

Language Development among Co-learning Agents

Viktor Gyenes

Eötvös Loránd University

Pázmány P. sétány 1/C, Budapest, Hungary, H-1117

gyenesvi@inf.elte.hu

András Lőrincz

Eötvös Loránd University

Pázmány P. sétány 1/C, Budapest, Hungary, H-1117

andras.lorincz@elte.hu

Abstract—We investigate the properties of coupled co-learning systems during the emergence of communication. Co-learning systems are more complex than individual learning systems because of being dependent on the learning process of each other, thus risking divergence. We developed a neural network approach and implemented a concept that we call reconstruction principle, which we found adequate for overcoming the instability problem. Experimental simulations were performed to test the emergence of both compositional and holistic communication. The results show that compositional communication is favorable when learning performance is considered, however it is more error-prone to differences in the conceptual representations of the individual systems. We show that our architecture enables the adjustment of the differences in the individual conceptual representations in case of compositional communication.

Index Terms—Language emergence, co-learning, reconstruction network, compositional communication.

I. INTRODUCTION

The emergence and evolution of communication has gained significant research attention in the past decade. Multi-agent systems provide a popular framework for computer simulations aiming to model the coordinated development of natural languages. The inherent property of the development of a coordinated communication system that *multiple agents* participate in it, poses extra difficulties to the algorithms aiming to model it.

When communication evolves, it should be the result of a *negotiation process* between many parties. During this process, certain new items are invented by individuals and accepted and learned by the others. Who invents things and who accepts them should neither be predefined nor one-sided. All agents take part in both of these tasks, that is, they teach and learn simultaneously. To let the whole process converge to a useful communication system, agents have to adapt to

This material is based upon work supported partially by the European Office of Aerospace Research and Development, Air Force Office of Scientific Research, Air Force Research Laboratory, under Contract No. FA-073029. This research has also been supported by an EC FET grant, the ‘New Ties project’ under contract 003752. Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the European Office of Aerospace Research and Development, Air Force Office of Scientific Research, Air Force Research Laboratory, the EC, or other members of the EC New Ties project.

each other, not only to the task to be learned; their learning depends on that of the others. The complexity of the problem is that learning concerns hidden variables different for each agent while learning is inherently coupled.

A great deal of work utilizes perceptually grounded agents both to model language development or acquisition (see e.g. [4], [11], [18], [16]) and language emergence motivated by natural language evolution [14], [9], [17], [15], [13]. Here, we consider language emergence, i.e. the construction of a new lexicon, but take a broader view: we consider the optimization of information transfer among the agents as a process of negotiation about a ‘language’. The conditions when language may emerge concern *environmental* conditions [3] and *learning* conditions. We consider the latter. Our work aims to enable the modelling of more realistic situations when both sets of conditions are important. Here, one of the motivations is to encapsulate the difficulties of *simultaneous learning* for agents that have different conceptual representations.

Here, we address this problem. We model the learning systems with the so called reconstruction networks. We provide a neural implementation of what we call *reconstruction principle*. We argue that it is efficient for making co-learning stable.

II. RELATED WORK

When modelling natural language evolution, it is a natural idea to involve knowledge transfer from generation to generation, like in the Iterated Learning Model of Kirby and colleagues [14], [9]: the new generation of language users learns the language from the previous generation and then the old generation is replaced by the new one. An interesting conclusion of the model is that the compositional nature of language might be the result of the learning bottleneck imposed when language has to pass from one generation to the other. Vogt [17] also builds on the Iterated Learning Model and combines it with language games [15] to model the emergence of compositional languages when agents aim to communicate about their observations. An interesting aspect of this work is that it deals with the conceptual representations of the agents upon which they build their language, which is closely related to the symbol grounding problem [5], and also the compositional nature of language. Smith [13] also considers the development of individual,

distinct meaning structures and examines its effect on the evolved language. All of these models apply learning from generation to generation, and thus teacher agents are fixed, so parallel learning does not occur. This way these models, in effect, avoid the problem of co-learning.

