Stochasticity as a Source of Innovation in Language Games

Luc Steels (1,2) and **Frédéric Kaplan (1,3)**
(1) Sony CSL - Paris - 6 Rue Amyot, 75005 Paris
(2) VUB AI Lab - Brussels
(3) LIP6 - UPMC - 4, Place Jussieu F-75252 Paris
E-mail: `steels@arti.vub.ac.be`

Abstract

Recent work on viewing language as a complex adaptive system has shown that self-organisation can explain how a group of distributed agents can reach a coherent set of linguistic conventions and how such a set can be preserved from one generation to the next based on cultural transmission. The paper continues these investigations by exploring the presence of stochasticity in the various aspects of lexical communication: stochasticity in the non-linguistic communication constraining meaning, the transmission of the message, and the retrieval from memory. We show that there is an upperbound on the amount of stochasticity which can be tolerated and that stochasticity causes and maintains language variation. Results are based on the further exploration of a minimal computational model of language interaction in a group of distributed agents, called the naming game.

Keywords: origins of language, evolution of language, self-organization.

1 Introduction

Exciting recent research in the origins and evolution of language (see overviews in [Hurford *et al.*, 1998] and [Steels, 1997c]) is showing that when language is viewed as a complex adaptive system, it becomes possible to understand how a set of distributed agents is capable to reach a shared set of conventions, even if there is no global controlling agency or prior design. The main mechanism responsible for the emergence of coherence is self-organisation: A positive feedback loop causes some naturally occurring variation to propagate and eventually dominate the population. This is similar to how a product comes to dominate a market in increasing-returns economics [Arthur, 1996], or how a group of social insects like an ant society can form a collective structure [Deneubourg, 1977]. In each of these cases, the system locks globally into specific choices based on positive feedback loops coupled to environmental conditions.

A coherent framework to study language as a complex adaptive system is to define populations of agents engaged in adaptive language games. Each game involves a linguistic as well as a non-linguistic interaction. The agents have feedback about success and failure and adapt so as to be more successful in future games. We have extensively experimented with a particular type of such a game, called the naming game, first introduced in [Steels, 1996b]. The game is played between a speaker and a hearer, randomly drawn from a population of agents. The speaker attempts to identify an object to the hearer, based on pointing and based on using a name. The game succeeds if the hearer guesses correctly the object chosen by the speaker. A speaker may create a new name when he does not have one yet. A hearer may adopt the name used by a speaker. Both monitor use and success and prefer in future games those names that had the highest score. This generates the desired positive feedback loop bringing the group progressively towards global coherence.

The naming game has been explored through computational simulations and is related to systems proposed and investigated by [MacLennan, 1991], [Werner and Dyer, 1991], and [Oliphant, 1996]. We have developed more complex variations of the game where the meaning consist of symbolic descriptions derived from discrimination games [Steels, 1997a]. The game has also been implemented on physically grounded mobile robotic agents [Steels and Vogt, 1997] and on vision-based robotic 'talking heads', watching dynamically evolving scenes [Steels, 1997b]. Of course in natural languages both the form and the meaning are vastly more complex than the atomic forms (words) and meanings (objects) used in the naming games discussed in this paper. However, the basic properties of naming games are independent of the complexity of the forms or the meanings.

The main topic of this paper is to explore what happens when stochasticity is introduced in language games. Stochasticity means that some aspects of the game exhibit unpredictable errors. It is caused by faults in production or perception, errors in guessing meaning from the context or from pointing, or malfunctioning of memory. We have experienced this stochasticity very strongly while grounding the language games on physical robots, but want to study theoretically its consequences through software simulation. In order to cope with stochasticity, perception delivers typically several possible forms and possible meanings with various degrees of confidence. The selection and the evaluation of these forms and meanings are determined by the Tolerance level and the Focus of the Hearer. Stochasticity, Tolerance and Focus interact and are important for explaining innovation and evolution. This paper focuses on stochasticity, whereas a companion paper explores the role of tolerance, focus and stochasticity in language change [Steels and Kaplan, 1998].

The rest of the paper has the following sections. First the naming game is defined. Then results are shown for the emergence of a set of conventions without any stochasticity. Next different sources of stochasticity are introduced: first in the extra-linguistic activities delimiting the context and the topic, second in terms of noise on the message being transmitted, and finally memory access. Some conclusions and suggestions for further work end the paper.

