# Blue Paper: A Research Roadmap for Developing Artificial Embodied and Communicating Agents

**Stefano Nolfi[1], Dario Floreano[2], José Halloy[3], Lars Holmquist[4], Vittorio Loreto[5], Eors Szathmáry[6], Ricard V. Solé[7], Luc Steels[8]**

[1]National Research Council, Institute of Cognitive Sciences and Technologies
[2]The Swiss Federal Institute of Technology Lausanne
[3]Université Libre de Bruxelles
[5]Physics Department, "La Sapienza" University
[6]Future Applications Lab, University of Goteborg
[6]Institute for Advanced Studies, Collegium Budapest
[7]University Pompeu Fabra
[8]Sony Computer Science Laboratory

## 1. Introduction

Understanding the evolution of communication and human language is one of the hardest problems in science and poses extremely hard challenges (Christiansen and Kirby, 2003). Indeed, human language is one of the most complex behaviours that we know, resulting from the dynamical interactions between a population of individuals dynamically interacting with their environment, and it continuously changes over time by adapting to changes in the environment and/or in the individuals. More generally, one of the main difficulties for understanding the evolution of communication and language derives from the need to address highly interdisciplinary issues such as how animal and human communication systems are structured and how they differ, how communication and language affect and are affected by the behavioural and cognitive capabilities of individuals, how humans acquire language during development, how evolution learning and culture interact, etc.

In this paper, we will approach this issue from an artificial perspective, that is, we will discuss how populations of artificial agents that interact autonomously (i.e. without human intervention) with the physical world and amongst themselves, can develop forms of communication of different complexity analogous to animal and human communication. Answering this question is important for both scientific and technological reasons. From a scientific point of view, understanding the conditions that might lead to the emergence of communication in a population of interacting embodied agents might shed lights on the evolution of animal communication and on the origin of human language. From a technological point of view, understanding how communication and language develop and evolve is probably a necessary prerequisite for developing technological artefacts that display the same flexibility and robustness of natural organisms. More specifically, in this paper we will provide a research roadmap for the emerging field that studies how embodied and communicating agents (ECAgents) can be developed and how they can self-organize a communication system that emerges as the result of the interaction between the agents and the agents and the environment (Steels, 1999, 2003; Cangelosi & Parisi, 2002; Kirby, 2002; Wagner et al., 2003; Nolfi 2005).

This paper is organized as follows. In the next section we will point out the complex adaptive nature of communication in embodied agents and the importance of relying on complex systems and self-organizing methods for developing embodied and communicating agents (ECAgents). In section 3, we will discuss collective intelligence methods that study of how coordinated behaviour might emerge in distributed systems in which individual agents only have access to local sensory or incomplete information. In section 4, we will discuss evolutionary methods in which the control system of the individual agents is evolved through an artificial evolutionary process analogous to natural evolution. In section 5, we will discuss social learning methods in which a ECAgents modify their internal structure and their interaction rules as a result of the outcome of their interactions. For each method, the basic idea, an exemplificative case, and the most promising research directions will be presented. In section 6, we will discuss the combination of evolutionary and social learning methods. In section 7, we will discuss how recent development from statistical physics and complex network theory might lead to the identification of universal properties of natural communication systems that might provide important insights on how ECAgents could be developed. Finally, in section 8 we will discuss the potential applications of these researches.

## 2. Communication and language as complex adaptive systems

Studying communication and language in embodied agents implies dealing with a complex adaptive system that involves several levels of organization (Keijzer, 2001; Nolfi, in press). This means that self-organizing principles and complex systems tools and methods might play a crucial role for the development of ECAgents (Steels, 2000). Moreover, this implies that the emergence of communication in ECAgents represents a great challenge and an ideal case study for complex system theory. In the next sections we will point out the adaptive complex system nature of individual and collective behaviour and of cultural processes (such as language) that arise from individual and collective behaviour. More specifically, we will show the importance of relying on self-organizing methods for developing ECAgents

### 2.1 Individual Behaviour

Behaviour is an emergent property resulting from the often non-linear interactions between an agent (natural or artificial), its body, and the external environment (including the social environment).

At any time step, the environmental structure and the agent/environmental relation co-determine the body and motor reaction of the agent that, in turn, co-determine how the environment and/or the agent/environmental relation vary. Sequences of these fine-grained interactions, occurring at a fast time rate, lead to an emergent property – behaviour – that extends over a significant larger time span than the interactions. Since interactions between the agent's control system, its body, and the external environment are non-linear (i.e. small variations in sensory states might lead to significantly different body and motor actions) and dynamical (i.e. small variations in the action performed at time $t$ might significantly impact later interactions at time $t_{+x}$), the relation between the rules that govern the interaction and the behaviour resulting from the interactions is very indirect. As a consequence, behavioural properties can hardly be predicted or inferred from an external observer even on the basis of a complete knowledge of the interacting elements and of the rules governing the interactions.

For the sake of completeness, beside extremely simple cases in which agents are reactive (i.e. their motor reaction is only determined by their current sensory state), individual behaviour results from the coupling between two dynamical processes: (1) the dynamical process occurring within the agent's control system and body, and (2) the dynamical process resulting from the interaction between the agent and the environment.

Moreover, behaviour is a multiply scaled phenomenon with different levels of organization and involves features occurring at different time scales (Nolfi, in press). Indeed, agent/environmental interactions occurring at a rate of milliseconds might lead to simple elementary behaviours that extend over a short time span (e.g. obstacle avoidance behaviours extending over hundreds milliseconds). Interaction between simple elementary behaviours (e.g. obstacle avoidance and target approaching) might lead to more complex behaviours that extend over longer time spans (e.g. navigation behaviours extending over seconds or minutes). Interactions occurring at a lower level of organization and extending over short time spans thus give rise to behavioural properties at higher levels of organization. In addition, processes occurring at higher levels of organization and extending over significant time spans might affect the interactions occurring at lower levels of organization (Keijzer, 2001; Nolfi, in press).

The overall picture thus is that of a multiple scaled phenomenon involving bottom-up and top-down relations between emergent properties occurring at different levels of organizations and occurring at different time rates.

## 2.2 Collective Behaviour

Collective behaviour is an emergent property resulting not only from the interactions between agents, their bodies, and the external environment but also among agents. In the case of collective behaviour, therefore, the fine-grained interactions occurring between the control system, the body, and the environment of each individual agent is only one of the factors that give rise to the individual behaviour of each agent. Agents in fact also interact between themselves, directly and indirectly (i.e. through physical interactions and through the environment).

As for individual behaviour, social behaviour is a multiply scaled phenomenon involving bottom-up and top-down relations between properties occurring at different time scales that emerge from the interaction between lower level properties. The process that leads to collective behaviour, however, is richer than the one leading to individual behaviour since it involves a much larger number of concurrent interactions. Indeed, in a social context, an external observer might typically identify individual behaviours, resulting from the interaction between the control system of an agent, its body, and the environment, and social behaviours, resulting from the interaction between agents and agent's individual behaviours. Both individual and social behaviour might involve different levels of organization that extend at different time scales. Moreover, in the case of collective behaviour, the layered structure might include complex high-level properties, such as culture or language, that change at an extremely slow time rate and extend over huge time spans.

## 2.3 Language and communication

Language and communication systems are high level behavioural properties that vary at a short time rate, extend over a long time span, and result from a large number of hierarchically organized and mutually interacting behavioural processes. These behaviours, that occur at lower levels of organizations and extend over shorter time spans, might include: communication interactions between individuals (e.g. dialogue behaviour in humans or dance behaviours in bees), collective behaviours (e.g. cooperative behaviours or shared attention behaviours), individuals behaviour (e.g. locomotion).

Although language and communication are properties of a population of interacting agents that emerge from the interaction between lower level properties, they can be conceptualised as independent entities that self-organize on the basis of processes involving their constituting elements and that constraint the lower level behavioural processes from which they result. An example of self-organizing processes occurring at the level of the communication system is the competitive process between alternative words that express the same meaning (Steels, 1999). In fact, although the permanence or the disappearance of synonymous words in the communication system ultimately depends on the characteristics of individual agents and on the effect of single agent/agent communicative interactions, the destiny of synonyms can be predicted on the basis of their relative frequency of use (i.e. on the basis of a property of the communication system). In other words, communication systems might self-organize and evolve according to dynamics that can be predicted (up to a certain extent) on the basis of their global organization.

As an example of how an existing communication system might constraint lower level behavioural processes is represented by the constraints that the communication system might impose on individual learning. During language acquisition, children do not only acquire a mechanisms to express meanings but also a structured way to categorize their environment that constrains their individual learning process beside language acquisition (Waxman & Markow, 1995; Nazzi & Gopnik, 2001).

### 2.4 Communication and language as adaptive systems

One fundamental characteristic of language and communication in natural organisms consists in their adaptive nature, i.e. in their ability to vary so to adapt to variation of the lower level properties (variation of the environmental, individual and social conditions). Indeed, communication and language systems vary as the result of three adaptive processes (genetic evolution, individual and social learning, and cultural evolution).

The adaptive nature of language and communication very likely plays a crucial role both from the point of view of the evolution of communication (i.e. for the emergence of communication systems from scratch or for the emergence of complex from simpler communication forms) and from the point of view of its functionality. One universal feature of human language is indeed the ability to convey a potentially infinite set of meanings and the ability to generate new words to express or indicate new entities that did not exist before.

From the point of view of building artificial systems, the attempt to develop ECAgents might be pursued by following two alternative methods: (a) a self-organization method in which some aspects of the agents are designed by the experimenter, but other important aspects are left unspecified and are developed through self-organizing methods, and (b) design methods in which the experimenter identifies and implements the characteristics that ECAgents should have in order to effectively communicate and to modify their communication system as a result of environmental, social, and individual variations.

In this paper we will especially focus on the former method since we believe that adaptation and self-organization are crucial pre-requisites for the possibility to develop artificial agents with the robustness, adaptability, and complexity of natural agents. Moreover, we believe that the complex adaptive system nature of behaviour and communication in embodied agents prevent the use of techniques based on explicit design in which the experimenter has to identify the rules that determine the interaction between each individual agents, its body, and the social and physical environment that lead to the desired emergent properties. The fact that the relations between the rules that determine the fine-grained interaction between the elements and the emergent behaviour resulting from the elements and their interactions are so complex and indirect, as we argued above, makes it impossible to infer the correct rules even on the basis of a clear and complete description of the problem (Funes et al., 2003; Nolfi, in press).

## 3 Collective Intelligence

Collective intelligence (or Swarm Intelligence) is a recent field that studies systems composed of a collection of agents that cooperate to achieve a common goal (Beckers *et al.* 1994; Bonabeau et al. 1999; Camazine et al., 2001; Dorigo & Stützle, 2004). In these systems, the problems are collectively self-solved in real time through the behaviour of the agents, which interact with each other and with the environment. Collective intelligence emphasizes aspects such as decentralization, local communication among agents, emergence of complex global behaviour from the interaction between agents exhibiting simple individual behaviours, and robustness.

For example, eusocial insects, ants, bees, wasps, termites, form colonies presenting a high social organization. Often such societies are described as “super-organisms” because of their high level of integration and division of labour including the reproductive task. The collective capabilities of these societies are impressive, they include sophisticated nest building, defence, food retrieval including some “agricultural” activities like “gardens” or “herd breeding” of other insects. Understanding the evolutionary emergence and the mechanisms underlying such sophisticated societies are still a scientific challenge (Wilson & Hölldobler, 2005). Obviously, the mechanisms regulating such integrated cooperation among individuals are genetically determined and genetics studies contribute to explain the evolutionary emergence of eusociality. Moreover, many experimental studies demonstrate that collective activities in such societies can also be based on self-organized behavioural mechanisms (Deneubourg *et al.*, 1987). These epigenetic mechanisms contribute to highly enhance the flexibility and adaptability of such societies (Camazine et al., 2001). Such behaviours are an example of functional self-organization that produces efficient or even optimal collective solution **(**Deneubourg *et al.*, 1983). Combined they provide the colony with a collective intelligence that reinforce the “super-organism” analogy as no individual posses a global view of the colony.

Self-organized behaviour represents collective behaviour that emerges at the level of the group from the numerous interactions within individuals and between the individuals and the environment. Moreover, the rules specifying the interactions within individuals and between the individuals and the environment are executed by using only local or incomplete information, without reference to the global pattern (i.e. without relying on global maps or global representations) (Detrain *et al.*, 1999; Camazine *et al.*, 2001).

The mechanisms that lead to self-organization in biological systems differ from those in physical systems in one important respect: the rules that determine the interactions within agents and within agents and the environment are influenced by genetically controlled properties. The properties emerging from the interactions thus are shaped by natural selection. By tuning the rules, selection thus has a way to shape the properties emerging from the interactions in order to achieve a given adaptive functionality.

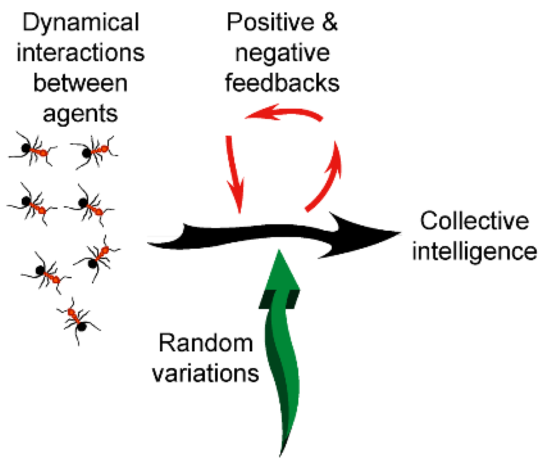


Figure 1. Core ingredients for collective intelligence.

According to current understanding, the core ingredients in collective intelligent systems are: (1) dynamical non-linear interactions between agents, (2) random variations, and (3) Positive and negative feedback mechanisms (Figure 1). In short, the dynamical interactions within agents and between the agents and the environment constitute the basis for the synthesis of emergent properties. Random fluctuations lead to innovations that might constitute potential effective solutions. Positive feedback mechanisms promote changes in the system by reinforcing the changes that that occur in the same direction of initial variations. Finally, negative feedback mechanisms keep the amplification process occurring as a result of positive feedback mechanisms under control. The combination and the interaction between these fundamental mechanisms might lead to the online selection of self-organized solutions. For example, the combination of these mechanisms might lead to a collective decision process in which an optimal or close to optimal decision is taken by the collection of agents despite no single agent has enough information to take the appropriate decision and despite agents are not 'aware' of the decision taken.

Collective intelligence methods are often used to model natural organisms. However, they can also be translated and applied in artificial systems in fields like collective robotics (Beckers *et al.* 1994; Dorigo et al., 2004), swarm intelligence and ant colony optimization (Bonabeau et al. 1999; Dorigo & Stützle, 2004).

Communication or interactions are essential to obtain cooperation between individuals. In insect societies semiochemical communication is very common. Pheromones are common chemicals used for communication within species (Wyatt, 2003). They are substances secreted by an individual and received by another individual of the same species in which they trigger a specific reaction or behaviour. In the next section we describe an exemplificative case that illustrates how such form of communication can be modelled in order to explain the epigenetic mechanisms regulating collective activities.

### 3.1 An exemplificative case: emergence of coherent behaviour and collective decisions on the basis of semiochemical communication.

Trail formation for food foraging in ants is a self-organizing property emerging from the interaction between the ants and between the ants and the environment mediated by chemical communication mechanisms (Detrain et al, 1999; Camazine *et al.* 2001). Here we will focus on how groups of ants belonging to different lines or strains can produce coordinated behaviour such us segregation or aggregation. Each group of ants has its own blend or pheromone trail and an ability to recognize its own pheromone and a partial ability to discriminate between its own pheromone and that of other groups.

