

# EVOLUTION OF COOPERATION IN ARTIFICIAL ANTS

THÈSE N° 3943 (2007)

PRÉSENTÉE LE 9 NOVEMBRE 2007

À LA FACULTÉ DES SCIENCES ET TECHNIQUES DE L'INGÉNIEUR

Laboratoire de Systèmes Intelligents

SECTION DE MICROTECHNIQUE

ÉCOLE POLYTECHNIQUE FÉDÉRALE DE LAUSANNE

POUR L'OBTENTION DU GRADE DE DOCTEUR ÈS SCIENCES

PAR

**Markus WAIBEL**

Dipl.-Ing. in Technical Physics, Vienna University of Technology, Autriche  
et de nationalité autrichienne

acceptée sur proposition du jury:

Prof. H. Bleuler, président du jury  
Prof. D. Floreano, Prof. L. Keller, directeurs de thèse  
Dr M. Chapuisat, rapporteur  
Prof. O. Holland, rapporteur  
Dr F. Mondada, rapporteur



ÉCOLE POLYTECHNIQUE  
FÉDÉRALE DE LAUSANNE

Lausanne, EPFL

2007



---

# Acknowledgements

---

[...] discoveries and improvements invariably  
involve the cooperation of many minds.

Alexander Graham Bell (1847-1922)

Many people have contributed to the work presented in this thesis. First I would like to thank my two supervisors, Prof. Dario Floreano and Prof. Laurent Keller, for making me part of this project and thus giving me the opportunity to dive into this amazing research field. I thank Laurent for his radiant motivation, endless patience and for sharing his clear mind, and Dario for his constant guidance, his trust and for being an inspiration. In addition, I would like to thank Dario for assembling the amazing group of people in the previous Autonomous Systems Laboratory and the current Laboratory of Intelligent Systems.

I would like to thank all my colleagues at the Laboratory of Intelligent Systems for providing the great, stimulating atmosphere that has made these years so enjoyable. In particular, I would like to thank Sara Mitri for her calm, persistent thinking and our endless discussions, Peter Dürri for his never-failing, contagious enthusiasm and his sharp wits, and Danesh Tarapore for the great time we shared in the office and for making Alice work. I would also like to thank Laurent's research group at the Department of Ecology and Evolution of the University of Lausanne, who have helped with their advice and criticism on countless occasions.

A number of people have enriched my academic experience and have helped me feel at home in Lausanne. I am very grateful to the Pavillon-Jaune-team, who made me feel at home from day one: Daniel Roggen, Stéphane Magnenat,

Francesco Mondada, Michael Bonani, Jean-Christophe Zufferey, Dominique Etienne, Jesper Blynel, Diego Federici and Claudio Mattiussi. My special gratitude goes to all movie-night participants, who are too numerous to name, for the many unforgettable memories.

In writing this thesis I have benefitted from comments and helpful discussion provided by a large number of people. I am particularly indebted to Peter Dürri and Sara Mitri for their truly tireless efforts in offering advice. I would also like to thank Danesh Tarapore, John Wang, Morgan Percy, Martijn Bosch, Claudio Mattiussi, Simon Harding, Daniel Marbach, Guy Theraulaz, Rob Hammon, Jean-Louis Deneubourg and the members of the thesis committee, Hannes Bleuler, Michel Chapuisat, Owen Holland and Francesco Mondada. In addition, I am indebted to Julien Chassot, Gilles Caprari, Stéphane Magnenat, Antoine Beyeler, Andres Perez-Urbe, Cyrille Dunant, Gintautas Narvydas and Danesh Tarapore for contributing to this project.

I would like to thank Dominique Etienne and Anouk Hein - nothing would have worked without you. I would also like to express my gratitude to the EPFL, the University of Lausanne, and the Swiss National Science Foundation for their support and for funding my research.

I also thank all my friends for sharing all aspects of life, my parents, Maria and Karl for their guidance, trust and patience, and my brother Roman and sister Sophie for their support.

Most importantly, I would like to thank Christine for all the things that can't be put into words.

*Lausanne, November 2007*

---

# Abstract

---

The evolution of cooperation is a fundamental and enduring puzzle in biology and the social sciences. Hundreds of theoretical models have been proposed, but empirical research has been hindered by the generation time of social organisms and by the difficulties of quantifying costs and benefits of cooperation. The significant increase in computational power in the last decade has made artificial evolution of simple social robots a promising alternative.

This thesis is concerned with the artificial evolution of groups of cooperating robots. It argues that artificial evolution of robotic agents is a powerful tool to address open questions in evolutionary biology, and shows how insights gained from the study of artificial and biological multi-agent systems can be mutually beneficial for both biology and robotics. The work presented in this thesis contributes to biology by showing how artificial evolution can be used to quantify key factors in the evolution of cooperation in biological systems and by providing an empirical test of a central part of biological theory. In addition, it reveals the importance of the genetic architecture for the evolution of efficient cooperation in groups of organisms. The work also contributes to robotics by identifying three different classes of multi-robot tasks depending on the amount of cooperation required between team members and by suggesting guidelines for the evolution of efficient robot teams. Furthermore it shows how simulations can be used to successfully evolve controllers for physical robot teams.

Keywords: Artificial evolution; multi-agent systems; social insects; evolutionary robotics; team composition; task allocation; division of labor; fitness allocation; cooperation; altruism.



---

# Zusammenfassung

---

Die Evolution von kooperativem Verhalten ist ein grundlegendes Problem der Biologie und der Sozialwissenschaften. Die empirische Überprüfung der zahlreichen theoretischen Modelle wird durch die Generationszeit sozialer Organismen und die Schwierigkeit Kosten und Nutzen sozialen Verhaltens zu quantifizieren erschwert. Der dramatische Anstieg der Rechnerleistung in den letzten zehn Jahren macht die künstliche Evolution von einfachen sozialen Robotern zur vielversprechenden Alternative.

Diese Dissertation beschäftigt sich mit der künstlichen Evolution von Gruppen kooperierender Roboter. Sie schlägt vor die künstliche Evolution von Robotern als mächtiges Werkzeug zur Bearbeitung offener Fragen der Evolutionsbiologie zu verwenden und zeigt, dass Erkenntnisse aus der Erforschung natürlicher und künstlicher Multi-Agenten-Systeme gleichzeitig einen Nutzen für die Biologie und die Robotik bringen können. Die vorliegende Arbeit leistet einen Beitrag zur Biologie, indem sie zeigt wie künstliche Evolution benutzt werden kann um Schlüsselfaktoren der Evolution sozialen Verhaltens in biologischen Systemen zu quantifizieren und liefert einen empirischen Test für eine zentrale Theorie der Evolutionsbiologie. Zusätzlich verdeutlicht sie den wesentlichen Einfluss der genetischen Architektur auf die Evolution effizienter kooperativer Gruppen von Organismen.

Die Arbeit leistet auch einen Beitrag zur Robotik, indem sie drei verschiedene Klassen von Multi-Roboter-Problemen identifiziert und diese anhand des unterschiedlichen Masses an Zusammenarbeit, das zur Lösung des Problems notwendig ist, unterscheidet. Ausserdem schlägt sie Richtlinien zur Evolution effizienter

Roboterteams vor und zeigt wie Computersimulationen genutzt werden können, um erfolgreiche Steueralgorithmen für reale Roboterteams zu erhalten.

Schlüsselwörter: Künstliche Evolution; Multi-Agenten-Systeme; Soziale Insekten; Robotik; Arbeitsverteilung; Arbeitsteilung; Kooperation; Altruismus.



---

# Contents

---

<b>Acknowledgements</b>	<b>i</b>
<b>Abstract</b>	<b>iii</b>
<b>Zusammenfassung</b>	<b>v</b>
<b>Contents</b>	<b>vii</b>
<b>1 Introduction</b>	<b>1</b>
1.1 Natural Selection and the Evolution of Cooperation . . . . .	1
1.2 Factors Influencing the Evolution of Cooperation . . . . .	2
1.2.1 Direct selection . . . . .	2
1.2.2 Indirect selection . . . . .	3
1.2.3 Group selection . . . . .	5
1.3 Empirical Research . . . . .	6
1.4 Modeling Cooperation . . . . .	7
1.4.1 Evolutionary Robotics . . . . .	9
1.5 Structure of the Thesis . . . . .	10
<b>2 A Quantitative Test of Hamilton's Rule</b>	<b>13</b>
2.1 State of the Art . . . . .	13
2.2 Materials and Methods . . . . .	15
2.2.1 Experimental setup . . . . .	15
2.2.2 Costs and benefits of altruism . . . . .	17
2.2.3 Relatedness and artificial evolution . . . . .	18
2.3 Results . . . . .	20

2.3.1	Groups with high relatedness . . . . .	21
2.3.2	Groups with intermediate relatedness . . . . .	23
2.3.3	Groups with low relatedness . . . . .	23
2.4	Discussion . . . . .	23
2.5	Conclusion . . . . .	25
<b>3</b>	<b>Division of Labor and Colony Efficiency in Social Insects</b>	<b>27</b>
3.1	Division of Labor and Colony Efficiency in Social Insects . . . . .	28
3.2	Materials and Methods . . . . .	30
3.2.1	Colony tasks . . . . .	30
3.2.2	Genetic architecture . . . . .	31
3.2.3	Environmental and internal perturbations . . . . .	32
3.2.4	Colonies and selection algorithm . . . . .	33
3.3	Results . . . . .	34
3.3.1	Colonies of unrelated individuals . . . . .	34
3.3.2	Colonies of highly related individuals . . . . .	37
3.3.3	Colonies with intermediate relatedness . . . . .	39
3.4	Discussion . . . . .	40
3.5	Conclusion . . . . .	44
<b>4</b>	<b>Factors in the Evolution of Multi-Agent Systems</b>	<b>47</b>
4.1	State of the Art . . . . .	48
4.2	Evolutionary Conditions . . . . .	52
4.3	Experimental Method . . . . .	56
4.3.1	Scenario . . . . .	56
4.3.2	Control and genetic architecture . . . . .	58
4.3.3	Collective tasks . . . . .	58
4.3.4	Evolutionary experiments . . . . .	59
4.4	Results . . . . .	60
4.4.1	Task 1 - Individual foraging . . . . .	60
4.4.2	Task 2 - Cooperative foraging . . . . .	60
4.4.3	Task 3 - Altruistic cooperative foraging . . . . .	60
4.5	Discussion . . . . .	63
4.5.1	Task 1 - Individual foraging . . . . .	63
4.5.2	Task 2 - Cooperative foraging . . . . .	65
4.5.3	Task 3 - Altruistic cooperative foraging . . . . .	69

---

4.6	Conclusion . . . . .	72
<b>5</b>	<b>Artificial Ants Scale Reality Gap: Cooperative Robot Foraging</b>	<b>75</b>
5.1	Control strategies for multi-robot systems . . . . .	75
5.2	Related work . . . . .	77
5.3	Experimental Setup . . . . .	79
5.3.1	Software platform . . . . .	80
5.3.2	Hardware platform . . . . .	80
5.3.3	Experimental method . . . . .	80
5.4	Results . . . . .	83
5.4.1	Performance of teams with evolved controllers . . . . .	83
5.4.2	Performance of teams with random controllers . . . . .	84
5.4.3	Performance of random controllers vs. evolved controllers.	84
5.5	Discussion . . . . .	87
5.6	Conclusion . . . . .	88
<b>6</b>	<b>Concluding Remarks</b>	<b>91</b>
6.1	Main Achievements . . . . .	91
6.2	Future work . . . . .	93
<b>A</b>	<b>Genetic Relatedness in Groups of Clones</b>	<b>95</b>
A.1	Relatedness in Groups Composed of Different Proportions of Clones	95
A.2	Genetic Relatedness with Competition Between Relatives . . . . .	96
A.3	Costs and Benefits of Cooperation . . . . .	97
<b>B</b>	<b>Supplementary Data for Chapter 2</b>	<b>99</b>
B.1	Evolution of Fitness and the Level of Altruism . . . . .	99
B.2	Statistics for Fitness and the Level of Altruism . . . . .	101
<b>C</b>	<b>The Effects of Colony Size and Reallocation on Colony Efficiency</b>	<b>103</b>
C.1	Colony Fitness for Large Colony Sizes . . . . .	104
C.2	Colony Fitness for Varying Proportions of Individuals . . . . .	105
<b>D</b>	<b>Supplementary Data for Chapter 4</b>	<b>107</b>
D.1	Performance in Task 1 - Individual Foraging . . . . .	108
D.2	Performance in Task 2 - Individual Foraging . . . . .	109
D.3	Performance in Task 3 - Altruistic Cooperative Foraging . . . . .	110

---

D.4	Corrected Proportion of Collected Tokens in Task 3 . . . . .	111
D.5	Corrected Proportion in Task 3, no Credit Assignment Problems . .	112
<b>E</b>	<b>Robot simulation</b>	<b>113</b>
E.1	Sensors . . . . .	114
E.1.1	Characterization . . . . .	114
E.1.2	Modeling . . . . .	115
E.2	Actuators . . . . .	117
E.2.1	Characterisation . . . . .	117
E.2.2	Modeling . . . . .	117
E.3	Interaction . . . . .	118
E.3.1	Characterization . . . . .	119
E.3.2	Modeling . . . . .	119
E.4	Verification . . . . .	119
E.4.1	Sensors and Motors . . . . .	120
E.4.2	Collisions . . . . .	120
E.4.3	Additional tests . . . . .	121
E.5	Conclusion . . . . .	121
	<b>Glossary</b>	<b>123</b>
	<b>Bibliography</b>	<b>130</b>
	<b>Curriculum vitæ</b>	<b>155</b>

---

# List of Figures

---

2.1	Experimental setup for altruistic foraging. . . . .	16
2.2	Robot sensors and control architecture. . . . .	18
2.3	Mean level of altruism $\bar{A}$ after 500 generations of artificial evolution. . . . .	22
3.1	Sub-linear foraging return functions. . . . .	31
3.2	Fitness for colonies of unrelated individuals, no perturbations. . . . .	35
3.3	Fitness for colonies of unrelated individuals with perturbations. . . . .	36
3.4	Fitness for colonies of highly related individuals, no perturbations. . . . .	38
3.5	Fitness for colonies of highly related individuals with perturbations. . . . .	39
3.6	Fitness for colonies of intermediate relatedness. . . . .	41
4.1	Approaches to the evolution of agent teams. . . . .	49
4.2	The four evolutionary conditions. . . . .	53
4.3	Experimental setup. . . . .	56
4.4	Micro-robot schematics and neural network architecture. . . . .	57
4.5	Fitness for task 1 - Individual foraging. . . . .	61
4.6	Fitness for task 2 - Cooperative foraging. . . . .	62
4.7	Fitness for task 3 - Altruistic cooperative foraging. . . . .	62
4.8	Proportions of collected tokens in task 3. . . . .	64
4.9	Fitness for task 1 without disparities and without CAP. . . . .	66
4.10	Fitness for task 2 without disparities and without CAP. . . . .	68
4.11	Fitness for task 3 without disparities and without CAP. . . . .	70
4.12	Effects of credit assignment problems on fitness for task 3. . . . .	71
5.1	Micro-robots in the hardware setup. . . . .	81
5.2	Evolution of best and mean team fitness. . . . .	83

---

5.3	Mean performance for evolved controllers. . . . .	85
5.4	Mean performance for random controllers. . . . .	86
B.1	Evolution of mean group fitness and mean level of altruism, part 1	99
B.2	Evolution of mean group fitness and mean level of altruism, part 2	100
C.1	Colony fitness for large colony sizes. . . . .	104
C.2	Colony fitness for varying proportions of reallocated individuals .	105
D.1	Fitness for task 1 without disparities and with CAP. . . . .	108
D.2	Fitness for task 2 without disparities and with CAP. . . . .	109
D.3	Fitness for task 3 without disparities and with CAP. . . . .	110
D.4	Corrected proportion of collected token with CAP for task 3. . . . .	111
D.5	Corrected proportion of collected tokens without CAP for task 3. .	112
E.1	Schematics of the micro-robot in simulation . . . . .	114
E.2	The linear sensor ray model . . . . .	116

---

## List of Tables

---

2.1	Mean levels of altruism $\bar{A}$ for different cost to benefit ratios $\frac{c}{b}$ . . . . .	21
3.1	Fitness average for different reallocation probabilities $P_r$ . . . . .	37
4.1	Suggested guidelines for the evolution of agent teams. . . . .	73
A.1	Corrected benefits $b$ and cost over benefit ratios $\frac{c}{b}$ for $\frac{c}{b} = r$ . . . . .	98
B.1	Mean group fitness $\bar{F}$ and mean level of altruism $\bar{A}$ . . . . .	101





---

# 1

---

## Introduction

Remember that we are picturing the animal as a robot survival machine with a pre-programmed computer controlling the muscles.

Richard Dawkins (*The Selfish Gene*, 1976)

In this chapter I will introduce the underlying questions that have motivated this thesis. I will present the background and concepts for our current understanding of the evolution of cooperation and discuss the key factors. In particular, I will highlight the challenges in studying the evolution of cooperation in biological groups such as multi-cellular organisms or social insect colonies, and in groups of artificial agents such as teams of software agents or robots. A more detailed introduction to each of these subjects is included in the pertaining chapters. A glossary at the end of this thesis serves as a quick reference and provides clarification for terms whose definition differs between research fields.

### **1.1 Natural Selection and the Evolution of Cooperation**

The evolution of cooperation is a fundamental and enduring puzzle in biology and the social sciences. If evolution is based on the survival of the fittest, how can it lead to cooperation? Individuals should, then, behave in ways that increase their own chances of reproduction and strive to reduce those of others. In many cases they do: Animals fight for mating opportunities and plants overshadow

each other in the struggle for light. Yet, cooperation is ubiquitous. It is found across all biological taxa and pervades all levels of biological organization, from genes to cells to organisms to societies. Cooperation has played a central role in the evolution of life: Higher organisms are the result of a succession of ever more complex replicating individuals created by combination of initially independent units (Maynard Smith and Szathmary [146]). Cooperation has acted as a driving force in the evolution of higher organisms by unlocking fitness benefits not available to single units. Many fascinating examples illustrate this process: The independent cells of the cellular slime mold *Dictyostelium* aggregate to form a slug that allows cells to migrate, covering large distances (Kessin [122], Strassmann et al. [211]). Individuals in many social insect species cooperatively build elaborate nests that protect them against adverse environmental conditions and predators (Holldobler and Wilson [104]). Wolves and many other predators hunt in packs and can subdue prey many times the size of an individual animal. Vervet monkeys elicit specific alarm calls to warn group members of different types of predators (Seyfarth et al. [202]).

Rather than the exception, cooperation seems to be a common and fundamental form of interaction between individuals. How is this fact compatible with our understanding of natural selection? A deeper understanding of this question requires a closer look at the factors that influence cooperation.

## 1.2 Factors Influencing the Evolution of Cooperation

Cooperation is a social behavior that increases the reproductive fitness of other individuals. This means that cooperative behaviors are those that increase the expected number of offspring of other individuals. Explanations for the evolution of cooperation fall into two categories, cases that lead to a reproductive fitness increase of the cooperator (“direct selection”) and cases that do not lead to such an increase (“indirect selection”).

### 1.2.1 Direct selection

It is easy to see why natural selection would lead to the evolution of cooperation if it increases the reproduction of cooperators. Several situations can lead to direct fitness benefits for cooperators. The simplest involve group activities where the

benefits of cooperation are shared among group members. Classical examples include cooperative hunting or cooperative nest building.

Although many authors have suggested that situations where cooperators gain fitness benefits have played a key role in driving the evolutionary transitions from the first simple replicators to higher organisms (Corning [56], Keller [117], Michod [150]), for many examples the evolution of cooperation cannot be explained by direct selection alone. One important reason is that cooperating groups are always at risk from exploitation by selfish individuals who benefit even more by not taking part in the group activity, yet reaping its benefits. For example, wolves hanging back during a dangerous group hunt should increase their chances of survival and thus be favored by selection. Cooperation of this form is often unstable and thwarted by selfish individuals (Maynard Smith [145]). This, along with the observation that many examples of cooperation involve social behaviors that do not increase the cooperators' fitness, explains why direct selection has received less attention compared to indirect selection (Hammerstein and Leimar [93]).

### 1.2.2 Indirect selection

It is more difficult to see how cooperation can evolve in the absence of direct fitness benefits to the cooperator. This is because selection of behavior that decreases an individual's fitness seems to run against the very idea of natural selection.

The key insight into this apparent dilemma goes back to J.B.S. Haldane in the 1930s, and was formalized by W. D. Hamilton in two landmark papers in 1964 (Hamilton [89, 90]). It has since been developed into a theory that forms much of the conceptual basis of our current understanding of social evolution and is now known under the name of "kin selection". Haldane's and Hamilton's key insights concern cooperative interactions that are directed towards relatives, i.e., towards individuals that are genetically similar to the cooperator. They realized that a gene for cooperative behavior can not only spread in a population by increasing the fitness of the cooperator, but also by increasing the fitness of a beneficiary carrying the same gene and thus perpetuating its own existence. This is because a gene will spread in a population if it can successfully create copies of itself, irrespective of whether copies are made by increasing the reproduction of its bearer, or by increasing the reproduction of other bearers of the same gene (related indi-

viduals). Cooperative interactions with relatives can thus indirectly increase the fitness of the cooperator by increasing its chances of spreading its genes. The cooperator's fitness gain obtained by helping related individuals is called indirect fitness.

Hamilton argued that interactions with related individuals can result in two ways. First, relatives might be distinguished from other individuals that are encountered. This can happen through transitory patterns of association such as being hatched in the same nest. It can also happen through kin recognition, i.e., when individuals can make a better than average estimate on their genetic similarity with other individuals. A special case that allows such an estimate are cooperative genes that are linked to a specific phenotypic trait (Keller and Ross [120]). Such genes are often called "green-beard" genes, a term coined by Dawkins [58] to describe a gene that will lead to a conspicuous phenotypic feature such as causing the bearer to have a green beard. In case of a perfect linkage between the gene coding for the cooperative behavior and the phenotypic trait, bearers of the gene can reliably direct their behavior towards other bearers of the same gene by recognizing this feature. Green-beard genes are an example for the importance of the underlying genetic architecture for the evolution of social groups. We return to the influence of genetic architecture on the evolution of social behavior in Chapter 3.

The second way to preferential interactions between related individuals suggested by Hamilton is limited dispersal from the natal group. This is because average relatedness between individuals in the same natal group is higher than between individuals in the whole population, and therefore non-dispersing individuals will interact more with kin. It has been argued that limited dispersal has played a key role in the evolution of cooperation in many different scenarios, from cooperative breeding in birds, mammals and insects to decreased virulence in parasites such as malaria. This idea is important because many social organisms are thought to lack the sensory or cognitive abilities necessary for kin recognition. However, recent research has warned that the importance of limited dispersal has likely been overestimated (see e.g., Queller [180], West et al. [234]). The reason is that limited dispersal can also lead to increased competition between related individuals, which counters or even negates the effects of increased relatedness. One way to see how this can happen is to think of an isolated group (i.e., no dispersal) of related individuals sustained by a limited

amount of resources. While a cooperative action can help a related neighbor produce more offspring, this increase in offspring must come at the cost of another related individual's offspring. In this case, the cooperative social behavior cannot lead to a net increase in the number of offspring and the competition between related individuals exactly cancels the reproductive benefit of cooperation.

This example reveals a conceptual difficulty of Hamilton's theory that has led to deep confusion and hindered progress in our understanding of the evolution of cooperation: To understand the spread of a cooperative gene by indirect selection, it is not enough to consider its presence in other members of the social group, but also whether it is present in competitors. A complete theory of kin selection needs to account for genetic similarity in a social group with respect to the genetic similarity to competitors. The effects of competition between relatives will be discussed in detail in Chapter 2. A brief outline of its mathematical formulation for a simple case can be found in Appendix A.

### 1.2.3 Group selection

To conclude the brief review of the theoretical foundation of today's understanding of the evolution of cooperation, we must discuss a related theoretical framework which goes by the name "group selection" or "multi-level selection".

Group selection goes back to an idea first proposed in the 1960s by Wynne-Edwards (Wynne-Edwards [243]), who argued that cooperation between individuals could evolve by selection acting between populations of individuals. According to his formulation, populations consisting of selfish individuals might overexploit their available resources and thus go extinct. The only populations to survive would be populations of individuals who did not necessarily have traits that were best for their own survival and reproduction, but which were good for the survival and reproduction of the population. While Wynne-Edwards' basic reasoning was correct, it soon became clear that this type of selection would only work under very restrictive conditions and did not have much practical importance. As numerous authors, including himself, pointed out in the 1980s this is because selection at the population level is much slower than selection at the individual level (Grafen [81]).

Wynne-Edwards' original idea was replaced by a different framework for competition between social groups called multi-level selection, which is centered around the idea that selection can be partitioned into different components acting

between groups at different levels. For example, selection may simultaneously operate at the level of genes, cells and organisms. We return to the importance of the level of selection for the evolution of cooperation in Chapter 4.

Numerous authors have pointed out that Hamilton's theory of kin selection and multi-level selection theory do not point to fundamentally different processes (e.g., Bourke and Franks [37], Keller [118], Lehmann et al. [126]). This is linked to the fact that a complete theory of kin selection must incorporate both the effects of genetic relatedness and of selection acting at all levels. Both theories have led to identical results in all considered cases, and modern versions of both theories are mathematically equivalent (Bourke and Franks [37], Lehmann and Keller [125]).

However, multi-level selection theory and its predecessors have created a substantial amount of controversy and semantic confusion (West et al. [232], Wilson [237]). Its history has been likened to a minefield (Bourke and Franks [37, p. 40]) and issues surrounding its correct application continue to cause misunderstandings and confusion to this day (compare, e.g., Alonso and Schuck-Paim [5], Wilson [240], Wilson and Hölldobler [241] with Foster et al. [75], or Traulsen and Nowak [218] with Lehmann et al. [126] for current controversies). Part of the reason may be found in the lack of underpinning empirical research.

### 1.3 Empirical Research

While hundreds of theoretical models have been proposed (for recent reviews see e.g., Lehmann and Keller [125], Sachs et al. [198]), it has proven very difficult to provide empirical data needed to match the theoretical models with empirically known examples of the evolution of cooperation. Many practical difficulties in studying evolution are due to the long generation times of social organisms. This makes the use of guided evolution, a process where the evolution of a trait can be followed by breeding organisms, impractical. Ideal empirical tests would include a direct, quantitative measure of all factors that affect the evolution of cooperation. Unfortunately, such tests would require detailed information on the lifetime interactions between individuals in a population as well as measures of costs and benefits associated with cooperation, and it has been pointed out that such empirical tests could be extremely difficult to carry out (West et al. [232]). These difficulties in conducting empirical tests of the evolution of cooperation

in natural systems have led many researchers to try to approximate biological systems using increasingly sophisticated models of cooperating agents (Axelrod [12, 14], Doebeli and Hauert [60]). In Chapter 2 I show how some of these limitations and difficulties can be overcome by using artificial robotic systems to model the evolution of complex social interactions.

## 1.4 Modeling Cooperation

Modeling of biological systems has a long history (Charnov [52], Kauffman [115], May [144], Pyke et al. [179]). Ideally, models capture the key features of a natural phenomenon and lead to new hypotheses and a deeper understanding. However, in many cases the key features to be included in a model need to be estimated by its designer (Husbands et al. [108], Jakobi et al. [112]). This is problematic because if a model does not capture the essential parameters it can lead to wrong results. It is important to note that models are always simplified, abstracted versions of reality, but considering salient features in isolation may fail to capture key properties of biological systems. In spite of these shortcomings, models have been essential in understanding many biological processes (Flake [69], Webb [228]).

Current modeling approaches to explaining the evolution of cooperation range in complexity from analytic mathematical descriptions to complex computational models (Flake [69]). The simplest models are analytic descriptions, which attempt to find analytical solutions for the conditions needed for the evolution of cooperation (see e.g., Frank [76], Hamilton [91], Queller [180]). While such models excel at describing and capturing essential properties of biological systems, their design requires a deep understanding of the underlying phenomena. In addition, this approach is ill-suited to describe the complete dynamics of system with many parameters (Flake [69]). The effects of individual variability and probabilistic events can be difficult to model analytically and are usually not included in these models.

Computational models can supplement and extend purely analytic models, especially for discrete phenomena (see, e.g., Maynard Smith [145]). Many authors have used simulations of very simple agents engaged in iterative games to study cooperation (see e.g., Axelrod [11], Axelrod et al. [13], Cohen et al. [54], Doebeli and Hauert [60], Hauert and Doebeli [97], Hauert and Stenull [98], Nowak [163], Nowak and Sigmund [164, 165], Riolo et al. [191], Sigmund and Nowak



[203]). In these models agent behavior is often reduced to a binary choice between cooperation and no cooperation. The appeal of such models is illustrated by the prisoner's dilemma, which captures the essential dilemma of the evolution of cooperation in a simple game, where two players (the prisoners) are given a choice to either cooperate or not cooperate (Axelrod [11]). These simple computational models are useful to study situations where the optimal behavior (often called an optimal "strategy") depends on the behavior of other agents. However, they often only consider a single isolated aspect of an agent and have been criticized for their limited predictive power, because it is often unclear how they relate to real world phenomena (Rowell et al. [196], West et al. [233]).

Some of these shortcomings are corrected by more complex agent-based models (Axelrod [12, 14], Axtell [15]). Agent-based modeling refers to simulation models that treat individuals as unique and discrete entities which have at least one property in addition to age that changes during their life cycle, e.g., weight or rank in a social hierarchy. This type of modeling has been used effectively in the field of ecology where individual variability between agents and their complex life-cycles cannot be neglected (see e.g., Grimm [86] for a review). However, in many cases large difficulties in linking simulation results to real world dynamics remain (Grimm [86]).

The significant increase in computational power in the last decade has allowed to further increase the complexity of computational models and to build richer and more complete models for the evolution of cooperation. The research in this thesis extends this work by modeling the evolution of simple artificial social organisms, such as teams of robotic agents in an evolutionary robotics setup. While this approach cannot overcome all shortcomings of current approaches, it has two advantages. A key advantage of modeling robotic agents is that results can be validated in the real world. This guarantees that models are robust against parameter variations and provides an important safeguard against computational artifacts. The latter point is of special importance for complex computational models which may depend on many parameters and specific initial conditions (Webb [227, 228]). A key advantage of using artificial evolution is the reduced number of model features that need to be estimated by a designer (Husbands et al. [108], Jakobi et al. [112]). This is important for phenomena that emerge from complex interactions of many levels of an agent's existence rather than resulting from a small number of salient features (Brooks [39, 40]). For example, it has



been stressed that an organism's individual or social behavior emerges from the interaction of all levels of an organism's existence including its morphology, sensory apparatus, cognitive abilities and genetic architecture (Keijzer [116], Pfeifer and Bongard [175]). Both advantages are combined in the use of evolutionary robotics, which creates emergent agent behaviors by combining robotic agents with artificial evolution.

### 1.4.1 Evolutionary Robotics

Evolutionary robotics is a relatively new technique for the automatic creation of robot controllers (Cliff et al. [53], Harvey et al. [95], Nolfi and Floreano [161]) and for the study of the evolution of adaptive behavior (Floreano et al. [72]). It uses artificial evolution, a process inspired by Darwinian evolution, to create control systems for robots. This is accomplished by encoding robot control systems in artificial genomes. First, an initial population of different, random genomes is created. These genomes are then evaluated by decoding the genomes into control systems, downloading them into a robot and monitoring the robot's performance on various tasks. Once all genomes have been evaluated, the best genomes are selected according to a predefined metric called a "fitness function". Selected genomes are reproduced by creating copies with the addition of some changes introduced by genetic operators such as mutation and/or cross-over. In most implementations the reproduced genomes are used to form the subsequent population, which is again evaluated, selected and reproduced. This process is repeated until a predefined criterion, such as a certain level of performance has been met (Nolfi and Floreano [161]).

