

The Evolution of Communication Systems by Adaptive Agents

Luc Steels

VUB AI Lab – Brussels
Sony Computer Science Laboratory – Paris
steels@arti.vub.ac.be

Abstract. The paper surveys some of the mechanisms that have been demonstrated to be relevant for evolving communication systems in software simulations or robotic experiments. In each case, precursors or parallels with work in the study of artificial life and adaptive behaviour are discussed.

1 Introduction

Almost since the beginning of research in Artificial Life and the Simulation of Adaptive Behaviour, there have been efforts to apply biological principles and the methodology of building artificial systems to understand the origins and evolution of communication systems with the complexity of natural languages, not only by abstract software simulations but also by experiments on situated embodied robotic agents operating in real world environments. In almost all these experiments, language is viewed as a complex adaptive system which emerges in a bottom-up fashion from local one-on-one interactions between situated embodied agents, and evolves and complexifies based on principles like cultural selection, structural coupling, and self-organisation. Rather than looking only at natural languages as they exist today, research in ‘artificial language evolution’ tries to evolve artificial languages with natural language-like properties – and thus explores the space of possible languages the same way artificial life explores the space of possible life forms [22]. Moreover the languages are not considered to be static. Attempts are made to have them evolve in ways that are similar to human language evolution.

This paper surveys some of this research which is relevant in several ways to the general questions posed by biologically-inspired agents research:

1. Communication is obviously a very important feature of higher animals, particularly humans. Indeed it has been argued that it is through the increasing power and needs of communication that cognition has been bootstrapped to human level intelligence in the first place. The study of communication and its complexification therefore fits within the general biological study of the ‘major transitions in evolution’ [27].

2. Research into the origins and evolution of language introduces a whole new approach towards the problem of getting autonomous robots to communicate in natural language with humans or each other. Instead of the traditional top-down approach used in most AI work on natural language processing, in which lexicon and grammar remain static, communication is seen as adaptive behaviour. The robot progressively acquires more complex forms of language-like communication – similar to a child which progresses from babbling to prelinguistic communication, then to simple forms of lexical language and finally full-blown grammar. Lexicon and grammar continue to change and develop throughout life. The emphasis on embodiment and situatedness naturally resonates strongly with work on behaviour-based robotics, particularly more recent work which tries to establish attention sharing, turn-taking and emotional communication [4].
3. As I will try to show in this paper, research on artificial language evolution has benefited greatly from adopting principles discovered in the study of the origins of biological complexity. But benefits could also flow the other way. Language is a very good domain to study how communication may self-organise and complexify. This topic is still only weakly understood in biology and is relevant to questions such as how do information codings arise in the brain, how do different organs within a body develop the necessary communication to coordinate their activity, or how has the genetic code evolved towards such great complexity.

Research on artificial language evolution is not only of interest for the study of autonomous adaptive agents but is receiving increased attention from other scientific disciplines interested in the evolution of language as well [14], [56].

The objective of this paper is to survey some of the work done so far with an eye on finding some of the principles that have proven to be relevant for simulating some form of language emergence. This survey is necessarily very brief and inevitably biased to my own research. I do not pretend that all researchers in the area subscribe to these principles. Moreover the complete problem is far from solved. Much remains to be discovered. Nevertheless, I believe it is relevant to occasionally perform this kind of synthesis to expose the gaps in our understanding and begin to exploit application opportunities. Complementary surveys and examples of research can be found in [7], [5], [15].

2 Language Games

Much of the success of research in artificial language evolution has come from framing the problem in terms of games. Before the advent of artificial life, game theory was already widely used in theoretical biology to study aspects of genetic evolution and animal cooperation [25]. Indeed, some of the early successes in Artificial Life have come from adopting this framework as a basis for studying complex dynamics and the evolution of complexity. For example, Lindgren [24], Kaneko, [17] and others studied the iterated prisoner's dilemma game and showed various evolutionary phenomena such as spatio-temporal chaos, co-evolution of

strategies, etc. Research in artificial language evolution has started to take off when iterated games were adopted as a way to study the emergence and evolution in language. One of the earliest examples is [28].

The mapping to language works as follows. Typically there is a population of agents (which can be static or dynamic). Two players (a speaker and a hearer) are randomly drawn from the population. The players engage in an interaction which is either a complete language game, i.e. a communication involving the real world, or an aspect of a language game, for example, only the exchange of sounds [9] or the exchange of a string together with what the string means [32], or well-formed sentences generated by an evolving grammar [13]. The players have their own private cognitive structures, like lexicons or grammars, which they use to play the game, just as they have their own strategy and memory in the case of iterated prisoner's dilemma games. There is an outcome of the game, for example, successful communication or successful imitation, and the behaviour of the players changes over time in such a way that success increases. The language phenomena are a side effect of repeated games. Language conventions are not put in from the start, and there is no central agency that controls how agents are supposed to act. Shared communication conventions must emerge from the distributed activity of the agents.

