The synthetic modeling of language origins

Luc Steels
Artificial Intelligence Laboratory
Vrije Universiteit Brussel
Pleinlaan 2, B-1050 Brussels, Belgium &
Sony Computer Science Laboratory - Paris
6 Rue Amyot, 75005 Paris, France
E-mail: steels@arti.vub.ac.be
In: Evolution of Communication Journal, 1(1).
John Benjamins, Amsterdam

September 5, 1997

Abstract

The paper surveys work on the computational modeling of the origins and evolution of language. The main approaches are clarified and some example experiments from the domains of the evolution of communication, phonetics, lexicon formation, and syntax are discussed.

1 Introduction

The paper surveys research in which software simulations and experiments with robotic agents are used to explore the viewpoint that language is a complex dynamical system. The main goal of the paper is to outline the approaches and show example experiments. Much more work needs to be done to arrive at a full-fledged theory of the origins of language and even about the work already done much more can be said than is possible in a single paper. Nevertheless, I hope to show that a new exciting approach to the study of the origins and evolution of language is taking shape.

The rest of the paper is in four parts. The next section clarifies the notion of a complex system and the multi-agent perspective. Section 3 discusses the relation between complex systems and the origins and evolution of language. Section 4 introduces some examples of research which the complex systems view enables and stimulates. The examples are selected from different areas of language: the evolution of communication itself, phonetics, the lexicon, and syntax. Some conclusions end the paper.

2 Background Notions

2.1 What are complex systems

Since the nineteen sixties, the self-organisational and chaotic behavior of complex systems has been investigated with great intensity. A *complex dynamical system*, or complex system for short, consists of a set of interacting elements where the behavior of the total is an indirect, non-hierarchical consequence of the behavior of the different parts. Complex systems differ in that sense from strictly hierarchical systems (such as most engineered artifacts like cars or computer programs) where the total behavior is a hierarchical composition of the behavior of the parts. In complex systems, global coherence is reached despite purely local non-linear interactions. There is no central control source. Typically the system is open: new elements are entering and leaving and/or energy is constantly supplied keeping the system out of equilibrium.

Almost all complex systems show three main types of behaviors depending on environmental parameters: equilibrium behavior where the system moves to a single steady state, self-organisation in which a so called dissipative structure emerges as long as the environmental conditions allow or enable it, and chaos, in which a form of extreme but nevertheless unpredictable order may appear. A classical example of a complex dynamical system is a fluid (more specifically the Bénard fluid) which is put in a dish and heated from below [**?**]. When the temperature difference is small, the fluid exhibits local fluctuations due to the heat. But these fluctuations die out and the fluid stays basically homogeneous. This is the equilibrium phase. When the temperature difference is larger, the fluid suddenly self-organises into spatially distributed hexagonal cells. This is the self-organisation phase.

When the temperature difference is still larger, the fluid enters a chaotic turbulent regime.

Within the set of complex dynamical systems a subclass can be distinguished, which is known as *complex adaptive systems*. Its main distinguishing feature is that the laws of the system are no longer constant. In a fluid system, the laws of nature determining the behavior of the fluid do not change. Different behavioral regimes are due to different environmental parameters. But in complex adaptive systems, both the behavior of individual elements and the nature of the interactions may change thus giving rise to a higher-order dynamics. Examples of complex adaptive systems are the economy [?], [?], genetic evolution [?], ecological systems, and social systems [?]. The past decades have seen intensive research in practically all natural, biological and social sciences guided by an awareness that complex (adaptive) systems are a key to understand a large part of the complexity found in nature [?].

2.2 Agent-based modeling

In parallel to these trends, a particular style of computational modeling has developed within artificial intelligence and its related disciplines. This style is known as *agent-based modeling* [?]. It is particularly but not exclusively being practiced in a field called *artificial life* [?]. Agent-based models analyse a domain in terms of active entities (agents) whose behavior is specified and programmed at a very fine-grained level. When agent populations are made to interact and evolve, we see a *multi-agent system* in which global properties emerge by local interactions. Current computer technology allows for experiments with hundreds of thousands or even millions of agents. For example, one large-scale experiment is attempting to model the mobility of the inhabitants of Los Angeles by programming individual entities like cars, houses, roads, people, etc. and studying emergent phenomena such as traffic jams and mobility patterns [?]. The agent-approach is being applied to a wide variety of problems in economics, ecology, population dynamics, cellular modeling, etc. and is the natural paradigm for studying these systems from the viewpoint of complex dynamics.

In Artificial Intelligence research, the multi-agent perspective has been mapped onto robotic agents, i.e. artifical agents which have a physical body and on-board sensors, actuators, computational and energy resources, etc. and 'live' in an ecosystem which includes other agents [?]. The goal of this

research is to study what cognitive properties are emergent due to the interaction between an agent and its real world environment and the agents among themselves [?]. This contrasts with the top-down approach of classical Artificial Intelligence research in which knowledge structures are analysed and programmed in by the experimenter.

3 Application to Natural Language

Modern trends in linguistics and cognitive science are focused almost exclusively on a single individual (idealised) speaker or hearer and have largely ignored dynamics, evolution and questions of origin. However, the complex adaptive systems approach and agent-based modeling and simulation is clearly relevant for language, particularly for two areas: understanding the origins of language, and understanding the evolutionary dynamics of whole languages.

3.1 Understanding the Origins of Language

Language is an obvious example of a complex phenomenon which has somehow emerged in the course of biological evolution. The field of complex systems has the origins of complexity in all its manifestations as prime target of investigation, and so it is quite natural to see whether theories, techniques, and methodologies can be transfered to explain the origins and evolution of language. Three approaches have been offered for explaining the origins of complexity in general: genetic evolution, self-organisation combined with adaptation, and genetic assimilation. Each of them can be applied to language. I first explain each approach and then the application.

3.1.1 Genetic evolution

The theory of evolution by natural selection is a well known (partial) explanation for the origin of biological complexity [?]. It entails the following aspects:

1. To get a build up in complexity, information about the structure of an organism must somehow be preserved. This task is accomplished by genes. The genes determine the physiological form and biochemical

functioning of the organism and its developmental process. In higher order organisms there is also the issue how behavior is preserved from one generation to the next. For this, it is proposed that the brain contains organs which embody specific innate behaviors. These organs are genetically encoded the same way other organs are encoded.

2. The preservation of information happens in genetic evolution because genes are copied from the parents to their offspring. There is assumed to be no transmission of acquired characteristics as Lamarcke, who originally introduced evolutionary thinking in biology, thought.

3. Novelty arises because the copying process may run imperfectly, resulting in mutations, and because genes from different parents are crossed over to yield a mixing of genes and hence possibly novel combinations.

4. What information gets preserved depends on the reproductive success of the indidivuals that carry the information. This reproductive success depends to a large extent on the fitness of the individual to cope with the environment in which it finds itself, and thus individuals appear adapted to the environment.