Bio-inspiration in robotics has a long history, dating back to Grey Walter's turtle robots (Holland [101]). A neurophysiologist, Grey Walter wanted to prove that complex looking behaviors did not require a large number of brain cells but that the secret of how the brain worked lay in the rich connections between neurons. His first two turtle robots were constructed between 1948 and 1949 and used a mere two electronic brain cells, which gave rise to a rich behavioral repertoire. The robots showed complex patterns of interactions with obstacles and "danced" when presented with mirror images of themselves or with another robot. They used phototaxis to autonomously return to a docking station to recharge their batteries. Walter's work has been credited as the creation of the first electronic autonomous robots and his turtle robots were one of the first ancestors of mod-

ern robotics, inspiring subsequent generations of robotics researchers (Holland [101], Pfeifer [174]).

In many ways evolutionary robotics builds on Walter's work. Evolutionary robotics typically uses simple neural networks to produce complex robot behaviors. It considers the robot and the environment as a tightly coupled dynamical system, and behavior is an emergent property of the interaction between the robot and the environment (Nolfi and Floreano [161]). Similar to Grey Walter's turtle robots, evolutionary robotics draws its key elements from the study of biological systems.

Evolutionary robotics has proven to be a successful alternative to programming robots manually when it is difficult to de-compose desired robot behaviors into simple behaviors to be implemented as separate modules of the control system (Baldassarre et al. [20], Quinn et al. [185]). This is because humans excel at designing very complex systems if they can be split up into small, modular subunits, but lack experience and intuition with complex behaviors (Braitenberg [38]). One class of robotic control problems that leads to particularly complex robot behavior is collective robotics (Husbands et al. [108]). This is because the behavior of multiple robots depends not only on interactions with the environment, but also on the behavior of other robots. As the number of interacting robots grows, or when robot behaviors become more sophisticated, the design of suitable control rules rapidly becomes very complex. This is especially true when agents are expected to coordinate or cooperate to collectively achieve a desired task. Chapter 4 discusses the application of evolutionary robotics to the control of teams of cooperating robots.

## 1.5 Structure of the Thesis

In this thesis I use artificial evolution to conduct selection experiments in simple artificial social organisms. Chapter 2 uses robotic experiments to conduct a quantitative test of Hamilton's rule. The results support current evolutionary theory and validate the feasibility and usefulness of this approach. Chapter 3 uses a computational model to study the influence of genetic architecture on the evolution of social behavior. It shows that the genotype to phenotype mapping can be a key factor for the evolution of efficient cooperation and division of labor in social insect colonies. The results also reveal that groups of interacting agents are

---

much more efficient when their behavior depends on the behavioral phenotypes of other group members. Chapter 4 reviews the current state of the art in the evolution of artificial multi-agent systems. It shows that multi-agent tasks can be divided into three categories by drawing on research of the evolution of cooperation in biological systems. The results demonstrate that genetic team composition and the level of selection significantly influence the performance of multi-agent systems. This chapter also provides a more in-depth discussion of the relationship between biological and robotic multi-agent systems. It concludes with suggested guidelines for the optimal combination of genetic team composition and level of selection for the artificial evolution of efficient multi-agent systems. Chapter 5 shows how agents evolved in simulation perform when tested on real robots. The results show that simulation is an essential tool for evolving controllers for real robots in spite of behavioral differences between simulation and the hardware platform. Chapter 6 summarizes the main achievements of this thesis and provides an outlook on future work.



---

# 2

## A Quantitative Test of Hamilton's Rule

---

The importance of Hamilton's work cannot be overstated - it is one of the few truly fundamental advances since Darwin in our understanding of natural selection.

Stuart A. West et al. (2007, [233])

Hamilton's theory of kin selection (Hamilton [91]) explains the evolution of cooperation using indirect selection, i.e., by considering the benefits of cooperation in groups of related individuals. In biology, this type of cooperation is known as altruistic cooperation or altruism. As pointed out in Chapter 1, it has proven very difficult to provide empirical data supporting Hamilton's theory and even 40 years after its introduction, no quantitative test of its role in evolution has been possible. Here we use a group of artificial social agents, implemented as autonomous foraging robots, to perform controlled selection experiments that are difficult or impossible to realize in biological organisms. We investigate the evolution of altruism in groups with different, well-defined levels of within-group relatedness, and across different costs and benefits for social actions. We illustrate the interactions of relatedness, costs and benefits for the evolution of altruism and provide the first quantitative test of Hamilton's rule.

### 2.1 State of the Art

Altruism takes a pivotal role in explaining the evolution of complex life. It is thought to have catalyzed major evolutionary transitions, such as the transition

from single-celled organisms to multi-cellularity (Maynard Smith and Szathmary [146]), and it is central to social interactions in all social species including humans. The framework used to understand altruism is built around Hamilton's rule (Equation 2.1), which provides a condition for the spread of a gene for a social action in a population (Bourke and Franks [37], Dawkins [59], Hamilton [91]), where we define a social action as a social behavior directed towards, or taking place between, members of the same species. It invokes three terms: the actor's reproductive cost  $c$  of the social action, the recipient's reproductive benefit  $b$  gained by the social action, and the genetic relatedness  $r$  between the actor and the recipient. According to Hamilton's rule, a gene for a social action will spread in a population if the relatedness  $r$  times the benefit  $b$  is greater than the cost  $c$ ,

$$r \cdot b - c > 0. \tag{2.1}$$

This rule explains how a gene will spread in a population if the number of individuals carrying the gene increases, irrespective of the fact whether this happens by increasing the reproduction of the bearer (actor), or by increasing the reproduction of other bearers of the same gene (related recipients).

Reproductive competition between relatives can influence the outcome of social interactions (Griffin and West [84], Taylor [214]). Since social groups in most social animals include relatives, a complete understanding of the evolution of altruism must consider both, the effects of kinship and the effects of competition. Recent theoretical models include the effects of competition and make predictions for the evolution of altruism (Frank [77], Queller [180], Appendix A).

Although the importance of empirical tests of these models has been acknowledged (Johnson and Gaines [114], West et al. [235]) and a methodology for such tests has been proposed (Oli [167]) such tests have been hindered due to practical difficulties. An ideal empirical test would include a direct measurement of all parameters that affect the evolution of altruism. However, even in the most primitive social species known (bacteria such as *Myxococcus* or *Pseudomonas* (Griffin et al. [85], Rainey and Rainey [187]) and social amoebae such as *Dictyostelium* (Strassmann et al. [211])) social interactions are complex and competition may vary in degree and over time and space. The only direct empirical test comes from a comparative study of 25 fig wasp species (West et al. [234]). Fig wasp taxa offer a unique population structure, with non-dispersing, wingless males confined to a single fruit and in fierce reproductive competition for winged females.

Relatedness between competing males varies greatly across species, owing to the number of females laying eggs into each fruit. By comparing levels of aggression in a range of average relatedness levels, the authors could confirm the detrimental effects of competition between relatives on the evolution of altruism. However, the difficulties in quantifying benefits and costs of social interactions did not allow a quantitative test of the theory.

Another way to perform an empirical test of Hamilton's rule is through selection experiments that carefully manipulate relatedness and competition during evolution. In an elegant example of such an experiment, Griffin et al. [85] separately varied competition and relatedness in a strain of bacteria that metabolize growth-enhancing iron as an altruistic trait. The authors also investigated the effects of competition between relatives and their data provided experimental proof that increased competition between relatives leads to lower levels of altruism. However, difficulties in quantifying benefits and costs for social actions again impeded a quantitative test.

Here we use a group of artificial social agents with precisely defined within-group relatedness for a quantitative analysis of the evolution of altruism. Social agents are implemented as small mobile robots. By independently varying within-group relatedness and the costs and benefits of social actions, we can quantify the relative importance of both factors for the evolution of altruism. Importantly, our setup allows us to study the evolution of simple artificial social organisms while retaining full control over the experimental parameters governing evolution.

## 2.2 Materials and Methods

Our experimental setup used a physics-based simulation of a group of micro robots (Magenat and Waibel [132], Magenat et al. [133]). The simulation was modeled on a parallel hardware setup (Appendix E), where social agents were implemented as small, autonomous mobile robots (Figure 2.1).

### 2.2.1 Experimental setup

We chose a collective foraging task to investigate the evolution of altruism, because foraging efficiency is a key factor for many biological social groups such as




Figure 2.1: (a) The software setup. 8 randomly positioned robots (black squares with arrows) searched for 8 randomly positioned tokens (small circles) and transported them into the target area (hatched area at bottom) under the white wall (the other three walls were black). (b) The same setup in hardware.

ant or bee colonies (Seeley [201], Traniello [217]), as well as an important application for collective robotic systems (Cao et al. [48]). Groups of 8 micro-robots and 8 food tokens were randomly placed in a  $50 \times 50 \text{ cm}^2$  arena (Figure 2.1). A robot could forage a food token by pushing it into a pre-defined area at one side of the arena marked by a white wall (the “nest”).

The robots (Caprari [49]) were small ( $2 \times 2 \times 4 \text{ cm}^3$ ), two-wheeled autonomous mobile robots equipped with three infrared sensors at the front, which could sense obstacles up to 3 cm away (Figures 2.1 and 2.2). An extension module adding a fourth infrared sensor with a range of up to 6 cm and a linear camera was mounted higher on the robot, overlooking food tokens.

Each robot’s behavior was controlled by a feed-forward neural network consisting of six sensory inputs, one bias input and six artificial neurons (Figure 2.2). Inputs and neurons were connected by synaptic weights representing the strength of the connection. The sensory inputs were given by the robot’s four infrared sensors providing distance information on tokens, walls and other robots, and the robot’s camera sensor providing visual information on the location of the target area. The camera sensor was split into two parts consisting of the left and



right part of the robot's field of view, respectively. It only provided binary color information (0 or 1) for the white or black walls. The bias input was connected to a constant input with value -1. Sensor inputs were scaled to a range of  $[-1; 1]$  and fed to the neurons. The activation of each of the 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). A single output neuron was used to control each, the left motor speed, the right motor speed, and a token sharing switch, which allowed robots to choose between the selfish or altruist payoffs at the moment of successful token foraging and which was ignored otherwise. The 33 genes of an individual robot each encoded an integer value from 0 to 255 (i.e., 8 bit resolution) which controlled the synaptic weights of the 33 neural connections. The genome given by the neural network's connection weights determined a robot's behavior.

### 2.2.2 Costs and benefits of altruism

When a robot had foraged a token, it could choose between either of two types of payoffs: A selfish payoff  $c$  that was only added to the robot's individual fitness, or an altruist payoff  $b$  that was awarded to each of the 7 other robots in the group. In other words, a robot had to forego its selfish payoff  $c$  to cooperate and thus paid a cost  $c$  for cooperating. At the same time the cooperating robot bestowed a benefit  $b$  on each group member except itself. This corresponds to a situation where individuals face a choice between investing resources into their own reproductive fitness or costly investment into that of their social group.

The evolution of altruism depends on the relative fitness gained for selfish versus altruistic social behavior, i.e. on the relative size of costs  $c$  versus benefits  $b$ . In all experiments presented here the size of costs  $c$  was arbitrarily fixed to  $c = 1$  and the ratio of costs to benefits  $\frac{c}{b}$  in Equation 2.1 was varied by varying the benefits  $b$  for altruistic social actions, except for  $\frac{c}{b} = 0$  where  $c = 0$ .

All evolutionary experiments presented here were conducted with 200 groups of 8 individuals each. Since individuals competed for reproduction with all members of the population, including those of their own group, our experiments resulted in a small amount of competition between related individuals from the same group. Since reproductive competition between related individuals can influence the outcome of social interactions (Griffin and West [84], Taylor [214]), cost to benefit ratios were corrected for this small effect (Frank [77], Queller [180],




Figure 2.2: (a) The robots were equipped with three infrared (IR) sensors at the front to detect obstacles and a linear camera to identify the white nest. A fourth infrared sensor overlooking food tokens allowed robots to distinguish tokens from walls and other robots. (b) Robot behavior was controlled by an artificial neural network. Inputs (black circles) were given by the robot's four infrared (IR) sensor activation values and two camera readings. Six neurons processed information (large white circles). Their firing threshold values were set by connections to a constant input value ("bias", small white circle). Two outputs were used to control the speed of the left and right wheel motors and a third output allowed the robots to choose whether to share a food token pushed into the nest. The 33 connection weights (lines connecting circles) were encoded to form the artificial genome (see text).

Appendix A).

### 2.2.3 Relatedness and artificial evolution

To investigate the role of relatedness, costs and benefits on the evolution of altruism, we used artificial evolution to create, test and select groups of robots. Since relatedness of biological groups varies widely, we compared 5 different levels of relatedness  $r = \{1; 0.75; 0.54; 0.25; 0\}$  spanning a wide range of relatedness values. Since the spread of genes for a social action depends on the relative sizes of costs  $c$ , benefits  $b$  and relatedness  $r$  linked by Hamilton's rule, we considered 5 different cost to benefit ratios  $\frac{c}{b} = \{1; r - 0.1; r; r + 0.1; 0\}$  with special focus on the predicted transition point  $\frac{c}{b} = r$ .

For the case of highly related individuals ( $r = 1$ ) each of the 200 groups was initially founded by a single individual with a random genome. This genome was duplicated 7 times to obtain 8 identical individuals per group. All groups were evaluated 10 times in the foraging task and the fitness of each individual was recorded. After evaluation one individual of all individuals of all groups was randomly selected with a probability proportional to its fitness (roulette wheel selection with replacement, Bäck [16]). A random portion of the genome of this individual was exchanged with that of another, different individual selected in the same fashion, with a “crossover probability” of 0.005. The resulting genome was subjected to mutation with a “mutation probability” of 0.005 per bit, i.e. 0.04 per gene and then duplicated 7 times to found a new group. This procedure was repeated 200 times to produce 200 new groups, each of which thus comprised 8 genetically identical (clonal) individuals. However, the genomes were always different between groups, hence leading to a within-group relatedness of 1.

To obtain groups with an average relatedness of  $r = 0.75$  we used the same procedure, but founded each group using two individuals. The first genome was duplicated 6 times to obtain a total of 8 genomes. This led to the formation of groups comprising two types of individuals with frequencies 7:1, resulting in an average relatedness of 0.75 (see Appendix A).

To obtain groups with an average relatedness of  $r = \frac{15}{28} \simeq 0.54$  we followed the same procedure, but initiated groups with three individuals. The first genome was duplicated 5 times to obtain groups with a 6:1:1 frequency distribution of the three genomes, corresponding to an average relatedness of  $\simeq 0.54$ .

To obtain groups with an average relatedness of  $r = 0.25$  we followed the same procedure, initiating groups with three founding individuals. Groups were composed of three types of clones in frequencies 3:3:2 as to result in an overall relatedness of 0.25.

To obtain groups with a relatedness of  $r = 0$  we followed the same procedure but initiated each group with 8 different individuals. Under this type of group formation individuals were, on average, not genetically more similar to individuals in their group than to individuals of other groups. Hence, the within-group relatedness was 0.

The genetic composition of groups thus differed from that of most animal groups in that some individuals were clones ( $r = 1$ ) rather than belonging to kin classes such as full siblings ( $r = 0.5$ ) or cousins ( $r = 0.125$ ). However, in the

absence of preferential interactions between kin, social evolution is influenced by the average group relatedness and not the exact relatedness between pairs of individuals. This is because genetic relatedness depends on interaction probabilities of genes (Bourke and Franks [37]) which in our model is equivalent to interaction probabilities of individuals. Our experimental setup prevented preferential interactions between individuals by randomizing starting positions (Section 2.2.1) and by using a neural network that did not allow individuals to memorize past interactions (Figure 2.2).

Overall, there were therefore five times 200 groups of 8 robots each for the five different within-group relatedness values. For each relatedness value groups were evaluated for the five corresponding cost to benefit ratios (with the exception of  $\frac{c}{b} = r - 0.1$  for groups with relatedness  $r = 0$  as the cost to benefit ratio cannot be negative). To evaluate groups we tested them in the foraging task for 10 separate 60-second evaluations. A fitness value  $F$  was calculated for each individual as the averaged sum of all its payoff values obtained through its own actions and altruistic actions of its group members. The level of altruism  $A$  was estimated for each group as the proportion of collected food tokens that was shared with group members,  $A = \frac{n_a}{n_a + n_s}$ , where  $n_a$  was the number of collected food tokens individuals shared and  $n_s$  the number of tokens individuals did not share. All selection experiments were repeated 20 times (20 independent replicates) and evolved for 500 generations for each experimental condition. To compare fitness and the level of altruism, we averaged the performance of the 200 groups over the last 10 generations of the 20 replicates. These 20 values per condition were compared with Student's  $t$ -tests. Data are first presented for groups of highly related individuals ( $r = 1$ ), next for groups with intermediate relatedness ( $r = 0.75$ ,  $r \simeq 0.54$  and  $r = 0.25$ ), and finally for groups with low relatedness ( $r = 0$ ).

## 2.3 Results

All experiments led to the evolution of foraging behavior (Appendix B). Figure 2.3 shows the mean level of altruism  $\bar{A}$  in groups (averaged over the last 10 generations) after 500 generations of artificial selection. As predicted by Hamilton's rule, high levels of altruism resulted for high levels of within-group relatedness relative to the cost to benefit ratio for token sharing ( $r \gg \frac{c}{b}$ ). Conversely, low levels of altruism were linked to low levels of within-group relatedness relative

to the cost to benefit ratio ( $r \ll \frac{c}{b}$ ). Note that at generation 0, the initial levels of altruism in groups were, by configuration, symmetrically distributed in the interval  $[0; 1]$ , with an expected mean level of altruism  $\bar{A} = 0.5$  (see Figures B.1 and B.2 in Appendix B).

To quantitatively assess the ratio  $\frac{c}{b}$  at the transition point  $\bar{A} = 0.5$  for a given within-group relatedness  $r$ , we used multiple linear regression on data at  $\frac{c}{b} = r - 0.1$ ,  $\frac{c}{b} = r$ , and  $\frac{c}{b} = r + 0.1$ . We estimated  $\frac{c}{b}$  from our data by intersecting the obtained regression line with the expected level of altruism at the transition point,  $\bar{A} = 0.5$ . We then analyzed the deviation of this estimate from the theoretical prediction ( $\frac{c}{b} = r$ ).

### 2.3.1 Groups with high relatedness

For groups with a relatedness of  $r = 1$ , low cost to benefit ratios ( $\frac{c}{b} = 0$ ) led to a very high proportion of shared food tokens (Table 2.1). In accordance with Hamilton's rule, our results indicate a transition from high to low levels of altruism when the within-group relatedness is equal to the cost to benefit ratio,  $r = \frac{c}{b} = 1$ . When  $r = \frac{c}{b}$  there was no selection for altruism, and after 500 generations of artificial evolution the mean level of altruism was not significantly different from its initial value  $\bar{A} = 0.5$  (one-sample  $t$ -test,  $df = 19$ ,  $P = 0.690$ ). High cost to benefit ratios ( $\frac{c}{b} = 1.1$ ) led to a very low proportion of shared food tokens. Multiple linear regression of data at  $\frac{c}{b} = 0.9$ ,  $\frac{c}{b} = 1$ , and  $\frac{c}{b} = 1.1$  led to an estimated transition point at  $\frac{c}{b} = 0.994$  for groups with  $r = 1$  (multiple linear regression, 95% confidence interval  $[0.768; 1.285]$ , slope of regression line  $-3.604$ ).

Table 2.1: Mean levels of altruism  $\bar{A}$  for different cost to benefit ratios  $\frac{c}{b}$  in groups with different within-group relatedness  $r$  (compare Figure 2.3).

	$\frac{c}{b} = 0$	$\frac{c}{b} = r - 0.1$	$\frac{c}{b} = r$	$\frac{c}{b} = r + 0.1$	$\frac{c}{b} = 1$
$r = 1$	0.983	0.875	0.478	0.110	0.478
$r = 0.75$	0.978	0.890	0.527	0.126	0.054
$r \simeq 0.54$	0.976	0.875	0.457	0.095	0.030
$r = 0.25$	0.959	0.901	0.469	0.061	0.021
$r = 0$	0.558	-	0.558	0.017	0.014

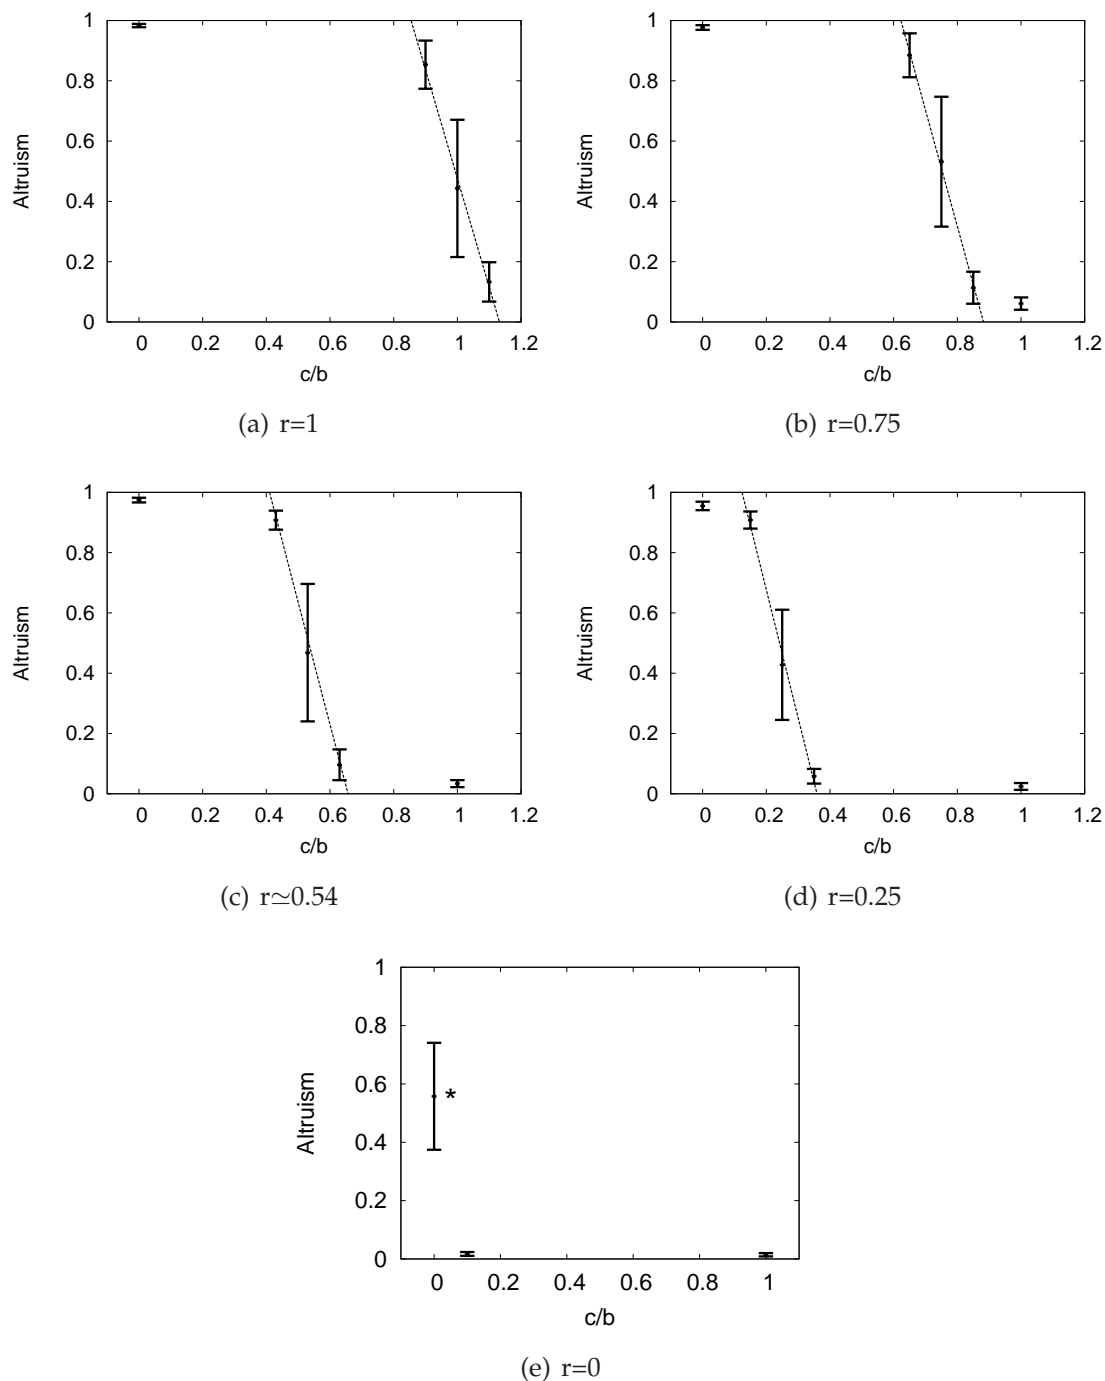


Figure 2.3: Altruism  $\bar{A}$  after 500 generations of artificial evolution for different relatedness values  $r$  and cost to benefit ratios  $\frac{c}{b}$ . Data shown are the mean and standard deviations of 20 replicates per condition. For  $r = \frac{c}{b} = 0$  (data point marked with a \* in graph (e)) six populations went extinct and the data point shown is an average of 14 experiments. The dashed line is the result of multiple linear regression on data for  $\frac{c}{b} = r - 0.1$ ,  $\frac{c}{b} = r$ , and  $\frac{c}{b} = r + 0.1$  (see text).

### 2.3.2 Groups with intermediate relatedness

For groups with intermediate relatedness values ( $r = 0.75$ ,  $r \simeq 0.54$  and  $r = 0.25$ ), low cost to benefit ratios ( $\frac{c}{b} = 0$ ) again led to a very high proportion of shared food tokens, and a similar pattern of transition from high to low levels of altruism persisted. For  $r = \frac{c}{b}$  artificial evolution for 500 generations did not affect the expected mean level of altruism  $\bar{A} = 0.5$  (one-sample  $t$ -tests,  $df = 19$ ,  $P = \{0.596; 0.452; 0.526\}$  for  $r = \{0.75; 0.54; 0.25\}$ ). As predicted by Hamilton's rule, a decrease in within-group relatedness led to a shift of the transition point towards lower ratios of  $\frac{c}{b}$ . High cost to benefit ratios  $\frac{c}{b} = 1$  again resulted in very low levels of altruism (Table 2.1). Linear regression led to estimated cost to benefit ratios of  $\frac{c}{b} = 0.753$  (95% confidence interval  $[0.603; 0.940]$ , slope  $-3.855$ ),  $\frac{c}{b} = 0.533$  ( $[0.430; 0.661]$ , slope  $-4.058$ ), and  $\frac{c}{b} = 0.242$  ( $[0.204; 0.287]$ , slope  $-4.250$ ) for relatedness values of  $r = 0.75$ ,  $r \simeq 0.54$  and  $r = 0.25$ , respectively.

### 2.3.3 Groups with low relatedness

When group members were unrelated (within-group relatedness  $r = 0$ ), low cost to benefit ratios ( $\frac{c}{b} = 0$ ) led to the extinction of 6 out of 20 replicates, i.e. in 6 out of the 20 replicates artificial selection could not be completed for the entire 500 generations, because all individuals in the population had a fitness of 0. Populations went extinct at generations 63, 90, 287, 374, 459, and 490 in these 6 replicates. The mean final level of altruism of the remaining 14 replicates at generation 500 for  $\frac{c}{b} = 0 = r$  was not significantly different from the initial value of  $\bar{A} = 0.5$  (one-sample  $t$ -test,  $df = 13$ ,  $P = 0.631$ ). High cost to benefit ratios ( $\frac{c}{b} = 1$ ) in this condition resulted in the lowest values of the level of altruism observed in all our experiments (Table 2.1). Due to the lack of data for  $\frac{c}{b} = r - 0.1$  no linear regression estimate was performed for this condition.

## 2.4 Discussion

Our results provide a clear experimental demonstration and a quantitative test of Hamilton's rule.

High within-group relatedness and low cost to benefit ratios resulted in high levels of altruism. This was because when cost to benefit ratios were much smaller than within-group relatedness, altruism was highly beneficial and led to large fit-



ness benefits for altruists. As the value of the cost to benefit ratio increased and approached the value of within-group relatedness, the size of benefits for altruism decreased. This led to lower selection pressure on altruistic behavior and higher behavioral variation. When cost to benefit ratios equalled within-group relatedness the reproductive success of a genotype did not depend on the level of altruism and the level of altruism was subject to random drift. This was because in this case the direct fitness benefits gained by selfish behavior were exactly equal to the indirect fitness benefits gained by increasing the reproduction of related group members,  $c = b \cdot r$ . When cost to benefit ratios exceeded within-group relatedness, altruism became costly. Increasing cost to benefit ratios led to increased costs for altruism and strong selection for selfish behavior. This explains the low levels of altruism for low within-group relatedness and high cost to benefit ratios.

Hamilton's theoretical model predicts an instantaneous transition between selection for and selection against altruism based on the inequality  $\frac{c}{b} < r$  (Appendix A). However, under natural (and, here, artificial) conditions traits vary, with the size of a trait's variation depending on the strength of selection. In our experiments we varied cost to benefit ratios by a constant absolute value ( $\pm 0.1$ ), which led to different relative changes for different within-group relatedness. Thus a 0.1 increase to the cost over benefit ratio corresponded to a 10% increase for  $r = 1$ , but to a 40% increase for  $r = 0.25$ . The effects of the resulting changes in selection pressure on the transition between high and low levels of altruism can be seen from the regression analysis. The slope of regression lines was highest for experiments with high within-group relatedness and decreased for experiments with lower within-group relatedness. Simultaneously variation in the level of altruism of data points for  $\frac{c}{b} = r - 0.1$  and  $\frac{c}{b} = r + 0.1$  decreased with decreasing relatedness.

Our results illustrate that even slight alterations in the benefits and costs of altruism that do not lead to significant changes in group fitness can translate into drastic shifts in the amount of altruism. For example, for  $r = 1$  a 10% increase or 10% decrease in the benefits of altruism in our setup resulted in an up to 83% increase or 77% decrease in the level of altruism.

When within-group relatedness and cost to benefit ratios were 0 ( $r = \frac{c}{b} = 0$ ) evolution led to the extinction of some populations. This was because in the absence of related recipients, genes could not spread by increasing the reproduc-



tion of other bearers of the same gene, i.e., there were no indirect fitness benefits. In this condition robots could only improve their relative fitness with respect to other individuals by opting for the selfish payoff rather than increasing the fitness of their competitors. At the same time, cost to benefit ratios of zero meant that there was no selfish payoff to actors ( $c = 0$ ) and therefore no direct fitness benefits. Robots could therefore not increase their own fitness directly and thus genes could not spread by increasing the reproduction of their bearer. However, this led to the collapse of altruism and fitness and thus resulted in the extinction of populations.

In our experiments competition between relatives was a negligible factor (Appendix B). However, previous work has shown that for some cases competition between relatives can have a large effect on the evolution of altruism (Griffin et al. [85]). The influence of this factor requires additional simulations and is a subject of further study.

