

Producer Biases and Kin Selection in the Evolution of Communication

How the Phylogenetic and the Adaptive Problems of Communication can be solved

Marco Mirolli and Domenico Parisi

Institute of Cognitive Sciences and Technologies, CNR, Roma, Italy
{marco.mirolli,domenico.pariis}@istc.cnr.it

Abstract. The evolution of communication requires the co-evolution of two abilities: the ability of sending useful signals and the ability of reacting appropriately to perceived signals. This fact poses two related but distinct problems, which are often confused the one with the other: (1) the phylogenetic problem regarding how can communication evolve if the two traits that are necessary for its emergence are complementary and seem to require each other for providing reproductive advantages; (2) the adaptive problem regarding how can communication systems that do not advantage both signallers and receivers in the same way emerge, given their altruistic character. Here we clarify the distinction, and provide some insights on how these problems can be solved in both real and artificial systems by reporting experiments on the evolution of artificial agents that have to evolve a simple food-call communication system. Our experiments show that (1) the phylogenetic problem can be solved thanks to the presence of producer biases that make agents spontaneously produce useful signals, an idea that is complementary to the well-known ‘receiver bias’ hypothesis found in the biological literature, and (2) the adaptive problem can be solved by having agents communicate preferentially among kin, as predicted by kin selection theory. We discuss these results with respect both to the scientific understanding of the evolution of communication and to the design of embodied and communicating artificial agents.

1 Introduction

The emergence of communication requires the co-evolution of both good speakers and good hearers. Good speakers are individuals that produce signals that systematically co-vary with situations that are relevant for survival and reproduction. Good hearers are individuals that react to signals in ways that are appropriate to the situations with which signals systematically co-vary. If hearers do not respond appropriately to the signals produced by the speakers, there is no reason for speakers to emit appropriate signals. In other words, good speakers without good hearers are useless because signals are not understood and reacted to appropriately. Conversely, if speakers do not produce the appropriate signals

in the appropriate circumstances, there are no useful signals for hearers to understand. In other words, good hearers without good speakers are meaningless because one cannot react appropriately to signals that do not co-vary with any relevant feature of the world.

The interactive nature of communication renders its emergence problematic for two related but independent reasons. The first reason has to do with the historical paths that can lead to a communication system: how might hearers appear if speakers are not already there sending useful signals? And, conversely, how might speakers appear if hearers are not already there understanding what they can say? More concretely: how can communication evolve if the traits necessary for its emergence - namely, good speaking and good hearing capabilities - taken in isolation are adaptively neutral in that they do not, by themselves, increase the reproductive chances of the individuals that possess them?

The second reason has to do with the adaptive advantage of communication itself. As a matter of fact, not all communication systems seem to provide (the same amount of) advantage to both actors of communication, that is, speakers and hearers. Generally speaking, with respect to adaptiveness, communicative systems can be broadly classified in three categories, according to who is advantaged by communicative interactions: only the speaker, only the hearer, or both. Communicative interactions that advantage both the actors of communication are typically those in which the speaker and the hearer have common interests and have to coordinate their behaviours (like in group hunting). Communicative interactions that advantage only the speaker are instances of manipulation (like in commands or in aggressive displays made for convincing the hearer to leave some contended resource to the speaker). Finally, communicative interactions that advantage only the hearer are typically those in which the speaker informs the hearer about some feature of the environment which is relevant for the latter (like in alarm or food calls). If communication clearly advantages both the actors of communicative interactions, its evolution poses only the above-mentioned co-evolutionary problem. But in those communication systems in which only one actor of communication is advantaged the problem of co-evolution generates an adaptive problem due to the altruistic character of these kinds of communication systems. Why should a speaker produce good signals if only the hearer is benefited by communication? Conversely, why should a hearer respond appropriately to signals if this provides benefits only to the speaker? How could such communication systems evolve?

According to Nikolas Tinbergen, one of the founders of modern ethology, ethological questions can be classified into four main categories (Tinbergen, 1963): (1) mechanistic (what are the mechanisms which determine the behavior?), (2) ontogenetic (what are the genetic and environmental factors that contribute to the behavior's development?), (3) functional (what are the effects of a given behavior on survival and reproduction?), and (4) phylogenetic (how did the behavior evolve during the history of the species?). The two aforementioned problems of communication are clearly related to two different Tinbergen's questions: the problem of co-evolution between signallers and receivers is a typical

phylogenetic problem, while the adaptive problem related to (apparently) altruistic communication system is related to the functional question.

The main goal of this chapter is to clarify the distinction between these two problems, and to provide some insights on how these problems can be solved in both real and artificial organisms. In order to do that, we address both problems within the same simulative set-up, that is constituted by a population of artificial agents which have to evolve a simple food-call communication system. The remainder of this chapter is structured as follows. In section 2 we first discuss the general approach that evolutionary biologists typically take with respect to communication and then we introduce in more details the adaptive and the phylogenetic problems. In section 3 we describe the experimental set-up with which we address these problems. In 4 we present the results of the basic simulation and discuss our hypotheses concerning the factors that can affect the emergence of communication. The next two sections are devoted to the test of the proposed hypotheses: section 5 with respect to the problem of altruism, section 6 with respect to the co-evolution of speakers and hearers. Finally, in section 7 we discuss the relevance of the presented work with respect to both the scientific understanding of communication and to the design of embodied and communicating artificial agents.

2 Two problems in the evolution of communication

2.1 The biological literature and the manipulation bias

The recent biological literature on communication has focused mostly on manipulative communication, due to the controversial but deeply influential work of Dawkins and Krebs (1978). Those authors gave a definition of communication which *identified* it with manipulative behavior. According to Dawkins and Krebs communication occurs “when an animal, the actor, does something which appears to be the result of selection to influence the sense organs of another animal, the reactor, so that the reactor’s behaviour changes to the advantage of the actor”. Though Dawkins and Krebs’ framework is the most comprehensive theoretical framework on communication that has been proposed so far, and though it has produced a number of interesting empirical predictions (Hauser, 1996), the identification of communication with manipulation seems to be unjustified both theoretically and empirically. The theoretical point is that the problem of the evolution of communication is symmetrical: pure adaptive reasoning implies that just as the signaller must get an advantage for sending information, the receiver must also get an advantage in order to respond to signals. Empirically, just as there are communication systems in which the result of the communicative interaction seems to benefit only the signaller, such as the aggressive displays produced when competing for some resources or begging calls produced by newborns for receiving extra feeding from parents, there are also communication systems which seem to provide benefit to both the actors of communication, such as mating calls or signals emitted during cooperative endeavors like group hunting, and communication systems which seem to benefit only the receiver,

such as the food or alarm calls produced by many species. Of course, we may discover that in these cases there is an advantage also for the sender: for example, many fleeing preys may confound the predator, or the sender may increase its social reputation for sending useful signals. But this is something which must, eventually, be *empirically demonstrated* for each particular case, not something which we can *establish a priori*, by including it in the definition of communication. Furthermore, exactly the same kind of reasoning can (and should) be done for manipulative communication systems: if these systems are stable it is most probable that also the receiver had some advantage in responding appropriately to signals: for example, responding to a threat signal by fleeing has the clear advantage of not getting hurt by engaging in a fight. This notwithstanding, Dawkins and Krebs' way of framing the problem of the evolution of communication as the evolution of manipulation has been informing most of both the theoretical and empirical biological literature (again, see [Hauser, 1996](#)).

