

Collective learning and semiotic dynamics

Luc Steels^{1,2} and Frederic Kaplan²

¹VUB AI Laboratory, Pleinlaan 2, 1050 Brussels,

²SONY Computer Science Laboratory, 6 Rue Amyot, Paris
steels@arti.vub.ac.be, kaplan@csl.sony.fr

Abstract. We report on a case study in the emergence of a lexicon in a group of autonomous distributed agents situated and grounded in an open environment. Because the agents are autonomous, grounded, and situated, the possible words and possible meanings are not fixed but continuously change as the agents autonomously evolve their communication system and adapt it to novel situations. The case study shows that a complex semiotic dynamics unfolds and that generalisations present in the language are due to processes outside the agent.

1 Introduction

In recent years it has become clear that the complex adaptive systems approach pioneered by Artificial Life research can fruitfully be applied to the study of the origins and evolution of language [9], particularly to the emergence of shared sound systems [3], the self-organisation of lexicons [7], [11], grounded word meaning [12], and the origins of grammar [4], [1], [5]. In all this research, the same mechanisms for the generation and maintenance of complexity are being used as exploited in other Artificial Life research, and a similar complex dynamics can be seen to emerge.

This paper focuses on *grounded* lexicons as they emerge from the local interactions of a group of distributed autonomous robotic agents, grounded in a real world physical environment through visual sensing. Consequently, and in contrast with other work so far, the meanings of words are no longer given as symbols supplied by a designer, nor is it assumed that hearers have perfect knowledge of what meaning is intended by a speaker. Rather the agents must autonomously infer the possible meanings of unknown words from their visual interpretation of the situations they encounter. Agents never get immediate feedback on whether they had the right meaning, only whether the communication was successful. Grounding introduces two additional difficulties: The ontology must be sufficiently robust to handle variations in the data. Anomalies in perception come from differences in the perception (for example one agent segmenting into different objects than the other). This causes the rise of additional ambiguities which need to be damped by the semiotic dynamics. We show that under these conditions, which are more realistic with respect to human natural language acquisition and use, a very complex semiotic dynamics is generated which nevertheless manages to self-organise into a successful communication system.

2 The Talking Heads experiment

The robotic setup used for the experiments in this paper consists of a set of ‘Talking Heads’ connected through the Internet. Each Talking Head features a Sony EVI-D31 camera with controllable pan/tilt motors for horizontal and vertical movement (figure 1), a computer for cognitive processing (perception, categorisation, lexicon lookup, etc.), a screen on which the internal states of the agent currently loaded in the body are shown, a TV-monitor showing the scene as seen through the camera, and devices for audio in- and output. Agents can load themselves in a physical Talking Head and teleport themselves to another Head by travelling through the Internet. By design, an agent can only interact with another one when it is physically instantiated in a body located in a shared physical environment. The experimental infrastructure also features a commentator which reports and comments on dialogs, displays measures of the ontologies and languages of the agents and game statistics, such as average communicative success, lexical coherence, average ontology and lexicon size, etc. For the experiments reported in this paper, the shared environment consists of a magnetic white board on which various shapes are pasted: colored triangles, circles, rectangles, etc.



Fig. 1. Two Talking Head cameras and associated monitors showing what each camera perceives.

The guessing game The interaction between the agents consists of a language game, called the guessing game. The guessing game is played between two visually grounded agents. One agent plays the role of *speaker* and the other one then plays the role of *hearer*. Agents take turns playing games so all of them develop the capacity to be speaker or hearer. Agents are capable of segmenting the image perceived through the camera into objects and of collecting various

sensory data about each object, such as the color (decomposed in RGB channels), average gray-scale or position. The set of objects and their data constitute a *context*. The speaker chooses one object from this context, further called the *topic*. The other objects form the *background*. The speaker then gives a verbal hint to the hearer.

The verbal hint is an utterance that identifies the topic with respect to the objects in the background. For example, if the context contains [1] a red square, [2] a blue triangle, and [3] a green circle, then the speaker may say something like "the red one" to communicate that [1] is the topic. If the context contains also a red triangle, he has to be more precise and say something like "the red square". Of course, the Talking Heads do not say "the red square" but use their own language and concepts which are never going to be the same as those used in English. For example, they may say "malewina" to mean [UPPER EXTREME-LEFT LOW-REDNESS]. The verbal hint is in this experiment assumed to be transmitted completely accurately.