## 2.5 Conclusion

Our experiments have allowed a quantitative test of Hamilton's rule in a robotic system. Similar tests in biological systems have been hindered due to practical difficulties in measuring relatedness, costs and benefits. Quantitative estimates of genetic relatedness of natural organisms commonly use pedigrees or genetic markers (Queller and Goodnight [181]). However, both methods are susceptible to biased estimates and often result in large sampling errors (Pamilo [170]). Here we avoided problems linked to *a-posteriori* estimation of genetic relatedness by artificially constructing social groups with precisely defined within-group relatedness. Similar results could be obtained by using guided evolution of social organisms but are only feasible in organisms with sufficiently short generation times (Griffin et al. [85], Velicer and Yu [221]).

Quantifying the costs and benefits of cooperative social actions in natural organisms is significantly more difficult. Such tests require detailed information on lifetime social interactions between individuals, which may vary over space and time (Johnson and Gaines [114], MacColl and Hatchwell [131], West et al. [235]). Here these requirements could be met by fixing the values of costs and benefits of social actions, and by restricting social interactions to well defined groups. It remains to be seen whether similar approaches can be used to overcome these

problems in biological systems (Oli [167]).

While the robotic model used here allowed us to avoid many of the practical difficulties encountered in biological systems, the social behaviors resulting from the evolution of robotic agents are inevitably much simpler than those found in their biological counterparts. The embodied nature of robotic agents nevertheless allows for richer and more plausible behaviors than purely computational models.

This study demonstrates that social behaviors such as altruism can evolve in groups of robots with simple neural networks. Our experiments indicate parallels between the evolutionary principles governing the evolution of biological and artificial social organisms. They also illustrate that knowledge transfer between the fields of evolutionary biology and evolutionary robotics can provide insights into the evolution of social behavior.

---

# 3

## Division of Labor and Colony Efficiency in Social Insects: Effects of Interactions between Genetic Architecture, Colony Kin Structure and Rate of Perturbations

---

[...] the brain of an ant is one of the most  
marvellous atoms of matter in the world, perhaps  
more so than the brain of a man.

Charles Darwin (*The Origin of Species*, 1859)

In the previous chapter we have discussed Hamilton's theory of kin selection, and we have shown that the evolution of sociality critically depends on the fitness benefits of social life. Hamilton's theory defines conditions for the spread of social genes, however it provides no information on the size of costs or benefits which determine the evolution of cooperative social actions. As pointed out in Chapter 1, the underlying genetic architecture can be an important factor that influences the costs and benefits of cooperation and therefore shapes the evolution of sociality. In this chapter we use a computational model to investigate how the genetic architecture affects fitness benefits in social insect colonies. The remainder of this chapter is based on a publication (Waibel et al. [222]).

### 3.1 Division of Labor and Colony Efficiency in Social Insects

Social insects (ants, wasps, bees and termites) take a special role in the study of the evolution of cooperation and altruism. This is because they provide the most advanced examples of social organization known, matched in scale and complexity only by human societies. Single colonies can consist of more than a million individuals and are marked by a high degree of specialization and by an efficient division of labor. Social insect societies base their success on a worker caste who forego their own reproduction and devote their whole lives to caring for the queen, constructing and protecting the nest, foraging for resources, and tending to the larvae. Such behavior is maximally altruistic: sterile workers do not leave any offspring of their own, but solely enhance the reproductive fitness of the queen. Other examples include the evolution of soldier castes with “kamikaze” weapons such as detachable stings (Seeley [201]) and exploding abdomens used in defense of the colony (Hölldobler and Wilson [104]).

It has been argued that the complex social organization of social insect societies is the cause of their ecological success (Hölldobler and Wilson [104]). Social insects compose about 15 percent of the entire animal bio-mass of most terrestrial environments (Wilson [238]), and this rate has been estimated to be as high as one third of the entire animal bio-mass in the Amazon rainforest and an astonishing 80 percent of the total insect bio-mass (Maynard Smith and Szathmáry [147]). However, out of the 750'000 insect species recognized by biologists only 2 percent are social (Hölldobler and Wilson [104]).

Colony level flexibility in response to external changes and internal perturbation is an essential feature of division of labor in social insects (Calabi [44], Robinson [193]). A colony has to perform a number of tasks, such as feeding the brood, foraging for resources, maintaining the nest, and defending the colony. Efficient allocation of individuals to these different tasks requires continuous dynamic adjustments in response to external changes such as risk of intra- and interspecific competition and amount of food available. The colony must also be able to respond to internal perturbations, such as specific mortality of individuals undertaking foraging or another task.

Considerable work has focussed on identifying the mechanisms which allow an efficient colony response to information that exceeds the sensory range or cog-

nitive capacity of any given individual (Bonabeau et al. [32]). The large majority of these models are built on the observation that individuals in a colony vary in their propensity to perform the various tasks. Experimental studies in bees, wasps and ants have revealed that workers vary genetically in their sensitivity to task stimuli (Fewell and Page Jr [67], O'Donnell [166], Page et al. [168]). Thus, within a colony workers from the same genetic background (same patriline or matriline) are more likely to perform similar tasks than are less related workers. To explain this association several response threshold models have been developed (Bertram et al. [26], Bonabeau [29], Gautrais et al. [78], Theraulaz et al. [215]). In these models, individuals vary genetically in the stimulus (threshold) for a given task at which they begin to perform that task. Only a few workers (those with low thresholds) will perform a task when the stimulus is very low. However, as the stimulus level increases, the thresholds of more individuals are exceeded and those workers begin performing the task.

While numerous models have investigated the mechanisms allowing efficient task allocation under situations where individuals in a colony vary in their response thresholds (Bertram et al. [26], Bonabeau [29], Gautrais et al. [78], Theraulaz et al. [215]), little attention has been devoted to the mechanisms responsible for variation in response thresholds. Most models assume that variation in task sensitivity is a consequence of additive effects of genotype and environment (Bonabeau et al. [31], Robinson et al. [194]) and that the distribution of genetic thresholds in the group is continuous and normally distributed. However, studies in the honey bee suggest that genotypic variation in tasks may involve a few major loci (Hunt et al. [107], Page et al. [169]) and also implicate genetic architectures with non-additive gene interactions (Rüppell et al. [197]). As pointed out by Bertram et al. [26], a better understanding of division of labor requires information on how genotypic variation relates to differences in intra-colony response threshold distributions. In a first step toward this goal, these authors developed a model and showed important effects on task allocation and colony behavior for differences in the number of loci and alleles controlling individual response thresholds.

The aim of the present study is to understand how different genetic architectures of threshold response may affect colony efficiency. While the work of Bertram et al. revealed that the genetic underpinning of response thresholds impinges on the pattern of task allocation, their static analytical model did not al-

low them to study the impact on colony efficiency, nor how variation in genetic architecture influences the evolutionary trajectory of division of labor. To address these issues we used artificial evolution to compare the performance of three simple genetic architectures underlying within colony variation in response thresholds of workers. We evolved five different types of colonies (colonies in which individuals are genetically identical ( $r=1$ ), colonies with intermediate relatedness values ( $r \simeq 0.25$ ,  $r \simeq 0.5$  and  $r \simeq 0.75$ ) and colonies formed by individuals randomly selected in the population ( $r=0$ )). The performances of the three genetic architectures and five types of colonies were compared under different rates of environmental perturbation. Our approach helps to bridge the gap between studies aiming at understanding the genetic basis of behavioral differences among colony members and evolutionary studies focusing on how individual differences in response thresholds may lead to efficient division of labor. Importantly, our approach also allows us to study task distribution when there are more than two tasks and how departure from an optimal distribution affects overall colony performance.

## 3.2 Materials and Methods

### 3.2.1 Colony tasks

We use a simple agent-based model to compare performance of colonies consisting of 100 individuals. Individuals can engage in five different tasks. We model sub-linear return functions typical for tasks as foraging or scouting using the five exponential functions

$$f_i = 1 - \exp(-ix_i), \quad (3.1)$$

where  $i = 1, 2, 3, 4, 5$  are the tasks and  $x_i$  is the percentage of the colony performing the task (Figure 3.1). Hence, the optimal allocation of workers to the five tasks has a non-trivial solution depicted in Figure 3.1. Importantly, the results of our simulations are not affected by the shape of payoff functions for each task, the general issue being that deviation from a given optimal task allocation results in decreased overall colony fitness.

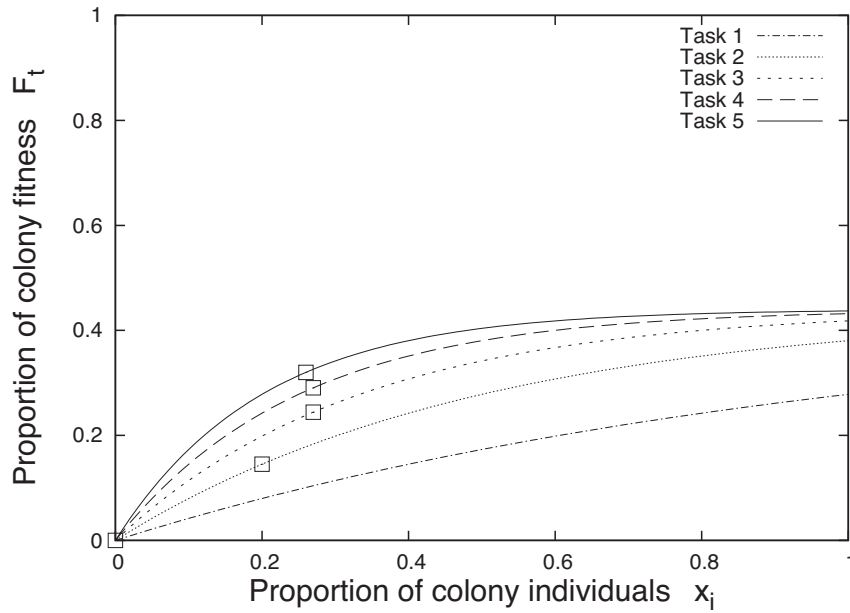


Figure 3.1: Proportion of the maximum colony fitness provided by each of the five tasks in function of the proportion of workers engaged in this task. Squares indicate the optimal allocation of workers between the five tasks (i.e., the distribution yielding maximum colony fitness).

### 3.2.2 Genetic architecture

We consider a simple situation where there are five genes ( $g_1 - g_5$ ), each encoding an integer value from 0 to 255 (i.e., 8 bit resolution) for one of the five tasks. We compare three mapping systems between genes and the behavioral phenotype of workers.

In the first genetic architecture (deterministic mapping), individuals engage in the task with the highest genetically encoded value  $\max(g_i)$ . This mapping system represents a situation where individual response thresholds are genetically determined with no environmental or social influence on individual responses. In the second genetic architecture (probabilistic mapping), the five genes influence the probability of engaging in a task. The corresponding gene value  $g_i$  is proportional to the probability  $P_i$  to engage in task  $i$ ,  $P_i = (\frac{g_i}{\sum_j g_j})$ . This system is similar to a situation where, in addition to a genetic component, the response thresholds of workers are also influenced by developmental noise or fluctuating environmental factors. Finally, in the third genetic architecture (dynamic mapping), the propensity of individuals to engage in each of the five tasks depends not only on

their genotype but also on the behavioral phenotypes of other colony members. Thus, individuals will perform the task  $i$  that maximizes the ratio  $\frac{g_i}{a_i}$ , where  $a_i$  is the number of colony members already working on task  $i$ , i.e.,  $\max(\frac{g_i}{a_i})$ . This system corresponds to the observed pattern of worker-worker inhibition where the likelihood of individuals to perform a given task is inversely proportional to the proportion of colony members already performing that task (Huang [105], Huang and Robinson [106], Wilson [239]). This system also mimics a situation where the stimulus for a given task decreases with increased number of individuals performing this task.

### 3.2.3 Environmental and internal perturbations

To study the ability of colonies to respond to perturbations in worker task allocation we studied evolution in environments with different degrees of perturbation. Colony life-span was divided into 100 time-steps. At each time-step we randomly selected one of the five tasks and with probability  $P_r$  removed all individuals performing this task to replace them with new individuals with genomes created from the same parents. The new individuals were then assigned to a given task according to their genome and genetic architecture. We used five different probabilities (0.0, 0.1, 0.2, 0.3 and 0.4) of perturbation per time-step. In other words, colony size was held constant over the 100 time-steps and, for each step, all individuals engaged in one of the five tasks (randomly selected) were replaced with probability  $P_r$ . Total colony fitness  $F$  was obtained by adding the fitness obtained at each time-step:

$$F = \sum_{t=1}^{100} F_t. \quad (3.2)$$

and colony fitness at each time-step ( $F_t$ ) was quantified as the sum of the payoff of each task  $i$ ,

$$F_t = \sum_{i=1}^5 f_i(x). \quad (3.3)$$

The minimum fitness of the colony is obtained when all 100 individuals engage in the least rewarding task during the entire colony life. Inversely, maximum fitness is obtained when there is an optimal ratio of individuals engaging in each of the five tasks. With the task payoffs chosen here the optimal solution is



at  $x_i = (0, 0.198, 0.267, 0.272, 0.262)$  as indicated by the squares in figure 3.1. For simplicity all fitness values are presented on a scale of 0-100, with 0 being the minimal possible fitness and 100 the maximum value.

### 3.2.4 Colonies and selection algorithm

We conducted selection experiments over 1000 generations in 100 colonies with five different levels of relatedness ( $r \simeq 0, 0.25, 0.5, 0.75$  and 1). Similar to the creation of groups in the previous chapter, colonies were again composed of different proportions of clones. However, since groups consisted of 100 individuals rather than 8 individual the exact relatedness values differed. For  $r=0$  the 100 colonies of unrelated individuals were initially formed by using 100 randomly generated genomes for each colony. The fitness of these colonies was compared and we randomly selected one individual of one of the 30 colonies with highest fitness. The genome of this individual was subjected to mutation (probability 0.2% per bit, i.e. 1.6% per gene) to form one individual of the next generation of colonies. This procedure was repeated 10'000 times to produce 100 new colonies, each consisting of 100 individuals. Under this type of colony formation individuals are, on average, not genetically more similar to individuals in their colony than to individuals of other colonies. Hence, the within-colony relatedness is 0. To construct colonies with highly related individuals ( $r=1$ ) we followed the same procedure but initiated each colony with only one individual which was duplicated 99 times after having been subjected to mutation. Each of the 100 colonies thus comprised 100 genetically identical (clonal) individuals. However, the genomes were always different between colonies, hence leading to a within-colony relatedness of 1. To obtain colonies with an approximate relatedness of  $\simeq 0.25$  we used the same procedure but started colonies with four founding individuals that were each clonally copied 24 times after having been mutated. Colonies thus consisted of four types of clones in equal frequency, leading to an overall relatedness of  $\simeq 0.25$  (as in Chapter 2, the precise relatedness can be calculated from equation A.2 in Appendix A). To produce colonies with relatedness  $\simeq 0.5$  we followed the same procedure but initiated colonies with two individuals. This led to the formation of colonies comprising two types of clones in equal proportions (i.e., a relatedness of  $\simeq 0.5$ ). Finally, colonies with a relatedness of  $r \simeq 0.75$  were created in the same way but using unequal proportions of the two types of clones. The proportion  $p$  of the most frequent clone was estimated so that the probability

of an individual to interact with another individual with the same genome was on average  $\simeq 0.75$  ( $r = p^2 + (1 - p)^2$ , see equation A.2 in Appendix A).

While constructing colonies out of different proportions of clones does not reflect the haplodiploid mode of reproduction present in real social insects, the model does accurately preserve interaction probabilities of genes.

Overall, there were therefore five groups of 100 colonies. The selection experiments were repeated 10 times over 1000 generations for each of these five groups. To compare performance of the three genetic architectures we averaged the performance of the 100 colonies over the first and last 10 generations in 10 separate simulation runs. These 10 values per genetic architecture were compared with Student's *t*-tests. Data are first presented for colonies of unrelated individuals, next for colonies of highly ( $r=1$ ) related individuals and finally for colonies of intermediate ( $r \simeq 0.25, 0.5$ , and  $0.75$ ) relatedness.

### 3.3 Results

#### 3.3.1 Colonies of unrelated individuals

When colonies did not experience perturbations (i.e., there were no instances of selective mortality of individuals performing a given task) there were only small differences between the three genetic architectures in performance during the 10 first generations of selection (Figure 3.2). The lower fitness values were obtained with the deterministic and probabilistic architectures. By their configurations, these two genetic architectures initially lead to a random distribution of individuals among the three castes and, accordingly, fitness values close to the value expected under such a distribution (expected value under random distribution of workers, 10'000 simulations: 90.86; dyn:  $\bar{F} = 92.10$ ; det:  $\bar{F} = 90.78$ ; prob:  $\bar{F} = 90.89$ ). During the first 10 generations, the dynamic mapping system ( $\bar{F} = 92.18$ ) led to a significantly higher fitness than the two other mapping systems (two-sample *t*-tests, dyn vs. det:  $\bar{F} = 91.91$ ,  $t = 4.25$ ,  $df = 18$ ,  $P < 0.001$ ; dyn vs. prob:  $\bar{F} = 91.09$ ,  $t = 33.31$ ,  $df = 18$ ,  $P < 0.001$ ) because this genetic architecture leads to a relatively equivalent distribution of workers to each of the five tasks, which results in a slightly higher overall fitness than a random distribution.

During the 1000 generations of selection there was an increase in colony fitness for each of the three genetic architectures (Figure 3.2). At the end of evolution




Figure 3.2: Mean fitness for colonies of unrelated individuals ( $r=0$ ) for three genetic architectures (dynamic: solid line; deterministic: dashed line; probabilistic: dotted line) without perturbations ( $P_r = 0.00$ ).

there was a significant difference in colony fitness between the three genetic architectures, with the highest performance achieved by the deterministic and the lowest with the probabilistic system (Table 3.1).

Evolution with selective removal of individuals engaged in a given task led to a significant reduction in colony fitness and a shift in the relative performances of the three genetic architectures (Figure 3.3). For each of the four levels of perturbation, the dynamic mapping system led to an initially greater colony fitness compared to the two other genetic architectures (dyn:  $\bar{F} = 91.74/91.67/91.51/91.36$ ; det:  $\bar{F} = 81.07/74.15/67.78/62.44$ ; prob:  $\bar{F} = 80.55/72.63/65.45/59.32$  for  $P_r = 0.10/0.20/0.30/0.40$ , all  $P < 0.001$ , Figure 3.3). During the 1000 generations of selection there was an increase in fitness for each of the three genetic architectures. However, there were still marked differences in performance with, in all cases, the dynamic system performing significantly better than the two others and the probabilistic system performing significantly worst (Table 3.1).

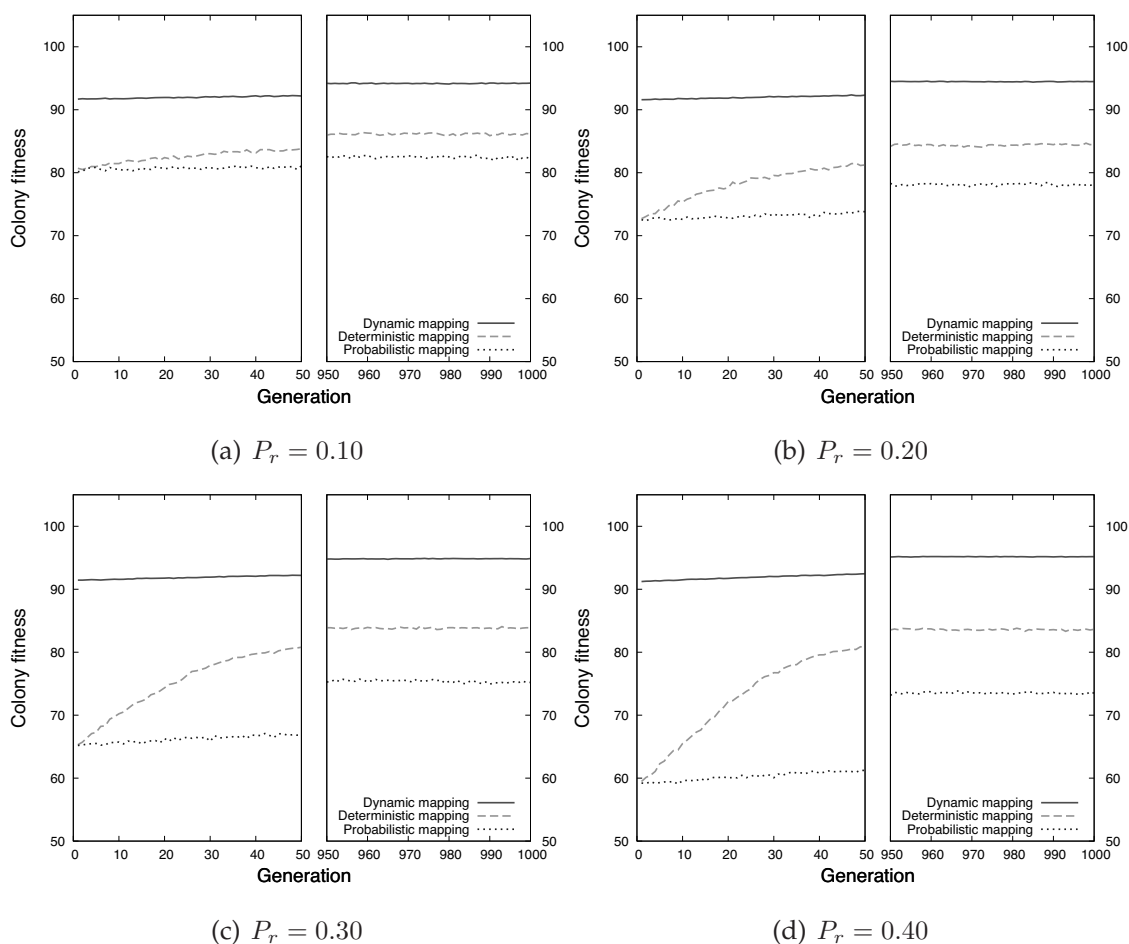


Figure 3.3: Mean fitness for colonies of unrelated individuals ( $r=0$ ) for three genetic architectures (dynamic: solid line; deterministic: dashed line; probabilistic: dotted line) for four degrees of perturbation ( $P_r = 0.10$ ,  $P_r = 0.20$ ,  $P_r = 0.30$ ,  $P_r = 0.40$ ).

	$r = 0$			$r = 1$		
	$\bar{F}$	$p$ -values		$\bar{F}$	$p$ -values	
dyn, $P_r = 0.00$	95.97	] < .001	] < .001	99.87	] < .001	] < .001
det, $P_r = 0.00$	98.14			21.97		
prob, $P_r = 0.00$	94.51			98.68		
dyn, $P_r = 0.10$	94.20	] < .001	] < .001	99.88	] < .001	] < .001
det, $P_r = 0.10$	86.11			21.95		
prob, $P_r = 0.10$	82.32			87.60		
dyn, $P_r = 0.20$	94.46	] < .001	] < .001	99.91	] < .001	] < .001
det, $P_r = 0.20$	84.50			21.96		
prob, $P_r = 0.20$	78.02			85.57		
dyn, $P_r = 0.30$	94.85	] < .001	] < .001	99.87	] < .001	] < .001
det, $P_r = 0.30$	83.83			21.93		
prob, $P_r = 0.30$	75.22			84.70		
dyn, $P_r = 0.40$	95.16	] < .001	] < .001	99.89	] < .001	] < .001
det, $P_r = 0.40$	83.55			21.95		
prob, $P_r = 0.40$	73.45			84.34		

Table 3.1: Fitness average of the last 10 generations ( $\bar{F}$ ) for the three genetic architectures in unrelated ( $r=0$ ) and highly related ( $r=1$ ) colonies with different reallocation probabilities  $P_r$ .  $P$ -values were calculated using two-sample  $t$ -tests ( $df = 18$ ).

### 3.3.2 Colonies of highly related individuals

The presence of highly related individuals ( $r=1$ ) within colonies had important consequences on the performances of the three genetic architectures. When colonies did not experience perturbations, the dynamic and probabilistic systems led to a very high fitness, while the deterministic system performed very poorly (Figure 3.4). At generation 1, the overall difference between the dynamic and probabilistic system was very small but significant, with a small advantage for the dynamic system (dyn:  $\bar{F} = 82.74$ ; dyn vs. prob:  $\bar{F} = 81.54$ ,  $t = 2.87$ ,  $df = 18$ ,  $P = 0.011$ ). The deterministic system performed poorly with a fitness significantly lower than for the two other systems ( $\bar{F} = 15.23$ , both  $P < 0.001$ ). During the first 10 generations, the difference between dynamic ( $\bar{F} = 96.22$ ) and the two other map-

ping systems was highly significant (two-sample  $t$ -tests, dyn vs. det:  $\bar{F} = 21.32$ ,  $t = 405.45$ ,  $df = 18$ ,  $P < 0.001$ ; dyn vs. prob:  $\bar{F} = 94.35$ ,  $t = 5.03$ ,  $df = 18$ ,  $P < 0.001$ ).

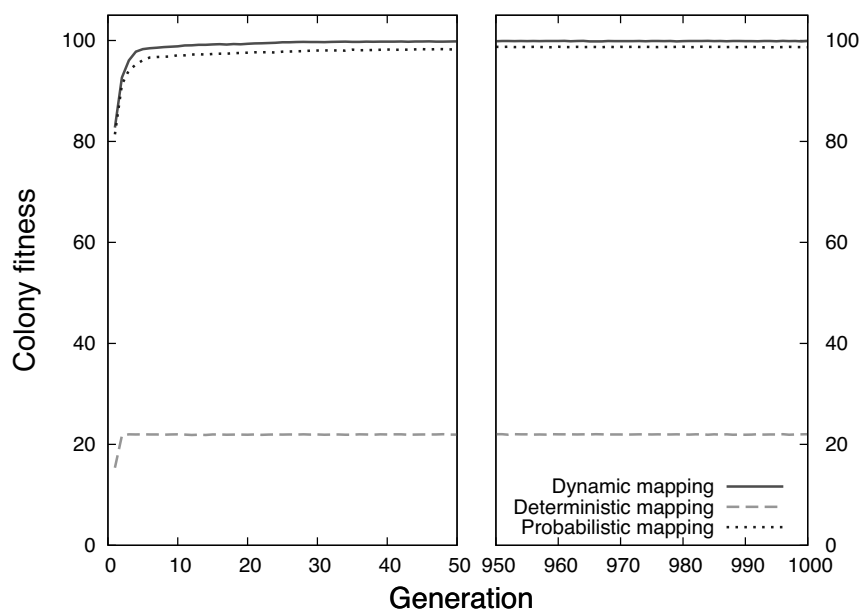


Figure 3.4: Mean fitness for colonies of highly related individuals ( $r=1$ ) for three genetic architectures (dynamic: solid line; deterministic: dashed line; probabilistic: dotted line) when there were no perturbations ( $P_r = 0.00$ ).

During the 1000 generations of selection there was an increase in colony fitness for each of the three genetic architectures (Figure 3.4). During the last 10 generations there was a significant difference in colony fitness between the three genetic architectures; the highest performances were achieved by the probabilistic and dynamic mapping systems, the latter performing slightly better than the former (Table 3.1). By contrast, the deterministic system performed poorly because all individuals within a colony had the same genome and therefore engaged in the same task.

Perturbations led to decreased fitness in the probabilistic system but not in the two other systems. For each of the four levels of perturbation, the dynamic mapping system led to an initially greater colony fitness compared to the two other mapping systems while, again, the deterministic system showed the worst performance (dyn:  $\bar{F} = 96.19/96.14/96.03/95.68$ ; det:  $\bar{F} = 21.36/21.31/21.33/21.37$ ; prob:  $\bar{F} = 83.09/80.64/79.51/77.65$  for  $P_r = 0.10/0.20/0.30/0.40$ , all  $P < 0.001$ , Figure 3.5). During the 1000 generations of selection there was an increase in fit-

ness for each of the three mapping systems. However, in the last 10 generations there were marked differences in performance with, in all cases, the dynamic system performing significantly better and the deterministic system performing worst (Figure 3.5, Table 3.1).




Figure 3.5: Mean fitness for colonies of highly related individuals ( $r=1$ ) for the three genetic architectures (dynamic: solid line; deterministic: dashed line; probabilistic: dotted line) for four degrees of perturbation ( $P_r = 0.10$ ,  $P_r = 0.20$ ,  $P_r = 0.30$ ,  $P_r = 0.40$ ).

### 3.3.3 Colonies with intermediate relatedness

At the end of the selection experiments, the performance of colonies with intermediate relatedness ( $r \simeq 0.25$ ,  $r \simeq 0.5$ ,  $r \simeq 0.75$ ) was intermediate between colonies

with unrelated ( $r=0$ ) and highly related ( $r=1$ ) individuals. This was true for each of the three genetic architectures and for each of the five levels of perturbation (Figure 3.6).

For each of the three relatedness values and each of the five levels of perturbation, the dynamic system performed better than the probabilistic system (all  $P < 0.001$ ) which itself invariably performed better than the deterministic system (Figure 3.6, all  $P < 0.001$ ). The difference in performance between the dynamic and probabilistic system was marked only when there were perturbations. By contrast, the dynamic and deterministic systems were more efficient than the deterministic system whatever the rate of perturbations.

### 3.4 Discussion

Our simulations revealed that the type of genetic architecture had very important effects on colony performance. When colonies consisted of unrelated individuals and were not subjected to perturbations, the three genetic architectures performed well, with a slight advantage for the deterministic mapping system, an intermediate performance for the dynamic and the lowest fitness for the probabilistic system. The relatively good performance of the deterministic system can be explained by the fact that when there are no perturbations and colonies comprise unrelated individuals, it is possible to select for a good ratio of genotypes specialized in each of the five castes. This is because this genetic architecture leads to a clear association between genotype and task so that frequency selection at the population level can lead to the optimal ratio of individuals specialized in each of the five tasks. The slightly lower performance of the dynamic system probably stems from the greater complexity of this system and the greater difficulty faced in keeping the optimal mix of genotypes within the population. Finally, the lowest performance of the probabilistic system can be explained by the fact that when colonies consist of completely unrelated individuals there is no benefit derived from increasing phenotypic variance. Rather, such variance leads to greater deviations from the optimal allocation of workers among the five tasks compared with the deterministic model.

Repeated perturbations had very different effects on the three genetic architectures. The performance of the dynamic mapping system was only minimally affected by the selective removal of workers performing a given task. The high



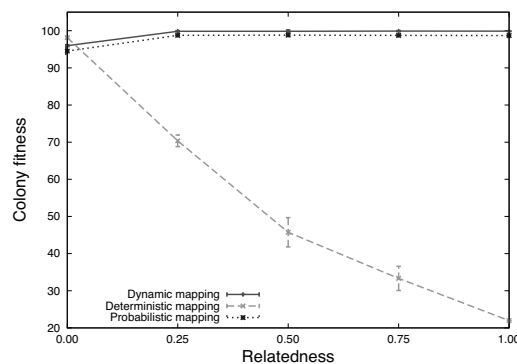

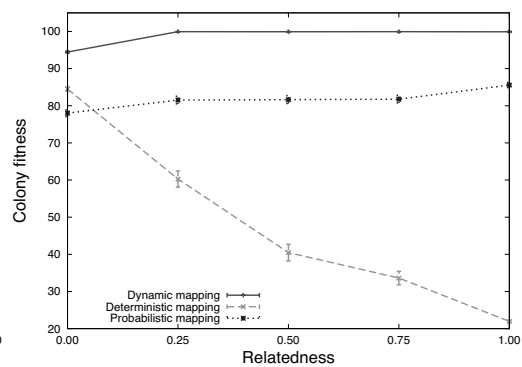


(a)  $P_r = 0.00$ (b)  $P_r = 0.10$ (c)  $P_r = 0.20$ (d)  $P_r = 0.30$ (e)  $P_r = 0.40$ 

Figure 3.6: Mean fitness values with standard deviations at generation 1000 for intermediate relatedness values (dynamic: solid line; deterministic: dashed line; probabilistic: dotted line) with five different degrees of perturbation ( $P_r = 0.00$ ,  $P_r = 0.10$ ,  $P_r = 0.20$ ,  $P_r = 0.30$ ,  $P_r = 0.40$ ).

resilience of this mapping system can be explained by the fact that when all workers doing a specific task were removed, there was a high stimulus for the replacement workers to engage in that task. As a result, colony perturbations led to almost no change in colony task allocation. By contrast, the two other genetic architectures were much more affected by the perturbations because the task choice of the replacement workers was genetically determined and independent of the type of workers removed. Hence, when workers performing a given task were removed, allocation of the new individuals almost invariably resulted in a sub-optimal number of workers performing that task. For both, the deterministic and probabilistic mapping systems, the decrease in performance was proportional to the frequency of perturbations.