2 The Naming Game Model

The Naming Game, as used in the present paper, is an enriched version of a model first presented in [Steels, 1996b]. We assume a set of *agents* $\mathcal{A}$ where each agent $a \in \mathcal{A}$ has contact with a set of *objects*. These objects constitute a set of *meanings* to be expressed $\mathcal{M} = \{m_1, ..., m_n\}$. All the experiments in this paper involve a population of 20 agents and 10 meanings. A *form* is a sequence of letters drawn from a finite alphabet. The agents are all assumed to share the same alphabet. A *lexicon* $\mathcal{L}$ is a time-dependent relation between meanings, forms and a score. Each agent $a \in A$ has his own set of forms $F_{a,t}$ and his own lexicon $L_{a,t} \subset \mathcal{M}_a \times F_{a,t} \times \mathcal{N}$, which is initially empty. An agent a is therefore defined at a time t as a pair $a_t =< F_{a,t}, L_{a,t} >$. There is the possibility of synonymy and homonymy: an agent can associate a single form with several meanings and a given meaning with several forms. It is not required that all agents have at all times the same set of forms and the same lexicon.

2.1 Operation of the Naming Game

The Naming Game is an interaction between a Speaker and a Hearer about a Topic in a given Context. The context consists of a set of objects and both the speaker and the hearer are assumed to be capable to identify meanings to distinguish the topic from the other objects in the context, using for example mechanisms as described in [Steels, 1996a].

2.1.1 Production. Let $C \subseteq \mathcal{M}$ with $\mathcal{M}$ the set of possible meanings. The meaning the speaker has associated with the topic is $m_s \in C$. He signals this topic using non-linguistic communication (such as through pointing). At the same time, the speaker retrieves from his lexicon all the associations indexed by m_s. This set is called the association-set of m_s. Let $m \in \mathcal{M}$ be a meaning, $a \in \mathcal{A}$ be an agent, and t a time moment, then the association-set of m is

$$A_{m,a,t} = \{< m, f, u > | < m, f, u > \in L_{a,t}\} \tag{1}$$

Each of the associations in this set suggests a form f_s to use for identifying m with a score $0.0 \leq u \leq 1.0$. The speaker chooses the association with the largest score and produces the form f_s which is part of this association to the hearer.

2.1.2 Transmission. Both linguistic and non-linguistic information are transmitted to the hearer. During the emission, transmission and reception phases, *stochasticity* (e.g. noise, unpredicable errors) can occur.

2.1.3 Comprehension. The Hearer perceives the linguistic and non-linguistic information. Because this information might have been altered during the transmission, the hearer must consider several possible forms and meanings and evaluates each of them. The form Focus FF and the Meaning Focus TF parameters determine the number of forms F_{cons} and meanings M_{cons} considered. These parameters indicate the maximum distance from the perceived information that the hearer is willing to consider:

$$M_{cons} = \{< m > | < m > \in \mathcal{M}, d(m, m') \leq TF\} \tag{2}$$

$$F_{cons} = \{f'\} \cup \{< f > | < f > \in \mathcal{F}, d(f, f') \leq FF\} \tag{3}$$

a. Meaning score. The hearer constructs a meaning-score $0.0 \leq s_m \leq 1.0$ for each possible meaning m in M_{cons} reflecting the likelihood that m is the meaning of the perceived topic m'. If there is absolute certainty, one meaning has a score of 1.0 and the others are all 0.0. If there is no non-linguistic communication, the likelihood of all meanings is the same. If there is only vague non-linguistic communication, the hearer has some idea what the topic is, but with less certainty. In our experiments, the distance $d(m', m)$ between the meaning of the perceived topic m' and the other meanings determines the meaning-score:

$$s_m = \frac{1}{1 + (\frac{d(m',m)}{\alpha})^2} \tag{4}$$

α is the tolerance factor for meaning perception.

b. Form score. The hearer constructs also a form-score $0.0 \leq s_f \leq 1.0$ for each form f of F_{cons}. The distance $d(f', f)$ between the perceived form f' and the considered form f gives a score

$$s_f = \frac{1}{1 + (\frac{d(f',f)}{\beta})^2} \tag{5}$$

β is the is the tolerance factor for form perception.

c. Decision matrix. For each form f_j in F, the hearer retrieves the association-set that contains it. He constructs a *decision-matrix* which contains for each meaning a row and for each form a column. The first column contains the meaning-scores s_{m_i}, the first row the form-scores s_{f_j}. Each cell in the inner-matrix contains the association-score for the relation between the object and the form in the lexicon of the hearer:

		f_1	f_2	...
		s_{f_1}	s_{f_2}	...
m_1	s_{m_1}	$s_{<m_1,f_1>}$	$s_{<m_1,f_2>}$	...
m_2	s_{m_2}	$s_{<m_2,f_1>}$	$s_{<m_2,f_2>}$	...
...	...	...	...	...

Obviously many cells in the matrix may be empty (and then set to 0.0), because a certain relation between a meaning and a form may not be in the lexicon of the hearer. Note also that there may be meanings identified by lexicon lookup which are not in the initial context C. They are added to the matrix, but their meaning-score is 0.0.