2.2 The phylogenetic problem

The influence of the idea of communication as manipulation is manifest in the treatment of the problem of co-evolution between speakers and hearers. A well known possible solution to this problem has in fact been given through the notion of the exploitation of the sensory (or receiver) bias ([Maynard-Smith and Harper, 2003](#); [Ryan, 1990](#); [Ryan and Rand, 1993](#)). The basic idea is that in order to manipulate the receiver, the speaker could exploit some behavioral bias. If some environmental feature triggers a specific behaviour in some organism that has an adaptive value for another one, the latter can produce signals that resemble that environmental feature in order to manipulate the behavior of the former for its own interests. This idea is certainly sound and can probably explain the evolution of several manipulative communication systems. But it seems much less adequate for explaining the emergence of communication systems which seem to advantage the hearer, like, for example, the alarm calls of vervet monkeys ([Cheney and Seyfarth, 1990](#); [Struhsaker, 1967](#)) or the food calls of honey bees ([Frisch, 1967](#); [Seeley, 1992](#)). It seems more reasonable that this kind of communication systems could emerge from an opposite mechanism, what we can call the *exploitation of a producer bias*. If organisms have some bias to produce behaviours that systematically co-vary with features of the environment which have an adaptive value, those behaviors can be used by other organisms as cues for their own purposes. In this way, the biased behaviors become signals, the organisms that exploit them become hearers and the individuals that produce them become speakers. Is the idea of a producer bias reasonable? Where could such producer biases come from? Some recent work in the artificial life field has provided a possible candidate: the signalling behavior of organisms could be biased in a positive way by the need for internally categorize experience in adaptive ways.

In an artificial life work similar to the one presented here, Cangelosi and Parisi ([1998](#)) were able to evolve a communication system which benefits only hearers in a population of artificial agents. Discussing their surprising result,

those authors suggested that even though in their set-up there was no advantage for speakers to produce useful signals, communication could evolve thanks to the link between communicative behavior and other cognitive abilities. The argument runs as follows. In order to behave adaptively, organisms need to internally categorize experiences in ways that are relevant for their survival. As a consequence, if communicative behaviors depend on internal representations of experience, which is true in Cangelosi and Parisi's simulations and is assumed to be true also in real organisms, then produced signals will tend to be correlated to adaptively useful categories. As a result, the evolution of good signallers might be explained as a by-product of their independently evolving ability to categorize the environment.

2.3 The adaptive problem

The view of communication as manipulation has influenced the biological literature also with respect to the adaptive problem of communication. In fact, any form of manipulation implies that there is an individual which gets manipulated, and which, therefore, provides advantages to another individual without receiving any benefit from doing so. As a result, the main questions that biologists have been asking about the evolution of communication concern adaptive factors: (a) why should hearers let themselves be manipulated by speakers? And, (b) is there anything that can guarantee the reliability of manipulative signals, so that they can be informative and hence useful also for the hearers? Theoretically, the answer to question (a) has been kin selection theory (Hamilton, 1964), while the answer to question (b) has been the handicap principle (Zahavi, 1975).

According to kin selection theory, an altruistic behavior can evolve if it is preferentially directed towards kin related individuals. In fact, in this case the individuals that are advantaged by the altruistic behavior will tend to share the same genes of the altruist, including the genes that determine the altruistic behavior itself. This kind of reasoning can be formalized: an altruistic trait can spread in a population as far as $c < br$, where c is the cost of the altruistic behavior for the individual who performs it, b is the benefit for the receiver, and r is the coefficient of kin-relatedness between the performer and the receiver of the behavior. Kin selection theory can explain, for example, the evolution of begging calls produced by new-borns for receiving extra feeding from parents: as far as the cost for the parent to provide un-necessary food is not greater than the benefit for the new-borns multiplied by the relatedness coefficient, which in sexual species is 0.5 between parents and siblings, the manipulative signalling behavior of the new-borns can spread in the population.

On the other hand, some forms of manipulative signalling behaviors, namely those in which the signaller tries to convince the receiver of its strength (like in mating or aggressive displays), can evolve even among non-kin because their reliability can be guaranteed by the handicap principle, according to which honest signalling can evolve if signals are costly. In fact, if signals are costly, in particular more costly for weak individuals than for strong ones (Grafen, 1990), then

they will be honest, because only really valuable individuals can afford to pay the costs of signalling.

Most of the computational works on the evolution of communication that have dealt with the adaptive problem have discussed one of these two hypotheses proposed by biologists: namely, that altruistic communication can evolve due to kin selection or to the handicap principle. For example, [Ackley and Littman \(1994\)](#) and [Oliphant \(1996\)](#) have shown that a communication system which was advantageous only for hearers could emerge only if the population was spatialized, in the sense that both communicative interactions and reproduction were local. The reason, according to both Ackley and Littman and Oliphant, is the following: if siblings tend to live near their parents and interactions are local, then communicating individuals will tend to be kind-related, and hence communication can emerge thanks to kin-selection. On the other hand, [Di Paolo \(1999\)](#) has criticized this interpretation of the role of space in computational models: by formally analysing kin relatedness in a model similar to that used by Ackley and Littman and Oliphant, he has shown that kin selection cannot fully explain the effects that spatialization has on the emergence of a communication system. Instead, spatial organization by itself, together with other factors such as discreteness and stochasticity, can play a major role in the emergence of cooperative communication. The role of spatial factors in the emergence of simple, altruistic signalling systems is also discussed in [Noble et al. \(Noble et al., 2001\)](#), together with other ecological factors such as the presence of noise. Noble et al. discuss also previous work by [Bullok \(1998\)](#) which demonstrated, through a simple artificial life model, how the handicap principle could lead to the evolution of honest signalling in a manipulative communication system. Finally, for a demonstration of how the genetic relatedness of interacting individuals and the level at which selection operates can affect the evolution of communication in groups of robots, see Chapter ??.

2.4 Disentangling the two problems

As discussed above, [Cangelosi and Parisi \(1998\)](#) suggested that the need to internally categorize the environment might be sufficient for evolving a communication system, even in a case, as their own, in which communication is altruistic. The idea of the emergence of spontaneous good signalling has been restated also by [Marocco et al. \(2003\)](#), who tried to evolve a communication system in a population of neural networks controlling a robot arm who had to categorize objects by physically interacting with them. But in that work the supposed cognitive pressure towards good communication was not sufficient for the evolution of communication. In fact, [Marocco et al.](#) had to introduce in their simulation also a very strong kin-selective pressure and had to make networks pre-adapt to solve the task before introducing inter-agent communication. Furthermore, notwithstanding those facilitations, a good communication system evolved only in 7 out of 10 replications of the simulation.

The work presented here is intended to shed light on the possible role of a 'cognitive pressure' to spontaneous good signalling with respect to the phyloge-

netic and the adaptive problems of communication. Do the need for categorizing experience in adaptive ways indeed constitute a drive toward good individual signalling behavior? What else, if anything, is needed for good communication to emerge at the level of the population? And do these factors, by themselves, render the communication system evolutionary stable even when communication is altruistic in that it advantages only receivers of signals?

3 Experimental set-up

The experimental set-up is a simplification of the ‘mushroom world’ (Cangelosi and Harnad, 2000; Cangelosi and Parisi, 1998; Parisi, 1997).¹ We simulate a population of artificial agents that are evolved for the ability to eat edible ‘mushrooms’ and avoid poisonous ones, and for the ability to minimize the movements made. Furthermore, agents can send signals to each other for communicating the quality of the encountered mushrooms.

3.1 The environment and the task

The environment is a one-dimensional cell world (Figure 1). In each trial an organism is placed in the start cell at the beginning of the corridor and a mushroom is placed in the last cell of the corridor (the mushroom cell). If the organism reaches the mushroom cell it eats the mushroom contained in the cell. The agents are provided with a sort of ‘camera’ with which they can perceive the visual pattern of the nearby mushroom, that is encoded as a vector of 10 bits. There are 420 different mushrooms, of which 210 are edible while 210 are poisonous.² The sensory system of the agents is very limited: an agent can see the mushroom only if they are in two adjacent cells. Hence, without communication, the only strategy which is available to the agents is to run through the entire corridor in each trial in order to reach the last but one cell and see whether the mushroom is edible or poisonous. On the other hand, if another agent that can perceive the mushroom communicates to the first agent whether the mushroom is edible or poisonous, the first agent can avoid spending its energy in order to run through the corridor in the trials with poisonous mushrooms. This confers an advantage to the first agent (the hearer) but not to the other agent (the speaker).

