THE EVOLUTION OF COMMUNICATION IN ROBOT SOCIETIES

THESIS SUBMITTED TO

ECOLE POLYTECHNIQUE FÉDÉRALE DE LAUSANNE

FOR THE DEGREE OF

DOCTOR OF PHILOSOPHY IN COMPUTER SCIENCE

BY

Sara Mitri

Thesis committee:

Prof. Babak Falsafi, committee president
Prof. Dario Floreano, thesis director
Prof. Laurent Keller, thesis co-director
Prof. Paolo De Los Rios, referee
Dr. Andy Gardner, referee
Prof. Stefano Nolfi, referee

Lausanne
October 2009

Acknowledgments

It is unfortunate that this thesis can only be attributed to one person, since so many people have contributed to its existence.

Most importantly, I owe this thesis to my two supervisors, Dario Floreano and Laurent Keller. Much of what I have learned and the ways in which I have changed over the last four years is due to their efforts. I thank Dario for his trust in me, for his guidance, and for teaching me to sell my work. I thank Laurent for his untiring feedback and availability, for teaching me to think and write scientifically, and for his support.

I would like to thank the Future and Emerging Technologies (IST-FET) program of the European Community and the Swiss National Science Foundation for funding the project.

For setting up the project and the experiments, and for much of the work with the robots, I thank Stéphane Magnenat. It would have been his name on this thesis had the topic intrigued him as much as me. I would also like to thank Kevin Frugier, Matthieu Bontemps and Alban Laflaquière for setting up the hardware demos and making the videos, and Walter Karlen and Antoine Beyeler for taking beautiful pictures of the robots.

I am particularly indebted to Steffen Wischmann, Markus Waibel, and Peter Dürr, who have always been available to have endless discussions on my results. Apart from their friendship and advice, they have helped me to listen to myself. For many interesting discussions on my work, I would like to thank Julien Hubert, Claudio Mattiussi, Daniel Marbach and Davide Marocco, as well as many people I've met at conferences, meetings and lab visits worldwide.

During the preparation of this thesis, Dario, Laurent, Markus, Pavan Ramdya, Peter and Steffen have given me valuable feedback, while Daniel has provided much help with the figures. I thank them all for taking the time to do so. I would also like to thank the members of the examination committee, Babak Falsafi, Andy Gardner, Stefano Nolfi, and Paolo De Los Rios for evaluating my thesis and attending the examination.

For the great work atmosphere, all the lunches and coffees at Sat, I thank the members of the Laboratory of Intelligent Systems. It has been an immense pleasure to be part of this group. I especially appreciate the help of Dominique Etienne, Anouk Hein, Filomena Jacquier and Michelle Wälti, who made administration so painless. I thank Stéphane, Julien, Peter, Zsuzsa Akos, Jean-Christophe Zufferey and Claudio Bruschini, who managed to share an overheated, blindingly sunny office with me, and Thomas Schaffter for his visits when I had no one to share an office with. I thank the Ants lab for the great weekends in Villars, and their valuable feedback, which always gave me a different perspective on my work.

Many friends have made Lausanne an amazing place to live. To those who taught me to ski, to ice-skate and to climb, who shared coffees, dinners, ceramics courses, cultural events, movienights, film nights and days by the lake with me, I thank you for making my life here so pleasurable. I would especially like to thank Peter for being a rock of support in difficult times. Over the last four years, I have shared my flat with 11 people in total. I thank them all for being my family during this time. I was always so happy to come home. I would also like to thank my friends in Egypt, particularly Lobna Abdel Hadi, Hussein El Buckley, Omar El Afifi and Shahira Selim for being by my side even though I was so far away. Last, but not least, I thank my parents, Safwat and Laila for giving me the freedom and support in all the decisions I have made to arrive where I am today.

Abstract

Communication is fundamental to life on earth. All social organisms, from bacteria to humans, use communicative signals to coordinate their behaviors with members of their own and other species. Despite its key role in social organization, many questions regarding the evolution of communication are yet to be answered. This is in part due to the difficulty of conducting experimental evolution on social species, and the challenges in experimentally manipulating and measuring signaling and response strategies in communicating organisms.

In this thesis, we circumvent these problems by using a system of experimental evolution with groups of foraging robots that could emit and perceive light to communicate. With this system, we have explored how communication can emerge and how different evolutionary conditions can determine the level of reliability of evolving signals.

Our experiments revealed that foraging robots initially produced inadvertent cues providing information to other robots about the location of the food. This resulted in increased foraging efficiency, and consequently, in competition near the food, which drove the co-evolution of signaling and response strategies. The reliability of the resulting communication system was found to depend on the level of relatedness between robots in a group and the level at which they were selected. Robots that were highly related or selected at the group level evolved reliable signals. In contrast, when relatedness between robots in a group was low and selection was acting at the level of the individual, robots were selected to suppress the inadvertent cues produced while foraging. However, because of the effect of mutations, these cues were never completely suppressed and some variability in signaling was maintained.

Because similar co-evolutionary processes should be common in natural systems, our findings explain why communicative strategies are so variable in many animal species when interests between them conflict. They also predict that relatedness will play an important role in the evolution of signal reliability in natural systems of communication. Additionally, our analyses have led us to devise a quantitative measure of signal reliability, which may be applied to measure reliability in natural systems of communication. The results of this study, together with an extensive review of the literature, illustrate how evolutionary robotic systems can be used to explore issues that cannot easily be studied experimentally with living organisms, and thus contribute to our understanding of biological systems.

Keywords: Evolution, communication, cues, information, relatedness, reliability, robots, signals.

Zusammenfassung

Kommunikation ist ein grundlegendes Element für Leben auf der Erde. Alle sozialen Lebewesen, von den Bakterien bis zu uns Menschen, verwenden kommunikative Signale, um ihr Verhalten zu koordinieren. Trotz dieser zentralen Rolle von Kommunikation für die Organisation sozialen Lebens, bleiben viele Fragen zu ihrer Evolution bis heute unbeantwortet. Ein Hauptgrund dafür ist die Schwierigkeit, evolutionäre Experimente mit sozialen Arten auszuführen und die verschiedenen Signal- und Reaktionsstrategien zu manipulieren und zu messen.

In der vorliegenden Dissertation überwanden wir dieses Problem mit Hilfe eines Systems zur experimentellen Evolution von Gruppen von Robotern. Die nach Nahrung suchenden Roboter konnten Lichtsignale erzeugen und wahrnehmen. Mittels dieses Systems konnten wir die evolutionäre Entstehung von Kommunikation, sowie den Einfluss verschiedener evolutionärer Bedingungen auf die Zuverlässlichkeit der evolvierten Signalen untersuchen.

Unsere Experimente haben gezeigt, dass Roboter zunächst unbeabsichtigte Indikatoren aussendeten, die von anderen Robotern als Informationsquelle genutzt werden konnten. Dies steigerte die Effizienz bei der Nahrungssuche und führte in der Folge zu erhöhter Konkurrenz bei der Futterquelle, was eine Koevolution von Signalen und Reaktionen in Gang setzte. Die Zuverlässigkeit der resultierenden Kommunikationssysteme hing einerseits vom Verwandschaftsgrad zwischen den Robotern und andererseits von der Selektionsebene ab. Roboter, die untereinander einen hohen Verwandschaftsgrad aufwiesen oder die der Gruppenselektion unterlagen, evolvierten zuverlässige Signale. Dagegen wurden Roboter, die einen niedrigen Verwandschaftsgrad aufwiesen und der Individualselektion unterlagen, selektiert, unbeabsichtigte Indikatoren bei der Nahrungssuche zu unterdrücken. Allerdings konnten diese Signale aufgrund der, durch Mutationen verursachten, hohen Variabilität nie volkommen unterdrückt werden.

Da wir erwarten können, dass ähnliche koevolutionäre Vorgänge in natürlichen Systemen ablaufen, erklären unsere Resultate, warum Kommunikationssysteme in vielen Tierarten, deren Interessen in Konflikt stehen, so variabel sind. Außerdem, können wir aus unseren Ergebnissen folgern, dass Verwandschaft einen bedeutenden Einfluss auf Signalzuverlässigkeit hat. Zusätzlich führte uns die Analyse unserer Experimente dazu, eine quantitatives Maß für Signalzuverlässigkeit auszuarbeiten, das auch in der Analyse natürlicher Kommunikationssysteme angewandt werden kann. Die Ergebnisse dieser Dissertation, sowie eine ausführliche Literaturübersicht, zeigen, wie evolutionäre robotische Systeme die Möglichkeit eröffnen, Themen zu untersuchen, die mit biologischen Lebewesen nur schwer experimentell zugänglich sind, und damit zum besseren Verständnis biologischer Systeme beitragen können.

Schlüsselwörter: Evolution, Kommunikation, Indikator, Information, Verwandschaft, Signalzuverlässlichkeit, Roboter, Signal.

Contents

1

Introduction

In this opening chapter, we introduce the topic of animal communication and the important open questions surrounding its evolution. We then outline the method we are using to explore these questions, and place it in its historical context of interdisciplinary work between biology and robotics. The chapter ends with an overview of the remaining chapters of this thesis.

Communication[1] appears to be ubiquitous and indispensable to all social organisms from humans to microbes. In these very different social organisms, communication plays a variety of different roles, from coordinating group activities, such as navigation or foraging, to determining social hierarchies, choosing mates or deciding on territorial boundaries. Moreover, communicative **signals** are transmitted over a wide range of channels, including visual, acoustic, chemical and tactile.

Because of the transient nature of many of these signals, they have left no fossil record, which makes it difficult to study how these various forms of communication evolved. Understanding how communication evolves in the first place, and once it has evolved, how evolutionary forces can mold it into one form or another is an important and fascinating problem: Important, because it may lead to a deeper and more general understanding of **social behavior** in a wide range of organisms; and fascinating, due to the complex dynamics that arise through the co-evolution of signalers, receivers and the roles they play in a given social interaction. These dynamics differentiate the problem of the evolution of communication from other types of social behavior that have been widely studied, such as **cooperation**.

1.1 Open questions on the evolution of communication

Because communicative signals do not fossilize, one is restricted to studying natural communication systems that are already in place. An important question that thus remains difficult to study is how communication systems emerge as a product of social interactions. The evolutionary origins of communication represent somewhat of a paradox, because it is unclear why an individual would evolve to signal if no receivers have evolved a **response** to the signal. Conversely, if no signal exists, there is nothing for individuals to evolve a response to.

To resolve this paradox, current theory suggests that signals may originate for reasons other than their communicative function, only to later evolve into signals (Maynard Smith and Harper, 2003). For example, the behavior of an injured bird can serve as a **cue** to a predator that the bird is vulnerable and a

---

[1] Terms written in bold font the first time they appear in the thesis are defined in the glossary.

good object of attack. If predators consistently respond to the bird's behavior when injured by attacking it, the bird might evolve to fake injury to distract a predator from attacking its offspring (Ristau, 1991). The bird's behavior thus evolves into a signal that alters the predator's attack behavior. In other words, it can be expected that communication systems originate through inadvertently produced **information** that evolves into a signal. In this thesis, we provide the first study demonstrating how this process can take place (see chapter 3).

Another key question in signaling theory that has been subject to much debate concerns the evolution of reliable signals (Searcy and Nowicki, 2005). Although researchers started out thinking that signaling systems would be overwhelmingly **honest**, it soon became clear that this assumption was overly simplistic (Johnstone, 1998b; Searcy and Nowicki, 2005). In 1978, Dawkins and Krebs argued that if mutant **deceivers** could achieve higher **fitness** than their honest conspecifics, they would quickly invade the population through their more numerous deceptive offspring. They concluded from this that deceptive communication should be more widespread in nature than had previously been thought (Dawkins and Krebs, 1978).

Since then, many animals have been shown to communicate deceptively. In an intriguing example, great tits have been shown to produce false alarm calls to scare away conspecifics and gain access to an overcrowded food source (Møller, 1988). Similar behavior has been observed in two species of fly-catching birds (Munn, 1986) and in tufted capuchin monkeys (Wheeler, 2009). It has also been found that ravens can actively mislead conspecifics into searching for food in empty caches to distract them from finding the actual food location (Bugnyar and Kotrschal, 2004). In these example, deceptive signals result in receiving individuals being worse off than if they had not perceived the signal. Many other signals, although not deceptive (i.e., leading to lower **performance** than if the signal had not been perceived), are **unreliable**, leading receivers to achieve lower performance than if the signal were random (i.e., not subject to selection). These include bluffing or exaggerating individual strength or quality (Adams and Caldwell, 1990; Candolin, 1999; Lailvaux et al., 2008), mimicry, which is common in plants and insects (Schiestl et al., 1999; Matsuura, 2006) and camouflage (Stevens and Merilaita, 2009).

Given all these examples, research has focused on understanding how informative communication can evolve at all if signalers stand to gain by deceit

(Johnstone, 1998b; Diggle et al., 2007a). A number of factors have been predicted to ensure the evolution and stability of honest communication. One of the first factors that have been argued to influence **signal reliability** is the **cost** of signal production (Zahavi, 1975, 1997; Grafen, 1990). In a mate-choice scenario, for example, where signals are used to indicate male quality to a choosy female, the "handicap principle" states that a signal of quality will be honest if signal production is less costly to higher quality individuals (Zahavi, 1975, 1997). In addition to sexual displays, the handicap principle also applies to threat displays that provide information regarding fighting ability (Andersson, 1980; Grafen, 1990; Számadó, 2000), or pursuit-deterrent signals emitted by prey to discourage predators from a pointless chase (Caro, 1994). Maynard Smith (1991) and Johnstone and Grafen (1992) extended the handicap principle to include interactions where a resource holder must decide whether to donate a resource to a potentially needy partner, a so-called "Sir Philip Sidney" game. The predictions of their models apply to various signals of need between parents and their offspring (Maynard Smith, 1991; Johnstone and Grafen, 1992; Bergstrom and Lachmann, 1998).

Another mechanism that is expected to play a role in maintaining honest communication is **relatedness** between communicating individuals (Brown and Johnstone, 2001; Keller and Surette, 2006; Diggle et al., 2007a). Relatedness is predicted to be particularly important when signaling constitutes an **altruistic** act that is costly to the signaler, but benefits the receiver. This is expected to be the case when signals provide information about the state of the signaling individual (Bergstrom and Lachmann, 1998) or about the state of the environment, such as alarm calls or signals used for collective foraging. In these scenarios, signaling individuals can increase their **inclusive fitness** by helping related individuals – that carry some of the signaler's genes – survive to the next generation by providing them with useful information (Hamilton, 1964; Keller and Surette, 2006; Lehmann and Keller, 2006; West et al., 2007a). This idea is supported by empirical data (Diggle et al., 2007b; Griesser, 2009), which suggest that relatedness is important in ensuring honest communication. However, it has been difficult to measure the level of honesty of signals in these systems, and to manipulate experimental conditions to fully test the effect of relatedness on signal evolution. In addition, the role of relatedness in the evolution of signals has been subject to very little theoretical research (Brown and Johnstone, 2001). Consequently, although relatedness has been found to play an important role in the evolution of social behavior more generally (Lehmann and Keller, 2006; West et al., 2007a), it is largely unclear how it would influence the evolution of communication.

Similarly, social evolution theory predicts that the level at which selection acts may also influence the evolution of cooperative behavior (Hamilton, 1975; Frank, 1998; Lehmann and Keller, 2006; West et al., 2007a). This is because selecting between individuals in a population leads to competition between them, whereas selection between groups reduces within-group competition (Frank, 1998). This effect has been shown experimentally by manipulating the level of selection in social organisms within and between evolving groups (Rainey and Rainey, 2003; Griffin et al., 2004). However, these mechanisms remain to be explored in the context of communication.

Two of the factors that have been central to our understanding of social evolutionary theory, relatedness and the level of selection, have thus remained underexplored in the context of communication, particularly in communication systems where sharing information is costly to the signaler and beneficial to the receiver. In this thesis we will attempt to shed some light on these issues. We hypothesize that both relatedness, as well as the level at which selection operates will influence the evolutionary dynamics of signal reliability in a social system where signaling constitutes a costly act of information sharing.

1.2 Studying the evolution of communication

A range of different approaches have been used to study the evolution of communication. One possibility is to conduct empirical research with living organisms, either by studying one particular system in depth, or by comparing communicative behavior between related taxa or between populations of the same species. Such observational work has inspired the formulation of many of the important theories of communication, such as the handicap principle discussed above (Zahavi, 1975), theories on receiver biases (Endler and Basolo, 1998), or on signal ritualization (Maynard Smith and Harper, 2003). Unfortunately, however, the produced theories can only be shown to hold by painstakingly collecting data from a variety of taxa. A powerful alternative is to conduct experimental evolution in species with an elaborate social organization (Griffin et al., 2004; Fiegna et al., 2006). This allows researchers to observe the evolutionary process in action, and thus to systematically test the proposed theories. Although such studies have been beneficial in addressing a variety of evolutionary questions, it remains difficult to manipulate certain parameters that are relevant to communication, such as signal costs, and to measure their effects. Furthermore, highly social species are not amenable to such experiments because they typically have long generation times and are difficult to breed in the laboratory.

Another approach involves building abstract models of the communication system by boiling it down to a few key components, thus avoiding the complexity of the real world (Murray, 2007). The question that this raises is how to determine these key components and which aspects of the real system to include in the model. Perhaps because the answer to this question is unclear, researchers have used a number of modeling techniques that differ largely in their level of abstraction. The more abstract the model, the easier it is to manipulate in order to study the effects of different parameters. In addition, it is straightforward to replicate results obtained using highly abstract models and to compare them to similar studies. On the down side, many of the factors that are typically neglected in such diluted models have been shown to play an important role in determining the evolutionary pathways of social systems. Interactions between individuals, for example, are generally not taken into account in abstract mathematical models (Judson, 1994).

Instead, such interactions are studied using **game theory** (Maynard Smith and Price, 1973), which is currently perhaps the most widespread approach to studying signaling theory (Johnstone, 1998b; Maynard Smith and Harper, 2003; Searcy and Nowicki, 2005). **Game-theoretical models**, in turn, rarely take environmental effects, such as noise into account and often ignore differences between individuals (Johnstone, 1998b; Rowell et al., 2006), which are all factors that have been shown to play an important role in shaping communication systems (Ryan and Wilczynski, 1991; Johnstone, 1994; Wollerman and Wiley, 2002).

One more layer of abstraction can be removed by using **individual-based models** (IBMs), which incorporate more details regarding individual differences, interactions between individuals, and between individuals and their environment (DeAngelis and Mooij, 2005). Such models have become widespread in studying the evolution of human language and animal communication (see section 1.3).

In this thesis, following the rationale that the complexity of real systems may be important in forming communication systems, we go one step further. We use mobile robots as model "organisms", and evolve their behavior using a system of artificial evolution (Holland, 1975). This technique can essentially be seen as an individual-based model in which physical properties of the system are taken into account. Although much of the work presented here is carried out using accurate physics-based computer simulations of these robots, we argue (mainly in chapters 5 and 6) that the use of robots as a base for these simulations makes significant contributions toward our understanding of communication systems and the forces that govern their evolution.

Many studies addressing evolutionary questions assume that evolution constitutes a process of optimization. It has been shown that this simplification neglects many properties of real evolutionary systems, such as genetic and phenotypic variation that are maintained through mutation-selection, which can have a profound influence on evolutionary dynamics (Houle, 1992; Lenormand et al., 2009) and thus on the evolution of social behavior (McNamara et al., 2004, 2008; Robinson et al., 2008). This view is supported by the results presented in in chapters 2, 3 and 4.

1.3 Individual-based models and the evolution of communication

The evolution of communication is one of the earliest subjects to be explored using IBMs. Over the last 20 years, IBMs have been used to explore a wide range of different questions (see Wagner et al. (2003); Nolfi (2005); Kirby (2006) for reviews). For example, in a study using simulated robots that could only perceive the environment using eight distance sensors, Quinn (2001) explored how dedicated communication channels can evolve between two individuals to coordinate their movement. Other studies have instead assumed that continuous communication channels, such as sound, already exist and have investigated how robots (or simulated robots) could evolve to coordinate the frequencies of emitted sounds to construct communication systems that allow them to solve a collective task (Marocco and Nolfi, 2006; Wischmann and Pasemann, 2006). This line of work raises questions on how continuous perceptual information can be categorized into discrete entities, and how a population of communicating individuals can then agree on a mapping between perceptual categories and discrete signals (i.e., “words”). A number of studies have explored this process of “symbol grounding” (Harnad, 1990) and vocabulary formation, often using robots (e.g., Billard and Dautenhahn (1999); Steels and Vogt (1997); Steels (1999)), but in some cases also with software models (e.g., Loula et al. (2008); Cangelosi (2001)). By understanding how such vocabularies emerge, researchers have been addressing questions regarding human language (Kirby, 2006). Similarly, other research has explored the role of learning in shaping communication systems (e.g., MacLennan and Burghardt (1993); Ackley and Littman (1994); Noble and Cliff (1996); Billard and Dautenhahn (1999); Mirolli and Parisi (2005)), and how linguistic units can organize themselves into syntactical structures (e.g., Batali (1994); Kirby (2000); Cangelosi (2001); Kirby et al. (2008); Steels (2005)).

A number of studies also tackle questions that are more directly relevant to this thesis. In particular, researchers have been interested in determining the conditions that allow for the evolution of stable systems of honest communication. In one of the earliest studies on this topic, Werner and Dyer (1991) showed how two types of software agents, blind mobile males and immobile sighted females, evolved communicative strategies that allowed them to locate each other in a simulated arena. In their study, both signaler and receiver were rewarded for an encounter, and thus had a common interest in solving the task. Similarly, in an experiment where agents could share information about their local environment, MacLennan and Burghardt (1993) showed that stable communication could evolve if both signaler and receiver were directly rewarded for successful information transfer.

To explore whether communication could evolve to be stable even if a signaler had no direct benefit in transmitting information to a receiver, Ackley and Littman (1994) and Oliphant (1996) conducted two independent studies, both suggesting that imposing a spatial structure on the population of agents (i.e., where an individual’s offspring were placed closed to it on a two-dimensional lattice) could lead to the evolution of stable communication. They argued that the spatial structure increased the likelihood that an individual would communicate with its kin than with unrelated individuals, thus selecting for cooperative behavior (Hamilton, 1964). However, because these studies did not systematically explore whether relatedness was indeed the reason for the evolutionary stability of the communication system, Di Paolo (1999) has argued that their observations may have been a product of other properties of the spatial setup (see also (Wakano et al., 2009), e.g.), rather than relatedness. In addition, theoretical work has shown that although population viscosity increases the probability of interacting with relatives, thus promoting cooperative behavior, this effect is cancelled out by an increase in local competition between individuals, which inhibits cooperation (Taylor, 1992; Griffin and West, 2002).

The role of relatedness on the stability of evolving communication systems was only explored much later by Mirolli and Parisi (2005), who conducted an experiment on foraging agents that could evolve to transmit information allowing other agents to distinguish edible from poisonous objects. The results of their study indicated that the average "language quality" of agents improved as the probability of an individual communicating with its siblings was experimentally increased. However, because of a lack of statistical analysis, their results can only be taken to be indicative, and require further testing. In addition, Ackley and Littman (1994) and Mirolli and Parisi (2005) claim to study **altruistic communication**, although there was no cost to producing signals.

1.4 In this thesis

Including this introductory chapter, the thesis is organized into six chapters. In the second through fourth chapter, we present the results of experiments carried out with our evolutionary robotic model of communication to explore the factors that shape the evolution of communication.

The second chapter introduces the model and tests the hypothesis stated in section 1.1 by exploring how relatedness and the level of selection influence the evolution of communication. Our findings indicate that qualitatively similar communication systems evolve if relatedness is high or robots are selected at the level of the group, i.e., competition between individuals in a group is low. However, when robots in a group are unrelated and the competition between them is high, the evolutionary process results in what looks like a deceptive communication system.

In chapter 3, we explore the evolved communication system in unrelated groups in more detail. The results of this analysis revealed that although robots did not evolve to deceive each other, they were concealing information from other robots by suppressing signal production. This strategy is found to be unreliable compared to the other communication strategies presented in chapter 2.