Based on the verbal hint, the hearer tries to guess what topic the speaker has chosen, and he communicates his choice to the speaker by pointing to the object. A robot points by transmitting in which direction he is looking. The game succeeds if the topic guessed by the hearer is equal to the topic chosen by the speaker. The game fails if the guess was wrong or if the speaker or the hearer failed at some earlier point in the game. In case of a failure, the speaker gives an extra-verbal hint by pointing to the topic he had in mind, and both agents try to repair their internal structures to be more successful in future games.

Agents start with no prior designer-supplied ontology nor lexicon. A shared ontology and lexicon must emerge from scratch in a self-organised process. The agents therefore not only play the game but also expand or adapt their ontology or lexicon to be more successful in future games.

The Conceptualisation Module Meanings are categories that distinguish the topic from the other objects in the context. The categories are organised in discrimination trees where each node contains a discriminator able to filter the set of objects into a subset that satisfies a category and another one that satisfies its opposition. For example, there might be a discriminator based on the horizontal position (HPOS) of the center of an object (scaled between 0.0 and 1.0) sorting the objects in the context in a bin for the category 'left' when $HPOS < 0.5$, (further labeled as [HPOS-0.0,0.5]) and one for 'right' when $HPOS > 0.5$ (labeled as [HPOS-0.5,1.0]). Further subcategories are created by restricting the region of each category. For example, the category 'very left' (or [HPOS-0.0,0.25]) applies when an object's HPOS value is in the region [0.0,0.25].

A distinctive category set is found by filtering the objects in the context from the top in each discrimination tree until there is a bin which only contains the topic. This means that only the topic falls within the category associated with that bin, and so this category uniquely filters out the topic from all the other objects in the scene. Often more than one solution is possible, but all solutions are passed on to the lexicon module.

The discrimination trees of each agent are formed using a growth and pruning dynamics coupled to the environment, which creates an ecology of distinctions. Discrimination trees grow randomly by the addition of new categorisers splitting the region of existing categories. Categorisers compete in each guessing game. The use and success of a categoriser is monitored and categorisers that are irrelevant for the environments encountered by the agent are pruned. More details about the discrimination game can be found in [10].

Verbalisation module The lexicon of each agent consists of a two-way association between forms (which are individual words) and meanings (which are single categories). Each association has a score. Words are random combinations of syllables. When a speaker needs to verbalise a category, he looks up all possible words associated with that category, orders them and picks the one with the best score for transmission to the hearer. When a hearer needs to interpret a word, he looks up all possible meanings, tests which meanings are applicable in the present context, i.e. which ones yield a possible single referent, and uses the remaining meaning with the highest score as the winner. The topic guessed by the hearer is the referent of this meaning.

Based on feedback on the outcome of the guessing game, the speaker and the hearer update the scores. When the game has succeeded, they increase the score of the winning association and decrease the competitors, thus implementing lateral inhibition. When the game has failed, they each decrease the score of the association they used. Occasionally new associations are stored. A speaker creates a new word when he does not have a word yet for a meaning he wants to express. A hearer may encounter a new word he has never heard before and then store a new association between this word and the best guess of the possible meaning. This guess is based on first guessing the topic using the extra-verbal hint provided by the speaker, and on performing categorisation using his own discrimination trees as developed thus far. These lexicon bootstrapping mechanisms have been explained and validated extensively in earlier papers [11].

The conceptualisation module proposes several solutions to the verbalisation module which prefers those that have already been lexicalised. Agents monitor success of categories in the total game and use this to target growth and pruning. The language therefore strongly influences the ontologies agents retain. The two modules are structurally coupled and thus get coordinated without a central coordinator.

Semiotic Dynamics We propose the notion of a *semiotic landscape* (which we also call RMF-landscape) to analyse grounded semiotic dynamics. The semiotic landscape is a graph, in which the nodes in the landscape are formed by referents (objects), meanings (categories) and forms (words), and there are links if the items associated with two nodes indeed co-occur (figure 2). The relations are labeled RM for referent to meaning, MR for meaning to referent, RF for referent to form, FR for form to referent, and FM for form to meaning and MF for meaning to form. The RMF-landscape in figure 2 (taken from the experiment to be discussed later) contains an example where the same object $O3$ is designated by two meanings $[G-0.25,0.5]$ and $[G-0.375,0.5]$. The first meaning, which is more

general, is expressed by two words "xu" and "fepi" and the second meaning by the word "pasi". Usually we see much more complex situations and complexity further increases when the same meaning is also used to denote other referents (which is obviously very common and indeed desirable). We track the changes




Fig. 2. A semiotic landscape represents the co-occurrences between referents, meanings and forms.

in the semiotic landscape by recording the actual verbal behavior of the agents while they engage in language interactions, more specifically, by collecting data on the co-occurrence of items such as the forms used with a certain referent or the meanings used with a certain form. Frequency of co-occurrence is represented in competition diagrams, such as the RF-diagram in figure 6, which plots the evolution of the frequency of the observed referent-form co-occurrences for a given referent in a series of games. Similar diagrams can be made for the FR, FM, MF, RM and MR relations.