Thus an experiment in artificial language evolution always has the following ingredients:

- A definition of an interaction protocol for the agents.
- A definition of the architecture of an agent (what cognitive structures are available, what input/output processing is done, how learning and language invention proceeds).
- An environment, possibly a real world environment if the agents are robotic.
- A set of measures which show that the language phenomena one is interested in indeed arise. For example, success in communication, growing size of the lexical repertoire, similar sound systems as in natural languages, grammatical structures, etc.

3 Genes versus Viruses

An important difference can be seen between a group of researchers that view language as the result of genetic evolution (e.g. [31]) and those which emphasise cultural evolution [42] (with various researchers taking an in-between position as well as in the case of Baldwinian evolution [5]). This mirrors disputes in linguistics between researchers like Chomsky and Pinker who defend nativist positions versus researchers like Lindblom, MacNeilage and Studdert-Kennedy [23] or Tomasello [51] who defend the social and cultural construction, learning, and evolution of language. All of these variations can be explored within the language game framework and so it is possible to compare the strength and weakness of each approach [45].

Genetic evolution is modeled by adopting the same framework as used in research on genetic algorithms [20]. The population is divided into different generations. A particular generation plays a set of games and individuals receive a score on how well they are doing in the game. The cumulative score determines the fitness of an agent. Then there is the creation of a new generation based on the previous one. The probability of having offspring is based on prior success in the game and the offspring inherits the linguistic knowledge of the parent(s), with possible mutations or recombinations. The computer simulations of [28] follow this framework. It has been clearly demonstrated by many simulations that this leads to the emergence of a shared set of conventions in a population. Mutation and recombination operators are the only way new structure can arise into the language. Language coherence arises because the same ‘language genes’ eventually spread in the total population.

In the alternative view, language conventions spread similar to the propagation of bacteria or viruses. Language evolution is viewed as driven by cultural (memetic) evolution rather than genetic evolution. In this case, there is no division into different generations, although it is still possible that there is a population flow, with agents entering and leaving. There is no fitness associated with the agents. Agents do not inherit anything from parents. The notion of offspring does not exist. Instead each agent adjusts his linguistic behaviour after each game in order to be better in the next game. Adjustment could mean: change the score in memory between a wordform and its meaning, invent a new word for a specific meaning, add a new sound to the phonetic repertoire, invent or adopt a new grammatical rule, etc. The computer simulations of [40], [32] or [9] and the robotic experiments in [49] or [47] all follow this particular framework. Again it has been demonstrated beyond doubt that cultural evolution also leads to the emergence of a shared set of conventions in a population. Language coherence now arises through self-organisation in the sense of complex systems theory [29] (as explained below).

There have been experiments which use a mixed form. For example, the simulations reported by Kirby [18] are structured like in genetic models. The population is divided into different generations and agents get a score which results in a fitness measure. But the language is learned by each generation from the previous one, as opposed to genetically inherited. In this learning process more structure (specifically more abstract rules) are introduced by the agents. So language does not evolve through mutation and recombination but in a cultural fashion. On the other hand, language coherence is still partly influenced by inheritance relations because success in the game influences whether an agent will have offspring or not.

4 Grounding

There have been important advances in robotics lately, largely due to adopting a behaviour-based approach [46] [34]. It has become more and more realistic to build robust autonomous robots which interact in real time with a dynami-

cally changing world. Behaviour-based robots use an architecture that couples sensing almost directly to actuating, de-emphasising complex internal symbolic representations. Moreover they include motivational and emotional parameters in deciding which action path to pursue. Together with rapid advances in mechanical and electrical engineering, the behaviour-based approach has led to complex mass produced pet robots such as Sony's AIBO and is leading to a new generation of humanoid robots [19]. All these advances are a tremendous opportunity for research on artificial language evolution because it becomes possible to implement language games on such robotic platforms and thus investigate fully grounded situated verbal communications between autonomous robots. Several researchers have been trying to do this [55], [1], [48], [36].

There are two key issues to be solved, known as the grounding issue and the bootstrapping issue. The grounding issue concerns the problem of relating the conceptualisations underlying a language utterance to the external world through a sensori-motor apparatus. Agents must implement the full semiotic cycle. That means, the speaker must perform the necessary pattern recognition and sensory processing on captured images and or types of sensory data, conceptualise the scene by categorising objects and events, verbalise this conceptualisation, and transmit it to the hearer. The hearer must decode the utterance, and confront the interpretation with his own conceptualisation of the sensory image.

