for Computational Linguistics, 2368–2378. <https://doi.org/10.18653/v1/P16-1224>
- [52] Barry Payne Welford. 1962. Note on a method for calculating corrected sums of squares and products. *Technometrics* 4, 3 (1962), 419–420.
- [53] Pieter Wellens. 2012. *Adaptive Strategies in the Emergence of Lexical Systems*. Ph.D. Dissertation. Vrije Universiteit Brussel, Brussels: VUB Press.
- [54] Pieter Wellens, Martin Loetzsch, and Luc Steels. 2008. Flexible word meaning in embodied agents. *Connection Science* 20, 2–3 (2008), 173–191. <https://doi.org/10.1080/09540090802091966>