3 Case Study

For real world environments, the set of possible referents is infinite, so the semiotic landscape is infinite. For purposes of analysis, we therefore need to restrict the possible environments and thus the possible referents artificially and then study the semiotic dynamics very precisely. For the present paper, we analyse a test run involving 20 agents and 8 objects, which means $4 \times C_8^4 = 280$ possible situations. The run starts with 4 objects (and hence 4 possible situations) and after every 5000 games a new object was introduced. During the final 15000 games, no new objects were introduced. The overall evolution of the dynamics is shown in figure 3. We see that success in communication or discrimination mounts quickly, drops after a new object is introduced, but increases again as the agents develop new concepts and words. We also see that it takes less and less time to absorb new objects, indicating that the language is not situation-specific. And that the system evolves towards maximum communicative and discriminatory success.

The main goal of this paper is to show how an analysis in terms of the semiotic dynamics aids to understand the evolution of the system. We examine one word, "fepi" which is general in the sense that it is clearly used for more than one object (O3 and O5) and in many different contexts (figure 4 left)




Fig. 3. The graph shows the communicative and discriminatory success for a series of 35000 language games.

Closer examination of the meanings (figure 4 right) reveals that "fepi" 'means' a particular shade of green, categorising a green-intensity [0.25,0.5] on the G channel of the RGB data. The category is labeled [G-0.25,0.5]. When examining the forms used to express this category (figure 5) we see that "fepi" has indeed emerged as dominant for this meaning, but that earlier on the word "xu" was dominant, which raises the question how "fepi" managed to overtake "xu" in the long run.

When inspecting in more detail the game traces, we see that "fepi" is created in game 328 by agent-3, playing the role of speaker, in order to refer to object *O3* using the meaning [G-0.25,0.5]. Agent-19 acquires at this point the same meaning for [G-0.25,0.5] as hearer. In one sense, we could say that agent-19 learns this meaning of "fepi" from agent-3 but that is not entirely accurate. Agent-19 constructed a possible meaning for "fepi" and this happened to be the same as the one used by agent-3, but this is accidental. This is a first important observation. Agents only indirectly learn the language from others. They construct a language which is compatible with the language used by others in the situation encountered - and in turn influence by their own use the language of other speakers. Compatibility between individual language systems occurs through the positive feedback between language use and communicative success.

As figure 6 (left) shows, "fepi" is not immediately successful. Instead "xu" wins the initial competition against "fepi" and several other words for designating *O3*. Typically, multiple ways to designate the same object develop but this synonymy is damped due to the lateral inhibition between the different forms competing for the same meaning.

"xu" has also the meaning [G-0.25,0.5] although there are some other meanings associated with "xu" as well, which are also distinctive for *O3*. Because




Fig. 4. Left: FR diagram showing the objects referred to by the word "fepi". After 10000 games "fepi" is consistently used for *O3* and *O5*. Right: FM diagram showing the meanings of "fepi".




Fig. 5. MF diagram showing the different words circulating in the population for expressing the concept [G-0.25,0.5].




Fig. 6. Left: RF diagram showing the different words being used for identifying O_3 . Right: RF diagram showing the different words used for O_5 .

agents only get feedback about reference and not about meaning, words are polysemous until disambiguated by situations in which its different shades of meaning are incompatible.

In game 5000, the arrival of a new object O_5 destabilises the association between "xu" and [G-0.25,0.5]. Closer examination reveals that O_5 's green value is a bit lighter (in the range [0.25,0.375]) than that of O_3 (which is in the range [0.375,0.5]), so that a more refined distinction is necessary if both objects occur in the same context. As seen from the RF-diagram in figure 6 left, "xu" is no longer used for O_3 . Instead the word "pasi" comes to dominate. "pasi" has indeed the more specific meaning [G-0.375,0.5]. At the same time we see from the RF-diagram for O_5 (figure 6 right) that the word "rimebi" dominates for designating O_5 . As expected "rimebi" has the second more specific meaning [G-0.25,0.375]. The more general word "xu" is still useful in contexts where the refined distinction is not necessary and so we would expect that "xu" continues to exist. However this is not the case. "xu" loses out completely and its role is taken over by "fepi". Why is this so? "xu" loses its strength because (1) the main meaning of "xu" is often not distinctive enough and therefore a game fails, and (2) other meanings competing for "xu" gain (as seen from figure 7) pulling down the original green-based meaning due to lateral inhibition. The weakening of "xu" opens the way to "fepi" which still carries the more general meaning of green and does not have competitors. We see from the RM diagram (figure 7 right) how first a general meaning coupled to "xu" is used for O_3 , then a more specific meaning coupled to "pasi" (after game 5000) and then again a more general meaning coupled to "fepi". The more general meaning increases because due to the arrival of several new objects, the more abstract green category becomes more useful again, even though the more specific meaning is still occasionally needed.