The grounding problem is an active field of research at the moment [12], [8] but there does not appear to be a simple straightforward solution, in the sense of a component that could be added to make a non-grounded agent grounded in external reality. Instead, grounding is a matter of setting up tight couplings between the behaviours of the agent and his environment on the one hand and the internal representations that are used on the other. It is the result of a total integrated process, in which adequate pattern recognition and image processing provides the ground work and adaptive categorisation algorithms (based on weighted decisions, nearest neighbor computation, discrimination trees, etc.) play key roles. In the case of language games, there is an additional complexity, namely the grounded representations constructed by the lower cognitive levels must be in tune with the language systems that verbalise or interpret these representations. Because both grounded representations and language are evolving systems, we need a way to coordinate them without a central coordinator or prior knowledge. I will argue below that the principle of structural coupling discussed below is relevant for this.

5 Linguistic Bootstrapping

The second issue for evolving grounded communication on embodied robots is how verbal communication itself can be bootstrapped. This is related to the general problem of the origins of communication which has also been studied in adaptive behaviour research [30]. Of course it is possible to pre-program the agents, in other words pre-program the game, but that would put into the robots

the processes we try to understand and explain. Instead we want to understand the process by which language gets bootstrapped, and empirical research shows that this is not an individualistic process. There is an important role for a ‘mediator’ that scaffolds the complexity, provides pragmatic feedback, and motivates learning [51].

As in the case of grounding, there is not a single magical trick to explain linguistic bootstrapping but many competences need to be integrated. Careful observations by developmental psychologists, following in the footsteps of Piaget and Bruner, have shown that ‘learning how to mean’ is a slow process which takes roughly 8 months starting from 6 months of age, and is estimated to involve as much as 50,000 interactions. The presence of a mediator is absolutely crucial. Verbal communication (initially with single words which only approximately sound like standard words) implies that (1) the speaker has an effect on the hearer (communicative effect), (2) the hearer interprets the speaker’s behaviour as communication (communicative inference), and (3) the speaker intends her behaviour to be communicative (intentional communication).

The observed developmental sequence is roughly as follows (see [11]):

1. Communicative effect: Infant acts (cries, kicks) => Caregiver reacts to these behaviours.
2. Communicative inference: Infant develops goal-directed behaviours (e.g. reach for toy while making sound) => Caregiver infers the intention and responds with appropriate behaviour. Caregiver also typically re-enforces the sounds and corrects.
3. Intentional communication: Infant realises power of communication and starts to use it deliberately. Communication includes vocalisation, eye contact, as well as gestures.
4. Upping the ante: The caregiver starts to require more precise vocalisations that resemble words used in the language.

Notice that the role of a caregiver as interpreter of behaviour is crucial, otherwise the infant cannot learn that vocalisations can have certain effects, and climb up the hill of more conventional and more complex language use.

So far there have been no convincing simulations of this developmental sequence although preliminary efforts have been going on in this direction [48]. It is obvious that there are many preconditions which are extremely difficult to realise on autonomous robots and which co-develop at the same time as language communication bootstraps. They include: localising and recognising other human beings, eye contact and gaze following, producing vocalisations (babbling), emotion recognition and production through sound, gesture tracking and interpretation, sharing attention with others to specific objects or actions, which implies segmentation, template matching and tracking, realising that actions can have causal effects, realising that to achieve an effect, the action needs to be performed that causes this effect, realising that a vocalisation is equivalent to such an action, adjusting a vocalisation so that it comes closer to a vocalisation heard by the caregiver, etc. Each of these competences has been the object of intense investigation lately by AI researchers, mostly in the context of humanoid

robotics research. The work of Breazeal [4] on emotional attention sharing and turn taking, Scassellati [37] on face identification and tracking, Oudeyer [33] on babbling and emotion expression, are some examples in this direction. Only when all these components can be integrated in a single system can we begin to simulate human-like linguistic bootstrapping.

6 Self-Organisation

We now return to the collective level. One of the key questions to understand how a communication system can arise, is how there can be coherence in the group, in other words how distributed agents without a central authority and without prior specification can nevertheless arrive at sufficiently shared language conventions to make communication possible. The genetic evolution hypothesis of language evolution ‘solves’ this problem by considering that successful language genes spread in the population, so after some time everybody shares a copy of the same most successful gene. However genetic evolution is extremely unlikely for most aspects of language (definitely for the lexicon, and even for grammar – there seems too much variation between languages to encode much if anything genetically [54]). However an alternative solution is available that could explain how coherence can arise in a cultural fashion, namely through self-organisation.

The concept of self-organisation (narrowly defined) has its roots in research in the fifties and sixties on certain types of chemical reactions such as the Belousov-Zhabotinsky reaction [29]. It then became generalised to many different types of systems not only physical but also biological [6] and even economical. Since the beginning of Artificial Life and Adaptive behaviour research, simulations of the self-organisation of ant paths, bird flocks, slime molds, pattern formation in morphogenesis, etc. have been common, with applications to collective robotics [10].

