

An architecture for evolving robust shared communication systems in noisy environments

Frédéric Kaplan, Luc Steels and Angus McIntyre

Sony CSL - Paris - 6 Rue Amyot, 75005 Paris

E-mail: { kaplan, steels, angus }@csl.sony.fr

Abstract

Our research aims to provide a means for independent entities to autonomously develop a set of shared conventions which will allow them to communicate with each other. The communication system thus developed needs to be efficient, robust, learnable and tolerant of noise occurring at all stages in the communication process. This paper proposes a system based on repeated interactions, coupled with learning procedures that allows shared communication systems to be developed even in the presence of noise.

1 Introduction

A number of researchers have already addressed the question of self-organization and propagation of lexical conventions in a population of artificial agents. The first reported research on the subject was carried out by [1] and [2]. Both demonstrate that a simple, genetically-encoded signalling system can emerge without central control through selection pressure acting on the system. MacLennan also showed that addition of a simple learning algorithm allowed communication conventions to propagate more quickly. These studies assumed what might be termed a "Saussurean" convention: that the same signal should be used for both reception and transmission of the message. [3] showed that such a convention could emerge naturally by genetic evolution, and [4] has identified constraints which help to determine the emergence of the convention. Experimental studies in the formation of form-meaning associations include both connectionist approaches, such as that of [5], and symbolic, for example [6]. In the latter study, lexical coherence was shown to emerge from a positive feedback loop based on preferential selection of words successfully used in the past.

All these results assume the ideal case of 'perfect' communication, involving – among other things – shared knowledge of the topic. This assumption is very unrealistic with respect to real-world systems where unpredictable errors (stochasticity) may occur throughout the communication process. Errors in signal production, transmission and reception due to inherent unreliability in the sender and receiver's signalling apparatus or in the transmission medium are an unavoidable feature of all real-world communication. Perceptual uncertainty can interfere with correct identification of the communication topic. Experiments carried out in SONY CSL Paris show that algorithms designed for use in conditions of perfect

communication are inefficient in 'noisy' environments. Introduction of stochastic elements into the communication process leads to an immediate drop in communicative success through – for instance – failure to recognise imperfectly-produced or -transmitted forms. For a real-world application, a mechanism capable of tolerating this kind of variation is essential.

2 The naming game model

This paper proposes an interaction model for developing shared conventions in a noisy environment. This model, termed the naming game, is an enriched version of one first presented by [6], redesigned in order to allow it to function when communication is imperfect. Through repeated formalised interactions, a set of agents adaptively construct a shared set of form-meaning associations. The present implementation of the naming game focuses on the association of atomic forms and meanings, but the results reported here, however, do not depend on this simplification, and the mechanisms described are also applicable to systems based on complex forms and meanings.

Using this architecture, shared conventions emerge through repeated interactions between agents. To represent form-meaning associations, each agent has a lexicon which is a time-dependent relation between meanings and forms coupled with a score representing the strength of the form-meaning association. Each agent's lexicon is initially empty; different agents may have different lexicons, and the lexicons support synonymy (use of different forms for the same meaning) and homonymy (different meanings for the same form). The formalised interaction – the naming game – involves a speaker and a hearer communicating about a topic in a given context, which consists of a set of objects. The speaker signals a topic to the hearer using non-linguistic means (such as pointing). At the same time, the speaker retrieves all forms associated with the topic meaning in his lexicon, selects

one form and produces it. The hearer perceives the linguistic and non-linguistic information generated by the speaker, considers the possible forms and meanings evoked by this information, and signals to the speaker the topic identified as a result of this consideration. If the topic signalled by the speaker agrees with that identified by the hearer, the game succeeds.