5. A coherent population of organisms of a certain type (species) arises due to the population dynamics of the underlying gene population, which stabilises on the genes appropriate for organisms in a particular environment.

This explanatory structure has been applied with a great deal of success to many aspects of biological organisms, even to complex organs like the eye. It is therefore not surprising that it has been applied to the language faculty by linguists [**?**] and biologists [**?**] alike. Various researchers engaged in computational modeling have followed up on this and attempted to construct computational models for the genetic evolution of language. The main steps in the genetic approach are as follows:

1. It is assumed first of all that there is a language organ inside the brain, which includes a language acquisition device (LAD) together with all the machinery for parsing and producing sentences in the language according to the grammar acquired by the LAD [**?**]. It is of course not claimed that the LAD completely determines a specific language.

Rather, the language organ embodies a set of universal principles and a set of parameters that must be set for a specific language. The acquisition of a specific language is thus explained by a developmental process which is influenced by the local context but driven by a 'bioprogram' [?]. This developmental process is compared to the growing of fingers on a hand as opposed to 'true' learning. Arguments for this position come from the claim that children do not get enough data (poverty of stimulus) and that no learning procedures have so far been found which are adequate to explain natural language acquisition [?]. They also come from the claim that creoles spontaneously arise when second generation immigrant children develop their own language [?]. It follows that language is basically preserved due to genetic transmission, which implies that there must be a set of language genes. Such language genes (or rather their absence or malfunction) have been claimed to be demonstrated in at least one family by Gopnik [?].

2. The preservation of language happens because language genes are copied from the parents to the offspring.

3. Linguistic novelty arises as in all genetically based systems: through mutation and recombination of language genes. Some linguists have claimed that a single catastrophic mutation has been responsible [?] whereas others have attempted to work out a progressive evolutionary pathway [?].

4. The shape of the grammar depends on the reproductive success of the individuals that carry the language genes and is therefore very indirect. This explains Chomsky's pessimism in looking for evolutionary explanations and his rejection of functional explanations for the nature of language. Language is viewed as autonomous and full of language-unique, idiosyncratic features, which could just as well have been otherwise [?]. This view has been echoed by biologists like Gould (see the discussion in [?], p. 384-393.

5. Linguistic coherence is assumed to arise in the same way as other species-common biological features (like having two eyes or two ears). The language genes have either a common origin due to copying, going

back ultimately to the same ancestors, or have spread the way other genes are known to have spread.

Before continuing, it should be pointed out that although these positions are very popular, particularly in the anglosaxon language and cognitive science research communities, they are by no means uncontroversial. Researchers trying to confirm Gopnik's language games have concluded that "the evidence (...) provides no support for the proposed existence of grammar-specific genes." [?], p. 930. The poverty of the stimulus argument has been challenged because empirical data show that there is no such poverty [?]. Learning procedures have been proposed and experimentally tested that are much more powerful than commonly assumed to be possible [?]. Some researchers in language dynamics have clearly demonstrated that the formation of creoles or other external influences have more to do with bilingualism, increased language contact and sociodynamics, than with return to a universal innate grammar [?]. The plausability of the genetic evolution explanation has also been argued to be incompatible with the limits on the speed of evolution and the increases in genetic storage capacity [?]. Most importantly, the exact nature of the LAD and its associated parsing and production apparatus remain to be found, although many hypotheses exist.

3.1.2 Adaptation and self-organisation

Another explanation for the origin of complexity, which has been studied most intensely in physics and chemistry, is self-organisation [?]. It is a process whereby a system of elements with only local interactions but strong positive feedback loops develops a global coherence in order to cope with an in- and outflux of energy or materials. In this case, there is clearly no genetic dimension and information is preserved by the maintenance of the self-organised structures. The phenomenon has been observed in a variety of open physico-chemical systems, including lasers or "Bhelouzow-Zhabotinsky" style chemical reactions. It has also been applied to understand certain non-genetic biological phenomena such as path formation in ant societies [?], pattern formation [?], and pre-RNA precursors to life [?].

The formation of a shared linguistic convention in a group of distributed agents can also be understood in terms of self-organisation [?]. The existing

natural variation in language behavior becomes canalised because the more members of a community adopt the same convention, the stronger the convention will become. This dynamics is only possible however if the agents change their behavior to be more in tune with the behavior of the rest of the community. Self-organisation therefore implies that we view a language community as a complex adaptive system. The change in behavior may be as simple as a change in the weight of neural network connections, as explored in neural network research [**?**], or it may be based on full-scale symbolic learning or analogy.

The application of the self-organisation/adaptive systems approach yields quite different hypotheses for the origins and maintenance of language compared to the genetic approach:

1. Information related to language is preserved in the memory of the individuals as opposed to the genes. There is no "language organ" although a more general substrate of cognitive, perceptual and motor abilities is innate. The language faculty has to make do with the capabilities provided by this general purpose substrate.

2. Language is preserved through behavioral imitation based on learning. It is transmitted in a cultural fashion as opposed to a genetic fashion.

3. Novelty arises due to various factors: The imitation/acquisition process may result in errors, no agent has a complete overview of the language used by all the others leading to overgeneralisation or overspecialisation, agents may want to differentiate themselves from the perceived norm, etc.

4. What information gets preserved, and thus what the shape is of a language, is driven by various selectionist criteria, including attempts to maximise communicative success, minimise cognitive processing and memory load, and be compatible with the limitations of the sensorimotor apparatus. Language can therefore be understood from a functional perspective and is not idiosyncratic, although there are of course 'historical accidents' that have become dominant due to social convention.

5. Language coherence does not arise genetically but is due to self-organisation. More concretely, there is a positive feedback loop between success in

> use and use. A word, sound, concept, or syntactic convention, will be preferentially used if its use has success. Hence the more success, the more use and the more use the more success.

The evolution and growth in complexity of language is thus a side effect of cultural transmission and adaptation. Hurford [?] has called this a *glossogenetic* evolution as opposed to the *phylogenetic* evolution proposed by genetically inclined linguists.

Note that a position on the question of origins and evolution has an impact on the nature of linguistic theory. The genetic approach championed by the Chomsky school rejects functionalism and sees the language faculty as a formal mechanism whose nature is largely arbitrary. The self-organisation/adaptation approach is based on the view that language is a device for communication and representation and therefore sees the specific form of language as the result of balancing physiological and functional constraints. Whereas the genetic approach resonates therefore the strongest with the generative, formalistic tradition in linguistics, the adaptive approach is more in line with cognitive [?] and functional grammar [?]. An adaptationist approach does not mean however that the source of variation is functional - the source of variation in genetic evolution is not functional either.

The adaptive systems approach to language origins and evolution adds two important aspects to research on neural network style or symbolic learning of language. First of all the population dynamics is considered to be crucial, as opposed to the learning by a single individual based on prepared examples. Second there is a constructive aspect, language users not only learn passively the language in their environment but continuously participate in its construction. Otherwise it is not possible to explain the origins of language *de novo*, or how new features ever enter into a language.