Self-organisation occurs when there is a system of distributed elements which all have a random behaviour in the equilibrium state. The system is then brought out of equilibrium, which is usually by the supply of energy in physical systems. A positive feedback loop becomes active, enforcing local fluctuations into coherent global behaviour. In the well-studied case of ant societies [10], an ant hits a food source in random exploration, and then returns to the nest depositing a pheromone. This attracts other ants, which enforce the chemical trail, attracting even more ants, etc. (the positive feedback effect). Progressively the whole group self-organises to a single path. When food is exhausted, no more pheromone is deposited and the chemical evaporates returning the system to a random exploration (i.e. equilibrium) stage. Self-organisation in this sense has now been studied extensively from the viewpoint of dynamical systems theory and a large body of mathematical models and techniques exist to describe it.

Around 1995, it became clear that this mechanism could also be applied to language evolution. It was first shown for lexicon formation (see [40], [32]) but then generalised to other aspects of language, including phonetics [9]. The application for the lexicon works as follows. Suppose speakers invent new words

for the meanings which they do not know how to express and listeners store the words used by other agents. In this case, agents will develop words for all meanings and adopt them from each other. However the lexicon will be very large. Many different words will be in use for the same meaning. But suppose now that a positive feedback is introduced between use and success: Agents keep a score for each word-meaning pair in their lexicon. When a game is successful the score of the word-meaning pair that was used increases, and that of competing word-meaning pairs is decreased (lateral inhibition). When a game fails, the score of the word-meaning pair is diminished. In interpreting or producing language, agents use the word-meaning pairs with the highest score. These dynamics indeed gives self-organisation towards a shared lexicon (figure 1). So it suffices to program the adaptive behaviour of individual agents in such a way that a positive feedback loop arises between use and success and self-organisation sets in.

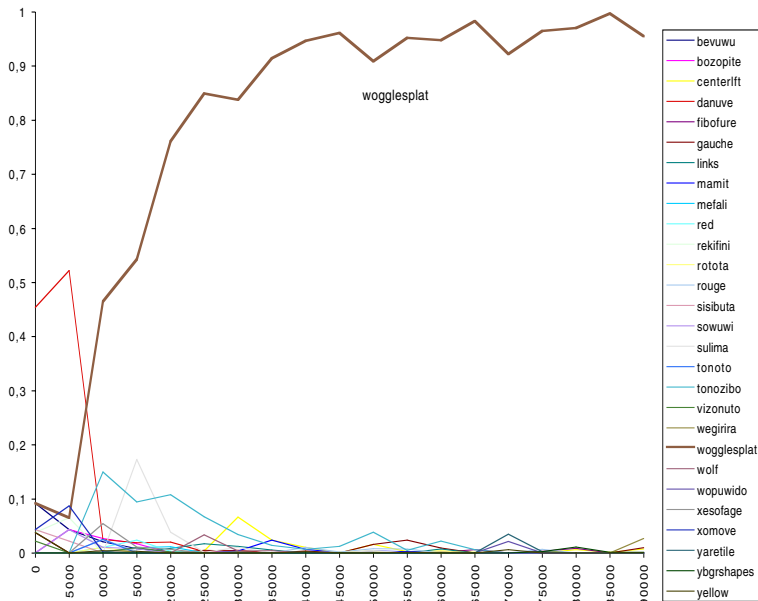


Fig. 1. This graph plots the usage rate of all possible words for the same meaning in 100,000 iterated language games played by a group of over 1000 agents. Initially many words are competing until one dominates due to a winner-take-all effect.

The adoption of self-organisation is a nice example where a principle from biology (in fact complexity science in general) could first be demonstrated in artificial life simulations and then transported into ‘artificial language evolution’.

7 Structural Coupling (Co-evolution)

Another key problem for artificial language evolution is how the different levels of language, which each have their own developmental trajectory, can become coordinated with each other. For example, how can the meanings underlying language become coordinated with the lexicon? There are profound differences between languages as far as their conceptualisations are concerned [50]. For example, the conceptualisation of the position of the car in “the car is behind the tree” is the opposite in most African languages compared to Western languages. The front of the tree is viewed as being in the same direction as the face of the speaker and hence the car is conceptualised as in front of the tree as opposed to behind the tree. Examples like this are not hard to find and they suggest that different human cultures invent their own ways to conceptualise reality and propagate it through language, implying a strong causal influence of language on concept formation (the Sapir-Whorf thesis) [3]. The same problem arises for the coordination between phonetics/phonology and the lexicon. The sound system of a language evolves independently, but this change creates effects on other language levels. For example, the loss of a case system in old English is generally attributed to phonetic effects which caused the case-markers at the end of words more difficult to perceive. Grammaticalisation processes commonly observed in natural language evolution [52] show that there is a strong interaction as well between lexicon and grammar. Typically certain lexical words become recruited for syntactic functions, they progressively lose meanings, become shorter, and may even disappear altogether so that the cycle of grammaticalisation restarts again.

