The Evolution of Grounded Communication for Robots

Luc Steels
SONY Computer Science Laboratory - Paris.
AI Laboratory Vrije Universiteit Brussel

February 3, 2003

Abstract

The computational and robotic modeling of language evolution is emerging as a new exciting subfield in cognitive science. The objective is to come up with precise operational models how communities of agents, equiped with a cognitive apparatus, a sensori-motor system, and a body, can arrive at shared grounded communication systems that have similar characteristics as human languages. Apart from its technological interest for building novel applications in the domain of human-robot or robot-robot interaction, this research is potentially relevant to the many disciplines interested in the origins and evolution of language.

1 Evolutionary Linguistics

Artificial Life developed in the early nineteen nineties as a field that investigates principles of living systems by building artificial systems that model some of their key aspects. For example, researchers have built operational models of pattern formation in biological systems (as observed on shells or in nest structures of insect societies) that generate similar structures as those found in nature [12]. These models thus give a causal insight in how these phenomena arise. The notion of a complex adaptive system, in which many independent elements dynamically coordinate their activity through self-organisation, has emerged as a key concept. Several other system concepts from biology, such as evolution by natural selection, structural coupling, order out of chaos, and network dynamics, could be operationalised and shown to be relevant to understand a wide range of natural phenomena.

The interest of this work is not only theoretical. Genetic evolution, as captured in artificial life models known as genetic algorithms, have made it possible to solve a large set of engineering problems by using the search power implicit in selectionist processes [22]. Also in robotics, the artificial life approach, which insists on a bottom up emergence of complexity and the role of collective dynamics in behavior, has had an important impact and is largely responsible for the impressive jump forward in autonomous robots accomplished the past decade [36]. More recently, Artifical Life models are beginning to impact research in biomolecular information processing [16].

Within the same methodological framework and pool of ideas, various researchers have started to look in a very novel way at language. They began to build computer simulations and robotic experiments in which artificial communication systems emerge, invented and learned by artificial agents. The objective is again twofold.

The first goal is to make precise models of how certain key properties of language, like compositionality, might have originated, and how and why languages continue to evolve the way they do. Quite naturally, language became viewed as a living system, that self-organises and continues to evolve through the collective dynamics of agents engaged in situated verbal interactions. The insights and apparatus for studying complex adaptive systems, as emerging from Artificial Life research, thus became relevant for the study of language as well. A new field, evolutionary linguistics, was born with a specific focus (understand the origins of language and meaning), a specific hypothesis (language is a complex adaptive system), and a specific methodology (construct artificial systems as a way to develop and test theories). An overview of earlier work in this field can be found in [31] and more recent work in [21]. Several collections of papers ([5], [8], [43]) provide additional source material.

Evolutionary linguistics has a secondary, more practical goal, namely to forge a new technology for communication between humans and robots, or among robots. Recently, tremendous advances in robotics and artificial intelligence have given us fully operational autonomous robots, even humanoid robots walking on two legs, with stereo vision, surround audition, real-time adapted dynamical trajectory planning, vision-based navigation, and many other features [23]. A new breed of 'pet' robots, such as the Sony dog-like AIBO robot, has been commercialised, whose primary aim is not to perform practical tasks but to interact and entertain people. The robots help to create a semi-artificial world, similar to the world of opera or puppet theatre, in which various forms of pretend play are spontaneously initiated by humans. Despite their playful appearance, these robots are tremendously complex, both from a hardware and software point of view.

2 Open Issues

The new generation of robots raises the question what kind of interactions are most appropriate and what forms of communication are best suited to support them. Three issues stand out. First of all, the traditional approach, which is based on pre-programmed communication protocols (as in a text processor or an industrial robot), cannot work, because these complex robots need to be capable of an open-ended set of possible interactions with the environment and human beings. Humans are known to negotiate shared conventions as part of ongoing interactions [9]. So the robots will always have to cope with new meanings, new variations in speech from unknown users, new words or shifts in the meaning of existing words, and new grammatical constructions.