To quantify the reliability of the signaling systems observed in chapters 2 and 3, we devise a measure of signal reliability in chapter 4, and use it to explore how gradually increasing within-group relatedness affects reliability. In addition, we discuss other methods that are widespread in studying reliability in signaling systems and show how the measure of reliability introduced in this chapter can contribute to more general comparisons between communication systems.

In chapter 5, we take a wider perspective and evaluate the use of robots not only in studying communication, but social behavior in general. By reviewing a number of studies in which robots – or simulations thereof – are used, we explore the advantages and disadvantages of robotic models and draw lessons from these studies to formulate an agenda for the future of the field.

Chapter 6 concludes the thesis with a summary of the main achievements, a broader perspective of its implications both for biology and robotics and an outlook towards future work.

2 Evolutionary Conditions for the Emergence of Communication

In this chapter[1], we will introduce the methodology used throughout this thesis, where we conduct experimental evolution on groups of foraging robots that could produce visual signals to provide information on the location of a food source. We then show how this system can be used to test the hypothesis formulated in chapter 1 that kin structures and the level of selection will influence the evolution of communication.

---

[1] This chapter is based on Floreano, D., Mitri, S., Magnenat, S., and Keller, L. (2007). Evolutionary conditions for the emergence of communication in robots. *Current Biology*, 17:514–519.

Although the evolution of signals has been studied in a range of scenarios where the signaler stands to gain from manipulating the receiver's behavior (Maynard Smith and Harper, 2003), it remains unclear why individuals would convey useful information to conspecifics if it incurs costs, yet no obvious benefits to them. Previous theoretical studies suggest that the kin structure of groups (Hamilton, 1964; Maynard Smith, 1991; Johnstone and Grafen, 1992) and the scale at which competition occurs (i.e., the level of selection) (Keller, 1999; West et al., 2002) will play a key role in determining whether individuals will evolve to share costly information. To determine the evolutionary conditions under which such communication systems can evolve, we explore the influence of these two factors in an artificial selection experiment involving groups of foraging robots that can use light to communicate.

2.1 Methods

2.1.1 Experimental setup

At the beginning of each experiment, a group of ten s-bot robots (Fig. 2.1B, Mondada et al. (2004)) was randomly placed in a 3m × 3m foraging arena that contained a food and a poison source, each placed at 1m from one of two opposite corners (see Fig. 2.1A). The 10cm-radius food and poison sources constantly emitted red light that could be seen by robots in the whole foraging arena. A circular piece of gray paper with a radius of 25cm was placed under the food source and a similar black paper under the poison source. A robot could thus only discriminate between the two sources using its floor sensors once it was driving over the colored paper. In addition to these sensors, the robots were equipped with two tracks that could independently rotate in both directions, a translucent ring around the body that could emit blue light, and a 360° vision system that could detect the amount and intensity of red and blue light (see Fig. 2.1B).

The robots had a sensory-motor cycle of 50ms during which they used a neural controller to process the visual information and the floor sensor input to set the direction and speed of the two tracks and control the emission of

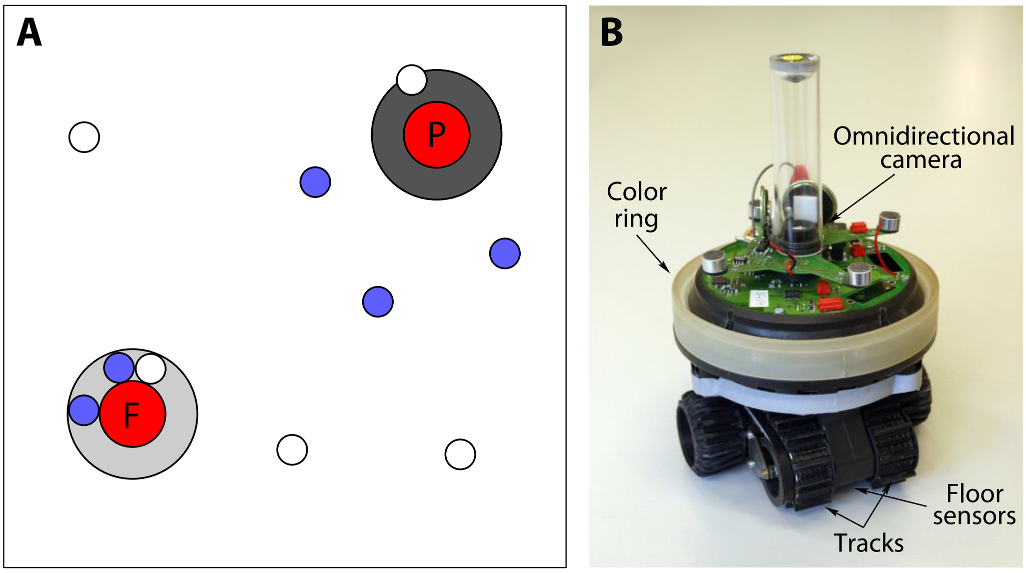

Figure 2.1: Experimental setup. (A) A food and poison source, both emitting red light in the square arena. Robots (small circles) can distinguish the two by sensing the color of the circles of paper using their floor sensors. (B) The 6cm-radius robot used for the experiments is equipped with two tracks to drive, an omni-directional (360°) vision camera, a ring of lights used to emit blue light and floor sensors to distinguish food and poison sources.

blue light accordingly during the next 50ms cycle. During each cycle, a robot gained one performance unit if it detected food with its ground sensors and lost one performance unit if it detected poison. For each group of ten robots, we conducted ten foraging trials, each lasting one minute. The performance of each robot at the end of a trial was computed as the sum of performance units obtained during that trial (1200 sensory motor cycles of 50ms). The robot's overall performance was quantified as the sum of performance units over all ten trials, normalized over the total number of time units (see Appendix A). Group performance was equal to the average performance of all robots in the group.

Under such circumstances, foraging efficiency can potentially be increased if robots transmit information on food and poison location. However, such communication may also incur direct costs to the signaler because it can result in higher robot density and increased competition and interference nearby the food[2]. Thus, although beneficial to other group members, signaling of a food location can constitute a costly act (Hamilton, 1964; Lehmann and Keller, 2006) because it decreases the food intake of signaling robots. This setup thus mimics the natural situation where communicating almost invariably incurs costs in terms of signal production or increased competition for resources (Zahavi, 1997).

2.1.2 Neural controller

The control system of each robot consisted of a feed-forward neural network with 11 input and three output neurons. Each input neuron was connected to every output neuron with a synaptic weight representing the strength of the connection (Fig. 2.2). One of the input neurons was devoted to the sensing of food and the other to the sensing of poison. Once a robot had detected the food or poison source, the corresponding neuron was set to 1. This value decayed to 0 by a factor of 0.95 every 50ms and thereby provided a short-term memory even after the robot's sensors were no longer in contact with the gray and black paper circles placed below the food and poison. Eight of the remaining neurons were used for encoding the 360° visual-input image, which was divided into four sections of 90° each. For each section, the average of the blue and red channels was calculated and normalized within the range of 0 and 1 such that one neural

---

[2] Spatial constraints around the food source allowed a maximum of eight robots out of ten to feed simultaneously and resulted in robots sometimes pushing each other away from the food.

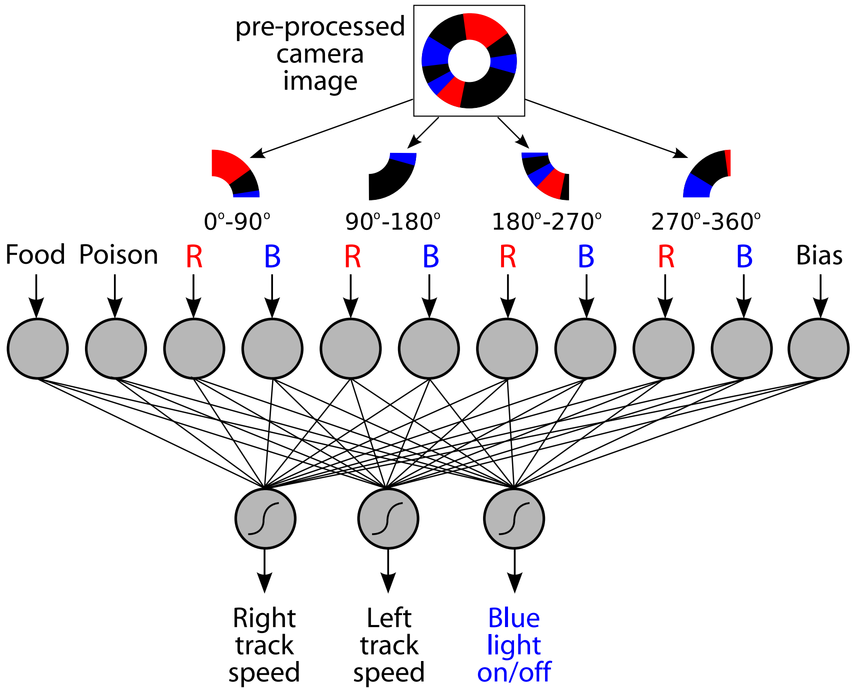

Figure 2.2: Neural network architecture. The first two input neurons are activated when feeding on either food or poison. The omnidirectional camera image is pre-processed to filter out red and blue channels, divided into sections and input to the neural network as fractions of red or blue in each section (between 0 and 1). Three output neurons with sigmoid, asymptotic activation receive weighted input from the 11 input units, which encode the speed of the tracks and whether to emit blue light.

input was used for the blue and one for the red value. Finally, a "bias" input neuron, constantly set to -1, was used to enable the robot to produce some behavior even if it perceived no sensory data.

The activation of each of the output neurons was computed as the sum of all inputs multiplied by the weight of the connection and passed through the continuous $tanh(x)$ function (i.e., their output was between -1 and 1). Two of the three output neurons were used for controlling the two tracks, where the output value of each neuron gave the direction of rotation (forward if > 0 and backward if < 0) and velocity (the absolute value) of one of the two tracks. The third output neuron determined whether to emit blue light, which was the case if the output was greater than 0.

The way in which a robot responded to environmental stimuli depended on the values of the 33 synaptic weights connecting the 11 input neurons with the three output neurons. These synaptic weights were encoded in artificial genomes (Fogel et al., 1990; Nolfi and Floreano, 2001). A genome, which consisted of a list of these weights, was assigned to each robot. Each synaptic weight was encoded in 8 bits, such that each weight could take one of 256 values that were mapped onto the interval $[-1, 1]$. The total length of the genetic string of an individual was therefore 8 bits × 11 input neurons × 3 output neurons (i.e., 264 bits).

2.1.3 Artificial selection

A population of 100 groups of 10 robots each was used to conduct selection experiments over 500 generations. In the first generation, the genomes of the robots were randomly generated genomes, corresponding to randomly wired neural controllers with no information about how to move and identify the food and poison sources. At each generation, the performance of the robots in the population was determined using the foraging task described in section 2.1.1, which allowed for the selection of the best robots, which were subjected to mutation, sexual reproduction, and recombination to form the population of robots for the next generation (Fig. 2.3) (Holland, 1975). The method of selection depended on the experimental treatment in question.

2.1.4 Experimental treatments

Studying why group members convey information when it incurs costs requires consideration of the kin structure of groups (Hamilton, 1964; Maynard Smith, 1991; Johnstone and Grafen, 1992) and the scale at which cooperation and competition occurs (level of selection) (Keller, 1999; West et al., 2002). We therefore chose two kin structures (low and high relatedness) and two levels of selection (individual- and group-level regimes). There were thus four treatments: high relatedness with group-level selection, high relatedness with individual-level selection, low relatedness with group-level selection, and low relatedness with individual-level selection (Fig. 2.4).

For each of the four treatments, selection experiments were repeated in 20 independent selection lines (replicates of populations with newly generated genomes), each consisting of 100 groups of 10 robots. In the individual-level selection treatment, we selected the best 20% of individuals from the population

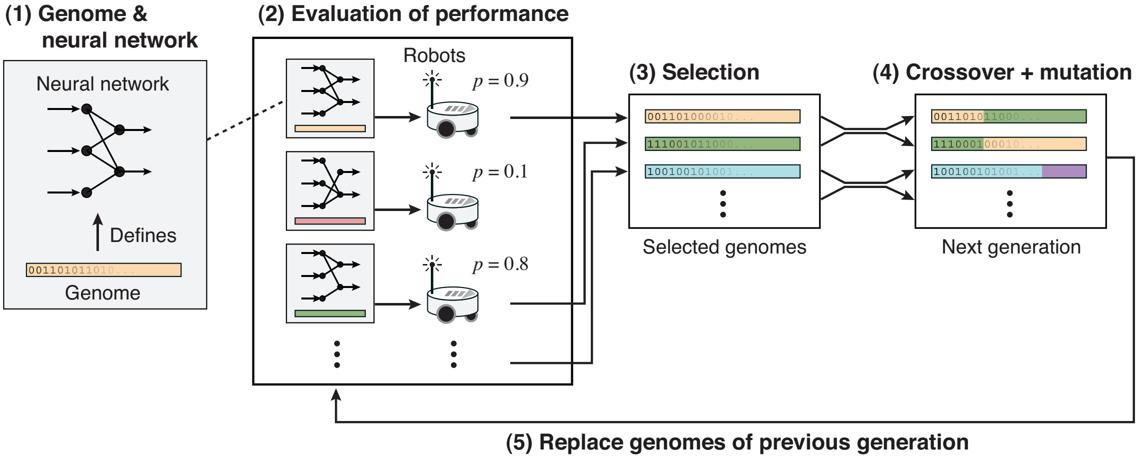

Figure 2.3: Illustration of the process of artificial selection in robots. The performance of each robot in the population (2) is calculated by testing the robot driven by the neural controller encoded in its genome (1) on a given task. The robots in a population are then ranked, selected (3), mutated and recombined (4) to form the genomes of the subsequent generation (5). Credit: Daniel Marbach.

of 1000 robots (Fig. 2.4). This selected pool of 200 robots was used for creating the new generation of robots. To form groups of related individuals $r = 1$, we randomly created (with replacement) 100 pairs of robots. A crossover operator was applied to their genomes with a probability of 0.2 at a randomly chosen point, and one of the two newly formed genomes was randomly selected and subjected to mutation (probability of mutation 0.01 for each of the 264 bits). The other genome was discarded. This procedure led to the formation of 100 new genomes that were each cloned ten times to construct 100 new groups of 10 identical robots. To form groups of unrelated individuals $r = 0$, we followed the same procedure but created 1000 pairs of robots from the selected pool of 200 robots. The 1000 new robots were randomly distributed among the 100 new groups.

In the group-level selection treatment, we followed exactly the same procedure as in the individual-level selection treatment, but the selected pool of 200 robots was formed with all of the robots from the best 20% of the 100 groups (Fig. 2.4).

Robots could communicate the presence of food or poison by producing blue light that could be perceived by other robots (light production was not costly).

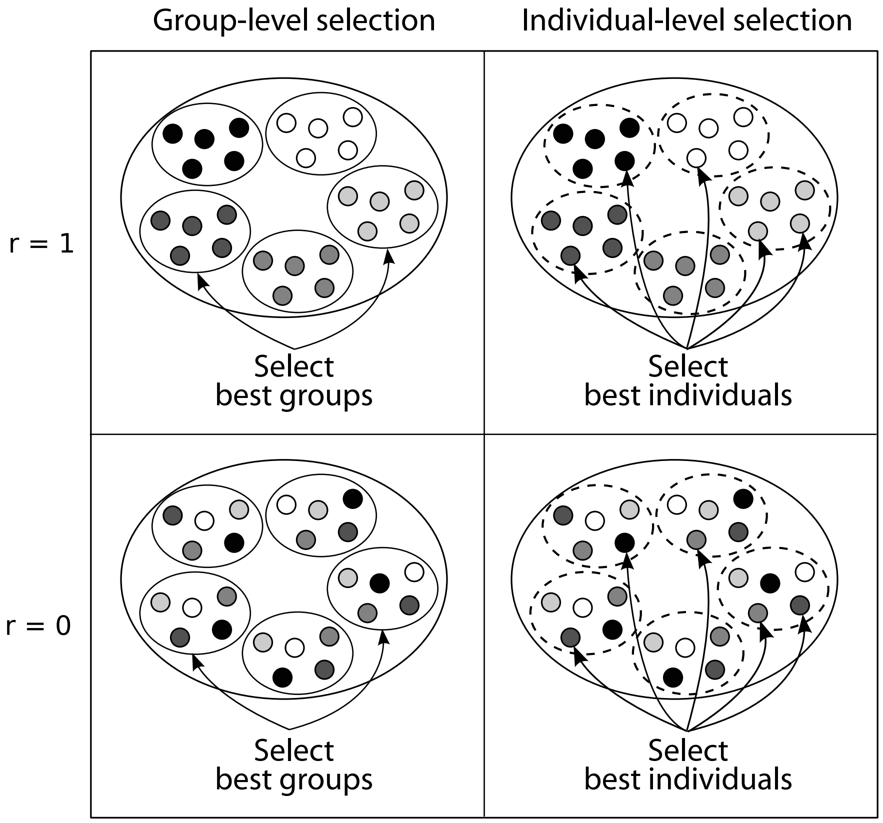

Figure 2.4: Illustration of the group composition and selection regime in the four treatments.

For each treatment, we determined whether communication evolved and quantified the benefits of communication by comparing group performance with control groups where robots were experimentally prevented from communicating (i.e., the blue lights were disabled).

All methods concerning measuring performance, signaling and response behaviors, as well as the tools used for statistical analysis are outlined in Appendix A.

2.1.5 Simulated vs. real robots

All experiments were conducted with a physics-based simulator that accurately models the dynamical properties of real robots. At the end of the experiments, we were able to successfully implement the evolved genome in real robots (Fig. 2.1) that displayed the same behavior observed in simulation, demonstrating that the physics-based simulations allowed us to mimic the behavior of real robots. A discussion of the merits and drawbacks of this approach is provided in chapter 5.

2.2 Results

2.2.1 Robots evolve to forage efficiently

In the control groups where robots could not emit blue light, foraging efficiency greatly increased over the 500 generations of selection (Fig. 2.5A). In each of the four experiments, robots evolved the ability to rapidly localize the food source, move in its direction, and stay nearby (more than half the robots found the food source within the first 30s). Both the degree of within-group relatedness and the level of selection significantly affected the overall performance of groups (Kruskal-Wallis test, $df = 79$, $P < 0.001$, Fig. 2.6). Groups where robots were highly related and subjected to group-level selection were more efficient than the three other types of groups (Mann-Whitney test, $df = 39$, all $P < 0.001$). The two treatments with individual-level selection led to intermediate performance values (non-significantly different from each other, $P = 0.39$, but different from the two other treatments, both $P < 0.001$). The lowest performance was achieved by robots in the low relatedness/group-level selection treatment with performances significantly lower than in all other treatments (all $P < 0.001$). This variation of

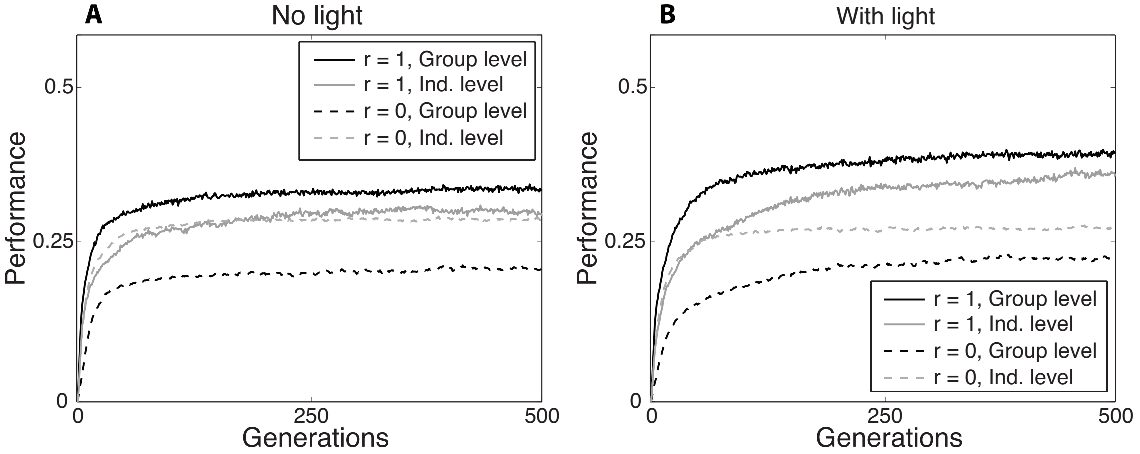

Figure 2.5: Mean performance (A) in control groups where robots could not emit blue light (20 replicates per treatment), and (B) in groups where robots could emit blue light (20 replicates per treatment).

performances in the control condition where robots could not emit blue light reflects differences in selection efficiency among the four treatments (Waibel et al., 2009).

In groups where robots could produce blue light, foraging efficiency also greatly increased over the 500 generations of selection (Fig. 2.5B). Importantly, the ability to emit blue light resulted in a significantly greater group efficiency compared to control experiments in three out of the four treatments (Fig. 2.6). An analysis of the robots' behavior revealed that this performance increment was associated with the evolution of effective systems of communication.

2.2.2 Cooperative signaling strategies

In groups of related robots with group-level selection, two distinct communication strategies evolved. In 12 of the 20 evolutionary replicates, robots preferentially produced light in the vicinity of the food, whereas in the other eight, robots tended to emit light near the poison (see Fig. 2.7 and 2.8). The response of robots to light production was tightly associated with these two signaling strategies, as shown by the strong positive association between the tendency of robots to be attracted to blue light and the tendency to produce light near the food rather than the poison source across the 20 replicates (Spearman's rank correlation test, $r_S = 0.74$, $P < 0.01$; see Fig. 2.7A). Overall, robots were positively attracted to

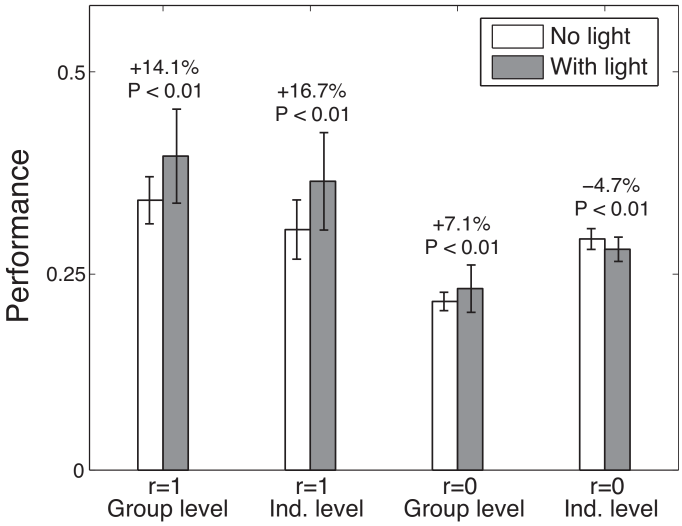

Figure 2.6: Comparison of mean ± SD performance of robots during the last 50 generations for each treatment when robots could versus could not emit blue light (20 replicates per treatment).

blue light in all the 12 replicates where they signaled in the vicinity of the food and repelled by blue light in seven out of the eight replicates where they had evolved a strategy of signaling near the poison. The communication strategy where robots signaled near the food and were attracted by blue light resulted in higher performance (mean ± SD: 0.43 ± 0.05) than the alternate strategy of producing light near the poison and being repelled by blue light (0.33 ± 0.03, Mann-Whitney test, $df = 7$, $P < 0.01$). This is probably because signaling near the food allows robots to signal in a more efficient, sustained way while they feed and because the food signal can easily be detected by other robots, even though the red light of the food is obscured by the robots feeding around it. Interestingly, once one type of communication was well established, we observed no transitions to the alternate strategy over the last 200 generations. This is because a change in either the signaling or response strategy would completely destroy the communication system and result in a performance decrease. Thus, each communication strategy effectively constitutes an adaptive peak separated by a valley with lower performance values (Wright, 1932).