A principle from biology has once again turned out to be helpful to understand how the co-evolution between different subsystems involved in language may be achieved. In the early nineteen seventies, Maturana introduced the concept of structural coupling and developed it further with Varela [26]: Given two adaptive systems operating independently but having a history of recurrent interactions in the same shared environment, a ‘structural congruence’ may develop under certain circumstances, so that they become coordinated without a central coordinator. It is important that each adaptive system acts as a perturber of the other, and, because they are adaptive, the perturbation leads to a structural change. Structural coupling has come out of attempts to understand certain biological phenomena, such as the development of multi-cellular systems or the coordination between organs. It is a system-level concept which has found application in areas ranging from physics to economics or social science. The concept is related to so called coupled maps [17] which are dynamical systems, for example systems of oscillators, where one subsystem acts as a context for the other.

The relevance of structural coupling to artificial language evolution has also become clear around 1995, particularly in the context of coordination between conceptualisation and lexicon formation [41], [16]. Both systems have to be adaptive: conceptualisation requires a mechanism that can generate new categories driven by the need for communication, for example, new distinctions may have to be introduced in order to refer to objects within a particular context. Lex-

icon formation is also adaptive because new words need to be invented or are being learned from others. Each system perturbs the other. The lexicon may push the conceptualisation system to develop new categories or categories that are also employed by other agents. The conceptualisation system occasionally comes up with categories that have not been lexicalised yet, so it perturbs the lexical system to make structural changes as well. Both systems have a history of interactions, not only in single agents but also in a group of agents. If the right structural coupling is set up, it can be shown that not only lexicons but also the conceptual repertoires underlying these lexicons can self-organise and become coordinated.

Figure 2 from [45] shows an example of this. In this experiment, the agents play language games about coloured stimuli (corresponding to the Munsell samples widely used in the anthropological literature). Given a set of samples, the hearer has to identify one of them based on a colour name provided by the speaker. The colour name assumes a certain categorisation of reality (for example green and blue colours) which the agents have to develop at the same time as they are developing from scratch a lexicon for naming these categories. Categorisation fails if the agent does not have a category in his repertoire that distinguishes the colour of the chosen sample from the other colours. For example, if there is a blue, green and red sample, and the blue one is chosen, then it will be necessary to have a colour category for blue which distinguishes blue from green and from red. In the experiment reported in [45] there is a structural coupling between the lexicon formation and concept formation processes, leading to progressive coherence of the categorial repertoires. If there is no such coupling and agents individually develop categories to distinguish samples, individual repertoires adequate for colour categorisation still develop but they are no longer similar. Figure 2 displays the evolution over time of category variance with (top graph) and without (bottom graph) structural coupling. The ratio between the two demonstrates how categorical similarity is drastically increased when there is a structural coupling.

8 Theory of Mind

The previous sections discussed mostly research in the domain of lexicon and concept formation. The problem of grammar has turned out to be much more difficult to crack and there is no consensus yet on how it should be done. In a series of intriguing simulations, Kirby and coworkers [18], [2] showed that in iterated games where agents from one generation learn grammars from the output from the previous generation, agents will choose a compositional as opposed to a non-compositional language because this overcomes the learning bottleneck, i.e. the problem that agents have to learn a language from limited data. In this case, learners (i.e. children) play a crucial role in shaping the future of a language. This approach has been confirmed by theoretical results of Nowak, et.al. [31].

But there is an alternative view, namely that grammar arises to optimise communication [44]. Speakers try to increase the chance of being understood

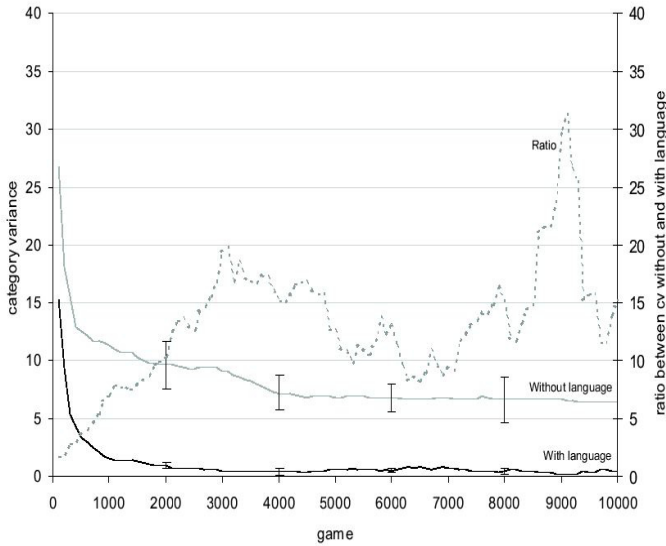


Fig. 2. The graph displays the variance between the emerging category sets used by a population of agents playing iterated language games, with (top) and without (bottom) a structural coupling between lexicon formation and category formation. The ratio between the two is displayed as well.

correctly by making additional aspects of meaning explicit and by minimising the processing that needs to be done by the hearer (and by themselves). Of course the grammatical rules that speakers introduce must still be learnable – otherwise they would not propagate in the population. Moreover in the adoption of rules used by others, a listener may overgeneralise, or a listener may overinterpret certain formal characteristics of an utterance to be carriers of meaning, whereas they were not intended to be so. This would also introduce additional structure and regularity as soon as the learner uses these rules in his own language production. Nevertheless the creative force in language evolution from this alternative perspective rests primarily with language producers.