The final state of an inner matrix cell of the score matrix is computed by the formula:

$$score_{m_i,f_j} = w_f.s_{f_j} + w_m.s_{m_i} + w_a.s_{<m_i,f_j>} \tag{6}$$

w_f is the weight of the form information, w_m is the weight of the non-linguistic information and w_l is the weight of the lexicon. In this paper, they are by default set at 1.0 and the score is then simply the sum of the three sources of information.

One meaning-form pair will have the best score and the corresponding meaning is the topic m_h chosen by the hearer. The association in the lexicon of this meaning-form pair is called the winning association. This choice integrates extra-linguistic information (the meaning-score), form ambiguity (the form-score), and the current state of the hearer's lexicon (the association-score).

2.1.4 Adaptation. The hearer then indicates to the speaker what topic he identified. In real-world language games, this could be through a subsequent action, like handing the topic to the hearer, or through another linguistic interaction. When a decision could be made and $m_h = m_s$ the game succeeds, otherwise it fails. The following adaptations take place by the speaker and the hearer based on the outcome of the game:

a. The game succeeds This means that speaker and hearer agree on the topic. To reenforce the lexicon, the speaker increments the score s of the association that he preferred, and hence used, with a fixed quantity δ. The hearer reenforces the winning association that has led to the right comprehension. Both decrement with δ the score of all the associations that share either the meaning or the form of the winning pair. 0.0 and 1.0 remain the lower and upperbound of s. These changes implement an excitation-exhibition dynamics similar to the one used in Kohonen networks, except that the change is constant.

b. The game fails There are several cases:

1. The Speaker does not know a form

   It could be that the speaker did not have an association covering the topic. In that case, the game fails but the speaker may create a new form f' and associate this with the topic m_s in his lexicon. This happens with a form creation probability p_c.

2. The hearer does not know the form.

   In other words there is no association in the lexicon of the hearer involving the form f_h of the winning association. In that case, the game ends in failure but the hearer may extend his lexicon with a form absorption probability p_a.

3. There is a mismatch between m_h and m_s.

   In this case, both speaker and hearer have to adapt their lexicons. The speaker decrements with δ the association (m_s, f_s) and the hearer decrements with δ the association (m_h, f_h)

2.2 Macroscopic variables

The naming game model can be viewed as a complex dynamical system. The agents have a certain local behavior (an agent can only interact with one single agent, not with all agents at the same time), which is determined by their internal lexicons. Behavior changes because agents adapt their lexicon. In order to 'see' the global order in the system, we need macroscopic variables. These macroscopic variables are invisible to the agents because no agent has a complete overview of the behavior of the group. The first such variable 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 can very well know a form but prefer not to use it himself.

In practice, an examination of the lexicons of the different agents shows a quite complex situation, so that it is non-trivial to extract what *the* shared language is. We determine the language of a single agent by translating his lexicon into a matrix where there is a row for every possible meaning, a column for every word, and the cells are filled by the association-scores, possibly 0.0. Based on this matrix it is possible to determine what the most preferred form is for naming a meaning, and thus what the preferred language is of the agent. Note that this represents the language for production, not for comprehension. Associations are not symmetrical. For example, in the matrix below, the agent prefers f_2 when m_1 needs to be expressed (and not f_1). When f_1 is heard, the same agent nevertheless expects m_1. The pair $< m_1, f_1 >$ is in the expected language but not in the produced language. The expected language includes the production language but not vice-versa. (See [Hurford, 1989] and [Oliphant, 1996] for a further exploration of the coordination between production and comprehension systems.)

	f_1	f_2
m_1	0.6	0.7
m_2	0.4	0.3

Given the preferred language for a single agent, it is straightforward to determine the language of the group as being the set of word-meaning associations that are preferred by most agents. The *coherence* of the language is equal to the average number of agents that prefer these most preferred word-meaning association.

3 Formation and Maintenance of Equilibrium States

We now investigate the behavior of naming games in the ideal case of closed populations of agents without any stochasticity. Agents take the information they perceive at face value ($TF = 0$ and $FF = 0$). $\delta = 0.2$ and new associations are created with an initial score of 0.2. $p_c = 0.1$ and $p_a = 1.0$.

For each experiment, it is instructive to look at the evolution of game success as well as coherence. Figure 1 shows a first simulation experiment involving a group of 20 agents. We see that very quickly coherence as well as average game success climbs up to both reach 100 %.

It is also instructive to look at the evolution of the average association-scores competing for the preferred expression of a particular form (or alternatively for the highest expectation). This is done through *competition diagrams* as the one shown in figure 2. The diagram shows that there is a winner-take-all situation. This is due to the positive feedback loop between score and use. The higher the score of an assocation, the more it is used, and the more its chances increase to be successful in further use. Such a winner-take-all situation takes place for every meaning so that a global shared lexicon emerges.