3.1.3 Genetic assimilation (the Baldwin effect)

As with many binary oppositions, it is also possible to find a synthesis, in this case between the genetic view and the adaptive systems view. This approach is taken by researchers that accept that adaptive forces are at work in the acquisition and formation of language, but that there is nevertheless a genetically specified language organ, which is the result of (indirect) progressive encoding of linguistic structures in the genetic material. The theoretical

framework explaining how this is possible has first been proposed by Baldwin and demonstrated in computational experiments by Hinton and Nowlan [?].

The main idea is to have a dual dynamics. There is genetic evolution but there is also individual adaptation. The combination gives an individual whose genetic make-up is already close and who can bridge the gap through adaptive learning, an advantage over an individual whose genetic make-up is further removed from the desired gene set or who has no such learning ability. An individual who needs less time for learning or adaptation has an even greater advantage (assuming the environment stays constant), so that the genes of this individual will have a higher chance of proliferating. It can be shown that certain characteristics which are initially acquired through cultural invention, propagation, and learning can thus gradually become compiled into innate structures.

The same scenario could be applied to language, and has been suggested by various researchers [?], [?]. Initially there would be almost no innate structures and the language is acquired/constructed through learning and adaptation. But, if the time of learning plays a role in selection (and if the language is not a moving target), organisms will be favored that already have part of the structure of the language encoded in their genes. The net effect would be an innate language acquisition device as well as adaptive/constructive capacities.

3.2 Language Dynamics

The multi-agent approach lends itself also very well to study a second set of issues in linguistics, namely those related to language dynamics [?]. The study of language dynamics focuses on what happens when two languages get in contact, when there is migration from one community to the territory of another one, when social structures break down or shift, when large-scale epidemic diseases wipe out large parts of a population, when a new culture with many new artefacts sweeps through an existing culture, etc. It is known that in all these cases, various dynamical phenomena can be observed, including language death, the birth of a new language (as in creoles), strong mutual influences, even replacement of large portions of the language by that of another one (for example, a take over of the lexicon while preserving the syntax, or a take over of the syntax while preserving the lexicon).

The multi-agent approach allows us to simulate large collections of agents

which each have their own ideolect and to study the propagation of new features or the changes due to language contact or change in social structure. What is remarkable is that we see behavioral regimes analogous to those observed in other complex systems. For example, under certain in- and outflux rates of agents and possible meanings, a self-organised coherence is seen, just like the Bénard fluid organises into hexagonal cells. But when these rates get out of balance, for example because new agents come in too quickly or existing agents leave too rapidly or too many new meanings need to be expressed, there is a disintegration of the language similar to a turbulent chaotic regime (see figure 1 from [**?**]). In between these two regimes, we see an increased variation, causing rapid (linguistic) evolution "on the edge of chaos". Similar phenomena of rapid evolution have also been observed in physico-chemical and biological systems which are far from equilibrium [**?**].

The study of language dynamics using concepts from complex systems is in itself extremely exciting but will not be further explored in the present paper.

4 Experiments

4.1 General Structure of synthetic experiments

It is possible to apply the multi-agent simulation methodology to all levels of language and explore each of the three approaches discussed above. The general structure of such synthetic experiments is however always the same. There is a population of 'entities' capable to engage in linguistic or cognitive behavior. For the rest of the paper I will call these entities *agents*. The population is typically open (new agents enter and others leave). To show that language is maintained despite this in- and outflux is already one of the main challenges of every experiment. In experiments involving a semantic component, there is also a population of situations in the world which can be the subject of a linguistic communication. To study issues of meaning creation and grounding, the agents are sometimes encapsulated in physical robots and the situations are then physical situations encountered by the robots operating in their environment. The set of possible situations is then open as well, as is the case for humans or animals.

The agents engage in two types of interactions. There are *linguistic inter-*

Figure 1: Each graph plots the evolution of a population and the running average of communicative success for 500,000 linguistic interactions. The population fluctuates between 50 and 90 agents. The agents develop a lexicon to refer to themselves so changes in the population also imply changes in the set of possible meanings. In the left graph, the population is able to cope with the change. In the right graph, the population changes much more quickly and communicative success drops occasionally back to zero before rebounding again.

actions, for example one agent produces a sound and another one perceives the sound, or one agent formulates a sentence and the other one tries to parse it. In the remainder of the paper, I will call such linguistic interactions *games*, partly in order to make the connection with evolutionary game theory [**?**]. In the adaptive approach, the linguistic rules of the agents may change based on the outcome of a game. In a purely genetic approach, the rules do not change as the result of a linguistic interaction. Instead, there are *genetic interactions* which involve the creation of offspring from one or two parents. Innate linguistic structures may change at this point due to mutation and recombination. In a genetic assimilation approach, a combination of adaptive changes and genetic changes is seen.

A simulation typically starts with a population of agents which engage in linguistic and/or genetic interactions. As a side effect of these interactions, new linguistic structures may be formed or structures may propagate. The exact nature of the interaction and the nature of the side effect is different for each experiment. There are however some important criteria which need to be respected to yield realistic models, including realistic models for natural language dynamics:

1. *The distributed systems constraint*: Multi-agent systems must be conceived as distributed systems. There should be no agents which have a complete view of the behavior of all the other agents and neither should there be agents which can completely control the behavior of all the others. Indeed, no single language user has a complete overview of his language nor can anyone determine by *fiat* what linguistic convention will have to be adopted by the whole population.

2. *The limited rationality constraint*: Agents should not have direct access to each other's internal states. They can only exchange information or cause change in others by interaction. This is clearly true for human language use. A language user cannot inspect or change the internal brain states of another language user.

3. *The open systems constraint*: There should be an influx and outflux of agents in the population and an influx and outflux of meanings. This is not only because of realism with respect to human populations and cultures. It is known in dynamical systems theory that this in-

and outflux is necessary to explain partly what kind of structures will emerge and why there is evolution in the structure.

The remainder of this section now gives examples of experiments for different aspects of language.

4.2 Evolution of communication

There has been a substantial amount of work addressing the question how communication itself may arise in a community of distributed agents. This work is more closely related to biology (more specifically ethology and ecology) than to linguistics because the communication systems studied have the characteristics of animal signalling systems rather than full-blown human languages: Only a very small fixed set of meanings can be expressed and there is no recursive combination of building blocks. Nevertheless these experiments show clearly how communication may evolve using the principles introduced earlier.

Some of the first experiments have been reported by MacLennan [**?**]. They have the following structure:

1. *Agent population*: There is a changing population of agents located in an environment. The agents can perceive aspects of their environment and potentially emit signals to communicate perceived states to others. An agent which correctly interprets a communication of another one has an increased chance of survival. New agents are obtained by the breeding of two existing agents. The probability of death or breeding depends on the fitness of the agents involved.