The key to the architecture is the way that it deals with stochasticity at all stages of the communication process (Figure 1). Stochasticity is modelled by operators affecting accuracy of production, transmission and perception. To cope with this potential noise, the hearer must consider a number of candidate forms and meanings and evaluate each (either in sequence or in parallel). The hearer constructs a meaning score for each possible meaning, reflecting the likelihood that a given meaning is intended. At the same time, a form score is computed for each candidate form, based on the distance between the actual perceived form and the candidate form. Using the associations between form and meaning stored in the lexicon, the hearer constructs a decision-matrix in which the form scores, meaning scores and the strength of each previously-recorded form-meaning association are combined as a weighted sum to produce a score for a possible association between each candidate form and each candidate meaning. The form-meaning association with the highest score is extracted from the matrix, yielding the hearer’s interpretation of the speaker’s communication. The use of the decision matrix allows the hearer to cope with noise introduced at any stage of the communication process. In order for shared conventions to emerge, both speaker and hearer must adapt their internal lexicons based on the result of the interaction, success leading to reinforcement of those associations that led to successful communication and weakening competing associations. Similarly, in the case of failure, the associations responsible for the miscommunication are weakened.

Other repair actions are possible. When failure occurred because one of the agents lacked a necessary form, the agent in question may extend its lexicon to include the required form. When failure was due to a mismatch between the intended and perceived meanings, both speaker and hearer decrement those form-meaning associations in their internal lexicons which were responsible for the failed communication. Repeated independent adaptation of the lexicon leads to the emergence of a robust but flexible shared communication system.

3 Experimental results

In this section we present the results of several experiments to illustrate the performance of this new architecture (a more complete set of experiments can be found in [7]). These experiments were carried out

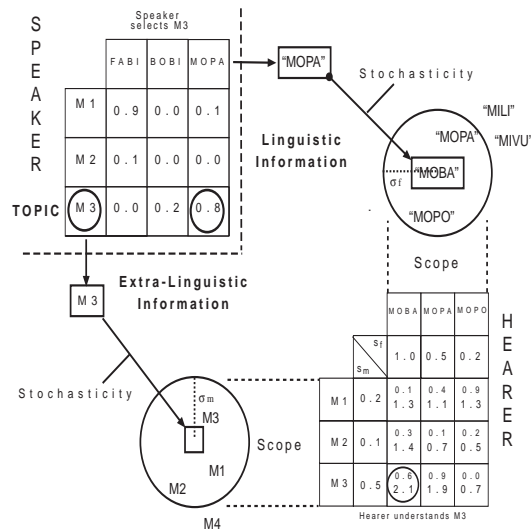


Figure 1: The Naming Game Model: (1) The speaker selects a topic M3 in the context and provides extra-linguistic information about this meaning (for instance by pointing). It also scans its lexicon and picks the form with the highest score: "MOPA". (2) The linguistic and extra-linguistic information are perceived by the hearer. The information may have been altered during transmission (the direction pointed may be shifted, "MOPA" may be perceived as "MOBA" or some other similar variant). (3) The hearer selects a set of possible forms and meanings close to the ones perceived, evaluates each of them and integrates them with its lexicon knowledge in a decision matrix in an attempt to determine the meaning intended by the speaker. (4) The game succeeds because the hearer has identified the correct meaning. Both speaker and hearer update their lexicons.

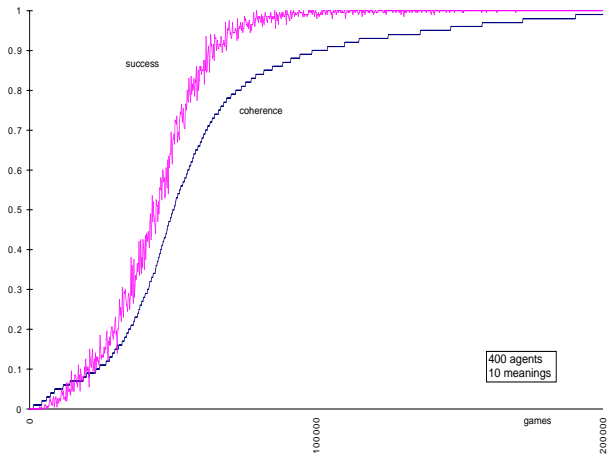


Figure 2: Evolution of average game success and coherence in a population of 400 agents for 10 objects. An equilibrium state is reached in which the agents gain an average success of 100% coupled with a high, stable coherence.