Existing inductive machine learning approaches, whether based on neural networks or symbolic machine learning [6], have not yielded anything like what is needed for this kind of fluent grounded open-ended communications with humans. These inductive approaches require large data sets which are simply not available here, because the robots will always encounter profoundly new situations and language expressions.

Evolutionary linguistics suggests an alternative solution. Robots could be programmed to have the capacity to participate in the invention of a shared communication system. They could constantly adapt their conceptualisations and modes of expression to cope with new challenges of the environment or new speakers. In contrast with statistical learning approaches, the robot should actively co-construct the language system, rather than passively induce regularities that exist. The language system is not fixed but emerges and is shaped and reshaped in language use, in the spirit of Hopper's emergent grammar [17].

Doing all this requires obviously deep insights into the creative side of language. For example, we have to understand better how analogy is used to expand the use of an existing construction for new purposes [15]. We also need insights in how a newly invented construct can become conventionalised and spread in a population.

Second, the communication needs to be embedded in the framework of social interaction. This topic has received increasing attention recently in the robotics literature, with applications in educational software and autistic therapy [10]. Social interaction relies on a wide variety of competences [3]. A partner in communication must be identified and tracked in real-time, some form of attention sharing must be established, turn-taking behavior must be initiated and maintained, each partner must tune into the dynamically evolving emotional and motivational state of the other, and partners must maintain an episodic memory which includes quite detailed models of others and how they might behave in certain circumstances. Although it is of course not feasible to build artificial systems that do all this in the same way as human beings, it is nevertheless possible to sufficiently simulate these various behaviors so that human beings interacting with the robot enter into the game.

Several researchers have recently produced very impressive demonstrations in this respect. One of the best known examples is the robot KISMET, built by Cynthia Breazeal and colleagues [3]. Note that although is possible to devise programs that recognise various emotional states from the speech signal or from facial expressions, and programs that animate an artificial face and modulate synthesised speech to reflect certain emotional states [28], this does not imply that these robots have such emotions in the same way as humans.

A third important issue concerns the grounding of meanings in the world. The robots must obviously be able to communicate about their environment as perceived and interpreted through their sensors. They must therefore go beyond pure symbol manipulation as in more traditional natural language processing applications. Two basic research issues need to be resolved. First, there is the problem of grounding itself, i.e. how concepts can be anchored through sensori-motor categories [13]. Second, there is the equally difficult issue how agents may conceptualise the world prior to communication. Here the insights of research in 'cognitive semantics' [24],[39] appears to be particularly pertinent, even though adequate formalisations and computational mechanisms for embodied cognitive grammars have been lacking so far (see [2] for recent attempts in this direction).

Embodiment, situatedness and social interaction allow the exploration of social and cultural learning as a new paradigm for shaping the cognitive capacities of artificial systems [40]. Social learning relies crucially on guessing what the other partner in the communication might possibly mean or already know as background context. It requires constant pragmatic feedback based on the tight interaction between speaker and hearer. Social learning allows potentially for much more effective learning, and may be the main explanation why children's word and grammar learning goes so fast. With social learning, as opposed to passive induction based on examples, the learner can make much better guesses about useful generalisations and gets feedback much more quickly on whether a generalisation goes in the right direction.

3 Evolutionary Language Games

Game theory has proven its usefulness in evolutionary biology and economics and has played a profound role in many artificial life experiments [25]. It is therefore not surprising that it has been widely adopted as a framework for studying the origins and evolution of lexicons and grammars in populations of agents, not only for performing large-scale computer simulations [38] but also for developing mathematical theories of language evolution [26].

A language game model consists of a population of agents. Each agent comes equiped with a cognitive apparatus, relevant to the aspects of language and meaning one wants to study. For example, this apparatus might feature components for parsing, producing, and learning a grammar, which would be necessary to study the emergence of grammar [?]. Or it might consist of an articulatory system, simulating the human vocal apparatus, an auditory perception component, and an associative memory relating speech percepts to vocal motor control programs. This is necessary for studying the emergence of human-like speech sounds in a population of agents [11].