To model ant behaviour and to investigate the possibility of developing artificial agents that display analogous capabilities a set of experiments in simulation have been conducted (Millor *et al.*, 2006). Two groups of agents are placed in an environment including 2 trails leading to two identical food sources equidistant from the nest (Figure 2). Agents' departure from the nest to the food source is based on the probability $\Phi$ ($S^{-1}$) to leave the nest per time unit (the flow of departure). The path chosen by the ant depend on the value and the type of pheromone concentration on each trail. Each group has its own blend of pheromone (e.g. a mix of trail pheromone and footprint hydrocarbons). Ants chose a trail on the basis of the relative concentration of pheromone on each trial and on the basis of the type of pheromone. All pheromone types attract the ants, however the type of pheromone that corresponds to the pheromone type produced by an ant has a higher attraction value with respect to other pheromone types. The parameter $\beta$ indicates the level of inter-attraction between pheromone types produced by other groups. After choosing a path, ants reach a food source, ingest food, and promptly return to the nest laying pheromone. Pheromones evaporate with a rate proportional to pheromone concentration and to a time constant (for details see, Millor *et al.*, 2006). These rules (i.e. the rules that determine the interaction within agents and between the agents and the environment) are based on empirical finding related to the behaviour of ants.

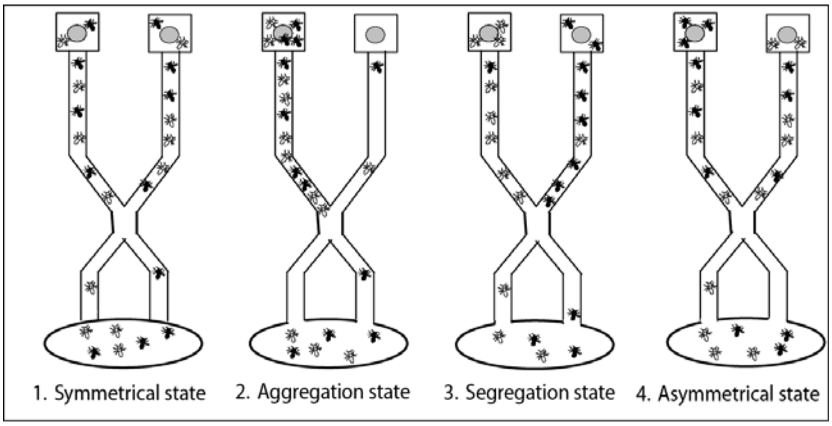


Figure 2. Two groups of agents placed in a environment with two trails leading to two identical and equidistant food sources. The grey and black individuals represent the ants of the two groups. The grey circles represent the food sources. The four picture represent schematically the four possible stationary states: (1) a symmetrical state in which the flow of both groups and the concentration of the two types of pheromones is equal on both trails, (2) an aggregation state in which both groups choose the same trail and the concentration of both types of pheromone is high in one trail and low on the other trail, (3) a segregation state in which the individuals of the two groups tend to choice different trails, (4) an asymmetrical state in which the two groups have a slight preference for two different trails but in which the level of segregation is low.

By varying the strength of inter-attraction (i.e. the probability to react differently to different pheromone types) $\beta$ and the flow of departure $\Phi$ ($S^{-1}$), one can observe that for low number of individuals (i.e. low flow of departure), the individuals distribute in a symmetrical way on the two trails, i.e. both branches are concomitantly exploited by both groups. As soon as the flow exceeds a threshold, a transition occurs, and the pattern switches to aggregative or segregative states depending on the $\beta$ value. For low value of inter-attractions (i.e. $\beta < 0.24$) between groups, the two groups might segregate or aggregate because the odour signal coming from the other strains plays only a little role on the amplification process in comparison with its own trail odour. For higher values of inter-attraction (i.e. $\beta > 0.24$) the two groups always aggregate.

The obtained results also show how a low level of discrimination between signals is sufficient to produce robust segregation or aggregation behaviour and different collective response occur without the need to change the individual behaviour or the communication system.

Notice how the collective behaviour displayed by the agents and its emergent properties (i.e. segregation, recruitment, collective decision) result from the following mechanisms:

(a) The probability to select the trail in which the concentration of pheromone (and in particular the concentration of the pheromone corresponding to agents own blend) is higher and the fact that, later on, pheromone is laid on the same trail. These mechanisms act as positive feedback mechanisms since small differences in pheromone concentration tend to be amplified as a result of the agent's behaviour.
(b) The evaporation of pheromone that acts as a negative feedback mechanism that limits the effect of the positive feedback mechanisms.
(c) Random fluctuations in agent's behaviour resulting from their probabilistic behaviour. In real ants, random fluctuations are due either to the ant's inability to always follow the trail and/or the noise in the chemical signal (resulting from the property of the ground) or to the ant's chemical detection capabilities.

### 3.2 A research road-map

Progresses in our understanding of how coherent collective behaviour arises in natural organisms and on the role of communication might significantly increase our ability to design a collection of artificial systems that exhibit cooperative and coherent behaviours. More specifically the identification of the mechanisms that lead to self-organization in natural organisms can allow us to derive a 'design for emergent' techniques (Pfeifer & Scheier, 1999), that is, techniques that might allow us to design artificial systems able to self-organize and exploit emergent properties (the role of evolutionary and learning rules in design and in the identification of design principles will be addressed in the next sections). In this section we identified some of these mechanisms (i.e. non-linear interaction rules, positive and negative feedback mechanisms, random variations). A deeper understanding of the role of these mechanisms and of their applicability to other contexts, and eventually the identification of other crucial mechanisms, might lead to a breakthrough in the field of collective intelligence in artificial systems.

For instance, the relation between the task complexity and the complexity of individuals and of their communication system is an open question. Similarly, what is the level or what are the levels at which artificial agents should be designed (hardware, micro-interaction rules, macro-interaction rules) and how one can relate formally and systematically different levels is an important open research issue.

A more specific research issue is the role that chemical communication (or other communication systems that share properties with chemical communication in insects) can have on the design or artificial systems. Chemical communication, in fact, has several specific characteristics that might be interesting from the point of view of developing artificial systems. In particular, chemical communication involves signals that: (1) last in time much longer (compared to sounds or light), (2) extend for different time spans (depending on the type of molecules involved), (3) by being released on a substrate are localized in space, (4) can be summed up to convey information related to signals produced by different agents, (5) allow more easily concurrent communication acts with respect to other types of communication media. Indeed, chemical communication shares more properties with artificial communication systems (e.g. multiplexing signals in electromagnetic radio communication) than with other natural communication media. Understanding the potentiality of chemical communication systems in artificial agents thus remains largely an open question. For an example of its application to establish coherent behaviour in groups including real insects and robots see (http://leurre.ulb.ac.be).

Another related promising research direction consists of progress in hardware technology for the detection and emission of chemical signals and/or of the development of new technologies that might allow the development of new communication media (e.g. based on radio or infrared communications) that include the main properties and advantages of chemical communication without necessarily involving a chemical medium.

Below we will discuss more specific research directions in line with these general objectives.

The exoskeleton of insects or cuticule is covered by lipids and hydrocarbon chains. These molecules have different roles like preventing desiccation or protecting from alien products or organisms. Preliminary evidence suggests that hydrocarbons on cuticule and postpharyngeal glands play communicative functions. For instance they are essential as a semiochemical recognition factor (Lenoir *et al.*, 1999). By semiochemical we mean chemicals that have a 'meaning' (semio) because they trigger specific behaviours when perceived by the receiver. In ant species, for example, they play a role in nest-mate recognition, i.e. they allow individuals of the same species to discriminate between individuals of the same or of other colonies. This discrimination capability, in turn, might allow individuals to display cooperative or aggressive behaviour with respect to individuals of the same or of a different colony, respectively. The acceptance or rejection of an individual might depend on the degree of overlapping between its own chemical signature and the chemical signature of the encountered individual.

One first important research direction concerns understanding the distribution of different chemicals at the level of the individual, at the level of the colony, and at the level of the species.

By measuring the relative abundance (frequency) of cuticular hydrocarbons in several species of ants and their occurrence rank in a colony, one can observe (Millor *et al.*, in prep.) that the relationship between the two quantities, expressed in a $\log_{10}$ scale, is linear (i.e. it adheres to Zipf law distribution in which the overall abundance of an hydrocarbon in a colony is inversely proportional to its relative frequency in the species). Since Zipf law distribution is common to many self-organizing systems and to communication systems (see also section 7), understanding the reasons, the functionality (if any), and the implication of this distribution might shed lights on chemical communication in insects and might provide insights on how effective and scalable artificial communication systems can be designed. Moreover, new studies are needed to verify whether these power law distributions characterize chemical distributions in other social or non-social insects.

One second promising research direction concerns understanding how chemical signals vary ontogenetically (i.e. during the lifetime of a colony) and evolutionarily. The mixture of molecules that determine the odour blend of an individual, in fact, is influenced by genetic, environmental and social factors. Environmental influences arise as a result of molecules collected from the environment. Social influences arise as a result of molecules exchanged with other individuals during physical contact. As a consequence, chemical signals might convey a rather rich amount of information related to the genetic characteristics of an individual and of his previous environmental and social experiences mediated by the individual's behaviour. Indeed, it has been experimentally shown that individual ants isolated from their colony for a few weeks loose the odour that identifies their colony and are then attacked by their former nest mates that perceive them as aliens (Lenoir *et al.*, 1999). For a first attempt to model ontogenetic variations of chemical signals see (Deneubourg *et al.* in prep.).

## 4. Evolutionary Methods

Evolutionary methods consist of the application of an artificial analogue of natural Darwinian evolution to the design of ECAgents (Cangelosi & Parisi, 2002; Kirby, 2002; Wagner et al., 2003; Nolfi, 2005).

The basic idea goes as follows (Figure 3). An initial population of different artificial genotypes, each encoding the control system (and eventually the morphology) of an agent, are created randomly. Each genotype is translated into a corresponding phenotype (i.e. into a corresponding agent) that is allowed to "live" (i.e. to move and interact with the external environment and with other agents) while its performance (*fitness*) with respect to a given task is automatically evaluated. Agents are placed in the environment and evaluated in groups that might be heterogeneous (i.e. might consist of agents with different characteristics corresponding to different genotypes) or homogeneous (i.e. might consist of identical agents obtained by copies of the same individual genotype). Next a new population is generated by allowing the genotype of the fittest agents to reproduce by generating copies of themselves with the addition of changes introduced by some genetic operators (e.g., mutations, crossover, duplication). This process is repeated for a number of generations until a team is born that satisfies the performance criterion (*fitness function*) set by the experimenter.

In this section we will restrict our analysis to cases in which the control system of the agents do not change during lifetime (i.e. agents that adapt phylogenetically only). Methods combining phylogenetic and ontogenetic adaptation will be discussed in Section 6.

To develop teams of agents able to solve a given task the experimenter has to determine:

1) A fitness function, i.e. a criterion for automatically evaluating the performance of a team of agents with respect to a given task.

2) The parameters that are fixed and are not subjected to the evolutionary process. These typically include: the parameters of the artificial evolutionary process (i.e. the size of the population, the mutation rate etc.); the sensory-motor system and body structure of the agents, and, in some cases, the architecture of the agents' controllers that typically consist of artificial neural networks.

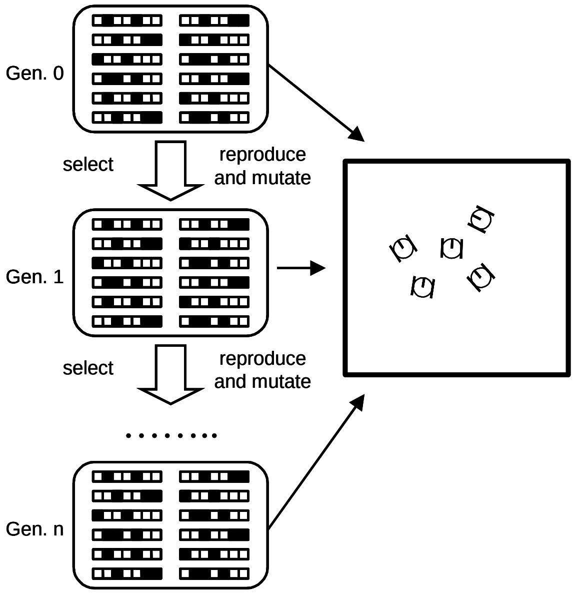


Figure 3. A schematic representation of an artificial evolutionary process. The stripes with black and white squares represent individual genotypes. The genotypes included in rectangular boxes indicate the genome of the population for different generations. The small wheeled robots placed inside a square represent a group of agents allowed to move and interact between themselves and with the external environment.

## 4.1 An exemplificative case: evolving ECAgents able to solve a collective navigation problem

Consider the case of four mobile robots placed in an arena that includes two target areas. The robots have to find and remain in the target area by equally dividing between the two (Marocco & Nolfi, 2004, 2005). Robots have a circular body and are provided with a simple sensory-motor systems and a neural controller (Figure 4) that allow them to move, produce a signal with an intensity varying in the range [0.0, 1.0], and gather information from their physical and social environment (including signals produced by other agents).

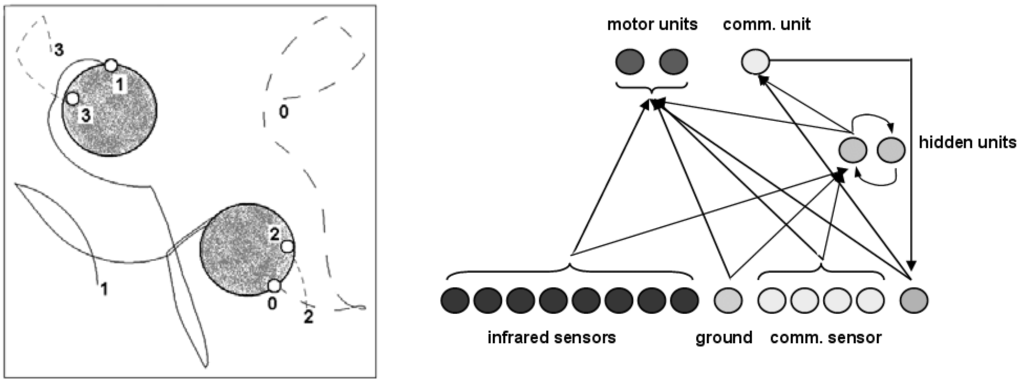


Figure 4. **Left**: The environment and the robots. The square represents the arena surrounded by walls. The grey circles represent the target areas. Lines inside the arena indicate the trajectory of the four evolved robots during a trial. The numbers indicate the starting and ending position of the corresponding robot (the ending position is marked with a white circle). **Right**: The neural controller evolving robots. The three motor neurons control the motors of the two corresponding wheels and the intensity of the signal produced. The sensory neurons encode the state of the robot infrared sensors, of the ground sensor (that detect whether the robot is placed on a target area or not), and the signal produced by other robots up to a certain distance from four corresponding directions.

By evolving the connection weights of the robots' neural controller for the ability to solve this collective navigation problem, one can observe that evolving robots rely on their ability to produce and detect other robots' signals to cooperate and coordinate in order to solve this collective navigation problem.

The analysis of the obtained results (Marocco & Nolfi, 2005) indicate that evolving robots develop: (a) an effective communication system, (b) an effective individual behaviour, (c) an ability to rely on different communication modalities and to autonomously select the modality that is appropriate to the current circumstances.