on the Babel simulation platform developed at Sony CSL Paris [8]. In order to study the global order of the system, we follow two macroscopic variables: communicative success and coherence. These variables are invisible to the agents because no agent has a complete overview of the behavior of the group. *Communicative success* quantifies the average success after n games. When average success approaches total success, this must mean that the conventions are sufficiently shared to speak of the emergence of a shared lexicon. But, because a form may have many meanings and the same meaning may be expressed by multiple forms, communicative success does not necessarily mean complete coherence. An agent may very well know a form but prefer not to use it itself. The language of the group is thus seen as being the set of word-meaning associations that are preferred by the largest number of agents. The *coherence* of the language is then equal to the average number of agents that favor the most preferred word-meaning associations.

3.1 Convergence towards equilibrium states

We first investigate the properties of naming games in the ideal case of closed populations of agents without any stochasticity. Figure 2 shows a first simulation experiment involving 400 agents naming 10 objects. We see that coherence and average communicative success both increase until they reach 100%.

The number of games necessary to reach total communicative success grows with the population size. Figure 3 shows several communicative success curves for different population sizes, using an x-axis scale expressed in games/agent (the total number of games divided by the population size). On average, around 150 games/agent are needed to reach the maximum




Figure 3: Several communicative success curves for different population size on a renormalised scale of games/agent. Complete communicative success is achieved even for large populations.

communicative success, which means that each agent has to play around 150 naming games to build a lexicon that will allow it to communicate reliably.

The same kind of studies can be carried out for the size of the object set (the number of meanings). Figure 4 shows that the number of games needed to reach total success grows linearly with set size for meaning sets smaller than 200 items. In this domain, each agent needs to play 10 to 15 naming games involving each meaning of the world to build a vocabulary that supports complete communicative success.

In the rest of the paper, the x-axis scale will be shown in games/agent.meaning. For practical reasons, the experiments only involve 20 agents naming 10 objects, but the results could be generalised to larger populations and meaning sets.

3.2 Resilience to population change

The lexicon built by the agents is resistant (to a certain extent) to changes in the population. This can be shown by introducing an in- and outflux in the population. The individual lexicons of any agents leaving the population are lost when they are removed (thus unique words contained only in those lexicons may be lost when the agent is removed). When new virgin agents enter, they have to acquire the language of the other agents in the group. They may occasionally create a new word (with a small probability, namely the word creation probability p_c) but this new word is generally unable to compete against the dominance of the existing preferred word. Acquisition of the existing language by a new agent happens without any addition or change to the model, as shown in figure 5 which also plots the language change. Change is quantified by comparing the state of the language at two time points and counting the number of preferred form-meaning pairs that changed. We see that




Figure 4: Several communicative success curves for different meaning set size on a renormalised scale of games/agent.meaning. The number of games needed to reach total success grows linearly for meaning sets smaller than 200 items.

the language changes rapidly in the beginning as the population moves towards total average game success, but that thereafter the language remains stable. Figure 5 shows what happens when the population is in a state of continual flux. As new agents come in, game success and coherence drop because the new agent has to acquire the language of the group. But if there are not too many agents coming in, the group will maintain a high success rate. More importantly, the language itself does not change at all. It is transmitted culturally from one generation to the next. When the rate of population renewal is too high, the language disintegrates, also illustrated in Figure 5. There is rapid language change because the new agents start to create new word-meaning associations, but these conventions cannot propagate through the population rapidly enough to become a stable part of the language.

3.3 Stochasticity in non-linguistic communication

Stochasticity in non-linguistic communication can be investigated by probabilistically introducing a random error in the perceived attributes of the topic. The properties of the meaning expressed can, for instance, be shifted by a fixed value. The probability is called the topic-recognition stochasticity E_T . Figure 6 shows the first results for an experiment exploring variations in E_T . When E_T is high (phase one), there is so much confusion that a language does not form at all. When E_T is decreased to 0.0 (phase two), a language starts to form quickly. This language maintains itself, even if E_T is again increased (third phase).