The agents typically interact with the environment through a sensorimotor system, which can take on various degrees of complexity. It can be as complex as a fully autonomous humanoid robot or as straightforward as a pan-tilt camera supplying a stream of visual images. In the case of robotic implementations, the environment is the real world itself, possibly including different actors and objects dynamically interacting with each other. In computer simulations, the environment may be artificially constructed.

Agents have scripts for playing language games and they take turns playing the role of speaker and hearer so that they build up competence both for interpreting and for producing language utterances. Of particular importance are mechanisms for handling breakdowns in the game: learning the meaning of a word never heard before, inventing or adopting by analogy a new word for a meaning that was never expressed before, make a new distinction, and so on.

Once games are well understood they can be the basis of robotic implementations, even though the agent's scripts to play a game have to be made much more flexible and open-ended (see box 1 based on [37]).

4 Lexical games

Figure 2 from [38] shows the result of a typical robotic experiment, in which a population of thousands of agents has played a series of lexical guessing games. Lexical games are games in which two agents use individual words or groups of words without syntax to draw attention to an object in the environment. The objective of these experiments is to show how such a population would be able to self-organise a shared lexicon as well as the perceptually grounded categorisations of the world expressed by this evolving lexicon.

The environment of this particular experiment (known as the 'Talking Heads Experiment') consists of an open-ended set of geometric figures pasted on a white board (figure 1). One figure is chosen randomly by the speaker as topic of the language game and the hearer has to guess which topic was intended, based on words supplied by the speaker. The robots use their pan-tilt cameras for visual sensing and pointing. At the start of a game, the speaker moves his camera towards a specific area of the white board, thus indicating roughly to which area attention should be paid. The hearer moves the camera to indicate which object he grasped from the transmitted utterance and the speaker uses it to resolve a communication breakdown. Whenever the hearer points to the wrong topic or signals failure in understanding, the hearer points to the topic he originally intended to indicate. Sharing attention and pragmatic feedback through non-verbal communication is absolutely necessary for bootstrapping a shared lexicon and categorial repertoire.

Two processes are required to play lexical language games. First the speaking agent must conceptualise the context in such a way that he finds

Figure 1: The 'Talking Heads' experiment featured pan-tilt cameras oriented towards a white board on which geometric figures were pasted. Agents used these cameras to play lexical language games drawing attention to a chosen topic.

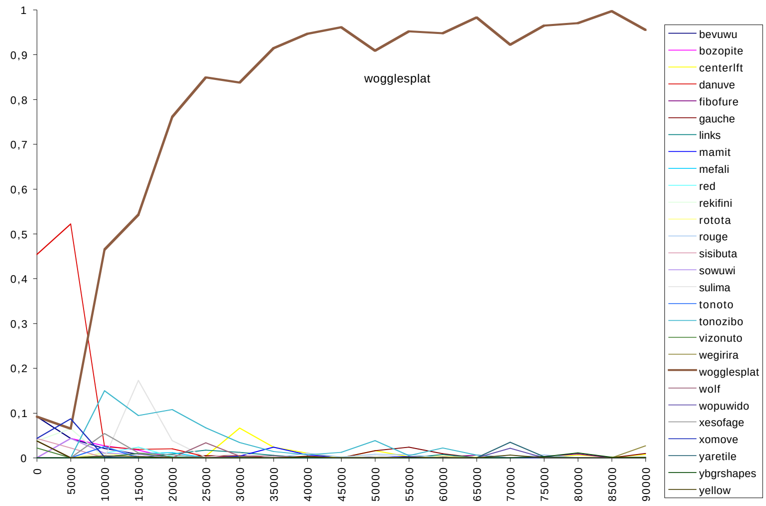

Figure 2: The x-axis displays the number of games and the y-axis the frequency with which a particular word is used for the same meaning. There is a winner-take-all effect due to the positive feedback between use and success.

a category or set of categories which distinguishes the topic from the other objects in the context. Thus if all objects have the same colour, it does not make sense to use colour distinctions. If the topic saliently stands out in size, then size is a good domain to use. If no distinction exists yet in this domain, or the distinction is too broad, a new more refined distinction can be introduced. It does not seem to make much difference whether a prototype-based categorisation [42], a neural network style radial basis function network [1], or some symbolic decision tree algorithm is used [30], as long as categories can be created and refined in an accumulative fashion.