The communication system that emerges in these experiments is based on 4-5 different signals that characterize crucial features of the environment, of the agents/agents relations, and agents/environmental relations (e.g. the relative location of a target area, the number of agents contained in a target area, etc.). These features, which have been discovered autonomously by the agents, are grounded in the agents' sensory-motor experiences. Used signals, therefore, do not only refer to the characteristics of the physical environment but also to those of the social environment constituted by the other agents and by their current state. Evolved individuals also display an ability to appropriately tune their individual and communicative behaviour on the basis of the signals detected (e.g. by approaching, avoiding, or exiting a target area, by modifying their exploratory behaviour, etc.). Indeed, the type of signals produced, the context in which they are produced, and the effect of signals detected constitute three interdependent aspects of the communication system that co-adapt during the evolutionary process and co-determine the 'meaning' and the efficacy of each signal and of the communication system as a whole.

The individual behaviour of evolved robots includes simple elementary behaviours that allow robots to avoid obstacles, explore the environment, and remain in target areas. Interestingly, the individual behaviour of the robots tends to be optimised (with respect to the possibility to obtain the best possible performance when signals produced by other robots are not available) despite the fact that robots are always evaluated in social conditions during the evolutionary process. This result might be explained by considering that the required signals might not always be available (even in normal conditions in which robots are allowed to communicate) since their availability depends on the physical location of the robots in the environment, which in turn depends on unpredictable events such as the initial position and orientation of the robots, or noise. In other words, optimised individual behaviour guarantees good performance even when required signals are not available. This tendency to optimise both individual and social behaviour lead to the development of control systems structured hierarchically according to a layered organization in which the individual abilities represent the most basic layer and communication/social abilities represent an higher level layer that modulate the lower level.

Individual and communicative behaviour tend to co-adapt. Individual behaviours are thus selected to maximize both individual performance (when signals from other robots are not available) and social performance (i.e. the performance that can be achieved by combining the individual and social capabilities of the robots).

The analysis of the evolutionary dynamics suggests that new individual capabilities might represent a crucial pre-requisite toward the development of new communication capabilities and vice versa. For example, the individual ability to explore the environment by entering and remaining into target areas represents a crucial pre-requisite for the development of an ability to produce a signal that attracts other robots toward the target area. On the other hand, the emergence of a social/communicative ability to avoid target areas that contain two robots and to exit from areas that contain more than two robots, represents a crucial pre-requisite for the development of better individual exploration strategies. In fact, highly effective exploration strategies provide an adaptive advantage only in combination with effective communication systems that allow the robots to avoid situations in which more than two robots are located in the same target area. This process in which progress in individual abilities might pose the basis for the achievement of progresses in communication abilities and vice versa might lead to an open ended evolutionary phases in which individuals tend to develop progressively more complex and effective strategies.

Evolved robots also exploit different communication modalities (e.g. mono-directional communication forms in which one robot acts as a ‘speaker’ and a second robot acts as a ‘hearer’, or bi-directional communication forms in which two robots concurrently influence each other through their signalling and/or motor behaviour) by selecting the modality that is appropriate to each specific communicative interaction. In some cases evolving individuals also engage in complex communication behaviours that involve three different robots that concurrently affect each other so to produce appropriate collective behaviours (e.g. so to push one of the three robots located inside the same target area out of the area). Finally, evolved robots exploit time varying signals that allows them to generate information that is not available to any single robot (e.g. information related to how many robots are located in a target area) and that serves different functions.

## 4.2 A research roadmap

Recent research in this area, including the example described in the previous section, clearly demonstrate that artificial evolutionary methods might lead to the development of simple animal-like communication forms. The identification of the potentialities and the limits of this method, however, is still an open question. In particular it is not clear whether the method illustrated in section 4.1 and 4.2 can scale up to more complex situations that might require the development of richer communication systems (that include a larger number of signals) or more complex communication systems (that include additional features such us compositionality or grammar).

From a scientific point of view, that is, from the point of view of understanding communication in natural organisms, evolutionary methods are an obvious candidate for modelling animal-like communication systems. The vast majority of animal-like communication (e.g. bee’s dances or vervet monkey’s alarm calls), in fact, are genetically transmitted and evolved. Since cultural evolution and transmission play a crucial role in language evolution, evolutionary methods can only be used to model the emergence of language or the transition from animal to human communication forms, in combination with other methods (see section 6).

From a technological point of view, that is, from the point of view of developing effective artificial agents, the potentialities and the limits of evolutionary methods are still largely unknown. In fact, differences between natural and artificial agents (related to their body structure, the media used to communication, the environmental structure, and the agents goals) might potentially lead to new forms of communication systems, which might for instance include characteristics of both animal and human-like communication systems.

In the next section we will discuss the methodological aspects that might be crucial for developing effective ECAgents without restricting our analysis to simple animal-like communication forms. In doing so, we will also try to identify the most promising research directions in this area.

### 4.2.1 Identifying suitable problems and domains

One fundamental issue for exploring the potentialities and limitations of evolutionary methods is the identification of suitable problems and agent/environmental structures for studying the emergence of communication in ECAgents.

As external observers we cannot formally determine whether a given group of agents situated in a given environment necessarily require a communication capability or not to solve a given problem. Each problem, in fact, typically admits several different solutions that are hard to imagine from the point of view of an external observer, even on the basis of a detailed description of the problem and the agent/environmental structure. However, we might follow at least one heuristic approach in which we try to identify the class of problems that, most likely, require communication. These classes, as we will see below, might also provide a taxonomy of communication forms (see also Oliphant, 1997; Nolfi, 2005).

A first class includes cases in which *agents have access to different sensory-motor information and in which they might increase their adaptive capabilities by having access to and exploiting other agents' information*. This class of communication behaviours can be described, from the point of view of an external observer, as *information exchange behaviours* in which one agent (the speaker) send signals that encode a given sensory-motor information and a second agent (the hearer) detects the signal and exploits it by modifying its successive motor behaviour accordingly. This type of communication behaviours might be further classified with respect to the type of information that is exchanged and exploited. These types might include: (a) information related to the current sensory-motor state experienced by an individual agent that, by being located in a given position and orientation or by being provided with different sensory-motor system from other agents, might have access to information to which other agents do not have access to, (b) information collected by an agent during its previous interaction with the environment that might not be available to other agents that have had different sensory-motor experiences, (c) physiological information related to the internal states of an individual agents that, by definition, is not directly available to other agents, (d) intentional information related to what an agent is going to do or when it will perform a specific action that, by being based on the internal state and the previous sensory-motor experiences of that agent, is not directly available and cannot be easily inferred by other agents.

A second class includes cases in which *agents might increase their adaptive ability by manipulating the behaviour of other agents.* From an external observer point of view, this form of communication behaviour might be described as a form of *manipulation behaviour* in which a first agent (the puppeteer) produces a signal that directly or indirectly encodes a command, and a second agent (the puppet) executes that command. This type of communication behaviour might be further classified with respect to the type of commands that are exchanged and exploited: (a) commands related to actions that should be executed by the puppet agent and that might not be directly executed by the puppeteer since the two agents are located in different relative positions or orientations or have different motor capabilities, (b) commands related to actions that are constrained by the previous execution of other specific preparatory actions that have been performed by the puppet but not by the puppeteer, (c) commands encoding actions that can only be performed in certain physiological conditions (e.g. when the necessary energy is available), (d) commands encoding information on when a given action that cannot be directly executed by the puppeteer, should be executed by the puppet agent.

Although a categorization based on the two classes described above is useful to characterize different communication forms, it should be stressed that the two classes are not necessarily mutually exclusives. Indeed, communication behaviours produced by a first agent that trigger a specific motor reaction in a second agent, and which have an adaptive functionality for both agents, can typically be interpreted both as an instance of information exchange behaviour and as an instance of manipulation behaviour.

A third and final class includes cases in which *agents should produce tightly coordinated behaviours that require a continuous bi-directional tuning of agents motor and/or signalling behaviour*. From an external observer point of view, these communication behaviours might be described as a form of *entrainment* in which agents mutually affect each other and in which the individual contribution of each agent on the resulting social behaviour cannot be separated (Di Paolo, 2000). This form of communication behaviour might be further classified with respect to the type of bi-directional interactions that might affect (a) agents' motor behaviour, (b) agents' communication behaviour, or (c) both agents' motor and communication behaviour. More specifically:

(a) bi-directional interactions that affect agents motor behaviour might be necessary to display joint coordinated motor behaviour (e.g. shared attention behaviour, or synchronized behaviours) that might require a continuous co-tuning of agents' motor behaviour.
(b) bi-directional interactions that affect agents communication behaviour might lead to more complex communication ability in which agents might tune their signalling behaviour on the basis of the current social context (e.g. by shaping the signals produced on the basis of the state of the agents that will detect them, by using signals to manipulate other agents' signalling behaviour, by taking turns [Iizuka & Ikegami, 2003a, 2003b]). These forms of bi-directional interactions affecting agents' communication behaviour might be required to reduce interferences in problems in which agents might detect the signals concurrently produced by several other agents or in problems in which agents might find themselves in situations in which they have to select the most relevant information to be communicated.
(c) combinations of the two classes described above.

### 4.2.2 Adaptive factors in the emergence of communication

The emergence of communication abilities requires the development of two complementary but independent abilities: an ability to produce signals that are useful (from the point of view of the signaller, the receiver, or both) and an ability to react to signals in a way that is useful (from the point of view of the signaller, the receiver, or both).

From an evolutionary point of view, a first important challenge consists in the fact that variations that lead to the production of a useful signal tend to be retained only if agents already have the complementary ability to react to that signal in a useful way, or vice versa, variations that lead to an ability to react to a signal in a useful way tend to be retained only if agents already have the complementary ability to produce the corresponding signal. In other words, adaptive variations that lead to the production of useful signals or to the exploitation of signals are adaptively neutral (and therefore might be lost) until the corresponding complementary condition is met. The evolutionary emergence of communication behaviour thus seems to be constrained by the possibility to exploit pre-existing traits that were neutral or that served different functions. One example of this process is the exploitation of the receiver bias (Maynard Smith and Harper, 2003) in which environmental cues that trigger a functional reaction in an animal (e.g. noise produced by a prey that trigger a chasing behaviour) might be exploited through the production of a signal (e.g. a signal that resemble the noise produced by the prey) that can serve a new function (e.g. triggering a chasing behaviour in conspecifics). Once this form of communication is established, the signal and the corresponding motor reaction might vary and assume new functionalities.

An interesting research direction thus consist in studying these type of processes and in identifying mechanisms that might favour the emergence of characteristics that, although they do not immediately provide an adaptive communicative advantage, might be later exploited by communication.

A second important issue is the identification of the level (i.e. individual level, group level, or eventually a combination of individual and group level) at which selection should operate. Indeed, while communication behaviours of shared interests tend to be evolutionary stable (Maynard Smith, 1991), communication behaviours that provide an advantage for the signaller or the hearer only tend to be unstable (when selection operates at the level of individuals) as a result of the conflict of interests between individuals. Indeed, populations including agents with a newly acquired characteristic that consists of the ability to produce a signal in a certain context that provides an adaptive advantage for the receivers (that react to that signal appropriately) tend to be invaded by mutant agents that keep their ability to exploit the signal produced by the other agents but do not produce the signal themselves. In this condition in fact, mutant's fitness might increase to the expenses of the other individuals. This will lead to an increase of mutants' individuals and consequently to the loss of the communication ability. For sake of completeness it should be noticed that conflict of interest could also arise when communication is advantageous for both the signaller and the receiver but when the costs/benefits ratio between the two differs.

The conflict of interests between individual might lead to cyclical dynamics (Batali, 1995; Mirolli & Parisi, 2005b, 2005c) in which: (1) a communication ability emerges, (2) the communication ability spreads in the population, (3) a mutant individual emerges, (4) the mutant individuals spread in the population so that the communication ability is destroyed, (5) as soon as a communication ability is discovered or re-discovered the process re-start from the first step. Although these cyclical dynamics have only been experimentally observed in experimental studies in which communication was only beneficial for the receiver, the same dynamics should be expected in cases in which communication is beneficial for the receiver only. In this latter case, in fact, the emergence of communication abilities that are only advantageous for the signaller, might be followed by the emergence of opportunistic mutant individuals that keep their ability to produce the signal but react inappropriately to the signal produced by other individuals.

Possible mechanisms that might reduce the effects of conflicts between individual and collective interests include:

(a) The use of teams of agents that are homogeneous (i.e. that share the same genetic characteristics (Baldassarre, Nolfi, and Parisi, 2002, 2003; Quinn et al., 2003; Marocco & Nolfi, 2004, 2005; Magnenat, Floreano, and Keller, 2005]). In this evolutionary schema, conflicts between individuals and collective interests cannot occur. The homogeneity between individual agents, however, also implies a homogeneity in agents' communication abilities which in turn might prevent the occurrence of a self-organization process at the level of the communication system in which, for example, alternative signals might compete for the same meaning.
(b) Selection schemas in which individual agents are selected on the basis of the performance of the group (Perez-Uribe, Floreano & Keller, 2003; Marocco & Nolfi, 2004, 2005; Magnenat, Floreano & Keller, 2005).
(c) Evaluation schemas in which individuals are evaluated in groups formed by genetically related individuals (Magnenat, Floreano, and Keller, 2005; Mirolli and Parisi, 2005a) or in which individuals have an higher probability to interact with their kin (e.g. spatially distributed selection schemas in which the probably to mate and to socially interact with other agents is proportional to the distance between individuals [Ackley & Littman, 1994; Oliphant, 1996]).
(d) The constraints introduced by the agents' internal categories developed to produce effective individual behaviours (Cangelosi & Parisi, 1998, Marocco, Cangelosi, & Nolfi, 2003, Mirolli and Parisi, 2005b, 2005c).
(e) The use of signals not only for social communication but also for "talking to oneself", that is as an aid to individual cognitive abilities (Mirolli and Parisi, 2005a);
(f) The presence of a cost on signalling (Grafen, 1990; Maynard Smith, 1991).
(g) Mechanisms that allow agents to identify opportunistic agents and to adopt adequate counter-measures that might reduce or eliminate the adaptive advantage of opportunistic behaviour.

#### 4.2.3 Identification and implementation of the necessary pre-requisites

From an evolutionary perspective the most straightforward way to approach the issue of how effective forms of interaction and communication can emerge is to start from simple but open-ended models that might lead to the emergence of progressively more complex forms of behavioural, cognitive and social capacities. After all, this is how these abilities emerged in natural life. This possibility, however, can reasonably be pursued only as a long-term research goal. On the short term, it is reasonable to assume that progresses might be achieved only by providing evolving agents with crucial characteristics that constitute pre-requisites for the development of communication abilities. Below we briefly illustrate some features that represent crucial pre-requisites for the emergence of communication and/or for the emergence of complex communication forms.

A first important pre-requisite is the availability of sensors and actuators that serve both individual and social/communicative functions. These two functions, in fact, cannot be completely separated since ECAgents can communicate either directly (e.g. through sound signals that mainly or only serve a communicative function) or indirectly (e.g. through the effect of their motor actions that serve also other non-communicative functions). An important role, in particular, might be played by sensory-motor structures that: (a) tend to spontaneously manifest internal states (e.g. facial expression), (b) produce permanent or long-lasting modification (e.g. pheromone), (c) allow pointing and detection of pointing, and (d) that might allow to identify specific individuals and detect what other individuals are doing.

A second important prerequisite concerns cognitive abilities such us: (a) the ability to form abstract categories by integrating sensory-motor information in time (Nolfi, 2002), (b) selective attention mechanisms that allow agents to focus on relevant sensory information, (c) mechanisms that allow agents to display a variety of different behaviours and to shape their behaviour on the basis of the current environmental and social circumstances.