2. *Innate structure*: The agents have signal-situation tables which they use to produce a signal and to decode a signal. The agents are born with a signal-situation table.

3. *Linguistic interaction*: One agent may produce a signal based on perceiving a situation and retrieving the corresponding signal from his table. Another agent perceives the signal, and by looking it up in his table, predict the situation which he cannot see otherwise. The interaction has been successful for the second agent if the predicted situation indeed coincides with the real situation.

4. *Adaptation*: MacLennan has done two types of experiments: experiments where changes in the signalling system are purely genetic, and other experiments where they are both genetic and cultural. In the latter case, an agent changes his situation-table when there has been a prediction failure, in such a way that this failure will no longer occur in the future.

5. *Source of novelty*: Novelty is due to the mutation and crossover operations acting on the genetically transmitted signal-situation tables.

6. *Genetic interaction*: A new agent is born with a crossover of the signal-situation tables of his two parents. Generations of agents partly overlap to allow cultural preservation as well.

The computer simulations show that the average fitness in the population (of size 100) increases 50 times due to communication based on genetic evolution alone. In other words, when communication gives an advantage, the genotype of the individuals will evolve through random mutation and combination so that it codes for communication. MacLennan also found that if a simple learning algorithm was added, fitness increases 150 times. This shows that learning is a more effective means to propagate linguistic conventions.

The MacLennan experiment makes a number of assumptions which other researchers have further investigated and often been able to weaken:

1. The experiments assumed that a successful communication has direct benefit, by design. Other researchers have weakened this assumption. For example, Werner and Dyer discuss an experiment where successful communication helps to find a mate [**?**]. Cangelosi and Parisi coupled successful communication to the warning for poisonous food resources [**?**]. More generally, it appears that the fitness measure can be anything that has an impact on future survival of the organism and its reproductive success.

2. The experiments assumed that communication aids in cooperation, which raises the question why agents would ever cooperate. Also this assumption has been weakened. First of all, there has been a lot of other work in the artificial life and economical literature on how cooperation itself may arise, even if the agents are not altruistic [**?**]. There

have also been experiments showing that communication may arise in situations where the agents have adversary relations [?], and that communication does not necessarily have to benefit only the information sender [?]. More generally, any form of interaction between agents is enriched when there is some form of communication.

3. All these experiments assume the 'Saussurean' convention that the same signal is used both for reception and emission of a message. Hurford [?] has been able to weaken this assumption and show that this convention itself will emerge in genetic evolution. Similar results have been obtained by Oliphant who has identified constraints under which the convention may or may not emerge [?].

In conclusion, we can say that it has been shown quite clearly in a variety of experiments that (limited) communication systems arise in groups of agents whenever communication is in some way important for the survival of these agents. Communication systems arise both in genetic and adaptive systems. The next question is how these simple signalling systems could ever have evolved into full-fledged languages. The first puzzle is how human communities have been able to develop shared but open-ended sound systems that through modular combination can form the basis of open-ended lexicons.

4.3 Phonetics and Phonology

Languages show an enormous variety of possible sounds [?], there is also variety between individual speakers of the same language and speaking contexts [?], and children or adults tend to learn the sound pattern available in their environment. Nobody has therefore ever proposed that the sounds of a specific language are innate, even though there are universal properties of human phonological systems. On the other hand, it has been proposed that humans are equipped with an innate perceptual and articulatory apparatus for recognising and producing distinctive features [?], [?]. Such distinctive features could therefore form a natural part of the universal language acquisition device, which is a part of the innate language organ. Attempts to synthesise sound systems using a purely genetic approach have however not been very successful (at least none have been reported in the literature to my knowledge). The main problem is that if a new sound (or a new distinctive

feature) originates in a single individual by genetic mutation, it does not give this individual any advantage. It is only when a sufficient number of individuals undergo the same mutation, which is exceedingly unlikely, that the shared sound is beneficial. The evolution of language differs in this sense drastically from the evolution of other biological features. In the case of language the environment is constantly changing and partially determined by the behavior of the agents.

On the other hand, many phoneticians who have reflected on the origins and evolution of language, such as Lindblom, MacNeilage and Studdert-Kennedy [?], have embraced an adaptive systems approach, in which a mapping between gestures and formant shapes evolves through cultural transmission and adaptation. Distinctive features then become emergent properties seen in retrospect by a descriptive linguist. The adaptation is driven by two types of selectionist criteria: perceptual constraints such as the limitations of the human ear, maximisation of distinctiveness and symmetrical balance, and articulatory constraints such as expressability, repeatability and energy minimisation. Several researchers have shown theoretically that these criteria are sufficient to constrain the kind of sound systems that occur in human languages [?], [?], [?], [?], et.al.

However these demonstrations do not yet show whether sound systems can be originated and acquired by local interaction between agents. This is where agent-based modeling and simulation comes in. A simulation experiment by de Boer and Steels [?] has successfully demonstrated the self-organisation of a sound system through adaptive imitation games. The experiment has the following structure:

1. *Agent population*: There is a population of agents and an in- and outflux of agents which is independent of their linguistic performance.

2. *Innate structure*: There is no innate phonetic knowledge. The agents have a synthetic articulator modeled after the human vocal tract and a perceptual apparatus that decomposes real-time signals into formants.

3. *Linguistic interaction*: One agent (the initiator) produces a sound or sound sequence from its repertoire, which is initially empty. The other agent (the replicator) attempts to imitate the sound, which implies that he is able to recognise the sounds produced by the initiator and

instantiate a gestural score that corresponds to the sounds. The initiator in turns interprets the sound produced by the imitator and gives a positive feedback when the imitation is deemed to be close enough.

4. *Adaptation*: Failure in imitation causes the imitator either to hypothesise a new sound, when the unrecognised sound is too distant from existing sounds, or to shift an existing sound slightly to get closer both in terms of perception and articulation to the sound just perceived.

5. *Source of novelty*: A new sound is created by a random combination of articulatory controls. There is outside pressure on the system to expand the sound repertoire, which could come from a lexicon formation process for example.

6. *Genetic interaction*: None.

An example of sound system evolution derived by playing a series of imitation games among a fluctuating group of about 20 agents is given in figure 2. A stable system is reached. Evolution continues if there is pressure on the system (e.g. from the lexicon) to create new vowels. A flux of the population does not destabilise the vowels as long as the in- and outflow is bounded. Such sound systems clearly show characteristics similar to human sound systems. Most importantly, they may expand or adapt when confronted with a population that has already evolved another sound system.

This experiment shows beyond doubt that (1) a shared sound system may emerge through self-organisation and adaptive imitation games, (2) no language-specific articulatory control or perceptual apparatus needs to be innate, and (3) the selectionist criteria proposed by phoneticians are indeed enough to derive natural sound systems in a bottom-up emergent way.

