Emergence of Grammatical Conventions in an Agent Population Using a Simplified Tree Adjoining Grammar

Christophe Allexandre and Andrei Popescu-Belis
Groupe Langage et Cognition, LIMSI ñ CNRS
popescu@limsi.fr

Abstract

This paper describes a population of communicating agents, rewarded for successful dialogs. Agents encode and decode messages about their environment using the TAG formalism. Experimental results show that lexical and word order conventions spread under suitable conditions.

1. Introduction

Modeling language development and use in a group of communicating synthetic agents may bring new insights to some problems of natural language processing. Such models have been proposed by MacLennan [1] [2] and by Dyer [3], using elementary, automata-like agents to study evolution of a common lexicon. Steels [4] advocates use of physical robots, and convincingly grounds lexicon and even grammatical categories in the agents' perception and categorization of the real world. Hashimoto and Ikegami [5] focus on emergent syntax from a formal point of view.

In this paper, we study emergence of words and of word order conventions among agents. We use a simplified version of the Tree Adjoining Grammar [6], in which derivation trees represent configurations of the environment (ß2), commented upon by the agents (ß3.1). Dialogs allow the linguistic conventions to spread (ß3.2) as shown by experimental results (ß4).

2. The Environment

Configurations of the environment are randomly generated by the control program. They consist in a logical form description of a block world made of objects (cubes, spheres, cylinders) with various characteristics (red, big, heavy), and involved in positional relations (on, in front of, left of) and comparison relations (greater, taller, heavier). We take advantage of the TAG formalism to describe situations in terms of *derivation trees* (DT), which account for semantic dependencies in a phrase. Here, they represent entities and their relations in a given situation of the environment (Fig. 2, top). The nodes and leaves of a derivation tree are concepts (@cube, @red, @on) and the branches show their combination: adjunction (continuous line) or substitution (dashed line).

3. The agents

The only goal of our agents is to exchange messages (strings) about their environment and update their linguistic knowledge depending on their successful understanding.

3.1. Linguistic knowledge

Initial knowledge of each agent is limited to the valence of the concepts, i.e., the number and labels of their arguments. Concepts are grouped in classes: objects (o), characteristics (c), positional (r1) and comparison (r2) relations. The agents have *a priori* elementary trees for each class: o, c, r1, r2 (Fig. 1). But, they have no initial knowledge of the concrete words for each concept (double quotes on Fig. 1) nor any preference for a given order among the elementary tree branches. For instance, r1 can have (O1, R1, O2) order, as well as (R1, O1, O2), etc. Lexical and syntactic parameters evolve through dialogs.

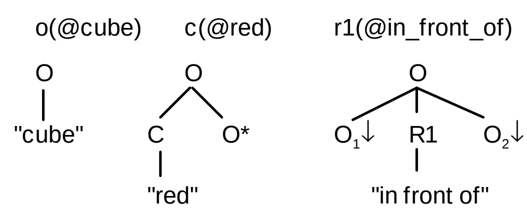

Figure 1. Example of elementary trees

Given a derivation tree DT corresponding to a situation, an agent is capable to derive the associated parse tree using its own elementary trees, and build a linear form LF (Fig. 2). Given an LF and the corresponding situation, an agent is also capable to use its lexicon and rules to *understand* the message, i.e. match the LF to any possible LF derived from (part of) the environment. The emission and reception capacities are used alternatively in dialogs.