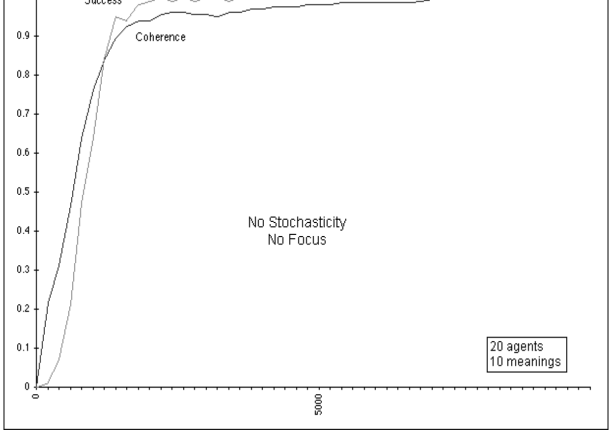

Figure 1: Evolution of average game success and coherence in a population of 20 agents for 10 objects. An equilibrium state is reached whereby the agents gain total average success and a high, stable coherence.

Once total game success is reached, the language does not change anymore. The only source of possible innovation is the introduction of new forms, which only happens when an agent does not have a form yet, or the progressive adoption of one form by the group, which stops as soon as a winner-take-all situation has been reached.

A language is even resistent (up to a certain degree) to changes in the population. This is investigated by introducing an in- and outflux in the population. When agents leave, they take their lexicons with them. 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 the word creation probability p_c) but this new word quickly gets damped against the dominance of the preferred word. Acquisition of an existing language by a new agent happens without any addition or change to the model, as shown in figure 3 which plots also 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 the language changes rapidly in the beginning as the population moves towards total average game success. Thereafter the language remains stable. Figure 3 shows what happens when a flux is introduced in the population. When new agents come in, game success and coherence drops 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 rate of success. More importantly, the language itself does not change at all. It is transmitted culturally from

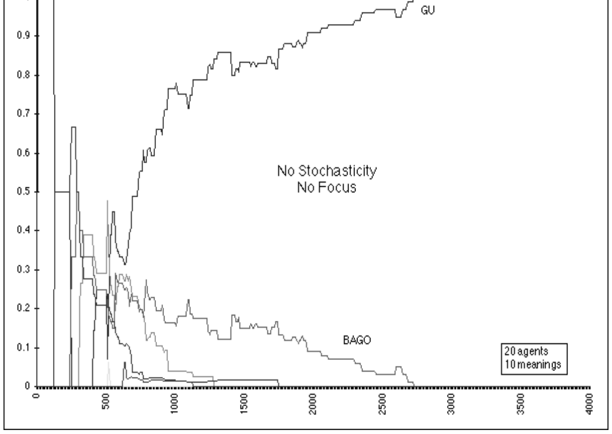

Figure 2: Competition diagram showing the competition between several forms for being the preferred way to express a certain meaning. The diagram plots the average renormalized score of the form-meaning associations of all agents for the same meaning. A winner-take-all situation emerges.

one generation to the next. When the rate of population renewal is too high, the language disintegrates, as also shown in figure 3. There is rapid language change because the new agents start to create new word-meaning associations, but these conventions cannot propagate fast enough in the population.

We will now look at the effect of stochasticty during three steps of the Naming Game: non-linguistic communication, form transmission, and memory access. To cope with stochasticity, agents have now a large focus and a standard tolerance level ($TF = 10$, $FF = 3$, $\alpha = \beta = 1$).

4 Stochasticity of Non-Linguistic Communication

In the results reported so far, it is assumed that non-linguistic communication is without error. This is clearly not always the case in real-word language interactions. Stochasticity in non-linguistic communication can be investigated by probabilistically introducing a random error in the perceived attributes of the topic. The object coordinates of the meaning expressed can, for instance, be shifted by a fixed value. The probability is called the topic-recognition stochasticity E_T. Figure 4 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

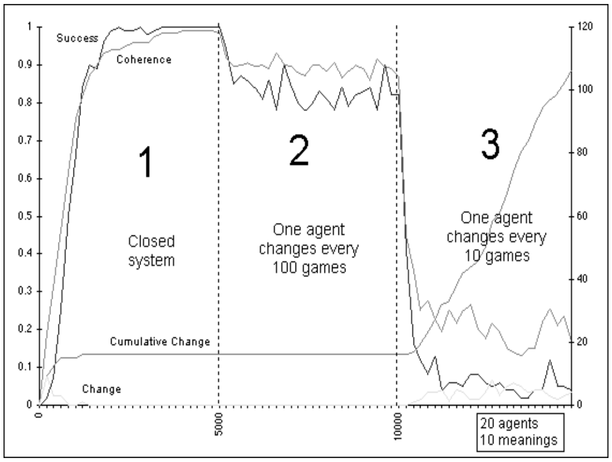

Figure 3: A language once formed remains stable even if there is an in- and outflow of agents in the population. This graph shows both language change and the average game success. In a first phase, the language forms itself in a closed population. In a second phase, an in- and outflow of agents (1 in/outflow per 100 games) is introduced, the language stays the same and success is maintained. In the third phase the flux is increased to 1 per 10 games and the language disintegrates. Average game success rapidly reaches very low levels.