3.2 The neural network

The agents’ behavior is controlled by a neural network that includes 10 visual input units, 2 communicative input units, 2 hidden units, 1 motor output unit

¹ The difference with respect to the original mushroom world is that the new world is one-dimensional rather than bi-dimensional. This simplification has been introduced in order to make the analysis of the results easier.

² The perceptual properties of edible mushrooms are encoded as all and only the patterns that differ by 4 bits from the prototype represented by ten 1s, whereas the perceptual properties of poisonous mushrooms are all and only the patterns that differ by 4 bits from the prototype represented by ten +1s.

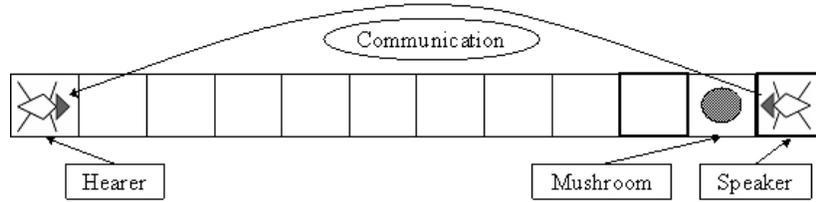


Fig. 1. The environment

and 2 communicative output units (Figure 2). The visual input units encode the perceptual properties of the mushroom that is in front of the agent, while when no mushroom is visible, their activation is 0. When an agent acts as a hearer, the activation of the 2 communication input units corresponds to the activation of the 2 communication output units of the agent that is near the mushroom and acts as speaker, while when an agent is playing the role of the speaker the activation of the communication input units is 0. There are only 2 hidden nodes (with bias), with hyperbolic tangent (\tanh) activation function, that are fully connected both with the input and with the output nodes. The locomotion of the organisms is controlled by a single output node with binary activation. If the activation is one, the organism moves to the next cell, otherwise, it stays still. Finally, there are two communication output nodes with \tanh activation function. The activation vector of these two nodes constitutes the signal that is transmitted by agents acting as speakers to agents acting as hearers.

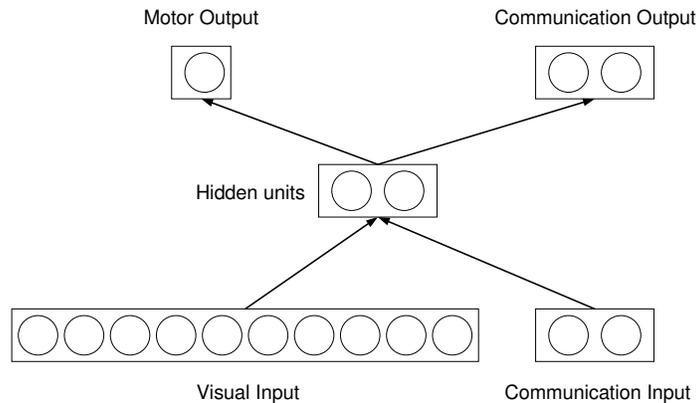


Fig. 2. The neural network. Arrows represent connections from all the neurons of the sending group to all the neurons of the receiving group

3.3 Individual life and the fitness formula

The agents can play the role of the speaker or of the hearer. In the former case the robots are placed in a cell adjacent to the mushroom, are allowed to produce a signal, do not receive signals, and are not allowed to move. In the latter case, the agents are placed at the start cell of the corridor (i.e. far from food), receive the signal produced by the speaker, and are allowed to move. Agents are evaluated only with respect to the behavior that they exhibit when they act as a hearer (i.e. for the ability to eat and avoid edible and poisonous mushrooms on the basis of visual information and on the basis of the signal they receive from speakers), while they are not rewarded on the basis of the behaviour that they produce when they act as speakers. The life of each individual (as a hearer) lasts 420 trials, one for each different mushroom. In each trial the agent which has to act as speaker is chosen randomly from among the remaining members of the population. The trial ends either if the hearer reaches the mushroom and eats it or after 11 input-output cycles, which implies that the hearer did not eat it. Fitness is calculated according to the following formula:

$$f(x) = \frac{an_e - bn_p - n_m}{norm} \quad (1)$$

where n_e and n_p are, respectively, the number of edible and poisonous mushrooms eaten by x , n_m is the number of movements made by x , a and b are two constants (set to 30 and 5, respectively), and $norm$ is the normalization factor (maximum possible fitness) which equals 4200.

3.4 The genetic algorithm

In each generation the population is constituted by 100 agents. The genotype of each individual encodes the connection weights and biases of its neural network as real values in the range $[4; +4]$. Reproduction is sexual and the population is divided by sex: in each generation there are 50 males and 50 females. For each offspring to be generated, one male and one female are selected for reproduction through the roulette method with windowing, and they reproduce sexually (with double-point crossover).³ After all the new individuals are created, we apply random mutations: each single weight of each individual has 0.1% probability to have its current value replaced by a new value randomly chosen in the range $[4; +4]$. Simulations run for 2000 generations, during which we calculate various statistics, the most important ones being average fitness and what we call 'communication system quality'.

³ 'Windowing' consists in reducing the fitness of each subject by the fitness of the worst subject, so that the fitnesses of all individuals are ≥ 0 . For an explanation of the details of the genetic algorithm, see (Mitchell, 1996) and (Mirolli and Parisi, 2005). In any case, the results reported in the next sections seem quite robust with respect to the details of the genetic algorithm: for example, runs with slightly different mutation rates or with asexual reproduction (i.e. without cross-over) provided qualitatively similar results.

3.5 Measuring communication system quality

For a communication system to be good, three conditions must apply: (a) objects belonging to the same category must be signalled in similar ways, (b) objects belonging to different categories must be signalled in different ways, and (c) all individuals must signal in the same way, meaning that the system is shared in the population. In the model, signals are two-dimensional vectors, which can be represented as points in a two-dimensional space, whose coordinates are the activations of the two communicative output units of speakers. Hence, the communication system of a population consists of two clouds of points: the points that represent the signals emitted by all the organisms in presence of all edible mushrooms and the points corresponding to the signals emitted in presence of all poisonous mushrooms. Condition (a) above is satisfied if the two clouds are small (meaning that mushrooms of the same category are signaled in similar ways). Condition (b) is satisfied if the two clouds are distant from one another (meaning that mushrooms belonging to different categories are signaled in different ways). Condition (c), which regards the sharedness of the communication system, is implicitly satisfied by the other two conditions, since the clouds represent the communication system of the entire population. Our communication system quality is calculated as the average between the two measures that make a communication system good: (1) the mean distance of the points of each cloud from the cloud's geometrical centre, and (2) the distance between the geometrical centres of the two clouds.

4 Cognitive, Genetic, and Adaptive Factors in the Evolution of Communication

In our simulation's scenario the emergence of a good communication system would provide an advantage only to hearers. A hearer who receive good signals could in fact use the received information for deciding whether to approach mushrooms signaled as 'edible' or to stay still in the case the received signal stands for 'poisonous', thus avoiding a waste of energy. In contrast, a speaker has no benefit whatsoever in producing good signals. On the contrary, since individuals compete for reproduction, producing good signals is a purely altruistic trait in that, by giving an advantage to competitors, it indirectly decreases the reproductive chances of the good signaller.

The results of 10 out of 50 replications of the simulation are shown in figure 3. The results are quite surprising: average fitness and communication system quality fluctuate cyclically between very low and very high values. This means that communication emerges continually but it is also continually disrupted. As this very unstable pattern is present in all the replications of the simulation, the mechanisms that generate it must be very strong. What are these mechanisms?