Cangelosi [2] uses artificial neural networks trained by a genetic algorithm to develop a language in an agent system that aims to differentiate between edible and poisonous food items and emphasizes that the evolution of language requires the parallel evolution of the ability of language understanding and production. He also considers the parallel development of input categorization and language. Hutchins and Hazlehurst aim to invent a shared lexicon [7] utilizing feedforward connectionist networks that model language learning agents.

The work of Oliphant and Batali [10] is very close to ours in the reconstruction aspect. They model the development of a stable coordinated communication system. They use a method that they call the ‘obverter’ procedure in which agents observe each other and try to maximize their chances to communicate successfully, instead of simply imitating the others. They provide mathematical considerations about the convergence of their method. The underlying idea is very similar to generative or reconstruction networks [1], which – in the field of vision – would claim that vision is inverse graphics [6].

Our architecture can be seen as a neural network implementation of the ‘obverter’ learner that also generalizes it for *compositional* internal representations and communication. Up to our best knowledge, no neural network approach has incorporated this idea, only ‘imitator’ approaches exist. Central to our methodology is the idea of Cangelosi that production and understanding must be maintained in parallel. Our framework enables the learners to have distinct conceptual representations. We investigate the properties of both compositional and non-compositional (holistic) communication systems. By compositionality, here we mean that an utterance aiming to express a thought composed of a set of concepts is composed of signals expressing the concepts. We treat the problem of co-learning, and restrict our methods to local Hebbian learning for the individual systems.

III. METHODS

This section details our network architecture and the learning methods applied. The general context of the learning is a signalling game, in which the networks observe inputs and the learning task is to co-develop a language (agree on a set of signals) to communicate the observations. This helps investigating communication related issues without being affected by other environmental factors, and gives us freedom to vary related parameters and test various settings.

A. Network Architecture

We model the agents by three-layer neural networks, the architecture is depicted in Fig. 1: the input layer of

the network receives the observation ($x \in R^m$), which is processed and an internal representation ($h \in \{0,1\}^n$) is formed. In our model, this transformation (G) represents the extraction of features (concepts describing the input), resulting in a *combinatorial* internal representation, whose components indicate the presence/absence of the features. The $x \rightarrow h$ transformation is modelled as follows. Each input x (belonging to a finite set for now) is assigned a vector $f_G(x) \in [0,1]^n$ of real values as if indicating the degree (or probability) of the presence of various features. The internal representation $h \in \{0,1\}^n$ is generated from $f_G(x)$ by rounding to 0 or 1. We have chosen this simple representation because our intention was to model the adjustment of concept formation; we do this by letting the values in $f_G(x)$ be tuned. A particular property of this characterization is that inputs are categorized into multiple categories: each input has a certain set of features, each feature representing being the member of a category (see Fig. 1).

Two other transformations govern the communication related behavior of the network. The network can generate an ‘utterance’ u from its internal representation by means of transformation Q . We let $u \in \{0,1\}^{2^n}$, so that the utterance may contain combinatorially many signals. Furthermore, the network also has another transformation, W , that we call ‘parsing’ or ‘understanding an utterance’, because it yields some internal representation based on the utterance received.

In our studies, we used two methods for generation and parsing. The first method performs the linear transformations Q and W that are followed by a nonlinearity σ rounding the values to 0 or 1: $u = \sigma(Qh)$, $h = \sigma(Wu)$. The other method implements the so called *reconstruction principle*. In this case, the network generates an utterance u_g for a given internal representation h such that when it is parsed (by the network itself) the resulting internal representation is closest to the original vector h . That is, the network tries to reconstruct the internal representation from its own intended output, and chooses an utterance that reconstructs the internal representation the best. The same principle is used for parsing, except the roles are changed: given an utterance u , the network chooses an internal state h , that when transformed back to an utterance, yields an utterance closest to u . Below, this idea is formalized for the generation of utterance u_g from internal representation h , and for creating an approximate \tilde{h} via parsing utterance u , respectively:

$$u_g = \operatorname{argmin}_u \| h - \sigma(Wu) \|_{L_2}^2 \quad (1)$$

$$\tilde{h} = \operatorname{argmin}_h \| u_g - \sigma(Qh) \|_{L_2}^2 \quad (2)$$

B. Optimization Method

The minimization tasks (1) and (2) are combinatorial optimization problems, since we restricted the vectors h and u to have only 0 – 1 components. To solve these problem, we use

