

Language Emergence in a Population of Artificial Agents Equipped with the Autotelic Principle

Miquel Cornudella

Sony Computer Science Laboratory Paris
6 rue Amyot, 75005
Paris, France
cornudella@csl.sony.fr

Thierry Poibeau

Laboratoire LATTICE-CNRS
1 rue Maurice Arnoux, 92120
Montrouge, France
thierry.poibeau@ens.fr

Abstract

Experiments on the emergence of a shared language in a population of agents usually rely on the control of the complexity by the experimenter. In this article we show how agents provided with the autotelic principle, a system by which agents can regulate their own development, progressively develop an emerging language evolving from one word to multi-word utterances, increasing its discriminative power.

1 Introduction

The evolution of communication has been a topic in artificial life since early 90s (Werner, 1991; Ackley and Littman, 1994). Short after that, a group of Alife researchers started to focus on the origins and emergence of human language-like communication systems through experiments with populations of artificial agents (Smith et al., 2003; Steels, 2003; Wagner et al., 2003). This line of research has shed light on the emergence of spatial terms and categories (Spranger, 2013), case systems (van Trijp, 2012), quantifiers (Pauw and Hilferty, 2012) or syntax (Kirby, 1999; Steels and Casademont, 2015). However, the success of these experiments usually relies on the control of complexity by the experimenter.

In order to let the agents manage complexity themselves it is necessary to provide them with a mechanism to regulate complexity in an autonomous way. Research in AI and robotics has explored systems that allow embodied agents to develop themselves in open-ended environments by means of error reduction (Andry et al., 2001), reinforcement learning (Huang and Weng, 2002), prediction (Marshall et al., 2004) or curiosity (Oudeyer et al., 2007; Kaplan and Oudeyer, 2007). This mechanisms are highly inspired by

psychological studies on the role of motivation (Hull, 1943; Skinner, 1953; White, 1959; Graham, 1996). Motivation can be defined as “to be moved to do something” (Ryan and Deci, 2000) and it is commonly divided in *extrinsic motivation*, when an activity is done to attain some separable outcome, and *intrinsic motivation*, when an activity is done for its inherent satisfactions.

This paper investigates the role of intrinsic motivation in language emergence. It presents an agent-based experiment where a population of artificial agents has to develop a language to refer to objects in a complex environment. In addition to mechanisms to invent and adopt words and syntactic patterns, agents are provided with an operational version of the Flow theory (Csikszentmihalyi, 1990) that enables them to self-regulate their development.

2 Flow Theory

The model of intrinsic motivation in a population of artificial agents used in this experiment is based on the Flow theory developed by the psychologist Csikszentmihalyi (1990). He studied what moves people to be deeply involved in a complex activity that does not present a direct reward. He called these activities *autotelic*, as the motivational driving force (*telos*) comes from the individual herself (*auto*).

Csikszentmihalyi states that in an autotelic activity there is a relation between *challenge*, how difficult a particular task is, and *skill*, the abilities a person requires to face that particular task. As a consequence of this relation, a person involved in an autotelic activity can experience three mental states: *boredom*, when the challenge is too low for the skills this person has, *flow*, when there is a balance between challenge and skills, and *anxiety*, when the challenge is too high for the available skills. The flow state produces an intense enjoyment in a person involved in an autotelic activ-

ity. The flow state is not static but in continuous movement, since the balance between challenge and skills creates the ideal conditions to develop skills. As a consequence this person becomes self-motivated, as she tries to stay in the flow state to experience this strong form of enjoyment.

3 Autotelic Principle

The autotelic principle is an operational version of the flow theory that provides agents with a system to self-regulate their development (Steels, 2004). It was first designed for developmental robotics (Steels, 2005) but it has also been used to study language emergence (Steels and Wellens, 2007). This principle proposes the balancing between challenge and skills as the motivational driving force in agents. Agents are therefore provided with mechanisms to set their own challenges and evaluate their performance to determine their *emotional state*. Depending on their emotional state, agents autonomously decide to increase their challenge (*boredom*), decrease it (*anxiety*) or continue with the current challenge to keep developing their skills (*flow*).

Challenges are defined as a specific configuration of a set of parameters. For example, parameters can be the number of objects or the number of properties of an object that agents can refer to. Challenges are formally represented as $\langle p_{i,1}, \dots, p_{i,n} \rangle$ in a multi-dimensional parameter space P , where $p_{i,j}$ corresponds to the configuration of the parameter j in the challenge i . Steels found advantageous to initialize the system with the lowest challenge configuration and grow in a bottom-up manner. There are no studies on the effect of a higher challenge configuration initialization in agents, but it will probably result in a slower development of skills.