The possibility to produce blue light also translated into higher performance in two other treatments: high relatedness with individual-level selection and low relatedness with group-level selection. In both cases, signaling strategies

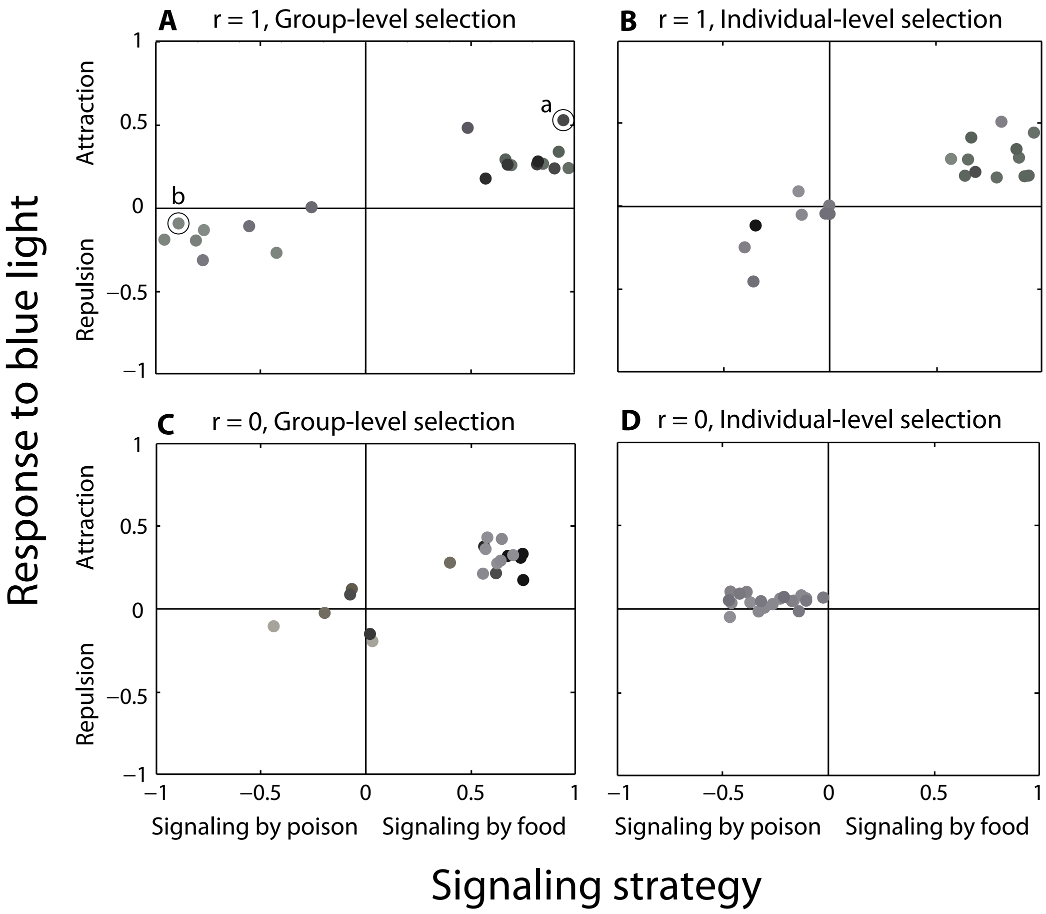

Figure 2.7: Relationship between signaling strategies and response to blue light in all four treatments (panels A to D). Each dot is the average for the 100 groups in one replicate after 500 generations of selection. Positive values for the signaling strategy indicate a tendency to signal close to the food, and negative values indicate a tendency to signal close to the poison. Positive values for the response to blue light indicate an attraction to blue light, and negative values indicate an aversion (see Appendix A for details). The darkness of the points is proportional to the mean performance. The two points “a” and “b” highlighted in panel (A) are examples of the two different signaling strategies of robots. The signaling strategies of the best group in these two experimental runs are shown in Figures 2.8A and 2.8B.

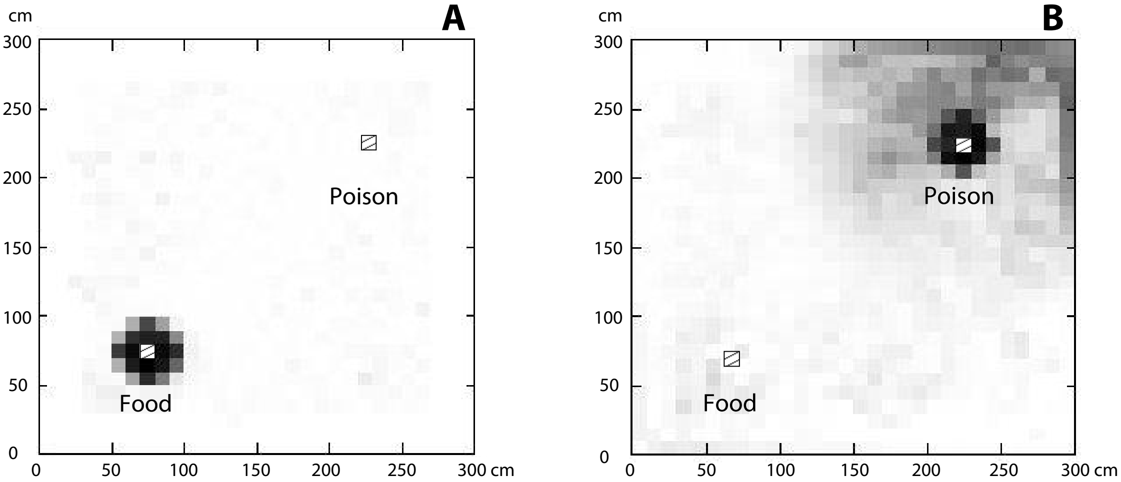

Figure 2.8: Spatial signaling frequency measured in each area of the arena for robots from two groups at generation 500. (A) The group was one where robots signal the presence of food (group with highest performance in population “a” in Figure 2.7A). (B) In this group, robots signal the presence of poison (group with highest performance in population “b” in Figure 2.7A). The darkness of each square is proportional to the amount of signaling in that area of the arena and is normalized by the highest amount of signaling in that particular trial.

evolved that were similar to those observed in the selection experiments with high relatedness and group-level selection (see Fig. 2.7B and 2.7C). There was also a strong positive correlation between the tendency to signal close to food and being attracted to blue light (high relatedness/individual-level selection: $r_S = 0.81$, $P < 0.01$; low relatedness/group-level selection: $r_S = 0.60$, $P < 0.01$). Moreover, in both treatments the strategy of signaling close to food yielded higher performance than the alternative poison-signaling strategy (both $P < 0.01$). However, when robots signaled near the poison, they were less efficient than in the treatments with high relatedness and group-level selection. In the case of high relatedness and group-level selection, robots signaled on average 82.3% of the time when detecting the poison, whereas the amount of poison signaling was only 18.3% (Mann-Whitney test, $df = 6$, $P < 0.001$) in groups with related individuals and individual-level selection and 24% ($P < 0.01$) in groups with low relatedness and group-level selection. Interestingly, the less efficient poison-signaling strategy permitted a switch to a food-signaling strategy in the last 200 generations of selection in three replicates for related robots selected at the individual level and in one replicate for low relatedness robots selected at the group level.

2.2.3 Signaling reduces foraging efficiency in unrelated robots

The only treatment where the possibility to communicate did not translate into a higher foraging efficiency was when groups comprised low-relatedness robots subjected to individual-level selection (Fig. 2.7D). In this case, the ability to signal resulted in a signaling strategy that led to a significant decrease in group performance compared to the situation where robots could not emit blue light. An analysis of individual behaviors revealed that in all replicates, robots tended to emit blue light when far away from the food. However, contrary to what one would expect, the robots still tended to be attracted rather than repelled by blue light (17 out of 20 replicates, binomial-test z-score = 3.13, $P < 0.01$). A potential explanation for this surprising finding is that in an early stage of selection, robots randomly produced blue light, and this resulted in robots being selected to be attracted by blue light because blue light emission was greater near food where robots aggregated. Emission of light far from the food would then have evolved to decrease competition near the food. This explanation is explored in chapter 3.

2.3 Discussion

Our results provide a clear experimental demonstration of how the kin structure and the level of selection jointly influence the evolution of cooperative communication. Under natural conditions, most communication systems are also costly because of the energy required for signal production or increased competition for resources resulting from information transfer about food location (Maynard Smith and Harper, 2003). Thus, cooperative communication is expected to occur principally among kin or when selection takes place at a group rather than an individual level. Consistent with this view, most sophisticated systems of communication indeed occur in animals forming kin groups as exemplified by pheromone communication in social insects (Wilson, 1971; Bourke and Franks, 1995) and quorum sensing in clonal groups of bacteria (Keller and Surette, 2006). Humans are a notable exception, but other selective forces such as direct and reputation-based reciprocity may operate to favor cooperation (Nowak and Sigmund, 2005) and costly communication.

This study demonstrates that sophisticated forms of communication including cooperative communication and communication that decreases performance of receivers can evolve in groups of robots with simple neural networks. Importantly, our results show that once a given system of communication has evolved, it may constrain the evolution of more efficient communication systems because it would require going through a stage where communication between signalers and receivers is perturbed. This finding supports the idea of the possible arbitrariness and imperfection of communication systems, which can be maintained despite their suboptimal nature. Similar observations have been made about evolved biological systems (Jacob, 1981), which are formed by the randomness of the evolutionary selection process, leading, for example, to different dialects in the language of the honey-bee dance (Von Frisch, 1967).

Finally, these experiments provide the first example in this thesis that the evolutionary principles governing the evolution of social behavior also operate in groups of artificial agents subjected to artificial selection, indicating that the use of evolutionary robotics may contribute to advancing the field of evolutionary biology.

3 Evolution of Communication when Interests Conflict

In chapter 2, we found that when robots were unrelated and competing for a chance to reproduce, they evolved a strategy where they were less likely to emit blue light by the food than by the poison, but were nevertheless attracted to blue light. In this chapter[1], we explore this unintuitive result, starting with a hypothesis that the robots' foraging behavior, which resulted in them clustering around the food, inadvertently provided information to other robots on the location of the food. This would explain the persistent attraction of robots to blue light. Through this analysis, we discuss the role of such cues in the evolution of communication systems where interests between signalers and receivers conflict.

---

[1] This chapter is based on Mitri, S., Floreano, D., and Keller, L. (2009). The evolution of information suppression in communicating robots with conflicting interests. *PNAS*, 106:15786-90.

Cues are thought to be common sources of information in nature. Indeed, in many species, individuals have been shown to monitor each other to decide how to behave (Giraldeau, 1997; Franks, 1999; West et al., 2000; Shuker and West, 2004; Tibbetts and Dale, 2004; Dall et al., 2005; Bshary and Grutter, 2006). For example, when foraging, simply observing the behavior of conspecifics can inform an animal about the location of a source of food (Buckley, 1997; Galef and Giraldeau, 2001). In many situations, producing inadvertent cues will also affect an individual's own fitness and should thus be under selection, with the consequence that cues providing inadvertent social information should evolve into signals. Importantly, selection on inadvertent cues may frequently take the form of decreasing the social information provided. An example of this would be birds living in a roost trying to hide information from other group members about a food source they have discovered (Bugnyar and Kotrschal, 2002).

Similarly, we may expect competing robots that produce inadvertent cues to be selected to decrease such information. In the following sections, we explore the hypothesis that the unrelated robots selected at the level of the individual described in the previous chapter were producing inadvertent cues, and that selection to reduce this information can explain the surprising findings of stable signaling that was detrimental to performance.

In addition to testing this hypothesis, this study represents a first exploration of the role of inadvertent social information on signal evolution. Relatively little attention has been given to this issue until now, since most research on the evolution of signals has focused on signaling as an independent behavior, decoupled from its social and behavioral context (Searcy and Nowicki, 2005).

3.1 Methods

The hardware and controllers of the robots, as well as the experimental setup are described in section 2.1. The selection method used in this chapter is identical to the experimental treatment where robots were unrelated and selected at the individual level, as described in section 2.1.4. The behavioral analysis of signaling and response behaviors of robots, in addition to the statistical tools used are presented in Appendix A. The measure of information is detailed in Appendix B.1.

3.2 Results

3.2.1 Inadvertent information

An inherent property of the foraging system used in this thesis is that blue light, even if emitted randomly, could provide inadvertent social information on food location. This is because in this physical setup, information is provided not only through patterns of light emission, but also through the robots' behavior. Thus, once robots evolve the ability to find food and stay nearby, their increasing density near the food source should translate into higher blue density near the food and a source of information for other robots in the arena. This was confirmed in an experiment where robots were constrained to produce light randomly (light was emitted with a probability of 0.5 for each unit of time). As robots became more efficient at finding and remaining near the food (Fig. 3.1), the concentration of blue light near food also increased, such that in all generations after generation 2, the intensity of blue light was significantly higher in the vicinity of the food than in the rest of the arena (Mann-Whitney test, $df = 39$, $P < 0.001$ for all generations). To quantify the amount of inadvertent information produced by the emission of blue light, we devised an index of information I (Shannon, 1948; Haldane and Spurway, 1954; Wilson, 1962, 1975; Harms, 2006), which varies between 0 when blue light is equally distributed in all directions relative to the direction of the food and 1 when light is always perceived in a predictable direction relative to the food (see Appendix B for a detailed description). The level of information rapidly increased over the generations (Fig. 3.2A) and robots became significantly attracted to blue light after generation 9 (average value between generation 9 and 500: 0.2 ± 0.03, two-sided sign test, $df = 19$, all $P < 0.001$, Fig. 3.2B).

This experiment revealed that, when emitted randomly, blue light was an inadvertent cue providing information on food location. Although light production was cost-free, sharing such information should be costly because it results in higher robot density and increased competition and interference nearby the food (i.e., spatial constraints around the food source allowed a maximum of 8 robots out of 10 to feed simultaneously and resulted in robots sometimes pushing each other away from the food, see section 2.1.1). Since selection occurred at the individual level and because there is no kin structure among the 100 groups of robots, selection should favor individuals concealing information on food loca-

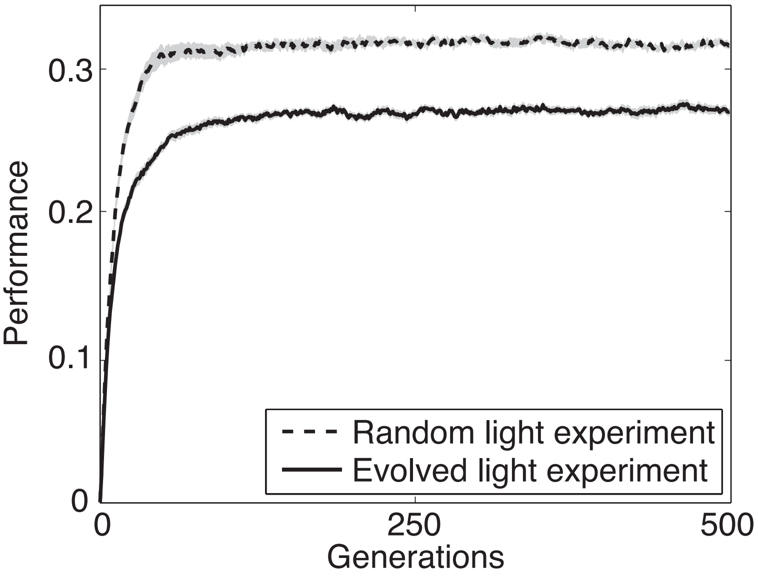

Figure 3.1: Mean and standard error (gray bands) of performance of robots in the 20 independent replicates over 500 generations in the experiments where robots were constrained to emit light randomly and the experiment where light emission could evolve.

tion (Hamilton, 1964; Wiley, 1983; Floreano et al., 2007). To test whether conflicts of interest between robots would affect blue light emission, we conducted a similar experiment with the difference that we allowed the rate of emission of blue light to evolve as part of the robots' behavior.

3.2.2 Evolving light emission

As in the previous experiment, the robots initially produced blue light randomly (gene values were random, such that the probability of light emission in any area of the arena was not different from 0.5 in the first 3 generations, two-sided sign test, $df = 19$, all $P > 0.1$, Fig. 3.3A). This random emission of blue light, together with the increase over generations in the density of robots near food, resulted in a rapid increase of information, the maximum ($I = 0.12$) being reached at generation 16 (Fig. 3.2A). However, because the increased information resulted in robots crowding around the food, robots were selected to decrease the rate of blue light emission (Fig. 3.3A). This decrease was significantly greater near the food than elsewhere, such that after 52 generations, robots became much less likely to produce light near food than near poison ($P < 0.01$ in all generations

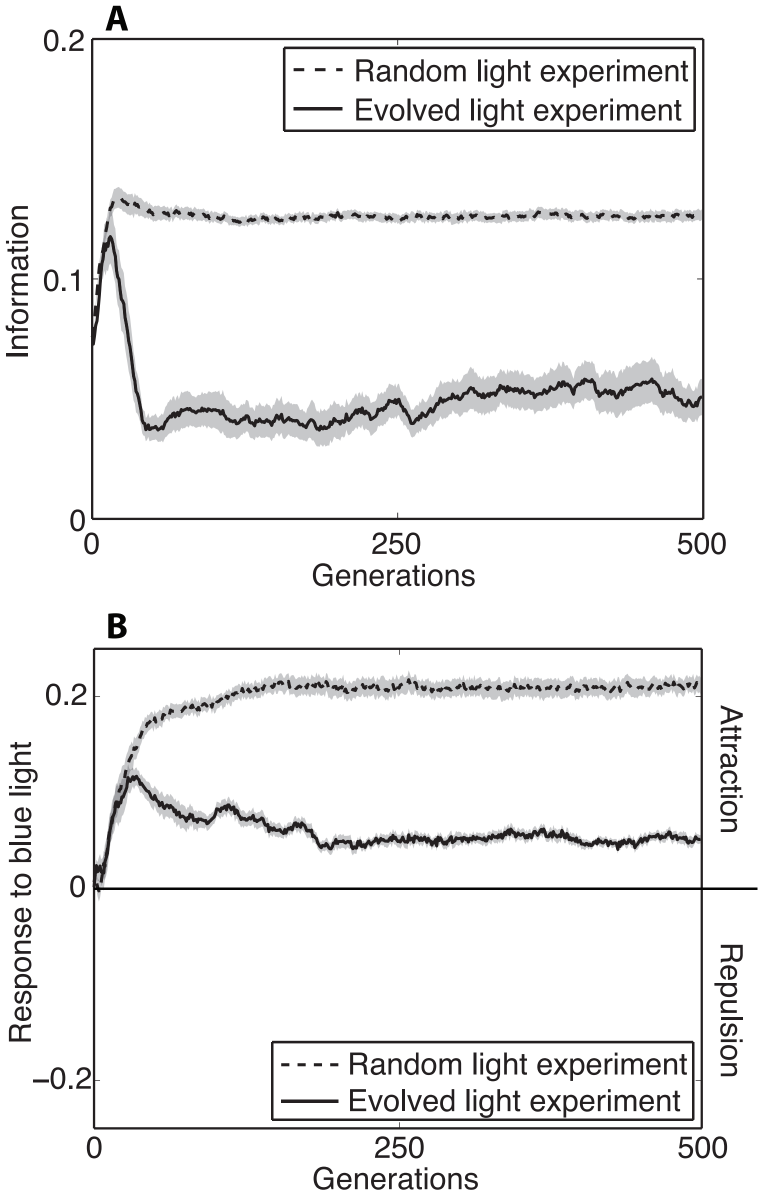

Figure 3.2: Information and response to blue light. Change over generations in (A) information content provided by blue light (i.e., the strength of the association between the direction in which robots perceived most light and the direction of the food, see Appendix B.1) and (B) the response to blue light, where positive values indicate attraction and negative values repulsion to blue light (see Appendix A). Both panels show the mean and standard error (gray bands) of the 20 independent replicates for both the random and evolved light emission experiments.

after generation 52, $df = 39$) or elsewhere ($P < 0.05$ in all but one of the generations 53-500, $df = 39$; Fig. 3.3A). Altogether, these changes in light emission strategy led to a drastic decrease in the amount of information provided by blue light intensity between generations 16 and 46 (Spearman's rank correlation test, $r_S = -0.997$, $P < 0.001$, Fig. 3.2A) such that, in all generations after generation 22, the level of information was significantly lower than in the experiment where blue light emission was random (all $P < 0.05$, $df = 39$, Fig. 3.2A).

Although selection was acting towards suppressing information on food location, the information content did not decrease to zero over the 500 generations of selection (Fig. 3.2A), resulting in robots remaining significantly attracted to blue light until generation 500 (average attraction between generations 36-500: 0.06 ± 0.01, two-sided sign test, $df = 19$, all $P < 0.001$, Fig. 3.2B). This somewhat surprising result can be explained by the fact that the strength of selection on light emission strategies depends on the level of information content and the robots' response to blue light. Thus, when the information content provided by blue light intensity is high, robots should be highly attracted to blue light and there should be a relatively important performance drop for robots emitting light near the food (i.e., strong selection pressure to reduce light emission by the food). By contrast, low information content should translate into a lower response of robots to blue light and a smaller performance reduction for robots that emit light near food (i.e., low selection pressure on reducing light by the food).

Support for suppression of information being impeded by the reduced strength of selection comes from the analysis of the response strategies of individual robots and their influence on the light emission strategies and performance over the 500 generations of selection. In both the experiments where light emission could evolve or was random, the level of attraction of robots to blue light rapidly increased during the first 36 generations of selection (Fig. 3.2B). However, while the level of attraction continued to increase when light production was random, it significantly decreased between generations 36 and 200 ($r_S = -0.8$, $P < 0.001$, Fig. 3.2B) in the experiment where the emission of light could evolve. This decrease in attraction to blue light, which resulted from the decrease in information content that occurred after generation 16 in this experiment (Fig. 3.2A), has important implications because the strength of selection on reducing the emission of light by food was positively correlated with the response of robots to

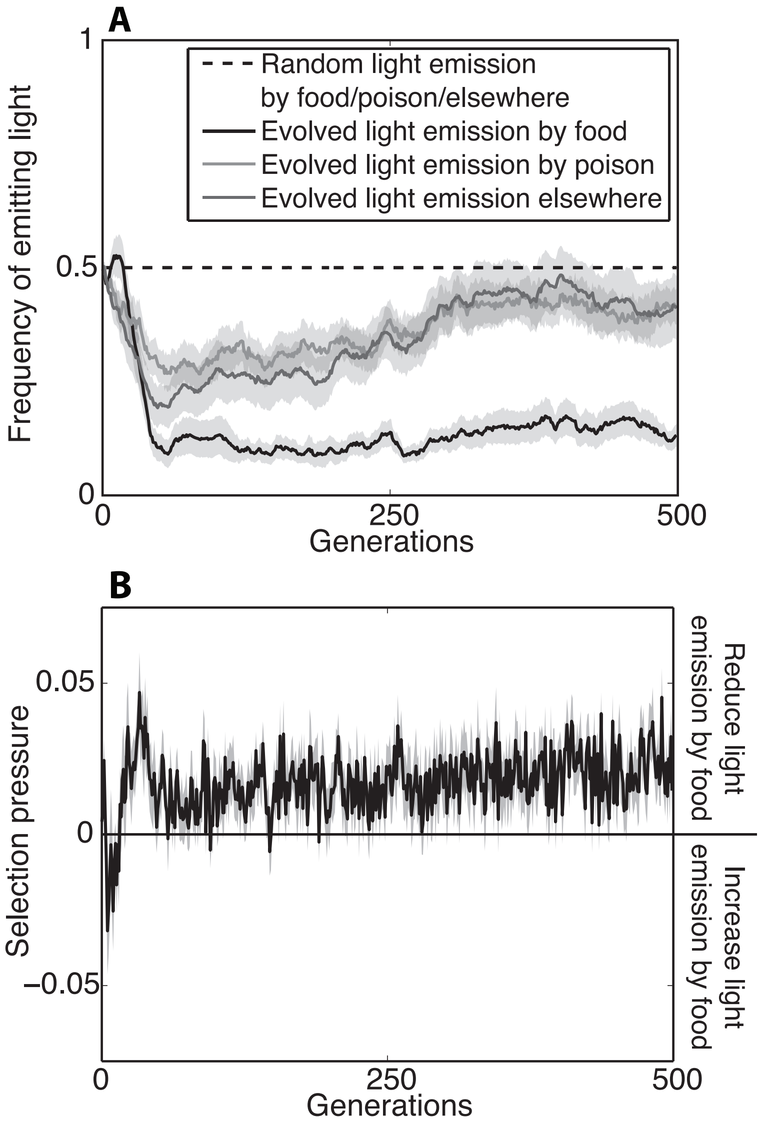

Figure 3.3: Light emission strategies and selection pressure. Change over generations in (A) the frequency of light emission in different areas of the arena (see Appendix A) for both the experiments where light was emitted randomly and evolved and (B) selection pressure to reduce emission of light by food (see Appendix A). Positive or negative selection pressure indicates that robots were selected to reduce or increase light emission near the food, respectively. Both panels show mean and standard error (gray bands) of the 20 independent replicates.

blue light ($r_S = 0.13 \pm 0.13$; correlation significant in 11 of the 20 selection lines, Fig. 3.2B, 3.3B). Thus, the reduced response to blue light after generation 36 led to weaker selection on reducing light emission by food (Fig. 3.3B) and a stable level of information being reached by mutation-selection after about 50-100 generations (Fig. 3.2A).