A similar experiment for deriving vowel systems through agent-based models has been reported by Berrah et.al. [**?**]. They have explored a genetic assimilation approach: Agents are selected based on how much effort they need to do to learn a phoneme. Berrah, et.al. also report convergence, however, agents are allowed to 'see' each others' internal brain states thus violating the limited rationality constraint.

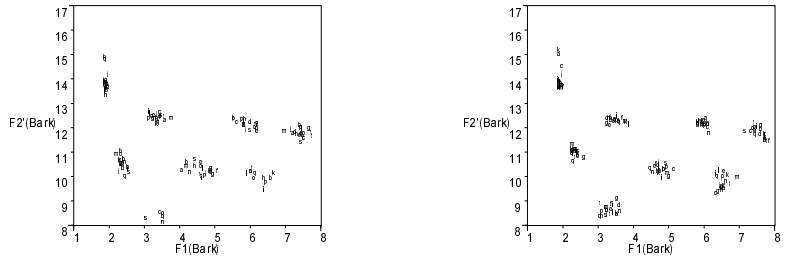

Figure 2: The figure shows the position of vowels in perceptual space for different agents (labeled a,b,c, etc.). The x- and y-axis correspond to the first and second formant respectively. The left system of 7 vowels originated from scratch. A new vowel is already emerging ath the left bottom. The right vowel system evolved from the left one after 1000 additional imitation games. The new vowel is confirmed and other vowels have moved slightly to maintain distinctive distance. All this happened despite of a steady in- and outflux of agents in the population.

4.4 The Lexicon.

Several simulation experiments have been reported in the connectionist literature showing how a neural network may acquire both a set of concepts, internally encoded as activations of hidden nodes, and a set of associated words for these concepts, by repeated exposure to pairs of situations and descriptions. A typical example has been reported by Regier [**?**] for the acquisition of spatial relations using structured connectionist networks and the back-propagation learning algorithm. A first experiment on real robots has been performed by Yanco and Stein [**?**]. Their robots use reinforcement learning to acquire the mapping between internal states/actions and signals. These experiments use an adaptive approach but do not yet show how a lexicon may originate, because the lexicon is provided externally by the experimenter. They also do now show how lexicons may expand based on the needs to lexicalise new meanings, because the set of meanings and agents is fixed.

One of the first experiments weakening these assumptions has been published by Hutchins and Hazlehurst [**?**]. It is prototypical for the adaptive/constructive approach to lexicon formation and meaning creation. Its structure is as follows:

1. *Agent population*: There is a population of agents and an in- and outflux of agents which is independent of their linguistic performance.

2. *Situation population*: There is a population of situations that the agents may encounter. In the Hutchins-Hazlehurst experiment, this situation population is however fixed.

3. *Innate structure*: There is no innate semantic nor lexical knowledge. The agents come equiped with a domain-independent connectionist network that analyses a situation into activation of nodes on a hidden layer. The hidden layer acts also as the verbal input/output layer (figure 3). The weights of these networks are initialised randomly.

4. *Linguistic interaction*: The agents engage in language games, taking turns as speakers and listeners. A speaker perceives a situation and emits a pattern of activation on the verbal output layer. The hearer perceives the same situation and uses the output produced by the listener as an additional input.

5. *Adaptation*: There are two types of adaptation. The first one is the auto-associator network which is adjusted by the speaker and the listener using the back-propagation method, until the visual output layer corresponds to the visual input layer. The second one is another associative network which is adjusted by the listener so that his hidden layer matches hidden layer of the speaker.

6. *Genetic interaction*: None.

Hutchins and Hazlehurst report that after a set of interactions, the agents generally converge on a shared lexicon, meaning that the patterns produced on the verbal input/output layer for the same situation are compatible among the agents. When an additional network is added mapping the hidden layer on an output layer, the internal representation (the distinctions encoded in the hidden layer) are no longer identical. There is communicative success despite differences in internal representations.

Steels [**?**], [**?**] has reported another series of experiments in adaptive lexicon formation which are similar to the Hutchins-Hazlehurst experiment: Agents have no innate semantic or lexical knowledge, but develop a joined lexicon and meaning repertoire through adaptive language games. Both the set of agents and the set of situations is open. In these experiments, the source of new meanings is not an auto-associative network but a selectionist mechanism that builds up features through discrimination games. The lexicon is not stored in the form of a connectionist network but as a set of associations between words, meanings, their use and success, so that ambiguity, homonymy, and synonymy can occur. A word-meaning pair that is experienced by the agent to be the most successful, is prefered. This induces a positive feedback loop (the more use the more success, and the more success the more use) giving rise to self-organised coherence.

The lexicon formation and meaning creation mechanisms proposed by Steels have been ported to physical robots [**?**], so that the situations are generated through the real-world interaction of the robot with the environment and the inputs for a language game are directly derived from sensory channels. The robots are autonomous 'Lego-vehicles' now widely used in experiments in "bottom-up" AI research (see figure 4). The robots engage in language games. When a robot decides to become the speaker, he seeks out another one, points (by orientation) to a topic of the conversation, which is one of the objects in the immediate environment, then conceptualises a

Figure 3: Connectionist network used in the Hutchins-Hazlehurst experiment. There is an auto-associatior network which attempts to reconstruct on the output layer the activation received on the input layer through the intermediary of a hidden layer. This hidden layer is associated with verbal input/output through an additional layer.

distinctive feature set distinguishing the topic from the other objects in the context, and verbalises this distinction. The hearer also conceptualises the same reality, parses the expression of the speaker, and matches his interpretation with his expectation. Positive feedback is given by the hearer when expectation and interpretation match. When any step in this game fails, adaptation occurs: new distinctions are created, new words are created, existing word-meaning pairs are revised, etc. The experiments show that a group of robots manages to develop a common set of distinctions and a common lexicon for verbalising them, all this without human intervention, and without innate linguistic or semantic knowledge nor genetic transmission.

4.5 Syntax

The domain of syntax (including morphology and morphosyntax) is the most complex area of language, particularly if we consider the production of syntactic structures driven by meaning and the parsing of syntactic structures into meaningful internal representations. It is therefore not surprising that there are fewer global experiments reported so far. Nevertheless, some early experiments focusing on partial aspects of syntax (e.g. word order) can be seen. These experiments have the same structure as the experiments discussed so far. By way of example, we discuss some experiments by Briscoe which explore the formation and propagation of grammatical constraints [**?**].

Briscoe explores grammars in the categorial grammar formalism. Possible categories and their combination constraints are encoded as feature structures which are organised in a lattice so that generalisations are captured through an inheritance hierarchy. Briscoe assumes the Chomskian Principles and Parameters framework. Agents are born with an extensive universal grammar (which specifies for example all the possible categories and their possible combinations), but there are parameters which have to be set based on examples from the language found in the environment. These parameters focus specifically on word order (SOV, VOS, etc.). The goal of the experiments is to show (1) that the values of parameters can be acquired through learning, and (2) that there is a progressive genetic assimilation, i.e. parameters become set and thus part of universal grammar.