Genetic relatedness significantly influenced the performance of the three genetic architectures. In contrast to the results with unrelated individuals, high relatedness ( $r=1$ ) always resulted in a very low fitness for the deterministic mapping system. This can be explained by this genetic architecture inducing all workers to perform the same task when they have an identical genome. Hence, the stable optimal solution was one that induced workers to engage in the most rewarding task and this is indeed the solution to which the selection experiments converged. This feature also explains why the performance of this mapping system was not significantly influenced by perturbations. Since the replacement workers had the same genome as the removed workers, they always performed the same job as those they were replacing, thus yielding no effect on colony fitness.

Unlike in the deterministic mapping system, high relatedness in the probabilistic and dynamic mapping systems translated into very high performances when there were no perturbations. In fact, the performance of these two systems was higher than in the situation of low relatedness. This can be explained by high colony relatedness allowing for a more efficient mode of selection of genotypes than in the situation where nest-mates were unrelated. In our experiments, the individuals that were selected to create the new generation of colonies were randomly chosen from the 30% of colonies with the highest fitness. Such a mode of selection is very efficient when nest-mates are highly related ( $r=1$ ) because individuals with a bad genome are unlikely to be selected. However, when individuals are not related, there is only very weak heritability of colony efficiency and very low selection for more efficient genomes.

Repeated perturbations led to a major decrease in fitness with the probabilis-

tic, but not with the dynamic mapping. The explanation for this difference is again because the dynamic mapping system allowed for the replacement of workers doing a particular task by new workers also performing that same task. The probabilistic mapping system on the other hand did not allow for a preferential replacement of workers doing the same task, hence resulting in a fitness decrease. Importantly, the dynamic mapping system allowed colonies to achieve fitness values close to the maximum value possible, regardless of the frequency of perturbations.

Colonies with intermediate relatedness values ( $r \simeq 0.25, 0.5$  and  $0.75$ ) invariably had performances between those of colonies with high and low relatedness. This was true for each of the three genetic architectures and for each of the five levels of perturbation. A comparison of colonies with different relatedness values showed that the performance of the deterministic system decreased sharply with increasing relatedness. The reason is again that task choice is entirely genetically determined and increased genetic similarity decreases the number of tasks performed by colony members. Thus, colonies consisting of only two types of individuals ( $r \simeq 0.5, r \simeq 0.75$ ) could only perform two out of the five tasks, while colonies with  $r \simeq 0.25$  performed the four most rewarding tasks. The effect of changes of relatedness was lower for the two other systems because they were more efficient at conducting all tasks, independent of the genetic similarity of colony members.

The dynamic system obtained fitness values close to the maximum and significantly higher than those of the probabilistic system. In both systems, increased relatedness again resulted in increased task performance, because in addition to frequency selection of specialists at the population level, relatedness also allowed the more efficient frequency selection at the colony level, with all colonies reaching near optimal solutions after 1000 generations of selection.

As in the case of high relatedness, repeated perturbations for intermediate relatedness values led to a severe performance drop for the probabilistic system and important effects for the deterministic system, but had no effect on the dynamic system because of the ability of the latter to preferentially replace missing workers.

### 3.5 Conclusion

The results of these simulations have important implications for our understanding of division of labor. First, and most importantly, they demonstrate that the underlying mechanisms responsible for the production of a behavioral phenotype from a given genotype have major effects on task partitioning and colony performance. Thus, under certain circumstances, there were up to five-fold differences in fitness between colonies, depending on the mapping system. Second, they demonstrate that the relative performance of the mapping systems varied greatly depending on the rate of colony perturbation and kin composition of colonies. Thus, while the deterministic mapping system performed very well with colonies containing unrelated individuals when there were no perturbations, it resulted in a very low fitness when colonies consisted of related individuals. Also, the rate of perturbation had important consequences on colony performance for both the deterministic and probabilistic mapping systems. By contrast, the dynamic mapping system was highly resilient to colony perturbations because it allowed for the replacement of workers performing a given task by new workers also performing this task. Our simulations also demonstrate that performance can be influenced by colony relatedness, in particular for the deterministic system. This is important, because in social insects there is important variation in relatedness both within and between species (Bourke and Franks [37], Crozier and Pamilo [57], Ross and Keller [195]).

Our model made two important assumptions. First, to simulate perturbations we removed all individuals of a randomly chosen task group. The removal of only a portion of the individuals doing a task does not qualitatively affect the results of the simulations as demonstrated by simulations where only 25, 50 or 75% of the individuals were removed (see Figure C.1 in Appendix C). Second, due to computational constraints, we conducted our experiments in colonies consisting of 100 individuals. However, additional simulations showed that colony size does not significantly change the relative performance of the three mapping systems. For example, there was very little difference in colony fitness between colonies of 100 and 10'000 individuals for each of the three genetic systems and for all of the levels of perturbation investigated (see Figure C.2 in Appendix C).

In this study we considered three genetic architectures. We selected these genetic architectures because they represent simple possibilities of how a genotype

can produce a behavioral phenotype. An important message emerging from their comparison is that high colony relatedness selects for higher phenotypic plasticity. Thus, a system such as our deterministic architecture which corresponds to a direct mapping between genotype and behavior leads to low colony fitness when nest-mates are highly related because the vast majority of individuals will engage in the same task. Under high relatedness, it is therefore better if task specialization also depends on the effects of other factors such as noise during development and age specific variations in response thresholds (Calderone and Page Jr [45], Huang and Robinson [106]).

Our simulations revealed that the dynamic system always performs better than the two others when there are perturbations. Under natural conditions, colonies of social insects frequently experience perturbations because of the selective death of individuals engaging in particular tasks or environmental changes requiring rapid adjustments of the number of workers engaged in various tasks. Hence, this should select for a genetic architecture allowing workers to respond to changes in perturbations in the distribution of individuals engaged in the different tasks and changes in colony needs. The frequency and magnitude of perturbations are probably influenced by many factors such as colony size, number of tasks performed by colony members, type of nest structure and type of food collected. Hence, it is likely that there is important variation among species in the frequency and magnitude of perturbations, possibly resulting in interspecific variation in the genetic architecture providing the best colony efficiency.

In conclusion, this study revealed that the type of mapping between genotype and individual behavior greatly influences the dynamics of task specialization and colony productivity. Our simulations also revealed complex interactions between the mode of mapping, level of within colony relatedness and frequency distribution. It is likely that the most successful mapping differs across species depending on their kin structure, risk of colony perturbation and degree of variation in colony needs.



---

# 4

## Genetic Team Composition and Level of Selection in the Evolution of Multi-Agent Systems

---

Many hands make light work.

John Heywood (English Playwright, 1497-1580)

In Chapter 2 we used evolutionary robotics as a research tool for biology. However, the application of evolutionary robotics to teams of robotic agents in particular and agent teams in general is itself an interesting and still largely unexplored research field. In addition, due to the parallels between the evolution of biological and artificial organisms, insights gained from the study of artificial and biological multi-agent systems can be mutually beneficial for both research fields.

In this chapter we are concerned with the evolution of efficient artificial agent teams. Chapter 3 discussed three important factors that can influence the efficiency of cooperation and division of labor. Here we focus our attention on the two central factors for the evolution of cooperation already mentioned in Chapter 2: Genetic relatedness and the level of selection. While both factors are intimately linked in most biological systems (see Section 1.2.2), they can be freely and independently varied in artificial systems. Here we exploit this fact to optimize the performance of artificial agent teams on three different classes of tasks. The remainder of this chapter is based on a publication currently pending peer review (Waibel et al. [223]).

## 4.1 State of the Art

Multi-agent systems (MAS) span a large number of research fields, from software agents to robotics, and play a key role in several industrial applications, such as ground and air vehicle control, supply chains or network routing. The design of control rules for multi-agent systems is challenging because agent behavior depends not only on interactions with the environment, but also on the behavior of other agents. As the number of interacting agents in a team grows, or when agent behaviors become more sophisticated, the design of suitable control rules rapidly becomes very complex. This is especially true when agents are expected to coordinate or cooperate to collectively achieve a desired task. Evolutionary computation has been advocated as an effective and promising strategy to generate control parameters and decision rules for collective agents (Baldassarre et al. [20], Nolfi and Floreano [161]).

In addition to the methodological issues of evolving agents that operate in isolation (Nolfi and Floreano [161]), the evolution of agent teams must address two major issues: (1) It must determine optimal team composition. Agents of a team may either share control rules (genetically homogeneous teams) or employ different ones (genetically heterogeneous teams). (2) It requires a suitable method for selective reproduction of desired team behavior. Selection may operate either on individuals (individual-level selection) or on teams (team-level selection). In the simplest case, one must decide between genetically homogeneous or heterogeneous teams, and between selecting agents at the individual or at the team level.

Figure 4.1 shows a sample of previous work on the evolution of multi-agent systems in robotics, combinatorial optimization, cellular automata, artificial life, genetic programming and others, plotted according to the chosen genetic team composition and level of selection. The majority of work uses genetically homogeneous teams, usually created from a cloned individual genome, with team selection. In some cases, authors created behaviorally heterogeneous agents out of a single team genome: Luke (Luke [128], Luke et al. [129]) decoded team genomes into six separate sub-teams with one or two identical players each. Other authors (Bongard [33], Hara and Nagao [94], Haynes and Sen [99], Miconi [151], Robinson and Spector [192]) decoded one team genome into different single agent genomes. Yet another approach was taken by work using distributed, embodied evolution



Genetic Team Composition	Homogeneous	[72], [156]	[3], [8], [9], [19], [20], [21], [22], [23], [24], [33], [36], [42], [43], [55], [62], [72], [87], [88], [94], [99], [100], [127], [128], [129], [130], [134], [151], [182], [185], [186], [190], [192], [199], [219], [220], [242], [245]
	Heterogeneous	[1], [2], [3], [27], [47], [68], [72], [136], [156], [182], [189], [204], [207], [208], [209], [224], [226], [229], [230]	[72], [130], [136]
		Individual	Team
Level of Selection			

Figure 4.1: A sample of approaches to the evolution of agent teams.

to evolve heterogeneous teams (Bianco and Nolfi [27], Ficici et al. [68], Simoes and Barone [204], Spector et al. [207, 208], Watson et al. [226]). In these cases selection and replication were entirely distributed among agents, with dynamics reminiscent of the replicator dynamics observed in bacterial evolution (Griffin et al. [85]) and game theoretic models (Hauert et al. [96]). In some cases, teams were evolved using a continuously updated gene-pool rather than separate gene-pools for subsequent generations (“steady state evolution”) (Reynolds [189], Werner and Dyer [229, 230]). Finally some authors have conducted more detailed comparisons of the influence of genetic team composition or level of selection alone: Martinoli (Martinoli [136]) also considered more complex methods of selection. Stanley et al. (Stanley et al. [209]) clustered genetically similar individuals into sub-teams

that shared fitness, which resulted in partially heterogeneous teams. Mirolli et al. (Mirolli and Parisi [156]) also compared partially heterogeneous teams. Quinn (Quinn [182]) evaluated individuals in different heterogeneous teams to create robust homogeneous teams.

In addition to work cited in Figure 4.1, some authors have used cooperative co-evolutionary algorithms (CCEAs; Potter and De Jong [176]) to evolve heterogeneous control rules for teams of agents (Blumenthal and Parker [28], Jim and Giles [113], Potter and De Jong [177], Wiegand et al. [236]). CCEAs are applied by decomposing problem representations into subcomponents and then creating a separate population of individuals for each subcomponent. This approach allows teams to be composed of specialized sub-groups and corresponds to the biological co-evolution of multiple species. In their basic form, CCEAs require the designer to manually decompose the multi-agent task, and thus to solve part of the optimization problem beforehand. Work that used machine learning techniques other than evolutionary computation (e.g., reinforcement learning) was not considered in this review.

Figure 4.1 suggests that the majority of current approaches to the evolution of multi-agent systems use genetically homogeneous teams evolved with team-level selection. Where the reasons for the choice of genetically homogeneous teams are made explicit, it is argued that homogeneous teams are easy to use (Baray [22], Trianni et al. [220]), require fewer evaluations (Luke et al. [129], Richards et al. [190]), scale more easily (Bryant and Miikkulainen [42]) and are more robust against the failure of team members (Bryant and Miikkulainen [42], Quinn et al. [184]) than heterogeneous teams. Many other approaches use genetically heterogeneous teams evolved with individual-level selection. Genetically heterogeneous teams are sometimes seen as providing more behavioral flexibility (Luke et al. [129]) and as allowing specialization (Baldassarre et al. [21], Bongard [34], Luke et al. [129], Quinn et al. [184]). The reasons for the choice of team-level or individual-level selection are rarely made explicit.

The terms “homogeneous team” and “heterogeneous team” used in the current literature cover many different aspects. It is important to note that while all agents in genetically homogeneous teams share the same genes, agents can nevertheless be behaviorally heterogeneous. This can happen when agents differentiate during their lifetime, for example due to varying initial conditions (Quinn et al. [185]), or due to developmental processes or learning (Yao [244]). This can

also happen when agents “activate” different parts of their genome, for example when each agent’s behavior is controlled by a different section of a single team genome (Bongard [33], Haynes and Sen [99], Miconi [151], Robinson and Spector [192]). In this case, agents can specialize on different functions, yet be genetically identical, just like specialized cells in a biological organism. Conversely, it is important to note that genetically heterogeneous teams are those in which agents are, on average, not genetically more similar to team members than to agents in the rest of the population (Queller [180], West et al. [234]). This means that teams resulting from embodied evolution or common versions of steady state evolution are usually genetically heterogeneous although these algorithms often generate multiple offspring from a single parent, resulting in genetically similar (but not identical) agents. In some cases, teams consist of clonal sub-teams (Luke [128], Luke et al. [129]) or of agents that share only part of their genome. Teams with agents that are, on average, genetically more similar (but not identical) to members of their team than to members of the rest of the population are termed “partially heterogeneous”. The effects of partial genetic heterogeneity on the evolution of agent teams are not yet fully explored in evolutionary computation (Mirolli and Parisi [156]), but have been deeply studied in biology (Hamilton [91], Lehmann and Keller [125]).

The choice of level of selection is rarely discussed explicitly. Some research has addressed the related issue of credit assignment for the evolution of multi-agent systems (Agogino and Tumer [4], Martinoli [136]). In the context of multi-agent systems, credit assignment is concerned with distributing fitness rewards among individual agents. Fitness distribution leads to credit assignment problems (Grefenstette [83], Minsky [155]) in many cooperative multi-agent tasks, because individual contributions to team performance are often difficult to estimate or difficult to monitor (Panait and Luke [171]). Selection is usually performed on the basis of accumulated individual or team fitness, which may be the result of many fitness rewards with different types of credit assignment. Therefore an optimal choice of level of selection is not only influenced by the type of task but also by the types of credit assignment used.

Genetic team composition and level of selection have long been identified as two important factors for the evolution of biological agent teams such as groups of genes, cells, individuals or other replicators (Hamilton [91], Keller [118]). In particular the evolution of altruism (West et al. [231]), in which agents cooperate

to increase team fitness in spite of an individual fitness cost to the cooperator, has received a lot of attention (Hamilton [92], Lehmann and Keller [125]). Here we define cooperation as a behavior that increases the fitness of other agents, and altruistic cooperation (altruism) as a behavior that increases the fitness of other agents and decreases the cooperator's fitness.

In this study, we compare the performance of robot teams evolved in four evolutionary conditions: genetically homogeneous teams evolved with team-level selection; genetically homogeneous teams evolved with individual-level selection; genetically heterogeneous teams evolved with team-level selection; and genetically heterogeneous teams evolved with individual-level selection. We evaluate the performance of robot teams evolved in these four evolutionary conditions for three classes of multi-robot tasks: a task that does not require cooperation; a task that requires cooperation but does not imply a cost for cooperators; and a task that requires altruistic cooperation, i.e., a task that implies an individual fitness cost for cooperators.

## 4.2 Evolutionary Conditions

The four possible combinations of genetic team composition and level of selection were formalized into four evolutionary algorithms (Figure 4.2). For the remainder of the paper we will use the terms "homogeneous" and "heterogeneous" to designate genetically homogeneous and genetically heterogeneous teams, respectively, and the terms "individual selection" and "team selection" to designate teams evolved with individual-level selection and team-level selection, respectively. We considered populations composed of  $M$  teams, each composed of  $N$  individuals. Population size and team sizes were kept constant across generations. At each generation, the old population was entirely replaced by a new population of offspring. Individuals' genomes were binary strings.

### **Algorithm 1 - Homogeneous teams, Individual selection**

Each of the  $M$  teams at generation 0 was formed by generating one random genome and cloning it  $N - 1$  times to obtain  $N$  identical robot genomes (clones) per team. Teams were evaluated in the task and an individual fitness determined for each of the  $N$  robots. For a new generation, each of the  $M$  new teams was

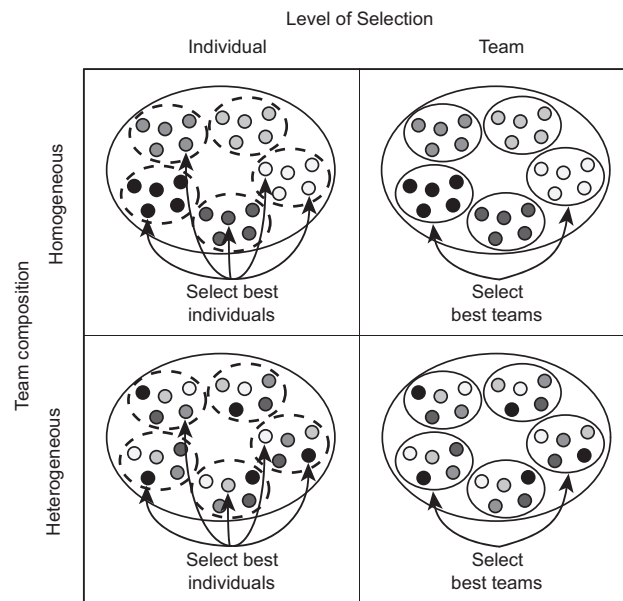


Figure 4.2: The four evolutionary conditions. A population (large oval) was composed of several teams (medium ovals), each of which was composed of several robots (small circles) evaluated together. Genetic team composition was varied by either composing teams of robots with *identical genomes* (homogeneous, identical shading), or *different genomes* (heterogeneous, different shading). Level of selection was varied by selecting *teams* (team selection), or selecting *individuals*, independent of their team affiliation (individual selection).

---

**Algorithm 1** Homogeneous teams, Individual selection
 

---

```

for each of  $M$  new teams do
  select two individuals from all old teams
  recombine their genomes to create one new genome
  mutate new genome
  clone new genome to obtain  $N$  genomes for new team
end for

```

---

created from two *individuals* selected among all individuals of all old teams in the population using roulette wheel selection. The two genomes of the selected individuals were recombined (crossover probability of 0.05) to produce one new genome. The resulting new genome was mutated by flipping the value of each bit with a probability of 0.05 and then cloned  $N - 1$  times to generate the  $N$  robot genomes of the new team. Teams evolved using this evolutionary condition were thus genetically *homogeneous*.

---

**Algorithm 2** Homogeneous teams, Team selection
 

---

```

for each of  $M$  new teams do
  select two old teams
  recombine their genomes to create one new genome
  mutate new genome
  clone new genome to obtain  $N$  genomes for new team
end for

```

---

**Algorithm 2 - Homogeneous teams, Team selection**

Each of the  $M$  teams at generation 0 was formed by generating one random genome and cloning it  $N - 1$  times to obtain  $N$  identical robot genomes (clones) per team. Teams were evaluated in the task, and for each team, a team fitness was determined as the sum of the individual fitnesses of all  $N$  robots. For a new generation, each of the  $M$  new teams was created from two old *teams* selected using roulette wheel selection. The two genomes of the selected teams were recombined (crossover probability of 0.05) to produce one new genome. The resulting new genome was mutated by flipping the value of each bit with a probability of 0.05 and then cloned  $N - 1$  times to obtain the  $N$  robot genomes of the new team. Teams evolved using this evolutionary condition were thus genetically *homogeneous*.

**Algorithm 3 - Heterogeneous teams, Individual selection**

Each of the  $M$  teams at generation 0 was formed by generating  $N$  random genomes. Teams were evaluated in the task and an individual fitness determined for each of the  $N$  robots. For a new generation, each of the  $N \times M$  new individuals was created from two *individuals* selected among all individuals of all old teams in

---

**Algorithm 3** Heterogeneous teams, Individual selection

---

```

for each of  $M$  new teams do
  for each of  $N$  new team members do
    select two individuals from all old teams
    recombine their genomes to create one new genome
    mutate new genome
    add new genome to new team
  end for
end for

```

---

the population using roulette wheel selection. The two genomes of the selected individuals were recombined (crossover probability of 0.05) to produce one new genome. The resulting new genome was mutated by flipping the value of each bit with a probability of 0.05. This process was repeated  $N \times M - 1$  times to form  $M$  new teams of  $N$  individuals each. In this evolutionary condition robots were not, on average, genetically more similar to team members than to robots in the rest of the population, and thus teams were genetically *heterogeneous*.

---

**Algorithm 4** Heterogeneous teams, Team selection

---

```

for each of  $M$  new teams do
  for each of  $N$  new team members do
    select two old teams
    randomly select two old team members
    recombine their genomes to create one new genome
    mutate new genome
    add new genome to new team
  end for
end for

```

---

**Algorithm 4 - Heterogeneous teams, Team selection**

Each of the  $M$  teams at generation 0 was formed by generating  $N$  random genomes. Teams were evaluated in the task, and for each team, a team fitness was determined as the sum of the individual fitnesses of all  $N$  robots. For a new generation, each of the  $N \times M$  individuals was created from two old *teams* selected using roulette wheel selection. Two genomes, each randomly selected among the mem-

bers of a selected team, were recombined (crossover probability of 0.05) to produce one new genome. The resulting new genome was mutated by flipping the value of each bit with a probability of 0.05. This process was repeated  $N \times M - 1$  times to form  $M$  new teams of  $N$  individuals each. In this evolutionary condition robots were not, on average, genetically more similar to team members than to robots in the rest of the population, and thus teams were genetically *heterogeneous*.

## 4.3 Experimental Method

### 4.3.1 Scenario

The experimental setup (Figure 4.3) consisted of a  $50 \times 50$  cm<sup>2</sup> arena with 10 micro-robots and two types of tokens, small and large. We chose to study a foraging task, because foraging combines several aspects of multi-agent tasks (distributed search, coordinated movement, transportation) and relates to many real-world

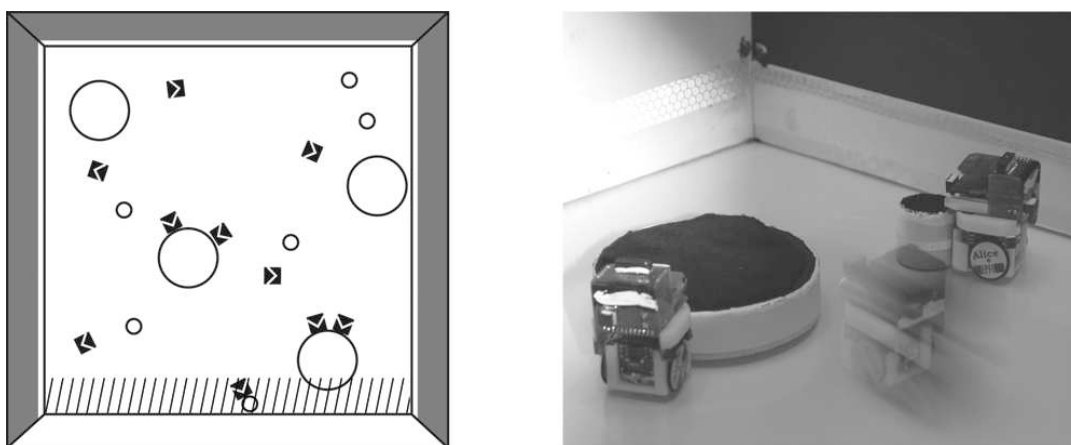


Figure 4.3: Left: The experimental setup for task 3, the altruistic cooperative foraging task. Ten micro-robots (black squares with arrows) searched for small and large tokens and transported them to the target area (hatched area at bottom) under the white wall (the other three walls were black). An identical setup was used in the other two tasks, except that the arena contained either only small tokens in task 1, or only large tokens in task 2. Right: Three micro-robots in task 3, the altruistic cooperative foraging task. The robot in the background could transport the small token by itself. The robot at the left could not transport the large token by itself and needed to wait for the arrival of a second robot.



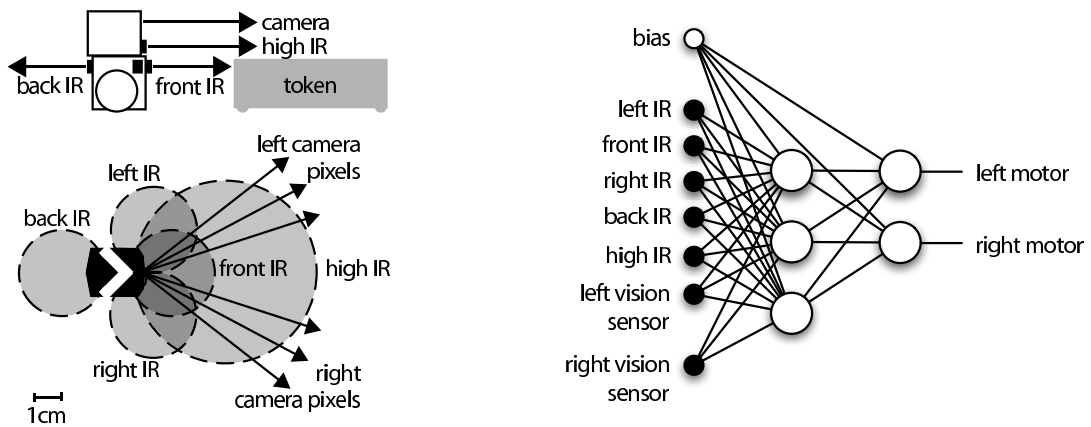


Figure 4.4: Left: Side and top-view schematics of a simulated micro-robot. The robot was equipped with four infrared (IR) distance sensors (three at the front, one at the back) to detect tokens, and a camera to identify the target area. A fifth infrared distance sensor (high IR) was mounted higher on the robot and thus overlooked tokens. This allowed robots to distinguish tokens from walls and other robots. Right: The neural network architecture, a feed-forward neural network with a single layer of three hidden neurons. Inputs were given by the activation values of five infrared (IR) sensors and two vision sensors with activation values computed from left and right camera pixels (see text).

problems (Balch [18], Cao et al. [48]). In addition, foraging is a wide-spread and well-studied behavior of many biological societies (Krieger et al. [124], Reyes Lopez [188], Traniello [217]).

Robots foraged tokens by transporting them into a 4 cm wide region at one side of the arena marked by a white wall. A single robot was sufficient to transport a small token. At least two robots were required to transport a large token, thus retrieval of large tokens required cooperation. Cooperating agents had to coordinate their behaviors to successfully align their positions before and during token transport.

The micro-robots (Caprari [49]) were small ( $2 \times 2 \times 4 \text{ cm}^3$ ), two-wheeled robots equipped with three infrared distance sensors at the front and one at the back, which could sense objects up to 3 cm away (Figure 4.4 left). An extension module with a fourth infrared distance sensor with a range of up to 6 cm and a linear camera were mounted higher on the robot, overlooking tokens but sensitive to other robots and walls.

### 4.3.2 Control and genetic architecture

Robots were controlled by a feed-forward neural network with a single layer of three hidden neurons (Figure 4.4 right) and a sigmoid activation function (*tanh*). The inputs were given by the activation values of five infrared sensors, two vision sensors, and a constant bias value of  $-1$ . Infrared sensor activation values were scaled in the range  $[-1; 1]$ . Vision sensors were an average of three equi-distally spread camera pixels spanning a field of view of  $18^\circ$ , for the left or right side of the image, respectively. The averages were thresholded to yield 0 for a white or 1 for a black arena wall. Using the average value of three pixels rather than a single pixel allowed a robust detection of the white foraging target area in spite of the presence of other robots in the field of view. The two output units were used to control the left and right wheel motors. The activation values in the range  $[-1; 1]$  were mapped into speeds in the range  $[-4; 4]$  cm/s, with speeds in the interval of  $[-2.5; 2.5]$  cm/s set to 0 because of unreliable motor response at low speeds.

The neural network connection weights were in the range of  $[-2; 2]$  and coded on 8 bits. The genome of one individual was thus  $8 \times 32$  bits long.

### 4.3.3 Collective tasks

We devised three types of foraging tasks that differed in the amount of cooperation required from agents.

#### Task 1 - Individual foraging

The arena contained 6 small tokens, which each awarded 1 fitness point to the foraging robot. This task did not require cooperation, because a single agent was sufficient to transport a small token.

#### Task 2 - Cooperative foraging

The arena contained 4 large tokens, which each awarded 1 fitness point to each team member, irrespective of its participation in the token foraging. This corresponded to a situation where the individual contributions to team performance were not known, i.e., a situation with credit assignment problems (Grefenstette [83], Minsky [155]), which is the case for many cooperative multi-agent tasks

(Panait and Luke [171]). This task required cooperation because it could not be accomplished by a single agent.

### **Task 3 - Altruistic cooperative foraging**

The arena contained 6 small and 4 large tokens. Small tokens each awarded 1 fitness point to the foraging robot and large tokens each awarded 1 fitness point to each team member, irrespective of their participation in the token foraging. In this task cooperation was costly for individuals, because individuals that did not cooperate always had higher fitness than their cooperating team mates. This meant that cooperators suffered a relative individual fitness cost and therefore this task required altruistic cooperation (Lehmann and Keller [125]).

#### **4.3.4 Evolutionary experiments**

Due to the large number of evaluations required for the evolution of robot behaviors, all evolutionary experiments were conducted using a physics-based 2D simulator (Magenat and Waibel [132]), which is available as part of an open evolutionary framework (Magenat et al. [133]). All simulation parameters, including robot size, shape, speed and weight, as well as collision dynamics, friction forces and sensor and actuator modalities, were based on the micro-robots described in Section 4.3.1.

We evolved teams of robots under the four evolutionary conditions separately for each of the three tasks, making a total of 12 experimental lines. Evolutionary experiments lasted for 300 generations. Twenty independent runs were performed for each experimental line. Populations consisted of 100 teams of 10 agents each. Each team was evaluated 10 times for three minutes with random token and robot starting positions and orientations. Fitness was averaged over the 10 evaluations.

To compare the efficiency of the four evolutionary conditions, we re-evaluated the best teams at generation 300 for 1000 times and compared their team fitness. Since the distributions of fitness values were unknown fitness differences were analyzed using the non-parametric Wilcoxon rank sum test. All fitness values were normalized for each task, with 0 being the minimal possible fitness and 1 the theoretical maximum value.