To test whether the stable level of production and attraction to light was affected by the mutation rate, we conducted an additional experiment with the only difference that between generations 250 and 500 we used a 100-fold lower mutation and crossing-over rate (mutation rate of 0.0001 per locus instead of 0.01; crossover rate of 0.002 instead of 0.2)[2]. The reduction resulted in a 33% decrease in the frequency of signaling near food (over the last 10 generations, mean ± SD: 0.08 ± 0.24 instead of 0.12 ± 0.1; Mann-Whitney test, $df = 39$, $P < 0.001$). Interestingly, however, the reduced emission of blue light near food did not translate into a decrease in the level of information ($df = 39$, $P = 0.48$) nor a decrease in attraction to blue light ($df = 39$, $P = 0.36$), because the decreased mutation rate also led to an increase in the average foraging efficiency of robots and thus a higher concentration of robots near food (0.91 ± 0.01 compared to 0.83 ± 0.01 with the regular rates, $df = 39$, $P < 0.001$). In other words, the lower rate of signaling combined with a higher concentration of robots at the food nevertheless generated a sufficient amount of information for the robots to remain equally attracted to blue light.

3.2.3 Within-population variation

An important consequence of the reduced selection pressure on light emission is that, at equilibrium, there was considerable individual variation in both the production and response to light (Fig. 3.4). While the majority (61.5%) of the robots never emitted light near food in the last 10 generations, there was high variation in light emission strategies with 11.2% of the robots emitting light over 50% of the time when near food (Fig. 3.4A). Similarly, the level of attraction of robots greatly varied with most robots exhibiting a low attraction to blue light,

---

[2] Since the lower mutation and crossover probabilities in the new treatment did not sustain sufficient diversity in the population to evolve good foraging behaviors (e.g., driving towards red light) in the initial generations, we used different mutation and crossing rates only after generation 250.

but 32.6% showing a negative attraction (i.e., repulsion) to blue light and 36.1% an attraction more than twice higher than the average (Fig. 3.4B). Furthermore, the within-population variance in attraction of robots to blue light was significantly higher in the last 10 generations where blue light production could evolve (0.15 ± 0.02) than when it was random (0.13 ± 0.01, $df = 39$, all $P < 0.01$).

These findings are interesting with regard to the discrepancy between theoretical predictions and previous empirical studies on the association between the strength of selection and phenotypic diversity. While theory suggests that lower selection and genetic drift should lead to higher within-population variation (Fisher, 1958; Falconer, 1981; Lynch and Hill, 1986), empirical studies failed to support this prediction (Houle, 1992). However, a problem with these empirical studies is that the strength of selection is assessed indirectly, for example by assuming greater selection on life-history traits than morphological traits (Houle, 1992). In our experiments, robots exhibited greater phenotypic variability in their response to blue light when light emission could evolve (i.e., when the level of information and strength of selection were low) than when light emission was fixed (i.e., when robots emitted light randomly, such that the level of information and strength of selection were higher). This is in line with theoretical predictions and supports the view that more controlled experiments are needed in studies with real organisms.

3.2.4 Performance reduction

Although these analyses provide an explanation for the attraction of robots to blue light despite the reduction in information, the paradox noted in the previous chapter, that allowing robots to produce blue light decreases performance, still remains. To explore this issue, we investigated the relationship between the robots' response strategies and individual performance. The maximum mean performance (0.29 ± 0.16) was achieved at a level of attraction of 0.09 but the large variation in response strategies meant that not all robots managed to maintain such high performance. While the performance of robots with response values around 0.09 (all values between -0.29 and 0.37, comprising 97.7% of all individuals) did not differ significantly from the maximum performance (Kruskal-Wallis test, $P > 0.05$, Fig. 3.5), the mean performance of the remaining 2.3% individuals was very low (0.01 ± 0.07), hence accounting for the lower overall performance of robots when blue light production was possible than in

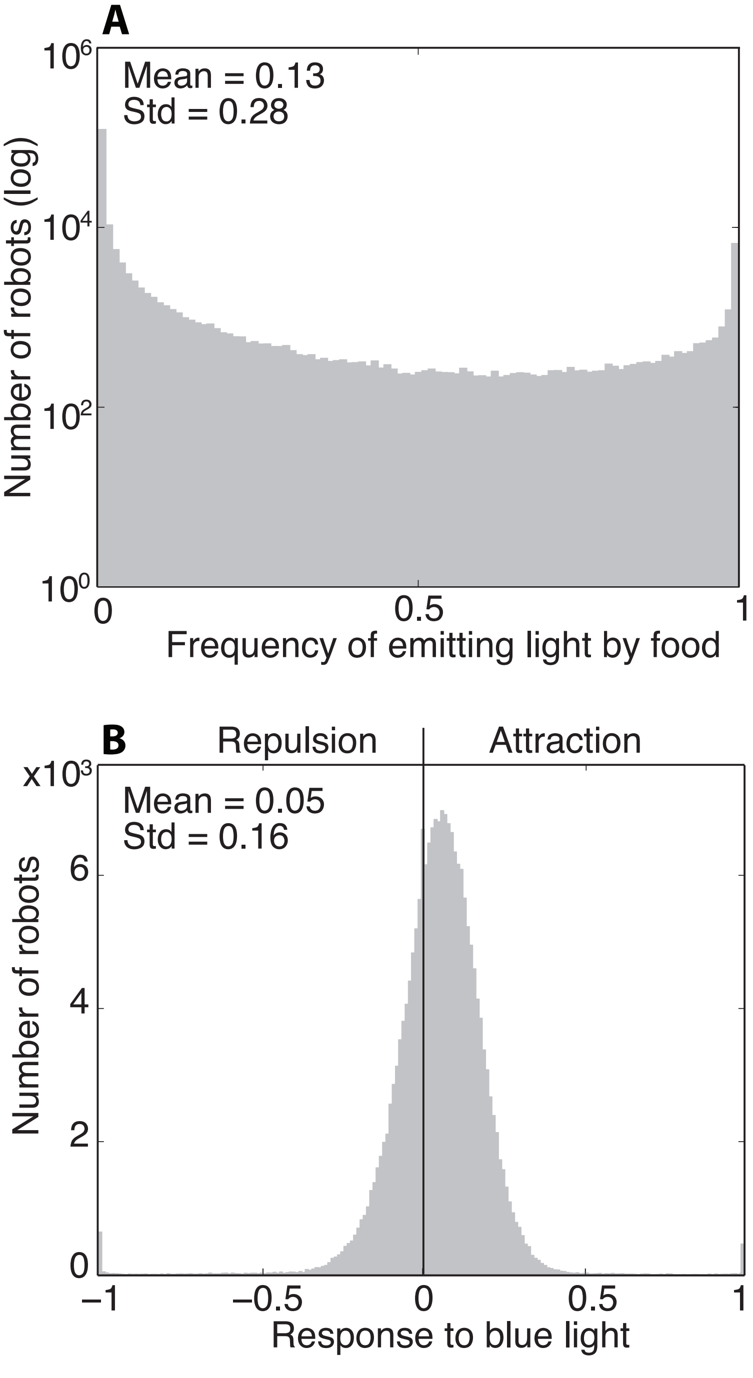

Figure 3.4: Variation in communicative strategies. Histograms of the frequency of blue light emission near the food in panel (A), and the response to blue light in panel (B) of the 1000 robots in each of the last 10 generations in all 20 independent replicates (200,000 robots in total) of the experiment where light emission was evolved. In panel (B), positive response values reflect attraction to blue light and negative responses repulsion. The vertical line marks the level at which robots exhibit no positive or negative response to blue light.

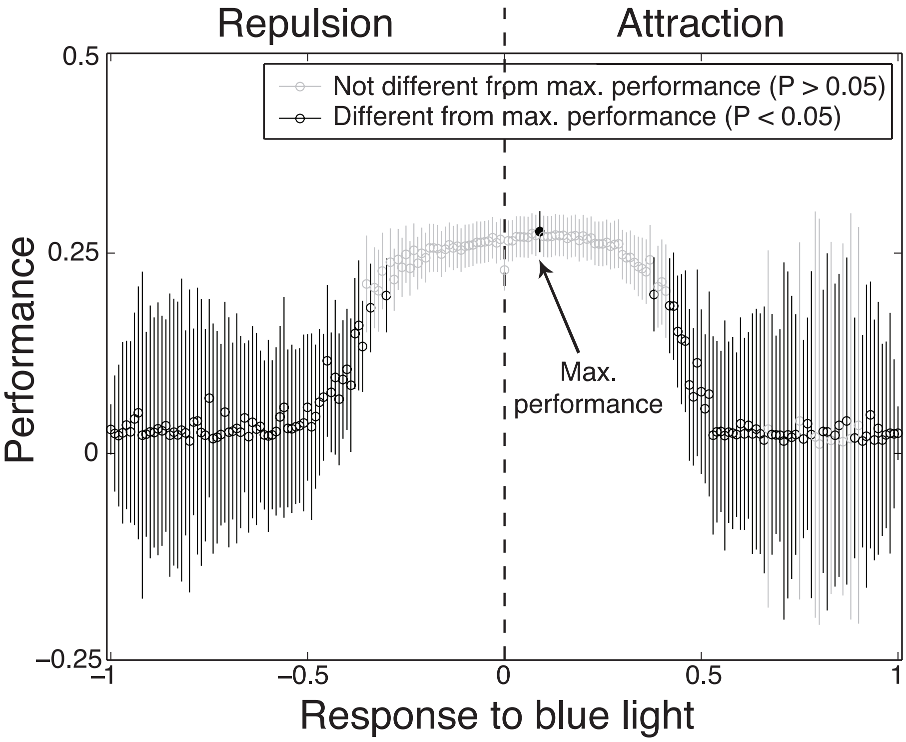

Figure 3.5: Performance of the 1000 robots in each generation from 491 to 500 in all 20 independent replicates (200,000 robots in total). Positive response values reflect attraction to blue light and negative responses repulsion. The vertical dashed line marks the level at which robots exhibit no positive or negative response to blue light. Circles show the mean performance of all individuals with a specific response. Two means are significantly different (Kruskal-Wallis test, $P < 0.05$) if their intervals (vertical lines) are disjoint, and are not significantly different ($P > 0.05$) if their intervals overlap. The maximum mean performance is at 0.09 (solid black point).

the control where it was prevented.

3.3 Discussion

3.3.1 The evolution of cues into signals

The complex dynamics between suppression of social information and selection pressure described in this study are likely to be general features of natural systems in which cues evolve into signals. When there are conflicts of interest between interacting individuals, those producing cues providing useful information to others should be selected to hide the information by interfering on the channel that carries the cue, thus resulting in signals carrying little information (Wiley, 1983; Hasson, 1994). However, complete information suppression might never be achieved, since a reduction in information will simultaneously reduce selection pressure on information reduction, resulting in a stable equilibrium, where information is not completely suppressed. Moreover, there might be constraints making some cues more difficult to suppress. For example, when a small nocturnal rodent moves on the forest floor, the sound emitted by brushing against the substrate provides a cue for an avian predator to detect it. Clearly, it will be impossible for the rodent to move without making any noise. But, by altering the way the rodent moves or by inducing it to avoid some noisy substrates natural selection can act so as to decrease the information provided to the predator. In the same way, it is likely that all inadvertent cues are to some extent subject to natural selection with a continuum in the degree by which the effect of natural selection is hindered by morphological, physiological or behavioral constraints.

In addition, our results suggest that considerable phenotypic variation within populations will be present at equilibrium when the degree of selection on communication strategies is weak. The magnitude of this variation should be dependent on the rate of mutation and complexity of the genetic architecture (Lande, 1976). Thus, considerable variation is to be expected in communication systems that rely on complex traits driven by many genes and a complex brain circuitry. Several empirical studies indeed report higher than expected intra-population variation in communication strategies when interests between parties conflict (Poulton, 1890; Whiteley et al., 1997; Dewitt et al., 1999; Bond and Kamil, 2002; Bond, 2007). For example, the great amount of polymorphism in the cryptic coloration in moths has fascinated naturalists for centuries, but explanations have remained elusive (Bond, 2007).

Finally, our study also underscores the importance of considering the coupling between communication and behavioral effects when studying communication systems (Pfeifer et al., 2007), which is rarely taken into account in current analytical and game-theoretical models (Johnstone, 1998a; Rowell et al., 2006). Evolutionary robotic systems implicitly encompass many behavioral components, such as the inadvertent production of information through foraging behavior, thus allowing for an unbiased investigation of the factors driving signal evolution (see chapter 5 for a detailed discussion). Moreover, the use of accurate robotic simulations is advantageous because it simultaneously preserves the mechanistic properties of hardware robotic models and allows one to conduct experimental evolution in large populations over many generations. The great degree of realism provided by evolutionary robotic systems thus provides a powerful tool for studies that cannot readily be performed with real organisms.

4

The Effect of Relatedness on Signal Reliability

Chapters 2 and 3 provided a detailed analysis of both reliable and unreliable signaling systems. However, determining the reliability of these signaling systems has been based on observations of performance measures and signaling strategies rather than directly measuring the reliability of signals. In this chapter[1], we introduce a quantitative measure of signal reliability and use it to explore the influence of varying within-group relatedness on the reliability of the evolving systems of communication.

---

[1] This chapter is based on Mitri, S., Floreano, D., and Keller, L. (In preparation).

Empirical observations in very different taxa suggest that honest and reliable communication frequently occurs among highly related individuals (Hamilton, 1964; Dunford, 1977; Cheney and Seyfarth, 1985; Diggle et al., 2007b; Boncoraglio et al., 2009). This is in line with current social evolutionary theory, which predicts that relatedness is a key component selecting for cooperative behavior (Hamilton, 1964; Lehmann and Keller, 2006; West et al., 2007b). However, because of the lack of quantitative measures of signal reliability in natural communication systems, it has proven difficult to test whether high relatedness is an important factor promoting reliable signaling.

To quantitatively explore how relatedness influences the evolution of signal reliability, in this chapter, we vary the level of relatedness within groups of robots over five different values (0, 0.25, 0.54, 0.75 and 1), and devise a measure of signal reliability to investigate how this index is influenced by within-group relatedness.

4.1 Methods

The hardware and controllers of the robots, as well as the experimental setup are similar to those used in the two previous chapters (see section 2.1). The only differences are that in this chapter, groups are made up of eight robots each, making a population size of 800, and the food and poison sources are smaller (radius of 8cm, instead of 10cm), i.e., robots still have to compete for space around the food. Selection is carried out at the individual level in this chapter (see section 2.1.4 for a description). The behavioral analysis of the signaling behaviors of robots, in addition to the statistical tools used are presented in Appendix A. The index of signal reliability is described in Appendix B.2.

4.1.1 Relatedness

To form a group of relatedness $r = 1$, as in chapter 2, one individual was randomly chosen from the pool of the 160 selected individuals (best 20% of population), recombined, mutated (according to probabilities given in section 2.1.4) and cloned 7 times, to make groups of 8 robots with identical genomes. At the other extreme, groups of unrelated robots ($r = 0$) were composed by repeatedly choosing 8 different genomes from the pool and assigning them to the same group.

For intermediate levels of relatedness, $r = 0.25$, $r = 0.54$ and $r = 0.75$, different proportions of clones were used to form groups of 8 robots (3:3:2, 1:1:6 and 1:7, for the three levels of relatedness, respectively, see Appendix B.3). Although this group composition differs from relatedness in most natural populations, social evolution should depend only on the average relatedness within a group, since the robots have no way of directing interactions towards specific individuals. 20 experimental replicates were conducted for each of the five levels of relatedness.

4.2 Results

4.2.1 Relatedness, performance and signal reliability

Over the 500 generations of selection, the performance of the robots increased at all levels of relatedness (Fig. 4.1). At the end of the selection experiment, the performance of robots was positively associated with within-group relatedness (Spearman's rank correlation test, $r_S = 0.76$, $P < 0.001$, Fig. 4.2A). The performance of the robots significantly increased between relatedness values 0 and 0.75 (Mann-Whitney tests, all $df = 39$, $P < 0.01$), but the performance did not differ significantly between relatedness values 0.75 and 1 ($df = 39$, $P = 0.88$).

To test whether these differences in performance were due to differences in signaling strategies, we conducted an additional experiment where robots were blind to blue light. In this experiment, blue light emission could evolve but it could no longer affect the robots' performance, as it could not be perceived. Consistent with the view that relatedness affected performance by altering the robots' communication strategy, there was no significant correlation between relatedness and performance in this experiment ($r_S = 0.17$, $P = 0.07$, see Fig. 4.2A).

To investigate how variations in relatedness affected the reliability of signaling, we devised an index of signal reliability, which consists of analyzing the inputs of each robot's visual system to establish whether robots perceived more blue light in a consistent direction with respect to the direction of the food (see Appendix B.2). This index could vary between 0 when blue light was equally distributed in all directions relative to the direction of the food (i.e., the signal is completely unreliable) to 1 when blue light was always perceived in a pre-

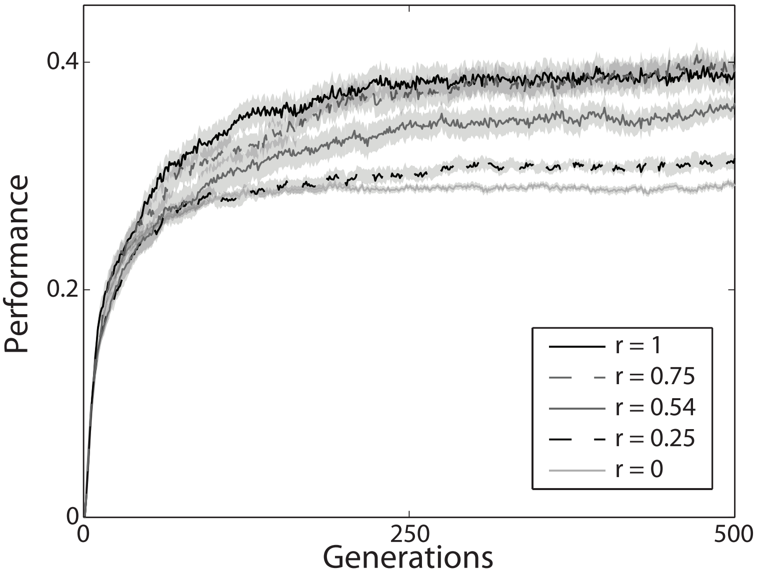

Figure 4.1: Mean and standard error (gray bands) of performance of robots in the 20 independent replicates over 500 generations at different levels of relatedness.

dictable direction relative to the food (i.e., the signal is completely reliable)[2]. The reliability index was significantly correlated with relatedness (Spearman's correlation test, $r_S = 0.59$, $df = 99$, $P < 0.001$, Fig. 4.2B) as well as with performance ($r_S = 0.76$, $df = 99$, $P < 0.001$).

To investigate how the signaling strategies of the robots changed with relatedness, we compared the frequency of blue light production in different areas of the arena (near the food, near the poison, and elsewhere in the arena) across relatedness treatments in the experiment where robots could perceive blue light. As relatedness increased, there was an increase in the frequency of signaling near the food ($r_S = 0.49$, $P < 0.001$) and a reduction in signaling frequency near poison ($r_S = -0.62$, $P < 0.001$) and elsewhere in the arena ($r_S = -0.29$, $P < 0.01$, Fig. 4.3).

---

[2] The index of signal reliability is similar to the index of information I presented in chapter 3. The difference is that the index of reliability only considers instances in which robots perceive blue light, whereas the measure of information also includes time-steps in which no blue light is perceived. See Appendix B for exact definitions.

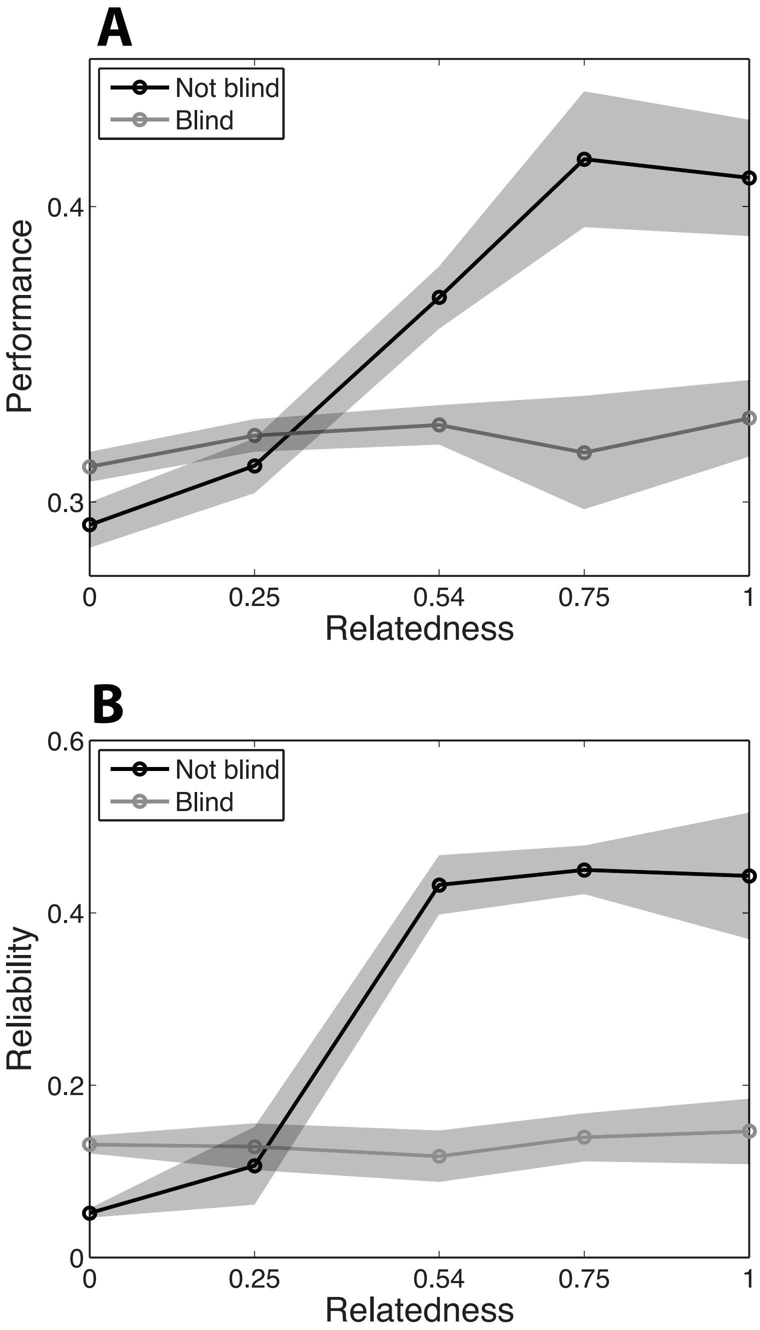

Figure 4.2: Performance (A) and signal reliability (B) at different levels of relatedness for robots that could perceive blue light compared to robots that were blind to blue light. Each point represents the median of 20 independent replicates, where the value of each replicate is the average reliability over all individuals in the last 20 generations. The gray bands represent the uncertainty about the median. Bands that do not overlap indicate that the medians differ at the 5% significance level.