The sudden drops in the communication system quality can be explained as a consequence of the conflict of interests between individuals that arises from the fact that communication advantages the hearer but not the speakers and from

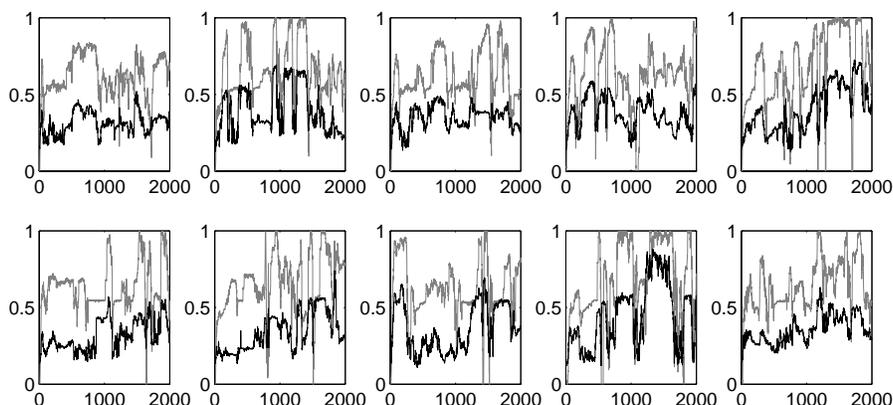


Fig. 3. Average fitness (gray line) and communication system quality (black line) of 10 replications of the simulation.

the fact that communicating individuals are not homogenous (as in the case of other artificial experiments described in the other chapters; on the issue of conflicts of interest between communicating individuals, see also Chapter ??). Once a good communication system has emerged in the population bad speakers will be selected against good speakers because they will act as cheaters: they take advantage of the good signals they receive while misleading other organisms by producing bad signals. The result is that cheaters rapidly invade the population with the consequence of a sudden decrease in the quality of the communication system. But why a good communication system should emerge in the first place or re-emerge after being disrupted? The reason lies in two factors that are constantly at work, pushing the communication system towards good quality.

The first factor is the ‘cognitive pressure’ suggested by Cangelosi and Parisi (1998). To solve their adaptive task, agents need to differentiate the behaviour they produce for different types of mushrooms. This can be accomplished by relying either on the heard signals or on one’s own ability to differentiate the visual patterns corresponding to eatable and poisonous mushrooms. Since signals are not always reliable, agents must be able to solve their problem by appropriately categorizing the perceived visual stimuli and by reacting consequently. In other words, in order for an agents to behave efficiently, its network’s connection weights must be such that mushrooms belonging to the same category elicit similar activation patterns in the hidden units (similar internal representations), while mushrooms belonging to different categories elicit different internal representations. But the way in which visual sensory states are transformed into internal states do not only affect how the agents react to such sensory states (i.e. whether they eat or not the food) but also the signals that they produce. This means that the development of an ability to produce different motor actions for

visual stimuli belonging to the two different categories and similar motor actions for the visual stimuli belonging to the same category also tend to produce an ability to produce different signals for visual stimuli belonging to different categories and similar signal for stimuli belonging to the same category (i.e. it spontaneously leads to the production of a informative signals).

However, this cognitive explanation of the spontaneous emergence of a good communication system is not enough. The cognitive pressure can in fact explain only the fact that each individual tends to produce good signals, that is, signals that systematically co-vary with the category of perceived mushrooms. It cannot explain why the communication system is shared in the population, that is, why different individuals tend to produce similar signals. But signals sharedness is necessary for good communication, and is indeed considered in our measure of the communication system quality. (Remember that the measure is taken at the populational level, considering the signals emitted by all individuals). In order to explain the sharedness of the communication system we need to introduce another factor, namely genetic convergence. In fact, if organisms have similar genes they will tend to have similar (categorical) internal representations and to produce similar signals. Genetic similarity can in turn be explained by genetic convergence due to selection. Only the fittest individuals reproduce in each generation. As a consequence, their genes will tend to spread in the population reducing genetic diversity.

In summary, the peculiar but very robust phenomenology of our simulations can be explained as the result of three factors which are constantly at work in the model:

1. the invasion of a well-communicating population by bad-speakers who cheat others by sending misleading signals
2. the cognitive pressure towards good individual signaling due to the need of agents to categorize mushrooms in an appropriate way
3. the genetic convergence of the population which tends to make signals shared

Factor 1 is due to the altruistic character of this kind of communication, and hence is related to the adaptive problem of communication. Factors 2 and 3 together determine what we call a ‘producer bias’ towards good signaling, which might constitute a possible solution to the phylogenetic problem of communication. In the next two sections we test our hypotheses about these three factors by comparing the results of this base-line simulation with those of three different control experiments.

5 The kin-selection simulation

5.1 Simulation

If in the base-line simulation the quality of the communication system really decreases because of the altruistic character of good signaling, then communication should not get disrupted if we add to the base-line simulation some adaptive factor that makes it possible for altruistic behaviours to evolve.

As discussed above, kin selection has already been claimed to play a role in other models of the evolution of communication (Ackley and Littman, 1994; Oliphant, 1996), but this hypothesis has never been tested directly. In fact, previous works assumed that kin selection was the at work due to the spatialization of the organisms interactions and reproduction, an assumption which has been criticized by Di Paolo (1999). Here we test the power of kin selection directly, through the manipulation of the interactions between individuals. In particular, we vary the probability that communicative interactions happen among close kin. The rationale for doing so is the following. Close kin tend to share the same genes. So, the higher the probability to speak to a close kin, the higher the chances that the altruistic behavior of producing good signals is preserved because it tends to confer an advantage to organisms with the same genes and hence with the same altruistic behavior.

The kin selection simulation runs as the base-line one, but for the following difference. In this simulation 50 out of the 100 individuals in each generation are males and 50 are females. For each breeding cycle, one male and one female are selected for reproduction and this pair generates two offspring, one male and one female, with different cross-over points and different mutations on connection weights. This breeding method makes sure that each organism of the population has at least one full sibling. So, we implement kin selection by varying the probability p with which in any trial the speaker is one of the full siblings of the hearer (with probability $1 - p$ the speaker is not one of the hearers full siblings, but it can be a half sibling).

5.2 Results

We run several different simulations with different values of p . Figure 4 compares the results of the base-line simulation with those of the kin-simulations with $p = 1\%$, 5% , 10% and 100% .

The results of the simulations are very clear: both the quality of the communication system and average fitness increase as the probability p of speaking to a close kin increase. Vice-versa, the lower this probability, the lower the communication system quality and the average fitness of the population. With $p = 0.01$, the results are practically the same as those of the base-line simulation. The fact that the results of the simulation with $p = 0.1$ are not very different from those with $p = 1$ should cause no surprise. After all, p determines only the probability that the speaker is a full sibling of the hearer, but even when this is not the case, the speaker can share the same genes of the hearer either because it is a half sibling or simply because of the genetic convergence of the population discussed above and demonstrated in the next section. In fact, augmenting the kin selection factor has just the effect of adding an additional selective pressure for the stabilization of the communication-output genes. In other words, kin selection simply strengthens a genetic pressure towards the emergence of communication which is already present in the base-line simulation.

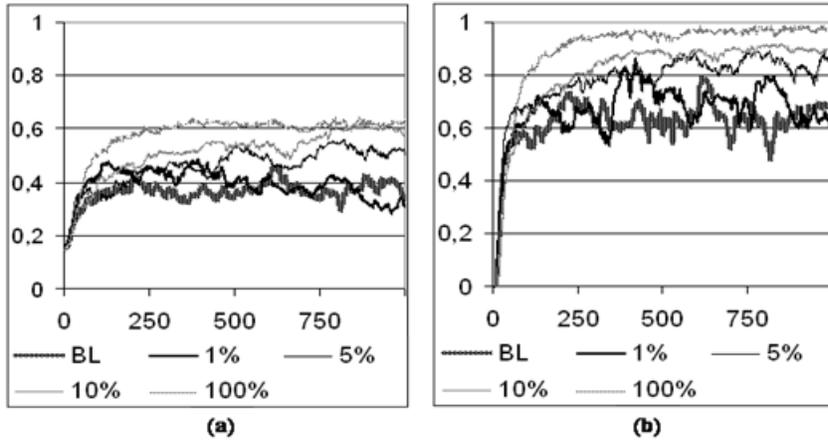


Fig. 4. Comparison between the base-line and of 4 kin selection simulations with different values of p (1%, 5%, 10% and 100%): (a) communication system quality and (b) average fitness. Average results of 10 replications of the simulations

6 The no-cognitive-pressure and no-communication simulations

In order to fully understand the complex evolutionary dynamics of our model, in (Mirolli and Parisi, 2008) we have provided a very detailed analysis of 200 generations of a single run of our simulation. For reasons of space, here we do not present these analyses. Rather, we test our hypotheses about the producer bias towards spontaneous good signaling by comparing the average results of 50 replications of the base-line simulation with the results of two control experiments, which we call the no-cognitive-pressure and the no-communication simulations.