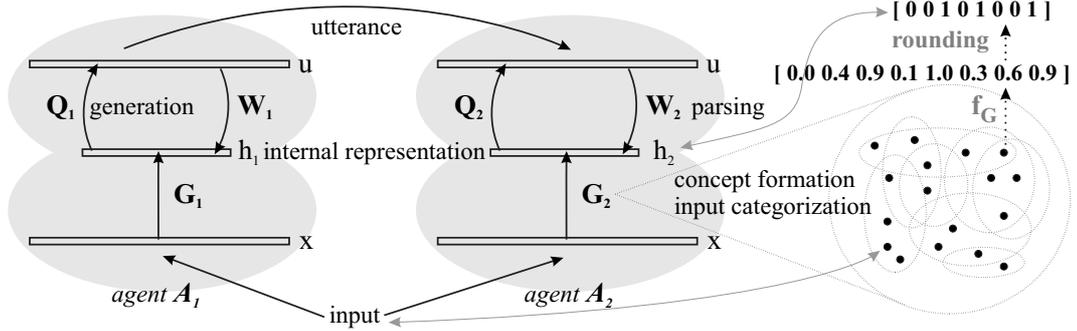


Fig. 1. **Network architecture.** Two agents, A_1 and A_2 , structure of the input set and the transformation to internal representation. For details, see text.

the cross-entropy (CE) method [12], which is a generic approach to combinatorial optimization. The CE method maintains a parameterized probability distribution, from which it iteratively randomly generates solution samples, evaluates them according to the cost function, and continuously updates the probability distribution, until convergence.

The CE algorithm nicely fits into the reconstruction network frame. We used the multi-dimensional Bernoulli distribution as the probability density function for generating random samples of n -dimensional 0–1 valued vectors. Initial guesses of the probability distribution are also provided by the linear transformations of the networks. The cost function is the reconstruction error. The original batch version of the CE method generates a population of random samples, and chooses the best p ($= 5$) percent, which is used to update the density function. We modified the method to make it online: the reconstruction error is considered as a Gaussian variable with given mean and standard deviation. We updated the distribution if the error fell into the best p percent of the most recent samples. The online CE algorithm of Table I finds \tilde{h} that minimizes $\|u - \sigma(Qh)\|_{L_2}^2$.

We note, that if the matrices Q and W are well tuned, then the initial guess for the probability density function becomes sharp, and the algorithm converges very quickly. In this case, the algorithm essentially behaves as a simple feedforward linear transformation. However, the whole combinatorial reconstruction algorithm is needed for proper training. We note too that the L_2 norm is equivalent to the L_1 norm for our case, because the vectors are 0–1 valued. The known property of the CE method that it easily finds *combinatorial* solutions is exploited during the learning of compositional communication.

C. Network Training

The training of the matrices Q and W is Hebbian and has certain ‘quasi-supervised flavor’: networks are presented with observations, from which they generate internal representations. One of the networks, say agent A_1 generates utterance u_1 , which is then sent to the another agent. That is, the

TABLE I
PSEUDO-CODE FOR THE CE RECONSTRUCTION ALGORITHM

$m = 0, d = 0$	// mean and deviation of errors
$\alpha, \beta \in [0, 1]$ ($= 0.1$)	// update rates
$q = 1.648$	// 95% percentile of normal distr.
$p = \frac{Wu}{\max(Wu)}$	// initial probability distribution
$\min = \infty$	// initial minimum value
until convergence or a fixed iteration count	
generate a random sample h from p	
$e = \ u - \sigma(Qh)\ _{L_2}^2$	// calculate reconstruction error
$m \leftarrow (1 - \beta)m + \beta e$	// update mean error
$d \leftarrow (1 - \beta)d + \beta(m - e)^2$	// update deviation of errors
if ($e < m - qd$)	// if error falls to the best 5%
$p \leftarrow (1 - \alpha)p + \alpha h$	// update probability distribution
end if	
if ($e < \min$)	
$\min \leftarrow e, \tilde{h} \leftarrow h$	// update current minimum
end if	
end	