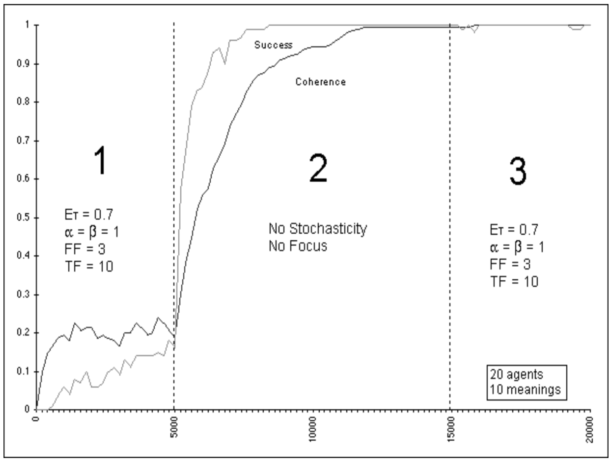

Figure 4: Exploration of variations in the stochasticity of non-linguistic communication. In the first phase stochasticity is high $E_T = 0.7$, a coherent language does not form. In the second phase stochasticity is absent, $E_T = 0.0$, a language forms. In the third phase stochasticity is increased again to $E_T = 0.7$. Communication can tolerate a high level of stochasticity, justifying linguistic communication complementary to non-linguistic communication.

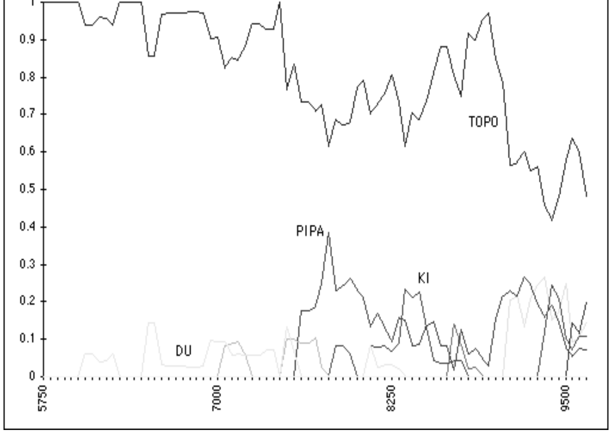

Figure 5: Example of competition diagram with $E_T = 0.7$. There is a more complex dynamics instead of an equilibrium winner-take-all situation. New associations enter and remain in the population, even though they have not yet been able to overtake the dominating word.

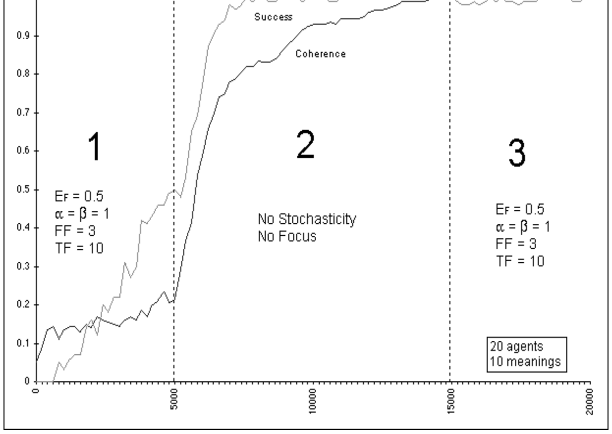

Figure 6: Exploration of variations in form stochasticity. In the first phase stochasticity is high $E_F = 0.5$. A language only slowly forms itself. In the second phase it is low $E_F = 0.0$, a language forms. Then $E_F = 0.5$. The language is resilient against a higher form stochasticity and average game success stays very high.

of reliability in non-linguistic communication at the initial phases of language formation, otherwise a language does not form. At the same time, it shows clearly that as soon as a language has bootstrapped itself, linguistic communication is capable to counteract the unreliability of non-linguistic communication.

We now investigate in how far the stochasticity of non-linguistic communication has an impact on language variation. Figure 5 shows a typical example of a competition diagram for a positive topic-recognition stochasticity. A language has already formed itself with a single winner (the form "topo") for the meaning being investigated. Stochasticity causes competition to arise, challenging - but not yet defeating - the dominating form-meaning association. Innovation is due to the fact that confusion about the topic may lead to a new form-meaning association which then starts to propagate.