A third important prerequisite is a genotype-to-phenotype mapping that ensures evolvability, i.e. the ability of random variations at the genotype level to sometimes produce improvement at the phenotype/behavioural level (Wagner & Altenberg, 1996; Nolfi and Floreano, 2000).

### 4.2.4 Identification and implementation of the necessary limits and constraints

Evolving agents have two possible ways to increase their adaptive capabilities: improving their individual behaviour or improving their social/communicative behaviour. Although the possibility to co-adapt both the individual and the social/communicative behaviours represents one of the most interesting characteristics of evolutionary methods, the competition between individual and social/communicative behaviour might lead to local minima in which social/communicative potentialities are not, or not fully, exploited. From an adaptive point of view, therefore, limits and constraints in the agents' sensory-motor structure or in the agent/environment interactions that limits the potentiality of individual behaviour might represent an important pre-requisite for the emergence of social/communication abilities in ECAgents.

More specifically, the emergence of communication forms that allow individual agents to indirectly access other agents' sensory information can only be expected in agents that have limited sensory capabilities. In other words, we should expect that, the larger the limitation of agents' sensory capacities are, the larger the potential adaptive advantages of social/communicative behaviours will be. Similarly, we should expect that, the larger the limitations of agents' motor capabilities are, the larger the adaptive advantage of social/communicative behaviours (that allow agents to indirectly control other agents' behaviour) will be. A similar role might be played by adaptive constraints that operate at the level of individual behaviour (e.g. the need to minimize the energy consumption, the constraints imposed by the current energy level, the need to minimize the time required to perform a given task, the need to accomplish several tasks at the same time, etc.)

Constraints that limit agents' social interactions or agents' sensory-motor capabilities might also favour the emergence of sophisticated communication abilities. For example, limits on the time rate with which agents are allowed to socially interact/communicate might favour the development of communication abilities based on abstract categories (i.e. the use of abstract signals that encode properties of sequences of sensory-motor states). Similarly, the impossibility to detect the relative position of two communicating robots might favour the emergence of agents able to communicate non-deictic information (e.g. spatial information that is independent from the signaller and receiver relative position.)

## 5. Social Learning Methods

In this section, we explore a different route than Evolutionary Methods: Social Learning. Instead of having an evolutionary fitness function, the emergence of more complex communication systems is driven by the social dynamics in a population of agents, that is, the need to maximize communicative success while at the same time minimizing the required resources (time, processing, memory, etc.). Every agent is therefore endowed with mechanisms to conceptualize what to say, produce and parse an expression, and interpret it back into reality. The agent also needs to have mechanisms to invent new categories and language conventions, adopt them from others, and co-ordinate conceptual and linguistic repertoires as a side effect of games. Although there can be a change in the population (new members coming in without any knowledge of existing conventions and other members going out taking their knowledge with them), the transmission (genetic or cultural) itself is not seen as the prime causal factor by which agents reach coherence or by which additional structure arises in the communication system.

Language is seen as a complex adaptive system that continuously has to cope with new challenges, meanings, etc. As human languages are capable of doing this, they are an important source of inspiration for the ECAgents design. The researcher takes a 'universal' feature of human language as the focus of his experiment and formulates a hypothesis that states which strategies are needed for this feature to emerge, and which mechanisms are used by these strategies. The strategies have to be found in the general cognitive domain of the agents, from where they can be 'recruited' by the language faculty. This recruitment operation is also known as the Recruitment Theory of the Origins of Language (Steels 2005a).

The experimental design typically involves at least two embodied robotic agents that are engaged in language games – routinised and measurable communicative interactions. The kind of games and the environment in which they are played have to be chosen in such a way that they provoke the need for the language feature that the researcher wants to observe. The experiments compare the communicative success and cognitive economy of communication systems with the recruitment of strategy Y and its mechanisms $X_1$, … $X_n$ to communication systems that don't use this strategy. If the communication system that uses strategy Y displays a significant improvement over the other communication system, the hypothesis is confirmed.

## 5.1 An Exemplificative Case: Evolving ECAgents able to Perform and Mark Perspective

A nice example of a universal language feature is the marking of perspective. All human languages allow speakers to adopt perspectives other than their own in conceptualizing the world, and mark this explicitly through words or grammatical constructions (for instance *my left* versus *your left*).

The hypothesis is that communication systems that use the strategy of perspective reversal will lead to more communicative success and expressive power, while at the same time reducing the cognitive effort of the agents, than communication systems that do not use perspective reversal. The hypothesized strategy is found in the general cognitive domain of the agents: the Egocentric Perspective Reversal system that enables an agent to reconstruct the other agent's perception of the world.

### 5.1.1 The Experimental Design

To provoke the need to express perspective, the agents have to communicate about events that they observe from a different angle. The experiment starts with a population of agents that are able to detect and track other robots and an orange ball. Two of the agents are selected as speaker and hearer and are instantiated in robots. They roam around freely in an unconstrained in-door environment, until one of them sees the orange ball, comes to a stop and searches for the other robot. The other robot also looks for the ball and stops when it sees it. The experimenter then pushes the ball with a stick so that it rolls for a short distance, for instance from left to right. This movement is tracked and analyzed by both robots and each one uses its own resulting perception as the basis for playing a language game.

The language game that is played by the agents is a guessing game: the agent that is selected as the speaker describes a ball-moving event to the hearer, either the last event or the event before that. Next the hearer parses the expression, guesses which event is compatible with the communicated meaning, and signals its choice to the speaker. The game is a success if the guess was the right one according to the speaker.

These series of language games were played and measured in four different scenarios: (a) the first time, the agents had to communicate about the observed events without a difference in perspective. In this rather unrealistic set-up, perspective plays no role in communication. Here, only two core mechanisms are applied: Categorization (based on Discrimination Trees) and a Bi-Directional Associative Memory for the lexicon. Then, (b) the same scenario was repeated, but this time introducing the difference in perspective. The next step (c) was to conduct the experiment again, but this time the agents were able to recruit perspective reversal. They were not able, however, to overtly mark perspective in their emergent language. And (d) finally, the same scenario was repeated again with the recruitment of perspective reversal and the overt marking of perspective in the linguistic system.

### 5.1.2 Results

In the first scenario, the speaker and the hearer shared the same world model and perspective did not play a role in communication. This experiment proved to be entirely successful with a communicative success of almost 100%. This scenario shows that the learning mechanism works and that the agents are capable of constructing a self-organized lexicon from scratch.

In the next scenario, each agent used its own perceptual data, which led to a different world model because the agents perceived the events from a different angle. The communicative success makes a drastic drop to a maximum of 60%. The only cases in which success is reached are those in which the viewpoints of the agents are sufficiently similar. This shows that embodiment automatically leads to different world models and the need for a more complex communication system.

From the third scenario on, the agents are able to recruit perspective reversal. This reconstruction of the other agent's viewpoint is never completely accurate and the emerging communication system should be able to cope with this problem. The experiment shows that, even when perspective is not overtly marked in the communication system, communicative success rises significantly with the use of the strategy of perspective reversal.

The fourth and final scenario is the same as the previous one, but this time the agents are able to overtly mark perspective in the emergent language. Marking perspective is not really needed for communicative success (the meaning can be inferred from the context), but it strongly reduces the cognitive effort. When the speaker specifies which perspective is adopted, the hearer knows immediately whether he should do a perspective transformation or not. This explains why the agents choose for the recruitment of perspective reversal and introducing new words for it in the lexicon.

A detailed overview of the experiment can be found in Steels, Loetzsch & Bergen (2005) and the supporting materials can be accessed at http://www.csl.sony.fr/perspective/.

## 5.2 A Research Roadmap

The case study presented in this section builds on a much larger body of research in which Social Learning Methods have been explored in the study of language, all following the same methodology. The following two subsections can be used as a guide through the most important achievements so far and as a starting point for new research directions.

### 5.2.1 Previous Research Efforts

Viewing language as a complex adaptive system has led to a whole avenue of new research directions. The first successful experiments involved lexicon formation. Within the framework of language games, it was shown how a population of autonomous agents was able to self-organize and to establish a set of linguistic conventions. By playing the so-called Naming Game, in which a bi-directional associative memory is recruited (Steels 1995, Baronchelli et al. 2005)), the agents were able to agree on ways to identify each other using names or spatial descriptions. Based on communicative success during social interactions, the agents continuously adapt their language inventory, which results in a positive feedback loop. The emergent communication systems are flexible enough to cope with new meanings and new members of the population without a performance loss.

At the same time, experiments were being conducted on meaning creation (Steels 1996). In these experiments, the 'grounding' challenge was being addressed. By recruiting Categorization, the agents were able to generate perceptually grounded distinctions by playing Discrimination Games. A new series of experiments combined the Discrimination Game with the Naming Game, which resulted in the Guessing Game (Steels 1997, 2003). Here, the agents create perceptually grounded categories and use language to share these categories. The experiments show how language and meaning creation co-evolve and continuously influence each other.

The research showed how coherence could be reached due to self-organization, but didn't explain yet why human languages keep changing all the time. New experiments were set up to explore the hypothesis that language change can be explained through the 'stochasticity' and noise observed in real-world natural language use (Steels and Kaplan 1998). These experiments showed that the self-organized lexicons of the agents do not evolve even if the population is changing. By introducing stochasticity[1], the emergent lexicons display constant innovation (new word-meaning associations), optimization of the language inventory and maintenance of variation in the population due to the tolerance agents need to exhibit in order to cope with stochasticity. All these research efforts culminated into the 'Talking Heads' experiments, which continued to focus on the co-development of conceptual and linguistic repertoires (see for instance Steels and Belpaeme 2005).

The Social Learning Methods can be expanded to all other domains of language as well. This is clearly shown in recent investigations into the origins of vowel systems and phonology (De Boer 2000, Oudeyer 2005). These experiments show that agents can self-organize vowel systems if they recruit a Bi-Directional Associative Memory to associate sound patterns with articulatory gestures. The emergent vowel systems even show the same statistical distribution as observed in human languages, thus providing a very strong claim for the Recruitment Theory of the Origins of Language.

Steady progress is also being made with regard to multi-word expressions (Van Looveren 1999), the origins of grammar (Steels 2005a) and the field of semantics (Steels and Bleys in press). These new directions lead to some exciting new projects that will be discussed in the next subsection.

### 5.2.2 Identifying and analysing minimal models

It is important to stress the need in this field of shared and general models to create a common framework where different research groups could compare their approaches and discuss the

---

[1] Note that stochasticity is an inherent feature of grounded language use, not of the agent's language capacity. This means that there is no recruitment of a certain strategy.

results. On the other hand the models should exhibit the extreme level of simplicity compatible with the desired phenomenology. This has several advantages. It could allow for discovering underlying universalities, i.e. realizing that behind the details of each single model there could be a level where the mathematical structure is similar. This implies, on its turn, the possibility to perform mapping with other known models and exploit the background of the already acquired knowledge for those models. (b) identifying the most suitable theoretical concepts and tools to attempt the solutions of the models. It is important to outline how the possibility to obtain analytical and general solutions for the models proposed could open the way to a positive feedback providing further inputs for understanding and designing new experiments and devices. In general there are two main questions one should be able to answer. On the one hand, we need to find the general laws that govern the semiotic dynamics of a particular system, for example, how the maximum number of words in use is related to the number of agents in the population. On the other hand, we need to find the explanation of these laws as a mathematical property of the dynamics. (c) coupled to the theoretical activity there should always be an experimental activity with a twofold aim. On the one hand the experiments, as well as the observation of the realities one is interested in, provide inputs for the modelling and the theoretical activity. On the other hand they represent the framework where the theoretical predictions are checked. The outcome of these experiments have to be compared with the theoretical and the numerical results and used to better focus the modelling and the theoretical approach. There should then exist a positive feedback mechanism between the theoretical and the experimental activities in order to make the progresses robust, well-understood and concrete.

A list of important issues which deserve consideration in a modelling and theoretical activity include:

- The role of the topology and the interplay between dynamics (interactions) and topology;
- The role of the system size (scalability);
- The role of specific strategies or features the agents are endowed with;
- The role of noise.

For a preliminary study in this direction that investigated the time require by agents to reach a consensus on the word used to indicate an object in a naming game see Baronchelli et al. (2005).

### 5.2.3 Fluid Construction Grammar

The current experiments are being conducted in Fluid Construction Grammar (FCG), a new grammar formalism that has been designed for doing experiments on the emergence and development of grammar. It captures the fluidity by which new grammatical constructions can enter or leave the language inventory and provides the agents with a much stronger framework to express more complex meanings. The implementation includes components that turn a semantic specification into a sentence via intermediary structures, and components for parsing a sentence into its semantic interpretation. In FCG, a construction always associates a semantic structure with a syntactic structure.

FCG uses many techniques from formal/computational linguistics, such as feature structures and unification, but the formalism and its processing components have a number of unique properties: all rules are bi-directional so that the same rules can be used for both parsing and production. They can also be flexibly applied so that ungrammatical sentences or meanings that are only partially covered by the language inventory can be handled without catastrophic performance degradation. Fluid Construction Grammar is already fully operational and experiments are being conducted on how grammar can be learnt within the context of situated embodied language games.

A lot of work still needs to be done, however, both from the engineering part of the formalism as from new language game experiments. Some of the current challenges feature the implementation of hierarchy in FCG (De Beule and Steels 2005), linking (Steels, De Beule and Neubauer 2005) and experiments on the predicate-argument structure, i.e. who is doing what to whom (Steels 2002). This still leaves numerous pathways unexplored: not only the many universal features of language that still need to be explained, but also combining all these strategies into a developmental implementation (similar to child language acquisition) and research on how agents themselves could discover suitable strategies to recruit in an epigenetic recruitment process.

The research has also reached a stage in which there is need for richer semantics. Language is a complex adaptive system, and this means that we have to accept this complexity. A lot of effort is therefore being put into a system for planning complex meanings and on their grammatical expression in Fluid Construction Grammar. The development of this system will have a direct impact on the design of FCG and will also lead to many new research topics. To be able to continue the steady stream of progress of the past ten years and to have a clear idea of which directions to explore, we keep on following the developmental stages proposed by Luc Steels (chapter 10, Steels & Nolfi, 2004).

## 6. Combining of Evolutionary and Learning Methods

In humans and in many other organisms language and communication systems evolve as a result of two adaptation processes – genetic evolution and learning – and of their interactions. As we will argue in this section, combining evolutionary and learning methods might be useful (or possibly necessary), also from the point of view of developing artificial ECAgents.

One first motivation for combining these two adaptive techniques is the possibility to co-evolve and co-adapt the communication system and the agents' physiological, cognitive, and behavioural characteristics. The characteristics of agents' control system (e.g. the architecture of agents' neural controller), their cognitive abilities (e.g. the ability to focus attention on relevant sensory information), and their behavioural abilities (e.g. the ability to produce specific behaviours that allow agents to collect information from the environment) are crucial pre-requisites for the emergence of communication. Part of or all these features can be developed through an artificial evolutionary process in which the neuro-physiological characteristics of the agents are encoded in an artificial genotype and in which individuals are selected on the basis of their behavioural capabilities. Combining an evolutionary process with a social learning process (e.g. a process in which agents develop a shared communication system through a social-learning method like that described in section 5) might allow to co-adapt agent's cognitive, physiological, and behavioural abilities and agent's communication abilities. Indeed, agents' cognitive, physiological, and behavioural capabilities that represent a pre-requisite for the emergence of a communication might later undergo changes triggered by genetic variations that are adaptive from a communication point of view.