output is not supplied externally but generated by one of the networks. Then each network has an internal representation-utterance pair and can use it to update its transformation matrices. The update is Hebbian, it uses the negative gradient of the squared reconstruction error. For agent A_i ($i=1,2$) we have:

$$\Delta Q_i = \varepsilon (u_1 - Q_i h_i) h_i^T = \varepsilon e_i^u h_i^T, \quad (3)$$

$$\Delta W_i = \varepsilon (h_i - W_i u_1) u_1^T = \varepsilon e_i^h u_1^T, \quad (4)$$

where $\varepsilon \in [0, 1]$ is some update factor, e_i^h and e_i^u ($i=1,2$) denote the errors at the internal representation and utterance level, respectively. Note that the vector u_1 is the same for each agent (i.e., transmission is distortion free), but vector h_i , the matrices Q_i and W_i are (may be) different. The matrices Q_i and W_i are initialized to have random values between 0 and 1.

Feature extraction can be tuned at the listener (agent A_2 in the present example), because internal representation h_2 is available and its estimation \tilde{h}_2 can be computed from the

utterance. Let us suppose that the input was x , then:

$$f_{G_2}(x) \leftarrow \theta \left(f_{G_2}(x) + \varepsilon(\tilde{h}_2 - h_2) \right), \quad (5)$$

where $f_{G_2}(x)$ is the vector containing the values that generates the internal representation h_2 for input x . Function θ clamps the values to $[0,1]$. We used this model for adjusting feature extraction in our illustrations. The key is that the error term $\tilde{h} - h$ is available in (5) in our frame.

IV. COMPUTER SIMULATIONS

We start with the simple case when networks develop the *same internal representation* ($h_1 = h_2$), merely to investigate language emergence independently from differences in the internal representations. In particular, we chose $G_1 = G_2 = I$, the identity transformation. Later we investigate the effect of different internal representations.

A. Test Scenarios

The following experimental scenarios were studied:

- We studied non-combinatorial ‘languages’, where the utterances were forced to have only one nonzero entries (signals), and also combinatorial ‘languages’, where utterances were let to have arbitrary combinations of nonzero entries.
- Besides using the reconstruction principle, we studied generation and parsing methods using simple linear transformations Q and W followed by a rounding non-linearity, i.e., without the reconstruction algorithm. In the non-combinatorial case, following the linear transformation the maximum valued component was set to 1, others to 0.
- We systematically varied the size of the internal representation, and the number of agents to see how the learning scales with these factors.

In each episode of learning, two random selected networks participated in communication. An input was selected randomly, and one of the networks generated an utterance to it, the other parsed it, and then both of them updated their transformations. To decide whether a consistent language had emerged, we defined a performance matrix: the relative frequency of the usage of each signal for each input was calculated from communications about the given input. We say that a consistent language developed, if all networks produced consistently the same signals for the same input. In the non-compositional case, each state was required to be denoted by a different signal. In the compositional case, we call a language consistent, if an utterance denoting an internal representation with certain features are composed of utterances referring to those features, and all the networks use the same combinations. We also recorded how often the parsing agent could reconstruct the same internal representation from the utterance it received as it generated from its observation ($\tilde{h} = h$) towards the end of a series of

communication episodes; this is the communication success rate (note that it happens to be 1 for a consistent language).

B. Results

First we tested how the combinatorial and non-combinatorial methods behave as a function of the size of the internal state and the number of agents. We evaluated the percentage of the runs when the method converged to a consistent language, and the average number of learning episodes that agents needed to reach that. We found that if the reconstruction principle was applied, learning reached a consistent state in 100% of the cases, both for combinatorial and for non-combinatorial languages. The number of episodes needed to reach an agreement is shown in Fig. 2. It can be seen that the combinatorial method needs reasonably smaller number of learning episodes as the state size and the agent count increases. We observed that when combinatorial solution was allowed then *combinatorial / compositional* language developed in all of the cases.