5 Stochasticity on Form

The second source of stochasticity is in the message transmission process. So far it is assumed that a message is produced perfectly by the speaker and received perfectly by the hearer. This is well known not to be true in the case of natural languages. Speakers make a large amount of errors and blur the pronounciation to minimize energy and maximise the number of sounds transmitted in a given period of time. Hearers have a very hard time to decode speech signals, simply because the speech signal is noisy and contains only hints for some sounds. Hearers are known to partially make up for it by expectations and knowledge about the language.

In the experiment reported earlier, several forms may already be triggered due to the large focus on form reception. These are all the forms that are at a certain distance from the form produced by the speaker. This uncertainty makes it less clear what form has been used but does not yet imply mistakes. 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 receive "mopa". The parameter controling this stochasticity is E_F, the form-recognition stochasticity: it is the probability that a character in the string of the form mutates.

Figure 6 shows results of experiments in varying this particular parameter. In the first phase $E_F = 0.5$ a language may eventually form itself but it would take a rather long time. $E_F = 0$ immediately causes the language to appear. In the third phase, we again increase the stochasticity. It is seen that the language is resilient. There are occasionally games that fail, but the language itself is not affected. As with human language users, the non-linguistic communication as well as expectations from 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.

Figure 7 investigates the impact of form-stochasticity on variation in the language. We see clearly that a positive form-recognition stochasticity $E_F = 0.3$ causes new forms to appear in the language. When $E_F = 0.0$, many of these forms disappear. Interestingly enough, competitors may still maintain themselves in the population. This is due to the large focus of the agents. "ludo" and "mudo" are words that are sufficiently close to each other that one group may have adopted one form and another

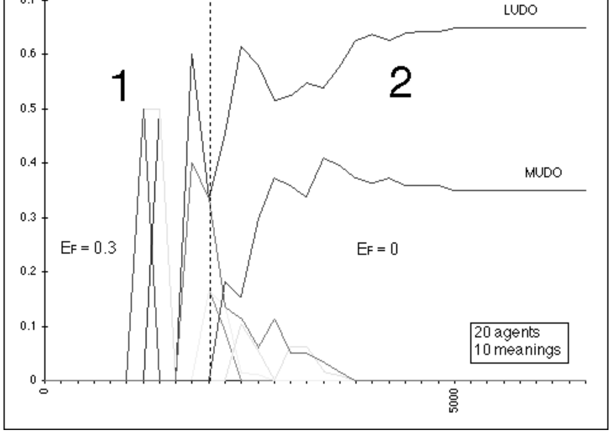

Figure 7: Competition diagram in the presence of form stochasticity. When $E_F = 0.3$ there is no clear winner-take-all situation as new forms are occasionally introduced resulting in new form-meaning associations. When $E_F = 0.0$ the innovation dies out although some form are still able to maintain themselves due to the large focus of the agents.

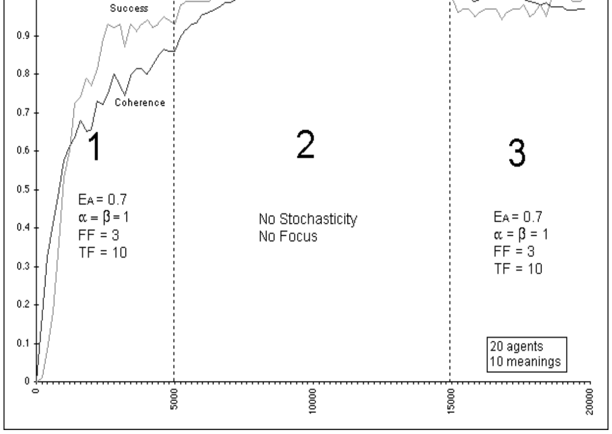

Figure 8: This figure shows the results in exploring memory malfunction. The first phase shows that a language has some difficulties forming for a positive memory-stochasticity ($E_A = 0.7$). In the second phase memory stochasticity is zero and language forms. The third phase shows resilience against positive memory stochasticity ($E_A = 0.7$)

slightly smaller group the other form. As uncertainty in form is tolerated, one group will always accept the form of the other even though they would not use exactly the same form themselves.

6 Stochasticity on Form-Meaning Associations

The final source of stochasticity comes from the utilisation of the lexicon. It is well known that biological systems occasionally malfunction even though there is globally a robust behavior. We hypothesise that this is also the case for memory. The form-meaning association retrieved from memory may not necessarily be the way that it was first stored. Thus the speaker could accidentally retrieve the wrong form for a particular meaning, or the hearer's memory system may suggest a form-meaning association which was never stored. These errors are modeled using a third stochastic operator based on a parameter E_A, the memory stochasticity, which alters the scores of the associations in the score matrix in a probabilistic fashion. Even scores that were zero could become positive. The higher the memory stochasticity, the more likely an association score changes.