The structure of the Briscoe experiments is as follows:

1. *Agent population*: There is a population of agents with an in- and

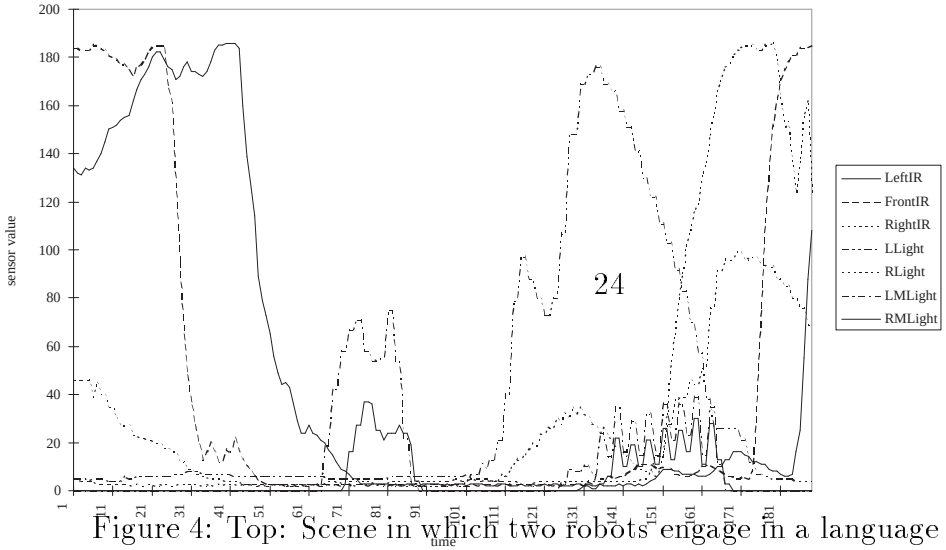

Figure 4: Top: Scene in which two robots engage in a language game that allows them to form a common repertoire of distinctions and a common lexicon for verbalising these distinctions. The bottom figure shows the sensory datastreams that are input to the conceptualisation module during the physical

outflux.

2. *Innate structure*: The agents are born with all the machinery for generating and analysing sequences of categories (sentences) based on a categorial grammar. The grammar itself is partly innate, in particular, it contains the kinds of categories that may occur, and some or all of the ordering in which categories can be combined.

3. *Linguistic interaction*: One agent generates a sequence of categories and another agent attempts to parse it based on his grammar. The interaction has been successful for the second agent if a parse tree could be constructed. The agents have a fitness which depends on their success in linguistic interactions, as well as the load on working memory and the expressability of their language.

4. *Adaptation*: An agent may set a parameter in his grammar when he was unable to parse a sentence. The new setting is based on a simple learning algorithm.

5. *Source of novelty*: There is no true source of novelty. The experiments are either started up with (adult) populations already containing different language distributions, or with randomly initialised populations. The experiments show that one type of language will dominate the population.

6. *Genetic interaction*: The most fit agents are allowed to generate offspring. The new agent obtains grammars of both parents after a crossover and possible mutation.

The experiments of Briscoe show that genetic assimilation indeed takes place. They also show that the genetically fixed parameters evolve towards the dominant language in the environment, say an SVO language. In other words, if the population starts out with a particular bias, this bias will become coded in their inherited parameter sets. Which language is dominant will depend, among other things, on the parsing complexity (working memory load) which plays a role in the fitness determination.

Several other interesting efforts in syntax can be reported. The role of functional constraints has been the subject of various genetic and/or cultural experiments in the origins and evolution of syntax. For example, Hawkins [**?**]

has developed a measure of complexity for parsing, which reflects how many units need to be retained in memory before decisions can be made about which constitutent structure underlies a particular sentence. He predicts on this basis the occurrence of word order universals. This measure has been shown in agent-based simulations by Kirby [?] to converge indeed to distributions that reflect universals found in language, even though not all features of language can be shown to have a functional explanation.

In another fascinating series of experiments Ikegami and Hashimoto [?] have investigated how a grammar formalism itself may become more complex, i.e. how formalisms of the power of regular phrase structure grammars may develop into context-free, context-sensitive and eventually unrestricted rewrite grammars. The experiments are also based on games in which one agent generates a sentence and another one attempt to parse it. Agents adjust their grammar by mutations or recombinations of grammar rules.

All the above experiments treat grammar as a formal device and do not embed the evolution of grammar within communicative acts (as the language games used in studying the origins of lexicons that were reported earlier). However, some first examples of semantically driven syntax origins using adaptation/self-organisation have been reported. One experiment has been conducted by Batali [?]. It involves agents that encode structured meaning into strings of characters and decode strings back to meaning, using recurrent neural networks. After an interaction, agents compare their meanings directly (a step which obviously needs to be revised to approach more realistic multi-agent conditions). If the speaker's intended meaning does not match with the hearer's, the hearer adjusts his grammar. After a sufficient number of interactions and adjustments, word order is seen to be systematically exploited for expressing meaning.

Another experiment in meaning-driven development of syntax has been reported by Steels [?]. The experimental setup which acts as the source of meaning, involves two robot "heads" which can track moving objects with a camera (figure 5). Moving objects are detected based on difference matching of two consecutive images and each head attempts to maintain the moving object in the center of vision by rotation. The heads watch a scene in which a robot is moving around. Occasionally, other objects come into view which are visually detected by distinctions against the background. Different properties of each image-object, such as the average greylevel, the position or changes in position of the angle of the head, the distance (correlated with the size of

the bounding box surrounding the object), etc. act as the sensory inputs to develop a repertoire of distinctions and then a lexicon for expressing these distinctions using methods discussed earlier [?]. The additional component responsible for syntax is a general structure-detecting engine that analyses objects (and hence possibly words) into hierarchical structures based on their properties and mutual constraints. The same engine is also capable to synthesise conglomerations of objects so that the constraints are again satisfied. This capability is not only relevant for language production but also for planning, for drawing objects, etc. Use and success of structures is monitored and the most successful structures are preferentially used, causing gradual self-organised coherence. Selectionist pressure comes from attempting to express as accurately as possible the recent (dynamic) scenes using a minimal amount of cognitive load both for parsing and producing. The experiments show a growing complexity in the syntax as needed to express more and more complex properties of the perceived objects and relations between them.

These various experiments in the origins of syntactic structure tackle pieces of a very complex puzzle that will all have to be put together to reach a satisfactory explanation for the multi-dimensional richness of natural language syntax. However it is quite clear that also in this domain, the paradigm of complex systems and multi-agent simulations is opening up a new tremendously exciting dimension in the study of language.