Fig. 7. Left: FM diagram showing the different meanings of XU. After game 5000, the meaning of "xu" becomes unclear and the word falls in disrepute. Right: RM diagram showing the meanings used to identify O3.

A similar picture is seen for the meanings used for O5 where "rimebi" and "fepi" are used depending on the degree of distinction required by the context.

4 Conclusion

The main goal of the paper was to analyse a case study of semiotic dynamics based on data recorded from a group of 20 distributed autonomous robotic agents playing language games about real world scenes through visual observations. We see that the word-meaning pairs active in a population show a complex picture. The agents have never exactly the same lexicon but each has their own ideolect. Moreover periods of heavy competition alternate with periods of relative stability. Stability occurs when one word temporarily manages to become dominant in a winner-take-all process. Semiotic dynamics therefore shows similarities to the dynamics exhibited by other types of complex adaptive systems such as punctuated equilibria in species evolution or evolution in adaptive cooperative games [6]. It illustrates the hypothesis that language is an evolving complex dynamical system which self-organises and gets transmitted in a cultural process.

5 Acknowledgement

This work was financed and carried out at the Sony Computer Science Laboratory in Paris. We are indebted to Toshi Doi and Mario Tokoro to be able to work in this superbly productive environment. The BABEL software tool under development by Angus McIntyre at SONY CSL-Paris has proved invaluable for efficiently exploring the issues discussed in this paper. We are also grateful

to Angus McIntyre and Joris Van Looveren from the VUB AI Laboratory for contributions to the specific software underlying this experiment.

References

1. Batali, J. *Computational Simulations of the Emergence of Grammar*. In: Hurford, J. et.al. [ed.] *Approaches to the Evolution of Language*: Edinburgh Univ. Press. Edinburgh, 1998.
2. Belpaeme, T., L. Steels and J. Van Looveren (1998) The construction and acquisition of visual categories. In: Birk, A. (ed.) *Workshop on learning robots*, Brighton. Springer-Verlag, Berlin.
3. De Boer, B. Generating vowel systems in a population of agents. In Phil Husbands and Inman Harvey, eds. *Fourth European Conference on Artificial Life*, Brighton, 1997. MIT Press, Cambridge, MA
4. Hashimoto, T. et Ikegami, T. (1996) Emergence of net-grammar in communicating agents. *BioSystems* 38 (1996) 1-14.
5. Kirby, S. Language Evolution without Natural Selection. From vocabulary to syntax in a population of learners. 2nd Evolution of Language Conference, London, 1998.
6. Lindgren, K. and Nordhal, M. Cooperation and Community Structure in Artificial Ecosystems. In: Langton, C. *Artificial Life: an overview*. MIT Press, Cambridge, MA
7. Oliphant, M. The dilemma of Saussurean communication. *Biosystems*, 1996, 37 [1-2], pp 31-38.
8. Quine, W (1960). *Word and Object*. MIT Press, Cambridge, MA.
9. Steels, L. (1997) The synthetic modeling of language origins, *Evolution of Communication Journal*, 1(1): 1-34.
10. Steels, L. (1997) Constructing and Sharing Perceptual Distinctions. In: van Someren, M. and G. Widmer (eds.) (1997) *Proceedings of the European Conference on Machine Learning*. Springer-Verlag, Berlin.
11. Steels, L. and Kaplan, F. Stochasticity as a source of innovation in language games. In Adami, C., Belew, R., Kitano, H. and Taylor, C. [eds.] *Proceedings of Artificial Life VI*, Los Angeles, June 1998, MIT Press, p. 368-376.
12. Yanco, H. and L. Stein (1993) An Adaptive Communication Protocol for Cooperating Mobile Robots. In: Meyer, J-A, H.L. Roitblat, and S. Wilson (1993) *From Animals to Animats 2. Proceedings of the Second International Conference on Simulation of Adaptive Behavior*. The MIT Press, Cambridge Ma. p. 478-485.