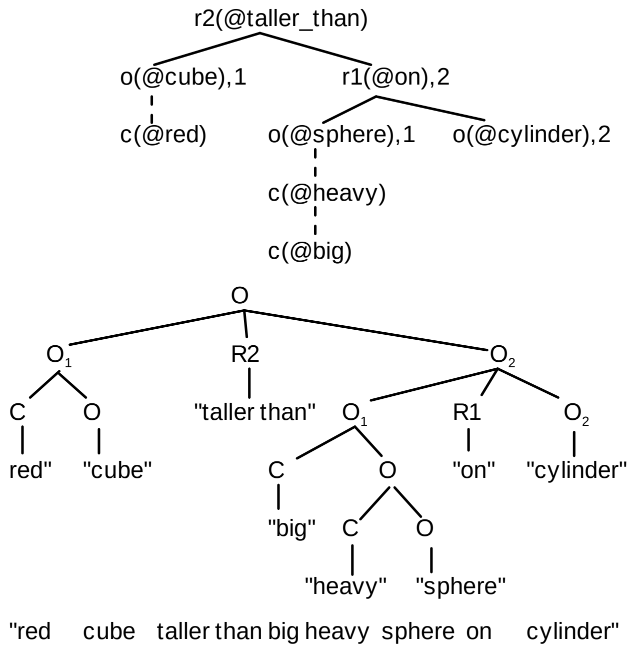

Figure 2. Derivation tree, parse tree and linear form

3.2. Dialogs between agents

The controller randomly selects two agents for a dialog, and gives them a common situation (DT). The sender uses its preferred lexicon and word order—or new ones, if it has none—to generate an LF which is sent to the receiver. The sender may encode the whole situation or only part of it. The receiver tries to decode the LF, using first its preferred parameters, then all parameters if necessary, and finally it may even guess one or several words.

The dialog's success depends on the receiver's comprehension. If the receiver understands only one (sub)situation using its preferred parameters, the dialog is a complete success. If it has several hypothesis, or it has used a non-preferred parameter, the dialog is a partial success. If the receiver can build no interpretation, the dialog is a failure.

In case of a complete success all the parameters used by the sender and the receiver see their fitness increase. If the success is partial, only the receiver increases its values; in case of a failure, the parameters are decreased. For each agent, the parameters whose fitness is too low are removed, and new ones are created randomly.

4. Results

The typical way to evolve linguistic conventions is to present agents with more and more complex situations (o, then o+c, then o+c+r1, etc.) For a group of five agents, an average of 500 dialogs is necessary to establish unique names for the objects. When characteristics are introduced, it takes ~1200 dialogs to establish unique order and names for them. It takes then ~9000 dialogs to spread r1-conventions, and the same for r2. With ten agents, ~1000 dialogs are necessary for o-conventions, ~2800 for c, ~50000 for r1 and the same for r2.

Senders can be authorized to describe only a part of the given situation, unknown to the receiver. For a population having already established linguistic conventions by incremental learning, experiments show that conventions still remain stable in this case, and allow the receiver to find out which situation is described.

When agents are presented from the beginning with situations combining o, c, r1 and r2 (non-incremental learning), they fail to develop a differentiated language and label all concepts with the same word. Thus, it seems that grammar development (lexicon and syntax) cannot bypass the one-word and two-words stages.

When a new agent is added to a stabilized population (five agents), it acquires the conventions after ~1500 dialogs. There is no need for specific incremental learning, as the new agent understands at the beginning only the simple descriptions (o, o+c), and then the more complex ones (o+c+r1+r2). But when a new population is added to a stabilized population of the same size, incremental learning becomes necessary.

In conclusion, the use of TAG formalism has lead to encouraging results. However, the dependence of the convergence rate on the complexity of the situations has still to be assessed through further experimentation.

5. References

[1] B. MacLennan, "Synthetic Ethology: An Approach to the study of communication", *Artificial Life II*, Santa Fe, New Mexico, 1992, pp. 631-658.

[2] J. Crumpton, *Simulated Evolution of Two Symbol Communication*, Master Thesis, University of Tennessee, Knoxville, 1994.

[3] G. Werner and M. Dyer, "Evolution of Communication in Artificial Organisms", *Artificial Life II*, Santa Fe, New Mexico, 1992, pp. 659-687.

[4] L. Steels, "The Origins of Syntax in Visually Grounded Robotic Agents", *International Joint Conference on Artificial Intelligence*, Nagoya, Japan, 1997, vol. 2, pp. 1632-1641.

[5] T. Hashimoto and T. Ikegami, "Emergence of Net-grammar in Communicating Agents", *European Conference on Artificial Life*, University of Sussex, Brighton, UK, 1997.

[6] A. AbeillÈ and Y. Schabes, "Parsing Idioms with a Lexicalized Tree Adjoining Grammar", *Fourth European ACL*, Manchester, UK, 1989, pp. 1-9.