A second motivation for combining evolutionary and learning techniques is the attempt to exploit the complementary characteristics of the two processes and the properties arising from their interactions. Although both genetic evolution and social learning are adaptive processes based on the same two fundamental mechanisms – variations and selection – they differ in important respects that make them especially suitable for different domains. One important difference, for instance, is the fact that genetic evolution is based on random variations that are introduced randomly and are retained or discarded on the basis of their effects on the long term. In social learning, instead, variations are produced on the basis of directional processes (e.g. an imitation learning procedure) and are kept or further modified on the basis of their effects on the short term. A second important difference is that in genetic evolution adaptive characters are transmitted as high fidelity copies from parents to offspring while, in social learning characters are not transmitted but rather generated through bi-directional social interactions and between many different individuals. These differences imply that genetic evolutionary methods might be more effective when: the effects of variations can be evaluated only over long periods of time, the number of characters to be transmitted is relatively small, and accuracy is more important than adaptability. On the contrary, social learning methods might be more effective when: the effects of variation can be evaluated in short time periods, the number of characters to be transmitted is high, and adaptability is more important than accuracy.

A third motivation is that the combination of evolutionary and learning methods might lead to useful dynamics that emerge from the interaction between these two adaptive processes. For example, in some cases characters discovered through learning might be assimilated in inherited characters later on, thus leading to more effective individuals that do not need to acquire these characters through learning anymore (Baldwin, 1896; Waddington, 1942). Or, for example, the combination of evolutionary and learning methods might lead to the emergence of characters, genetically encoded, that are not adaptive by themselves but that favour the acquisition of adaptive characters through learning (Nolfi & Floreano, 1999). These pre-disposition to learn might consists of biases or abilities that favour agents individual learning capabilities (e.g. a “curiosity” bias that might increase agents’ chances to be exposed to useful learning experiences) or agents’ social learning abilities (e.g. an ability to produce behaviour that are easy to imitate or that can be easily be transformed so to match other agents behaviour).

## 7. Universalities and their implications for the origin of communication and language

The emergence of language is one of the major transitions in evolution (Smith & Szathmary, 1997). Understanding the origins of communication and language thus implies understanding the essential pre-requisites that lead to this transition. One way to approach this problem involves identifying these crucial pre-requisites, model and implement them in a population of artificial agents that is left free to self-organize, and observe whether an homologous transition occurs. The analysis of the language or communication system that emerge from agents interactions (if any) and the comparison of the effects of different initial conditions might allow us to shape our initial intuitions and to modify our models so to progressively increase our understanding of the phenomenon and our ability to replicate it on an artificial system. This is the approach followed in the methods described in the previous 4 sections. A different, complementary approach consists in trying to identify general properties of existing language and communication systems or of the dynamics that can be observed in language acquisition processes. The identification of these universal properties and the identification of methods with which these properties can be analysed might provide important indications on the essential pre-requisites that have lead to the emergence of language in natural systems, and to new methods for analysing artificial systems and for comparing natural and artificial systems. Universals also help to formulate old and new questions in a much more quantitative way, thus allowing proper testing of tentative theories on a scientific basis.

In this section we will discuss how recent development from statistical physics, complex network theory, and evolutionary theory might lead to the identification of universal properties of natural language and communication systems. Moreover, we will show how the same methods and techniques can be applied to computer mediated communication systems (e.g. Flickr, Connotea, Citeulike, Delicious).

## 7.1 Linguistic networks

Words in human language interact within sentences in non-random ways and allow humans to construct a large variety of sentences from a limited number of discrete units. The co-occurrence of words within sentences reflect language organization in a subtle manner which can be described in terms of complex networks in which words are represented as nodes and relations between words are represented as links. The type of relations represented with links might vary depending on the extension and objectives of the study (Solé et al., 2005) and might consists of syntactic relations (Ferrer i Cancho et al., 2003) co-occurrence relations (i.e. two words are considered to be connected if they appear after each other at least in a given sentence, see Ferrer i Cancho and Solé, 2001; Yun and YaoYao, 2005; Solé et al., 2006) or semantic relations (i.e. words are related to each other through semantic associations such as synonymy, hypernymy, meronimy or antonimy; see Sigman and Cecchi, 2001).

All linguistic networks (either encoding syntactic, semantic, and co-occurrence relations) analysed up to now exhibit a Zipf's law distribution and a small-world property. Briefly stated, a Zipf's law distributions (Solé, 2006) means that, if we take all the words and order them by rank from the most common to the rarest, the frequency (number of occurrence) decays inversely with their rank (see Figure 5). The small world property means that the average distance between two words d (i.e. the average minimum number of jumps to be made from an arbitrary word to another) is very small compared to network size (Solé et al, 2005) and close to what should be expected from a random graph. Instead, the local organization, as described by the number of triangles found in the graph, strongly departs from randomness. Moreover, both syntactic (Ferrer I Cancho et al., 2003) and co-occurrence (Solé et al., 2006) networks analysed so far also display the dissassortative mixing feature, which implies that strong connected units are weakly connected between them while they tend to link to elements with few connections.

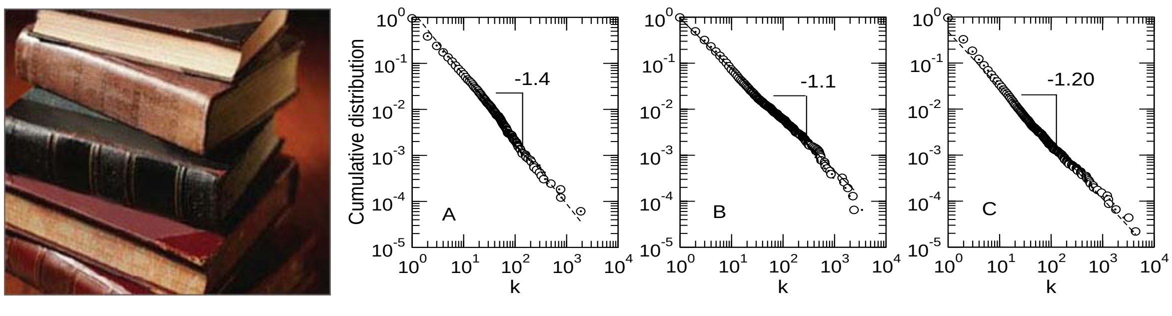


Figure 5. Co-occurrence networks can be obtained by using written corpuses, by considering two words as linked if they appear next to each other at least within a given sentence. The plots show the cumulative frequency of words having k links. They all follow a power law distribution, with most words having a few links and a small number of hubs being connected to many words. From left to right, basque, English and Russian are shown (Solé et al., 2005).

The same technique might be applied to analyse language development in children (see Figure 6). In particular, by using available corpus data one might explore how the topology of the linguistic network change as children learn their maternal language and the relation between changes in the network topology and changes in the children linguistic abilities (Yun and YaoYao, 2005, Corominas et al., in preparation).

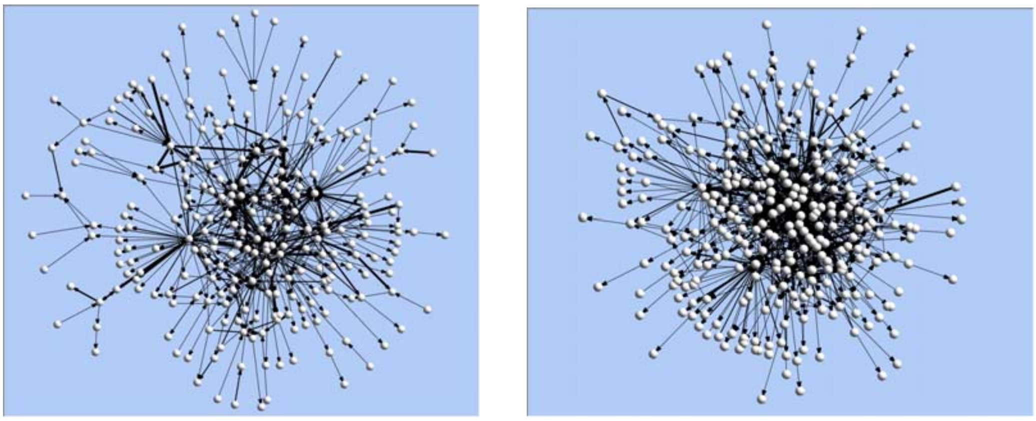

Figure 6. Networks constructed on the basis of child spoken language. The left and right pictures correspond to the linguistic network of the same child at two different stages of language acquisition. Nodes and links encode words and syntactic relations between them (Solé et al., 2005). The analysis of the topological evolution of these webs offers a very good, quantitative window to many relevant questions on language ontogeny and how it relates to cognition.

This research constitutes the first steps towards a theory of language networks. Further research based in other linguistic units (beyond words) or other type of relations might provide important new insights. In particular, new studies should consider the origins of language architecture under a dynamical perspective and how fundamental constraints might shape them (Solé, 2005). In this context, recent work using language games has show that the topology of interactions among agents might have a non-trivial role during the process of lexical evolution (Corominas and Solé, 2005). This work suggests that the pattern of agent-agent interactions might strongly influence the emergent language webs. Moreover, the analysis of emerging language networks in ECAgents might offer a unique opportunity to explore the architecture of artificial language webs and how they emerge.

## 7.2 Social networks

Communication and language always involve a population of interacting agents. In natural systems each agent does not interact with everybody else with the same probability: a social structure is at work, with non-homogeneous patterns of interactions among individuals. The different probability with which agents interact can be conveniently represented by social networks in which agents are represented as nodes and interaction probabilities as links.

Social networks based on empirical data, by providing a global pattern of systems where many agents interact locally, can allow us to identify universal properties of these system and correlations between topologies of interaction and global characteristics of these systems (e.g. correlation between the topology of the interaction and the characteristics of the communication system that emerge from the interactions). For example (Valverde & Solé (2004) have shown how some network topologies in complex systems (such as the Internet) can forbid a priori some undesirable results in systems constituted by agents that exchange messages related to a shared world. Some simple choices in the ways information is shared and transmitted can have dramatic effects on the system's performance in terms of communication.

In artificial systems, network theory might provide important insights on how the topology of the network might affects the properties emerging from the interacting agents and lead to the design of better artificial systems. In most of the cases, in fact, artificial systems assume that agents interact with equal probability with anybody else. The design of more realistic models needs to consider more restrictive social networks and the inspiration from the properties of real social networks is crucial.

Another interesting research direction concern, besides studying how a given topology affects the collective behaviour of a population of agents, what type of topologies can spontaneously emerge from a population of agents that have to solve a given problem. In this context, it has been recently shown that some unexpected commonalities are found in communities of social insects and groups of open source developers (Valverde et al, 2006). In both types of systems, an emergent, fluid hierarchy of interactions is developed through time, based on different types of information exchanges. The final result is a weighted network of interacting individuals following non-trivial scaling laws, with a few agents dominating the interactions and a majority of them exhibiting weak interactions with others. Once again, universal rules of (self-) organization seem to be at work. Understanding their origin and their importance in maintaining a reliable information exchanges will be a fascinating task for the future.

## 7.3 Emergence of language as an instance of major evolutionary transitions

The origin of language represents one of a number of 'major transitions' have been identified in the evolutionary literature (Maynard Smith & Szathmáry, 1995; Szathmáry & Maynard Smith, 1995). They include: (i) the transition from replicators to protocells; (ii) the origin of chromosomes; (iii) the appearance of the genetic code (translation); (iv) the origin of the eukaryotic cells; (v) the appearance of sex (meiosis and syngamy); (vi) the origin of multicellular eukaryotes (complex development and epigenetic inheritance); (vii) the origin of animal societies; (viii) the origin of natural language and human society.

These transitions have a number of features that are sufficiently common as applied to the list; namely: (1) the emergence of higher level units of evolution out of units at the lower level; (2) the appearance of novel inheritance systems; (3) the division of labour/combination of functions; (4) contingent irreversibility.

The problem of the origin of language is of course of the outmost importance (e.g. Christiansen & Kirby, 2003) This transition happened only in one lineage and is therefore unique, and in this respect has the status of the origin of the genetic code and the emergence of eukaryotic cell (Maynard Smith & Szathmáry, 1995). It has led to the emergence of a novel inheritance system (Jablonka & & Szathmáry, 1995) which opened up the possibility for cumulative cultural evolution and the beginning of history (Maynard Smith & Harper, 2003). Moreover, this novel inheritance system has allowed for the emergence of a much more complex society than anytime before: human society rests on negotiated division of labour and the collaboration of large non-kin groups (Avital & Jablonka, 2000).

The universalities of major evolutionary transitions and the specificities of the transition that lead to the emergence of language suggest important hints for identifying the fundamental properties that ECAgents should have and for devising the conditions in which ECAgents might emerge from initially non-communicating agents. For instance, the key role played by local interaction on the dynamics of selection and on the stability of cooperation in major transitions suggest that this aspect should not be oversimplified in designing ECAgents.

Moreover, universalities in major evolutionary transitions suggest that evolvable genotype-to-phenotype mapping might play a key role in the emergence of communication and language. Novel inheritance systems always appeared in the context of pre-existing ones; furthermore, specific evolution of the latter co-triggered the appearance of the former. Finally, another key role might be played by the possibility to observe the emergence or higher level units (e.g. institutions) able to reduce the variance of fitness within groups.

## 7.4 Semiotic dynamic on the WWW

Eight years after the first vision of the Semantic Web by Tim Berners-Lee a set of new semantically-enabled applications is swiftly shaping the next generation of the World-Wide Web (the so-called WEB 2.0). One of the forces driving this change is a distributed classification paradigm known as “*collaborative tagging*”, which has been successfully deployed in web application designed to organize and share diverse online resources such as web pages, academic references, photographs and music. This new paradigm, which is quickly gaining impact in large-scale information systems, is named folksonomies. In applications like Flickr, Connotea, Citeulike, Delicious, etc. people no longer make passive use of online resources - they take on an active role and enrich resources with semantically meaningful information. Such information consists of terminology (or “tags") freely associated by each user to resources and is shared with users of the online community. Web users interact with a collaborative tagging system by posting content (*resources)* into the system, and freely associating text labels (*tags*) with that content, as shown in Figure 7.

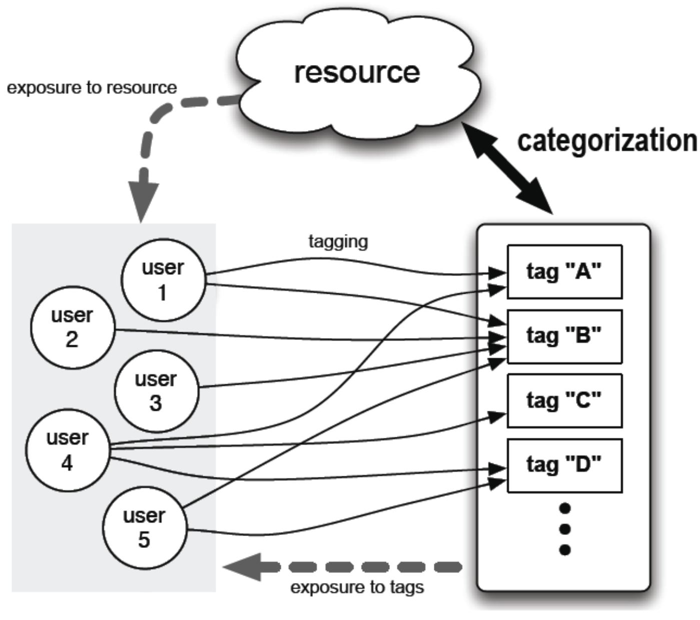


Figure 7. Schematic description of the collaborative tagging process: web users are exposed to a resource and freely associate tags with it. Their interaction with the system also exposes them to tags previously entered by themselves and by other users. The aggregated activity of users leads to an emergent categorization of resources in terms of tags shared by a community.