Recent experiments [44] have shown examples how all this might work. The first important step is to view natural language as an inferential coding system [39], which means that the sender assumes that the receiver is embedded in the same context and is intelligent enough to infer commonly known relevant facts about the current situation. The message is therefore incomplete and cannot be interpreted without the context. This contrasts with Shannon-like pure coding systems where the message carries all the meaning that the sender wants to transmit. Inferential coding systems can transmit much more information with fewer means, however, there is a risk of misunderstanding and there is a risk that the hearer has to do more work than he is willing to do to interpret the message.

This is why grammatical elements (as well as additional lexical elements) get introduced.

In the experiment reported in [44], the speaker simulates the understanding of his own utterance as part of language production and detects potential difficulties. The experiments focus on case grammar, which centers around case markers that help to express the relations of objects with respect to events (as in ‘He gave her the ball’ versus ‘She gave him the ball’). It is possible to communicate without explicating these event-object relations, and often they can be inferred from the present situation and context. But most languages have developed grammatical tools to express event-object relations to minimise the risk that roles get confused. For example, English uses word order, German or Latin use case affixes, Japanese uses particles, etc. In the experiment, agents detect where ambiguity or uncertainty arises and repair it by introducing additional (case)markers. The hearer assumes that unknown elements of the utterance are meaningful and are intended to help in interpretation. When the hearer can construct an interpretation, this helps to figure out the possible meaning of unknown utterances.

The main mechanism to simulate these processes is to introduce a subsystem to infer how the listener will interpret a sentence in a particular context, which amounts to a kind of ‘theory of mind’ of the other. The growing complexity of robots and the rise of humanoid robots makes this more feasible because these robots are much more situated and therefore have more information available than is relevant to sustain a grounded communication [38]. Moreover the speaker can use himself as a model to predict how the other agent will react.

9 Further Evolution

Language not just self-organises once, but evolves, and sometimes very rapidly [21] – which is one of the reasons why it is implausible that language evolution is entirely genetic. Even without population change and throughout the life time of an individual, new words are introduced, meanings of words shift, grammatical rules change, the phonetics undergoes change, etc. When human populations with mixed languages are put together and change rapidly, creoles may form which recruit elements from source languages but re-invent many grammatical phenomena like expression of tense, aspect, mood, case systems, reflexivity, determiners, etc.

Evolution requires variation and selection. These can easily be mapped onto language evolution. As soon as there is a distributed set of agents which each evolve their own communication system, variation is inevitable. An individual’s language behaviour is affected by past developmental history: what environments were encountered, with which other agents most interactions took place, what random choices were made. Additional variation may come from the inevitable stochasticity in language communication: errors in the transmission or reception of the spoken signal, errors in pragmatic feedback, processing errors in parsing and production. There are many selective forces at work, ranging from physi-

ology (particularly important for constraining the kinds of speech signals that can be produced and the kinds of sensori-motor data that is available for conceptualisation), the environment, the ecology (what are important distinctions), cognitive constraints (memory, processing speed, learning speed), the dominating conventions adopted by the group, and the specific communicative tasks that need to be realised. A language system is never going to be optimal with respect to all these constraints. For example, sometimes parts of words are no longer pronounced to make the utterance shorter but this may lead to a loss of information (such as case marking) which then gives rise to grammatical instability that needs to be solved by the re-introduction of markers or by a shift to another kind of system [53].

10 Conclusions

The paper has presented a number of principles that are being explored by a growing group of researchers to explore artificial language evolution. This field attempts to set up systems of autonomous distributed agents that self-organise communication systems which have properties comparable to human natural languages. The agents are either software agents or fully embodied and situated robotic agents. The relevance to adaptive behaviour research is twofold: On the one hand the study of language evolution provides insight into a number of processes generating complexity in communication systems. These processes appear similar to mechanisms generating complexity in other areas of biology. Self-organisation, structural coupling, level formation and cultural selection are examples. On the other hand, the study of how complex communication has evolved is giving new ways to implement open-ended grounded communication with autonomous robots, and to simulate the epigenetic development of cognition.

Discussion of general principles is risky but at the same time necessary because it is only at this level that bridges between fields, particularly between biology and evolutionary linguistics, can be established. Moreover I want to emphasise again that much remains to be discovered. The principles reported here are far from complete and need to be explored in many more case studies.