Agents can estimate their skills by measuring their *performance*. Performance is measured taking into account an overall estimation of the interaction (if they have succeed or failed) and specific performance measures for each component used. Components are subsystems of the agent that are responsible for specific tasks, such as selecting a topic, conceptualise it into a meaning predicate or formulate an utterance given a meaning predicate. For example, in a communicative challenge the conceptual component has a performance measure of how well the resulting conceptualisation discriminates the topic or the language component

a measure that evaluates if it could formulate an utterance covering the conceptualisation.

Agents also keep track of how *confident* they are to succeed on the challenge they have posed to themselves. The confidence in a challenge is related to the skills agents require to deal with that challenge. In a challenge where agents have to come up with names for objects, the development of a lexicon increases its communicative success and the confidence in being able to cope with the challenge.

Agents are constantly alternating between the operational and the shake-up phases. The operational phase takes place when the challenge parameters are fixed. The agent explores this configuration and tries to develop its skills to reach a certain level of performance. The shake-up phase occurs when the performance and confidence measures are stable. Agents employ this measures to determine how the challenge parameters should be adjusted. If the performance and confidence measures are low, agents perceive that they are in an anxious state and decrease the challenge parameters. Alternatively, when both performance and confidence measures are high, agents enter a boredom state and increase the challenge parameters.

4 Experiment configuration

The aim of this experiment is to show how a population of artificial agents provided with the autotelic principle develop a shared language without any control on the complexity by the experimenter. Agents play a *language game*, which consist in situated communicative interactions between two agents of a population (Steels, 2012). These agents are randomly selected from a population of ten agents. One of them assumes the role of *speaker* and the other the role of *hearer*.

4.1 World

In the experiment, agents share a world, which consist of ten different scenes. Each scene is composed of two objects and a spacial relation between them, such as *close*, *far* or *left of*. Objects are characterised by three feature-value pairs: *prototype* (e.g.: chair, box, table), *color* (e.g.: green, blue, purple) and *shape* (e.g.: round, hexagonal, square). Objects and scenes are unique, but a particular feature-value can be shared by two or more objects. In an interaction speaker and hearer share the same context, which consist of a randomly se-

lected scene from the world.

4.2 Language game

The specific language game that agents play is called *multi-word guessing game*. The speaker selects a topic from the context of the interaction, based on his current communicative challenge. It conceptualises this concept into a meaning predicate and uses its language component to formulate an utterance which is transmitted as text to the hearer. The hearer tries to comprehend the utterance and construct hypotheses about the topic. If the hearer has only one hypothesis, it points to the interpreted topic. If the hypothesis corresponds to the topic, the speaker gives positive feedback and the interaction ends. On the other hand, if the hypothesis does not correspond to it, the speaker gives negative feedback to the hearer and points to the intended topic. When the hearer has multiple hypotheses, it signs to the speaker that it could not identify the topic. The speaker then gives feedback by pointing to the intended topic. The interaction is a success only when the hearer has one hypothesis about the topic that corresponds with the topic selected by the speaker. In all other cases, the result of the interaction is a failure.

4.3 Challenges

Agents refer to one or two objects in the scene, and minimally express the *prototype* of the object(s). Apart from the prototype, agents can refer to one or more properties of objects or to the relation between them. The challenge configuration is therefore based on two parameters: the number of properties agents refer to and if the relation is expressed or not. Challenges have a *confidence* value between 0.0 and 1.0, initialised at 0.0. After each interaction, speaker and hearer update their confidence value with a score obtained computing the average between the result of the interaction (success or failure) and the performance evaluation of the components used by the agent. The update score has a low value (between 0.008 and -0.032) to provide agents enough time to develop the skills necessary to cope with the challenge.

The challenge level one (refer only to prototypes of objects) is set as the initial challenge. In the experiment agents can adjust the challenge configuration up to level four: refer to up to two objects expressing three of their properties or to relations between objects.

Challenge Level	Properties	Relation
1	0	0
2	1	0
3	2	0
4	3	0
	2	1

Table 1: Challenge levels.

4.4 Mechanisms

Agents are equipped with *conceptualisation* and *interpretation* mechanisms to map between the world model and meaning predicates that refer to it. For example, a blue table is conceptualised into ($blue(x), table(x)$). Agents start without any form-meaning mappings (also called *constructions*). This mappings will emerge during interactions by using three mechanisms: *diagnostics*, *repairs* and *alignment*.