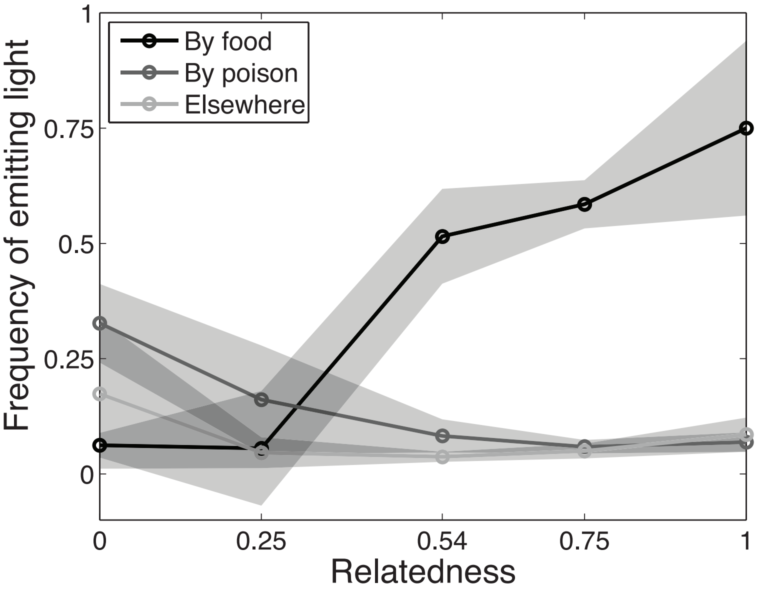

Figure 4.3: Frequencies of blue light emission in different areas of the arena at different levels of relatedness. Each point represents the median of 20 independent replicates, where the value of each replicate is the average reliability over all individuals in the last 20 generations. The gray bands represent the uncertainty about the median. Bands that do not overlap indicate that the medians differ at the 5% significance level.

4.2.2 Signal frequency vs. signal reliability

Because measuring signal reliability is difficult under natural conditions, researchers have typically used signal production frequencies as a proxy for signal reliability in a variety of contexts (e.g., signals of need (Redondo and Castro, 1992; Cotton et al., 1996), sexual signals (Zahavi, 1975) and cooperative signals (Wauters and Richard-Yris, 2003)). To explore whether signal production frequencies are good indicators of signal reliability in our system, we analyzed the frequency of blue light emission in the different treatments. There was no significant difference in light emission frequencies across the five levels of relatedness (Kruskal-Wallis test, $df = 99$, $P = 0.14$). Furthermore, within each relatedness treatment, we found no consistent association between signal frequency and performance. Across the 20 replicates the correlation between the frequency of signaling and performance was negative for relatedness 0 and 0.54 (Pearson's correlation coefficient $\rho = -0.44$ at $r = 0$ and $\rho = -0.53$ at $r = 0.54$, $P < 0.05$), positive for relatedness 1 ($\rho = 0.89$, $P < 0.001$) and not significant for relatedness 0.25 and 0.75 ($P > 0.05$, Fig. 4.4A). Because a good measure of signal reliability is expected to correlate positively with performance, these data show that signal frequency is not a good estimate of signal reliability in our experiments.

By contrast, a similar analysis between the index of reliability and performance revealed a significant positive correlation at all levels of relatedness greater than 0 ($r = 0.25$ to $r = 1$, all 4 Pearson's correlation coefficients $\rho > 0.63$, $P < 0.01$, Fig. 4.4B). At $r = 0$, light emission was unreliable and resulted in low performance in all experimental replicates and thus, no significant correlation between the two measures was observed ($\rho = -0.01$, $P = 0.97$).

4.3 Discussion

In this chapter, we have shown that increasing relatedness within groups results in an increase in the reliability of the robots' evolved signaling strategies. These findings can be explained by **kin selection** theory (Hamilton, 1964). Because unrelated robots were selected to behave selfishly rather than altruistically, they were selected to avoid providing reliable information on the location of food, as this would result in crowding by the food, increased competition and thus decreased individual performance (chapter 3). Conversely, high relatedness be-

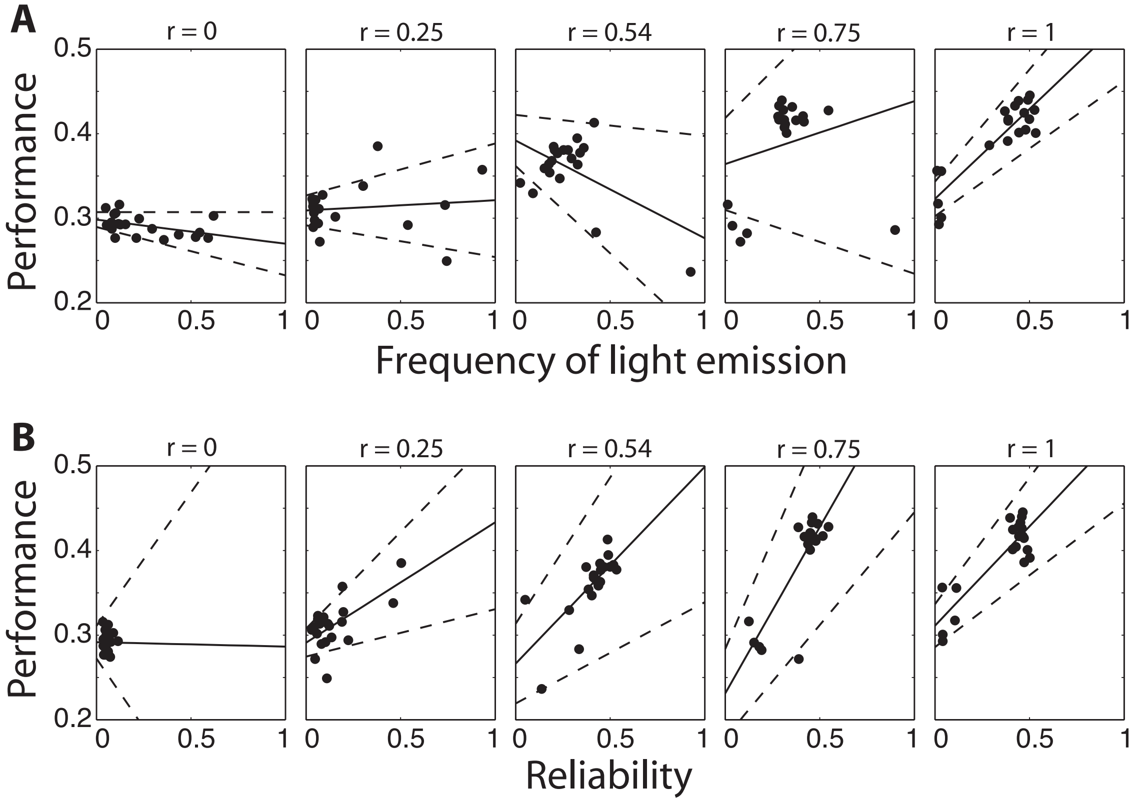

Figure 4.4: The frequency of blue light emission vs. performance (A) and signal reliability vs. performance (B) of robots at different levels of relatedness. Each point represents the average over the last 20 generations for one experimental replicate. Solid lines indicate results of a linear regression, while dashed lines show 95% confidence intervals.

tween robots selected for the evolution of reliable signals that allowed related individuals to easily locate the food (chapter 2).

Higher relatedness has also been shown to lead to more reliable signals in natural communication systems (Hamilton, 1964; Dunford, 1977; Cheney and Seyfarth, 1985; Diggle et al., 2007b; Boncoraglio et al., 2009). For example, chicks that beg for food from their parents produce less intense begging displays in groups of siblings than in mixed groups (Boncoraglio et al., 2009). Similarly, bacterial cells that are closely related produce more signal molecules that help the group to coordinate cooperative behaviors than strains of unrelated bacteria (Diggle et al., 2007b). In these two examples, it is notable that two opposing trends in signal frequency are taken to indicate reliable signaling. This apparent contradiction stems from the difference in the benefits of signaling in the two scenarios. In the case of the chicks, louder signaling will lead to larger rewards for the individual signaler and consequently less for its nest-mates, whereas in the case of the bacteria, stronger signaling tends to increase the benefits for the group rather than for the signaling individual alone. The individual microbe instead gains more by signaling less and reaping the benefits of the collective behavior of its conspecifics.

The context in which signals are produced and perceived, such as the way in which benefits gained through signaling are distributed among individuals, are thus expected to largely influence the frequency of signal production and its relation to signal reliability. Nevertheless, many studies use signal production frequencies to estimate signal reliability (see Searcy and Nowicki (2005) for a review). An alternative is to measure signals at the receiving end, as we have shown in this chapter, resulting in a measure of signal reliability that is independent of such contextual elements. This is, of course, an easier task to accomplish using robots, where the experimenter can access the robots' sensory inputs. However, technical advances should allow the measurement of physiological aspects of signal perception in living organisms (Cummings and Partridge, 2001; Cummings, 2007; Ryan, 2007; Stuart-Fox et al., 2007), making quantitative measurements of signal reliability possible in natural systems.

5

Using Robots to Model Social Behavior

The approach of using robots to study social behavior is relatively novel. It is thus important to carefully consider whether the use of robots really can lead to relevant contributions to understanding social behavior, which would not have been possible with other, more traditional approaches. In this chapter[1] we give an overview of the different approaches to studying social behavior and discuss the merits and drawbacks of the use of robots by reviewing related work, including the experiments described in this thesis, where robots are used to study social behavior.

---

[1] This chapter is based on Mitri, S., Wischmann, S., Floreano, D., and Keller, L. (In preparation).

Robots are machines that have sensors with which to perceive the world around them, actuators to manipulate it, and control programs that allow them to translate perception into action. Even with a very simple control program, the behavior of a robot can look uncannily life-like (Braitenberg, 1984). Furthermore, having many of these robots interact together in "robot societies" has revealed emergent behavior akin to many self-organizing collective behaviors observed in nature: cooperative transport (Kube and Bonabeau, 2000; Groß and Dorigo, 2004; Mondada et al., 2004; Waibel, 2007), object sorting (Holland and Melhuish, 1999; Melhuish et al., 2006) and even communication (Nolfi (2005); Ampatzis et al. (2008), chapters 2, 3 and 4 in this thesis). Although much of this work was initiated as a way of solving engineering problems by taking inspiration from natural collective systems, it has become clear that contributions can also be made in the opposite direction, leading to progress in the fields of sociobiology, ecology and evolutionary biology (Floreano and Mattiussi, 2008).

Social behavior is inherently difficult to model. This is firstly because of the complexity that arises from modeling many individuals simultaneously and secondly because of the importance of their local interactions with each other and with their environment (Judson, 1994). One way of understanding how groups of animals behave is to build robots that are in some ways similar either to a specific animal or to animals in general (Holland and McFarland, 2001; Webb, 2001, 2009), to program their individual behavior, and then to analyze the collective behavior resulting from allowing them to interact. If the collective or social behavior is similar to that of real animals, testable hypotheses can be proposed to empirically investigate whether the animals may be using similar rules. Alternatively, one could imagine allowing the robots to interact directly with real animals to explore how the animals will react to different stimuli from their "artificial conspecifics".

Although such ideas may sound appealing, it is necessary to carefully consider whether it really is worthwhile to use robots to understand collective behavior in living organisms. Our aim in this chapter is to discuss this issue by comparing robot models to other, more traditional modeling techniques, reviewing studies that have used robots to study social behavior and providing guidelines on how to exploit the potential power of this approach.

5.1 Modeling social behavior

A number of different approaches are used to study social behavior, which vary on a scale of **situatedness** (see Fig. 5.1). By situatedness we refer to the extent to which individuals are embedded in an environment that can be sensed and modified by those individuals (Varela et al., 1991; Clark, 1996). At one extreme, abstract mathematical models (also known as analytical, classical or minimal models (Grimm and Railsback, 2005)) aim to boil down a system to its minimal components to explore the effects of what are considered to be key parameters on its dynamics (e.g., the effects of the interaction rates between a population of predators and their prey on changes in the sizes of the two populations (Lotka, 1925; Volterra, 1926)). These models are at the low end of the situatedness scale (left of Fig. 5.1) because they do not model individuals nor their environment explicitly, but instead assume that they will not have a significant influence on the conclusions obtained using the model (Murray (2007)).

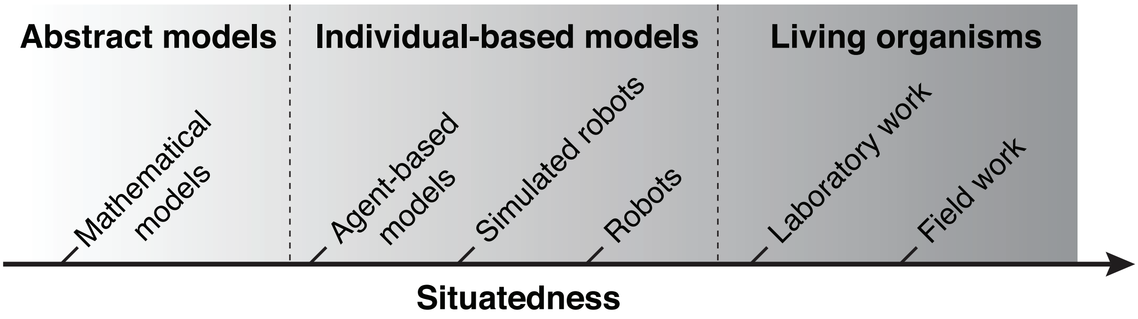

Figure 5.1: Different approaches to studying social behavior on a scale of situatedness, i.e., the extent to which individuals are embedded in an environment that they can sense and modify (Varela et al., 1991; Clark, 1996).

At the other end of the situatedness spectrum (right of Fig. 5.1), researchers conduct experiments using living organisms. Field work, conducted in the organisms' natural habitats can be seen as the highest level of situatedness, since the effect of all the different properties of the organism's natural environment are included in the observed social behavior. For example, studying frogs' mating calls in their natural environment will take the influence of environmental noise into account, which may have an impact on their vocalizations (Wollerman and Wiley, 2002). This is a challenging approach, however, because the experimenter has very little control over these environmental factors, and cannot easily disentangle all their effects from the parameters of interest. A partial solution to this problem is to conduct studies of social systems in the laboratory (Fig. 5.1), where the environment is controlled. By doing this, the experimenter must make some simplifying assumptions as to which aspects of the organism's natural habitat are important for the particular question at hand. More generally, when choosing the level of situatedness at which to study a given problem, a trade-off exists between an increasing ease in manipulating and understanding the system, and a simultaneous increase in the risk of removing elements of the real system (living organism in its natural habitat) that may be crucial to its resulting behavior.

To balance these trade-offs, researchers have resorted to a type of modeling that is intermediate on the scale of situatedness: individual-based models (IBMs). Compared to living organisms, IBMs are easier to manipulate and understand. At the same time, they take differences between individuals (e.g., age or size) and their interactions into account, which have proven to have important consequences on collective and social behavior, and are generally not included in abstract mathematical models (Judson, 1994; Grimm and Railsback, 2005; DeAngelis and Mooij, 2005). For these reasons, IBMs have largely been accepted as part of the toolbox for modeling social systems from small groups to populations and ecosystems (DeAngelis and Mooij, 2005).

Robotic models can be seen as more situated IBMs (Fig. 5.1). Robots allow abstract individuals to leave their simplified, artificial world and inhabit the physical world, with all its complexity. It is important, however, to explore whether anything is to be gained from this increase in situatedness.

5.2 The pleasures and perils of social robots

The main advantage of using robotic organisms over agent-based models is that fewer assumptions need to be made regarding the environmental properties of the model. This is because the laws of physics are included in robotic models "for free". A direct consequence of this, is that it is more likely that experiments using robots will lead to surprising outcomes, because some unexpected property of the physical world – that would intuitively not have been included in the equivalent agent-based model – has an important influence on the results.

This point is nicely illustrated by the experiments presented in chapter 3, where we have shown that the spatial properties of the environment and the robots' visual system were important in the evolution of the signaling system. This is because the visual effects of robots clustering around the food provided cues to other robots, which had an important influence on the resulting dynamics of the communication system. Another example is provided by two studies on cockroach aggregation behavior using robots (Garnier et al., 2008) and software agents (Jeanson et al., 2005), respectively. Garnier and colleagues found that their robotic model resulted in more accurate predictions of the cockroaches' behavior than the agent-based model because the computer simulation did not allow for robots to hide behind objects or each other. This effect proved to be crucial to the resulting collective behavior (Garnier et al., 2008). Other studies have shown that friction caused by collisions (Krieger et al., 2000; May et al., 2006) also plays an important role in shaping social behavior. Although in some cases it may be possible to simulate these physical properties in a simpler model, it is difficult to predict *a priori* which factors of the physical model will influence the resulting social behavior. The use of robots thus allows the experimenter to discover these factors.

A second advantage of using physical robots, as opposed to agent-based models, is that they can interact with living organisms. A growing number of researchers are using such **mixed models**, where robots are used to infiltrate animal societies by playing back behavioral sequences. By analyzing the animals' responses to the robots, the experimenter can gain insights into the animals' behavioral codes. Although he did not use robots, Tinbergen (1948) designed one of the earliest experiments of this type, in which static models of sticklebacks with different visual properties were used to explore how real sticklebacks reacted to visual cues from conspecifics. More recently, as technology has become more advanced, affordable and accessible, dummies (such as Tinbergen's sticklebacks) have been replaced by robots. Robots provide an advantage in such studies because they enable the researcher to disentangle the effects of different mechanistic components on the resulting social behavior. Thus, one of the main uses of robots in mixed models has been to explore the role of different components of a signal on the resulting communication system (Michelsen et al., 1992; Patricelli et al., 2002; Göth and Evans, 2004; Martins et al., 2005; Fernández-Juricic et al., 2006; Halloy et al., 2007; Rundus et al., 2007; Taylor et al., 2008;

Partan et al., 2009).

Because robots need to work in the real world, however, their use often poses many technical challenges. Consider, for example, the problem of building a robot that can visually distinguish between its peers. While there are many straight-forward ways to deal with this problem using software agents, this is a considerable challenge for the state-of-the-art in robotics. Although some of these problems may provide insights into the biological species being modeled, this is often not the case (Webb, 2009). In addition, surprising outcomes may be a result of artifacts that are specific to the robots and do not have parallels in the natural world. Another reason why one might decide against using robots is because they are expensive to build. This is particularly problematic when studying collective behavior because construction costs limit the number of robots a modeler can use. One is therefore limited to study systems of small group sizes, or where the size of the group is expected to be irrelevant to its dynamics. A related problem is that working with robots can be extremely slow (Holland and McFarland, 2001). Because robots live in real-time, experiments cannot be accelerated, as they might when using agent-based models, for example. This problem particularly affects studies using artificial evolution where hundreds of robot controllers need to be tested over many generations (Nolfi and Floreano, 1998; Marocco et al., 2003; Floreano et al., 2007; Waibel, 2007; Ampatzis et al., 2008; Mitri et al., 2009).

Many of the disadvantages of using robots can be avoided by building physics-based software simulations of the robots, which has become possible because of the recent rise in computational power and simulation technology (Waldner, 2008). We place these software simulations of robots at an intermediate level on the situatedness scale between agent-based models and real robots (Fig. 5.1). Although at first sight, such simulations appear to be equivalent to using agent-based models, the crucial difference is that they are based on already-existing physical robots, and simulate them as accurately as possible (see e.g., Fig. 5.2A). This can be ensured by systematically comparing results obtained using a small number of real robots and the simulation until the latter proves to be a reliable substitute for the former, as we have done in this thesis (see section 2.1.5). However, simulations similarly suffer from a scaling problem, since the necessary computations grow exponentially with the number of robots (Martinoli and Easton, 2002; Waldner, 2008).

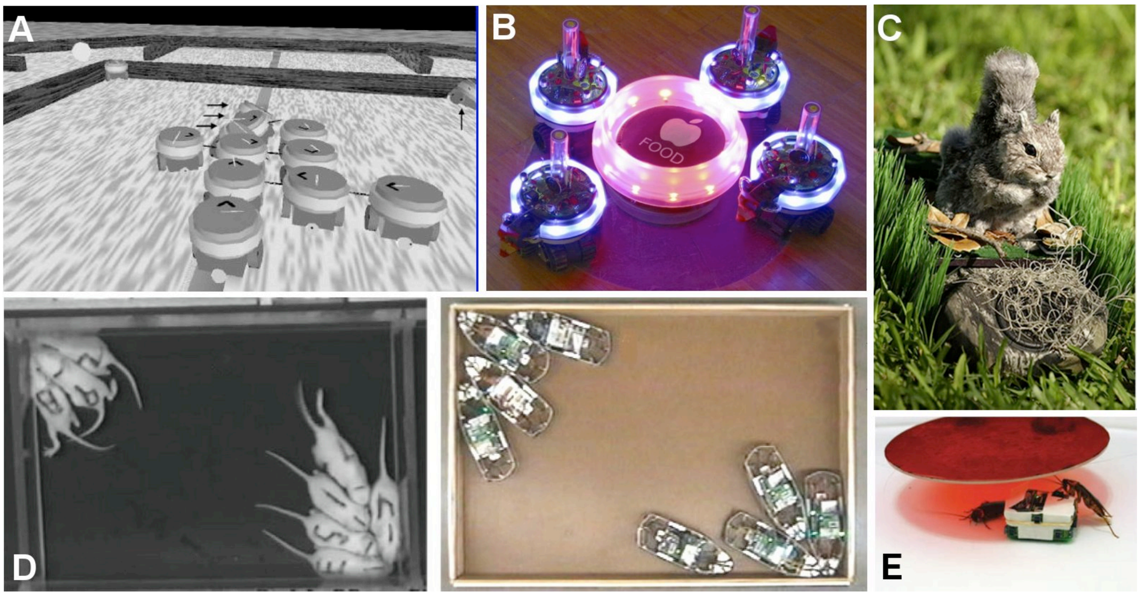

Figure 5.2: Robots used to study social behavior. A. Simulated robots to study coordination of collective behavior (Baldassarre et al., 2006). B. Foraging robots used in this thesis. C. Robotic squirrel built to disentangle the effects of different communication channels in squirrels (Partan et al., 2009). D. Modeling rat pup aggregation behavior using robots (May et al., 2006). E. Robots used to explore decision-making in cockroaches (Halloy et al., 2007).

5.3 Three ways to study social behavior using robots

To understand how researchers make use of robotic models, we have reviewed articles that use robots or simulated robots to study questions regarding social behavior in nature. We suggest that these studies can be classified into three categories depending on the goals of the study (Table 5.1).

Firstly, robotic models can be used to test the validity of established theories that have been difficult to test empirically. This is because robotic models (or even simulations of robots) are more situated, and thus implicitly include more environmental properties than mathematical or agent-based models. An example of a theory that has proven difficult to test quantitatively in living organisms is Hamilton's rule, which predicts when individuals in a group will evolve to behave cooperatively, as a function of the average relatedness in the group, and the costs and benefits of the cooperative act (Hamilton, 1964). Although tests of the theory have been conducted in living organisms by manipulating relatedness in colonies of bacteria (Griffin et al., 2004), fig wasps (West et al., 2001) and ants (Sundström et al., 1996), the costs and benefits of cooperative acts have been difficult to precisely quantify and vary in these systems. Alternatively, Waibel (2007) conducted a test of Hamilton's rule using groups of robots that evolved through artificial selection, where relatedness, costs and benefits could be varied. These experiments have confirmed that the quantitative predictions of the theory hold even with the noise and complexity of the real world. In another example, Michelsen et al. (1992) used a robotic bee to test a long-standing theory regarding the honeybee dance language. With their robot, the researchers could isolate the effects of different components of the waggle dance on the bees' foraging behavior, thus confirming that abstract information on distance and direction of a source of food were carried by the waggle dance. Other such studies that test existing theories are listed in the first column of Table 5.1.