## 4.4 Results

### 4.4.1 Task 1 - Individual foraging

Successful foraging behavior evolved for all four evolutionary conditions (Figure 4.5). After 300 generations of artificial evolution, heterogeneous teams evolved with individual selection collected all 10 tokens in most evaluations and achieved fitness values close to the maximum value achievable. These fitness values were higher than those of homogeneous teams evolved with individual selection and homogeneous teams evolved with team selection (Wilcoxon rank sum test,  $df = 38$ ,  $P < 0.001$  and  $P < 0.006$ , respectively). Performance of homogeneous teams evolved with individual selection and homogeneous teams evolved with team selection did not differ significantly ( $P = 0.337$ ). Heterogeneous teams evolved with team selection performed significantly worse than all other evolutionary conditions (all three  $P < 0.002$ ).

### 4.4.2 Task 2 - Cooperative foraging

Successful foraging behavior evolved for all four evolutionary conditions (Figure 4.6). The experiments with a cooperative task led to a change in the relative performance of the four evolutionary conditions. The highest performance was now achieved by homogeneous teams evolved with individual selection and homogeneous teams evolved with team selection. Performance of homogeneous teams evolved with individual selection and homogeneous teams evolved with team selection did not differ significantly ( $P = 0.839$ ), but was significantly higher than that of heterogeneous teams evolved with individual selection and heterogeneous teams evolved with team selection (all four  $P < 0.001$ ). Performance of heterogeneous teams evolved with individual selection and heterogeneous teams evolved with team selection did not differ significantly ( $P = 0.365$ ).

### 4.4.3 Task 3 - Altruistic cooperative foraging

Successful foraging behaviors evolved for all four evolutionary conditions (Figure 4.7). Homogeneous teams achieved significantly higher fitness values than heterogeneous teams (all four  $P < 0.001$ ).

Performance of homogeneous teams evolved with individual selection and

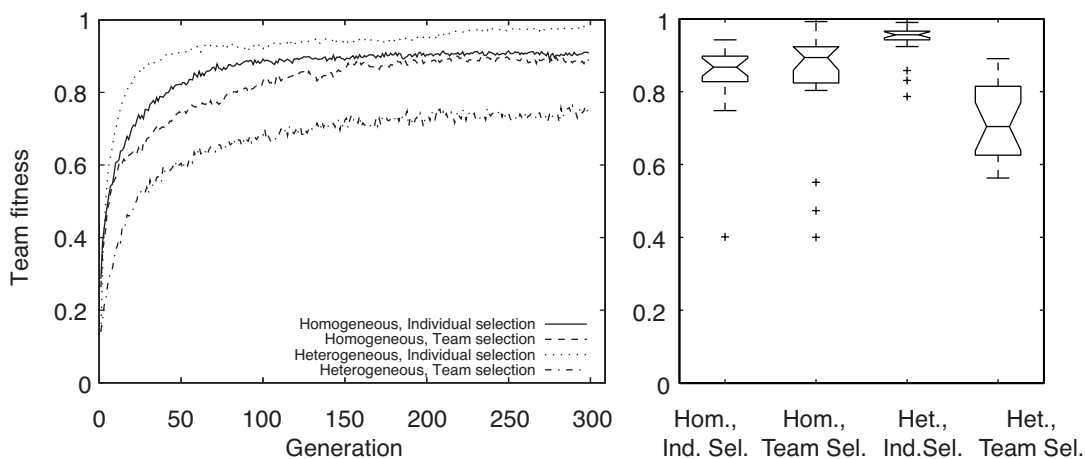


Figure 4.5: Task 1 - Individual foraging. Left: Evolution of the best team fitness averaged over the best teams in 20 independent evolutionary runs over 300 generations. Right: The best team at generation 300 of each of the 20 independent experiments per evolutionary condition and per task was evaluated 1000 times. The mid line in the box is the median, while the box represents the upper and lower quartile above and below the median. The bars outside the box generally represent the max and min values, except when there are outliers, which are shown as crosses. We define outliers as data points which differ more than 1.5 times the interquartile range from the border of the box. The notches represent the uncertainty in the difference of the medians for box-to-box comparison. Boxes whose notches do not overlap indicate that the medians differ at the 5 % significance level (McGill et al. [148]). In this task, which did not require cooperation, heterogeneous teams evolved with individual selection performed best, followed by homogeneous teams evolved with individual selection and homogeneous teams evolved with team selection. Heterogeneous teams evolved with team selection performed significantly worse than all other evolutionary conditions.

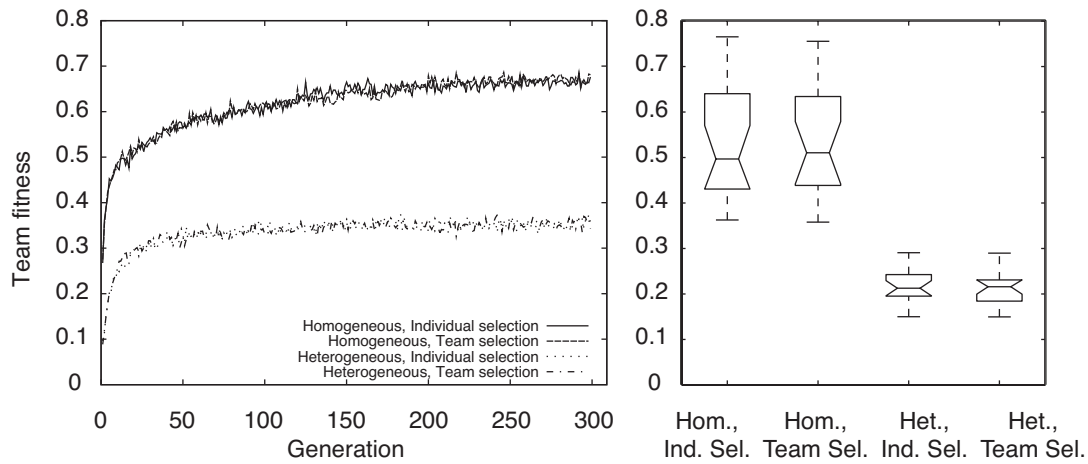


Figure 4.6: Task 2 - Cooperative foraging. Left: Evolution of the best team fitness averaged over the best teams in 20 independent evolutionary runs over 300 generations. Right: The best team at generation 300 of each of the 20 independent experiments per evolutionary condition and per task was evaluated 1000 times. Homogeneous teams performed significantly better than heterogeneous teams. Boxplot explanations see Figure 4.5.

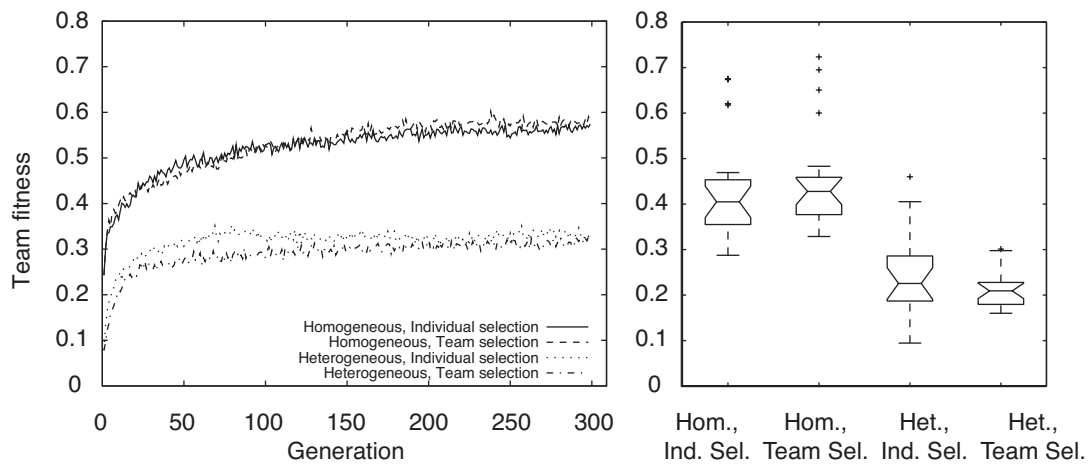


Figure 4.7: Task 3 - Altruistic cooperative foraging. Left: Evolution of the best team fitness averaged over the best teams in 20 independent evolutionary runs over 300 generations. Right: The best team at generation 300 of each of the 20 independent experiments per evolutionary condition and per task was evaluated 1000 times. Homogeneous teams performed significantly better than heterogeneous teams. Boxplot explanations see Figure 4.5.

homogeneous teams evolved with team selection did not differ significantly ( $P = 0.310$ ). Performance of heterogeneous teams evolved with individual selection and heterogeneous teams evolved with team selection did not differ significantly ( $P = 0.490$ ).

However, the four evolutionary conditions resulted in different foraging strategies in this task (Figure 4.8): While homogeneous teams evolved with individual selection and homogeneous teams evolved with team selection as well as heterogeneous teams evolved with team selection collected a significantly higher proportion of large tokens than small tokens (all three  $P < 0.001$ ), heterogeneous teams evolved with individual selection collected a significantly higher proportion of small tokens than large tokens ( $P < 0.001$ ). In comparison to the other three evolutionary conditions, heterogeneous teams evolved with individual selection collected the significantly highest proportion of small tokens (all three  $P < 0.001$ ), but the significantly lowest proportion of large tokens of all four evolutionary conditions (all three  $P < 0.003$ ).

## 4.5 Discussion

The three types of tasks resulted in significant performance differences for teams evolved under the four evolutionary conditions.

### 4.5.1 Task 1 - Individual foraging

In the individual foraging task, heterogeneous teams evolved with individual selection led to the best team performance. A possible reason could be the disparities in genome evaluation in homogeneous and heterogeneous teams. For a team size of  $N$  agents, heterogeneous teams evaluated  $N$  times more genomes than homogeneous teams. This was because each heterogeneous team consisted of  $N$  different genomes, whereas homogeneous teams consisted of  $N$  identical genomes. On the other hand, homogeneous teams evaluated each genome  $N$  times more often than heterogeneous teams. This was because each team evaluation evaluated an identical genome  $N$  times. Our results suggest that higher evaluation accuracy may have been less important than a larger number of different genomes in this task. The larger number of genomes may have allowed heterogeneous teams to discover solutions faster than homogeneous teams, which could explain the steep




Figure 4.8: Task 3 - Altruistic cooperative foraging. The plot shows the average proportion of the six small tokens and four large tokens collected by the best teams at generation 300 for each of the 20 independent experiments and for each of the four evolutionary conditions. Heterogeneous teams evolved with individual selection pursued a different foraging strategy than teams of the other three evolutionary conditions, collecting very few large tokens but most small tokens.

initial fitness increase. It may also have allowed heterogeneous teams to discover better solutions than homogeneous teams, which could explain the higher final fitness obtained with this evolutionary condition. To test whether these disparities in genome evaluation caused the high team performance of heterogeneous teams evolved with individual selection we performed a set of additional experiments (see below).

Homogeneous teams evolved with team selection and homogeneous teams evolved with individual selection performed similarly. This was because with roulette wheel selection, the probability of a team to be selected was the same as the sum of the probabilities of each individual team member to be selected. Since all team members of homogeneous teams shared the same genome, selection probabilities for a given genome were equal for both homogeneous evolutionary conditions. It should be noted, however, that this is not necessarily true for other types of selection. Selection mechanisms where the fitness of a genome is not di-



rectly proportional to its probability to be selected (e.g., truncation or rank-based selection) may lead to differences in the number of selected individuals with a given genotype and consequently affect the relative performance of homogeneous teams evolved with individual selection and homogeneous teams evolved with team selection. In these cases individual selection may select for genomes that lead to higher maximum but lower average individual performance.

Heterogeneous teams evolved with team selection performed worse than all other evolutionary conditions. This was because, unlike all other three evolutionary conditions, this evolutionary condition did not allow a direct link between the performance of a genome and its probability to be selected. Instead, selection of good genomes could only happen indirectly, by selecting those teams that contained a better mix of genomes than other teams. Since good genomes could be part of bad teams or bad genomes part of good teams, selection for good individual genomes was inefficient. This explains the slow initial fitness increase and the lowest final fitness of heterogeneous teams evolved with team selection.

To test the hypothesis that the high team performance of heterogeneous teams evolved with individual selection was caused by disparities in genome evaluation, we performed a set of additional experiments. First, we evolved homogeneous teams in the same task, but used only 1 evaluation per team rather than 10 evaluations (Figure 4.9 (a)). Second, we evolved heterogeneous teams in the same task, but used only 100 agents per population rather than 1000 agents (Figure 4.9 (b)). In this set of experiments homogeneous and heterogeneous teams therefore evaluated the same number of genomes and had the same number of evaluations per genome.

Without disparities in genome evaluation, heterogeneous teams evolved with individual selection performed similarly to homogeneous teams evolved with individual selection and homogeneous teams evolved with team selection (all three  $P > 0.597$ ). Heterogeneous teams evolved with team selection performed worse than all other evolutionary conditions (all three  $P < 0.001$ ), because the efficiency of selection was not affected by the changes in genome evaluation.

## 4.5.2 Task 2 - Cooperative foraging

In the cooperative foraging task, homogeneous evolutionary conditions dramatically outperformed heterogeneous evolutionary conditions, with the best fitness




Figure 4.9: Task 1 - Individual foraging without disparities in genome evaluation. (a) Homogeneous teams evolved with 1 evaluation per team (instead of 10) and (b) heterogeneous teams evolved with 100 agents per population (instead of 1000). Heterogeneous teams evolved with individual selection performed similarly to homogeneous teams evolved with individual selection and homogeneous teams evolved with team selection in this task. Boxplot explanations see Figure 4.5.

values in homogeneous teams up to 70% higher than in heterogeneous teams.

Heterogeneous teams evolved with individual selection performed poorly in this task. A possible reason are disparities in genome evaluation between homogeneous and heterogeneous teams. Another possible reason is that selection of good genomes could only happen indirectly in this task, which may have led to inefficient selection just as in heterogeneous teams evolved with team selection in task 1 (Section 4.5.1). This could have been because fitness in this task was assigned to all team members, irrespective of their participation in the token foraging.

Homogeneous teams evolved with individual selection and homogeneous teams evolved with team selection performed similarly, because selection proba-

bilities for a given genome were again equal for both homogeneous evolutionary conditions.

Heterogeneous teams evolved with team selection performed poorly in this task. This was again because selection was inefficient.

To test the hypothesis that the differences in performance of heterogeneous teams evolved with individual selection and homogeneous teams evolved with individual selection and team selection were caused by disparities in genome evaluation or by the fitness assignment to all team members, we performed two sets of additional experiments. First, we again corrected for the disparities in genome evaluation. However, correcting for this factor alone did not eliminate the performance differences (see Figure D.1 in Appendix C). Second, we performed experiments where we again corrected for the disparities in genome evaluation and where fitness was only assigned to team members that participated in the token foraging. In these experiments, each of the 4 large tokens awarded 5 fitness points to each of the two transporting robots, rather than 1 fitness point to each of the 10 team members. This second additional set of experiments therefore corresponded to a situation where the individual contributions to team performance were known, i.e., a situation without credit assignment problems.

Without the disparities in genome evaluation and without credit assignment problems, heterogeneous teams evolved with individual selection outperformed heterogeneous teams evolved with team selection ( $P < 0.001$ ). This was because selection of good genomes could now happen directly, which allowed for efficient selection. However, the performance of heterogeneous teams evolved with individual selection remained lower than that of homogeneous teams evolved with individual selection and homogeneous teams evolved with team selection ( $P < 0.001$  and  $P < 0.002$ , respectively, Figure 4.10). A possible reason is that heterogeneous teams had to solve a more complex optimization task than homogeneous teams. Successful cooperation in heterogeneous teams required individuals to evolve behaviors to coordinate their actions with  $N - 1$  different team members, while individuals in homogeneous teams only had to evolve behaviors to coordinate with a single type of team member. In other words, homogeneous teams led to a smaller search space because all team members were per definition identical, and thus only a subset of the total number of possible team compositions was considered in these teams. Furthermore, individuals in heterogeneous

teams were not just different in a team, but team members changed from one generation to the next. Both factors may have hindered the evolution of cooperative behavior in heterogeneous teams.

The performance of homogeneous teams evolved with individual selection and homogeneous teams evolved with team selection did not differ significantly ( $P = 0.441$ ) in this second additional set of experiments.

Heterogeneous teams evolved with team selection performed worse than all other evolutionary conditions due to inefficient selection (all three  $P < 0.001$ ).




Figure 4.10: Task 2 - Cooperative foraging without disparities in genome evaluation and without credit assignment problems. (a) Homogeneous teams evolved with 1 evaluation per team (instead of 10) and (b) heterogeneous teams evolved with 100 agents per population (instead of 1000). The performance of heterogeneous teams evolved with individual selection was higher than the performance of heterogeneous teams evolved with team selection, but did not reach that of homogeneous teams in this task. Boxplot explanations see Figure 4.5.

### 4.5.3 Task 3 - Altruistic cooperative foraging

As for the previous two tasks, successful foraging behavior evolved for all four evolutionary conditions. Team performance in the altruistic cooperative foraging task was systematically lower than in the cooperative foraging task. This may seem surprising because the larger number of tokens in the arena increased the total number of fitness points available. A possible reason is that the increased number of tokens led to more clutter in the arena, which made successful token transport more difficult.

Heterogeneous teams evolved with individual selection again performed poorly compared to homogeneous teams. Possible reasons are disparities in genome evaluation and inefficient selection for the foraging of large tokens because fitness points gained from large tokens were assigned to all team members, irrespective of their participation in the token foraging.

Homogeneous teams evolved with individual selection and homogeneous teams evolved with team selection performed similarly, because selection probabilities for a given genome were again equal for both homogeneous evolutionary conditions.

Heterogeneous teams evolved with team selection performed poorly in this task because of inefficient selection.

To test the hypothesis that the differences in performance of heterogeneous teams evolved with individual selection were caused by disparities in genome evaluation or by the fitness assignment to all team members, we performed two sets of additional experiments similar to those described in Section 4.5.2. First, we again corrected for the disparities in genome evaluation. However, correcting for this factor alone did not eliminate the performance differences (see Figure D.3 in Appendix C). Second, we again performed experiments where we corrected for the disparities in genome evaluation and for credit assignment problems.

Without the disparities in genome evaluation and without credit assignment problems heterogeneous teams evolved with individual selection outperformed heterogeneous teams evolved with team selection ( $P < 0.001$ ). This was because selection of good genomes could again happen directly, which allowed for efficient selection. However, the performance of heterogeneous teams evolved with individual selection remained lower than that of homogeneous teams evolved with individual selection and homogeneous teams evolved with team selection

( $P < 0.015$  and  $P < 0.003$ , respectively, Figure 4.11). This may have been because heterogeneous teams had to solve a more complex optimization task.

In this second additional set of experiments, the performance of homogeneous teams evolved with individual selection and homogeneous teams evolved with team selection did not differ ( $P = 0.133$ ).

Heterogeneous teams evolved with team selection performed worse than all other evolutionary conditions due to inefficient selection (all three  $P < 0.001$ ).




Figure 4.11: Task 3 - Altruistic cooperative foraging without disparities in genome evaluation and without credit assignment problems. (a) Homogeneous teams evolved with 1 evaluation per team (instead of 10) and (b) heterogeneous teams evolved with 100 agents per population (instead of 1000). The performance of heterogeneous teams evolved with individual selection was higher than the performance of heterogeneous teams evolved with team selection, but did not reach that of homogeneous teams. Boxplot explanations see Figure 4.5.

Importantly, the altruistic cooperative foraging task led to the evolution of a different foraging strategy in heterogeneous teams evolved with individual selection than in the other three evolutionary conditions (Figure 4.8). A possible reason is that cooperation to collect large tokens now implied a cost for individuals.

To test this hypothesis we performed additional experiments with this evolutionary condition. First we repeated the experiments with a setup identical to that of task 3, i.e., with 1000 agents per population and 10 evaluations per team, but with known individual contributions to large token foraging, i.e., a situation without credit assignment problems. Each of the 4 large tokens awarded 5 fitness points to each of the two transporting robots, rather than 1 fitness point to each of the 10 team members. Then, at generation 300, we changed the fitness assignment and assumed unknown individual contributions to large token foraging, i.e., a situation with credit assignment problems. Each of the 4 large tokens awarded 1 fitness point to each team member, irrespective of their participation in token foraging.

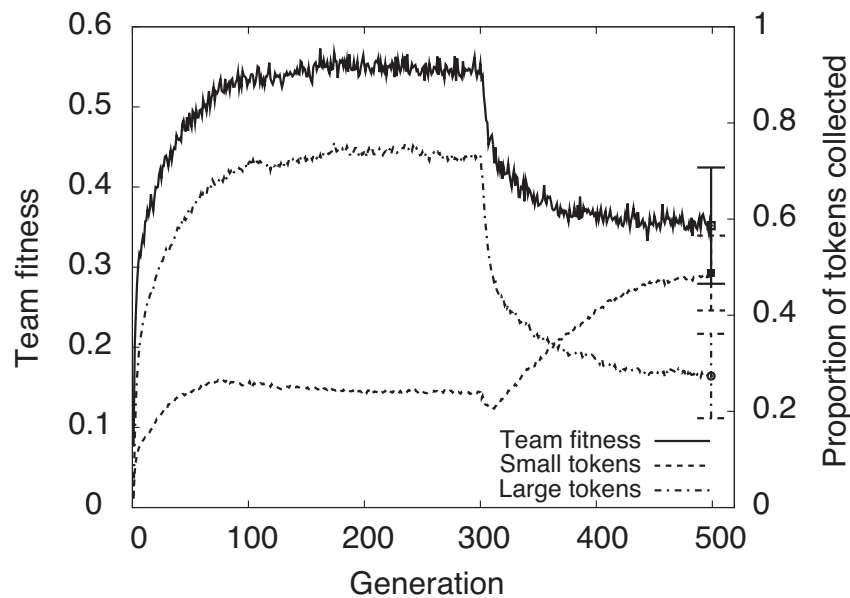


Figure 4.12: Task 3 - Altruistic cooperative foraging in heterogeneous teams evolved with individual selection. For the first 300 generations, individual contributions to the cooperative foraging of large tokens were known (no credit assignment problems). From generation 300 onward individual contributions to the foraging of large tokens were presumed unknown (credit assignment problems). The introduction of credit assignment problems led to the rapid collapse of cooperation and a decrease in team fitness.

This change in fitness assignment resulted in a drastic change in foraging strategy (Figure 4.12). While at generation 300 heterogeneous teams evolved with individual selection collected a significantly higher proportion of large tokens

than small tokens ( $P < 0.001$ ), at generation 500 they collected a significantly lower proportion of large than small tokens ( $P < 0.001$ ). As a direct result of this change, team performance decreased significantly ( $P < 0.001$ ) between generation 300 and generation 500. This was because after the introduction of credit assignment problems, fitness points gained from large tokens were assigned to all team members, and therefore individuals collecting small tokens gained a fitness advantage over their team mates. This led to the selection of individuals that foraged for small tokens and resulted in fewer and fewer individuals foraging for large tokens. The observed drop in team fitness also implies a drop in average individual fitness. This illustrates that fitness is a relative measure of performance and therefore evolution selects for performance increase relative to the performance of competitors, rather than for absolute performance. The simplicity of the neural network controllers did not allow individuals to accurately discriminate large and small tokens, which explains the incomplete collapse of large token foraging.

In contrast, the foraging strategy in homogeneous teams evolved with individual selection and in homogeneous teams evolved with team selection was not affected by the costs implied in large token foraging (see Figures D.4 and D.5 in Appendix C). This was because relative fitness differences between team members could not have an influence on the selection of genomes when individuals were genetically identical.

Foraging strategy in heterogeneous teams evolved with team selection was not affected by the costs implied in large token foraging (see Figures D.4 and D.5 in Appendix C). This was because relative fitness differences between team members did not have an influence on selection of genomes when selection acted at the level of the team.

## 4.6 Conclusion

This study provides an experimental demonstration of how the choice of genetic team composition and level of selection influences the performance of multi-agent systems. We have identified three different classes of multi-agent tasks depending on the amount of cooperation required between team members. Our results demonstrate that different combinations of genetic team composition and level of selection lead to significant performance differences. No combination



achieved optimal performance in all three classes of tasks.

The main results of this study are summarized in Table 4.1. The symbol “√” indicates a suitable method, “○” a method with potential drawbacks and “×” a method that can not be recommended.

Table 4.1: Suggested guidelines for the choice of genetic team composition and level of selection

	Hom., Ind. S.	Hom., Team S.	Het., Ind. S.	Het., Team S.
Single agent tasks, no credit ass. problems	○	○	√	×
Cooperative agent tasks, no credit ass. problems	○	√	○	×
Altruistic cooperation, no credit ass. problems	○	√	○	×
Single agent tasks, credit ass. problems	○	√	×	×
Cooperative agent tasks, credit ass. problems	○	√	×	×
Altruistic cooperation, credit ass. problems	○	√	×	×

In tasks that did not require cooperation, heterogeneous teams evolved with individual level selection achieved the highest team performance. Team heterogeneity allowed to evaluate a high number of different genomes in parallel, and individual selection allowed efficient selection of good genomes. However, these teams performed poorly in tasks that required cooperation and in tasks with credit assignment problems.

For multi-agent tasks that required cooperation, the highest team performance was achieved by homogeneous teams. These teams led to efficient cooperation between team members and they were not affected by credit assignment problems or costs associated with cooperation. Our results suggest that homogeneous teams are a safe choice when the requirements for agent cooperation are difficult to estimate. Compared to heterogeneous teams, homogeneous teams evaluate

less genomes, which may result in premature convergence to sub-optimal solutions. Our experimental results indicate that a simple way to prevent this problem is to use populations made of a large number of homogeneous teams.

Heterogeneous teams evolved with team selection were inefficient at selecting for good performance in all three classes of tasks and can not be recommended.

This study has two notable limitations. First, it did not address tasks that require specialization and division of labor. There is evidence that behavioral heterogeneity can lead to significant performance advantages for such tasks (Balch [17], Bongard [34], Potter et al. [178], Quinn [182], Tarapore et al. [213], Waibel et al. [222]). Also, our study did not consider teams with intermediate genetic similarity. Biological research has shown that such teams can overcome individual fitness costs of cooperation (Griffin and West [84], Griffin et al. [85]), thus combining the best of both worlds, enhanced genetic diversity with readiness to adopt altruistic behaviors. A good understanding of those conditions will require significant further research.

---

# 5 Artificial Ants Scale Reality Gap: Cooperative Robot Foraging

---

In theory, there is no difference between theory  
and practice. In practice, there is.

Albert Einstein (1879-1955)

In the previous chapter we used evolutionary robotics to optimize the performance of a team of robots. All optimization experiments were carried out in simulation, because the large number of evaluations required by the evolutionary process would have been too time-consuming in hardware. This chapter focuses on the application of evolutionary robotics to physical robot teams. It first reviews current approaches for the design of control strategies for physical robot teams and the position of evolutionary robotics in this context (Section 5.1). This is followed by an overview of the current state of the art in using evolutionary robotics for physical robot teams (Section 5.2). The software and hardware setup used in the experiments in this chapter is detailed in Section 5.3. In Section 5.4.1 we use simulation to evolve controllers for robot teams and transfer the best evolved controllers to the hardware setup. Section 5.4.2 presents a second set of such transfer experiments to gauge the potential of evolving physical robot teams without the use of simulation. An analysis and discussion of the results in Sections 5.5 and 5.6 concludes this chapter.

## 5.1 Control strategies for multi-robot systems

Collective robotics holds the promise of overcoming many limitations of single robot systems and has a wide range of applications (see Arai et al. [10], Cao

et al. [48], Dudek et al. [63, 64], Farinelli et al. [66], Panait and Luke [171], Parker [172], Stone and Veloso [210], Sycara [212] for reviews of multi-robot systems and their applications). Potential advantages of multi-robot systems are numerous and include concurrent vs. subsequent task performance, simultaneous sensing and action in multiple places, task dependent re-configurability, inherent robustness, collective power increase, smooth task allocation, and easy scalability (Bonabeau et al. [30], Holland et al. [103], Keller and Chapuisat [119], Martinoli [136], Mondada et al. [157]).

Control of robot teams has proven to be a difficult task (Mataric et al. [138]). Centralized control is appealing because it corresponds to the hierarchical nature of human reasoning (Baldassarre et al. [19], Bonabeau et al. [30]). However, it requires continuous communication between all robots (or all robots and a centralized controller), with a bandwidth and range related to the complexity and the size of the team. This can severely limit the scale of such systems. In addition, a centralized controller leads to comparably less robust systems than distributed control systems.

This is why most current attempts to control teams of robots use distributed control schemes, which fall into two broad categories: “Top down” architectures simplify control tasks by designing separate modules for subproblems (e.g. communication, decision making, task execution) and combining them using individual and global rules (Parker [172]). Due to their modularity and clear structure, these collective systems are relatively easy to understand and design. Further advantages include good compatibility with standard engineering toolkits (e.g. integration of optimized image processing algorithms) and comparatively easy analysis and debugging. On the down side, such modular solutions often have high demands on sensor quality and processing power, which reduce their appeal for collective systems with a large number of units (Panait and Luke [171], Parker [172]).

“Bottom up” approaches on the other hand try to avoid problem decomposition and instead attempt to find holistic solutions by exploiting invariants of the collective system and its interactions with the environment (Nolfi and Floreano [161]). Such simple solutions to apparently complex problems abound in nature (see e.g. Anderson et al. [7] for a review), which has raised hopes for success of similar strategies in robotics and other engineering applications (Beckers et al. [25], Bonabeau et al. [30], Camazine et al. [46], Dorigo et al. [62], Hol-

land and Melhuish [102], Nolfi and Floreano [161]). However, this class of solutions typically relies on complex interactions between individual robots based on local rules and local information, and can not easily be decomposed into functional units (Bonabeau et al. [30]). Global effects of individual decisions are often unpredictable and difficult to understand, and exploiting potentially useful self-organizing properties of such systems with classic explicit design methodologies seems impractical (Baldassarre et al. [19]).

An important example for a bottom up approach is evolutionary robotics, which allows the automatic creation of autonomous robot controllers (Floreano [70], Harvey et al. [95], Nolfi and Floreano [161]) and eliminates the need for task decomposition into simple basic behaviors or modules. By selecting desired behaviors according to a predefined metric (“fitness function”), suitable controllers emerge as the best (“fittest”) members of a pool of possible solutions.

Evolutionary robotics has been shown to provide solutions for a number of single robot tasks (see Nolfi and Floreano [161] for a review). It has been argued that evolutionary robotics may also offer advantages for the design of multi-robot systems. Evolution is a slow, but powerful way of adaptation and may find unintuitive ways to overcome the high amount of noise and interference characteristic of multi-robot systems (Nolfi and Parisi [162]). In addition, evolutionary robotics can produce unexpectedly simple solutions with minimal hardware requirements by exploiting hidden features of the software or hardware platform (Floreano et al. [71], Martinoli [135], Zufferey et al. [246]). This is important because many key advantages of distributed robotic systems, such as simultaneous action and sensing in multiple places or robustness, are linked to team size and simpler solutions with low hardware requirements are more cost efficient.

## 5.2 Related work

The roots of collective evolutionary robotics go back to attempts to generate controllers for teams of artificial birds or fish. In a seminal work, Reynolds [189] evolved controllers for a team of agents (*boids*) that fled from a predator in an environment with obstacles. This work was taken up and continued by Ward et al. [224], whose *e-boids* were capable of displaying schooling behavior in a 2D environment while successfully avoiding a predator and collecting food. Werner and Dyer [230] reported interesting results on a similar task, but with a co-evolved

predator. Zaera et al. [245] were able to evolve aggregation and dispersal for fish in a more realistic 3D simulation. Spector et al. [207] have investigated the evolution of altruistic food sharing in flocks in a population of flying agents. This historically relevant body of work has exclusively focussed on computer simulation of swarm behavior in idealized conditions.