5 Conclusions

The origins and evolution of language is one of the major unresolved problems of science, despite a long history of research on the subject [?],[?]. There is at the moment a growing research effort to model and synthesise processes underlying the origins and evolution of language. This paper has attempted to circumscribe and illustrate these efforts. Three main approaches are seen: genetic evolution, cultural adaptation, and genetic assimilation. Various experiments have already been reported that show inklings of how various aspects of language may have come into existence and how it continues to evolve and adapt. There is no doubt that the coming decade will see increasingly more sophisticated models, scaled up computer simulations, and intriguing robotic experiments, all leading to a solid scientific theory on the origins and evolution of language.

Figure 5: The top figure shows the 'artificial eye' of a robotic head capable to track a moving object. The bottom figure shows the (smoothed) dynamically varying properties of perceived object-images which act as input to the language games played by two heads watching and commenting on dynamic

6 Acknowledgement

The writing of this paper was sponsored by the Sony Computer Science Laboratory in Paris, and stimulated by Sherman Wilcox. Bart de Boer (VUB AI Lab) , Angus McIntyre (Sony CSL), Paul Vogt (VUB AI lab) and Tony Belpaeme (VUB AI Lab) have collaborated on the language experiments in Brussels and Paris and have helped to sharpen my insights through a continuous discussion of the issues raised in the paper. I also thank the many members of the 'computational' origins of language community for sending me drafts of their papers and providing further references. Further comments to improve and complete this paper are greatly appreciated.

References

[1] Arthur, B., et.al. (ed.) (1996) The Economy as a Complex Adaptive System. Santa Fe Institute Series on Complexity. Addison-Wesley, Menlo Park Ca.

[2] Batali, J. (1995) Small Signaling Systems can Evolve in the Absence of Benefit to the Information Sender. Submitted for publication.

[3] Batali, J. (1997) Computational Simulations of the Emergence of Grammar. To appear in Hurford, J., et.al. (1997).

[4] Berrah, A. H. Glotin, R. Laboissiere, P. Besiere and J-L Boe (1996) From Form to Formation of Phonetic Structures: An evolutionary computing perspective. In: Fogarty, T. and G. Venturini (eds.) ICML '96 workshop on Evolutionary Computing and Machine Learning. Bari. pp. 23-29.

[5] Bickerton, D. (1990) Language and Species. University of Chicago Press, Chicago.

[6] Boe, L-J., J-L. Schwartz and N. Valle (1995) The Prediction of Vowel Systems: Perceptual Contrast and Stability. In: Keller, E. (ed.) (1995) Fundamentals of Speech Synthesis and Speech Recognition. John Wiley, London. pp. 185-213.

[7] Carre, R. and M. Mrayati (1993) Vowel transitions, vowel systems and the distinctive region model. In: Soring, C. et.al. (eds) (1993) Levels in Speech Communication: Relations and Interactions. Elsevier, Amsterdam. pp. 73-89.

[8] de Bourcier, P. and M. Wheeler (1994) Signalling and Territorial Aggression: an investigation by means of synthetic behavioural ecology. In: Cliff, D. et.al. (1994) Proceedings of the third Simulation of Adaptive Behavior Conference. The MIT Press, Cambridge Ma. p. 462-472.

[9] Briscoe, T. (1997) Language Acquisition: The Bioprogram Hypothesis and the Baldwin Effect. To appear in Hurford, J., et.al. (1997).

[10] Cangelosi, A. and D. Parisi (1996) The Emergence of a "Langage" in an Evolving Population of Neural Networks. Submitted for publication.

[11] Chomsky, N. (1975) Reflections on Language. Pantheon, New York.

[12] Chomsky, N. (1976) On the nature of language. In: Harnad, S., H. Steklis and J. Lancaster (eds.) Origins and Evolution of Language and Speech. Annals of the New York Academy of Science, New York. Vol 280 pp. 46-57.

[13] Chomsky, N. and M. Halle (1968) The Sound Pattern of English. The MIT Press, Cambridge Ma.

[14] Ferber, J. (1995) Les systemes multi-agents: vers une intelligence collective. InterEditions, Paris.

[15] De Boer, B. (1997) Emergent Vowel Systems in a Population of Agents. Submitted for publication to ECAL-97.

[16] Daelmans, W, G. Durieux and S. Gillis (1994) The Acquisition of Stress: A Data-Oriented Approach. Computational Linguistics. (1994) p. 421-451.

[17] Demolin, D., M. George and A. Soquet. (1996) The role of self-organisation in the emergence of phonological systems. Submitted for Publication.

[18] Dennett, D. (1995) Darwin's Dangerous Idea. Evolution and the Meaning of Life. Simon and Schuster, New York.

[19] Pasteels, J.M. and J-L Deneubourg (1987) From Individual to Collective Behaviour in Social Insects. Birkhauser, Basel.

[20] Dik, S. (1980) Studies in Functional Grammar. Academic Press, London.

[21] Epstein, J.M. and R. Axtell (1996) Growing Artificial Societies. Social Science from the Bottom up. The MIT Press, Cambridge Ma.

[22] Gopnik, M. (1994) Impairments of tense in a familial language disorder. Journal of Neurolinguistics, Vol 8, 109-133.

[23] Hashimoto, T. and T. Ikegami (1996) Emergence of net-grammar in communicating agents. BioSystems 38 (1996) 1-14.

[24] Hawkins, J. (1992) Innateness and Function in Language Universals. In: Hawkins and Gell-Man (1992) p. 87-120.

[25] Hawkins, J. and M. Gell-Mann (eds.) (1992) The Evolution of Human Languages. SFI Studies in the Sciences of complexity, Proc. Vol. X. Addison-Wesley Pub., Menlo Park Ca.

[26] Hinton, G. and S. Nowlan (1987) How learning can guide evolution. Complex Systems, Vol 1, 495-502.

[27] Holland, J. (1993) Hidden Order: How Adaptation Builds Complexity. Addison-Wesley Pub., Menlo Park Ca.

[28] Hurford, J. (1989) Biological evolution of the Saussurean sign as a componnet of the language acquisition device. Lingua, 77. pp. 187-222.

[29] Hurford, J. (1990) Nativist and functional explanations in language acquisition. In: Roca, I. (ed) (1990) Logical Issues in Language Acquisition. London, Foris Publications. pp. 85-136.

[30] Hurford, J., C. Knight and M. Studdert-Kennedy (eds.) (1997) Evolution of Human Language. Edinburgh Univ. Press. Edinburgh.

[31] Hutchins, E. and B. Hazelhurst (1995) How to invent a lexicon. The development of shared symbols in interaction. In: Gilbert, N. and R. Conte (eds.) Artificial societies: The computer simulation of social life. UCL Press, London.

[32] Kaneko, K. (1995) Chaos as a Source of Complexity and Diversity in Evolution. In: Langton, C. (ed) (1995) Artificial Life. An Overview. The MIT Press, Cambridge Ma. p. 163-177.

[33] Kaufman, S. (1995) The Origins of Order: Self-Organization and Selection in Evolution. Oxford University Press, Oxford.