Diagnostics are a set of processes by which agents can identify problems during formulation (when agents go from a meaning predicate to an utterance) and comprehension (when agents reconstruct the meaning predicate from an input utterance). In the experiment agents can identify *unknown meanings*, *unknown words*, *unsolved word orders* and *referent problems*.

Repairs are strategies used by agents to solve diagnosed problems. For example, an unknown meaning can be solved by the speaker with a repair that creates a new word for that meaning, or an unknown word can be solved by the hearer with a repair that uses the feedback of the speaker to identify which meaning corresponds to that word. Notice that the later is only possible when the hearer can unambiguously deduce the meaning of the unknown word. Unsolved word orders and referent problems appear when agents start to build multi-word utterances. This problem can be solved by creating grammatical constructions that introduce constraints on how properties and prototypes are ordered when formulating and comprehending multi-word utterances.

There is a competition of form-meaning mappings (both lexical and grammatical) during the emergence of a shared language. This competition occurs either when multiple forms refer to the same meaning or when one word can express several meanings. Each mapping has a score between 0.0 and 1.0 and is initialised at 0.5. Alignment is a mechanism that guides the choice of which con-

structions agents use based on the score of their constructions. The scores of the mappings used by the speaker and hearer are updated after each interaction. When a form-meaning mapping gets a score of 0.0 is deleted from the construction inventory of the agent. The alignment used in this experiment follows the dynamics of *lateral inhibition* (De Vylder and Tuyls, 2006).

When there is communicative success, both speaker and hearer align, which means that they increase the scores of the mappings used by 0.1 and decrease its competitors by 0.1. Note that the mapping competitors for the speaker are those constructions that express the same meaning, while mapping competitors for the hearer are those that contain the same form. When there is communicative failure, the alignment differs for speaker and hearer. If the speaker has formulated one word utterance, it decreases the score of the construction used by 0.1. The hearer aligns only when the intended topic by the speaker is among its hypotheses. It increases the score of the constructions used by 0.1 and decreases the score of its form competitors by 0.1. In all other cases agents are not able to identify what caused the communicative failure and do not align.

5 Experimental results

The results of ten experimental runs for a population of ten agents equipped with the autotelic principle are shown in Figure 1.




Figure 1: This graph shows communicative success (left y-axis) and the average confidence on challenge level (right y-axis) in a population of 10 agents equipped with the autotelic principle.

Agents start with an empty construction inventory and with the challenge of emerging a shared language for prototypes. They develop it rapidly, increasing their confidence on the first challenge

up to its maximum value around interaction 2000. Note that the communication success starts to drop before the average confidence value in the population has come to its maximum. This is due to the fact that some agents have already reached the highest confidence score and therefore they have moved to the next challenge.

Communicative success and the speed at which agents gain confidence decreases at this point, as agents begin to refer also to the color and shape of objects. Agents have to agree now on form-meaning mappings to refer to color and shape and grammatical constructions to manage reference problems in multi-word utterances. Communicative success and confidence in challenge levels two and three grow steadily until they reach its maximum value around interaction 5500. The population has reached the maximum level of confidence for the first three challenge levels and start to address challenges of level four. The communicative success slightly diminishes at this point due to the fact that agents have to agree on how to refer to relations. By interaction 9000 all agents have reached the maximum confidence for each challenge.

There are differences on the percentage of communicate success that agents are able to reach for each challenge level. These differences are due to the fact that some topic descriptions are ambiguous. The discriminative power of an utterance increases when agents refer to more properties of objects or the relation between them. This accounts for the differences observed on Figure 1, where agents reach a higher percentage of communicative success once they have agreed on how to refer to properties and relations.

The results obtained show that a population of agents equipped with the autotelic principle manage to autonomously increase the complexity of a shared language through recurrent interactions. Agents succeed in progressively develop their communicative skills when trying to stay in a state of flow. As a result, agents reach a higher communicative success in their interactions, as they can successfully refer to more informative topic descriptions which are less ambiguous.

Acknowledgments

Miquel Cornudella is partially supported by a CIFRE grant (agreement no. 2013/0730). The authors wish to thank their colleagues for their feed-

back and support, particularly Remi van Trijp and Paul Van Eecke.