Since its beginnings, considerable progress in the simulation of robotic agents has been made. Current simulations (Chaimowicz et al. [51], Gerkey et al. [79], Michel [149], Minar et al. [154], Noda [160], Pettinaro et al. [173]) allow to simulate multiple interacting robotic agents with a high degree of realism. However, most implementations are still too slow or too inaccurate to allow efficient evolutionary robotics experiments. These difficulties arise from the fact that the successful evolution and transfer of controllers to physical robots must meet two opposing goals: First, the simulation must be sufficiently fast to allow for the evolution of efficient robot controllers. This is especially challenging for the evolution of robot teams, because the number of potential interactions  $I$  (e.g., collisions or social behaviors) increases with the square of the number of robots  $N$ , i.e.,  $I(N) = O(N^2)$ . Second, it must be sufficiently accurate. This is challenging because the evolutionary process typically exploits inaccuracies of the simulation such as perfect object geometries (e.g., circles, squares) and simulation artifacts such as discrete time steps (Floreano et al. [71], Jakobi [109, 110]). These inaccuracies seem particularly large for the simulation of collisions.

Although there are several examples of successful transfer of evolved controllers for single robot tasks (see e.g., Jakobi [111], Jakobi et al. [112], Miglino et al. [152], Smith [205], Zufferey et al. [247]), only very few examples of successful transfer of evolved multi-robot teams exist. Quinn et al. [183, 184, 185] reported on the evolution of role allocation and team work for a team of three Khepera robots. Using infrared sensors, robots assembled into a chain with a leader and two followers and collectively moved away from their starting location. Fitness was calculated as a function of distance covered and team cohesion. Robot collisions were penalized by freezing robot position and orientation if their movement would lead to a collision. Since the controllers used infrared sensors to coordinate movement, they were evolved in a specifically tailored simulation using accurate sensor modeling. In addition, the physical robots were covered with a white reflective box to simplify sensor modeling.

As part of the SWARMBOT project (Mondada et al. [158]), Ampatzis et al. [6]

used robots to study the transfer of controllers evolved in simulation to physical robots. The authors evolved two robots in a phototaxis/anti-phototaxis task consisting of a light source obstructed by obstacles. Robots were evolved to recognize whether a path to the light source was open or obstructed. Both robots could communicate their findings using visual cues. After testing various evolved controllers on the physical robots, the authors transferred a single hand-picked controller. Forty tests with this controller resulted in correct obstacle recognition in all cases. Again, robot collisions were penalized, in this case by terminating the simulation if robots came within 2.5 cm of a collision more than 3 times. Unfortunately no details on the type of simulation used were given.

Floreano et al. [72] used robot teams to study the evolution of communication in a foraging task. Robots used light and omnidirectional vision as a communication channel. The authors evolved robot controllers in simulation for 500 generations and subsequently transferred them to a team of physical robots. Results showed a qualitatively similar behavior to that of the simulated robots, however no quantitative data for the robot performance in the hardware setup was given. The simulations used (Magenat and Waibel [132]) allowed for one-point collisions between circular robots and a circular immobile foraging target, however collision outcomes were not essential for foraging success.

Unlike previous work, the experiments presented here purposely minimized hardware adaptations (e.g., simplified robot shapes (Floreano et al. [72], Quinn et al. [183, 184, 185])) and common software shortcuts (e.g., eliminating agents on wall collisions or *ad-hoc* defined “poor behaviors” (Ampatzis et al. [6], Nolfi and Floreano [161], Zufferey et al. [247])) that could limit the range of evolved behaviors. The foraging task studied in our experiments also explicitly included collisions and made them a central, essential component for foraging success.

### 5.3 Experimental Setup

The experiments presented in this chapter used a cooperative altruistic foraging scenario similar to the one described in Section 4.3.3 of Chapter 4 with 6 micro robots, 8 small and 2 large tokens. The experimental arena and token payoffs as well as the robots, their control architecture and their genetic encoding were identical to those described in Section 4.3. Based on the results in Chapter 4, we chose a homogeneous team composition with team level selection.



### 5.3.1 Software platform

Simulations were conducted in a physics-based 2D simulator (Magenat and Waibel [132]), available as part of an open evolutionary framework (Magenat et al. [133]). All simulation parameters were modeled based on the hardware platform (see below). Simulations were designed to be sufficiently fast to allow an extensive number of multi-agent team experiments (300 – 500× as fast as real time on a state-of-the-art desktop computer).

### 5.3.2 Hardware platform

The hardware setup consisted of micro-robots (Caprari [49], Caprari and Siegwart [50]) and an experimental arena (Figure 5.1). The arena was equipped with a camera tracking system to identify arena positions of robots and tokens in the arena, an infra-red communication system to remotely program and drive robots, and a magnetic x-y table to automatically reposition tokens in the arena (not visible in Figure 5.1). All sub-systems were fully integrated to allow automatic and unsupervised evaluation of multiple robot teams for up to 15 hours (i.e., for the duration of robot battery autonomy). During automated experiments, approximately equal amounts of time were dedicated to the evaluation of robot teams and to the automatic repositioning of robots and tokens to random starting positions. The automatically collected data was verified using direct supervision, log files, and video recordings of all experiments.

### 5.3.3 Experimental method

We performed two sets of experiments. The first set used simulation to evolve controllers for a team of robots and subsequently transferred the best evolved controllers to a hardware platform. The second set of experiments tested randomly generated controllers in hardware without prior evolution in software.

#### Evolved controllers

For this first set of experiments, controllers were evolved for 150 generations. Evolution used the corresponding algorithm described in Section 4.2 with a population of  $M = 200$  teams of  $N = 6$  robots each. Each team was evaluated 25 times






Figure 5.1: Top image: Three micro-robots with a small food token. Bottom image: Micro-robots during an initial, exploratory transfer experiment with 10 robots per team (results not shown). The three robots at the front cooperate to push a large token towards the foraging goal (white wall visible in the left bottom corner). In the background another robot transports a small token. The other robots drive around the arena looking for tokens to forage.

for 90 seconds each, using random token and robot starting positions and orientations. All evolutionary experiments were independently replicated 40 times.

At the end of evolution we compared the performance of the best evolved teams in simulation to their performance on the hardware platform by transferring the evolved controllers to physical robots. For comparison we selected the team with the highest fitness at generation 150 of each of the 40 replicates. Since these best teams were likely to have benefitted from favorable robot starting positions in the initial 25 evaluations, all 40 teams were re-evaluated in simulation using a new set of random token and robot starting positions and orientations. Evaluations were repeated for 25 times for 120 seconds each in both the simulated foraging task and after transfer on the hardware platform. Overall, we therefore performed 1'000 foraging evaluations (25 evaluations for 40 teams) of 120 seconds each in the hardware platform and compared them to an identical number of evaluations in the software simulation.

### **Random controllers**

For this second set of experiments, we performed 20 replicates of experiments where we generated a random controller for one team of 6 robots. Random controllers were created by generating a random number in the range  $[0; 255]$  for each of the 32 connection weights of the robot's neural network control architecture (Section 4.3.2). Performance of each robot controller in simulation and on the hardware platform was again compared by repeating evaluations for 25 times for 120 seconds for each of the 20 teams in both the simulated foraging task and on the hardware platform. Overall, we therefore performed 500 foraging evaluations (25 evaluations for 20 teams) of 120 seconds each in the hardware platform and compared them to an identical number of evaluations in the software simulation.

### **Analysis**

For analysis, we compared the mean fitness values, the mean proportion of collected small tokens, and the mean proportion of collected large tokens of the 25 evaluations per team. For comparability and statistical analysis identical experimental parameters (experiment duration, number of evaluations, number of robots and tokens, etc.) were used for evaluations in simulation and on the hardware platform. The fitness values and the proportions of small and large tokens

collected per team were compared using Student's  $t$ -tests. All fitness values were normalized, with 0 being the minimum possible fitness and 1 the theoretical maximum fitness value.

## 5.4 Results

### 5.4.1 Performance of teams with evolved controllers

Figure 5.2 shows the foraging performance for the 40 best teams and the population average of 40 independent experiments over 150 generations of artificial evolution in simulation.

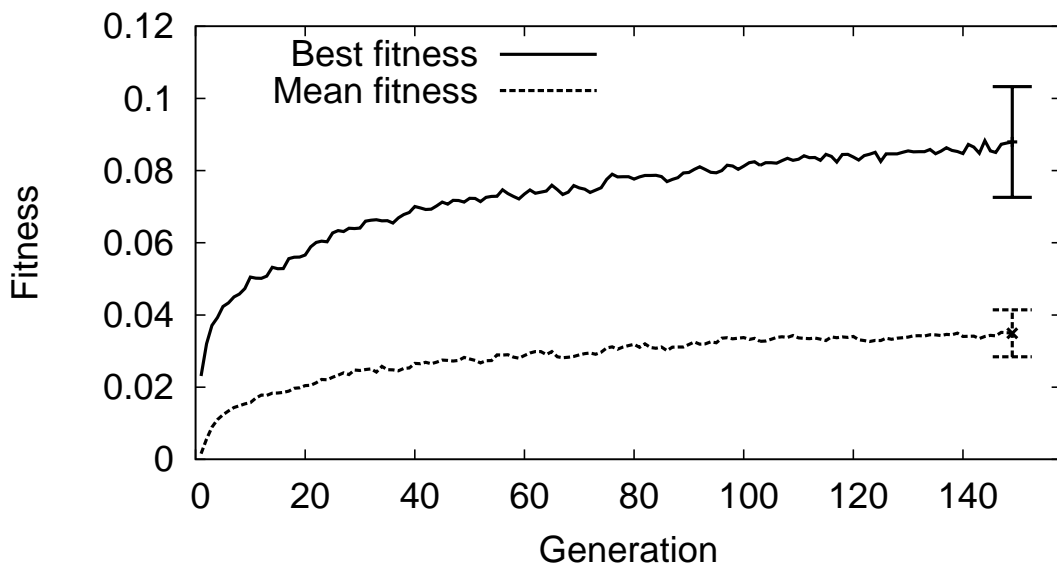


Figure 5.2: Evolution of the fitness of the best teams (“Best fitness”) and of the population average (“Mean fitness”) for 40 independent evolutionary runs over 150 generations of evolution in simulation. The error bars show standard deviations at the end of simulation.

Transfer of the 40 best teams in simulation at generation 150 and evaluation on the hardware platform led to important changes in the proportions of collected

small and large tokens and in team fitness (Figure 5.3). In the majority of cases (27 out of 40 teams) transfer led to a reduction of fitness. The overall, significant fitness loss (simulation versus hardware platform:  $-35.4\%$ , paired  $t$ -test,  $df = 39$ ,  $P = 0.019$ ) was caused by the performance drop in the collection of small tokens ( $-73.0\%$ ,  $P < 0.001$ ) after transfer. Large token collection on the other hand increased significantly after transfer ( $+355.0\%$ ,  $P < 0.001$ ).

### 5.4.2 Performance of teams with random controllers

The evaluation of the 20 randomly generated teams in simulation led to the collection of some small tokens in 11 out of the 20 teams, but did not result in the collection of large tokens in any team (Figure 5.4). Transfer and evaluation of the 20 randomly generated teams on the hardware platform led to important changes in the proportions of collected small and large tokens and in team fitness. In half of the cases (10 out of 20 teams) transfer led to a fitness increase. However, there was no significant overall change in fitness (simulation versus hardware platform:  $+40.4\%$ , paired  $t$ -test,  $df = 19$ ,  $P = 0.515$ ). Transfer from simulation to the hardware platform did not significantly affect the proportion of small tokens collected ( $+87.5\%$ ,  $P = 0.254$ ). Evaluations did not result in the collection of large tokens in any team in simulation and transfer did not result in a significant increase in large token collection ( $n/a$ ,  $P = 0.072$ ).

### 5.4.3 Performance of teams with evolved controllers versus teams with random controllers

To evaluate the usefulness of evolving controllers in simulation we compared the performance of the evolved controllers to that of the random controllers. In simulation, the fitness and the number of small tokens collected in teams with random controllers was significantly lower than that of evolved teams (random controllers in simulation versus evolved controllers in simulation:  $-89.9\%$  and  $-95.9\%$ , respectively; two-sample  $t$ -tests,  $df = 58$ , both  $P < 0.001$ ). The proportion of large tokens collected by teams with random controllers did not differ significantly from that collected by evolved teams when tested in simulation ( $P = 0.206$ ). On the hardware platform, the fitness of teams with random controllers was significantly lower than that of evolved teams (random controllers on the hardware platform versus evolved controllers on the hardware

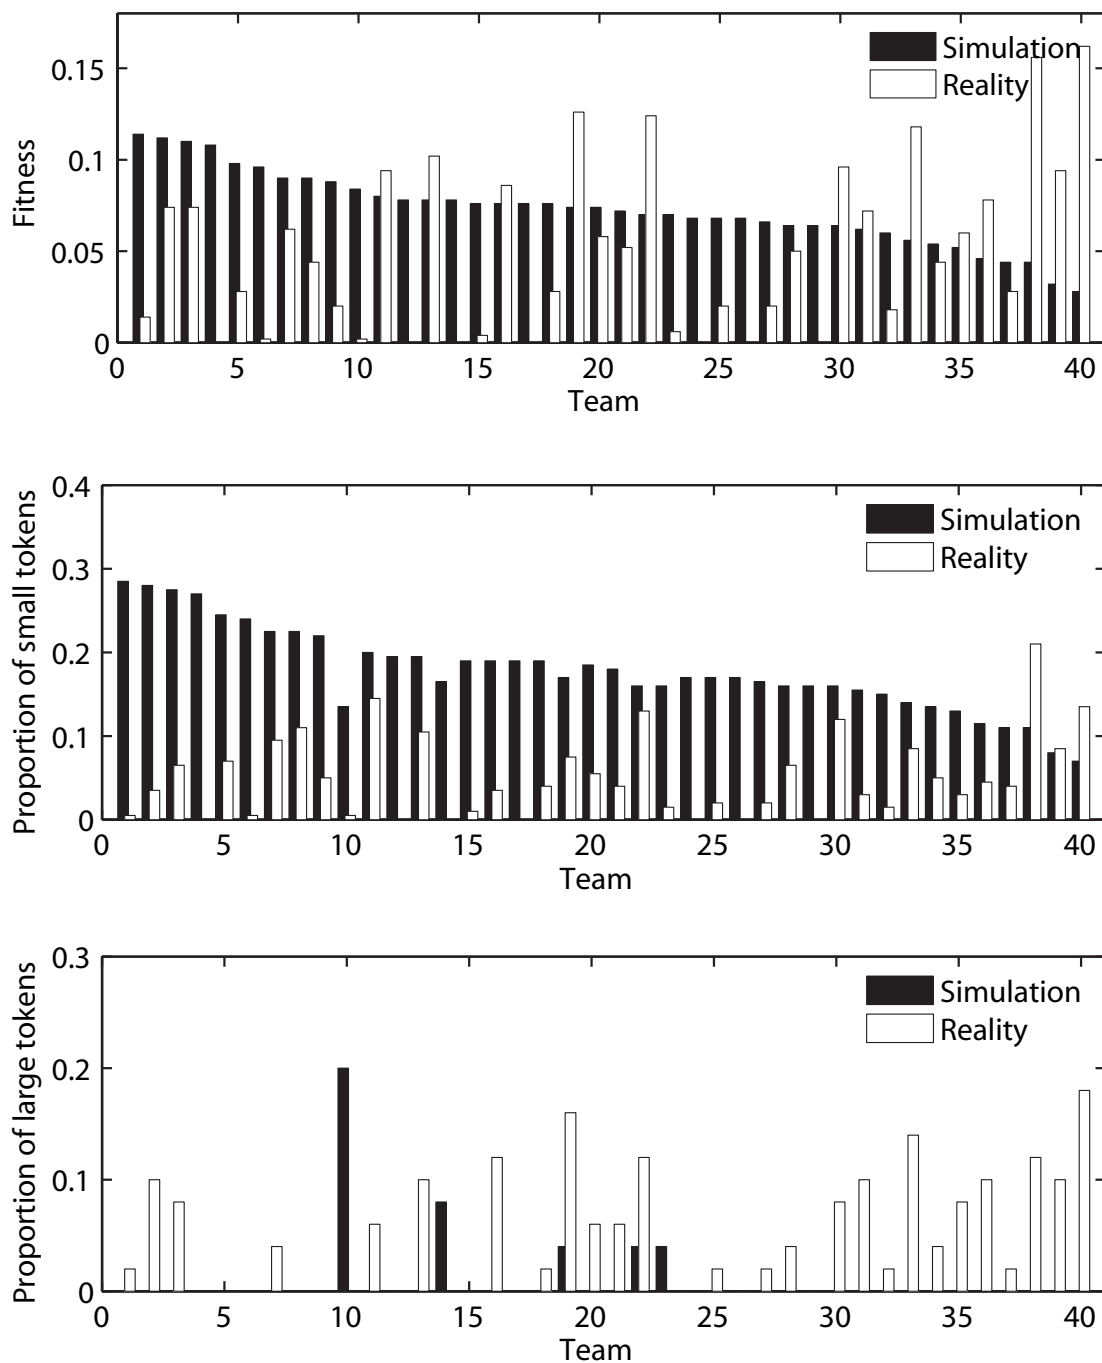


Figure 5.3: Mean fitness and proportions of collected small and large tokens of the best teams at generation 150 for 40 independent replicates of the experiment (25 trials per team). The black bars show the performance of teams in simulation, white bars the performance after transfer to the physical robots. For better readability teams were sorted by their mean fitness in simulation.

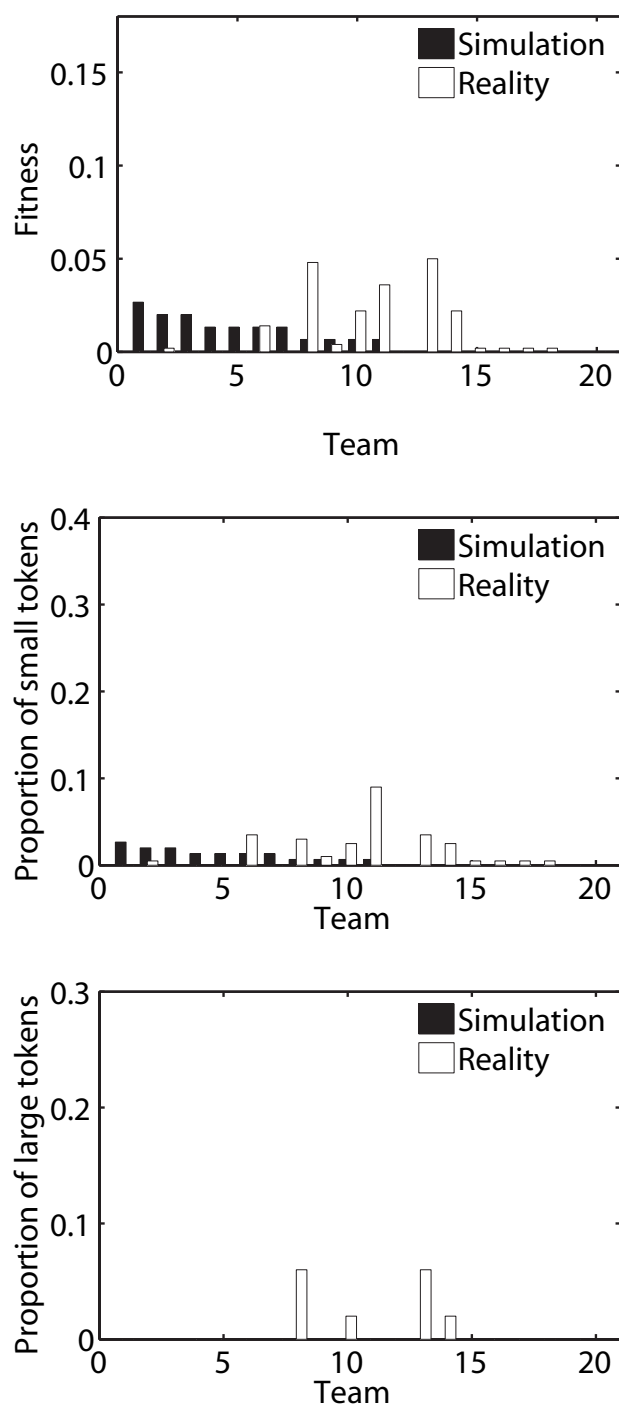


Figure 5.4: Mean fitness and proportions of collected small and large tokens of randomly generated controllers for 20 independent replicates of the experiment (25 trials per team). The black bars show the performance of teams in simulation, white bars the performance after transfer to the physical robots. For better readability teams were sorted by their mean fitness in simulation.

platform:  $-79.6\%$ ,  $P < 0.001$ ). The proportion of small and large tokens collected in teams with random controllers was significantly lower than that of evolved teams tested in hardware (random controllers on the hardware platform versus evolved controllers on the hardware platform:  $-73.0\%$  and  $-84.0\%$ , respectively; both  $P < 0.003$ ).

## 5.5 Discussion

When evolving controllers in simulation, successful foraging behavior evolved in all 40 experimental replicates during the 150 generations of artificial evolution. Overall, the robots only collected a small proportion of all tokens during foraging. This was due to the relatively short evaluation time (compare results in Chapter 4).

The evaluation of the foraging strategies evolved in simulation on the hardware platform resulted in an overall fitness decrease caused by the reduction of the number of small tokens collected. The same evolved foraging strategies led to a significant increase in the number of large tokens collected after transfer. A closer qualitative analysis of the evolved controllers transferred to the physical robots was conducted using video recordings. It revealed pronounced differences in the outcomes of collisions between robots in simulation and the outcomes of similar collisions on the hardware platform. Due to robot and token geometry, such collisions were difficult to avoid for robots that cooperatively pushed large tokens (see Figure 5.1). On the hardware platform collisions between robots had highly unpredictable outcomes because of the complex shapes of the robots which featured protruding sensors and connectors (Figure 5.1). The outcomes of collisions between physical robots often depended on small changes in robot behavior. In simulation the outcomes of such collisions were approximated by introducing a probability to stop robot movement. In spite of efforts to create a model that could predict collision outcomes for all robot behaviors (Appendix E), no accurate model could be built given the requirements for simulation speed to allow artificial evolution in software.

The comparison of the performance of random controllers in simulation and on the hardware platform did not lead to significant differences in fitness or the proportion of small and large tokens collected.

The comparison of controllers evolved in simulation and transferred to the hardware platform with random controllers showed no clear pattern. Random controllers fared worse than evolved controllers for the collection of small tokens and obtained lower fitness values in both simulation and reality. However, this comparison did not show significant differences for the proportion of large tokens collected in simulation. This is likely to be linked to the generally low proportion of large tokens collected in simulation rather than intrinsic limitations of evolved controllers (compare results in Chapter 4).

Our experiments also illustrated the high amount of time required for real robotic experiments when compared to experiments in simulation. The fully automated hardware setup described in Section 5.3.2 completed 1'500 foraging evaluations corresponding to 50 hours of continuous robot evaluation in less than 12 days, allowing for robot repositioning, recharging and occasional hardware problems. In comparison, the evolution of robot controllers used in the transfer experiments amounted to a total of 781 days of continuous robot evaluation, but took less than a day to complete using twenty computers.

## 5.6 Conclusion

Previous work on single robot systems has shown that transfer of robot controllers evolved in simulation to a hardware platform can lead to significant discrepancies in robot performance and behavior (Jakobi et al. [112], Miglino et al. [152, 153]). Our results suggest that these discrepancies may be similarly large in multi-robot systems.

Our observations suggest that the discrepancies observed in this study were due to the difficulties in accurately modeling the outcomes of collisions. Indeed, we are not aware of successful transfers of evolved controllers for tasks that go beyond instant one-point collisions, such as kicking a soccer ball (Smith [205]). Current work that relies on physical interactions often uses hand coded controllers in combination with evolution to avoid such discrepancies (Marco Dorigo, personal communication).

Our experiments show an example where the necessary tradeoff between simulation accuracy and simulation speed led to discrepancies between foraging behavior predicted by simulation and observed on the hardware platform. This reiterates the importance of validating controllers evolved for robotic applications



on a hardware platform (Ampatzis et al. [6]).

While simulation did not lead to high behavioral correspondence in the foraging experiments presented here, it has been shown to be sufficiently accurate for some multi-robot applications (Quinn et al. [185]). At the start of this project none of the available simulators met the requirements for simulation accuracy and simulation speed for this study, which led to the development of a custom-built solution (Appendix E). Recent years have seen a dramatic increase in the use of physics simulators. Rapid progress in robot simulation promises to reduce the discrepancies between simulation and reality by allowing for faster simulation of more accurate models (Michel [149], Pettinaro et al. [173]). Future work will show whether improved simulation will allow to overcome the limitations revealed by this work and if simulation in robotics will follow the trend set by the aircraft and automobile industry, where simulation has largely replaced prototypes and real-world experiments.



---

# 6

---

## Concluding Remarks

---

The world is its own best model.

Rodney Brooks (*Elephants don't play chess*, 1990)

### 6.1 Main Achievements

The work in this thesis indicates how parallels between the evolution of biological and artificial social organisms can lead to useful synergies between the research fields of evolutionary biology and evolutionary robotics. On the one hand the use of artificial social organisms presents an important step towards richer and more complete models for the evolution of cooperation. Most current modeling approaches for the study of the evolution of cooperation are limited to modeling frequency changes of isolated, predefined processes. The use of artificial evolution allows agent behavior to emerge from a large range of possible behaviors constrained by a minimal set of assumptions. This raises the prospect of modeling social behaviors as emergent phenomena, which result from an organism's genetic, cognitive and morphological features embedded in a rich (social) environment.

On the other hand this work illustrates how evolutionary robotics can benefit from the parallels to evolutionary biology by building on the broad spectrum of knowledge gained from the study of biological social organisms. Our experiments provide clear evidence that biological theory can equally apply to the evolution of robotic agents. They also show how biological insights can guide the design of multi-agent systems and lead to more efficient robot teams.

In Chapter 2 the evolution of artificial social organisms was used to test a fundamental part of biological theory for the evolution of cooperation, namely Hamilton's theory of kin selection. Our approach allowed us to overcome critical limitations in biological systems and to quantify the influence of genetic relatedness and cost to benefit ratios on the evolution of cooperation. Problems linked to the *a-posteriori* estimation of genetic relatedness in biological social groups could be sidestepped by artificially constructing social groups with precisely defined within-group relatedness. The even more daunting task of correctly estimating lifetime costs and benefits of social interactions could be avoided by fixing the values of costs and benefits of social actions, and by restricting social interactions to well defined groups.

In Chapter 3 we used the evolution of simple artificial agents to study the importance of the genetic architecture for the efficiency of cooperation and division of labor in social insects. This chapter revealed that the genotype to phenotype mapping can be a key factor for the success of biological teams. The comparison of three types of mapping showed that groups of interacting agents are much more efficient when their behavior depends on the behavioral phenotypes of other group members.

In Chapter 4 we used insights gained from the study of biological teams to identify three distinct classes of multi-agent tasks. By varying two key factors for the evolution of cooperation identified by biologists, we could compare the efficiency of teams and suggest guidelines for the evolution of agent teams for all three classes of tasks.

Chapters 2, 3 and 4 used simulation to evolve agent teams. Chapter 5 tackled a critical aspect for the applicability of evolutionary robotics to physical robot teams by transferring evolved teams from simulation to a hardware platform. Our experiments showed that artificial evolution in physics-based simulations can result in efficient physical robot teams. The experiments on the hardware platform constituted an important step towards evolving robot teams directly in hardware.

The behaviors resulting from the evolution of artificial social organisms in this thesis were inevitably much simpler than those found in their biological counterparts. Future work will show how the evolution of artificial social organisms can lead to far more complex social behaviors.

## 6.2 Future work

This thesis has shown the effects of several key parameters, including genetic relatedness and the costs and benefits of social actions, on the evolution of cooperation. Recent research has emphasized that other factors such as competition between relatives can significantly alter Hamilton's original predictions (Queller [180], West et al. [234]). While the experimental setup presented here included the effects of competition between relatives, it did not directly address this issue. We are currently working on extending our setup in order to study the effects of varying levels of competition between relatives on the evolution of cooperation.

Chapter 5 uncovered significant differences between behavior in simulation and reality. Evolution in simulation is appealing because of the large amount of time required for evolution in hardware and the high degree of hardware robustness required to withstand extensive evaluations on a robotic platform (Ficici et al. [68], Floreano and Urzelai [74], Mataric and Cliff [137], Thompson [216], Watson et al. [225, 226]). However, "the world is its own best model" (Brooks [41]), and can only be approximated by simulation. This is important because of the potential shortcomings of computer simulations (see Section 1.4) and because the usefulness of evolutionary robotics as a design tool for robot applications critically depends on its applicability to physical robot teams. An alternative approach to the transfer of teams evolved in simulation (Chapter 5) is to conduct artificial evolution directly on the hardware platform. This can be achieved by either conducting all evaluations directly on the hardware platform (Nolfi and Floreano [161]), by continuing the evolutionary process after transferring robot controllers to the hardware platform, or by alternating between evaluations in simulation and on the hardware platform. Additional work is needed to gauge the potential of these approaches for multi-agent systems.

Individual and social behaviors of robotic agents in this work emerged as a subset of a large, but deliberately limited range of possible behaviors. For behavioral studies of the evolution of social behavior such restrictions may be undesirable because they incorporate assumptions that restrict the total range of possible behaviors. Recent work has shown promising methods to reduce all three types of restrictions (Bongard et al. [35], Dürr et al. [65], Mattiussi [139], Mattiussi et al. [142]).

The experiments presented in this thesis have demonstrated knowledge transfer between the fields of evolutionary biology and evolutionary robotics. They have revealed synergies, which have led to significant contributions in both fields. Although our explorations were limited to interactions between a small number of artificial organisms with comparably simple morphologies and simple behaviors, they have allowed to test a critical part of biological theory for the evolution of cooperation in a rich model. This constitutes an important first step to show the potential of evolutionary robotics for the study of social behavior. By drawing on biological insights, this work has also significantly extended our understanding of the application of evolutionary robotics to robot teams.

We expect the cooperative interaction of these two complementary research fields to evolve over time and to lead us to a deeper understanding of the evolution of social behavior in animals and in robots.

---

# A

---

## Genetic Relatedness in Groups of Clones

---

This appendix provides the mathematical background for Chapters 2 and 3. It first explains within-group relatedness in groups of clones. Next it discusses how competition between relatives affects the evolution of cooperation and shows how its effects can be integrated into Hamilton's rule. It then shows how the effects of competition between relatives are accounted for in our experiments and concludes with a quantitative estimate of the corrections.