Figure 8 shows the impact of memory stochasticity on language formation. When E_A is positive, language formation is more difficult, although progress can be seen. We see also that coherence and success can be maintained even if memory is malfunctioning and yielding spurious association scores. This experiment demonstrates again that the overall language system is fault tolerant because it maximises information from three sources: non-linguistic communication, form recognition, and form-meaning conventions. The better a language is established, the more resilient it is to use in difficult circumstances.

Also in this case, we see continued language innovation due to stochasticity. This is illustrated clearly in a competition diagram shown in figure 9, running for the same simulation as figure 8. One form (”pi”) is dominant for the meaning being investigated. When the memory stochasticity becomes positive, the competition intensifies and new words (”te”, ”lavi”) enter.

7 Combination of Stochasticity

The different forms of stochasticity combined lead to innovation in different areas as seen in figure 10 and figure 11.

8 Conclusion

This paper has investigated the effect of stochasticity on linguistic and non-linguistic communication, as it unfolds in a population of distributed agents playing adaptive language games. This was done for the three main components of a language game: the non-linguistic communication, which constrains the set of possible meanings, the message itself, and the use of the lexicon. Each of these sources of stochasticity is realistic from the viewpoint of real world language use. Human users (as well as robots) cannot be expected to guess accurately the possible meaning of an utterance purely based on non-linguistic means. The message is often errorful due to the inherent unreliability of sound-based message transmis-

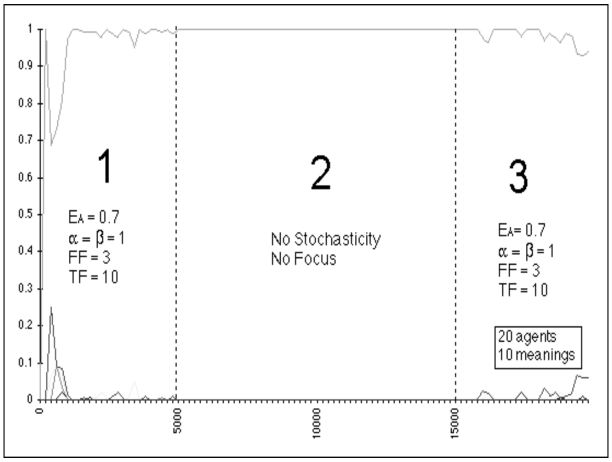

Figure 9: Typical competition diagram. First there is a positive association stochasticity $E_A = 0.7$. The language is slowly forming itself. We still see a rich competitive dynamics between the different form-meaning pairs even if one is dominating. This innovation dies out when $E_A = 0$. When $E_A = 0.7$ new associations get into the system and variation is maintained.

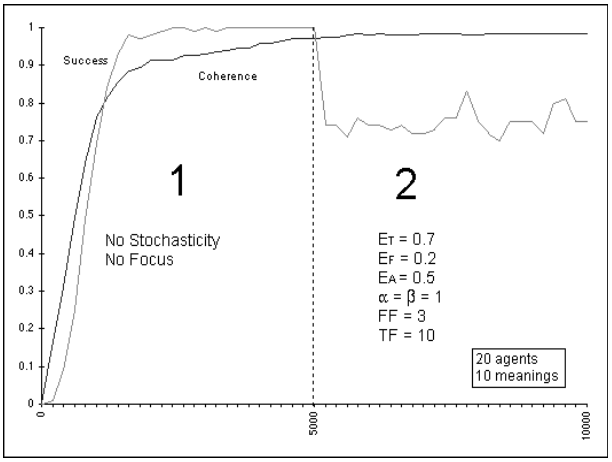

Figure 10: Average success and coherence. The three different forms of stochasticity are introduced in phase 2. We can see that average success drops because incompatibilities between non-linguistic and linguistic communication and the introduction of new forms. Notice that the language coherence remains unaffected.

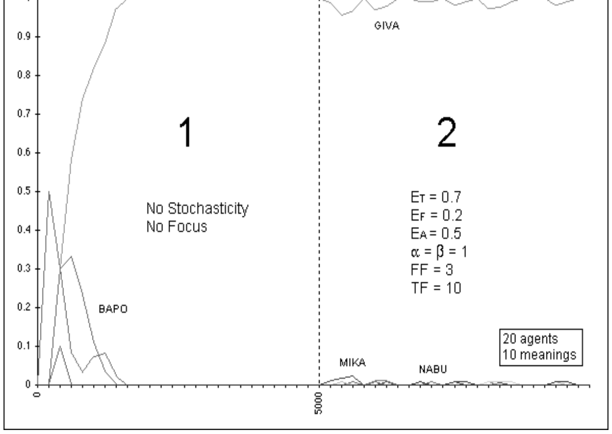

Figure 11: Competition diagram for one meaning during the same experiment as in 10. We see that there is a winner-take-all situation in the first phase and then an attack by new forms ("mika", "nabu", etc.) even though they never manage to overtake the existing word ("giva").

sion. The brain, as many biological systems, may have unreliable components but nevertheless show global fault tolerance.