A second way in which robots are used is to model a particular social behavior with a general question in mind, and to explore whether the implicit inclusion of physical properties in robotic models will reveal novel aspects concerning the collective behavior in question. Results of such studies can then be used to formulate new, empirically testable hypotheses (see "exploratory" studies in Table 5.1). For example, Krieger et al. (2000) used a swarm of small, ant-like robots to explore how collective behavior and division of labor may influence foraging

	Testing	Exploratory	Gap-filling
Robots or simulated robots	Floreano et al., 2007 (ch. 2) Garnier et al., 2008 Marocco et al., 2003 Waibel, 2007 Chapter 4	Baldassarre et al., 2006 Belpaeme & Birk, 1997 Birk & Wiernik, 2002 **Di Paolo, 2000** **Holland & Melhuish, 1999** Krieger & Billeter, 2000 **Krieger et al., 2000** **Kube & Bonabeau, 2000** Marocco et al., 2003 Marocco & Nolfi, 2006 Mataric, 1993 McFarland, 1994 Mitri et al., 2009 (ch. 3) **Nolfi & Floreano, 1998** Quinn, 2001 Rubenstein et al., 2009 Steels & Vogt, 1997 Steels, 1998 Vogt, 2000 Wischmann et al., 2006a Wischamnn et al., 2006b	Groß et al., 2008 **May et al., 2006** **Melhuish et al., 2006** Mitri et al., 2009 (ch. 3)
Mixed models	**Göth & Evans, 2004** Halloy et al., 2007 Michelsen et al., 1992 Taylor et al., 2008	Böhlen, 1999 **Fernández-Juricic et al., 2006** Kubinyi et al., 2004 Takanishi et al., 1998	**Fernández-Juricic et al., 2006** **Göth & Evans, 2004** **Martins et al., 2005** **Ord & Stamps, 2008** Partan et al., 2009 **Patricelli et al., 2002** **Reaney et al., 2008** Rundus et al., 2007

Table 5.1: The articles in this table were chosen because they report on research using robots or simulated robots (based on a physical robot), and explicitly state that they aim to understand social behavior in living organisms. One representative article was chosen when numerous articles were found that all had similar conclusions. Citations in bold have generated new, concrete, testable hypotheses.

efficiency. Although their original intention was not to study the effects of different group sizes, they found that as the number of robots in a group grew beyond a certain size, the robots' foraging efficiency began to drop. They thus hypothesized that increased interference between individuals through crowding may be responsible for the decrease in efficiency and that such effects may also limit group size in social insect colonies – a conclusion that could hardly have been drawn from a model that does not contain the physical constraints of moving bodies in the real environment.

A third and final use of robots is to provide evidence that may explain an observed natural phenomenon for which previous research had provided contradictory or no evidence. We will designate studies of this type as "gap-filling", since they fill a gap in scientific knowledge (see Table 5.1). We differentiate such studies from studies of the "testing" type, in which well-established theories that have already been modeled at lower levels of situatedness are tested in a robotic scenario. By contrast, "gap-filling" studies test novel hypotheses that constitute open questions. For example, ground squirrels have been found to repeatedly raise their tails at threatening rattlesnakes, even in a dark experimental room, where the visual component of the signal was completely absent. This led researchers to hypothesize that there might be an infrared component to the squirrels' signal. By building a robotic squirrel that emitted such infrared signals and observing the snakes' reaction to it, Rundus et al. (2007) provided evidence that this was indeed the case. Similarly, the study presented in chapter 3 has provided a novel hypothesis, supported by experimental evidence that the variability in natural communication strategies when interests between communicating individuals conflict, may be explained by the interplay between genetic variability introduced through mutations and weak selection pressure on signaling and response behaviors.

Some patterns stand out in the list in Table 5.1. Firstly, it is notable that the majority of articles ($\sim$72%) that report on using robots or simulated robots to study collective behavior are of the exploratory type. Although the power of exploratory studies lies in formulating concrete, testable hypotheses as a result of using robots, we found that only few studies of this type ($\sim$24%) generated such novel hypotheses. By contrast, we found that hypotheses were often generated in the two other types of studies (testing and gap-filling studies).

Another observation is that gap-filling studies were pre-dominantly mixed models ($\sim$73%), and that half of the mixed models were of the gap-filling type. This is likely to be because mixed models are often designed by researchers already working on a living system, who then turn to robots to address a specific question that they could not have answered otherwise, such as the role of different signal components on response strategies (see previous section). In contrast, studies involving robots or simulated robots alone seemed to either be used to test theories or to address more general, exploratory questions.

5.4 Lessons learned: when and how to use robots

Models can result in scientific contributions if they accomplish one of three things: either allow the testing of an existing theory, provide a novel explanation to a puzzling biological phenomenon, or generate novel testable hypotheses. Among the studies covered in this review, many can be considered to have made such contributions to our understanding of social systems. These studies have shown how some theories can stand the test of reality (e.g., Michelsen et al. (1992), Floreano et al. (2007) (chapter 2), Waibel (2007), chapter 4), they have addressed questions that were considered difficult or even impossible to answer (e.g., Rundus et al. (2007); Taylor et al. (2008)), and they have inspired us by proposing alternative, non-intuitive answers to long-standing questions (e.g., May et al. (2006); Reaney et al. (2008)). However, some studies – although potentially interesting – have so far contributed little to our understanding of social behavior. What lessons can we take home from these studies? What are the promising directions to follow? And how can we avoid the pitfalls?

Of the three types of studies that we have considered in this review (Table 5.1), testing and gap-filling studies have been found to contribute the most to the field (in the terms defined above). Although exploratory studies can and do make important contributions, this is only the case if researchers use the findings of the model to generate relevant hypotheses and ensure that they lead to concrete follow-up studies. Furthermore, they are risky endeavors, since there is no guarantee that any interesting findings will be made, and because they often set out to answer broad questions, which are more difficult to answer (Jacob, 1977). We therefore suggest that research is more likely to lead to fruitful outcomes if robots are used to test existing theories, which are difficult or impossible to test using living organisms. Understanding whether theoretical predictions of social behavior hold at all levels of situatedness, independently of various environmental factors, is important both for theoretical and empirical biologists. In addition, we suggest that robots be used to answer precise questions, for which they provide a clear advantage over simpler modeling methods.

In section 5.2 we have suggested that robots may be preferred over agent-based models in two types of studies. Firstly, when physical features of the real world, such as sensory noise, friction, visual or spatial effects, are expected to play an important role in shaping the social behavior in question. And secondly, when the question of the study is specifically about the influence of such physical properties on social behavior. At first sight, it may seem that we are arguing that robots are restricted to questions regarding mechanistic aspects of social behavior and cannot be used to answer questions regarding its ultimate (evolutionary) causes. However, a potential use of robots would be to study the *interplay* between effects at the behavioral level that may be highly dependent on physical factors, and effects at the level of evolutionary dynamics. Mayr (1961) and Tinbergen (1963) were perhaps the first to argue for the synergistic benefits of a complementary approach addressing both proximate (life-time) and ultimate (evolutionary) causes. Our experiments in chapter 3 provide a good example of this complementary approach, by showing that inadvertent information produced by the spatial distribution of signaling robots had an important influence on the evolutionary dynamics of the communication system. If the behavior of robots, the spatial properties of their environment and their perceptual mechanisms can influence evolutionary dynamics, we must suppose that the mechanistic properties of a system may indeed be important in understanding the bigger evolutionary picture. Robotic models may thus play a key role in addressing ultimate questions in sociobiology.

This complementary approach can be taken one step further by using mixed models to study the interplay between behavior and evolutionary processes. A nice illustration of this idea is provided by a study by Bond and Kamil (2002), in which blue jays searched for digital moths on computer monitors. The digital moths that were not pecked by the birds survived to subsequent generations, thus allowing the researchers to determine how selection shapes the coloration of real moths. Although this study did not involve any robots, we believe it may inspire similar studies using mixed models where the robots' controllers or morphology can evolve over time. Such studies provide a unique opportunity to manipulate behavioral processes of living organisms and to explore how they affect the evolution of social systems.

As robots become more affordable and easier to construct and manipulate, we predict that the use of robots in studies of social behavior will increase. In this chapter, we have shown how robotic models can address questions that are difficult to answer with more traditional methods, test well-established theories in a realistic setup or uncover interesting phenomena that have gone unnoticed. However, this research is only possible if collaborations between researchers from different disciplines are formed, if roboticists focus on well-formulated questions to which the robots can make important contributions and biologists recognize problems for which robotics may provide a solution.

6

Conclusions

Communication is the glue that holds societies together.

Carl Zimmer, 2008

By applying experimental evolution to groups of robots that could use communication to more efficiently solve a foraging task, this thesis has resulted in a number of contributions that may help in better understanding the forces that govern the evolution of communication in nature. Moreover, throughout this work, we have drawn some lessons regarding the different approaches to studying this question. In this chapter we will summarize the work presented in previous chapters, discuss the overall significance of our findings and how they can inspire future research, both in theoretical and empirical studies.

6.1 Original contributions

6.1.1 Understanding the evolution of communication

The work presented in this thesis has illustrated how co-evolution between signalers and receivers can lead to the organization of stable communication systems. It has shown that for such systems to evolve, communicating individuals do not require complex cognitive skills. Indeed, our robots were equipped with very simple, feed-forward neural networks with no hidden layers and a total of 33 connection weights. These simple "brains" not only allowed robots to use a range of signaling strategies with varying levels of reliability, and the corresponding response strategies, but they also made it possible for robots to forage efficiently in a noisy physical environment.

In the introductory chapter, we formulated the hypothesis that both relatedness within groups of robots, and the level at which robots are selected (group or individual) would influence the evolution of their communication strategies. In chapter 2 we have shown that this is in fact the case. Groups in which robots were clones evolved signaling strategies which reliably indicated the location of the food source, while the robots simultaneously evolved to respond appropriately to these signals. Communication was thus an adaptive behavior that increased the performance of these highly related robots. These findings confirm predictions (Keller and Surette, 2006; Diggle et al., 2007a) that communicating individuals that share identical genes by descent would evolve to produce reliable signals, because this allows them to achieve higher inclusive fitness, thus making them more likely to be selected for subsequent generations. Additional experiments presented in chapter 4 confirmed this finding by directly measuring signal reliability and showing that reliability increases gradually with increasing within-group relatedness.

Similar adaptive communication strategies were observed in the experiments presented in chapter 2 where robots were selected at the level of the group. This is because selection at the level of the group reduces competition within the group, since the likelihood of survival depends not only on one's own performance, but also on the performance of other group members (Frank, 1998; Rainey and Rainey, 2003; Griffin et al., 2004).

These initial experiments in chapter 2 also revealed that having neither of these two mechanisms (high relatedness or group-level selection) to allow for honest communication led to the evolution of communication strategies that decreased the performance of the robots compared to a situation in which no communication was possible. A detailed analysis of this case was carried out in chapter 3, revealing that through their foraging behavior, robots inadvertently transmitted information on the location of the food and were thus selected to reduce light emission by food and instead produce signals in other areas of the arena. Essentially, this amounted to a suppression of the inadvertently produced information, which made it more difficult for other robots to find the food.

This shows that deception is not the only outcome to be expected when interests conflict (Dawkins and Krebs, 1978). Instead, simply suppressing information that inadvertently provides cues to other individuals may be a way to alter signaling strategies in the presence of competition. These strategies can be likened to information suppression observed in groups of conspecifics that conceal information from one another, for example, by not producing vocalizations on finding a source of food (Hauser, 1992; Laidre, 2006). Similarly, many prey are selected to provide as little information as possible to their predators regarding their location, and thus evolve to blend into their environment (Stevens and Merilaita, 2009). The work presented in chapter 3 thus constitutes a first study of the role of inadvertent information in the evolution of communication. Although such cues are common in natural communication systems (Giraldeau, 1997; Danchin et al., 2004; Bonnie and Earley, 2007), they may have been overlooked in models of communication because these models rarely couple signaling behaviors with other behaviors. Our study thus shows that the ensemble of an individual's behaviors may carry more information than signaling alone, and must therefore be taken into account when analyzing communicative behavior.

Another interesting observation made in the study in chapter 3 was that complete suppression of information was never achieved. Instead, the system stabilized at an equilibrium where little information was transmitted and response to the information was weak. This was surprising, since one would have expected robots to evolve to completely eliminate any information on food location. However, this "suboptimal" equilibrium was found to be a result of the interplay between variation introduced through mutations and weak selection pressure on information suppression and the response to signaling. Initially, robots were selected to suppress the inadvertent information produced by their random signals, which weakened selection on attraction to the signals. In turn,

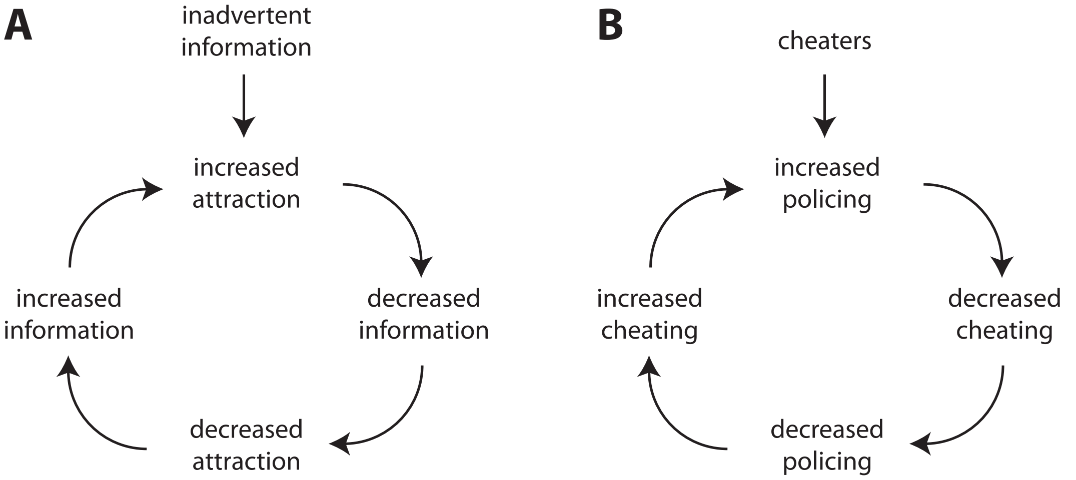

Figure 6.1: The interplay between variation and weak selection pressure on (A) information suppression and response to signals leading to the equilibrium observed in chapter 3, may have parallels in (B) societies where cheaters are suppressed through policing.

this lowered the strength of selection on information suppression, because producing information no longer had such dire consequences on a signaler's performance. The subsequent rise in information resulted in an increase in the response to the signal, leading again to increased selection for information suppression (Fig. 6.1A). The resulting equilibrium also led to high variation both in signaling and response strategies. These findings may provide a possible explanation to the long-standing question of high intra-population polymorphism observed in signaling systems where interests between signalers and receivers conflict (Poulton, 1890; Whiteley et al., 1997; Dewitt et al., 1999; Bond and Kamil, 2002; Bond, 2007). In addition, such equilibria may not be restricted to systems in which information is suppressed, but may apply to other co-evolutionary systems in which interests conflict, such as systems in which cheaters are punished through policing. If policing leads to a reduction in the amount of cheating, it is expected to be subject to weak selection, resulting – similarly to the observations made in chapter 3 – in an equilibrium where both policing and cheating are maintained (Fig. 6.1B). In fact, the introduction of variation into a population of agents through mutations has been shown to result in stable equilibria, in which policing and cheating are maintained (Foster and Kokko, 2006; McNamara et al., 2008).

6.1.2 Modeling the evolution of communication

Many of the findings discussed in the previous section, which have contributed to our understanding of the evolution of communication can be attributed to the modeling approach used in this thesis.

In the introductory chapter, we argued that modeling evolutionary processes, as opposed to simply calculating optimal strategies, could lead to important insights. Throughout this thesis, we have found this to be the case. Firstly, in chapter 2, two reliable signaling strategies evolved. In some experimental replicates, robots evolved to signal by the food and to be attracted to blue light, thus increasing overall performance. However, an alternative strategy was found in other experimental replicates, where blue light was used to signal the presence of poison, while robots were repulsed by blue light, thus allowing them to find the food. Although this strategy was less efficient than the food signaling strategy, once this poison strategy had evolved, the robots rarely switched to the more efficient food signaling strategy because switching entailed crossing a “valley” of low performance. These findings may explain the existence of the large variety of signaling strategies in nature, although they may not all be equally efficient (Von Frisch, 1967). A second finding that depended on the stochasticity of our evolutionary model, was the low-information equilibrium and high polymorphism discussed in chapter 3 and in the previous section. Simply optimizing the strategy of signalers would have led to no signaling at all. Instead, mutations allowed for the persistence of information and the observed equilibrium.

A second important feature of our approach is, of course, the use of robots. Many of the contributions in this thesis would not have been made if we had used a more abstract model. This is not because it is impossible to model the important aspects of the robotic model mathematically or in software. In fact, our results are based on software simulations of the physical robots. Rather, we believe that many of the features of the simulation were included because we were aiming to precisely model physical robots, and not because we expected them to affect the evolving communication system. For example, the findings of chapter 3 that the robots’ behavior produced inadvertent information regarding the location of food was simply a side-effect of the spatial features of the arena, and of the visual systems of the robots. Importantly, these features are not only specific to the robots, but are also common features of real communication systems (Danchin et al., 2004; Dall et al., 2005). Similarly, the measure of signal reliability introduced in chapter 4 would not have been necessary in a simpler model. It would have perhaps been sufficient to study signal frequencies, as is done in many other models (Searcy and Nowicki, 2005). Our model has thus revealed that in order to integrate the effects of individual behavior into a measure of signal reliability, reliability must be measured from the perspective of the receiver (Rowe, 1999). This realization has indeed been taken into account in some early studies characterizing communication systems (Haldane and Spurway, 1954; Wilson, 1962; Conant and Steinberg, 1973; Fuchs, 1976; Lenoir, 1982), but has been neglected in studies on signal reliability (Searcy and Nowicki, 2005).

Our experiments have thus revealed that evolutionary processes, spatial properties of the environment, and perceptual effects of using vision were all important in our model and in the obtained results. These results echo what we have found through the literature review presented in chapter 5, where we concluded that robots are useful when addressing questions in which mechanistic properties of the individuals and their interactions with the environment are expected to influence the resulting social behavior and its evolutionary dynamics.

6.2 Outlook and future directions

The experiments presented in this thesis were not designed to reflect any specific communication system in nature. The decision to use visual communication, for example, was based on the capabilities of the robots' hardware, rather than a desire to specifically study visual communication systems. However, it would be interesting to see whether changes in the communication channel, for example to using chemical communication, would influence the evolutionary outcome of the system. This might be expected because in chemical communication systems, individuals perceive their own signals, which is not necessarily the case in visual communication. Pheromones can thus be used as a form of memory to mark one's own path, or to remain where the food was found, and not only to the benefit of other individuals. Changing evaporation rates could then be used as a means to regulate the strength of this effect.

Another possibility would be to study the effect of increasing the number of food patches in the arena. This would change the dynamics of the foraging behavior, because robots would then no longer be attracted to areas with the highest concentration of robots, but rather with intermediate concentrations, indicating that the location of a food source that is not yet overcrowded. It is unclear how these changes would influence the dynamics of the communication system. The effects of food distribution would be interesting to explore because they have been known to play an important role in shaping natural foraging strategies (MacArthur and Pianka, 1966; Giraldeau and Caraco, 2000).

A third direction would be to explore the role of signal cost on signal evolution. As discussed in chapter 1, cost has been an important factor in signaling theory, and is expected to play an important role in many natural communication systems (Maynard Smith and Harper, 2003; Searcy and Nowicki, 2005). In addition, it would be easy to integrate into our current model by simply subtracting performance points from robots that produce signals.

In addition to exploring the dimensions of the model further, this thesis has provided a number of hypotheses that can be tested empirically. For example, it would be interesting to explore whether the stable equilibrium observed in chapter 3 is in fact mirrored in natural systems, both in communication and in other co-evolutionary systems where interests conflict (Fig. 6.1), such as policing, or sanctioning in mutualistic interactions (Foster and Kokko, 2006; McNamara et al., 2008). Our results predict that in all these systems, an equilibrium of low levels of the "suppression" and "cheating" behaviors are expected, with high variation in both. Another testable prediction is that signal reliability increases with relatedness, as observed in chapter 4. It would be interesting to apply the measure of reliability outlined in chapter 4 to a natural communication system and to study whether changes in kin structures will show qualitatively similar changes in the level of signal reliability as we have found in this work.

As the quote by Carl Zimmer at the beginning of this chapter suggests, communication is a fundamental ingredient to any social behavior. It is thus essential to understand communication in order to fully grasp the mechanisms driving the evolution of social systems. With the experimental work presented in this thesis, we have taken some small steps toward this goal. Moreover, we hope to have convinced the reader that there is much potential in using mechanistic models to explore questions regarding social evolution, and that such approaches will play their role in completing the big picture of communication and social behavior.

A

Data Analysis

In this appendix, we outline the way in which performance, signaling strategies, response strategies and selection pressure are computed in the experiments discussed in chapters 2, 3 and 4. We also outline the statistical methods used for comparisons between experimental treatments in these chapters.

A.1 Performance

Performance was calculated for each robot as the average performance over 10 trials. In each trial, performance was calculated as the number of time-steps t_f the robots spent by the food f minus the number of time-steps t_p spent by the poison p, normalized by the total number of time-steps in a trial T. Performance could thus vary between -1 and 1.

A.2 Signaling

The total frequency of light emission S was computed as the mean number of time-steps L robots spent emitting light divided by the mean number of time-steps in a trial T. This measure is used in chapter 4 (Fig. 4.4A).

The frequency of light emission s_o when in the vicinity of object o, where $o \in O = \{f, p, n\}$ (f stands for food, p for poison, and n for elsewhere in the arena) was computed as the mean number of time-steps l_o robots spent emitting light in the vicinity of object o divided by the mean number of time-steps t_o they spent in the vicinity of o:

$$s_o = \frac{l_o}{t_o} \tag{A.1}$$

A robot was considered in the vicinity of food or poison if touching the paper disc placed under the food or poison. Otherwise, the robot was counted as being elsewhere in the foraging arena. This measure is used in chapters 3 and 4 (Fig. 3.3A, 3.4A and 4.3).

The signaling strategy s used in chapter 2 (Fig. 2.7) was quantified for each robot by estimating the average frequency of signaling near food f and poison p. This was quantified by:

$$s = s_f - s_p \tag{A.2}$$

The value of s can therefore vary from -1 to 1, with a value of -1 indicating that robots signaled only when near the poison and a value of 1 that signaling occurred only when near the food. A value of 0 would indicate that robots were not more likely to signal near food or poison.

Finally, in chapter 3, we refer to the concentration of blue light $\hat{s}_o$ in the three areas o of the arena, which was computed as:

$$\hat{s}_o = \frac{l_o}{a_o} \left(\sum_{i \epsilon O} \frac{l_i}{a_i} \right)^{-1} \tag{A.3}$$

where a_o is the proportion of the space of the arena occupied by object o ($a_f = a_p = 0.024, a_n = 0.953$).

A.3 Response

The level of response to blue light b was measured by placing each robot at a distance of $x = 35$cm and $y = 35$cm from a second stationary robot emitting blue light, and recording the position of the moving robot relative to its original position after 10 cycles. A decrease in both dimensions ($x < 35$ and $y < 35$) was considered to be attraction, an increase in both dimensions ($x > 35$ and $y > 35$) as repulsion, whereas other outcomes were not considered (see Fig. A.1). This test was performed N times for each robot, and the response b to blue light was calculated as:

$$b = \frac{i - j}{N} \tag{A.4}$$

where i is the number of attractions and j the number of repulsions.

Four tests per robot ($N = 4$) were used to generate the data in chapter 2 (Fig. 2.7), ten tests ($N = 10$) to generate the data in Fig. 3.2B, and 100 tests ($N = 100$) for the data in Fig. 3.4B and 3.5.

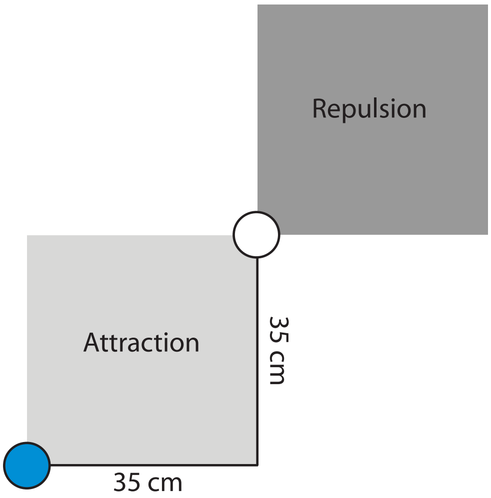

Figure A.1: The response of the white robot to the robot producing blue light was recorded as attraction if the robot moved in the light gray box and as repulsion if it moved in the dark gray box (nothing was recorded if the robot moved in the white space after the 10 cycles). This test was repeated N times for each robot.

A.4 Selection pressure

Selection pressure was computed separately for each of the 20 experimental replicates in chapter 3 (Fig. 3.3). It was equal to Spearman's rank correlation coefficient between signaling frequency by the food s_f and performance p of the 1000 individuals in the population of robots in each replicate at each generation.