[34] Kirby, S. (1996) Function, Selection, and Innateness: The Emergence of Language Universals. Ph.D. Thesis. University of Edinburgh.

[35] Kirby, S. (1997) Fitness and the Selective Adaptation of Language. In: Hurford, J., et.al. (ed.).

[36] Labov, W. (1994) Principles of Linguistic Change. Volume 1: Internal Factors. Blackwell, Oxford.

[37] Ladefoged, P. and I. Maddieson (1996) The Sounds of the World's Languages. Blackwell, Oxford.

[38] Langacker, R.W. (1986) Foundations of Cognitive Grammar. Stanford University Press, Stanford.

[39] Langton, C. (ed.) (1995) Artificial Life. An overview. The MIT Press, Cambridge Ma.

[40] Liljencrants, L. and B. Lindblom (1972) Numerical Simulations of Vowel Quality systems: The role of perceptual contrast. Language, 48, pp. 839-862.

[41] Lindgren, K. (1991) Evolutionary phenomena in simple dynamics. In: Langton, C., et.al. (1991) Artificial Life II. Addison-Wesley, Redwood Ca. pp. 295-312.

[42] MacLennan, B. (1991) Synthetic Ethology: An approach to the study of communication. In: Langton, C., et.al. (1991) Artificial Life II. Addison-Wesley Pub. Cy, Redwood City Ca. pp. 603-631.

[43] Meinhardt, (1982) Models of Biological Pattern Formation. Academic Press, London.

[44] Nagel, K, S. Rasmussen, C. Barret (1996) Network Traffic as a Self-organised Critical Phenomenon. Working Paper 96-8-65. Santa Fe Institute New Mexico.

[45] Nichols, J. (1992) Linguistic Diversity in Space and Time. The University of Chicago Press, Chicago and London.

[46] Lieberman, P. (1991) Uniquely Human: The Evolution of Speech, Thought and Selfless Behavior. Harvard University Press, Cambridge Ma.

[47] Lindblom, B., P. MacNeilage, and M. Studdert-Kennedy (1984) Self-organizing processes and the explanation of phonological universals. In: Butterworth, G., B. Comrie and O. Dahl (eds.) (1984) Explanations for Language Universals. Walter de Gruyter, Berlin. pp. 181-203.

[48] Maynard-Smith, J. (1983) Evolution and the Theory of Games. Cambridge University Press, Cambridge.

[49] Maynard-Smith, J. and E. Szathmary (1994) The major transitions in evolution. Freeman Spektrum, Oxford.

[50] McClelland, J.L., and Rumelhart, D.E. eds. 1986. Explorations in Parallel Distributed Processing. The MIT Press/Bradford Books, Cambridge Ma.

[51] Arends. J, P. Muysken and N. Smith (1995) Pidgins and creoles: an introduction. Benjamins, Amsterdam.

[52] Nicolis, G. and I. Prigogine (1993) Exploring Complexity. Piper, Berlin.

[53] Oliphant, M. (1996) The dilemma of Saussurean communication. Biosystems, 37 (1-2), pp. 31-38.

[54] Prigogine, I. and I. Stengers (1984) Order Out of Chaos. Bantam Books, New York.

[55] Pinker, S. (1994) The language instinct. Penguin Books, London.

[56] Pinker, S. and P. Bloom (1992) Natural language and natural selection. In: Barkow, J, L. Cosmides and J. Tooby (eds.) (1992) The Adapted Mind: Evolutionary Psychology and the Generation of Culture. Oxford University Press, Oxford. pp. 451-493.

[57] Pullum, G. (1996) Learnability, Hyperlearning, and the Poverty of the Stimulus. manuscript.

[58] Regier, T. (1995) A Model of the Human Capacity for Categorizing Spatial Relations. Cognitive Linguistics, 6-1, pp. 63-88.

[59] Romaine, S. (1988) Pidgin and Creole Languages. Longman, London.

[60] Steels, L. (ed.) (1995) The Biology and Technology of Intelligent Autonomous Agents. Nato ASI Series F, vol 144. Springer-Verlag, Berlin.

[61] Steels, L. (1996a) Self-organising vocabularies. In: Langton, C. (ed.) Proceedings of Artificial Life V. Nara.

[62] Steels, L. (1996b) Emergent Adaptive Lexicons. In P. Maes, Mataric, M., Meyer, J-A., Pollack, J., and Wilson, S. (eds.) (1996) From Animals to Animats 4: Proceedings of the Fourth International Conference on Simulation of Adaptive Behavior. Cambridge, Ma: The MIT Press/Bradford Books.

[63] Steels, L. (1996c) Synthesising the origins of language and meaning using co-evolution, self-organisation and level-formation. In: Hurford, J. et.al. (ed.) (1997) Evolution of Human Language. Edinburgh University Press, Edinburgh.

[64] Steels, L. (1997a) 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.

[65] Steels, L. (1997b) Language Learning and Language Contact. In: Daelemans, W. (ed.) Empirical approaches to Language Learning. ECML Prague.

[66] Steels, L. (1997c) The origins of syntax in visually grounded robotic agents. In: Pollack, M. (ed.) (1997) Proceedings of the 15th International

Joint Conference on Artificial Intelligence. The AAAI Press, Menlo Park, Ca.

[67] Steels, L. and R. Brooks (1995) The Artificial Life Route to Artificial Intelligence. Building Embodied Situated Agents. Lawrence Erlbaum, New Haven.

[68] Steels, L. and P. Vogt (1997) Grounding adaptive language games in robotic agents. Submitted to ECAL-97.

[69] Steels, L. and A. McIntyre (1997) Spatially Distributed Naming Games. Submitted to ECAL-97

[70] Thomason, S.G., T. Kaufman (1988) Language Contact, Creolization, and Genetic Linguistics. University of California Press, Berkeley.

[71] Thomas, D. A. (1995) Music and the origins of language. Theories from the French Enlightment. Cambridge University Press, Cambridge.

[72] Trabant, J. (1996) Thunder, Girls, and Sheep, and Other Origins of Language. In: Trabant, J. (ed.) (1996) Origins of Language. Collegium Budapest Workshop Series No. 2. pp. 39-69.

[73] Vargha-Khadem, F., K. Watkins, K. Alcock, P. Fletcher and R. Passingham (1994) Praxic and nonverbal cognitive deficits in a large familility with a genetically transmitted speech and language disorder. Proc. Natl. Acad. Sci. USA Vol 92, pp 930-933.

[74] Werner, G. and M. Dyer (1991) Evolution of Communication in Artificial Organisms. In: Langton, C., et.al. (ed.) Artificial Life II. Addison-Wesley Pub. Co. Redwood City, Ca. p. 659-687.

[75] Wittgenstein, L. (1974) Philosophical Investigations. Translated by G. Anscombe. Basil Blackwell, Oxford.

[76] Worden, R. (1995) A Speed Limit for Evolution. Journal of Theoretical Biology 176, pp. 137-152.

[77] 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