The basic unit of information in a collaborative tagging system is an **entry**, that is a **(user, resource, {tags})** triple, with the largest online communities comprising hundreds of thousands of users and millions of resources. Users are exposed both to the resources and to the tags already existing within the system, and freely associate tags with newly entered resources. At the global level the set of tags, though freely determined, evolves in time and leads towards patterns of terminology usage that are shared by the entire user community. Hence one observes the emergence of a loose categorization system -- commonly referred to as "*folksonomy*" -- that can be effectively used to navigate through a large and heterogeneous body of resources.

New web applications hinged on collaborative tagging (such as *del.icio.us* or *Flickr*) fall precisely in this perspective and they can be regarded as cases of Semiotic Dynamics at play: the emergence of a folksonomy exhibits dynamical aspects also observed in human languages, such as the crystallization of naming conventions, competition between terms, takeovers by neologisms, and more. The development of those tagging systems on the web has expanded that scope of those models of convention to systems that have been created spontaneously by a community of web users in a more natural context. These systems represent an ideal object of study for Semiotic Dynamics, that is the field that studies how populations of humans or agents can establish and share semiotic systems, typically driven by their use in communication. Indeed, these systems allow researchers to easily analyse not only the folksonomies that emerge from the computer mediated interaction between agents, but also its evolutionary dynamics.

The first data analysis that can be performed concerns the quantitative statistical analysis of the emergent metadata, as for instance the set of user-defined tags in folksonomy including the frequency distribution of the metadata (tags/keywords distributions). It is easy to foresee the observation of non-Gaussian statistics, i.e. the appearance of fat-tailed distribution, similar to the power law observed in other domains (see section 3.3, 7.1, Newman [2005]). Due to the dynamical nature of the systems in study, the time evolution of the distribution of its first moments have to be explicitly considered. Examining this sort of distributions could give a better indication of whether folksonomy converges on terms and fosters consensus, if and how the vocabulary of tags grows/scales as the number of users grows, and if the distribution flattens or narrows, perhaps indicating less or more agreement. Time evolution and time auto-correlation of the frequency of the typical or of the most popular tag surely represent interesting observable data. This kind of analysis, borrowing concept and methods from Probability Theory and Time Series Analysis, can clearly provide insight for the design of ECAgents. For instance (Cattuto, Loreto and Pietronero, 2005a and 2005b), on studying the frequency-rank distribution of tags co-occurring with the selected one, one finds a heavy-tailed behavior, which is the mark of human activity, as well as an emergent hierarchy of tags. We have introduced a stochastic model embodying two main aspects of collaborative tagging: (i) a fundamental multiplicative character closely related to the idea that users are exposed to each other's tagging activity, (ii) a notion of memory - or aging of resources - in the form of a heavy-tailed access to the past state of the system. Remarkably, our simple modelling is able to account quantitatively for the measured frequency-rank properties of tag association, with a surprisingly high accuracy. This is a clear indication that collaborative tagging is able to recruit the uncoordinated actions of web users to create a predictable and coherent semiotic dynamics at the emergent level.

Another analysis that might provide insights on the characteristics of these systems concern the topology of interacting elements. The raw data from the real systems can be used to define a very complex dynamically evolving network. Different type of networks can be modelled and analysed by encoding as nodes different elements (Mika, 2005): **users/actors**, **resources/instances** (web addresses, pictures, scientific citations, etc...), **tags/concepts.** The participation of users in folksonomies also allow to define social networks in which, for example, users are linked through the use of specific tags. Preliminary observations (Cattuto, unpublished) indicate that these type of social networks are not homogeneous but include distinct groups or communities in which nodes belonging to the same community are highly connected while nodes belonging to different community are only sparsely connected. However much remains to be done to describe the structure of these network in quantitative terms and to analyse the rules, if any, that regulate the variation in network topologies.

## 8. Applications

Some of the properties of ECAgents technologies act in favour of potential applications and domains, while others rather counteract. Properties that have a positive influence on applications include the capability to evolve behaviours, cooperation patterns and languages from scratch, even when no a priori knowledge has been programmed into the system. An important property that might exhibit a negative influence, include uncertainty with regard to the quality of such evolved properties, and the "black box" nature of evolutionary systems in general, which means that application developers can not always guarantee that a system provides the correct solution to a certain problem.

As argued in Steels *et al.* (2004), taken together these properties make ECAgents suitable for applications where an unpredictable behaviour is in fact valued (like entertainment and art) and less suitable for applications where a consistent and transparent system is more desirable (like security and information retrieval). This is supported by the fact that the only commercial ECAgent-like products we could initially find were in the area of entertainment, in particular computer games. An example is the computer game *The Sims* and *Creatures.* However, while entertainment continues to be a relevant field, it is also desirable to expand the application areas of ECAgents into other domains. As the technology evolves and new insights are gained, this is becoming increasingly possible. In the following, we will discuss some areas where ECAgents technologies could potentially be applied in the future.

### 8.1 Everyday Robot Applications

Everyday robots are becoming a reality. Entertainment robots, service robots and intelligent communicating devices have started to find a niche in our everyday environment (Kaplan 2005a, Ljungblad and Holmquist 2005, Holmquist 2005). ECAgents technologies permit to envision how populations of such devices can autonomously create adapted communication systems. Such agents would neither need a central coordinator nor a prior repertoire of common concepts. Using ECAgents technologies, machines could bootstrap shared categorical systems and associated conventionalized signals (Steels and Belpame 2005, Steels and Kaplan 2002) as well non-verbal communication systems (Hafner and Kaplan 2005). These are crucial features for several kinds of situations:

**Autonomous devices in open environments**: In many real world situations, the number of situations to communicate is not known in advance. ECAgents technologies permit to negotiate shared conventional communication systems for new situations as they occur.

**Autonomous devices in remote environments**: In the case of exploration of remote environments (e.g. planetary exploration), the use of a central coordination system may not be possible. ECAgents technologies permit distributed co-ordination among agents in the population and collective agreement.

**Population of heterogeneous autonomous devices**: In the case where the population is made of devices of several kinds, ECAgents technologies may still permit distributed negotiation of a common communication system because such technologies do not assume any kind of "telepathy" between devices but just external signals.

**Mixed human-robots populations**: Situations where both groups of humans and robots communicate with one another are also tackled by particular ECAgents technologies and have already been demonstrated in large-scale experiments, like the “Talking Heads experiment” (Steels 1999, Kaplan 2001).

In all these situations, the agreement on shared communication systems permit the exchange of “experiences” between machines and from humans to machines. In some ways, such devices become capable of bootstrapping cultural systems of their own, and of accumulating know-how generation after generation. This creates a positive network effect: the more communicating devices there are, the more important this “accumulation” effect is.

It is important to understand that such technologies are not limited to robots in the strict sense, but can be extended to a large variety of smart objects, communicating devices or elements of ambient intelligence. In the near future, robots are likely to blend into a more general class of everyday objects capable of some forms of awareness and equipped with communication functions. ECAgents technologies are likely to be extremely relevant in this new technological ecology (Kaplan 2005b, Ljungblad and Holmquist 2005).

The experimental results mentioned above indicate that this kind of technology can be efficiently used in the various kinds of situations we mentioned. However, there is still a significant gap between research prototypes and actual effective real world applications. Significant progress in the fundamental research issues described in the previous sections are crucial in order to develop real world applications.

## 8.2 Towards realizing collective robotic applications

Progress in ECAgents research might also have a significant impact on collective robotics research (i.e. the attempt to build cooperative robots that can accomplish useful tasks that a single robot could not possibly do [Cao et al., 1997))

The largest amount of work in collective robotics so far has focussed on how to develop a decentralized approach to control a multi-robot system without explicit communication among the robots. The hypothesis is that a decentralized, non-communicating system should scale more easily with the number of robots (Kube and Zhang, 1993). Cooperation here emerges according to a principle where a robot's action is determined or influenced by the consequences of another robot's previous action, similar to the phenomenon of *stigmergy* (Bonabeau et al.,1999). Experiments with robots have successfully exploited this principle to perform a number of tasks such clustering objects (Holland and Melhuish, 1999), coordinated navigation and objects’ transport (Kube and colleagues, 2000; Baldassarre et al., 2005), collective pulling of a stick object (Ijspeert et al., 2001 after Martinoli and Mondada, 1995). An exception to this approach is the work by Gerkey and Mataric (2002), who developed a model where robots cooperate explicitly and with purpose, often through task-related communication. The authors claim that intentional cooperation is better suited than emergent cooperation for tasks that humans would like robots to perform. They presented a novel method of dynamic task allocation for multi-robot systems based on simple auctions to allocate tasks. The results indicated that the system could take into account various changes in the environment.

Some researchers have explored the idea of providing robots with learning capabilities to coordinate robot interactions and improve team performance. However, learning in physically embedded robots is known to be a difficult problem, due to sensor and actuator uncertainty, partial observability of the robot's own environment, and dynamic properties of the environment, especially when multiple learners are involved (Mataric, 1998). Moreover, most learning approaches used in cooperative robotics do not appear to be scalable, because they imply for each learning robot a state-space growth exponential in the number of team members. Touzet (2000) recently proposed to exploit robot awareness, defined as the perception of the locations and actions of other robots, in order to improve cooperative learning. However, in real environments (as opposed to simulations) this approach requires a reliable radio communication and rich information on the state of each robot in the team.

To summarize, the prevailing approach in collective robotics is to provide robots with a set of predefined algorithms and observe the team performance by varying environmental variables and/or team size. There is almost no experimental work addressing the emergence of communication in groups of autonomous robots and how this affects their performance. On the basis of this analysis, we think that there are a number of questions that need to be addressed in the context of communicating robots.

**Environmental situations.** Since explicit communication is computationally and technically expensive for autonomous robots, it is necessary to devise situations where it can provide an advantage with respect to behavioural self-organization described above. For example, communication may bring an advantage when it extends the sensory range of robots by allowing individuals to access information gathered by other robots.

**Communication technology.** Communication can be supported by wireless technologies, such as WI-FI and Bluetooth, or by signals such as sound and body appearance as in nature. Wireless technologies are essentially digital and less affected by environmental conditions (small obstacles, light conditions, acoustic noise), but do not carry directional information. However, the analog nature and directionality of bio-inspired signalling is more suitable for studying the emergence of communication in animals and may provide added information when coupled to wireless technologies. For example, a robot could receive a signal by a neighbouring robot and identify it by a specific pattern of colour LED.

**Cognitive abilities.** The emergence of complex forms of communication might require complex cognitive abilities such us the ability to develop abstract categories, the ability to identify and remember interactions with other robots, etc.

## 8.3 Ubiquitous computing devices

Ubiquitous computing is the notion of distributing computational capabilities into the environment (Weiser 1991). As already mentioned, potential robot applications goe beyond traditional industrial robots and even self-locomoting robots such as the AIBO (Ljungblad and Holmquist 2005, Holmquist 2005). *Tangible interfaces,* i.e. physically embodied computing devices, have been a trend in human-computer interaction in recent years; however, so far this work has been concentrating mostly on tangible input, with output mainly being visual and auditory (Ishii and Ullmer 1997). Forlizzi (2005) introduced the notion of *robotic products,* which are actuated physical computing devices but not necessarily robots in the traditional sense. Such robotic products can take the shape of everyday artefacts such as furniture (Gemperle et al 2003). Much current work in robotics has the potential to dovetail with developments in human-computer interaction, making entirely new forms of actuated input and output possible (Holmquist 2005).

This trend is not specific for ECAgents but provided as a more general observation; however, it clearly shows that the potential of embodied agents goes beyond traditional robots and into everyday environments, where physically embodied ECAgents technologies might be incorporated without being thought of as “robots”. There is a potential commercial market where robotic technologies are combined with evolutionary algorithms and other features of

ECAgents to create robotic products with emergent properties. To develop useful and relevant applications based on ECAgents technologies, questions that need to be addressed include:

**Appropriate technical platforms for ECAgents.** Mobile phones and other hand-held devices are increasingly popular computing platforms. These already provide opportunities for running distributed evolutionary applications, in particular when using ad-hoc wireless networking. However, they currently allow very little physical input or output. Research in this area includes increased emotional input to mobile devices using haptics and gestures (Sundström et al 2005). Commercial activity in this area is increasing, and some mobile devices already come equipped with additional tactile or gestural input, such as Samsung's gesture-recognition phones SPH-S4000 and SCH-S400. Other sensors can easily be incorporated into mobile devices, such as Smart-Its (Holmquist et al 2004). There is a need to explore which platforms are most suitable for future ECAgents applications, in particular with regard to commercial viability (e.g. mass-market devices vs. high-end and expensive hardware) and versatility (applications using specialized hardware vs. one hardware platform for many different applications).

**Communication modalities for ECAgents and humans.** In future ECAgents-based applications, it will be necessary to allow not only agent-to-agent communication, but also agent-to-human communication and even human-to-human communication mediated by agents. In simulations, it is easy to let agents communicate directly, and in the real world, agents might communicate using wireless transmissions. However, such transmissions are invisible to humans, which could be a disadvantage in applications, since it is impossible for users to understand when, if and about what agents are communicating. Furthermore, if agents are to evolve alongside humans, they need communication channels that are equally visible to humans and agents. This might mean audio signals (modulated by frequency, volume, etc.); light signals (modulated by color, intensity, location, etc.) or other shared modalities. The Swarm-Bots (Mondada et al., 2004) and the AIBO$^{TM}$ (Sony Corporation) already have the capability to communicate using light and sound, but these communications are not always understandable for humans. Using such modalities also introduces the problem of registration – for instance, picking up communication using microphones or cameras is much less reliable than a direct radio link. Research is thus needed to devise and test appropriate communication channels for both agents and humans.

**Anticipating user expectations of ECAgents.** As already discussed, the behaviour of ECAgents can be unpredictable an opaque. It can also show surprising signs of "intelligence", as behaviors and languages evolve in anticipation of user actions and changes in the environments. This is an interesting challenge to application design. Results from the field of human-computer interaction indicates that it can be problematic to anthropomorphize intelligent agents too much, as it leads to higher expectations towards the agent's "intelligence", and subsequent disappointment as the system fails to meet those expectations (Shneiderman and Maes 1997, Cassell et al 2000). This is even more pronounced in humanoid robots. If the robot becomes too much like a human, yet not quite human enough, the appearance of the robot enters the "uncanny valley" where its features can make users feel uncomfortable (Mori 1970, Reichard 1978, diSalvo et al 2002). In ECAgents, this problem is further accentuated since the behaviour and skills of an agent may change over time. Therefore, the physical design of embodied agents has to be balanced so that users expect the right kind of interaction and features from the agent. This might mean that the appearance of agents evolves over time or that they are given shapes which remind users of real creatures with similar properties (for instance insects).

## 8.4 P2P Systems

The use of peer-to-peer systems for distribution and sharing of media has exploded in recent years. The legality of such systems has sometimes been questionable, but currently there are several initiatives to introduce digital rights management, payment systems and other measures to make P2P into a legitimate channel for music, movies and other material (e.g. *Snocap* for payment of P2P-distributed materials). While the trend started on stationary workstations on fixed Internet connections, several systems for mobile P2P are also under development (e.g *Simpay* for superdistribution and payment in mobile phone networks).