References

- David H Ackley and Michael L Littman. 1994. Altruism in the evolution of communication. *Artificial life IV*, pages 40–48.
- Pierre Andry, Philippe Gaussier, Sorin Moga, Jean-Paul Banquet, and Jacqueline Nadel. 2001. Learning and communication via imitation: An autonomous robot perspective. *Systems, Man and Cybernetics, Part A: Systems and Humans, IEEE Transactions on*, 31(5):431–442.
- Mihaly Csikszentmihalyi. 1990. *Flow: The psychology of optimal experience*. Harper and Row, New York.
- Bart De Vylder and Karl Tuyls. 2006. How to reach linguistic consensus: A proof of convergence for the naming game. *Journal of Theoretical Biology*, 242(4):818 – 831.
- Sandra Graham. 1996. Theories and principles of motivation. *Handbook of educational psychology*, 4:63–84.
- Xiao Huang and John Weng. 2002. Novelty and reinforcement learning in the value system of developmental robots. In *Lund University Cognitive Studies*, pages 47–55.
- Clark Leonard Hull. 1943. *Principles of behavior: an introduction to behavior theory*. Appleton-Century.
- Frédéric Kaplan and Pierre-Yves Oudeyer. 2007. The progress-drive hypothesis: an interpretation of early imitation. *Models and mechanisms of imitation and social learning: Behavioural, social and communication dimensions*, pages 361–377.
- Simon Kirby. 1999. Syntax out of learning: The cultural evolution of structured communication in a population of induction algorithms. In D. Floreano, J.D. Nicoud, and F. Mondada, editors, *Advances in Artificial Life: Proceedings of the 5th European Conference on Artificial Life*, pages 694–703, Berlin. Springer.
- James B. Marshall, Douglas Blank, and Lisa Meeden. 2004. An emergent framework for self-motivation in developmental robotics. In *Proceedings of the 3rd international conference on development and learning, Salk Institute, San Diego*, volume 10.
- Pierre-Yves Oudeyer, Frédéric Kaplan, and Verena Vanessa Hafner. 2007. Intrinsic motivation systems for autonomous mental development. *Evolutionary Computation, IEEE Transactions on*, 11(2):265–286.
- Simon Pauw and Joseph Hilferty. 2012. The emergence of quantifiers. *Experiments in Cultural Language Evolution*, 3:277.
- Richard M Ryan and Edward L Deci. 2000. Intrinsic and extrinsic motivations: Classic definitions and new directions. *Contemporary educational psychology*, 25(1):54–67.
- Burrhus Frederic Skinner. 1953. *Science and human behavior*. Simon and Schuster.
- Kenny Smith, Simon Kirby, and Henry Brighton. 2003. Iterated learning: A framework for the emergence of language. *Artificial Life*, 9(4):371–386.
- Michael Spranger. 2013. Evolving grounded spatial language strategies. *KI-Künstliche Intelligenz*, 27(2):97–106.
- Luc Steels and Emília Garcia Casademont. 2015. Ambiguity and the origins of syntax. *The Linguistic Review*, 32(1):37–60.
- Luc Steels and Pieter Wellens. 2007. Scaffolding language emergence using the autotelic principle. In *Artificial Life, 2007. ALIFE'07. IEEE Symposium on*, pages 325–332. IEEE.
- Luc Steels. 2003. Evolving grounded communication for robots. *Trends in cognitive sciences*, 7(7):308–312.
- Luc Steels. 2004. The autotelic principle. In Fumiya Iida, Rolf Pfeifer, Luc Steels, and Yasuo Kuniyoshi, editors, *Embodied Artificial Intelligence*, volume 3139 of *Lecture Notes in Computer Science*, pages 231–242. Springer Berlin Heidelberg.
- Luc Steels. 2005. The emergence and evolution of linguistic structure: from lexical to grammatical communication systems. *Connection Science*, 17(3-4):213–230, December.
- Luc Steels. 2012. *Experiments in Cultural Language Evolution*. Advances in interaction studies. John Benjamins Publishing Company.
- Remi van Trijp. 2012. Not as awful as it seems: Explaining german case through computational experiments in fluid construction grammar. In *Proceedings of the 13th Conference of the European Chapter of the Association for Computational Linguistics, EACL '12*, pages 829–839, Stroudsburg, PA, USA. Association for Computational Linguistics.
- Kyle Wagner, James A. Reggia, Juan Uriagereka, and Gerald S. Wilkinson. 2003. Progress in the simulation of emergent communication and language. *Adaptive Behavior*, 11(1):37.
- Gregory M Werner. 1991. Evolution of communication in artificial organisms, artificial life ii. In *Proceedings of the Second International Conference of Artificial Life*, pages 659–687.
- Robert W White. 1959. Motivation reconsidered: the concept of competence. *Psychological review*, 66(5):297.