6.1 Simulations

The no-cognitive-pressure simulation runs exactly as the base-line one except for the architecture of organisms' neural networks (figure 5b). The difference is that the neural networks of the no-cognitive-pressure simulation have two groups of hidden units (composed by 2 neurons each). The first group receives connections from both the visual and the communicative input units and send connections to the motor output units. The second group receives connections from the visual input units and sends connections to the communicative output units. In this way we artificially avoid any impact of categorization on communicative behaviour: since the communicative output does not depend on the same neurons on which

the motor output does, with the new architecture there cannot be any cognitive pressure towards good signal production.⁴

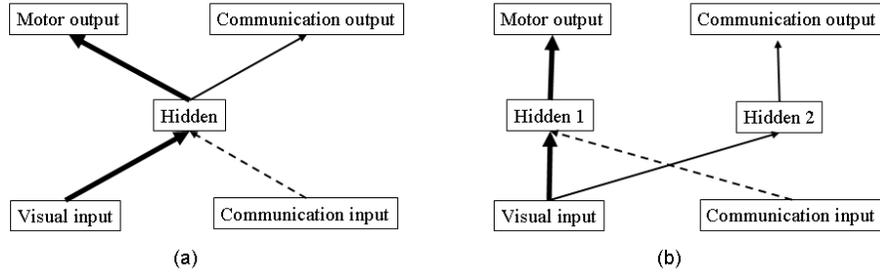


Fig. 5. Comparison between the neural network used in the base-line and no-communication simulations (a) with that used in the no-cognitive pressure simulation (b). Connection weights are grouped according to their roles in communication: non-communicative (thick arrows), communicative-output (thin arrows) and communicative-input (broken arrows).

Organisms of the no-communication simulation have the same neural network as the one used in the base-line simulation (figure 5a) but in this case communication is not allowed. In other words, in this simulation the organisms never act as speakers. Consequently, tested individuals never receive any communicative input. Under such conditions the only way for an individual to behave correctly is to move towards the mushroom until the mushroom is close enough to be clearly perceived. At that point, the organism can decide whether to eat the mushroom or refrain from doing so according to the perceived mushroom’s quality. As a consequence, the maximum fitness that can be obtained by the individuals of the no-communication simulation is 0.55 because these individuals must always pay the cost of moving towards mushrooms. This cost could be avoided, when mushrooms happen to be poisonous, by organisms that receive reliable signals (in the simulations with communication). Note that in the no-communication simulation the communicative weights (those that link the communicative input and output units with the hidden units) are never used. Hence, they are adaptively neutral and subject to genetic drift. Nonetheless, we can still test what signals are produced by the networks located near to a mushroom and, consequently, measure the communication system quality of the no-communication simulation even if the communication system is never used.

⁴ We also ran a simulation in which the visual input is directly connected to the communication output, with results qualitatively similar as those of the reported no-cognitive-pressure simulation.

6.2 Results

Figure 6 shows the results of the 3 simulations (base-line, no-communication, and no-cognitive-pressure) in terms of communication system quality (a) and average fitness (b). The results, averaged over 50 replications of each simulation, clearly support the hypothesis regarding the cognitive pressure towards the emergence of a good communication system. The communication system quality in the no-communication simulation is quite high (about 0.5, see figure 6a), in particular, significantly higher than in the base-line simulation. This means that the need to categorize mushrooms and the genetic convergence of the population are sufficient to produce a good communication system even in a population in which communication is not allowed and the communicative-output weights are subject to genetic drift. Furthermore, the fact that the quality of the communication system in the no-communication condition is considerably higher than in the base-line condition (about 0.35) means that, if signals are used, there is a direct selective pressure *against* good speaking, due to the altruistic character of communication.

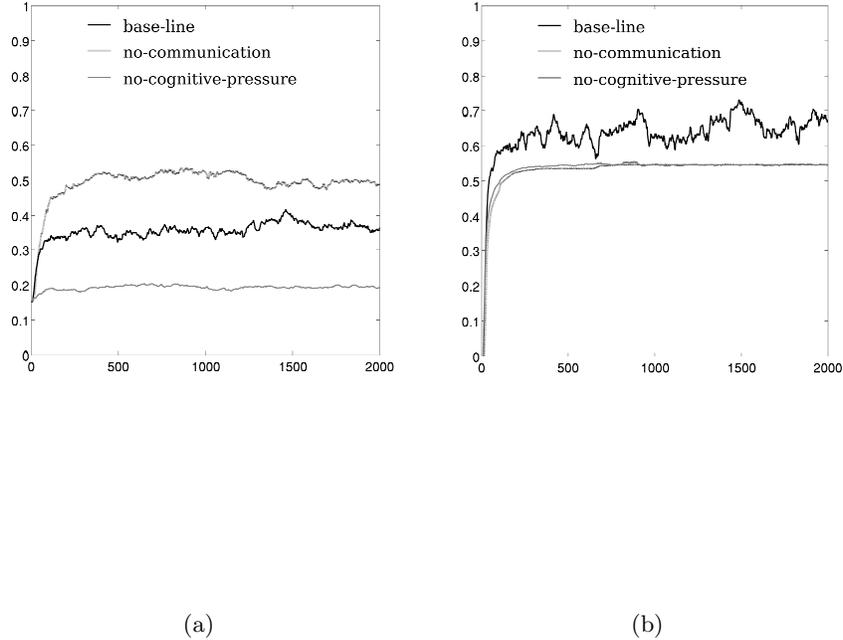


Fig. 6. Communication system quality (a) and average fitness (b) of the base-line, no-communication and no-cognitive-pressure simulations. Average results of 50 replications.

The results of the no-cognitive-pressure simulation confirm our analysis. If we prevent any influence of cognition on signal production by manipulating the architecture of the neural network, a good communication system never emerges. This results in an average communication system quality of about 0.2, which is considerably lower than that of the base-line condition.

The results in terms of average fitness (figure 6b) just confirm this analysis. In both the no-communication and the no-cognitive-pressure simulations fitness reaches the maximum value reachable by organisms which cannot get any benefit from communication: in the first case because communication is not allowed, in the second case because signals are useless and cannot be exploited. Since the organisms of the base-line simulation can sometimes take advantage of communication, their average fitness is a little (but not much) higher.

Since the genetic convergence of the population is necessary, according to our hypothesis, for explaining the sharedness of the spontaneously emerging communication system, an interesting statistic to take is the genetic variance of our population. But since our explanation postulates different genetic pressures on different sets of genes, we wanted to determine whether those different genetic pressures would result in different genetic variances. Consequently, we divided our system (the genome of the evolving population) into three sub-systems corresponding to the three distinct functional roles that different sets of genes play with respect to communication: the non-communicative genes, the communicative-input genes, and the communicative-output genes (see figure 5). Genetic variance for each gene (connection weight) x (V_x) is calculated, standardly, according to the following formula:

$$V_x = \frac{\sum_{i=1}^N (\bar{w}_x - w_x^i)^2}{N} \quad (2)$$

where N is the number of individuals in the population (i.e. 100), \bar{w}_x is the mean value for weight x , and w_x^i is the weight x of the i_{th} individual. Consequently, the variances (V^k) for the three categories ($k \in K$) of weights (non-communicative, communicative-input, and communicative-output) are calculated as follows:

$$V^k = \frac{\sum_{x=1}^{N^k} V_x}{N^k norm} \quad (3)$$

where N^k is the number of weights of the k_{th} category (i.e. 28, 4 and 4 for non-communicative, communicative-input, and communicative-output weights, respectively), V_x is the variance on weight $x \in k$ and $norm$ is the normalization factor, so that V^k is in $[0;1]$ (since weights are bound in $[-4, 4]$ the theoretical maximum variance for each gene is 16, obtained when half of the population has a value of -4 and the other half a value of 4).