This experiment shows that there must be a minimum level of reliability in non-linguistic communication during the initial phases of language formation

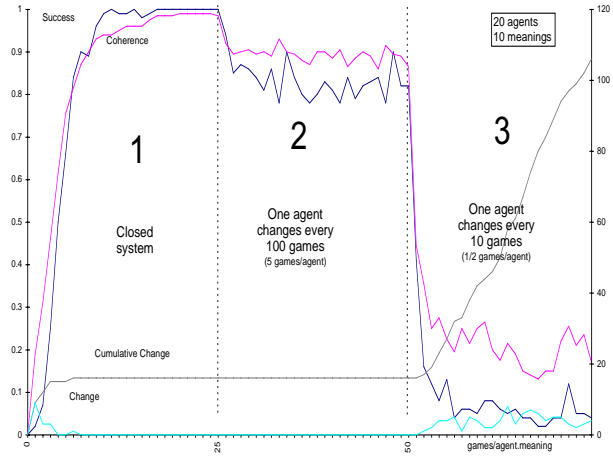


Figure 5: Once formed, a language remains stable even if there is an in- and outflow of agents in the population. This graph shows both language change and the average game success. In the first part, the language forms itself in a closed population. During a second phase, an in- and outflow of agents (1 agent is replaced every 100 games) is introduced, the language remains the same and success is maintained. In the third phase the turnover is increased to 1 agent every 10 games and the language disintegrates. Average game success rapidly falls to very low levels.

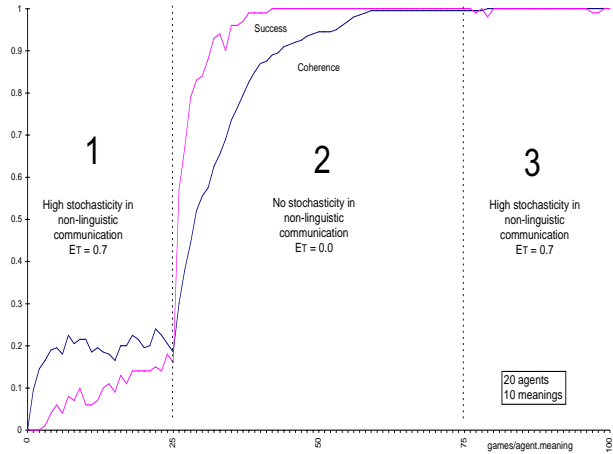


Figure 6: Exploration of variations in the stochasticity of non-linguistic communication. In the first phase stochasticity is high $E_T = 0.7$, a coherent language does not form. In the second phase stochasticity is absent, $E_T = 0.0$, a language forms. In the third phase stochasticity is increased again to $E_T = 0.7$. Communication can tolerate a high level of stochasticity, justifying linguistic communication complementary to non-linguistic communication.




Figure 7: Exploration of variations in form stochasticity. In the first phase stochasticity is high $E_F = 0.5$. A language only forms slowly. In the second phase it is low $E_F = 0.0$, and a language forms. When form stochasticity is reset so that $E_F = 0.5$, the language proves resilient under higher form stochasticity and average game success stays very high.

for a language to form. Once the language has bootstrapped itself, however, linguistic communication is able to overcome the unreliability of non-linguistic communication.

3.4 Stochasticity in form transmission

We now introduce a second stochastic operator that causes a transformation of the form transmitted. For example, the speaker may produce "moba" but the hearer may perceive "mopa". The parameter controlling this stochasticity is E_F , the form-recognition stochasticity: it is the probability that a character in the string of the form will mutate (string mutation can be considered analogous to phonetic distortions introduced by ambient noise or production errors in real world communication).

Figure 7 shows the results of experiments varying this particular parameter. In the first phase $E_F = 0.5$ a language may eventually form but takes rather a long time. $E_F = 0$ causes the language to appear immediately. In the third phase, we again increase the stochasticity. The language is seen to be resilient. Games may occasionally fail, but the language itself is not affected. As with human language users, the combination of non-linguistic communication and expectations based on the lexicon partially offset the problems in determining what form has been used. These experiments clearly show that once a language has formed, it counterbalances errors in message transmission.