However, the true potential of peer-to-peer might not lie so much in the actual distribution of files as in the recommendation and discovery of new material. By its very nature, P2P is a shared effort, where all users contribute material from their own sources. They also potentially seed the system with personal taste knowledge and other information that might be converted to meta-data that could help users better find what they what. But the problem is that due to this bottom-up approach distribution there are no shared descriptive languages or ontologies. This is where properties of ECAgents technologies are very relevant. If techniques from the ECAgents project were to be applied to P2P networks, they might help to evolve automatic classifications and categories to guide users to interesting material. This could lead to more efficient recommendations and, coupled with the appropriate payment models and digital rights management, to new business opportunities for the European content and telecom industries. For an example of an application of this ideas in the area of mobile music sharing, see Håkansson et al. (2005) and Jacobsson et al. (2005, 2006).

To take this work further, the following areas are of particular interest:

**Evolution of shared descriptive languages.** In the Talking Heads experiment (Steels 1999), agents evolved a common language to denote features in their environment. Such a language could also be evolved to describe files in a P2P system. This would help agents to negotiate among themselves and automatically recommend content to users. Furthermore, it would be particularly interesting if this language was shared not just among agents, but also understood by the human users. In this case the "language" might be represented more abstractly, for instance in colors, visual patterns, positions on a scale, or other ways that would be readable for a human user.

**New forms of content.** Music has been the most popular media for P2P applications, but theatrical movies and television shows are also being extensively distributed. There are possibilities for evolving communication languages and ontologies for many other forms of content including scientific papers, digital photographs, current news items, blog entries, and so on. If the methods developed for evolution of languages is only based on sharing and usage patterns (like in the Push!Music system, Jacobsson et al. 2006) they might in some cases be directly converted to other media forms with minimal changes. More likely, however, is that new techniques will have to be developed for new content types, since the usage and sharing patterns are different. For instance, whereas a music file is used over a long period and not always shared frequently, a hot news item or blog entry item might be shared among many users in a very short time period and then not seen again.

**Calibrating user acceptance.** Automatic recommendations are always somewhat problematic, as they are never 100% percent correct. Evaluations show that humans give more exact recommendations; on the other hand, automatic recommendations introduce more surprises (Swearingen and Sinha 2001). If automatic recommendations are to large a part of a system, users might even find that control is taken away from them and they are unable to influence their own media consumption. This is particularly relevant for ECAgent technologies, as the workings behind a certain evolved system are not transparent. Therefore, the best way to adequately test the success of ECAgents-based recommendations in P2P systems is through user testing. Such tests should preferably be in naturalistic settings rather than in a lab to ensure that the results are realistic.

## 9. Conclusions

The computational and robotic synthesis of language evolution is emerging as a new exciting field of research. The objective is to come up with precise operational models of how communities of agents, equipped with a cognitive apparatus, a sensory-motor system, and a body, can develop a shared grounded communication system. Such system might have similar characteristics to animal or human language.

In this paper we have tried to explicit the theoretical foundations of this research area, the methods that can be used to achieve its objectives, and the most promising research directions. Moreover, we have discussed some potential application areas.

With respect to the theoretical foundations, we pointed out that one crucial aspect is the realization of the complex adaptive system nature of behaviour in general terms and of communication behaviours in particular. The given nature of behaviour and communication behaviour implies the need to rely on complex system methods and techniques to model this phenomenon and to synthesize it in artificial embodied agents. We view language and communication as a living system, that self-organize and evolve through the collective dynamics of interacting agents and that consists of a multi-scale phenomenon involving bottom-up and top-down interactions between different levels of organization.

In addition to explicit the intrinsic nature of communication and language from a general point view, we pointed out the need to further characterize this phenomenon by identifying its universal properties and the elements that represent crucial pre-requisites for the emergence of communication and language. More specifically we tried to make the hypothesis behind the experimental works in this area explicit and we indicated the methods that can be used to identify crucial pre-requisites for the emergence of communication.

With respect to the methods that can be used to develop ECAgents we discussed collective intelligence methods, evolutionary methods, and social learning methods. All these methods share important properties that make them suitable to study the emergence of communication in embodied agents. In particular, in all these methods: (a) a population of agents and a set of rules that govern the agent/agent and agent/environmental interactions is defined, (b) the agents are allowed to “live” by interacting with the environment, and (c) the collective behaviour emerging from the interactions and eventually the development of a shared communication system is observed. Moreover, all methods rely on positive and negative feedback mechanisms, and random variation and/or noise. As example of positive feedback mechanisms, consider the collective intelligent model described in section 3 in which a given concentration of a certain pheromone type tend to attract especially ants that produce that pheromone type, or the case of the social learning model described in section 5 in which differences in the frequency of use of synonymous word within a population tend to be amplified by agents’ tendency to adopt words that led to successful communications). As an example of random variation and/or noise, consider the effect of mutation operators in evolutionary models described in section 4 or the effect of the different perceptual scene experienced by two interacting agents in social learning methods described in section 5).

However, these methods also differs in important respects that make them suitable to different circumstances. More specifically, evolutionary and social learning methods differ from collective intelligence methods since the former methods includes mechanisms for modifying the free parameters that regulate how agents interact with the physical and social environment on the basis of a measure of success or failure while the latter method does not.

Moreover, social learning methods are based on stereotyped social interactions (i.e. language games) that, by providing a measure of success or failure in the short term, allow agents to modify their free parameters after each social interaction. In evolutionary learning methods, instead, feedback on agents' success of failure are available only in the long term. These differences implies that, in general terms, collective intelligence methods are especially suitable to study the transition from non-communicating to communicating agents and the role of self-organization processes that might lead to the spontaneous emergence of communication abilities. Evolutionary methods are especially suitable to study the emergence and the evolution of communication in population of agents that interact freely between themselves and with the environment and to study the co-adaptation of agents' motor and communication abilities. Social learning methods are especially suitable to study the emergence of complex communication systems and human language in population of agents that already posses complex cognitive, behavioural, and social abilities.

We believe that the study of the emergence and evolution of communication in such a broad sense (from simple coordinated interaction to complex human-language like communication abilities) through the use of different and complementary methods combined with quantitative analysis of existing natural and artificial communication systems will lead to the identification of universals properties and to a cross-fertilization of researches often conducted in isolation by different scientific communities.

## References


Ackley D.H. and Littman, M. L. (1994) Altruism in the evolution of communication. In R.A. Brooks & P. Maes (Eds.), Artificial Life IV: Proceedings of the International Workshop on the Synthesis and Simulation of Living Systems, Cambridge, MA: MIT Press.

Albert R., Jeong, H. and Barabási, A. (1999) Diameter of the World Wide Web Nature 401, 130.

Albert, R., Jeong, H. and Barabási, A. (2000) Error and attack tolerance of complex networks, Nature 406, 378-382.

Anderson, C. (2002) Self-organization in relation to several similar concepts: Are the boundaries to self-organization indistinct? *Biological Bulletin*, **202**, 247-255.

Avital, E. & Jablonka, E. (2000) *Animal Traditions: Behavioural Inheritance in Evolution.* Cambridge, Cambridge University Press.

Baldassarre G., Nolfi S. & Parisi D. (2002). Evolving mobile robots able to display collective behaviours. In C.K. Hemelrijk (Ed.), *Proceedings of the International Workshop on Self-Organisation and Evolution of Social Behaviour*. Zurich, Switzerland: Swiss Federal Institute of Technology.

Baldassarre G., Nolfi S., Parisi D. (2003). Evolving mobile robots able to display collective behaviour. *Artificial Life*, 9: 255-267.

Baldassarre G., Trianni V., Bonani M., Mondada F., Dorigo M., Nolfi S. (2005). Coordinated motion in groups of physically connected robots. *Technical Report.* Rome: Institute of Cognitive Science and Technology, CNR.

Baldwin J.M. (1896). A new factor in evolution. *American Naturalist*, 30: 441-451.

Barabási, A-L. (2003). *Linked: How Everything Is Connected to Everything Else and What It Means for Business, Science, and Everyday Life* , Plume Books.

Baronchelli, A. Felici, M., Caglioti, E., Loreto, V. & Steels, L. (2005). Sharp Transition towards Shared Vocabularies in Multi-Agent Systems. physics/0509075.

Bollobás, B. (1998) *Modern graph theory*, Springer, New York.

Bonabeau E. M. Dorigo & G. Theraulaz (1999). *Swarm Intelligence: From Natural to Artificial Systems.* New York, NY: Oxford University Press.

Camazine S., Deneubourg J.L., Franks N., Sneyd J., Bonabeau E.& Theraulaz G. (2001) Self-Organization in Biological Systems Princeton University Press .

Cangelosi A. & Parisi D. (1998). The emergence of a ‘language’ in an evolving population of neural networks. *Connection Science* 10(2): 83-97.

Cangelosi A. & Parisi D. (Eds.) (2002). *Simulating the Evolution of Language*. London: Springer-Verlag

Cao Y.U., Fukuaga A. and Kahng A. (1997), Cooperative Mobile Robotics: Antecedent and Directions, *Autonomous Robots* 4, 7-27.

Cassell, Justine, Joseph Sullivan, Scott Prevost and Elizabeth Churchill. *Embodied Conversational Agents.* MIT Press, 2000.

Cattuto C., Loreto V., Pietronero L., (2005a), Semiotic Dynamics and Collaborative Tagging, (submitted to *Nature*).

Cattuto C., Loreto V., Pietronero L., (2005b), Semiotic Dynamics in Online Social Communities, submitted to *Eur. Phys. J. C.*

Chomsky, N. (1995). *The Minimalist Program* Cambridge, Mass: MIT. Press.

Christiansen M.H., Kirby S. (2003). Language Evolution. Oxford, U.K.: Oxford University Press.

Corominas, B. and Solé, R. V. (2005) Network topology and self-consistency in language games Journal of Theoretical Biology (in press).

De Beule J. and Steels L. (2005). Hierarchy in Fluid Construction Grammar. In: Furbach, Ulrich (ed.), *Proceedings of the 28th Annual German Conference on AI, KI 2005, Lecture Notes in Artificial Intelligence (vol. 3698)*, p. 1-15, Berlin Heidelberg: Springer-Verlag.

Deneubourg J.L., Pasteels J.M. & Verhaeghe J.C. (1983). Probabilistic behaviour in ants: a strategy of errors? Journal of theoretical Biology, 105, 259-271.

Detrain C., Deneubourg J.L. &. Pasteels J.M. (1999) Decision-making in foraging by social insects In: Information processing in social insects. Detrain C., Deneubourg J.L. et J.M. Pasteels (eds), Birkhauser Verlag, Basel, 331-354

Di Paolo E.A. (2000). Behavioural coordination, structural congruence and entrainment in a simulation of acoustically coupled agents. Adaptive Behaviour 8:1. 25-46.

DiSalvo, CF, F Gemperle, J Forlizzi, SB Kiesler. All robots are not created equal: the design and perception of humanoid robot heads. *Symposium on Designing Interactive Systems,* ACM Press, 2002

Dorigo M., Trianni V., Sahin E., Gross R., Labella T. H., Baldassarre G., Nolfi S., Deneubourg J.-L., Mondada F., Floreano D., Gambardella L. M. (2004). Evolving self-organizing behaviors for a swarmbot. Autonomous Robots, 17(2-3): 223-245.

Dorigo, M. & Stützle, Th. (2004). Ant Colony Optimization, MIT Press.

Dorogovtsev, S.N. and Mendes, J.F.F. (2003) *Evolution of Networks: from biological nets to the Internet and WWW* Oxford U. Press, Oxford

Feller, W. (1971). *An Introduction to Probability Theory and its Applications, Volume 2*. New York: Wiley.

Ferrer i Cancho, R. (2003). Language: universals, principles and origins. PhD Thesis.

Ferrer i Cancho, R. and Sole, R. V. (2003) Least effort and the origins of scaling in human language. PNAS, 100:788-791.

Ferrer i Cancho, R. and Servedio, V. D. P. (2005). *Can simple models explain Zipf's law for all exponents?* , Glottometrics 11, 1.

Ferrer i Cancho, R. and Solé, R. V. (2001) The Small-World of Human Language Proceedings of the Royal Society of London B 268, 2261-2266

Ferrer i Cancho, R.; Koehler, R. and Solé, R. V. (2004) Patterns in syntactic dependency networks Physical Review E 69, 32767.

Forlizzi, J. Robotic products to assist the aging population. *ACM <interactions>,* Volume 12 , Issue 2, March+April 2005.

Funes P., Orme B. & Bonabeau E. (2003). Evolving emergent group behaviours for simple humans agents. In P. Dittrich & J. T. Kim (Eds.) Proceedings of the Seven European Conference on Artificial Life .

Gemperle, Francine, Carl DiSalvo, Jodi Forlizzi, & Willy Yonkers. The Hug: a New Form for Communication. *Proceedings of the 2003 conference on Designing for user experiences*, ACM Press.

Gerkey, B.P. and Mataric, M. (2002), Sold!: Auction methods for multi-robot coordination, *IEEE Transactions on Robotics and Automation*, Special Issue on Multi-robot Systems, 2002

Golder, S. and Huberman, B. A. (2005), *The Structure of Collaborative Tagging Systems* (http://www.hpl.hp.com/research/idl/papers/tags/)

Grafen A. (1990). Biological signals as handicaps. *Journal of Theoretical Biology*, 174: 217-222.

Greenberg J.H. (1963) *Universals of Language*. London: MIT Press.

Hafner and Kaplan 2005 Hafner, V. V., Kaplan, F. (2005), Learning to interpret pointing gestures: experiments with four-legged autonomous robots, in Wermter, S., Palm, G., Elshaw, M. (Eds.) Biomimetic Neural Learning for Intelligent Robots. Intelligent Systems, Cognitive Robotics, and Neuroscience. Series: Lecture Notes in Computer Science. Subseries: Lecture Notes in Artificial Intelligence, Vol. 3575. Springer Verlag

Håkansson M., Jacobsson M. and Holmquist L.E. (2005) Designing a Mobile Music Sharing System Based on Emergent Properties. In *Proceedings of AMT 2005, the 3rd International Conference on Active Media Technology*, May 2005, Takamatsu, Japan.

Haken, H. (1988). Information and Self-organization: A Macroscopic Approach to complex Systems, Berlin: Springer-Verlag.

Holland O.E. and Melhuish C. (1999), Stigmergy, self-organisation, and sorting in collective robotics, *Artificial Life*, 5:2, 173-202.

Holmquist, L.E. The Robots are Coming, *ACM <interactions>*, May-June 2005, ACM Press.

Iizuka H. and Ikegami T. (2003a). Adaptive Coupling and Intersubjectivity in Simulated Turn-Taking Behaviours. In Banzahf et al. (Eds.), Proceedings of ECAL 03, Dortmund: Springer Verlag.

Iizuka H. and Ikegami T. (2003b). Simulating Turn-taking Behaviors with Coupled Dynamical Recognizers. In R.K. Standish, M.A. Bedau and H.A. Abbass (Eds.), MIT, Proceedings of Artificial Life VIII, Cambridge, MA: MIT Press.

Ijspeert A.J., Martinoli A., Billard A., Gambardella, L.M. (2001), Collaboration through the exploitation of local interactions in autonomous collective robotics: the stick pulling experiment, *Autonomous Robots*, Vol. 11:2, 149-171.

Ishii, H. and Ullmer, B., Tangible Bits: Towards Seamless Interfaces between People, Bits and Atoms, in *Proceedings of Conference on Human Factors in Computing Systems CHI '97)*, (Atlanta, March 1997), ACM Press, pp. 234-241.

Deneubourg J.L, S. Goss, J.M. Pasteels, D. Fresneau, & J.P. Lachaud (1987). Self-organization mechanisms in ant societies (II): learning in foraging and division of labor. In From individual to collective behavior in social insects, Eds J.M. Pasteels & J.L. Deneubourg. Experientia Supplementum, 54, Birkhaüser, Bâle, 177-196.