Second, the selected category must be associated with a word. This word must be chosen for maximum communicative success. Hence agents should keep a score of their success for each form-meaning pair in the lexicon and choose the one with the maximum score. If a word-meaning pair fails in the game, its score should be decreased. If it succeeds, the score should increase and that of competing word forms, i.e. words associated with the same meaning, should decrease. It has been shown, both in various computer simulations and in mathematical proofs, that such a dynamics progressively results in the acquisition of the lexicon in a group (see for example [27]). Moreover, by tightly coupling the processes that make and maintain categorial distinctions and the processes that lexicalise them, the conceptual repertoires of the agents become coordinated without prior innate determination, nor central control, nor telepathy.

Figure 2 displays the evolution in the frequency with words are used to express a particular meaning in the Talking Heads experiment. We see that there is first a struggle in which different words compete until the population settles on a single dominant word. So there is indeed the self-organisation of a shared lexicon. The winner-take-all effect is due to positive feedback loop between use and success. The more agents prefer a particular word, the more they will use this word and the more success the word will have. This will further increase its score and progressively all agents will adopt it. This type of self-organisation has been widely studied in artificial life research and illustrates nicely how fundamental principles from biology are carried over to understand aspects of language.

Figure 3, taken from an experiment focusing on the evolution of a shared lexicon for colour categories [35], shows how the coupling of concept formation and language results in a shared categorial repertoire. In a first experiment, the agents each individually developed categories for successfully discriminating the topic from other objects in the context. In this case the similarity of their categorial repertoires was purely dependent on the structure of the environment. In a second experiment, the agents not only developed categories but also used them in a lexical guessing game. Whether a category continued to remain in an agent's repertoire not only depends on success in discrimination but also in communication. Due to this structural coupling the categorial repertoires of the agents become coordinated.

5 The Origins of Grammar

Evolutionary linguistics can contribute to the discussion on the origins of grammar because it offers tools for formalisation and a systematic experimental methodology for testing any kind of theory one wishes to examine. The subject of the origins of grammar disappeared from the research agenda when interest in historical linguistics waned in the beginning of the 20th century, but has recently become popular again, in line with a general scientific interest in questions of origins in biology and natural sciences like cosmology. The renewed interest is illustrated by a series of exciting conferences (starting with [18]), and widely discussed articles [19] and books [7].

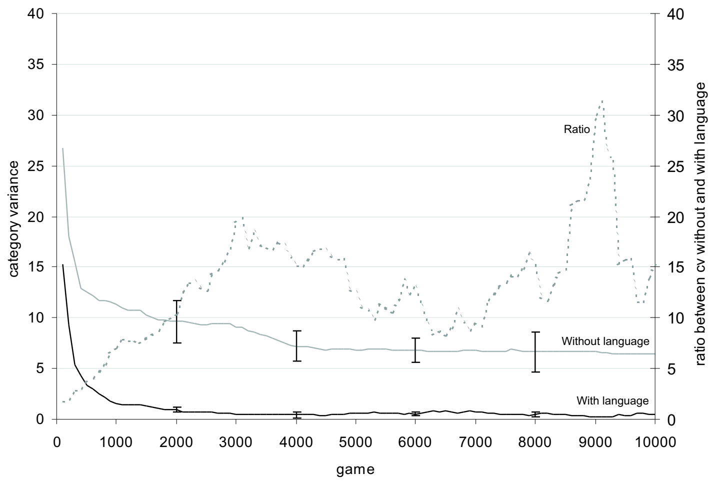

Figure 3: This figure compares the category variance (y-axis) between autonomously evolving agents with and without playing language games (x-axis). The ratio between the two clearly demonstrates how categorial similarity drastically increases when category formation is coupled to language.