Acknowledgement. I am indebted to members of the Artificial Intelligence Laboratory in Brussels who have been involved in several experiments reported in this document, financed by the Belgian Government GOA grant, in particular, Tony Belpaeme (financed by the Belgian FWO) who implemented the colour naming and categorisation experiments. I am also indebted to members of the ‘origins of language’ group at the Sony Computer Science Laboratory in Paris, particularly Frederic Kaplan who implemented the language games on the AIBO, and Pierre-Yves Oudeyer who has focused on experiments in evolving phonetic systems.

References

1. Billard, A. and K. Dautenhahn (2000) Experiments in social robotics: grounding and use of communication in autonomous agents. *Adaptive behaviour* vol. 7:3/4.
2. Brighton, H. (2002). *Compositional Syntax from Cultural Transmission*. *Artificial Life*, 8(1)
3. Bowerman, M. and S. Levinson (2001) *Language acquisition and conceptual development*. Cambridge University Press, Cambridge.
4. Breazeal, C. (1998) *A Motivational System for Regulating Human-Robot Interaction*, Proceedings of AAAI98, Madison, WI.
5. Briscoe Edward J. (ed.) (2002) : *Linguistic Evolution Through Language Acquisition: Formal and Computational Models*, Cambridge, UK, Cambridge University Press.
6. Camazine, S. J.-L. Deneubourg, N. Franks, J. Sneyd, G. Theraulaz, and E. Bonabeau (2001) *Self-Organization in Biological Systems*. Princeton University Press, Princeton.
7. Cangelosi, A. and D. Parisi (eds.) (2001) *Simulating the Evolution of Language*. Springer-Verlag, Berlin.
8. Cohen, P., et.al. (2001) *Proceedings of the AAAI Spring Symposium on Grounding*. AAAI Press, Anaheim, Ca.
9. De Boer, B. (1997) *Self-Organisation in Vowel Systems through Imitation*. In: P. Husbands and I. Harvey (eds.) *Proceedings of the fourth European Conference on Artificial Life*, Cambridge Ma. (1997)
10. Deneubourg, J.-L., et.al. (1993) *Self-organisation and life: from simple rules to global complexity*. Proceedings of the Second European Conference on Artificial Life, ULB. Brussels.
11. Harding, C. G. (1983) *Setting the Stage for Language Acquisition: Communication Development in the First Year*. In: Golifkoff, R. (1983) (ed.) *The Transition from Prelinguistic to Linguistic Communication*. Lawrence Erlbaum Ass. Hillsdale NJ. p. 93–113.
12. Harnad, S. (1990) *The Symbol Grounding Problem*. *Physica D* 42: 335–346.
13. Hashimoto, T. and Ikegami, T., (1996), *Emergence of net-grammar in communicating agents*, *BioSystems*, 38, 1–14, 1996.
14. Hurford, J., C. Knight and M. Studdert-Kennedy (1998) *Approaches to the Evolution of Language: Social and Cognitive bases*. Cambridge University Press, Cambridge. pp 405–426.
15. Hurford, J. (2001) *Expression / Induction Models of Language Evolution: Dimensions and Issues*. In: Briscoe, T. (2001) *Linguistic Evolution through Language Acquisition: Formal and Computational Models*, Cambridge University Press.
16. Ikegami, T. and M Taiji (1999) *Imitation and Cooperation in Coupled Dynamical Recognizers*. In: Floreano, D. J. Nicoud and F. Mondada (eds.) *Advances in Artificial Life (ECAL 99)*. Lecture Notes in Computer Science.
17. Kaneko, K. and J. Suzuki. (1994) *Evolution to the edge of chaos in imitation games*. In: *Artificial Life III*, the MIT Press, Cambridge Ma. pp. 43–54.
18. Kirby, S. (1999) *Function, Selection and Innateness: The Emergence of Language Universals*. Oxford University Press, Oxford.
19. Kuroki, Y., T. Ishida, J. Yamaguchi, M. Fujita, T. Doi (2001) *A Small Biped Entertainment Robot*. In: *Proceedings of the IEEE-RAS International Conference on Humanoid Robots*. Waseda University, Tokyo. p. 181–186.