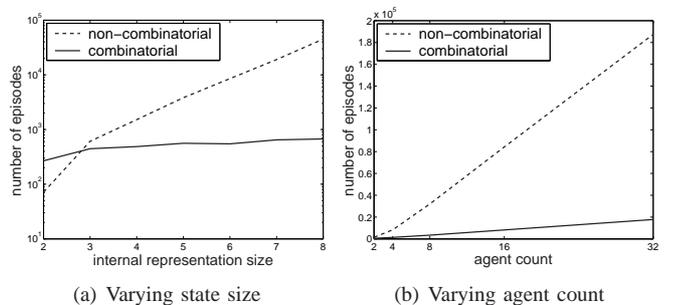


Fig. 2. **The effect of varying representation size and agent number.** (a) the y axis is on a logarithmic scale: slope is slow for compositional languages, but it is exponential for holistic ones as a function of the size of the internal representation (b) the y axis is linear; the runs were done with an internal representation size of 4

We also investigated how the learning changes when the reconstruction principle was not applied. Surprisingly, in the non-combinatorial case, this method was never sufficient to develop a consistent language. For the combinatorial case, Fig. 3 shows the results. It can be seen that without the reconstruction principle the development of the language becomes inefficient with the size of internal representation. The effect is similar if the agent count is increased (not shown here). The ratio of consistent languages drops drastically. Furthermore, the application of the reconstruction principle also has an intriguing effect on the number of signals used by the agents. Theoretically an n -component state can be communicated using the combination of n signals. This lower bound have been reached when the reconstruction principle was applied, but it was significantly exceeded when not. Agents used many ‘synonyms’ in this case.

To see how learning behaves when agents have different internal representations, we have explicitly generated

feature sets (i.e., G transformations) for a finite number of inputs. Differences between agents' G transformations were systematically introduced. In this case, totally consistent language can not develop, agents can not agree because the categorization of the observations differs. We let learning run for a sufficiently large number of episodes, and during the last 1000 episodes, we evaluated what fraction of the communication episodes were successful. That is, we counted the ratio of cases when the parsing network was able to reconstruct the same internal state from the utterance as it actually developed from the input. We found that communication efficiency drops as the differences in the internal representation increase (Fig. 4a) as expected.

However, when adjusting feature values (concept formation) utilizing (5) was enabled after an initial number of communication tuning episodes, learning process became convergent again with *compositional* languages. Consistent communication developed in all of the cases when the initial differences were not too large (Fig. 4b).

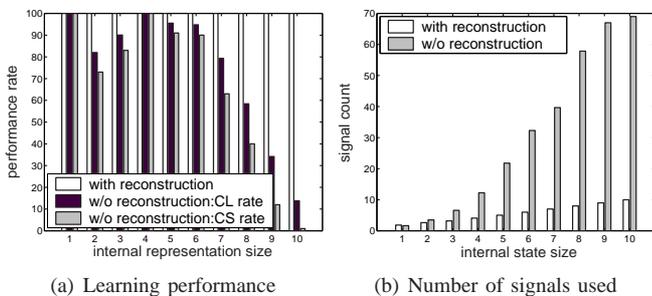


Fig. 3. **Using and not using the reconstruction principle.** CS: communication success, CL: consistent language. (a) sharp drop for larger state sizes and (b) strong increase in the number of developed signals without reconstruction. Results are similar for increasing agent count.

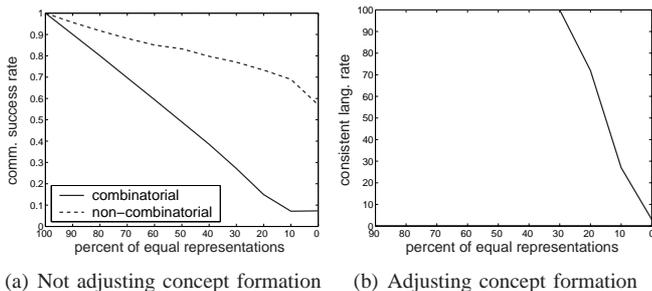


Fig. 4. **The effect of differences in the internal representations.** Size of representation: 6. (a) drop of performance is more serious in the case of compositional languages. (b) learning becomes successful for compositional communication with probability 1 if initial differences are not too large.