Linguistics itself is divided on the question whether grammatical languages evolved primarily by some key genetic mutations [7] or by a process of cultural evolution in which existing cognitive capabilities (like categorisation, associative memory, analogy making, etc.) were recruited to support an increasingly more complex system for communication and representation. We find the same division in the computational and robotic models of evolutionary linguistics.

Some researchers [26] have examined the hypothesis that human language is essentially due to genetic mutations that gave rise to a Universal Grammar, which was then transmitted vertically from one generation to the next. Others [4] have explored the idea that genetic assimilation has played a key role, in other words that the brain progressively became optimised to engage in culturally evolved grammatical behaviors. Still others [34], [20] try to show that grammar is the result of a cultural evolution. Some of this work relies heavily on recent work in historical linguistics on how new grammatical constructions come up in a language and erode again due to grammaticalisation processes [41]. Although hypotheses have been formulated about what cognitive processes could drive these empirically observed grammaticalisation phenomena [14], no concrete operational models exist yet - partly because this requires detailed models of linguistic agents, with the mindboggling integration of many competences, including vision, language processing, and learning, and models of the collective behavior and population dynamics of linguistic communities. The advances in robotics and the tools and methods of evolutionary linguistics are beginning to make it possible to simulate these grammaticalisation processes and thus develop very concrete models on how grammatical subsystems, such as a tense-aspect-modality system, a grammar for case, a system of determiners, etc. can arise (see box 2 based on [34]).

6 Artificial Languages

The large-scale simulation of evolving languages and conceptual repertoires introduces several novel elements into linguistic modeling. First of all it assumes populations of speakers, whose members do not have exactly the same linguistic knowledge. 'The' grammar or 'the' lexicon does not exist as such, but can only be characterised by observing the most common usage of words and constructions at a particular point in time (as in figure 2). This contrasts with more traditional generativist theories which model a single, idealised speaker, who is assumed to know and speak the language perfectly.

Second, the language game ensures that all aspects of a verbal interaction are taken into account, including pragmatics (environmental context and discourse), and semantics (including conceptualisation prior to verbalisation). This contrasts with models which only look at syntax and assume that there is no profound interaction between choosing what to say and choosing how to say it.

Third, an individual's knowledge of the language is not considered to be static, but in constant evolution. Every verbal interaction changes the linguistic knowledge state of an agent, particularly if the communication did not succeed. Evolution of the language as a whole is crucially dependent on natural variation in the linguistic behavior of the population and some of these variants are culturally selected and become dominant.

Many empirically inclined linguists are initially upset that these simulations produce artificial languages which exhibit perhaps certain similarities to phenomena also observed in natural languages but are not at all equal to existing human natural languages. Indeed this will never be possible. Human natural languages are the evolutionary consequence of an enormous number of factors, most of which cannot be modelled. Instead, evolutionary linguistics should be seen as a new approach to theoretical linguistics. It is similar to theoretical biology, which attempts to discover the general principles underlying the evolution of biological structures and functions without ever being able to predict how a species is going to evolve further. However it can formalise general principles and show what consequences they have.

7 Conclusions

Evolutionary linguistics has recently emerged as a new subfield, studying the origins and evolution of language through computational and robotic simulations. The field draws on ideas, tools, and techniques both from Artificial Life research and from Computational Linguistics and Artificial Intelligence, but is at the same time strongly rooted in the study of human natural language. Research is necessarily slow given the enormous complexity of setting up interesting experiments, particularly if they are carried out with autonomous robots. Nevertheless, solid results are beginning to emerge contributing to our understanding of the origins of human speech systems, lexicons and grammars. These results have typically been achieved by adopting the view that language is a complex adaptive system, in which shared systematicity emerges and continues to evolve dynamically by local interactions between language users.

Although evolutionary linguistics is primarily intended to contribute to our theoretical understanding of language - and hence complementary to empirical work - there is an important spin-off, namely the design and implementation of communication systems for robots. These systems would not be based on prior design, nor on the statistical induction of language regularities from a large set of examples, but on an active participation in the construction and negotiation of adequate communication conventions with natural language like properties.