### A.1 Relatedness in Groups Composed of Different Proportions of Clones

Let a group of  $n$  individuals be composed of  $k$  types of clones with respective frequencies  $x_i \in [0; 1], i = 1 \dots k$ . Given random interactions between individuals in a group, the probability that an individual of type  $i$  interacts with a clone is then

$$P_i = \frac{nx_i - 1}{n - 1}. \quad (\text{A.1})$$

The average probability to encounter a clone corresponds to the average within-group relatedness  $r$  and is then given by

$$r = \sum_{i=1}^k x_i \cdot P_i, \quad (\text{A.2})$$

with

$$\sum_{i=1}^k x_i = 1. \quad (\text{A.3})$$

In our experiments all groups consisted of 8 individuals. To create groups with a relatedness of  $r = 1$  we composed groups entirely of individuals with the same genotype (clones). To create groups with an average relatedness of  $r = 0.75$  we used two different types of clones, A and B. Each group contained both types of clones with proportions  $A : B = 1 : 7$ . Using equation A.2, the resulting average relatedness in these groups was thus  $r = 0.75$ . To create groups with an average relatedness of  $r \simeq 0.54$  we composed each group of three types of clones in proportions  $6 : 1 : 1$ , which led to  $r = \frac{15}{28} \simeq 0.54$ . To create groups with an average relatedness of  $r = 0.25$  we again composed each group of three types of clones, but this time using proportions  $3 : 3 : 2$ , which resulted in an average relatedness of  $r = 0.25$ . To create groups with a relatedness of  $r = 0$  we composed groups of genetically different individuals, so that individuals were, on average, not genetically more similar to individuals in their group than to individuals in other groups.

## A.2 Genetic Relatedness with Competition Between Relatives

In its basic form, Hamilton's rule is given by

$$rb - c > 0, \quad (\text{A.4})$$

i.e., individuals will be selected for altruistic behaviors if the fitness benefit  $b$  to the beneficiary times the genetic relatedness  $r$  between actor and beneficiary is greater than the actor's fitness cost  $c$  of the altruistic behavior.

However, if individuals compete for reproduction with all members of their own group, the beneficial effects of altruism towards relatives are reduced by the amount of benefits that increase competition. Numerous authors have proposed enhanced versions of Hamilton's rule to incorporate the effects of competition between relatives (Frank [77], Grafen [82], Queller [180]). Here we follow Queller's



approach (Queller [180]) with Hamilton's rule rewritten as:

$$rb - c - r_c(b - c) > 0, \quad (\text{A.5})$$

where  $r$  is the focal individual's average relatedness to a random individual in its group, and  $r_c$  is the focal individual's average relatedness to a random individual in its competitive neighborhood. The term  $b - c$  is the general fitness increase of the competitive neighborhood resulting from the altruistic act. For a constant population size ("inelastic population regulation", West et al. [235]), the fitness increase of relatives in the competitive neighborhood  $r_c(b - c)$  corresponds exactly to the increase in competition between relatives.

Alternatively, the effects of competition between relatives can be integrated in the relatedness term. Relatedness coefficients measure genetic similarity as deviations from the population mean. By rearranging equation A.5 we can rewrite Hamilton's rule as

$$\frac{r - r_c}{1 - r_c} b - c > 0. \quad (\text{A.6})$$

This corresponds to the regression definition of relatedness (Queller [180]), with within-group relatedness  $r$  measured with respect to the baseline relatedness in the competitive neighborhood  $r_c$ .

The spatial scale at which competition occurs can be conveniently expressed by a single parameter  $a$  according to Frank [77], which links the relatedness  $r$  measured with respect to the group, to the relatedness  $r_c$  measured with respect to the competitive neighborhood

$$r_c = ar. \quad (\text{A.7})$$

Note that if the focal individual is unrelated to its competitors ( $r_c = 0$ ) or if competition is global ( $a = 0$ ) the basic form of Hamilton's equation holds. Conversely, if the focal individual is equally related to the beneficiaries of its altruism and to its competitors ( $r = r_c$ ,  $a = 1$ ), no altruism can evolve.

### A.3 Costs and Benefits of Cooperation

To correct for competition between relatives in our model, we adjusted the payoffs  $b$  distributed to group members. Using equation A.5, the choice of perform-

ing a selfish or an altruistic action will have no effect, if

$$rb - c - r_c(b - c) = 0, \quad (\text{A.8})$$

or if

$$b = \frac{c(1 - r_c)}{r - r_c} = \frac{c(1 - ar)}{r - ar}. \quad (\text{A.9})$$

To conduct experiments with  $\frac{c}{b} = 1, r - 0.1, r, r + 0.1$  and 0 for groups with different relatedness values  $r$ , we adjusted the benefits  $b$  according to equation A.9, while keeping costs  $c$  as well as the within-group relatedness  $r$  and the relatedness in the competitive neighborhood  $r_c$  constant, except for  $\frac{c}{b} = 0$  where  $c = 0$ . The resulting corrections to benefits  $b$  and cost to benefit ratios  $\frac{c}{b}$  for our experimental setup with 200 groups of 8 individuals each can be found in Table A.1.

Table A.1: Corrected benefits  $b$  and cost over benefit ratios for  $\frac{c}{b} = r$ .

	Benefits $b$	Corrected benefits $b$	Correction	Costs to benefits $\frac{c}{b}$	Corrected $\frac{c}{b}$	Correction
$r = 1$	1	1	0%	1	1	0%
$r = 0.75$	$\simeq 1.333$	$\simeq 1.335$	$\simeq 0.110\%$	0.75	$\simeq 0.749$	$\simeq -0.110\%$
$r \simeq 0.54$	$\simeq 1.866$	$\simeq 1.870$	$\simeq 0.204\%$	$\simeq 0.536$	$\simeq 0.535$	$\simeq -0.204\%$
$r = 0.25$	4	$\simeq 4.013$	$\simeq 0.330\%$	0.25	$\simeq 0.249$	$\simeq -0.330\%$
$r = 0$	7	7	0%	0	0	0%

---

# B

---

## Supplementary Data for Chapter 2

---

This section shows the evolution of fitness and the level of altruism for the experiments presented in Chapter 2. Section B.2 provides data including statistics for the mean fitness and the mean level of altruism after 500 generations of artificial selection over 20 replicates.

### B.1 Evolution of Fitness and Level of Altruism




Figure B.1: Evolution of the mean group fitness (a) and the mean level of altruism (b) during 500 generations of artificial selection for the 20 independent replicates per condition. Fitness values are normalized so that the maximum fitness of selfish groups is 1. The maximum fitness of altruistic groups is  $\frac{b}{c}$  times higher. Statistics can be found in Table B.1. (Figure continues on next page.)

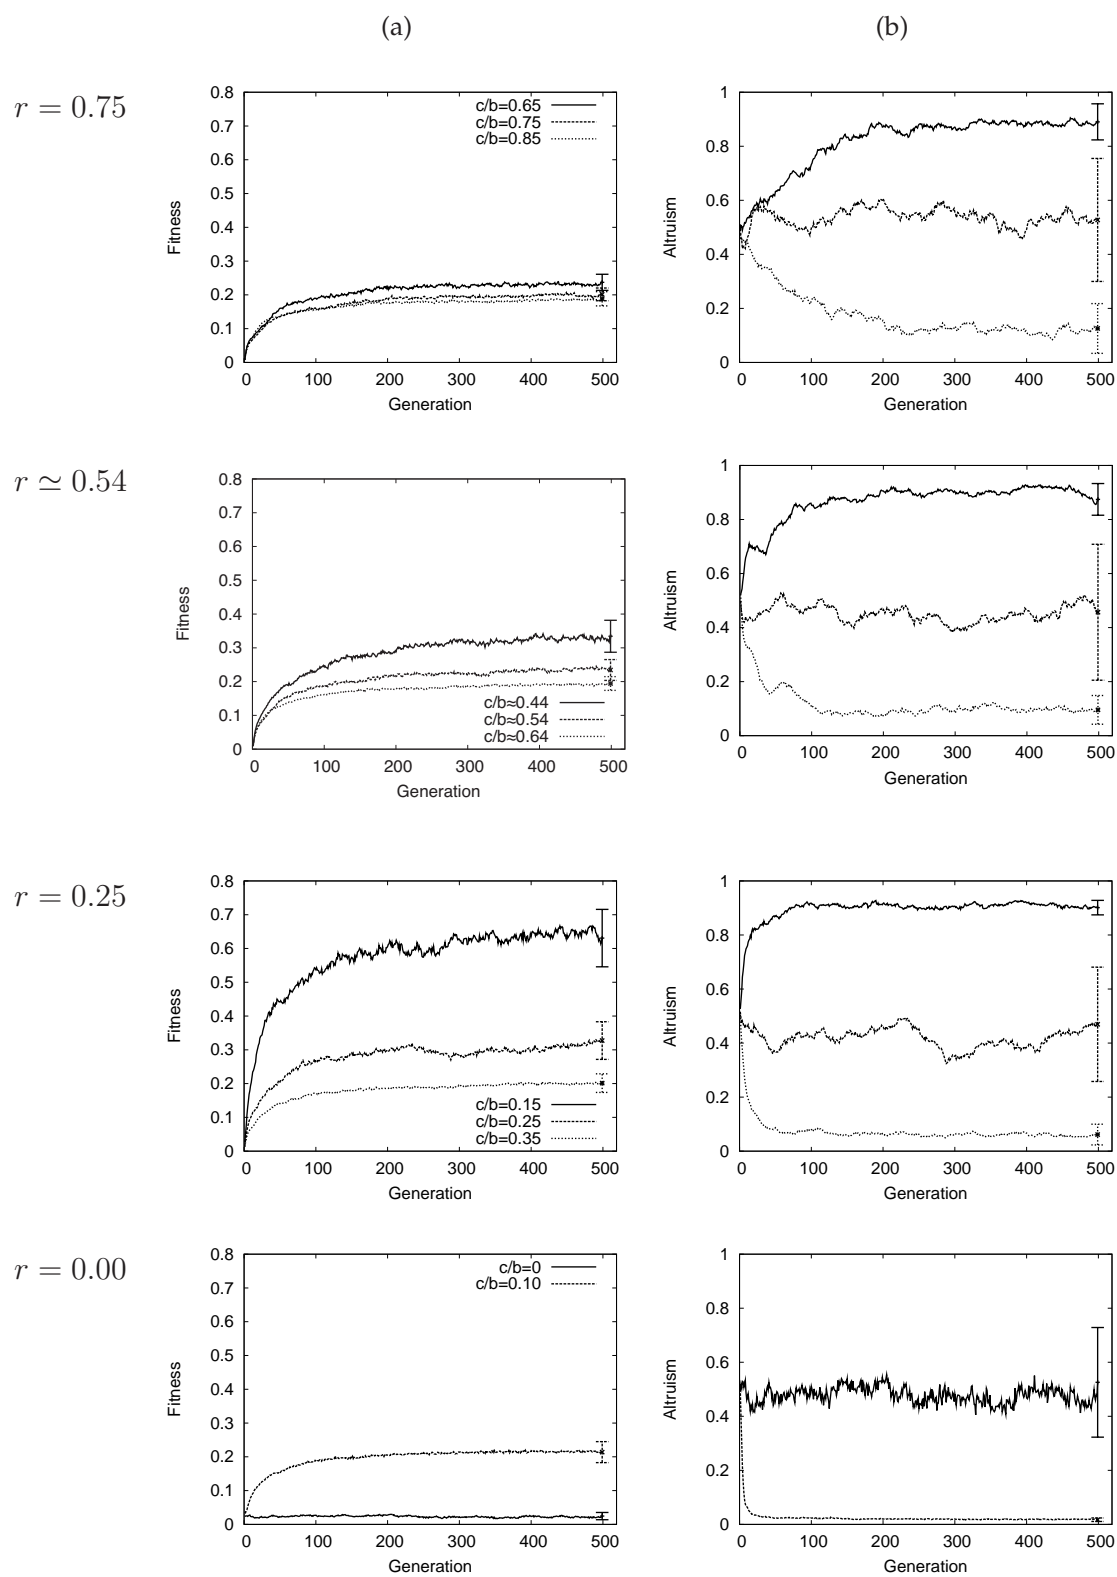


Figure B.2: (Figure continued from previous page.)

## B.2 Statistics for Fitness and the Level of Altruism after 500 Generations of Artificial Selection.

	Mean fitness			Mean level of altruism		
	$\bar{F}$	p-values		$\bar{A}$	p-values	
$r = 1, \frac{c}{b} = 0.90$	0.198	] < .001 ] < .305	] < .004	0.875	] < .001 ] < .001	] < .001
$r = 1, \frac{c}{b} = 1.00$	0.166			0.478		
$r = 1, \frac{c}{b} = 1.10$	0.174			0.110		
$r = 0.75, \frac{c}{b} = 0.65$	0.232	] < .001 ] < .012	] < .001	0.890	] < .001 ] < .001	] < .001
$r = 0.75, \frac{c}{b} = 0.75$	0.199			0.527		
$r = 0.75, \frac{c}{b} = 0.85$	0.186			0.126		
$r \simeq 0.54, \frac{c}{b} \simeq 0.44$	0.328	] < .001 ] < .001	] < .001	0.875	] < .001 ] < .001	] < .001
$r \simeq 0.54, \frac{c}{b} \simeq 0.54$	0.236			0.457		
$r \simeq 0.54, \frac{c}{b} \simeq 0.64$	0.191			0.095		
$r = 0.25, \frac{c}{b} = 0.15$	0.646	] < .001 ] < .001	] < .001	0.901	] < .001 ] < .001	] < .001
$r = 0.25, \frac{c}{b} = 0.25$	0.312			0.469		
$r = 0.25, \frac{c}{b} = 0.35$	0.200			0.061		
$r = 0, \frac{c}{b} = 0.00$	0.026	] < .001*		0.558	] < .001*	
$r = 0, \frac{c}{b} = 0.10$	0.214			0.017		

Table B.1: Mean group fitness  $\bar{F}$  and mean level of altruism  $\bar{A}$  of the last 10 generations of evolution for different cost to benefit ratios  $\frac{c}{b}$  and different within-colony relatedness values  $r$ . Mean values between experimental conditions were compared using two-sample  $t$ -tests ( $df = 38$ ) on the 20 independent replicates per condition. Due to missing data points for  $r = \frac{c}{b} = 0$  (tests marked with a \*),  $t$ -tests with this condition were performed on the 14 remaining data points ( $df = 26$ ; see Sections 2.3 and 2.4).



---

# C The Effects of Colony Size and Reallocation on Colony Efficiency

---

This section contains the results of additional simulations to confirm the robustness of our results with respect to two important model parameters. Section C.1 demonstrates that our results do not lead to marked differences in relative performance for variations in colony size. Section C.2 shows that variations in the proportion of reallocated individuals do not significantly affect our results.

## C.1 Colony Fitness for Large Colony Sizes




Figure C.1: Mean fitness values (10 runs) for the first 50 and last 50 generations in colonies with 100 or 10'000 individuals. Values are given for three different reallocation probabilities  $P_r$  and each of three genetic architectures (dynamic: red solid line, deterministic: green dashed line and probabilistic: blue dotted line). Increased colony size does not lead to marked differences in relative performance of the three genetic architectures.



## C.2 Colony Fitness for Varying Proportions of Reallocated Individuals

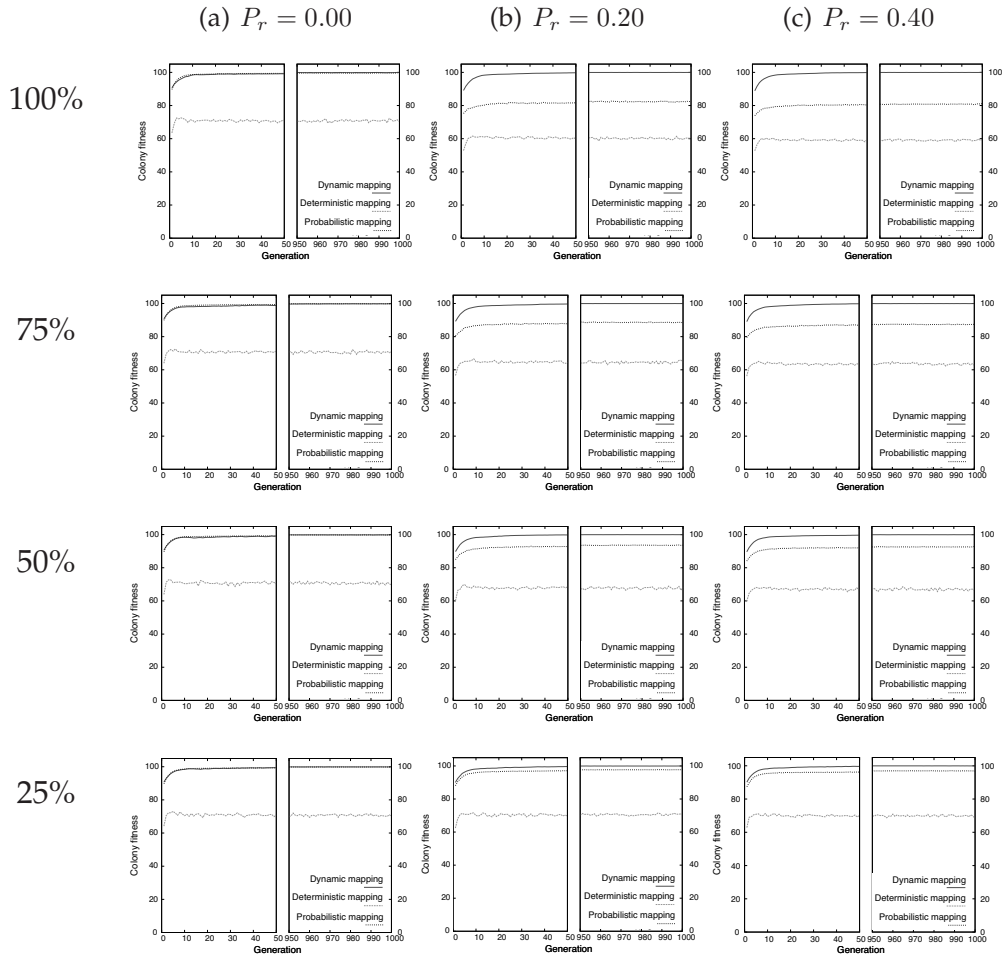


Figure C.2: Effect of the proportion of individuals reallocated (100%, 75%, 50% and 25%) on mean colony fitness. Mean fitness values (10 runs) are given for the first 50 and last 50 generations in colonies with 100 individuals for three different reallocation probabilities  $P_r$  and each of three genetic architectures (dynamic: red solid line, deterministic: green dashed line and probabilistic: blue dotted line). When a lower proportion of individuals is reallocated to a new task, differences between the three genetic architectures decrease but always remain highly significant (all  $P < 0.001$ ).



---

# D

---

## Supplementary Data for Chapter 4

---

This section contains additional data for simulations without disparities in genome evaluations (i.e., homogeneous teams with a single evaluation per team (instead of 10) and heterogeneous teams with 100 agents per population (instead of 1000)) and with credit assignment problems (CAP) for all three foraging tasks. Sections D.4 and D.5 show data for the foraging strategy in the altruistic foraging task (task 3) for simulations without disparities in genome evaluation and with/without credit assignment problems.

## D.1 Task 1 - Individual Foraging: Performance without disparities in genome evaluation and with credit assignment problems




Figure D.1: Task 1 - Individual Foraging without disparities in genome evaluations and with credit assignment problems. (a) Homogeneous teams with a single evaluation per team (instead of 10) and (b) heterogeneous teams with 100 agents per population (instead of 1000). Performance of homogeneous teams was higher than that of heterogeneous teams (all four  $P < 0.001$ ). Performance of homogeneous teams evolved with individual selection and homogeneous teams evolved with team selection did not differ significantly ( $P = 0.797$ ). Performance of heterogeneous teams evolved with individual selection and heterogeneous teams evolved with team selection did not differ significantly ( $P = 0.882$ ). For boxplot explanations refer to Figure 4.5.

## D.2 Task 2 - Cooperative Foraging: Performance without disparities in genome evaluation

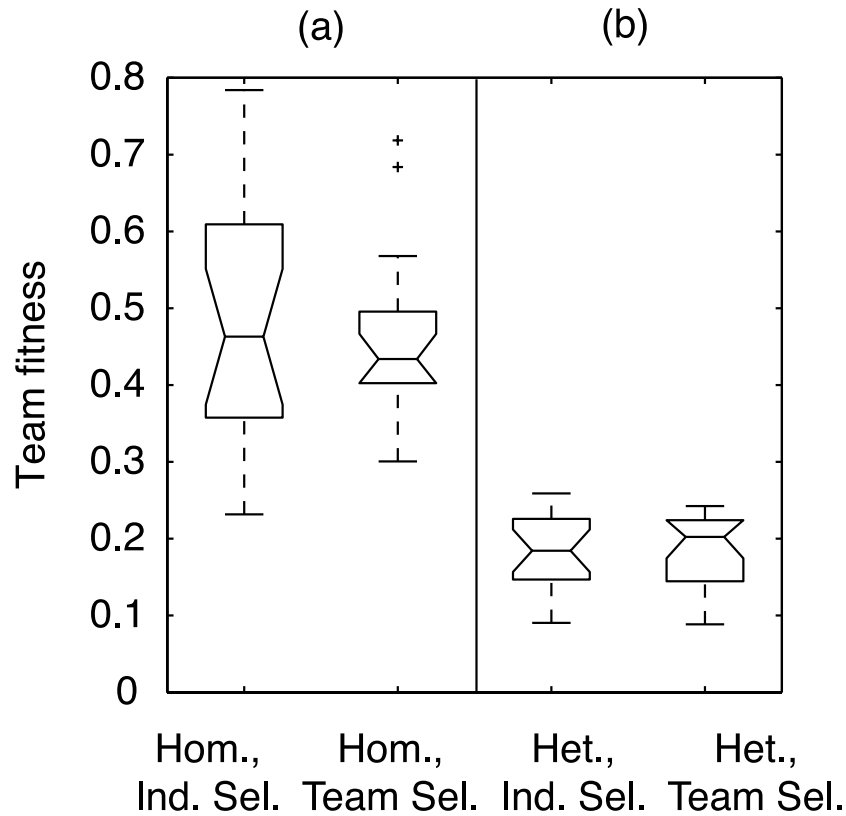


Figure D.2: Task 2 - Cooperative Foraging without disparities in genome evaluations. (a) Homogeneous teams with a single evaluation per team (instead of 10) and (b) heterogeneous teams with 100 agents per population (instead of 1000). The performance of heterogeneous teams with individual fitness assignment does not reach that of homogeneous teams (both  $P < 0.001$ ). Performance of heterogeneous teams evolved with individual selection and heterogeneous teams evolved with team selection did not differ significantly ( $P = 0.989$ ). Performance of homogeneous teams evolved with individual selection and homogeneous teams evolved with team selection did not differ significantly ( $P = 0.525$ ). Performance of heterogeneous teams evolved with team selection was significantly lower than that of homogeneous teams (both  $P < 0.001$ ). For boxplot explanations refer to Figure 4.5.

### D.3 Task 3 - Altruistic Cooperative Foraging: Performance without disparities in genome evaluation




Figure D.3: Task 3 - Altruistic Cooperative Foraging. (a) Homogeneous teams with a single evaluation per team (instead of 10) and (b) heterogeneous teams with 100 agents per population (instead of 1000). The performance of heterogeneous teams with individual fitness assignment does not reach that of homogeneous teams (both  $P < 0.001$ ). Performance of heterogeneous teams evolved with individual selection and heterogeneous teams evolved with team selection did not differ significantly ( $P = 0.076$ ). Performance of homogeneous teams evolved with individual selection and homogeneous teams evolved with team selection did not differ significantly ( $P = 0.756$ ). Performance of heterogeneous teams evolved with team selection was significantly lower than that of homogeneous teams (both  $P < 0.001$ ). For boxplot explanations refer to Figure 4.5.

## D.4 Task 3 - Altruistic Cooperative Foraging: Proportion of tokens collected without disparities in genome evaluation




Figure D.4: Task 3 - Altruistic Cooperative Foraging without disparities in genome evaluation, i.e. homogeneous teams with a single evaluation per team (instead of 10) and heterogeneous teams with 100 agents per population (instead of 1000). Foraging strategies of all evolutionary conditions were not affected by disparities in genome evaluation (compare Figure 8 in the paper). Homogeneous teams evolved with individual selection, homogeneous teams evolved with team selection and heterogeneous teams evolved with team selection collected a significantly larger proportion of large tokens than small tokens (all  $P < 0.01$ ), whereas heterogeneous teams evolved with individual selection collected a significantly smaller proportion of large tokens than small tokens ( $P < 0.001$ ).

## D.5 Task 3 - Altruistic Cooperative Foraging: Proportion of tokens collected without disparities in genome evaluation and without credit assignment problems




Figure D.5: Task 3 - Altruistic Cooperative Foraging without disparities in genome evaluation and without credit assignment problems, i.e. homogeneous teams with a single evaluation per team (instead of 10) and heterogeneous teams with 100 agents per population (instead of 1000). Foraging strategy in homogeneous teams evolved with individual selection, in homogeneous teams evolved with team selection and heterogeneous teams evolved with team selection was not affected by disparities in genome evaluation and credit assignment. Teams in these three evolutionary conditions again collected a significantly larger proportion of large tokens than small tokens (all  $P < 0.001$ ). However, credit assignment affected heterogeneous teams evolved with individual selection, which now collected a significantly larger proportion of large tokens than small tokens ( $P < 0.001$ ).



---

# E

---

## Robot simulation

The evolutionary robotics experiments presented in this thesis used a custom-built, discrete, 2D, physics-based simulation (Magnenat and Waibel [132]). Experiments used a fixed step size of 125 ms corresponding to a maximum displacement of 5 mm per time step for the maximum robot speed of 4 cm/s. On collisions objects were sequentially de-interlaced to their geometric boundaries according to conservation of energy, momentum and angular momentum and accounting for friction. Typical speed-ups for the experiments presented here were between 200 - 500 times real-time (depending on the neural network controller) on an Intel Pentium 4 with 2.80 GHz. The simulation is freely available online at <http://lis.epfl.ch> as part of an open evolutionary framework (Magnenat et al. [133]).

The following sections give an overview of key aspects of the simulation. Rather than providing a detailed account of the modeling process, it focusses on the methodology used to adapt simulation parameters to the hardware platform.

## E.1 Sensors

The micro-robot has two types of infrared sensors differing in range and a linear camera. Only five of the seven infrared sensors were modeled in our simulations. Regarding the camera, out of a total of 102 pixels, only 6 pixels in the central region of the linear camera were used (Figure E.1).




Figure E.1: Left: A schematic of the micro-robot in simulation (white octagonal polygon). The octagonal shape is a compromise between an exact geometric blueprint of the robot and minimal requirements for accurate token transport modeling (Section E.3.2). Infrared sensors (black and grey rectangles) were modeled using three rays per sensor. The schematic shows the ray bounding circles (grey circles) used for rapid detection of intersections with detectable objects. The grid-points (black and white crosses) mark 10 of the grid locations used for sensor characterization for objects such as the small token (black circle) shown in the figure. All grid locations intersecting the ray bounding circles were sampled (Section E.1.1). The black square on top of the robot marks the location of the robot camera. Right: The camera combined three rays to each side of the robot into two camera input values for error-robust sensor readings (Section E.1.2).

### E.1.1 Characterization

Sensor values for both types of infrared sensors were collected by displacing small tokens, large tokens and micro-robots (for 4 micro-robot orientations) on

a 2 dimensional grid  $10 \times 10$  mm in front of the robot. For each position, sensor values for both types of sensor were collected for all grid locations that led to an intersection between the objects and the robot's sensor bounding circles. In addition, activation values for the black and white arena wall were measured at 5 mm intervals and for - 45, 0 and +45 degree angles for both types of infrared sensor. Independent measurements of 6 robots were collected. This large number of measurements was needed to characterize the influence of at least three different factors associated with the measurements: First, individual infrared sensors vary in sensitivity. Second, sensor response shows a large dependance on both the horizontal and vertical inclination of the sensor, which can not be completely eliminated in hardware. Third, manual positioning of the micro-robots on the grid leads to statistic variation in the x/y position and angle.

Aperture and inclination angle of the linear camera were measured for 5 robots. In addition, camera response to the black and white walls of the arena was tested at different positions and orientations in the arena to determine the influence of varying light intensity and of light reflections.

### E.1.2 Modeling

Sensor simulation is one of the most computationally expensive operations required in our experiments. To minimize computational cost, we implemented a linear ray model for each infrared sensor (Figure E.1). Note that while ray tracing is more computationally expensive than the use of a lookup table, the former eliminates problems linked to combining activation values of multiple tokens in the vicinity of the sensor ("sensor occlusion", Quinn et al. [184]). Sensor occlusion becomes increasingly problematic as robot or token density increases. In our case, alignment of multiple robots to collectively push a large token critically depended on correct sensor values, which made a lookup table approach ineffective.

Each sensor was modeled using three rays. The response value  $R$  of each sensor ray was computed using a linear ray model of the form

$$R_i = S + M \cdot (\gamma \cdot e^{-\delta \cdot d_i} - \eta \cdot e^{-\theta \cdot d_i} + \zeta) \quad (\text{E.1})$$

where  $d_i$  was the distance to the closest obstacle of ray  $i$ , the greek letters  $\gamma$ ,  $\delta$ ,  $\eta$ ,  $\theta$  and  $\zeta$  were parameters of the model and  $M$  and  $S$  were a multiplicative and an offset factor, respectively (Figure E.2).




Figure E.2: The linear ray model for a single ray (full line) and its noise envelope (dashed lines). The final sensor response function was calculated as the weighted sum of the individual ray models for each of the three rays and subject to a running average (see text).

The weighted sum of the three ray values  $R_{1-3}$  was used to obtain the total sensor response at time step  $t$ ,  $A_t$ , which only depended on distance and was independent of surface reflectiveness or incident angle.

$$A_t = \sum_{i=1}^3 w_i \cdot R_i \quad (\text{E.2})$$

To find an optimal set of the parameters  $M$ ,  $S$ ,  $\gamma$ ,  $\delta$ ,  $\eta$ ,  $\theta$ ,  $\zeta$ ,  $w_1$ ,  $w_2$  and  $w_3$  used in Equations E.1 and E.2 as well as optimal opening angles for the three sensor rays and the sensor orientations we computed the least-square fit between our model and measured sensor data for all grid points and in all tested situations (small and large food token, micro-robot, walls) using a standard evolutionary algorithm available from our software framework. The resulting linear ray model is shown in Figure E.2.

The resulting approximated sensor model did not account for sensor noise, which was modeled separately. Three types of noise were used. First, a constant white noise and, second, a proportional noise were added to the total sensor response value at each time step. Third, a random offset noise was added to the sensor at the beginning of each trial to account for systematic differences in sensor readings.

In addition to these three types of noise, the total sensor response was sub-

jected to a running average, with the sensor activation  $A$  at time step  $t + 1$

$$A_{t+1} = P_s \cdot A_t + (1 - P_s) \cdot A_{t+1}, \quad (\text{E.3})$$

where  $P_s$  was a parameter for this running average, which reduced effects of discrete simulation time steps. The value of  $P_s$  was estimated by comparing sensor readings during test cases (see Sections E.2 and E.3) in simulation and reality.

Similar to the infrared sensors, the linear camera used equally spaced rays (see Figure E.1). However, each ray returned a 1 bit value depending on the color of obstacle seen. Since the only purpose of the camera was determination of the direction of the white target area, all tokens and walls except the target area were painted black. The 6 camera ray values were mapped to 2 values by combining the three ray values from each side using a logical OR function. In other words, if any of the three left (three right) rays hit the target area, the corresponding sensor value was set to 1. Since camera fidelity in hardware was high and the number of errors small compared to the amount of noise introduced by robots blocking each other's view of the target area no noise model for the camera was introduced.

## E.2 Actuators

The micro-robots have two rubber rimmed wheels driven by step motors at a maximum speed of 4 cm/s. Motors do not have encoders and were driven in open loop control.

### E.2.1 Characterisation

Basic motor characteristics were determined by measuring robot movement in a variety of simple movement tasks (linear displacement, circles with varying radii) for different speeds during free movement as well as under load (while pushing a small food token).