V. DISCUSSION

We observed that at the beginning of the learning, the networks have synonymous signals for denoting components of

the state. First '*they learn*' to understand each others' signals, and later '*they refine*' their dictionaries to single common signals for any given component. When the reconstruction principle is not in effect, this negotiation is not successful and the number of signals often increases. However, when reconstruction is utilized, negotiation is accomplished by adaptation to signals used by the other parties.

The obverter learning procedure [10] applies the same idea as our reconstruction principle. In [10] they prove that the best strategy for agents is to produce utterances that maximize the chance of other agents understanding it. They argue that agents do not have access to what others would understand, so it seems a good idea to produce utterances that the agent itself would understand well. This idea is exploited in our reconstruction network. Our algorithm goes beyond the ideas described in [10], because we can handle compositions, and we can work with individuals having distinct conceptual representations. The necessity of an obvious feature of our model, that production and understanding are dependent on each other and evolve simultaneously, has been emphasized by Cangelosi [2].

In [8] Kirby argues that compositional languages emerge due to the learning bottleneck effect of linguistic knowledge transfer from generation to generation. He claims that compositional languages are favored because they are easier to pass to the next generation since fewer observations are enough to learn them because of their compressed nature. Our simulations indicate that there is another reason why compositional languages are favored, namely that they are easier to agree upon. Nonetheless, the reason is the same as that of Kirby; their compressed nature enables faster negotiation, since only the signals referring to components (instead of their combinations) need to be agreed on. Actually, we have observed, that relatively few categorization samples are enough to agree on a consistent compositional language (about 20 samples in case of an 8-component representation).

Smith [13] and Vogt [17] both use discrimination games, by which agents develop a categorial representation of observations. When experimenting with the effect of different representations, Smith comes to a conclusion that the overall success of communication seems to be directly related to the amount of shared meaning structure in the agents. This conjecture is strengthened by our results, and is also present in the conclusions of Cangelosi. However, [2] only deals with holistic communication in a scenario where there is an evolutionary pressure for agents to develop *similar* internal representations. We had shown that holistic communication is more resistant to representational differences. The underlying reason is probably the compactness and generalizing capability of compositional communication: the misunderstanding of utterances generalizes across similar internal representations. To let successful communication emerge even in the case of different internal representations, we adjusted the

representations themselves. The adjustments are based on the differences between the representations induced by the utterance and the one generated from the agent's own observation. In case of a compositional language, the utterance is a projection of how the other agent categorizes the observation, and this information can be used to alter an agent's own categorization. This might turn out to be an important feature of compositional languages, since holistic languages lack this information.

It must be noted, that we did not aim to model the development of compositional *syntax*, only that of a compositional *lexicon*. Syntactic structures (e.g. the order of signals) could compress information even further. Instead of letting $u \in \{0,1\}^n$, the utterance should be transformed into a sequence of signals (words), organized by syntactic rules. Our formalism treats the problem of *which signals to use* to express a given internal representation (meaning), and defers the problem of sequencing the set of signals. However, using syntax to organize the signals might also effect the selection of signals, since some aspects of meaning can be expressed by pure syntax, like word order, thus further compressing the utterance.

The use of the powerful cross-entropy stochastic optimization for the individual examples allowed us to restrict learning to a local Hebbian rule derived from the reconstruction principle, which is an important feature of our approach. Furthermore, no other method reported 100% success even for many agents.

The potentials of any approach to language emergence should be examined in full-fledged challenges: language development / improvement may be a must in various cases, including the problem of distributed sensory and intelligent computational units experiencing different slices of the world. Consider the following situation. Given is a world where agents aim to reach their goals, e.g., they aim to survive, solve a given task alone or in groups. At a certain time, they observe different portions or aspects of the world either because of their locations, or because they have different sensors, or both. Thus they may build different internal representations about the world alike to the *six blind men* in the famous *Indian legend*. Suppose that the task of these blind men is travelling. Communication would reduce the partially observed nature of the world, and it might enable gathering novel knowledge. How could they combine their limited knowledge and find out the high value of the elephant as an unexpected conveyance means that fits their purposes? It would be useful to develop a common language and to explain their sensory information in the common terms. We think that *intentional* coordinated exploration is necessary to establish common concepts and to make common discoveries. We believe that this is the underlying reason for the emergence of language during evolution. If this is true, then we are barely scratching the surface of the full problem.