8 Questions for further investigation

Evolutionary linguistics has clearly opened up a new avenue for the study of language, but most of the work remains to be done. A large set of very deep open problems has been generated, which were not on the research agenda before. Some of them can perhaps be handled by further elaboration or application of techniques that have already been discovered. But most of them are more fundamental and will require major conceptual advances as well as jumps forward in technology.

1. Clearly appropriate language games are crucial for bootstrapping language competence but how do the games themselves emerge? How can a pattern of behavior get established and become conventionalised? Solving this problem will require looking beyond the individual, which so far tends to be the main focus of cognitive science in general and cognitive neuroscience in particular, and incorporating many more insights from the social sciences.

2. Research in cognitive semantics has yielded intriguing descriptions of the kinds of conceptualisations humans employ for language. But much remains to be done to explain for specific semantic domains (for time and space, event-structure, determination, etc.) how such conceptualisations can arise and how do they get shared across human beings.

3. Operational models for the origins of natural language like grammars are almost non-existent. No convincing simulations of grammaticalisation phenomena, which are so common in the formation and emergence of human languages, have been shown yet. One of the major difficulties is that grammaticalisation processes are strongly grounded in many aspects of human culture and embodiment and these aspects are very difficult to incorporate in operational models.

4. Although intriguing simulations (not reported in this paper) could be shown for the origins of individual speech sounds [11], major work remains to be done to explain phonemic coding and other universal tendencies in the speech sounds of the world.

References

[1] Belpaeme, T. (2001) Simulating the formation of colour categories. In: Nebel, B. eds. (2001) Proceedings of IJCAI-01. Morgan Kauffman Pub. San Francisco. p. 393-398.

[2] Bergen, B. and N. Chang (2002) Embodied Construction Grammar. TR-02-004. International Computer Science Institute, Berkeley Ca.

[3] Breazeal, C. and B. Scassellati (2002), Robots that imitate humans in Trends in Cognitive Science, vol. 6, pp. 481-487.

[4] Briscoe, T. (2002a) Grammatical Acquisition and Linguistic Selection. In: Briscoe, T. (ed.) (2002) Linguistic Evolution through Language Acquisition: Formal and Computational Models. Cambridge University Press, 2002

[5] Briscoe, T. (ed.) (2002b) Linguistic Evolution through Language Acquisition: Formal and Computational Models. Cambridge University Press, 2002

[6] Broeder, P. and J. Murre (2000) Models of Language Acquisition. Inductive and Deductive Approaches. Oxford University Press, Oxford.

[7] Calvin, W. and D. Bickerton (2001) Lingua Ex Machina. MIT Press, Cambridge Ma.

[8] Cangelosi, A. and D. Parisi (eds.) (2001) Simulating the Evolution of Language. Springer Verlag, London, 2001.

[9] Clark, H. H., and Brennan, S. A. (1991). Grounding in communication. In: Resnick, L.B., J.M. Levine, and S.D. Teasley (Eds.). Perspectives on socially shared cognition. Washington: APA Books.

[10] Dautenhahn, K., A. Bond, L. Canamero, B. Edmonds (eds.): Socially Intelligent Agents - Creating Relationships with Computers and Robots, Kluwer Academic Publishers.

[11] Bart de Boer (2000) Self organization in vowel systems. Journal of Phonetics 28 (4), pp. 441465

[12] Deneubourg, J-L., et.al. (1993) Self-organisation and life: from simple rules to global complexity. Proceddings of the Second European Conference on Artificial Life, ULB. Brussels.

[13] Harnad, S. (1990) The Symbol Grounding Problem. Physica D 42: 335-346.

[14] Heine, (1997) The cognitive foundations of grammar. Oxford University Press, Oxford.

[15] Holyoak, B. and K. Kokinov (Eds.), The analogical mind: Perspectives from cognitive science . Cambridge, MA: MIT Press.

[16] Hotani, H., R. Lahoz-Beltra, B. Combs, S. Hameroff, and S. Rasmussen. Liposomes, Microtubules, and Artificial Cells. Nanobiology, Vol 1, 1, 1992, p. 61-74.