4 Potential applications

In this section, we present two possible applications of the naming game architecture. Neither has, as yet,

been the subject of a research project.

4.1 Speech interaction with devices

The naming game can be embedded in consumer electronic devices to enable interaction through speech. Of course in this case the necessary speech modules need to be integrated for recognising and producing speech sounds. The role of the system proposed in this paper is to allow users to teach their own vocabulary for controlling devices, so that the interface becomes truly adapted whatever the source language.

4.2 Adaptive protocols in the Internet

Agents need a shared set of conventions to communicate. These conventions can be determined for a particular domain and take the form of norms that all agent designers respect. Major standardization initiatives, such as the definition of the KQML language [9], have made it reach world-wide consensus on agent interaction protocols. But for continuously growing and not centrally controlled networks such as the Internet, global standards are difficult to define and maintain. Once these conventions are fixed, they cannot evolve further. This lack of adaptivity may turn out to be an important drawback for agents interacting in an open environment where new situations and requirements can arise. One potential application of the present research is the definition of *adaptive protocols* that allow agents to collectively build a shared set of conventions [10]. With this technology agents would continuously adapt in order to communicate with one another, without the need of any central controlling agency.

5 Conclusions

The architecture presented in this paper allows shared communication systems to be developed even in the presence of noise (i.e. errors in production, transmission or reception of signals). This is an essential requirement for real-world applications, where perfect communication cannot be assumed. The use of a weighted decision matrix allows noise-tolerant communication and learning, while reinforcement learning techniques establish a stable set of shared conventions.

Initial embodiment of the architecture is in the form of a software system using symbolic representations. The nature of the representations used by the system also makes it suitable for implementation as part of an embedded system based on purpose-built hardware.

6 Acknowledgement

This research was carried out at the Sony Computer Science Laboratory in Paris. We are indebted to Mario Tokoro and Toshi Doi for the opportunity to

work in this superb research environment. And to Mario Tokoro in particular for suggesting that the role of stochasticity might play an important role in making the linguistic system more robust.

References

- [1] B. MacLennan. Synthetic ethology: An approach to the study of communication. In C. Langton, editor, *Artificial Life II*, Redwood City, Ca., 1991. Addison-Wesley Pub. Co.
- [2] G. M. Werner and M. G. Dyer. Evolution of communication in artificial organisms. In C. G. Langton, C. Taylor, and J.D. Farmer, editors, *Artificial Life II, Vol.X of SFI Studies in the Sciences of Complexity*, Redwood City, Ca., 1991. Addison-Wesley Pub.
- [3] J. Hurford. Biological evolution of the saussurean sign as a component of the language acquisition device. *Lingua*, 77:187–222, 1989.
- [4] M. Oliphant. The dilemma of saussurean communication. *Biosystems*, 1–2(37):31–38, 1996.
- [5] E. Hutchins and B. Hazlehurst. How to invent a lexicon: the development of shared symbols in interaction. In N. Gilbert and R. Conte, editors, *Artificial Societies: The Computer Simulation of Social Life*. UCL Press, 1995.
- [6] L. Steels. Self-organizing vocabularies. In C. Langton, editor, *Proceeding of Alife V*, Nara, Japan, 1996.
- [7] L. Steels and F. Kaplan. Stochasticity as a source of innovation in language games. In C. Adami, R. Belew, H. Kitano, and C. Taylor, editors, *Proceedings of Artificial Life VI*, pages 368–376, Los Angeles, June 1998. MIT Press.
- [8] A. McIntyre. Babel: A testbed for research in origins of language. In *Proceedings of Coling-ACL 98*, pages 830–835, Montreal, August 1998.
- [9] Arpa Knowledge Sharing Initiative. Specification of the kqml agent-communication language. July 1993.
- [10] L. Steels. The origins of ontologies and communication conventions in multi-agent systems. *Journal of Agents and Multi-Agent Systems*, 1(1), 1997.