VI. CONCLUSIONS

To conclude, a neural network approach was developed that effectively implements ideas about the development of a combinatorial communication system. It applies the reconstruction principle that makes the negotiation process of the parties become convergent in the case of identical meaning structures. Experiments show that communication is sensitive to the differences in the conceptual representation, and compositional communication has the advantage that it carries a potential to adjust it.

ACKNOWLEDGMENT

We are grateful to István Szita for his valuable contribution to the development of the online cross-entropy method.

REFERENCES

- [1] D. H. Ballard, G. E. Hinton, and T. J. Sejnowski. Parallel visual computation. *Nature*, 306:21–26, 1983.
- [2] A. Cangelosi and D. Parisi. The emergence of a language in an evolving population of neural networks. *Connection Science*, 10(2):83–97, 1998.
- [3] B. Comrie. From potential to realization: an episode in the origin of language. *Linguistics*, 38 (5):989–1004, 2000.
- [4] P. F. Dominey and J. F. Boucher. Developmental stages of perception and language acquisition in a perceptually grounded robot. *Cognitive Systems Research*, 6:243–259, 2005.
- [5] S. Harnad. The symbol grounding problem. *Physica D*, 42:335–346, 1990.
- [6] B. K. P. Horn. Understanding image intensities. *Artificial Intelligence*, 8:201–231, 1977.
- [7] E. Hutchins and B. Hazlehurst. How to invent a lexicon: the development of shared symbols in interaction. In G. N. Gilbert and R. Conte, editors, *Artificial Societies: The computer simulation of social life*. UCL Press, London, UK, 1995.
- [8] S. Kirby. Learning, bottlenecks and infinity: A working model of the evolution of syntactic communication. In K. Dautenhahn and C. Nehaniv, editors, *Proc. of the AISB '99 Symp. on Imitation in Animals and Artifacts*, 1999.
- [9] S. Kirby and J. Hurford. The emergence of linguistic structure: An overview of the iterated learning model. In A. Cangelosi and D. Parisi, editors, *Simulating the Evolution of Language*, chapter 6, pages 121–148. Springer Verlag, London, 2002.
- [10] M. Oliphant and J. Batali. Learning and the emergence of coordinated communication. *Newsletter of the Center for Research in Language*, 11(1), 1997.
- [11] D. Roy and A. Pentland. Learning words from sights and sounds: A computational model. *Cognitive Science*, 26:113–146, 2002.
- [12] R. Y. Rubinstein and D. P. Kroese. *The Cross-Entropy Method*. Information Science and Statistics. Springer Verlag, 2004.
- [13] A. D. M. Smith. Establishing communication systems without explicit meaning transmission. volume 2159 of *LNCS*, pages 381–390, 2001.
- [14] K. Smith, S. Kirby, and H. Brighton. Iterated learning: A framework for the emergence of language. *Artificial Life*, 9(4):371–386, 2003.
- [15] L. Steels. Grounding symbols through evolutionary language games. In A. Cangelosi and D. Parisi, editors, *Simulating the Evolution of Language*, chapter 10, pages 211–226. Springer Verlag, London, 2002.
- [16] L. Steels and F. Kaplan. Aibo's first words, the social learning of language and meaning. *Evolution of Communication*, 4 (1):3–32, 2001.
- [17] P. Vogt. Iterated learning and grounding: From holistic to compositional languages. In S. Kirby, editor, *Proc. of Lang. Evol. Comp. Workshop*, pages 76–86, Vienna, A., 2003.
- [18] C. Yu and D. H. Ballard. A multimodal learning interface for grounding spoken language in sensory perceptions. *ACM Transactions on Applied Perception*, 1:57–80, 2004.