A.5 Statistical analysis

In all experiments presented in this thesis, 20 replicates were conducted for each experimental treatment.

To compare various measurements (e.g., performance, light emission frequencies, response to blue light, etc.) between treatments over many generations, we calculated average values of all individuals in the population (1000 in chapters 2 and 3 and 800 in chapter 4) for each of the 20 replicates per treatment at each generation. The resulting 20 values per treatment, per generation were used to describe data (mean ± standard deviation or standard error) and were compared with non-parametric (Kruskal-Wallis and Mann-Whitney) tests as some of the data did not follow a normal distribution (Kolmogorov-Smirnov test).

When analyses were carried out at the end-point of the evolutionary experiments, population average values were again averaged over the last g generations for each of the 20 replicates per treatment ($g = 50$ in chapter 2, $g = 10$ in chapter 3 and $g = 20$ in chapter 4), resulting in 20 values per treatment. These values were again compared using non-parametric tests.

Within-population variance was computed for each of the 20 replicates by taking the average of the standard deviations in the population in each of the last g generations ($g = 10$ in chapter 3, where variance was used). Non-parametric (Mann-Whitney) tests were used to compare the resulting 20 values across experiments.

B Information, Reliability and Relatedness

In this appendix, we provide details on the methods used to compute measures of information, reliability and within-group relatedness. These measures are used in chapters 3 and 4.

B.1 Information

To quantify the information content transmitted by blue light, we analyzed the inputs of each robot's omni-directional visual system to establish whether robots perceived more blue light in the direction of the food. This was done by (i) ranking the four quadrants of the robot's visual system by the amount of light perceived at each time-step until the robot reached the food (i.e., the quadrant with the largest amount of blue light was assigned rank $q = 1$, etc., where $q \in Q = \{1, 2, 3, 4\}$) and (ii) computing the ratios p_{Iq} of time-steps where the food was located in each of the four quadrant ranks q (Fig. B.1, white bars).

The amount of information on food location provided by blue light emission was then calculated using the Shannon entropy (Shannon, 1948) (uncertainty) of the frequency distribution $X_I = \{p_{I1}, p_{I2}, p_{I3}, p_{I4}\}$ using the following equation:

$$H(X_I) = -\sum_{q \in Q} p_{Iq} log p_{Iq} \tag{B.1}$$

This entropy value H ranges from $H_{min} = H(X_{min}) = 0$, when there is no uncertainty on food location (e.g., blue light is perceived only in the direction of the food, $X_{min} = \{1, 0, 0, 0\}$) to $H_{max} = H(X_{max}) = 1.39$, when uncertainty on the location of food is maximal, (i.e., there is no association between blue light intensity and food location, $X_{max} = \{0.25, 0.25, 0.25, 0.25\}$). Accordingly, the in-

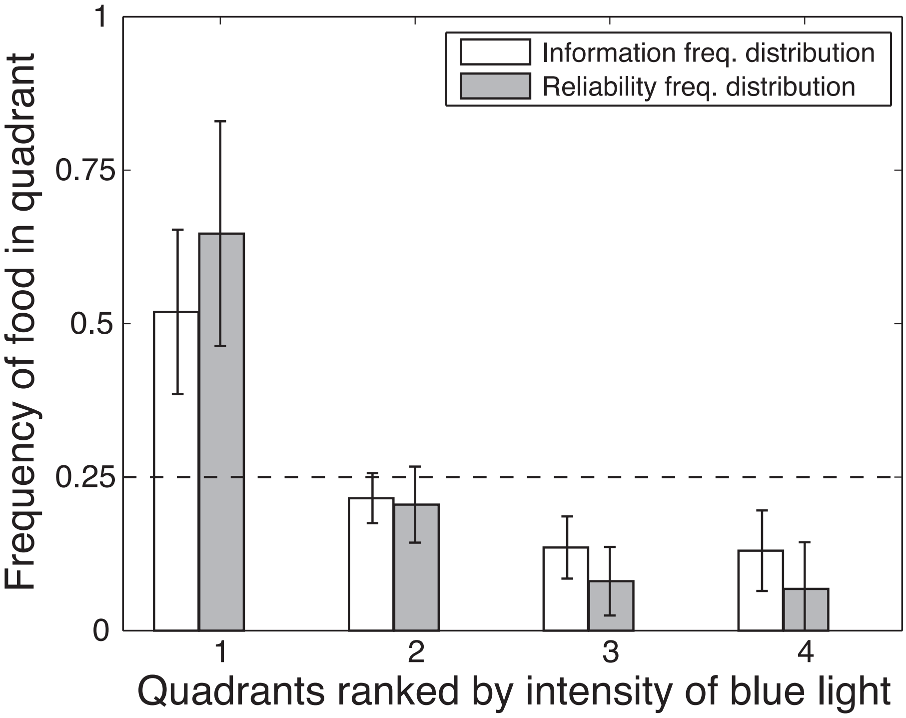

Figure B.1: Association between the intensity of blue light perception and food location given as the proportion of time steps where the food was located in the quarter of the 360° camera image with greatest, second greatest, third greatest and lowest intensity of blue light. The bars represent means and the error bars standard deviations of the 20 independent replicates from generations 491 to 500 for the frequency distribution used to compute information I (white, $I = 0.16 \pm 0.09$) and reliability R (gray, $R = 0.36 \pm 0.17$) for the experiments presented in chapter 4 when relatedness was 1 and robots were not blind to blue light. The horizontal dashed line gives the expected value when there is no association between intensity of blue light in the quadrants and food location (i.e., $I = 0$ and $R = 0$). Maximum information and reliability ($I = 1$ and $R = 1$) would be achieved if one of the 4 columns would be at frequency 1 and the other three at frequency 0.

formation content I is defined as the difference between the maximum entropy H_{max} and entropy $H(X_I)$, normalized by H_{max}, thus yielding values of I between 0 (when the intensity of blue light provides no information of the location of the food source) and 1 (when the intensity of blue light perceived provides perfect information on the location of the food source at all time-steps):

$$I(X_I) = \frac{H_{max} - H(X_I)}{H_{max}}. \tag{B.2}$$

B.2 Reliability

To measure signal reliability R, we followed a similar procedure as with information, except that instances in which a robot could perceive no blue light were excluded in the computation of the probability distribution. Instead, to compute reliability, we (i) ranked the four quadrants of the robot's visual system by the amount of light perceived at each time-step *in which the robot could perceive light* until the robot reached the food and (ii) computed the ratios p_{Rq} of time-steps where the food was located in each of the four quadrant ranks q (Fig. B.1, gray bars).

The level of signal reliability was then calculated using the Shannon entropy of the frequency distribution $X_R = \{p_{R1}, p_{R2}, p_{R3}, p_{R4}\}$ using the following equation:

$$H(X_R) = -\sum_{q \in Q} p_{Rq} log p_{Rq}, \tag{B.3}$$

calculating the difference between the maximum entropy H_{max} and $H(X_R)$, and normalizing it by H_{max} to obtain values of R between 0 (when all perceived signals were equally distributed in all four directions with respect to the food) and 1 (when all perceived signals were in the same direction with respect to the food):

$$R(X_R) = \frac{H_{max} - H(X_R)}{H_{max}}. \tag{B.4}$$

B.3 Relatedness

To create a group of size n with a level of relatedness r, we composed the group of k different types of clones, with $k \leq n$. The frequency x_i of each type of clone i could vary between $1/n$ and 1, such that the sum of x_i for all values of $i = 1, .., k$

r	k	x_i, $i = 1..k$
0	8	$\frac{1}{8}; \frac{1}{8}; \frac{1}{8}; \frac{1}{8}; \frac{1}{8}; \frac{1}{8}; \frac{1}{8}; \frac{1}{8}$
0.25	3	$\frac{3}{8}; \frac{3}{8}; \frac{2}{8}$
0.54	3	$\frac{6}{8}; \frac{1}{8}; \frac{1}{8}$
0.75	2	$\frac{7}{8}; \frac{1}{8}$
1	1	$\frac{8}{8}$

Table B.1: Number k and proportions x_i of clones used to compose groups of 5 different levels of relatedness r used in chapter 4.

was equal to 1. For each type of clone i, we computed the probability that a signal emitted by an individual of type i would be perceived by its clone:

$$P_i = \frac{nx_i - 1}{n - 1}. \tag{B.5}$$

By averaging P_i over all types of clones in a group, we obtain relatedness r:

$$r = \sum_{i=1}^{k} x_i \cdot P_i. \tag{B.6}$$

The numbers of different types of clones k and their frequencies x_i are given in table B.1 for the relatedness values used in chapter 4 (0, 0.25, 0.54, 0.75 and 1). Note that a group size of 8 was chosen for these experiments (instead of 10 in chapters 2 and 3) because it made it possible to construct colonies with relatedness levels that were evenly distributed between 0 and 1. A relatedness level of 0.54 was used because it was the nearest value to 0.5 that was achievable with this approach and a group size of 8.

Glossary

Much of the terminology used in this thesis has previously been defined in the literature. This glossary is provided to clarify which definitions we subscribe to. We also define a number of terms that are either rarely explicitly defined, or have a definition that is specific to this thesis.

Altruism. "A behavior which is costly to the actor and beneficial to a recipient; [..] cost and benefit are defined on the basis of the lifetime direct fitness consequences of a behavior" (West et al., 2007b).

Altruistic communication. A behavior or phenotype that alters the behavior of another organism or other organisms that is costly to the actor, and beneficial to the receiver(s), and which has evolved because of its beneficial effect on the receiver(s).

Communication. A behavior or phenotype that alters the behavior of another organism or other organisms, which has evolved because it is beneficial to either one or both of the participants (Wilson, 1975).

Cooperation. "A behavior which provides a benefit to a recipient, and which is selected for because of its beneficial effect on the recipient" (West et al., 2007b).

Cue. A behavior or phenotype that alters the behavior of another organism or other organisms, which is beneficial to the receiver(s), but has not evolved because of that effect (Keller and Surette, 2006; Scott-Phillips, 2008).

Deceptive signal. A *signal*, which has evolved because of its negative effect on the receiver(s) compared to other individuals that do not receive the signal. Deceptive signaling is also commonly referred to as "coercion" (Diggle et al., 2007a; Scott-Phillips, 2008).

Fitness. "The ability of organisms [..] to survive and reproduce in the environment in which they find themselves" (Orr, 2009). In the context of this thesis, it is equivalent to performance.

Game-theoretical model. See *game theory*.

Game theory. "A mathematical formalism used to [..] find the optimal strategy for one player to use when [its] opponent also plays optimally" (Flake, 2000).

Honest signal. A *signal*, which has evolved because of its positive effect on the receiver(s) compared to other individuals that do not receive the signal.

Inclusive fitness. *Fitness* from personal survival and reproduction plus *fitness* component received from effects on the survival and reproduction of relatives (Foster et al., 2006).

Individual-based model. Models that describe individuals as discrete, autonomous entities (Huston et al., 1988).

Information. Generally defined as "a reduction in uncertainty" (Scott-Phillips, 2008). For the mathematical definition used in this thesis, see Appendix B.1.

Kin selection. "Process by which traits are favoured because of their beneficial effects on the fitness of relatives" (West et al., 2007b).

Mixed model. Study in which a robot or multiple robots are made to interact with a living organism or organisms.

Performance. In the context of this thesis, it refers to the efficiency of an individual at accomplishing a task or a number of tasks during its lifetime.

Relatedness. Generally defined as the "genetic correlation among individual loci or organisms" (Foster et al., 2006). In the context of this thesis, it denotes the probability of a *signal* receiver being the signaler's clone, see Appendix B.3.

Reliability. In the context of this thesis, probability that a received *signal* will have a positive effect on the receiver(s) compared to other individuals that receive a *cue*. For the mathematical definition used in this thesis, see Appendix B.2.

Response. A behavior that evolved because of the effect of a *signal*.

Signal. A behavior or phenotype that alters the behavior of another organism or other organisms, which evolved because of that effect, and that is effective because the receiver's *response* has also evolved. (Keller and Surette, 2006; Scott-Phillips, 2008). Note that signals can be *honest* or *deceptive*, but differ from *cues*.

Signal cost. Reduction in *fitness* as a result of *signal* production.

Signal reliability. See *reliability*.

Situatedness. The extent to which individuals are embedded in an environment that can be sensed and modified by those individuals (Varela et al., 1991; Clark, 1996)

Social behavior. Any interaction that involves two or more individuals from the same or different species.

Unreliable signal. A *signal*, which has evolved because of its negative effect on the receiver(s) compared to other individuals that receive a *cue*.

Bibliography

Ackley, D. and Littman, M. (1994). Altruism in the evolution of communication. In Brooks, R. and Maes, P., editors, *Proceedings of the Fourth International Workshop on the Synthesis and Simulation of Living Systems*, pages 40–48, Cambridge, MA. MIT Press.

Adams, E. S. and Caldwell, R. L. (1990). Deceptive communication in asymmetric fights of the stomatopod crustacean Gonodactylus bredini. *Animal Behaviour*, 39:706–716.

Ampatzis, C., Tuci, E., Trianni, V., and Dorigo, M. (2008). Evolution of signaling in a multi-robot system: Categorization and communication. *Adaptive Behavior*, 16:5–26.

Andersson, M. (1980). Why are there so many threat displays? *Journal of Theoretical Biology*, 86:773–781.

Baldassarre, G., Parisi, D., and Nolfi, S. (2006). Distributed coordination of simulated robots based on self-organization. *Artificial Life*, 12:289–311.

Batali, J. (1994). Innate biases and critical periods: Combining evolution and learning in the acquisitions of syntax. In Brooks, R. and Maes, P., editors, *Proceedings of the Fourth International Workshop on the Synthesis and Simulation of Living Systems*, pages 160–171, Cambridge, MA. MIT Press.

Belpaeme, T. and Birk, A. (1997). A real-world ecosystem featuring several robot species. In Almeida E Costa, F., Rocha, L., Costa, E., and Harvey, I., editors, *The European Conference on Artificial Life ECAL97*, Berlin. Springer.

Bergstrom, C. T. and Lachmann, M. (1998). Signaling among relatives. III. Talk is cheap. *PNAS*, 95:5100–5105.

Billard, A. and Dautenhahn, K. (1999). Experiments in learning by imitation – grounding and use of communication in robotic agents. *Adaptive Behavior*, 7:415–438.

Birk, A. and Wiernik, J. (2002). An N-player prisoner's dilemma in a robotic ecosystem. *Robotics and Autonomous Systems*, 39:223–233.

Böhlen, M. (1999). A robot in a cage-exploring interactions between animals and robots. In *IEEE International Symposium on Computational Intelligence in Robotics and Automation*, volume 12, pages 214 – 219, Monterey, CA. IEEE Press.

Boncoraglio, G., Caprioli, M., and Saino, N. (2009). Fine-tuned modulation of competitive behaviour according to kinship in barn swallow nestlings. *Proc R Soc Lond Biol*, 276:2117–2123.

Bond, A. B. (2007). The evolution of color polymorphism: Crypticity, searching images, and apostatic selection. *Annual Review of Ecology, Evolution, and Systematics*, 38:489–514.

Bond, A. B. and Kamil, A. C. (2002). Visual predators select for crypticity and polymorphism in virtual prey. *Nature*, 415:609–613.

Bonnie, K. E. and Earley, R. L. (2007). Expanding the scope for social information use. *Animal Behaviour*, 74:171–181.

Bourke, A. F. G. and Franks, N. R. (1995). *Social Evolution in Ants*. Princeton University Press, Princeton, NJ.

Braitenberg, V. (1984). *Vehicles: Experiments in synthetic psychology*. MIT Press, Cambridge, MA.

Brown, S. P. and Johnstone, R. A. (2001). Cooperation in the dark: Signalling and collective action in quorum-sensing bacteria. *Proc R Soc Lond Biol*, 268:961–965.

Bshary, R. and Grutter, A. S. (2006). Image scoring and cooperation in a cleaner fish mutualism. *Nature*, 441:975–978.

Buckley, N. J. (1997). Spatial-concentration effects and the importance of local enhancement in the evolution of colonial breeding in seabirds. *The American Naturalist*, 149:1091–1112.

Bugnyar, T. and Kotrschal, K. (2002). Observational learning and the raiding of food caches in ravens, Corvus corax: Is it "tactical" deception? *Animal Behaviour*, 64:185–195.

Bugnyar, T. and Kotrschal, K. (2004). Leading a conspecific away from food in ravens (Corvus corax)? *Animal Cognition*, 7:69–76.

Candolin, U. (1999). The relationship between signal quality and physical condition: Is sexual signalling honest in the three-spined stickleback? *Animal Behaviour*, 58:1261–1267.

Cangelosi, A. (2001). Evolution of communication and language using signals, symbols, and words. *IEEE Transactions on Evolutionary Computation*, 5:93–101.

Caro, T. M. (1994). Ungulate antipredator behaviour: Preliminary and comparative data from African bovids. *Behaviour*, 128:189 – 228.

Cheney, D. L. and Seyfarth, R. M. (1985). Vervet monkey alarm calls: Manipulation through shared information? *Behaviour*, 94:150–166.

Clark, A. (1996). *Being There: Putting Brain, Body, and World Together Again*. MIT Press, Cambridge, MA.

Conant, R. C. and Steinberg, J. B. (1973). Information exchanged in grasshopper interactions. *Information and Control*, 23:221–233.

Cotton, P. A., Kacelnik, A., and Wright, J. (1996). Chick begging as a signal: Are nestlings honest? *Behavioral Ecology*, 7:178–182.

Cummings, M. and Partridge, J. (2001). Visual pigments and optical habitats of surfperch (Embiotocidae) in the California kelp forest. *Journal of Comparative Physiology A: Sensory, Neural, and Behavioral Physiology*, 187:875–889.

Cummings, M. E. (2007). Sensory trade-offs predict signal divergence in Surfperch. *Evolution*, 61:530–45.

Dall, S. R. X., Giraldeau, L.-A., Olsson, O., McNamara, J. M., and Stephens, D. W. (2005). Information and its use by animals in evolutionary ecology. *Trends in Ecology & Evolution*, 20:187–193.

Danchin, E., Giraldeau, L.-A., Valone, T. J., and Wagner, R. H. (2004). Public information: From nosy neighbors to cultural evolution. *Science*, 305:487–491.

Dawkins, R. and Krebs, J. R. (1978). *Animal Signals: Information or Manipulation?*, chapter 10, pages 282–309. Blackwell Scientific Publications, Oxford.

DeAngelis, D. L. and Mooij, W. M. (2005). Individual-based modeling of ecological and evolutionary processes. *Annual Review of Ecology, Evolution, and Systematics*, 36:147–168.

Dewitt, T. J., Sih, A., and Hucko, J. A. (1999). Trait compensation and cospecialization in a freshwater snail: Size, shape and antipredator behaviour. *Animal Behaviour*, 58:397–407.

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 *Advances in Artificial Life*, volume 167 of *Lecture Notes in Computer Science*, pages 504–513. Springer, Berlin / Heidelberg.

Di Paolo, E. A. (2000). Behavioral coordination, structural congruence and entrainment in a simulation of acoustically coupled agents. *Adaptive Behavior*, 8:27–48.

Diggle, S. P., Gardner, A., West, S. A., and Griffin, A. S. (2007a). Evolutionary theory of bacterial quorum sensing: When is a signal not a signal? *Philos Trans R Soc London Biol*, 362:1241–1249.

Diggle, S. P., Griffin, A. S., Campbell, G. S., and West, S. A. (2007b). Cooperation and conflict in quorum-sensing bacterial populations. *Nature*, 450:411–414.

Dunford, C. (1977). Kin selection for ground squirrel alarm calls. *The American Naturalist*, 111:782 – 785.

Endler, J. A. and Basolo, A. L. (1998). Sensory ecology, receiver biases and sexual selection. *Trends in Ecology & Evolution*, 13:415–420.

Falconer, D. S. (1981). *Introduction to Quantitative Genetics*. Longman, New York.

Fernández-Juricic, E., Gilak, N., Mcdonald, J. C., Pithia, P., and Valcarcel, A. (2006). A dynamic method to study the transmission of social foraging information in flocks using robots. *Animal Behaviour*, 71:901–911.

Fiegna, F., Yu, Y.-T. N., Kadam, S. V., and Velicer, G. J. (2006). Evolution of an obligate social cheater to a superior cooperator. *Nature*, 441:310–314.

Fisher, R. A. (1958). *The Genetical Theory of Natural Selection*. Dover Publications, New York.

Flake, G. W. (2000). *The Computational Beauty of Nature: Computer Explorations of Fractals, Chaos, Complex Systems, and Adaptation*. MIT Press, Cambridge, MA.

Floreano, D. and Mattiussi, C. (2008). *Bio-Inspired Artificial Intelligence: Theories, Methods, and Technologies*. MIT Press, Cambridge, MA.

Floreano, D., Mitri, S., Magnenat, S., and Keller, L. (2007). Evolutionary conditions for the emergence of communication in robots. *Current Biology*, 17:514–519.

Fogel, D., Fogel, L., and Porto, V. (1990). Evolving neural networks. *Biological Cybernetics*, 63:487–493.

Foster, K. R. and Kokko, H. (2006). Cheating can stabilize cooperation in mutualisms. *Proc R Soc Lond Biol*, 273:2233–2239.

Foster, K. R., Wenseleers, T., and Ratnieks, F. L. W. (2006). Kin selection is the key to altruism. *Trends in Ecology & Evolution*, 21:57–60.

Frank, S. A. (1998). *Foundations of Social Evolution*. Princeton University Press, Princeton, NJ.

Franks, N. R. (1999). Information flow in the social domain: how individuals decide what to do next. In Detrain, C., Deneubourg, J. L., and Pasteels, J. M., editors, *Information Processing in Social Insects*, pages 101–114. Birkhäuser Verlag, Basel, Switzerland.

Fuchs, S. (1976). An informational analysis of the alarm communication by drumming behavior in nests of carpenter ants (Camponotus, Formicidae, Hymenoptera). *Behavioral Ecology and Sociobiology*, 1:315–336.

Galef, B. G. and Giraldeau, L.-A. (2001). Social influences on foraging in vertebrates: causal mechanisms and adaptive functions. *Animal Behaviour*, 61:3–15.

Garnier, S., Jost, C., Gautrais, J., Asadpour, M., Caprari, G., Jeanson, R., Grimal, A., and Theraulaz, G. (2008). The embodiment of cockroach aggregation behavior in a group of micro-robots. *Artificial Life*, 14:387–408.

Giraldeau, L.-A. (1997). The ecology of information use. In R., K. J. and Davies, N. B., editors, *Behavioural Ecology: An Evolutionary Approach*, pages 42–68. Blackwell, Oxford.

Giraldeau, L.-A. and Caraco, T. (2000). *Social Foraging Theory*. Princeton University Press, Princeton, NJ.

Göth, A. and Evans, C. S. (2004). Social responses without early experience: Australian brush-turkey chicks use specific visual cues to aggregate with conspecifics. *Journal of Experimental Biology*, 207:2199–2208.

Grafen, A. (1990). Biological signals as handicaps. *Journal of Theoretical Biology*, 144:517–546.

Griesser, M. (2009). Mobbing calls signal predator category in a kin group-living bird species. *Proc R Soc Lond Biol*, 276:2887–2892.

Griffin, A. and West, S. A. (2002). Kin selection: Fact and fiction. *Trends in Ecology & Evolution*, 17:15–21.

Griffin, A. S., West, S. A., and Buckling, A. (2004). Cooperation and competition in a pathogenic bacterium. *Nature*, 430:1024–1027.

Grimm, V. and Railsback, S. F. (2005). *Individual-based Modeling and Ecology*. Princeton University Press, Princeton, NJ.

Groß , R. and Dorigo, M. (2004). *Evolving a Cooperative Transport Behavior for Two Simple Robots*, pages 305–316. Springer, Berlin.

Groß, R., Nouyan, S., Bonani, M., Mondada, F., and Dorigo, M. (2008). *Division of Labour in Self-organised Groups*, volume 5040 of *Lecture Notes in Computer Science*, pages 426–436. Springer, Berlin.

Haldane, J. B. S. and Spurway, H. (1954). A statistical analysis of communication in Apis mellifera and a comparison with communication in other animals. *Insectes Sociaux*, 1:247–283.