Figure 7a shows the variances of non-communicative and communicative-output weights of the base-line simulation. There is indeed a considerable genetic convergence since both variances are very low, which explains the sharedness of the spontaneously emerging communication system. But the communicative-output weights have a much higher variance (about 0.1) than

the non-communicative weights (about 0.01). The explanation is to be found in the selective pressure against good signal production due to the altruistic character of communication. The cheaters which invade the population when the communication system is good are those organisms that, due to mutations, have communicative-output weights that are different from those on which the population has converged and which, consequently, let them produce misleading signals. In other words, the variance of the communicative-output weights is higher than the variance of the non-communicative weights because each time the population has genetically converged and the quality of the communication is high, there is a strong and reliable pressure to possess communicative-output genes that differ from those of the population in order to cheat competitors.

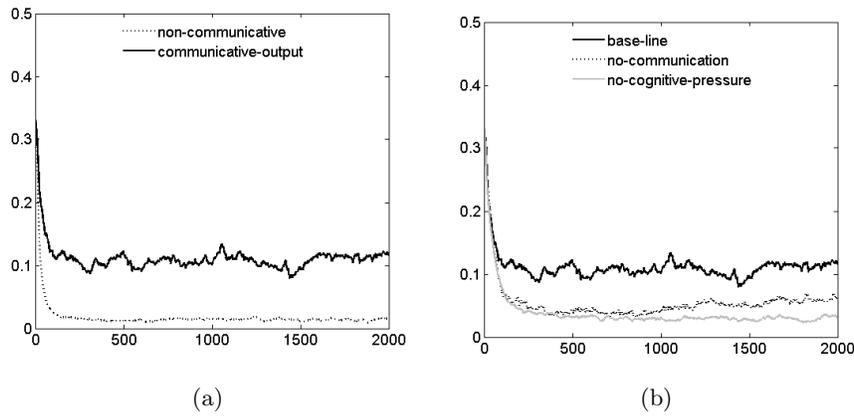


Fig. 7. Average variances on the non-communicative and communicative-output weights of the base-line simulations (a) and average variances on the communicative-output weights of the base-line, no-communication and no-cognitive-pressure simulations (b). Average results of 50 replications of the simulations.

This analysis is further confirmed by the comparison between the variances on the communicative-output genes in the three conditions: base-line, no-communication, and no-cognitive-pressures (figure 7b). The three different values reflect the different pressures on those genes in the three conditions. In the no-cognitive-pressure simulation there is no spontaneous tendency to produce a good communication system, and the low variance of communicative-output weights (about 0.03) is explained by the genetic convergence of the population. The fact that the variance of communicative-output weights in the no-communication condition is higher (about 0.06) and tends to grow after reaching its minimum is due to the fact that in that simulation there is no pressure whatsoever on the communicative-output genes, so they are subject to genetic drift. On

the contrary, in the no-cognitive-pressure simulation there is some mild pressure to generate bad signals. Nonetheless, the variance of the communicative-output genes is still higher in the base-line simulation than in the no-communication simulation. This confirms the hypothesis that when communication is used there is a reliable selective pressure against good speakers, resulting in a pressure against the genetic convergence on the communicative-output weights: mutations on those genes that generate cheaters are rewarded by selection and spread in the population.

7 Discussion

As discussed in the introduction, the fact that communication requires two independent capabilities, speaking (producing useful signals) and hearing (understanding perceived signals), poses two related general problems: a phylogenetic problem, and an adaptive problem. In this section we discuss our results with respect to both these problems.

7.1 The Producer Bias Hypothesis

The phylogenetic problem is very general problem posed by all kind of communication systems, and can be stated as follows: what evolutionary paths can lead to the emergence of communication systems, since speaking and hearing abilities are complementary and hence adaptively neutral if they are not both present at the same time? The biological literature has provided a partial answer to this question with the concept of a receiver (or sensory) bias: if some stimuli trigger in some organism a behavioral response that advantages another organism, the latter can produce signals that resemble those stimuli in order to manipulate the behavior of the former. In this way, the problem of co-evolution is by-passed through the exploitation of behavioral biases in the receivers which are assumed to be present for reasons other than communication. There are several reasons for receiver biases to be present. In particular, they can be both the results of selective processes independent from communication (for example, detection of prey), and non-selected, incidental traits which are simply consequences of how the sensory systems and the brain of a species are formed (see [Endler and Basolo, 1998](#); [Guilford and Dawkins, 1991](#)). Since its introduction in the biological literature on the evolution of communication the concept of a receiver bias has inspired a lot of both empirical (e.g. [Basolo, 1995](#); [Proctor, 1992](#); [Ryan et al., 1990](#)) and theoretical/computational research (e.g. [Arak and Enquist, 1995](#); [Ryan et al., 2001](#)).

In the present chapter we have argued for the complementary idea of the exploitation of a producer bias, which has been so far overlooked in the biological literature on the evolution of communication. In fact, though ideas similar to that of producer bias have appeared in the literature under the label of ‘intention movements’ (i.e. movements necessarily preceding an action; see, e.g., [Krebs](#)

and Dawkins, 1984; Tinbergen, 1952), these ideas have not received enough theoretical attention and have not inspired any substantial empirical research. The reason for this is, we suppose, twofold: first, as discussed in the introduction, the idea of producer biases is fundamentally related to *informative* communication, while the influential conception of communication of Dawkins and Krebs has led the biological literature to focus on *manipulative* communication; second, ideas similar to the producer bias have mostly been discussed in relation to the adaptive problem of communication (in particular in relation to the handicap principle: see, for example, Lotem et al., 1999; Noble, 1998), while they have never been explicitly and clearly proposed as a possible general solution to the phylogenetic problem of communication.

But as the exploitation of receiver biases is an important mechanism which can trigger the evolution of manipulative communication systems, the exploitation of producer biases can play an analogous role in the historical emergence of communication systems which (principally) benefit hearers, like alarm and food calls. Furthermore, the simulations presented in this paper suggest also one possible mechanism which can lead to producer biases: namely, the cognitive pressure due to the need for organisms to internally represent stimuli according to adaptively relevant categories. In fact, as our simulations have shown, good categorization of stimuli can result in the spontaneous production of good signals which can be thus be exploited by hearers.

But is the hypothesis of a cognitive pressure towards good signalling behavior plausible? Does it hold only for our simulations and related artificial systems? Or is it also a possible candidate for explaining the evolutionary emergence of real communication systems? We think that the hypothesis is indeed plausible and that it should be taken in consideration when trying to explain the evolutionary emergence of animal communication systems, in particular those which seem to benefit principally receivers. After all, the only assumption which needs to be made for our hypothesis to work is the link, in the organisms' brains, between communicative and non-communicative behaviors. In fact, if such a link exists, then the need for individuals to categorize the environment in adaptive ways will generate a bias towards the production of useful signals. And this assumption seems quite reasonable: what should signals reflect if not what an organism have in their mind/brain?

Note that notwithstanding our use of a representationalist jargon, the hypothesis of possible cognitive pressures towards good signaling does not depend on taking a representationalist stance in the representationalists vs. dynamicists debate. Even the most radical enthusiast of the dynamical approach to cognition and the most critic or skeptic on the use of the notion of representation for explaining adaptive behavior (Beer, 2003; Brooks, 1991; Harvey, 1996) acknowledges that an organism's behavior depends also on internal (neural) dynamics. Hence, the same single assumption is needed, from a dynamical systems point of view, for accepting the possibility of a possible cognitive pressure towards good signaling: namely, the assumption that signaling behavior is internally linked to the brain structures that govern also other non-communicative behaviors. In

fact, for an organism's behavior to be adaptive, different adaptive conditions will be correlated with different internal dynamics, which in turn will tend to be reflected in different signaling behavior. But this is exactly what the hypothesis on a cognitive pressure towards good signalling states: that produced signals will tend to spontaneously reflect adaptively relevant situations due to the need for organisms to cope adaptively with their environment.