Jablonka, E. & Szathmáry, E. (1995). The evolution of information storage and heredity. *Trends Ecol. Evol.* **10**, 206-211.

Jacobsson M., Rost M. and Holmquist L.E. (2006) When Media Gets Wise: Collaborative Filtering with Mobile Media Agents. To appear in *Proceedings of IUI 2006, the 10th International Conference on Intelligent User Interfaces*, February 2006, Sydney, Australia.

Jacobsson M., Rost M., Håkansson M. and Holmquist L.E. (2005) Push!Music: Intelligent Music Sharing on Mobile Devices. In Adjunct Proceedings of UbiComp 2005, the 7th International Conference on Ubiquitous Computing, September 2005, Tokyo, Japan.

Jinyun; K. and YaoYao, (2005) A study of language acquisition from a network perspective. (Unpublished work)

Kaplan, F. (2001). La naissance d'une langue chez les robots, Hermes, 2005.

Kaplan, F. (2005a). Les machines apprivoisées, Vuibert, 2005

Kaplan, F. (2005b). Everyday robotics: robots as everyday objects, Proceedings of Soc-Eusai 2005, Grenoble, France.

Keijzer F. (2001). *Representation and behaviour*. Cambridge, MA: MIT Press.

Kirby S. (2002). Natural Language from Artificial Life. *Artificial Life*, 8(2):185-215.

Krieger M., Billeter, J-B., and Keller, L. (2000), Ant-like task allocation and recruitment in cooperative robots, *Nature* 406, 992-995.

Kube, C.R. and Bonabeau, E. (2000), Cooperative transport by ants and robots, *Robotics and Autonomous Systems*,Vol. 30, Issue 1-2, 85-101.
Kube, C.R. and Zhang, H. (1993) Collective Robotics: From Social Insects to Robots, *Adaptive Behavior*, 2(2), 189-219.
Lenoir A., Fresneau D., Errard C.& Hefetz A. (1999) - Individual and colonial identity in ants: the emergence of the social representation concept. Dans *Information processing in social insects*, C. Detrain, J.L. Deneubourg et J.M. Pasteels eds, Birkhauser, Bâle, 219-237.
Liljeros F, Edling CR, Amaral LA, Stanley HE, Aberg Y. (2001) The web of human sexual contacts. Nature; 411:907-908.
Ljungblad and Holmquist 2005 Designing robot applications for everyday use, Proceedings of Soc-Eusai, Grenoble, France.
Magnenat, S., Floreano, D. and Keller, L. (2005) Effects of levels of selection and genetic relatedness on emergence of communication. *Proceedings of the National Academy of Sciences, USA*. Submitted.
Marocco D. & Nolfi S. (2004). Emergence of Communication in Embodied Agents: Co-Adapting Communicative and Non-Communicative Behaviours. In A. Cangelosi et al. (eds.), Modeling language, cognition and action: Proceedings of the 9th Neural Computation and Psychology Workshop. Singapore: World Scientific.
Marocco D. & Nolfi S. (2005). Emergence of Communication in Embodied Agents Evolved for the Ability to Solve a Collective Navigation Problem. Technical Report. Roma, Italy: Institute of Cognitive Science and Technology, CNR.
Marocco D., Cangelosi A., Nolfi S. (2003), The emergence of communication in evolutionary robots. Philosophical Transactions of the Royal Society London - A, 361: 2397-2421.
Martinoli, A. and Mondada, F. (1995), Collective and cooperative group behaviours: Biologically inspired experiments in robotics. In O. Khatib & J.K.Salisbury (Eds.), *Proc. of the Fourth Int. Symp. on Experimental Robotics*, 3-10.
Mataric, M. (1998), Coordination and Learning in Multi-Robot Systems, *IEEE Intelligent Systems*, Mar/Apr., 6-8.
Maynard Smith J. (1991). Honest Signalling: The Philip Sidney Game. *Animal Behaviour*, 42: 1034-1035.
Maynard Smith, J. & Harper, D. (2003) *Animal Signals.* Oxford University Press, Oxford.
Maynard Smith, J. & Szathmáry, E. (1995) *Major Transitions in Evolution*, Freeman: Oxford.
Melcuck, I. (1989). *Dependency grammar: Theory and practice*. New York, University of New York.
Mika, P. (2005). Ontologies are us: A unified model of social networks and semantics. In Proc. Of ISWC-2005, Galway, Ireland, Springer (2005).
Millor J., Amé J.M., Halloy J. & Deneubourg J.L. (2006) Individual discrimination capability and collective decision-making. Journal of theoretical Biology (in press).
Milo, R.; Shen-Orr, S.; Itzkovitz, S.; Kashtan, N.; Chklovskii, D.and Alon, U. (2002). Network Motifs: Simple Building Blocks of Complex Networks Science, 298:824-827.
Mirolli M. & Parisi D. (2005a) How can we explain the emergence of a language which benefits the hearer but not the speaker?, Connection Science, 17(3-4): 325-341.
Mirolli M. & Parisi D. (2005b) Production bias and genetic convergence in the evolution of communication. *Technical Report*. Roma, Italy: Institute of Cognitive Science and Technology, CNR.
Mirolli, M. & Parisi, D. (2005c) Evolutionary Dynamics in the Emergence of Communication. *Technical Report*. Roma, Italy: Institute of Cognitive Science and Technology, CNR.
Mondada F., Pettinaro G., Guignard A., Kwee I., Floreano D., Denebourg J-L, Nolfi S., Gambardella L.M., Dorigo M. (2004). Swarm-bot: A new distributed robotic concept. Autonomous Robots. vol. 17 (2-3), pp. 193-221.
Mori, M. Bukimi no tani [the uncanny valley], *Energy*, 7:33-35, 1970.
Nazzi T., Gopnik A. (2001). Linguistic and cognitive abilities in infancy: When does language become a tool for categorization? Cognition, 80:303-312.
Newman, M. E. J. (2005). Power laws, Pareto distributions and Zipf's law. Contemporary Physics 46, 323-351.

Nicolis S.C., Detrain C., Demolin D. & J.L. Deneubourg.(2003) Optimality of collective choices: a stochastic approach. Bulletin of Mathematical Biology,65, 795-808.

Nolfi S. & Floreano D. (1999). Learning and Evolution, Autonomous Robots, 7(1): 89-113.

Nolfi S. & Floreano D. (2000). *Evolutionary Robotics: The Biology, Intelligence, and Technology of Self-Organizing Machines.* Cambridge, MA: MIT Press/Bradford Books.

Nolfi S. (2002). Evolving robots able to self-localize in the environment: The importance of viewing cognition as the result of processes occurring at different time scales. Connection Science (14) 3:231-244.

Nolfi S. (2005). Emergence of communication in embodied agents: co-adapting communicative and non-communicative behaviours. *Connection Science*, (17) 3-4: 231-248.

Nolfi S. (in press). Behaviour as a complex adaptive system: On the role of self-organization in the development of individual and collective behaviour. *ComPlexUs*.

Nowak, M. A. and Krakauer, D. C. (1999). The evolution of language. Proc. Natl. Acad. Sci. USA 96, 8028.

Oliphant M. (1997). Formal Approaches to Innate and Learned Communication: Laying the Foundation for Language. *PhD thesis*. San Diego, USA: Department of Cognitive Science, University of California, San Diego.

Oliphant, M. (1996) The dilemma of Saussurean communication. Biosystems, 37(1-2): 31-38.

Oudeyer, Pierre-Yves (2005). The Self-Organization of Speech Sounds. *Journal of Theoretical Biology, 233(3)*, p. 435-449.

Parker, L.E. (1998), ALLIANCE: An architecture for fault tolerant multirobot cooperation, *IEEE Trans. on Robotics and Automation* 14 (2), 220-240.

Perez-Uribe A., Floreano D & Keller L. (2003). Effects of Group Composition and Level of Selection in the Evolution of Cooperation in Artificial Ants. *Lecture Notes in Computer Science*, 2801:128-137.

Pfeifer R., Scheier C. (1999). Understanding Intelligence. Cambridge, MA: MIT Press.

Quinn M., Smith L., Mayley G. & Husbands P. (2003) Evolving controllers for a homogeneous system of physical robots: Structured cooperation with minimal sensors. *Philosophical Transactions of the Royal Society of London, Series A: Mathematical, Physical and Engineering Sciences* 361:2321-2344.

R. Beckers, O.E. Holland & J.L.Deneubourg (1994). From local actions to global tasks: stigmergy and collective robotics. In Proceedings of ALIFE IV, Eds R.A. Brooks & P. Maes, MIT Press, Cambridge (Mass).

Radford, A. (1990) Syntactic Theory and the Acquisition of English Syntax: the nature of early child grammars of English, Blackwell. Oxford.

Reichard,J. *Robots: Fact, Fiction, and Prediction.* Penguin Books, 1978.

Shneiderman, B. & Maes, P. (1997). Direct manipulation vs. interface agents: Excerpts from debates. *Interactions,* (Dec. 1997), 42-61.

Sigman, M. and Cecchi, GA (2002) Global organization of the Wordnet lexicon. PNAS, 99(3):1742-1747. 2001.

Solé, R, V. (2005) Syntax for free? Nature 434, 289.

Solé, R, V.; Corominas Murtra, B.; Valverde, S. and Steels, L. (2005) Language networks: their structure, function and evolution. Santa Fe Institute Working Paper 05-12-042 (submitted to Trends in Cognitive Sciences)

Solé, R, V.; Ferrer i Cancho, R.; Montoya, J. M. and Valverde, S. (2002) Selection, Tinkering and Emergence in Complex Networks Complexity 8, 20-33.

Solé, R. V. (2006) Scaling Laws in Human Language. In: *Power Laws in the Social Sciences*. C. Cioffi-Revilla (editor). Cambridge U. Press. To appear.

Solé, R. V. and Pastor-Satorras, R. (2002) Complex networks in genomics and proteomics. In: Handbook of Graphs and Networks, S. Bornholdt and H.G. Schuster, eds. pp. 147-169. John Wiley-VCH. Berlin.

Steels L. (1995). A Self-Organizing Spatial Vocabulary. *Artificial Life 2(3)*, p. 319-332.

Steels L. (1996). Perceptually Grounded Meaning Creation. In: Tokoro, M. (ed.), *Proceedings of the Second International Conference on Multi-Agent Systems, ICMAS-96*, p. 338-344, Menclo Park, CA: AAAI Press.

Steels L. (1997). Constructing and Sharing Perceptual Distinctions. In: Van Someren, Maarten and Widmer, Gerhard (eds.), *Proceedings of the 9th European Conference on Machine Learning*, p. 4-13, Berlin: Springer-Verlag.
Steels L. (1999). *The Talking heads experiment. Volume 1. Words and meanings.* Laboratorium, Antwerpen. Limited Pre-edition.
Steels L. (2000). Language as a Complex Adaptive System. In Schoenauer, M. (Ed.), *Proceedings of PPSN VI: Lecture notes in Computer Science*, pages 17-26, Berlin: Springer-Verlag.
Steels L. (2002). Simulating the Evolution of a Grammar for Case. Presented at the 4th International Conference on the Evolution of Language, Harvard University.
Steels L. (2003). Evolving grounded communication for robots. *Trends in Cognitive Science*. 7(7): 308-312.
Steels L. (2005a). The Recruitment Theory of the Origins of Language. Submitted.
Steels L. (2005b). The Role of Construction Grammar in Fluid Language Grounding. Submitted.
Steels L. and Belpaeme T. (2005). Coordinating Perceptually Grounded Categories Through Language. A Case Study for Colour. *Behavioral and Brain Sciences, 24(6).*
Steels L. and Bleys J. (in press). Planning What to Say: Second Order Semantics for Fluid Construction Grammars. In: Bugardin Diz, A. and J. Santos Reyes, (eds.), *Proceedings of CAEPIA '05. Lecture Notes in AI*, Berlin: Springer-Verlag.
Steels L. and Kaplan F. (1998). Stochasticity as a Source of Innovation in Language Games. In: C. Adami, R. Belew, H. Kitano and C. Taylor (eds.), *Artificial Life VI*, Los Angeles: MIT Press.
Steels L. and Kaplan F. (2002). Bootstrapping grounded word semantics. In Briscoe,T. Linguistic evolution through language acquisition: formal and computational models, Cambdrige University Press.
Steels L. and Kaplan F. (2002). Spontaneous lexicon change, in Proceedings of COLING-ACL, 1243 (Montreal).
Steels L. and Nolfi S. (Eds.) (2004). *The White Paper on Target Problems and Grand Challenges for Evolving Embodied and Communicating Agents*. Unpublished Volume. Paris, France: SONY-CSL. Available from: http://ecagents.org/imgs/D1-1.pdf.
Steels L., De Beule J. and Neubauer N. (2005). Linking in Fluid Construction Grammars. *Proceedings of BNAIC.* Transactions of the Belgian Royal Society of Arts and Sciences, Brussels.
Steels L., Loetzsch M. and Bergen B. (2005). Why Humans Mark Perspective. Submitted.
Steyvers, M. and Tenenbaum, J. B. (2005) The large-scale structure of semantic networks: statistical analyses and a model of semantic growth. Cognitive Science, 29(1).
Sundström, Petra, Anna Ståhl, and Kristina Höök (2005) eMoto - Affectively Involving both Body and Mind, *Extended abstract CHI'05,* Portland, Oregon, USA.
Swearingen, K, R Sinha. Beyond Algorithms: An HCI Perspective on Recommender Systems. *ACM SIGIR 2001 Workshop on Recommender Systems,* ACM Press, 2001.
Szathmáry, E. & Maynard Smith. J. (1995). The major evolutionary transitions. *Nature* **374,** 227-232.
Touzet, C. (2000). Robot awareness in cooperative mobile robot learning, *Autonomous Robots* 8 (1), 87-97.
Valverde, S. and Solé, R.V. (2004) Internet's Critical Path Horizon. European Physics Journal 38(2), 245-252.
Valverde, S., Theraulaz, G., Gautrais, J., Fourcassié, V. and Solé, R, V. (2006). Emergent behaviour in agent networks: self-organization in wasp and open source communities. IEEE Intelligent Systems (to appear 2006)
Valverde, S.; Ferrer, R. and Solé R. V. (2002) Scale free networks from optimal design Europhysics Letters 60, 512-517.
Van Looveren J. (1999). Multiple Word Naming Games. In: *Proceedings of the 11th Belgium-Netherlands Conference on Artificial Intelligence*, Maastricht: Universiteit Maastricht.
Vicksek T. (2001). A question of scale. Nature 411, 421.
Vicksek T. (2002). The Bigger Picture, Nature 418, 131.
Waddington C.H. (1942). Canalization of development and the inheritance of acquired characters. *Nature*, 150: 563-565.
Wagner G.P. & Altenberg L. (1996) Complex adaptations and the evolution of evolvability. *Evolution*, 50:967-976.

Wagner K., Reggia J.A., Uriagereka J. & Wilkinson G.S. (2003). Progress in the simulation of emergent communication and language. *Adaptive Behavior*, 11(1):37-69.

Watts D.J. and Strogatz S.H. (1998) Collective dynamics of 'small-world' networks, Nature, 393:440-442.

Waxman S.R., Markow D.B. (1995). Words as invitations to form Categories: Evidence from 12- to 13-month-old infants. Cognitive Psychology, 29:257-302

Weiser, M. The Computer for the 21st Century. Scientific American, 265(3):94–104, September 1991.

Wilson E.O. & Hölldobler B. (2005). Eusociality: Origin and consequences. *Proc. Natl. Acad. Sci. USA.* 102: 13367-13371.

Wyatt T.D. (2003). Pheromones and Animal Behaviour. Oxford University Press.