In each of the cases that were studied, the effect of stochasticity was similar and can be summarised as follows:

1. There are upper bounds on the amount of stochasticity that can be present during the initial phases of language formation. When the stochasticity is too high a language cannot self-organise.

2. Once a language has established itself, stochasticity for one component is partially counterbalanced by the other components. If the message is scrambled, non-linguistic communication and expectations from the lexicon can make up for it. If non-linguistic communication is unreliable or absent, linguistic communication can suffice. If memory is malfunctioning, clues from the environment may counterbalance.

3. Stochasticity introduces and maintains variation in the language. There is no longer a clear winner-take-all situation, whereby the language stays in an equilibrium state, even in a changing population. Instead, there is a rich dynamics where new forms appear, new associations are established, and the domination pattern of associations changes. The different sources of stochasticity each innovate in their own way: Topic stochasticity introduces new form-meaning associations for existing forms. Form stochasticity introduces new forms and hence potentially new form-meaning associations. Memory stochasticity shifts the

balance among the form-meaning associations competing for the expression of the same meaning. All of these sources of stochasticity are clearly observed in real natural language use.

A complementary paper explores how stochasticity and tolerance are essential for explaining language change [Steels and Kaplan, 1998].

9 Acknowledgement

The research described in this paper was financed and conducted at the Sony Computer Science Laboratory in Paris. The simulations presented have been built on top of the BABEL toolkit developed by Angus McIntyre [McIntyre, 1998] of Sony CSL. Without this superb toolkit, it would not have been possible to perform the required investigations within the time available. We are also indebted to Mario Tokoro of Sony CSL Tokyo for continuing to emphasise the importance of stochasticity in complex adaptive systems.

References

[Arthur, 1996] B. Arthur, editor. *The Economy as a Complex Adaptive System.* Santa Fe Institute Series on Complexity. Addison-Wesley, Menlo Park, Ca, 1996.

[Deneubourg, 1977] J-L. Deneubourg. Application de l'ordre par fluctuations à la description de certaines étapes de la construction du nid chez les termites. *Insectes Sociaux*, 24(2):117–130, 1977.

[Hurford *et al.*, 1998] J. Hurford, C. Knight, and M. Studdert-Kennedy, editors. *Evolution of Human Language.* Edinburgh University Press, Edinburgh, 1998.

[Hurford, 1989] J. Hurford. Biological evolution of the saussurean sign as a component of the language acquisition device. *Lingua*, 77:187–222, 1989.

[MacLennan, 1991] 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.

[McIntyre, 1998] A. McIntyre. Babel: A testbed for research in origins of language, 1998. Submitted for Publication.

[Oliphant, 1996] M. Oliphant. The dilemma of saussurean communication. *Biosystems*, 1–2(37):31–38, 1996.

[Steels and Kaplan, 1998] L. Steels and F. Kaplan. Explaining language evolution, 1998. Submitted for Publication.

[Steels and Vogt, 1997] L. Steels and P. Vogt. Grounding adaptive language games in robotic agents. In I. Harvey and P. Husbands, editors, *Proceedings of the 4th European Conference on Artificial Life*, Cambridge, MA, 1997. The MIT Press.

[Steels, 1996a] L. Steels. Perceptually grounded meaning creation. In M. Tokoro, editor, *Proceedings of the International Conference on Multi-Agent Systems*, Cambridge, Ma, 1996. The MIT Press.

[Steels, 1996b] L. Steels. Self-organizing vocabularies. In C. Langton, editor, *Proceeding of Alife V*, Nara, Japan, 1996.

[Steels, 1997a] L. Steels. Constructing and sharing perceptual distinctions. In M. van Someren and G. Widmer, editors, *Proceedings of the European Conference on Machine Learning*, Berlin, 1997. Springer-Verlag.

[Steels, 1997b] L. Steels. The origins of syntax in visually grounded robotic agents. In M. Pollack, editor, *Proceedings of the 15th International Joint Conference on Artificial Intelligence*, Los Angeles, 1997. Morgan Kauffman Publishers.

[Steels, 1997c] L. Steels. The synthetic modeling of language origins. *Evolution of Communication Journal*, 1(1):1–34, 1997.

[Werner and Dyer, 1991] 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.