20. Koza, J. (1992) *Genetic Programming: on the programming of computers by means of natural selection*. MIT Press, Cambridge Ma.
21. Labov, W. (1994) *Principles of Linguistic Change. Volume 1: Internal Factors*. Blackwell, Oxford.
22. Langton, C. (ed.) (1989) *Artificial Life*. Addison-Wesley, Reading Ma.
23. Lindblom, B, P. MacNeilage and M Studdert-Kennedy (1984) Self-organizing processes and the explanation of language universals. In: Butterworth, B, et.al. *Explanations for language universals*. Walter de Gruyter and Co, pp. 181–203.
24. Lindgren, K. and M. Nordahl (1994) Cooperation and Community Structure in Artificial Ecosystems. *Artificial Life Journal* 1(1).
25. Maynard Smith, J. (1982) *Evolution and the theory of games*. Cambridge University Press, Cambridge.
26. Maturana, H. and Varela, F. (1998) *The Tree of Knowledge (revised edition)*. Shambhala Press, Boston.
27. Maynard Smith, J. and E. Szathmary (1995) *The Major Transitions in Evolution*. Oxford University Press, Oxford.
28. MacLennan, B. (1992) 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.
29. Nicolis, G. and I. Prigogine (1988) *Exploring Complexity*. Freeman and Co. New York.
30. Noble, J (1999). Cooperation, conflict and the evolution of communication. *Adaptive behaviour*, 7(3/4), 349–370.
31. Nowak, M., J. Plotkin and J. Vincent (2000) The evolution of syntactic communication. *Nature*, vol. 404 (30/03/2000), pp. 495–498.
32. Oliphant, M. (1997) *Formal Approaches to Innate and Learned Communication: Laying the Foundation for Language*, Ph.D. Thesis, University of California San Diego, Cognitive Science Department.
33. Oudeyer P-Y. (2001) Emotional Interactions with Humanoids Using Speech. In: *Proceedings of the IEEE-RAS International Conference on Humanoid Robots*. Waseda University, Tokyo. pp. 17–24.
34. Pfeifer, R. and R. Scheier (2000) *Understanding Intelligence*. The MIT Press, Cambridge Ma.
35. Pinker, S. (1994) *The Language Instinct. The New Science of Language and Mind*. Penguin, Harmondsworth.
36. Roy, D. (2001) Learning Visually Grounded Words and Syntax of Natural Spoken Language. *Evolution of communication* 4(1).
37. Scassellati, B. (1998) Eye Finding via Face Detection for a Foveated, Active Vision system. In *Proceedings of AAAI-98*. AAAI Press Books.
38. Scassellati, B. (2002) *Foundations for a Theory of Mind for a Humanoid Robot*. Ph.D. Thesis, MIT.
39. Sperber, D. and D. Wilson (1986), *Relevance: Communication and Cognition*. Harvard University Press, Cambridge Ma.
40. Steels, L. (1996) Self-Organising Vocabularies. In: Langton, C. *Proceedings of Alife V*. The MIT Press, Cambridge.
41. 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.
- Steels, L. (1998) The origins of syntax in visually grounded agents. *Artificial Intelligence* 103 (1998) 1–24.

42. Steels, L. (2001) Language Games for Autonomous Robots. IEEE Intelligent systems, September/October 2001, p. 16–22.
43. Steels, L. (2001) Social Learning and Verbal Communication with Humanoid Robots. In: Proceedings of the IEEE-RAS International Conference on Humanoid Robots. Waseda University, Tokyo. pp. 335–342.
44. Steels, L. (2002) Computer simulations of the origins of case grammar. Fourth Evolution of Language Conference, Harvard, Cambridge Ma.
45. Steels, L. and T. Belpaeme (2003) Computer Simulations of Colour Categorisation and Colour Naming. Submitted to BBS.
46. Steels, L. and R. Brooks (eds.) (1994) The artificial life route to artificial intelligence. Building situated embodied agents. Lawrence Erlbaum, Hillsdale, NJ.
47. Steels, L., F. Kaplan, A McIntyre and J. Van Looveren (2001) Crucial factors in the origins of word-meaning. In Wray, A., editor, *The Transition to Language*, Oxford University Press. Oxford, UK, 2002.
48. Steels, L. and F. Kaplan (2001) AIBO's first words. The social learning of language and meaning. *Evolution of Communication* 4(1).
49. Steels, L. and P. Vogt (1997) Grounding Adaptive Language Games in Robotic Agents. In Harvey, I. et.al. (eds.) *Proceedings of ECAL 97*, Brighton UK, July 1997. The MIT Press, Cambridge Ma.
50. Talmy, L. (2000) *Toward a Cognitive Semantics: Concept Structuring Systems (Language, Speech, and Communication)* The MIT Press, Cambridge Ma.
51. Tomasello, M. (1999). *The Cultural Origins of Human Cognition*. Harvard University Press.
52. Traugott, E. and Heine, B. (1991) *Approaches to Grammaticalization*. Volume I and II. John Benjamins Publishing Company, Amsterdam.
53. Van Kemenade, A (1987) *Syntactic Case and Morphological Case in the History of English*. Forist Publications, Dordrecht.
54. Vogel, P.M. and B. Comrie (eds.) (2000) *Approaches to the Typology of Word Classes (Empirical Approaches to Language Typology, 23)* Cambridge University Press, Cambridge.
55. Vogt, P. (2001) Bootstrapping grounded symbols in minimal autonomous robots. *Evolution of Communication* 4(1).
56. Wray, A.,et.al. (ed.) (2002): *The Transition to Language*, Oxford, UK, Oxford University Press.