The producer bias hypothesis is also empirically testable. One way to test it is the following. Take an animal capable of sending signals. Train it to send a signal in at least two situations which differ from the point of view of the animal's survival (in presence of edible vs. poisonous items, as in our simulation, or for asking two significantly different kinds of things, like food vs. a partner for mating). Our hypothesis of the cognitive pressure on good signalling predicts that the animal will tend to *spontaneously differentiate* between the signals emitted in the two different situations.

It is of the most importance that a first confirmation of this prediction can indeed already be found in recent empirical experiments on Japanese monkeys performed by [Hihara et al. \(2003\)](#). In these experiments a monkey was trained to produce a generic call in two different conditions: (a) for receiving food and (b) for receiving a tool with which it could reach a distant food. Surprisingly, without any reward for doing so, the trained monkey started to spontaneously differentiate the sounds emitted in the two conditions. The authors of the experiments interpret their striking finding as the result of a change, in the monkey, from emotional to intentional vocalizations. But the spontaneous differentiation of monkey's calls can be explained, less speculatively, by the simple fact that in the two conditions the monkey's brain was in two different states: one corresponding to the need for a piece of food and the other corresponding to the need of a tool with which to reach the food. As predicted by our hypothesis, different internal states, be they 'conscious' (as Hihara et al. suggest), 'intentional' or 'emotional', tend to produce different (communicative) behaviors. To put it shortly: what you have in mind will influence what you say. The signals an organism produce will tend to reflect the way in which the organism categorizes its experience.

Furthermore, we would like to note that the hypothesis that communication might phylogenetically emerge thanks to the exploitation of producer biases is far more general than the hypothesis of the cognitive pressure towards spontaneous good signalling. In fact, though we think that the need to adaptively categorize the environment does constitute a strong mechanism for producer biases to appear, there is no reason for producer biases to have exclusively cognitive origins. On the contrary, there are good reasons to suppose that other, non-cognitive, mechanisms can equally generate producer biases. Indeed, it is reasonable to think that most emotional communication might have evolved thanks to producer biases: since emotional states are associated with specific internal bodily changes (e.g. changes in heart rate, blood pressure...), the visible consequences of these changes can be exploited by 'receivers' to infer the emotional state of

the ‘producer’, thus resulting in the evolution of an emotional communication signal through a non-cognitive producer bias.

With respect to the endeavor of designing embodied and communicating artificial agents, our research has at least two important implications. The first one is related to the importance of studying the evolution of communication together with the evolution of other, non-communicative behaviors (see Chapter ??). The results of the simulations presented in this chapter add, to the theoretical claim that communication can only be understood if it is grounded in the environmental and social needs of communicating organisms, another, practical, reason for making communication emerge from the needs posed by other non-communicative tasks. The reason is that the need for internally categorize experiences in adaptive ways for purposes other than communication can play a major role in solving the phylogenetic problem of communication. As our non-cognitive pressure simulations demonstrate, if the evolution of producer biases towards communication is not allowed, it is much more difficult to evolve a communication system in artificial organisms. And if the only thing that agents have to do is to communicate, the presence of producer biases is prevented.

Second, the favorable effect that having to represent one’s environment in adaptive ways can have on the emergence of communication can show up only if the control systems of agents permit this to happen. This is the case for neural networks but not for other kinds of control system used in the adaptive behavior literature on the evolution of communication such as look-up tables (e.g. [Di Paolo, 1998](#); [MacLennan and Burghardt, 1993](#); [Steels, 1996](#)). Neural networks, as the control systems of real organisms, that is brains, process information in a parallel and distributed way, and the same processing units are involved when perceiving different situations and for producing different behaviors. It is the link between communicative and non-communicative behavior that tends to result in producer biases towards communication. If the control system of an artificial agent is a look-up table, in which the relationships between each possible perception and each possible action are completely independent the ones from the others, the emergence of a producer bias would be very unlikely, and only due to chance. On the contrary, neural networks make the interdependencies between different behaviors explicit, thus providing a systematic mechanism for the emergence of producer biases.

This implies two different views of ‘modules’ in control systems. In a neural network, and very likely in the brain itself, modules are purely functional but structurally they can include and use subsets of units (neurons) that are also included in and used by other modules with different functions. On the other hand, in non-neural computational models such as look-up tables modules are both functionally and structurally distinct entities. If cognitive pressures and producer biases really exist and have an important role in explaining the evolutionary emergence of communication, neural models appear to be more appropriate than non-neural models to capture the evolutionary emergence of communication.

7.2 Adaptive Factors

While the phylogenetic problem of communication is completely general, the adaptive problem hold only for those communication systems in which there is a conflict of interests between the two actors of communicative interactions. Cooperative communication systems in which both senders and receivers are equally advantaged by signalling do not pose any particular adaptive problem. On the other hand, manipulative communicative systems like aggressive displays and informative communication systems like alarm or food calls pose an adaptive problem due to the fact that either receiver (in manipulation) or senders (in information transfer) seem to act altruistically.

From the above discussion we can conclude that the hypothesis of producer biases towards spontaneous good signaling represents a possible solution to the phylogenetic problem of communication. But can ‘producer biases’, by themselves, solve also the adaptive problem of communication, as suggested by (Cangelosi and Parisi, 1998)? Contrary to what happens in the simulations by Cangelosi and Parisi, in our model the emerged communication system is continually disrupted by the invasion of the population by cheaters. The reason is that a population of altruistically communicating individuals is always subject to be invaded by cheating individuals which benefit from the received signals but do not reciprocate (or *lie*, thus misleading competing individuals), *unless some other factor does not prevent this from happening*. In (Mirolli and Parisi, 2008) we have provided a possible explanation about why the communication system in the simulations of Cangelosi and Parisi turn out to be evolutionary stable: in short, the stability of the communication system was due to the peculiar mutation operator used in that work, that prevented the invasion of the population by cheaters.

The most classical solutions to the problem of altruism in the evolution of communication are kin selection (Hamilton, 1964) and sexual selection through the handicap principle (Bullock, 1998; Grafen, 1990; Zahavi, 1975) (another classical explanation of altruistic behaviors is reciprocal altruism Axelrod, 1981; Trivers, 1971, but as far as we know this has never been used in the contest of altruistic communication).

Our kin selection simulations clearly shows that the probability that the speaker and the hearer share the same genes directly affects the stability of a communication system which benefits the hearer but not the speaker. In particular, the higher the probability that the participants to communicative interactions are closely kin-related, the higher the stability of the evolving communication system. This is perfectly in line both with classical kin-selection theory and with other recent works on the evolution of communication in embodied agents in which good signalling has been reported to emerge only when communicating agents are clones (e.g. Marocco et al., 2003; see also Chapter ??; indeed, evolutionary robotics experiments dealing with the evolution of cooperative behaviors in general and communication in particular typically use clones, as those presented in Chapters ?? and ?? of this volume).

In [Mirolli and Parisi \(2005\)](#) we have shown that the invasion of a communicating population by cheaters can be prevented not only by kin-selection, but by two other possible factors: (1) the social learning of the communication system through docility ([Simon, 1990](#)) and (2) the use of signals not only for communication but also for individual purposes (see also [Mirolli and Parisi, 2006, res](#)) such as as memory aids. While kin selection can account for the evolution of altruistic forms of communication like the alarm and food calls of many animals, these other two factors might have played a role in the evolution of human language, and, in particular, in the altruistic uses of language such as for communicating information about the external environment, for which language seems to be particularly well suited (see Chapter ??).