Halloy, J., Sempo, G., Caprari, G., Rivault, C., Asadpour, M., Tâche, F., Saïd, I., Durier, V., Canonge, S., Amé, J. M., Detrain, C., Correll, N., Martinoli, A., Mondada, F., Siegwart, R., and Deneubourg, J. L. (2007). Social integration of robots into groups of cockroaches to control self-organized choices. *Science*, 318:1155–1158.

Hamilton, W. D. (1964). The genetical evolution of social behaviour. *Journal of Theoretical Biology*, 7:1–52.

Hamilton, W. D. (1975). *Innate social aptitudes of man: An approach from evolutionary genetics*, pages 133–157. Malaby Press, London.

Harms, W. F. (2006). What is information? Three concepts. *Biological Theory*, 1:230–242.

Harnad, S. (1990). The symbol grounding problem. *Physica D*, 42:335–346.

Hasson, O. (1994). Cheating signals. *Journal of Theoretical Biology*, 167:223–238.

Hauser, M. D. (1992). Costs of deception: Cheaters are punished in rhesus monkeys (Macaca mulatta). *PNAS*, 89:12137–9.

Holland, J. H. (1975). *Adaptation in Natural and Artificial Systems: An Introductory Analysis With Applications to Biology, Control, and Artificial Intelligence*. University of Michigan Press, Ann Arbor, MI.

Holland, O. and McFarland, D. (2001). *Artificial Ethology*. Oxford University Press, New York, NY.

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

Houle, D. (1992). Comparing evolvability and variability of quantitative traits. *Genetics*, 130:195–204.

Huston, M., DeAngelis, D., and Post, W. (1988). New computer models unify ecological theory. *BioScience*, 38:682 – 691.

Jacob, F. (1977). Evolution and tinkering. *Science*, 196:1161–1166.

Jacob, F. (1981). *Le Jeu des Possibles*. Fayard, Paris.

Jeanson, R., Rivault, C., Deneubourg, J.-L., Blanco, S., Fournier, R., Jost, C., and Theraulaz, G. (2005). Self-organized aggregation in cockroaches. *Animal Behaviour*, 69:169–180.

Johnstone, R. A. (1994). Honest signalling, perceptual error and the evolution of 'all-or-nothing' displays. *Proc R Soc Lond Biol*, 256:169–175.

Johnstone, R. A. (1998a). Efficacy and honesty in communication between relatives. *The American Naturalist*, 152:45–58.

Johnstone, R. A. (1998b). *Game theory and communication*, pages 94–117. Oxford University Press, Oxford.

Johnstone, R. A. and Grafen, A. (1992). The continuous Sir Philip Sidney game: a simple model of biological signalling. *Journal of Theoretical Biology*, 156:215–234.

Judson, O. P. (1994). The rise of the individual-based model in ecology. *Trends in Ecology & Evolution*, 9:9–14.

Keller, L., editor (1999). *Levels of Selection in Evolution*. Princeton University Press, Princeton, NJ.

Keller, L. and Surette, M. G. (2006). Communication in bacteria: An ecological and evolutionary perspective. *Nature Reviews Microbiology*, 4:249–258.

Kirby, S. (2000). *Syntax without natural selection: How compositionality emerges from vocabulary in a population of learners*, pages 303–323. Cambridge University Press, Cambridge, UK.

Kirby, S. (2006). Natural language from Artificial Life. *Artificial Life*, 8:185–215.

Kirby, S., Cornish, H., and Smith, K. (2008). Cumulative cultural evolution in the laboratory: An experimental approach to the origins of structure in human language. *PNAS*, 105:10681–10686.

Krieger, M. J. B. and Billeter, J.-B. (2000). The call of duty: Self-organised task allocation in a population of up to twelve mobile robots. *Robotics and Autonomous Systems*, 30:65–84.

Krieger, M. J. B., 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*, 30:85–101.

Kubinyi, E., Miklósi, A., Kaplan, F., Gácsi, M., Topál, J., and Csányi, V. (2004). Social behaviour of dogs encountering AIBO, an animal-like robot in a neutral and in a feeding situation. *Behavioural Processes*, 65:231–239.

Laidre, M. E. (2006). Manipulation without mind-reading: Information suppression and leakage during food discovery by mandrills (Mandrillus sphinx). *Behaviour*, 143:365–392.

Lailvaux, S. P., Reaney, L. T., and Backwell, P. R. Y. (2008). Dishonest signalling of fighting ability and multiple performance traits in the fiddler crab Uca mjoebergi. *Functional Ecology*, 23:359–366.

Lande, R. (1976). The maintenance of genetic variability by mutation in a polygenic character with linked loci. *Genetics Research*, 26:221–235.

Lehmann, L. and Keller, L. (2006). The evolution of cooperation and altruism: A general framework and a classification of models. *Journal of Evolutionary Biology*, 19:1365–1379.

Lenoir, A. (1982). An informational analysis of antennal communication in the ant myrmica Rubra L. during trophallaxis. *Behavioural Processes*, 7:27–35.

Lenormand, T., Roze, D., and Rousset, F. (2009). Stochasticity in evolution. *Trends in Ecology & Evolution*, 24:157–165.

Lotka, A. J. (1925). *Elements of Physical Biology*. Williams and Wilkins Company.

Loula, A., Gudwin, R., El-Hani, C. N. A., and Queiroz, J. A. (2008). Emergence of self-organized symbol-based communication in artificial creatures. *Cognitive Systems Research*.

Lynch, M. and Hill, W. G. (1986). Phenotypic evolution by neutral mutation. *Evolution*, 40:915–935.

MacArthur, R. H. and Pianka, E. R. (1966). On optimal use of a patchy environment. *American Naturalist*, 100:603.

MacLennan, B. and Burghardt, G. M. (1993). Synthetic ethology and the evolution of cooperative communication. *Adaptive Behavior*, 2:161–187.

Marocco, D., Cangelosi, A., and Nolfi, S. (2003). The emergence of communication in evolutionary robots. *Philos Trans R Soc London A*, 361:2397–2421.

Marocco, D. and Nolfi, S. (2006). Self-organization of communication in evolving robots. In Rocha, L. M., Yaeger, L., Bedau, M., Floreano, D., Goldstone, R., and Vespignani, A., editors, *Proceedings of the Tenth International Conference on Artificial Life*, pages 199–205, Cambridge, MA. MIT Press.

Martinoli, A. and Easton, K. (2002). Modeling swarm robotic systems. In Siciliano, B. and Dario, P., editors, *Proceedings of the Eighth International Symposium on Experimental Robotics*, pages 297–306, Berlin. Springer.

Martins, E. P., Ord, T. J., and Davenport, S. W. (2005). Combining motions into complex displays: Playbacks with a robotic lizard. *Behavioral Ecology and Sociobiology*, 58:351–360.

Mataric, M. J. (1993). Kin recognition, similarity, and group behavior. In *Proceedings of the Fifteenth Annual Conference of the Cognitive Science Society*, pages 705–710, Philadelphia, PA. Lawrence Erlbaum Associates.

Matsuura, K. (2006). Termite-egg mimicry by a sclerotium-forming fungus. *Proc R Soc Lond Biol*, 273:1203–1209.

May, C. J., Schank, J. C., Joshi, S., Tran, J., Taylor, R. J., and Scott, I.-E. (2006). Rat pups and random robots generate similar self-organized and intentional behavior. *Complexity*, 12:53–66.

Maynard Smith, J. (1991). Honest signalling: The Sir Philip Sidney game. *Animal Behavior*, 42:1034–1035.

Maynard Smith, J. and Harper, D. (2003). *Animal Signals*. Oxford University Press, Oxford.

Maynard Smith, J. and Price, G. R. (1973). The logic of animal conflict. *Nature*, 246:15–18.

Mayr, E. (1961). Cause and effect in biology: Kinds of causes, predictability, and teleology are viewed by a practicing biologist. *Science*, 134:1501–1506.

McFarland, D. (1994). Towards robot cooperation. In Cliff, D., Husbands, P., Meyer, J.-A., and Wilson, S. W., editors, *From Animals to Animats 3: Proceedings of the Third International Conference on Simulation of Adaptive Behavior*, pages 440–444, Cambridge, MA. MIT Press.

McNamara, J. M., Barta, Z., Fromhage, L., and Houston, A. I. (2008). The coevolution of choosiness and cooperation. *Nature*, 451:189–192.

McNamara, J. M., Barta, Z., and Hou, A. I. (2004). Variation in behaviour promotes cooperation in the Prisoner's Dilemma game. *Nature*, 428:745–748.

Melhuish, C., Sendova-Franks, A. B., Scholes, S., Horsfield, I., and Welsby, F. (2006). Ant-inspired sorting by robots: The importance of initial clustering. *J R Soc Interface*, 3:235–242.

Michelsen, A., Andersen, B. B., Storm, J., Kirchner, W. H., and Lindauer, M. (1992). How honeybees perceive communication dances, studied by means of a mechanical model. *Behavioral Ecology and Sociobiology*, 30:143–150.

Mirolli, M. and Parisi, D. (2005). How can we explain the emergence of a language that benefits the hearer but not the speaker? *Connection Science*, 17:307–324.

Mitri, S., Floreano, D., and Keller, L. (2009). The evolution of information suppression in communicating robots with conflicting interests. *PNAS*, 106:15786–15790.

Møller, A. (1988). False alarm calls as a means of resource usurpation in the great tit Parus major. *Ethology*, 79:25–30.

Mondada, F., Pettinaro, G. C., Guignard, A., Kwee, I., Floreano, D., Deneubourg, J.-L., Nolfi, S., Gambardella, L. M., and Dorigo, M. (2004). Swarm-bot: A new distributed robotic concept. *Autonomous Robots*, 17:193–221.

Munn, C. A. (1986). Birds that "cry wolf". *Nature*, 319:143–145.

Murray, J. D. (2007). *Mathematical Biology: I. An Introduction*. Springer, Berlin.

Noble, J. and Cliff, D. (1996). On simulating the evolution of communication. In Maes, P., Mataric, M., Meyer, J.-A., Pollack, J., and Wilson, S. W., editors, *Proceedings of the Fourth International Conference on Simulation of Adaptive Behavior*, Cambridge, MA. MIT Press.

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

Nolfi, S. and Floreano, D. (1998). Co-evolving predator and prey robots: Do "arm races" arise in artificial evolution. *Artificial Life*, 4:311–335.

Nolfi, S. and Floreano, D. (2001). *Evolutionary Robotics. The Biology, Intelligence, and Technologyof Self-organizing Machines*. MIT Press, Cambridge, MA.

Nowak, M. A. and Sigmund, K. (2005). Evolution of indirect reciprocity. *Nature*, 437:1291–1298.

Oliphant, M. (1996). The dilemma of Saussurean communication. *Biosystems*, 37:31–38.

Ord, T. J. and Stamps, J. A. (2008). Alert signals enhance animal communication in "noisy" environments. *PNAS*, 105:18830–18835.

Orr, H. A. (2009). Fundamental concepts in genetics: Fitness and its role in evolutionary genetics. *Nature Reviews Genetics*, 10:531–539.

Partan, S. R., Larco, C. P., and Owens, M. J. (2009). Wild tree squirrels respond with multisensory enhancement to conspecific robot alarm behaviour. *Animal Behaviour*, 77:1127–1135.

Patricelli, G. L., Uy, J. A. C., Walsh, G., and Borgia, G. (2002). Male displays adjusted to female's response. *Nature*, 415:279–80.

Pfeifer, R., Lungarella, M., and Iida, F. (2007). Self-organization, embodiment, and biologically inspired robotics. *Science*, 318:1088–1093.

Poulton, E. B. (1890). *The Colours of Animals: Their Meaning and Use, Especially Considered in the Case of Insects.* Appleton, New York.

Quinn, M. (2001). Evolving communciation without dedicated communication channels. In Sosík, P. and Kelemen, J., editors, *The Sixth European Conference on Advances in Artificial Life*, pages 357–366, Berlin. Springer.

Rainey, P. B. and Rainey, K. (2003). Evolution of cooperation and conflict in experimental bacterial populations. *Nature*, 425:72–74.

Reaney, L. T., Sims, R. A., Sims, S. W. M., Jennions, M. D., and Backwell, P. R. Y. (2008). Experiments with robots explain synchronized courtship in fiddler crabs. *Current Biology*, 18:R62–R63.

Redondo, T. and Castro, F. (1992). Signalling of nutritional need by magpie nestlings. *Ethology*, 92:193–204.

Ristau, C. A. (1991). Aspects of the cognitive ethology of an injury-feigning bird, the piping plover. In Ristau, C. A., editor, *Cognitive Ethology: The Minds of Other Animals*, pages 91–126. Lawrence Erlbaum, London.

Robinson, G. E., Fernald, R. D., and Clayton, D. F. (2008). Genes and social behavior. *Science*, 322:896–900.

Rowe, C. (1999). Receiver psychology and the evolution of multicomponent signals. *Animal behaviour*, 58:921–931.

Rowell, J. T., Ellner, S. P., and Reeve, H. K. (2006). Why animals lie: How dishonesty and belief can coexist in a signaling system. *The American Naturalist*, 168:E180–E204.

Rubenstein, M., Sai, Y., Chuong, C.-M., and Shen, W.-M. (2009). Regenerative patterning in swarm robots: Mutual benefits of research in robotics and stem cell biology. *The International journal of developmental biology*, 53:869–81.

Rundus, A. S., Owings, D. H., Joshi, S. S., Chinn, E., and Giannini, N. (2007). Ground squirrels use an infrared signal to deter rattlesnake predation. *PNAS*, 104:14372–6.

Ryan, M. J. (2007). Sensory ecology: See me, hear me. *Current Biology*, 17:1019–1021.

Ryan, M. J. and Wilczynski, W. (1991). Evolution of intraspecific variation in the advertisement call of a cricket frog (Acris crepitans, Hylidae). *Biological Journal of the Linnean Society*, 44:249–271.

Schiestl, F. P., Ayasse, M., Paulus, H. F., Löfstedt, C., Hansson, B. S., Ibarra, F., and Francke, W. (1999). Orchid pollination by sexual swindle. *Nature*, 399:421.

Scott-Phillips, T. C. (2008). Defining biological communication. *Journal of Evolutionary Biology*, 21:387–395.

Searcy, W. A. and Nowicki, S. (2005). *The Evolution of Animal Communication: Reliability and Deception in Signaling Systems*. Princeton University Press, Princeton and Oxford.

Shannon, C. E. (1948). A mathematical theory of communication. *The Bell System Technical Journal*, 27:379–423,623–656.

Shuker, D. M. and West, S. A. (2004). Information constraints and the precision of adaptation: Sex ratio manipulation in wasps. *PNAS*, 101:10363–10367.

Steels, L. (1998). The origins of syntax in visually grounded robotic agents. *Artificial Intelligence*, 103:133–156.

Steels, L. (1999). *The Talking Heads Experiment*. VUB Artificial Intelligence Library, Brussels, Belguim.

Steels, L. (2005). The emergence and evolution of linguistic structure: From lexical to grammatical communication systems. *Connection Science*, 17:213–230.

Steels, L. and Vogt, P. (1997). Grounding adaptive language games in robotic agents. In Husbands, P. and Harvey, I., editors, *Fourth European Conference on Artificial Life*, Cambridge, MA. MIT Press.

Stevens, M. and Merilaita, S. (2009). Animal camouflage: Current issues and new perspectives. *Philos Trans R Soc London Biol*, 364:423–427.

Stuart-Fox, D., Moussalli, A., and Whiting, M. J. (2007). Natural selection on social signals: Signal efficacy and the evolution of chameleon display coloration. *The American Naturalist*, 170:916–930.

Sundström, L., Chapuisat, M., and Keller, L. (1996). Conditional manipulation of sex ratios by ant workers: A test of kin selection theory. *Science*, 274:993–995.

Számadó, S. (2000). Cheating as a mixed strategy in a simple model of aggressive communication. *Animal Behaviour*, 59:221–230.

Takanishi, A., Aoki, T., Ito, M., Ohkawa, Y., and Yamaguchi, J. (1998). Interaction between creature and robot: Development of an experiment system for rat and rat robot interaction. *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 3:1975 – 1980.

Taylor, P. D. (1992). Altruism in viscous populations - an inclusive fitness model. *Evolutionary Ecology*, 6:352–356.

Taylor, R. C., Klein, B. A., Stein, J., and Ryan, M. J. (2008). Faux frogs: Multimodal signalling and the value of robotics in animal behaviour. *Animal Behaviour*.

Tibbetts, E. A. and Dale, J. (2004). A socially enforced signal of quality in a paper wasp. *Nature*, 432:218–222.

Tinbergen, N. (1948). Social releasers and the experimental method required for their study. *The Wilson Bulletin*, 60:6 – 51.

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

Varela, F., Thompson, E., and Rosch, E. (1991). *The Embodied Mind: Cognitive Science and Human Experience*. MIT Press, Cambridge, MA.

Vogt, P. (2000). *Lexicon Grounding on Mobile Robots*. PhD thesis.

Volterra, V. (1926). Variazioni e fluttuazioni del numero d'individui in specie animali conviventi. *Mem R Accad Naz dei Lincei Ser VI*, 2:31–113.

Von Frisch, K. (1967). *The Dance Language and Orientation of Bees*. Harvard University Press, Cambridge, MA.

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

Waibel, M. (2007). *Evolution of Cooperation in Artificial Ants*. PhD thesis.

Waibel, M., Keller, L., and Floreano, D. (2009). Genetic team composition and level of selection in the evolution of cooperation. *IEEE Transactions on Evolutionary Computation*, 13:648–660.

Wakano, J. Y., Nowak, M. A., and Hauert, C. (2009). Spatial dynamics of ecological public goods. *PNAS*, 106:7910–4.

Waldner, J.-B. (2008). *Nanocomputers and Swarm Intelligence*. Wiley-ISTE, London.

Wauters, A. and Richard-Yris, M.-A. (2003). Maternal food calling in domestic hens: Influence of feeding context. *Comptes Rendus Biologies*, 326:677–686.

Webb, B. (2001). Can robots make good models of biological behaviour? *Behavioral and Brain Sciences*, 24:1033–1050.

Webb, B. (2009). Animals versus animats: Or why not model the real iguana? *Adaptive Behavior*, 17:269–286.

Werner, G. M. and Dyer, M. G. (1991). Evolution of communication in artificial organisms. In Langton, C., Taylor, C., Farmer, D., and Rasmussen, S., editors, *Artificial life II*, pages 659–687.

West, S. A., Griffin, A. S., and Gardner, A. (2007a). Evolutionary explanations for cooperation. *Current Biology*, 17:R661–R672.

West, S. A., Griffin, A. S., and Gardner, A. (2007b). Social semantics: Altruism, cooperation, mutualism, strong reciprocity and group selection. *Journal of Evolutionary Biology*, 20:415–432.

West, S. A., Herre, E. A., and Sheldon, B. C. (2000). The benefits of allocating sex. *Science*, 290:288–290.

West, S. A., Murray, M. G., Machado, C. A., Griffin, A. S., and Herre, E. A. (2001). Testing Hamilton's rule with competition between relatives. *Nature*, 409:510–3.

West, S. A., Pen, I., and Griffin, A. S. (2002). Cooperation and competition between relatives. *Science*, 296:72–75.

Wheeler, B. C. (2009). Monkeys crying wolf? Tufted capuchin monkeys use anti-predator calls to usurp resources from conspecifics. *Proc R Soc Lond Biol*, 276:3013–3018.

Whiteley, D. A. A., Owen, D. F., and Smith, D. A. S. (1997). Massive polymorphism and natural selection in Donacilla cornea (Poli, 1791) (Bivalvia: Mesodesmatidae). *Biological Journal of the Linnean Society*, 62:475–494.

Wiley, R. H. (1983). *The evolution of communication: Information and manipulation*, volume 2, pages 156–189. Blackwell Scientific Publications, Oxford.

Wilson, E. O. (1962). Chemical communication among workers of the fire ant Solenopsis saevissima (Fr. Smith) 2. An information analysis of the odour trail. *Animal Behaviour*, 10:148–158.

Wilson, E. O. (1971). *The Insect Societies*. Harvard University Press, Cambridge, MA.

Wilson, E. O. (1975). *Sociobiology: The New Synthesis*. The Belknap Press of Harvard University Press, Cambridge, MA.

Wischmann, S., Hülse, M., Knabe, J. F., and Pasemann, F. (2006). Synchronization of internal neural rhythms in multi-robotic systems. *Adaptive Behavior*, 14:117–127.

Wischmann, S. and Pasemann, F. (2006). The emergence of communication by evolving dynamical systems. In Nolfi, S., Baldassarre, G., Calabretta, R., Hallam, J. C. T., Marocco, D., Meyer, J. A., Miglino, O., and Parisi, D., editors, *Ninth International Conference on Simulation of Adaptive Behavior*, pages 777–788, Berlin. Springer.

Wollerman, L. and Wiley, R. H. (2002). Background noise from a natural chorus alters female discrimination of male calls in a Neotropical frog. *Animal Behaviour*, 63:15–22.

Wright, S. (1932). The roles of mutation, inbreeding, crossbreeding and selection in evolution. In Jones, D. F., editor, *Proceedings of the VI International Congress of Genetics*, pages 356–366.

Zahavi, A. (1975). Mate selection – a selection for a handicap. *Journal of Theoretical Biology*, 53:205–14.

Zahavi, A. (1997). *The Handicap Principle: A Missing Piece of Darwin's Puzzle*. Oxford University Press, New York.

Curriculum Vitæ

I was born on the 19th of March, 1981 in Cairo, Egypt, where I lived until the age of 23. For 12 years, I attended the Deutsche Schule der Borromäerinnen, a German school in Cairo from which I graduated at the top of my class in 1999. In the same year, I began my studies in computer science (major), mathematics (minor) and art (minor) at the American University in Cairo, which I completed in 2003 with highest honors. My graduation project, entitled "Applying Particle Swarm Optimization to the Egyptian Board Game Seega" resulted in two conference publications. At my graduation ceremony, I was awarded the Ahmed H. Zewail Prize for Excellence in Sciences and Humanities.

In 2004, I spent 4 months as an intern at the Fraunhofer Institut für Autonome Intelligente Systeme in Sankt Augustin in Germany, where I worked on building computer vision algorithms for an autonomous robot, which led to three publications. Following this brief internship, I began a Master's course in 2004 on "Cognitive Science and Natural Language" at the informatics department of the University of Edinburgh, where I specialized in Natural Language, Language Engineering and Intelligent Robotics. Both my tuition fees and living expenses during this time were covered by a Shell Centenary Scholarship, a prestigious award given to two applicants from Egypt per year. I completed the program in 2005 with highest honors, and a Master's thesis entitled "The Co-evolution of Language and Behaviour in Autonomous Robots".

Since October, 2005, I have been a Ph. D. candidate at the Ecole Polytechnique Fédérale de Lausanne, where I have been carrying out the research presented in this thesis under the supervision of Prof. Dario Floreano at the Laboratory of Intelligent Systems and Prof. Laurent Keller at the Department of Ecology and Evolution at the University of Lausanne. My research has led to 2 publications in prestigious journals, 2 articles that are currently in preparation, and has attracted substantial media coverage. Through my research, I have acquired a range of academic skills, including data analysis, academic writing, which I hope are reflected in this thesis, and giving academic talks. I was awarded the fourth prize for my presentation at the Annual European Meeting for PhD Students in Evolutionary Biology. In addition to these research activities, I have supervised a number of student projects, and have acted as a teaching assistant and co-lecturer in Prof. Dario Floreano's course on Bio-inspired Artificial Intelligence.

Together with two colleagues, I have organized a workshop on the evolution of social behavior, which was held in Lisbon, Portugal, in September, 2007, and attracted around 30 participants. This workshop led to the publication of a special issue, published in 2008 in the journal Biological Theory, which I co-edited.

In 2008, I was approached by the Museum of Zoology in Lausanne to set up an exhibit on evolutionary robotics as part of an exhibition for the year of Darwin. With the help of two colleagues, we have completed the exhibit, which is due to start at the end of October, 2009 and will run for 12 months.

For the future, I hope to continue my research on studying the evolution of social behavior, for which I would like to use a combination of mathematical models, IBMs and living organisms. In particular, I am interested in changing my study organisms from robots to microbes.