[17] Hopper, (1991) Emergent Grammar. In: Traugott, E. and Heine, B. (1991) Approaches to Grammaticalization. Volume I and II. John Benjamins Publishing Company, Amsterdam, 1991.

[18] Hurford, J., C. Knight and M. Studdert-Kennedy (1998) Approaches to the Evolution of Language: Social and Cognitive bases. Cambridge University Press, Cambridge.

[19] Jackendoff, R. (1999). Possible stages in the evolution of the language capacity. Trends in Cognitive Sciences, 3(7):272–279.

[20] Kirby, S. (2002a) Learning, Bottlenecks and the Evolution of Recursive Syntax. In: Briscoe, T. (ed.) (2002b) Linguistic Evolution through Language Acquisition: Formal and Computational Models. Cambridge University Press, 2002. Chapter 6.

[21] Kirby, S. (2002b) Natural Language from Artificial Life. Artificial Life Journal, 8(2) p. 185–215.

[22] Koza, John R. 1992. Genetic Programming: On the Programming of Computers by Means of Natural Selection. Cambridge, MA: The MIT Press.

[23] 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.

[24] Langacker, R. (1991) Foundations of Cognitive Grammar. 2 vols. Stanford University Press, Stanford.

[25] Langton, C. (ed.) (1995) Artificial Life. An Overview. MIT Press, Cambridge Ma.

[26] Nowak, M. A., J.B. Plotkin and V.A. Jansen (2000) The evolution of syntactic communication. Nature 404 495-498.

[27] Oliphant, M. and J. Batali (1997) Learning and the emergence of coordinated communication. The newsletter of the Center for Research in Language. Vol 11, 1, 1997.

[28] Oudeyer, P-Y. (2002) The production and recognition of emotions in speech: features and algorithms. International Journal of Human Computer Interaction, 2002.

[29] Scassellati, B. (1998) Imitation and Mechanisms of Joint Attention: A Developmental Structure for Building Social Skills on a Humanoid Robot, in C. Nehaniv, ed., Computation for Metaphors, Analogy and Agents, Vol. 1562 of Springer Lecture Notes in Artificial Intelligence, Springer-Verlag, 1998

[30] Steels, L. (1996b) Perceptually grounded meaning creation. In: Tokoro, M. (ed.) (1996) *Proceedings of the International Conference on Multi-Agent Systems.* AAAI Press, Menlo Park Ca. pp. 338-344.

[31] Steels, L. (1997) The Synthetic Modeling of Language Origins. Evolution of Communication Journal 1(1), 1-35. (1997)

[32] Steels, L. (2001) Language Games for Autonomous Robots. IEEE Intelligent systems, September/October 2001, p. 16-22.

[33] 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.

[34] Steels, L. (2002) Simulating the evolution of a grammar for case. In: Proceedings of the Fourth Evolution of Language Conference, Harvard. (to appear)

[35] Steels, L. and T. Belpaeme (2003) Computer simulations of color categorisation and color naming. Behavioral and Brain Sciences (to appear).

[36] Steels, L. and R.A. Brooks (1995) (eds.) The Artificial Life Route to Artificial Intelligence: Building Embodied Situated Agents. Lawrence Erlbaum Associates, Inc., Hillsdale, NJ, 1995.

[37] Steels, L. and F. Kaplan (2001) AIBO's first words. The social learning of language and meaning. Evolution of Communication 4(1).

[38] 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.

[39] Talmy, L. (2000) Toward a Cognitive Semantics: Concept Structuring Systems (Language, Speech, and Communication) The MIT Press, Cambridge Ma.

[40] Tomasello, (1999) The Cultural Origins of Human Cognition. Harvard University Press, Cambridge Ma.

[41] Traugott, E. and Heine, B. (1991) Approaches to Grammaticalization. Volume I and II. John Benjamins Publishing Company, Amsterdam.

[42] Vogt, P. (2001) Bootstrapping grounded symbols by minimal autonomous robots. Evolution of Communication 4(1).

[43] Wray, A. (ed) (2002) The Transition to Language. Oxford University Press, Oxford, 2002.