Bibliography

- Ackley, D. H. and Littman, M. L. (1994). Altruism in the evolution of communication. In Brooks, R. A. and Maes, P., editors, *Artificial Life IV: Proceedings of the International Workshop on the Synthesis and Simulation of Living Systems*, pages 40–48, Cambridge, MA. MIT Press.
- Arak, A. and Enquist, M. (1995). Conflict, receiver bias and the evolution of signal form. *Philosophical Transactions of the Royal Society, Series B*, 349:337–344.
- Axelrod, R. (1981). The evolution of cooperation. *Science*, 211(4489):1390–1396.
- Basolo, A. L. (1995). Phylogenetic evidence for the role of a pre-existing bias in sexual selection. *Proceedings of the Royal Society of London B*, 259(1356):307–11.
- Beer, R. D. (2003). The dynamics of active categorical perception in an evolved model agent. *Adaptive Behavior*, 11(4):209–243.
- Brooks, R. A. (1991). Intelligence without representation. *Artificial Intelligence Journal*, 47:139–159.
- Bullock, S. (1998). A continuous evolutionary simulation model of the attainability of honest signalling equilibria. In Adami, C., Belew, R., Kitano, H., and Taylor, C., editors, *Proceedings of the Sixth International Conference on Artificial Life*, pages 339–348.
- Cangelosi, A. and Harnad, S. (2000). The adaptive advantage of symbolic theft over sensorimotor toil: Grounding language in perceptual categories. *Evolution of Communication*, 4:117–142.
- Cangelosi, A. and Parisi, D. (1998). The emergence of a language in an evolving population of neural networks. *Connection Science*, 10(2):83–97.
- Cheney, D. L. and Seyfarth, R. M. (1990). *How monkeys see the world: Inside the mind of another species*. University of Chicago Press, Chicago.
- Dawkins, R. and Krebs, J. R. (1978). Animal signals: information or manipulation? In Krebs, J. R. and Davies, N. B., editors, *Behavioural Ecology: An Evolutionary Approach*, pages 282–309. Blackwell Scientific Publications, Oxford.
- Di Paolo, E. (1998). An investigation into the evolution of communication. *Adaptive Behavior*, 6(2):285–324.
- Di Paolo, E. A. (1999). A little more than kind and less than kin: the unwarranted use of kin selection in spatial models of communication. In Floreano, D., Nicoud, J., and Mondada, F., editors, *Advances in Artificial life. Proceedings*

- of the 5th European Conference on Artificial Life, pages 504–513, Lausanne. Springer-Verlag.
- Endler, J. A. and Basolo, A. L. (1998). Sensory ecology, receiver biases and sexual selection. *Trends in ecology & evolution*, 13(10):415–420.
- Frisch, K. v. (1967). *The Dance Language and Orientation of Bees*. Harvard University Press, Cambridge, Mass.
- Grafen, A. (1990). Biological signals as handicaps. *Journal of Theoretical Biology*, 144:517–546.
- Guilford, T. and Dawkins, M. S. (1991). Receiver psychology and the evolution of animal signals. *Animal Behaviour*, 42:1–14.
- Hamilton, W. D. (1964). Genetic evolution of social behavior. *Journal of Theoretical Biology*, 7(1):1–52.
- Harvey, I. (1996). Untimed and misrepresented: connectionism and the computer metaphor. *AISB Quarterly*, 96:20–27.
- Hauser, M. D. (1996). *The Evolution of Communication*. MIT Press, Cambridge, MA.
- Hihara, S., Yamada, H., Iriki, A., and Okanoya, K. (2003). Spontaneous vocal differentiation of coo-calls for tools and food in japanese monkeys. *Neuroscience Research*, 45:383–389.
- Krebs, J. R. and Dawkins, R. (1984). Animal signals: Mind-reading and manipulation. In Krebs, J. R. and Davies, N. B., editors, *Behavioral Ecology: An Evolutionary Approach*, chapter 15, pages 380–402. Blackwell Scientific Publications, Oxford, second edition edition.
- Lotem, A., Wagner, R., and Balshine-Earn, S. (1999). The overlooked signaling component in non-signaling behavior. *Behavioral Ecology*, 10:209–212.
- MacLennan, B. and Burghardt, G. M. (1993). Synthetic ethology and the evolution of cooperative communication. *Adaptive Behavior*, 2(2):161–187.
- Marocco, D., Cangelosi, A., and Nolfi, S. (2003). The emergence of communication in evolutionary robots. *Philosophical Transactions to the Royal Society A: Mathematical, Physical and Engineering Sciences*, 361(1811):2397–2421.
- Maynard-Smith, J. and Harper, D. G. (2003). *Animal Signals*. Oxford University Press.
- Mirolli, M. and Parisi, D. (2005). 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. and Parisi, D. (2006). Talking to oneself as a selective pressure for the emergence of language. In Cangelosi, A., Smith, A., and Smith, K., editors, *The Evolution of Language: Proceedings of the 6th International Conference*

- on the Evolution of Language*, pages 214–221. World Scientific Publishing.
- Mirolli, M. and Parisi, D. (2008). How producer biases can favour the evolution of communication: An analysis of evolutionary dynamics. *Adaptive Behavior*, 16(1):27–52.
- Mirolli, M. and Parisi, D. (InPress). Towards a vygotskyan cognitive robotics: The role of language as a cognitive tool. *New Ideas in Psychology*.
- Mitchell, M. (1996). *Introduction to Genetic Algorithms*. MIT Press, Cambridge, MA.
- Noble, J. (1998). Tough guys don't dance: Intention movements and the evolution of signalling in animal. In Pfeifer, R., Blumberg, B., and Kobayashi, H., editors, *Proceedings of the Fifth International Conference on the Simulation of Adaptive Behavior*, Cambridge, MA. MIT Press.
- Noble, J., Di Paolo, E. A., and Bullock, S. (2001). Adaptive factors in the evolution of signalling systems. In Cangelosi, A. and Parisi, D., editors, *Simulating the Evolution of Language*, pages 53–78. Springer, Heidelberg.
- Oliphant, M. (1996). The dilemma of saussurean communication. *Biosystems*, 37(1-2):31–38.
- Parisi, D. (1997). An artificial life approach to language. *Brain and Language*, 59:121–146.
- Proctor, H. C. (1992). Sensory exploitation and the evolution of male mating behaviour: a cladistic test using water mites (acari: Parasitengona). *Animal Behaviour*, (44):745–752.
- Ryan, M. J. (1990). Sexual selection, sensory systems and sensory exploitation. *Oxford Surveys in Evolutionary Biology*, 7:157–195.
- Ryan, M. J., Fox, J. H., Wilczynski, W., and Rand, A. S. (1990). Sexual selection for sensory exploitation in the frog *Physalaemus pustulosus*. *Nature*, 343(6253):66–67.
- Ryan, M. J., Phelps, S. M., and Rand, A. S. (2001). How evolutionary history shapes recognition mechanisms. *Trends in Cognitive Sciences*, 5(4):143–148.
- Ryan, M. J. and Rand, A. S. (1993). Sexual selection and signal evolution: the ghost of biases past. *Philosophical Transactions of the Royal Society of London: Biological Sciences*, 340:187–195.
- Seeley, T. D. (1992). The tremble dance of the honey bee: message and meanings. *Behavioral Ecology and Sociobiology*, 31:375–383.
- Simon, H. A. (1990). A mechanism for social selection and successful altruism. *Science*, 250:1665–1668.
- Steels, L. (1996). Emergent adaptive lexicons. In Maes, P., Mataric, M., Meyer, J., Pollack, J., and Wilson, S., editors, *From Animals to Animats 4: Pro-*

ceedings of the Fourth International Conference On Simulation of Adaptive Behavior, pages 562–567, Cambridge, MA. The MIT Press.

Struhsaker, T. T. (1967). Auditory communication among vervet monkeys (*ceropithecus aethiops*). In Altmann, S. A., editor, *Social Communication among Primates*. University of Chicago Press, Chicago.

Tinbergen, N. (1952). Derived activities: Their causation, biological significance, origin and emancipation during evolution. *Quarterly Review of Biology*, 27:1–32.

Tinbergen, N. (1963). On aims and methods of ethology. *Zeitschrift fur Tierpsychologie*, 20:410–433.

Trivers, R. L. (1971). The evolution of reciprocal altruism. *Quarterly Review of Biology*, 46:35–57.

Zahavi, A. (1975). Mate selection. a selection for a handicap. *Journal of Theoretical Biology*, 53:205–214.