### E.2.2 Modeling

Basic motion was modeled based on the robot geometry, i.e., on the size of the wheel base and the wheel diameters. The motor transfer function, i.e. how motor

commands translated into actual motor movement, on the one hand, and motor response to different loads on the other hand were measured. Inversely proportional motor noise was added in accordance with measured values ( $\pm 5\%$  for 4 cm/s to  $\pm 15\%$  for 2.5 cm/s). Reproducibility of speeds  $< 2.5$  cm/s was very low, so they were excluded by implementing a stepwise transfer function in simulation and on the robot PIC. This function set all speeds in the range of  $[-2.5; 2.5]$  cm/s to 0 and left all other speeds unaffected.

### E.3 Interaction

For a number of reasons, adequate 2D collision modeling in evolutionary robotics is very challenging and has not previously been tackled. One main difficulty is the tendency of evolved controllers to exploit limit cases and boundary effects, and to rely on small details specific to the model implementation. In many cases during the development of our simulation we observed evolved controllers that optimized performance by relying on, and fine tuning to simulation artifacts caused by discrete time steps and other abstractions such as simplified robot geometry. The simplest way of avoiding such problems is to penalize evolved controllers by stopping the robot on collision, aborting simulation or eliminating “offending” individuals from the gene-pool. However, while such simplifications help in obtaining viable results, simply excluding possibilities inevitably limits the space of possible solutions and, thus, directly undermines one of the key advantages of evolutionary robotics: Exploitation of physical invariants that often results in surprisingly simple and robust behaviors.

Another challenge results from the generally low reproducibility of collision results in reality. Everybody who has played a game of pool can readily appreciate the rapid increase of complexity resulting from multiple collisions of even simple objects like smooth spheres. Exact simulation of the dynamics of collisions involving multiple robots with protruding connectors and infrared sensors was not feasible within the time frame of this project and would not have met the requirements for simulation speed.

A third difficulty specific to our setup was static friction in connection with token foraging in reality. In spite of a number of measures to enhance reproducibility, such as a glass arena surface, metal contact points and fine tuning using lead weights, the combined noise of variations in individual motor power, wheel fric-

tion, token orientation and small changes in robot-robot and robot-disk contact points resulted in high variation of collision outcomes in reality. Individual collision outcomes in simulation were difficult to model and only statistically accurate collision outcomes could be achieved.

### **E.3.1 Characterization**

All basic parameters such as dimensions, mass of tokens and robots as well as their geometries were quantified by measurement. Due to their high variability physical properties including static and dynamic friction between objects and the ground as well as between robots, tokens and walls were not measured directly but rather observed for typical situations and adjusted to closely match those cases.

### **E.3.2 Modeling**

Early tests showed that accuracy of collision modeling using highly simplified robot geometries was insufficient. In spite of numerous attempts, more faithful geometric models (e.g., featuring protruding sensors) also failed to capture essential details of robot-token collisions. The final robot shape used in simulation was obtained using a heuristic approach. The octagonal shape shown in Figure E.1 led to good correspondence for most collisions and also proved stable in subsequent evolutions. While single collision outcomes in the final simulation were fairly different, the degree of error and stochastic variation as well as overall behavior in simulation and reality were qualitatively similar.

## **E.4 Verification**

The sensor and motor models were verified using two methods. First, differences between sensor readings were assessed between simulation and the hardware platform for a number of static positions. Second, a series of test cases were used to study the infrared sensor and camera model, the motor model, the friction model, and their interaction.

### E.4.1 Sensors and Motors

The motor model was verified by comparing performance in simple movement tasks in simulation and reality. In addition, the interplay between sensor and actuator modeling was verified using two dynamic test cases:

- **Wall following:** The robot used its infrared sensor to follow the wall at different distances. This test case verified proximity readings for the infrared sensors. Varying motor speeds allowed us to simultaneously check the motor noise model.
- **Camera hysteresis:** The robot moved toward a corner between a white and a black arena wall using a hand coded controller. This controller described a pre-defined arc towards one side until it lost sight of the corner (i.e., until one of its two camera inputs changed its value) and then executed a similar arc towards the other side. The amplitude of the hysteresis was used to verify the correspondence of camera update cycles.

Apart from discovering potential problems concerning the interplay of sensors and actuators, these test cases also allowed us to study the effect of discrete simulation time steps in more complex behaviors. This was especially important with respect to the relatively slow camera refresh rate of only 4 Hz.

### E.4.2 Collisions

To evaluate correspondence of typical situations in simulation and reality another series of test cases was used:

- **Robot pushing a small token:** This is one of the basic behaviors required for successful completion of the foraging task. Varying robot motor speeds while pushing in a straight line allowed us to improve estimated static and dynamic friction thresholds. Pushing in an arc allowed us to adapt the shape of the simulated robot.
- **Robot pushing a small token along a wall:** This is a frequently evolved behavior and it was used in the same way as the previous point.
- **Single robot pushing a large token:** Token friction was tuned as to make transport of a large food token by a single robot very difficult, resulting in a large advantage for cooperation of two or more robots.



- Two robots pushing a large token: The second basic situation required for foraging.
- Robot pushing robot: Due to the robot's high friction rubber wheels, a robot could not be pushed by another robot, unless the robot was already moving in the same direction. Due to the symmetry of interactions in the simulation, this case was very difficult to model and presents a major difference between simulation and reality.

### E.4.3 Additional tests

In addition to these basic test cases, a number of recurrent situations were studied in detail.

- Robot is stuck pushing a small token into a corner: In many cases, robots push food tokens into a corner. Due to a combination of discrete time step simulation and noise, simulated robots frequently freed the obstacle and themselves.
- Multiple robot-robot collisions: A typical strategy found in evolved robot controllers was to first collect tokens by pushing them towards the nearest wall and then follow the wall, forming a chain ("traffic jam"). This allowed robots to transport multiple small and large tokens towards the target area. Due to the high amount of friction between a robot, other robots and walls and the high number of collisions this scenario was a useful test bed for simulation.
- Additional hand-coded strategies: In addition to the test cases described above, we hand-coded controllers for a number of simple, Braitenberg [38] behaviors such as obstacle avoidance, small and large disk pushing or target area homing using the camera.

## E.5 Conclusion

The simulation built as part of this project and subsequently tuned to the specific requirements for the robot foraging experiments used throughout this work proved to be fast and sufficiently accurate to evolve robot behaviors qualitatively

similar in software and hardware. However, as shown in Chapter 5, the transfer of the robot controllers evolved in simulation to the physical micro-robots led to large discrepancies in the number of tokens foraged.

As observed during parameter tuning using the test cases described in Sections E.2 and E.3, the main difficulties encountered were linked to collisions between robots and tokens, walls and other robots. In spite of numerous, iterative efforts with various simulated friction models, different robot geometries and simulation step sizes, no adequate solution could be found.

A reduction of the large variation in collision outcomes seems to be a key component of successful modeling of this foraging task. Potentially large improvements lie in simplified robot geometries and homogeneous surface materials to reduce the amount and variability of surface friction.

---

# Glossary

---

The research presented in this thesis addresses topics in biology, evolutionary computation and robotics. Identical terms are used in these historically different research fields, but in many cases their definitions differ between and even within research fields (West et al. [232]). For example, the term “cooperation” has been widely defined as a social behavior that is beneficial to the actor and the recipient (Bourke and Franks [37], Lehmann and Keller [125]), but is also used more generally for behaviors that are beneficial to the recipient but can be either beneficial or costly for the actor (Foster et al. [75], Maynard Smith and Szathmary [146], Sachs et al. [198]). In addition, the interdisciplinary nature of this thesis means that its readers have different backgrounds, which further increases the potential for semantic confusion. To avoid misunderstandings and conceptual difficulties I have made an effort to define key terms when they are first introduced. This glossary gives a summary to serve as a quick reference and lists good starting points for further reading. It also provides a concise picture of the relationship between key topics of this work.

**Actor** Focal *individual* which performs a *social behavior*.

**Altruism, altruistic cooperation** *Social behavior* that increases the *fitness* of other *individuals* and decreases the actor’s *fitness* (West et al. [231]).

**Artificial genome** In the context of *evolutionary computation*, artificial genomes are used as candidate solutions and encode the parameters that are optimized by an *evolutionary algorithm*. All experiments in this thesis use a binary genome composed of bits which are interpreted as weights of an *artificial neural network*, however many other methods for mapping an artificial

genome to a *phenotype* exist (Mattiussi et al. [140], Mattiussi and Floreano [141]).

**Artificial neural network** A network of simple processing elements (neurons) which can exhibit complex global behavior, determined by the connections between the neurons and neuron parameters (sometimes called “weights”). Artificial neural networks (sometimes called “neural networks”) are commonly used in artificial intelligence and *evolutionary robotics*. See, for example, Figure 2.2 and Section 2.2.1.

**Artificial social organism** See *organism*.

**Beneficiary** An *individual* that benefits from a *cooperative social behavior*.

**Benefit** See *Hamilton’s rule*.

**Competition** In the context of this thesis, reproductive competition, i.e. equivalent to Darwin’s struggle for existence. The contest for *selection* and *reproduction*.

**Competitive neighborhood** The scale at which intra-specific *competition* takes place.

**Conflict** *Competition* between lower level units of selection, leading to a disruption of the functioning of the group. Also see *level of selection*.

**Cooperation** *Social behavior* that increases the *fitness* of other *individuals*. See *social behavior*.

**Cost** See *Hamilton’s rule*.

**Credit assignment, credit assignment problems** In the context of *evolutionary computation*, a single or usually many fitness rewards to one or more individuals. The distribution of *fitness* rewards leads to credit assignment problems (Grefenstette [83], Minsky [155]) in many cooperative multi-agent tasks, because individual contributions to group performance are often difficult to estimate or difficult to monitor.

**Crossover** See *genetic operators*.

**Direct fitness** Fitness gained through the production of offspring. See *inclusive fitness*.

**Diversity** In the context of *evolutionary computation*, variation in the *genotypes* or *phenotypes* of a group of *individuals* or of a *population*. A variety of different measures exist (Mattiussi et al. [143]).

**Evaluation** In the context of *evolutionary computation*, the testing of a candidate solution to determine its *fitness*.

**Evolution** In the context of evolutionary biology, change in *genetic* traits between *generations*. In the context of *evolutionary computation*, the automatic creation of a system using an *evolutionary algorithm*. The term is used interchangeably in both contexts throughout this thesis.

**Evolutionary algorithm** Subset of *evolutionary computation* that uses *genetic operators* on a *population* of candidate solutions to solve optimization problems or to study principles of biological evolution (Goldberg [80]). Evolutionary algorithms were inspired by the evolution of biological organisms and rely on similar mechanisms to find candidate solutions. Candidate solutions play the role of *individuals* in a *population* and are represented by *artificial genomes*. *Individuals* are evaluated on the task and are assigned a quality (*fitness*) according to their performance. Some *individuals* of the *population* are then *selected* for *reproduction*, which uses *genetic operators* to form new *individuals*. Repeated *evaluation*, *fitness assignment*, *selection* and *reproduction* results in the evolution of the *population* and leads to improved candidate solutions to the optimization problems.

**Evolutionary computation** Evolutionary computation refers to meta-heuristic optimization algorithms, such as *evolutionary algorithms*, genetic programming (Koza [123]), particle swarm optimization (Kennedy and Eberhart [121]) or ant colony optimization (Dorigo and Stuetzle [61]), which use iterative processes to solve optimization problems.

**Evolutionary robotics** The application of *evolutionary computation* to robotics (Nolfi and Floreano [161]).

**Fitness** In the context of *evolutionary computation*, the performance of candidate solutions in an *evolutionary algorithm*. It is important to note that evolution

selects for *fitness* increase relative to the *fitness* of competitors, rather than for an absolute increase in *fitness*. In the context of evolutionary biology, the overall fitness of an organism is usually defined as a mathematical expectation  $p \cdot e$ , where  $p$  is the probability that an organism at the egg stage will reach adulthood, and  $e$  is the expected number of offspring that the adult organism will have. However, other definitions are used when needed. For a detailed discussion, see Sober [206].

**Fitness landscape** The *fitness* of *genomes* can be represented as a surface in an often high dimensional space known as the fitness landscape. *Genomes* are coordinates in the landscape, and the fitness of these genetic strings represents the “height” of the surface at the corresponding coordinates. *Genotypes* which are very similar are said to be “close” to each other, while those that are very different are “far” from each other.

**Gene** In the context of *evolutionary computation*, a part of the *genome* that encodes a parameter which is evolved.

**Generation** In the context of *evolutionary computation*, a single round of *evaluation*, *selection* and *reproduction*.

**Genetic architecture** The *genotype-phenotype mapping* between *individual's genes* and their behavioral *phenotype*.

**Genetic operators** Genetic operators are processes used in *evolutionary computation*. They usually mimic the processes of biological evolution. The most commonly used genetic operators include the *selection* operator, which selects individuals for *reproduction* according to their *fitness*, the crossover operator, which exchanges parts of the genomes of two individuals, and the mutation operator, which introduces random variation in the *individual's genome*.

**Genome** See *artificial genome*.

**Genotype** The genotype contains the genetic material or “instructions” required to build a biological organism or an artificial system. See *phenotype*.

**Genotype-phenotype mapping** Process by which the *genotype* is decoded into the corresponding *phenotype* (e.g., a biological organism or an artificial sys-

tem). In biological organisms the *phenotype* is usually created from the *genotype* via a complex developmental process. *Evolutionary computation* usually uses a simple, direct (“one-to-one”) mapping to create the *phenotype*.

**Group selection** In the context of evolutionary biology, group selection refers to the idea that *genes* can spread in a *population* because of the *benefits* they bestow on groups, regardless of the *fitness* of *individuals* within that group. Although *selection* at the group level is theoretically possible, it is usually a weak force compared to *selection* between individuals. This is because *selection* at the individual level proceeds much faster than selection at the group level.

**Hamilton’s rule** Rule that governs the spread of a *gene* for a *social action* in a population, due to W. D. Hamilton (Hamilton [91]). It invokes three terms: the actor’s reproductive cost  $c$  of the *social action*, the recipient’s reproductive benefit  $b$  gained by the *social action*, and the genetic *relatedness*  $r$  between the actor and the recipient. Hamilton’s rule states that a *gene* for a *social action* can spread in a *population* if  $rb > c$ . This is because a *gene* will spread if it can successfully create copies of itself, irrespective of the fact whether copies are made by increasing the *reproduction* of the bearer (actor), or by increasing the *reproduction* of other bearers of the same *gene* (*related* recipients).

**Hamilton’s theory** Theory dealing with aspects relating to *Hamilton’s rule*. It is referred to in many ways. Hamilton called it “inclusive fitness theory”, but it is more often referred to as “kin selection”.

**Helping** *Social behavior* that increases the recipient’s fitness (Lehmann and Keller [125]). This includes *cooperative* behavior and *altruistic* behavior. See *social behavior*.

**Inclusive fitness** In the context of evolutionary biology, fitness encompasses conventional Darwinian fitness with the addition of the consequences of *altruistic* or selfish behaviors. Inclusive fitness of an individual is divided into two components: “direct fitness” and “indirect fitness”. Direct fitness is gained through *reproduction*, i.e. the production of offspring, and indirect fitness through aiding the reproduction of non-descendent relatives. By *helping* a close relative *reproduce*, an *individual* still passes on its own *genes* to the next *generation*, albeit indirectly. See *Hamilton’s rule*.

**Indirect fitness** Fitness obtained by aiding the *reproduction* of relatives. See *inclusive fitness*.

**Individual** In the context of *evolutionary computation*, a member of the *population* of candidate solutions on which *evolutionary algorithms* operate.

**Kin selection** Another name for *Hamilton's theory*.

**Level of selection** The unit containing genetic information that is visible to *selection* (sometimes termed "vehicle"; Dawkins [58]), which is usually different from a *gene* (Keller [118]). *Selection* typically acts at the level of individuals and possibly higher-level units, whereas only the genetic information contained inside those units is *replicated* and endures over *generations*. This important distinction has spawned several decades of debate and still continues to cause misunderstandings and confusion (compare, e.g., Foster et al. [75]).

**Multi-level selection** Rather than acting at a single level, *selection* usually operates at multiple levels. The theory concerned with the transition between different levels of selection and their interaction is called multi-level selection theory. See also *level of selection*.

**Mutation** See *genetic operators*.

**Natural selection** Process of genetic change over time resulting from Darwin's three conditions for evolution: genetic variation, genetic heritability and genetic selection for reproduction ("struggle for existence"). Also see *selection*.

**Neural network** See *artificial neural network*.

**Organism** In the context of this thesis we define an organism (in Greek "organon" = instrument) broadly as a living complex adaptive system. The term *artificial social organism* used throughout this thesis to describe interacting software and robotic agents describes agents that are arguably complex and adaptive. However, both definitions of life and organisms differ widely between sources and are topics of debate (Maynard Smith and Szathmáry [146, chap. 2], Murphy and O'Neill [159], Schrödinger [200]).



**Phenotype** The phenotype is the observable part of a biological organism or of an artificial system that results from the decoding of the *genotype*.

**Population** In the context of *evolutionary computation*, the pool of *individuals* from which *individuals* are selected for reproduction.

**Recipient** An *individual* that benefits from a *cooperative social behavior*.

**Relatedness, Genetic relatedness** A measure of genetic similarity. A definition is complex and has been the topic of much debate. In *kin selection* theory it is usually defined as a regression coefficient (see for example Queller and Goodnight [181]). A more detailed definition can be found in Sections 2.2.3 and Appendix A.

**Reproduction** In the context of *evolutionary computation*, the replication of a selected candidate solution.

**Reproductive fitness** See *fitness*.

**Run** In the context of this thesis, a single of usually many replicates of an evolutionary experiment.

**Selection** In the context of *evolution*, *individuals* (or other *units of selection*) are said to be selected if they contribute offspring to the succeeding *generation*. In *evolutionary computation*, selection is performed based on *fitness* using a *genetic operator*. See also *level of selection*.

**Selfishness** A *social behavior* which is *beneficial* to the *actor* and *costly* to the *recipient*.

**Social action, social behavior** In the context of this thesis, a social behavior is a behavior directed towards, or taking place between, members of the same population. Social actions or social behaviors can be categorized according to their consequences for the actor. Additional categorizations that also take into account consequences for the receiver are not discussed here (see for instance West et al. [231]). Social actions can be beneficial, i.e. increase *direct fitness*, or costly, i.e. decrease *direct fitness*. Social actions that increase the actor's *direct fitness*, i.e. when the cost for the social action  $c < 0$  (see

*Hamilton's rule*), are called cooperative. Social actions that decrease the actor's *direct fitness*, i.e. when the cost for the social action  $c > 0$ , are called altruistic (Lehmann and Keller [125]).

**Trial** In the context of this thesis, an *evaluation* consists of several instances of the same experiment called trials and conducted with different initial conditions to reduce the effect of random events.

**Unit of selection** See *level of selection*.

---

# Bibliography

---

- [1] Agah, A. and Bekey, G. A. (1997). Phylogenetic and ontogenetic learning in a colony of interacting robots. *Autonomous Robots*, 4:85–100.
- [2] Agah, A. and Tanie, K. (1997). Robots playing to win: evolutionary soccer strategies. In *Proceedings of the 1997 IEEE International Conference on Robotics and Automation*, volume 1, pages 632–637.
- [3] Agogino, A., Stanley, K., and Miikkulainen, R. (2000). Online interactive neuro-evolution. *Neural Processing Letters*, 11(1):29–38.
- [4] Agogino, A. and Tumer, K. (2004). Efficient evaluation functions for multi-rover systems. *Lecture Notes in Computer Science*, 3102:1–11.
- [5] Alonso, W. J. and Schuck-Paim, C. (2002). Sex-ratio conflicts, kin selection, and the evolution of altruism. *Proc. Natl. Acad. Sci. U.S.A.*, 99:6843–6847.
- [6] Ampatzis, C., Tuci, E., Trianni, V., and Dorigo, M. (2005). Evolving communicating agents that integrate information over time: a real robot experiment. In *CDROM Proceedings of the Seventh International Conference on Artificial Evolution (EA 2005)*. Springer Verlag, Berlin, Germany.
- [7] Anderson, C., Theraulaz, G., and Deneubourg, J. L. (2002). Self-assemblages in insect societies. *Insectes Soc*, 49:99–110.
- [8] Andre, D. (1995). The automatic programming of agents that learn mental models and create simple plans of action. *IJCAI-95: Proceedings of the Fourteenth International Joint Conference on Artificial Intelligence*, 1:741–747.

- [9] Andre, D. and Teller, A. (1999). *RoboCup-98: Robot Soccer World Cup II*, volume 1604, chapter Evolving Team Darwin United, pages 346–351. Springer Berlin, Heidelberg.
- [10] Arai, T., Pagello, E., and Parker, L. E. (2002). Editorial: Advances in multi-robot systems. *IEEE Transactions On Robotics and Automation*, 18(5):655–661.
- [11] Axelrod, R. (1984). *The Evolution of Cooperation*. Basic Books, New York, NY, USA.
- [12] Axelrod, R. (1997). *The Complexity of Cooperation - Agent-Based Models of Competition and Collaboration*. Princeton University Press.
- [13] Axelrod, R., Riolo, R. L., and Cohen, M. D. (2002). Beyond geography: Cooperation with persistent links in the absence of clustered neighborhoods. *Personality and Social Psychology Review*, 6:341–346.
- [14] Axelrod, R. a. (2005). A guide for newcomers to agent-based modeling in the social sciences. *Handbook of Computational Economics*, 2.
- [15] Axtell, R. (2000). Why agents? on the varied motivations for agent computing in the social sciences. In *Proceedings of the Workshop On Agent Simulation: Applications, Models and Tools*. Argonne National Laboratory, IL.
- [16] Bäck, T. (1996). *Evolutionary Algorithms in Theory and Practice: evolution strategies, evolutionary programming, genetic algorithms*. Oxford University Press, New York, NY.
- [17] Balch, T. (1998). *Behavioral Diversity in Learning Robot Teams*. PhD thesis, Georgia Institute of Technology.
- [18] Balch, T. (1999). The impact of diversity on performance in multi-robot foraging. *International Conference On Autonomous Agents, Seattle, Washington, United States*, pages 92–99.
- [19] Baldassarre, G., Nolfi, S., and Parisi, D. (2003a). Evolution of collective behavior in a team of physically linked robots. In Cagnoni, S. E. A., editor, *Applications in Evolutionary Computing*, pages 581–592. Berlin: Springer Verlag.
- [20] Baldassarre, G., Nolfi, S., and Parisi, D. (2003b). Evolving mobile robots able to display collective behavior. *Artificial Life*, 9:255–267.

- 
- [21] Baldassarre, G., Parisi, D., and Nolfi, S. (2003c). Coordination and behavior integration in cooperating simulated robots. In Schaal, S., Ijspeert, A., Billard, A., Vijayakumar, S., Hallam, J., and Meyer, J.-A., editors, *From Animals to Animats 8: Proceedings of the VIII International Conference On Simulation of Adaptive Behavior*, pages 385–394, Cambridge (MA). MIT Press.
- [22] Baray, C. (1997). Evolving cooperation via communication in homogeneous multi-agent systems. *Intelligent Information Systems, 1997. IIS'97. Proceedings*, pages 204–208.
- [23] Baray, C. (1998). Effects of population size upon emergent group behavior. *Complexity International*, 6.
- [24] Bassett, J. K. and De Jong, K. A. (2000). Evolving behaviors for cooperating agents. In Ras, Z. W. and Ohsuga, S., editors, *Foundations of Intelligent Systems: 12th International Symposium, ISMIS 2000, Proceedings.*, pages 157–165. Springer Berlin / Heidelberg.
- [25] Beckers, R., Holland, O., and Deneubourg, J. (1994). From local actions to global tasks: Stigmergy and collective robotics. In *Proceedings of Artificial Life IV*, volume 181, page 189.
- [26] Bertram, S. M., Gorelick, R., and Fewell, J. H. (2003). Colony reponse to graded resource changes: an analytical model of the influence of genotype, environment, and dominance. *Theor. Popul. Biol.*, 64:151–162.
- [27] Bianco, R. and Nolfi, S. (2004). Toward open-ended evolutionary robotics: Evolving elementary robotic units able to self-assemble and self-reproduce. *Connection Science*, 16(4):227–248.
- [28] Blumenthal, H. J. and Parker, G. B. (2004). Co-evolving team capture strategies for dissimilar robots. In *Proceedings of the AAAI 2004 Symposium On Artificial Multiagent Learning, Washington*.
- [29] Bonabeau, E. (1998). Fixed response thresholds and the regulation of division of labor in insect societies. *Bull. Math. Biol.*, 60:753–807.
- [30] Bonabeau, E., Dorigo, M., and Theraulaz, G. (1999). *Swarm Intelligence - From Natural to Artificial Systems*. Oxford University Press, New York, NY.

- [31] Bonabeau, E., Theraulaz, G., and Deneubourg, J. L. (1996). Quantitative study of the fixed threshold model for the regulation of division of labour in insect societies. *Proc. R. Soc. Lond., B, Biol. Sci.*, 263:1565–1569.
- [32] Bonabeau, E., Theraulaz, G., Deneubourg, J. L., Serge, A., and Camazine, S. (1997). Self-organization in social insects. *Trends Ecol. Evol.*, 12:188–193.
- [33] Bongard, J. C. (2000a). *Genetic Programming*, volume 1802/2004, chapter The Legion System: A Novel Approach to Evolving Heterogeneity for Collective Problem Solving, pages 16–28. Springer Berlin / Heidelberg.
- [34] Bongard, J. C. (2000b). Reducing collective behavioural complexity through heterogeneity. In et al, M. B., editor, *Artificial Life VII: Proceedings of the Seventh International Conference*, pages 327–336.
- [35] Bongard, J. C., Zykov, V., and Lipson, H. (2006). Resilient machines through continuous self-modeling. *Science*, 314:1118–1121.
- [36] Botee, H. and Bonabeau, E. (1998). Evolving ant colony optimization. *Advances in Complex Systems*, 1(2/3):149–159.
- [37] Bourke, A. F. G. and Franks, N. R. (1995). *Social Evolution in Ants*. Princeton University Press, Princeton, NJ.
- [38] Braitenberg, V. (1986). *Vehicles: Experiments in Synthetic Psychology*. Bradford Books and MIT Press, Cambridge, Massachusetts.
- [39] Brooks, R. (1991). Intelligence without Representation. *Artificial Intelligence*, 47(1-3):139–159.
- [40] Brooks, R. A. (1986). A robust layered control system for a mobile robot. *Robotics and Automation, IEEE Journal of [legacy, pre-1988]*, 2(1):14–23.
- [41] Brooks, R. A. (1990). Elephants dont play chess. *Robotics and Autonomous Systems*, 6:01.03.2015.
- [42] Bryant, B. and Miikkulainen, R. (2003). Neuroevolution for adaptive teams. In *Proceedings of the 2003 Congress on Evolutionary Computation, CEC '03*, volume 3, pages 2194–2201.

- 
- [43] Bull, L. and Holland, O. (1997). Evolutionary computing in multiagent environments: Eusociality. *Proceedings of the Annual Conference On Genetic Programming*, pages 347–352.
- [44] Calabi, P. (1988). *Behavioral Flexibility in Hymenoptera: A Re-Examination of the Concept of Caste*, pages 237–258. J. C. Trager (Ed.), Leiden: Brill Press.
- [45] Calderone, N. W. and Page Jr, R. E. (1988). Genotypic variability in age polyethism and task specialisation in the honey bee, *apis mellifera*. *Behav. Ecol. Sociobiol.*, 22:17–25.
- [46] Camazine, S., Franks, N., Sneyd, J., Bonabeau, E., Deneubourg, J., and Theraula, G. (2003). *Self-Organization in Biological Systems*. Princeton University Press.
- [47] Cangelosi, A. and Parisi, D. (1998). The emergence of a ‘language’ in an evolving population of neural networks. *Connection Science*, 10:83–97.
- [48] Cao, Y. U., Fukunage, A. S., and Kahng, A. B. (1997). Cooperative mobile robotics: Antecedents and directions. In *Expanded Version of 1995 IEEE/RSJ IROS Conference Proceedings*, pages 7–27.
- [49] Caprari, G. (2003). *Autonomous Micro-Robots: Applications and Limitations*. PhD thesis, Ecole polytechnique fédérale de Lausanne (EPFL).
- [50] Caprari, G. and Siegwart, R. (2005). Mobile micro-robots ready to use: Alice. In *Proceedings of IEEE/RSJ International Conference On Intelligent Robots and Systems*.
- [51] Chaimowicz, L., Campos, M., and Kumar, V. (2001). Simulating Loosely and Tightly Coupled Multi-Robot Cooperation. 5th Brazilian Symposium on Intelligent Automation (SBAI). Canela, RS, Brazil.
- [52] Charnov, E. L. (1976). Optimal foraging, the marginal value theorem. *J. Theor. Biol.*, 9(2):129–136.
- [53] Cliff, D., Harvey, I., and Husbands, P. (1993). Explorations in evolutionary robotics. *Adaptive Behavior*, 2:73–110.
- [54] Cohen, M. D., Riolo, R. L., and Axelrod, R. (2001). The role of social structure in the maintenance of cooperative regimes. *Rationality and Society*, 13:5–32.



- [55] Collins, R. J. and Jefferson, D. R. (1992). Antfarm: Towards simulated evolution. In Langton, C. G., Taylor, C., Farmer, J. D., and Rasmussen, S., editors, *Artificial Life II*, pages 579–601, Redwood City, CA. Addison-Wesley.
- [56] Corning, P. (1996). The co-operative gene: on the role of synergy in evolution. *Evolutionary Theory*, 11:183–207.
- [57] Crozier, R. H. and Pamilo, P. (1996). *Evolution of Social Insect Colonies: Sex Allocation and Kin Selection*. Oxford University Press, New York, NY.
- [58] Dawkins, R. (1976). *The Selfish Gene*. Oxford University Press, New York, NY.
- [59] Dawkins, R. (1999). *The Extended Phenotype: The Long Reach of the Gene*. Oxford University Press, New York, NY.
- [60] Doebeli, M. and Hauert, C. (2005). REVIEWS AND SYNTHESSES: Models of cooperation based on the Prisoner’s Dilemma and the Snowdrift game. *Ecology Letters*, 8(7):748–766.
- [61] Dorigo, M. and Stuetzle, T. (2004). *Ant Colony Optimisation*. MIT Press, Cambridge, Massachusetts.
- [62] Dorigo, M., Trianni, V., Scdilahin, E., Gross, R., Labella, T. H., Baldassarre, G., Nolfi, S., Deneubourg, J. L., Mondada, F., Floreano, D., and Gambardella, L. M. (2004). Evolving self-organizing behaviors for a swarm-bot. *Autonomous Robots*, 17(2-3):223–245.
- [63] Dudek, G., Jenkin, M., Milios, E., and Wilkes, D. (2002). *A Taxonomy of Multirobot Systems*, chapter 1, pages 1–26. AK Peters, Ltd., 1st edition.
- [64] Dudek, G., Jenkin, M. R. M., Milios, E., and Wilkes, D. (1996). A taxonomy for multi-agent robotics. *Autonomous Robots*, 3:375–397.
- [65] Dürr, P., Mattiussi, C., and Floreano, D. (2006). Neuroevolution with Analog Genetic Encoding. In *Parallel Problem Solving from Nature - PPSN iX*, volume 9 of *Lecture Notes in Computer Science*, pages 671–680.
- [66] Farinelli, A., Iocchi, L., and Nardi, D. (2004). Multirobot systems: a classification focused on coordination. *IEEE Transactions On Systems, Man and Cybernetics, Part B*, 34:2015–2028.



- 
- [67] Fewell, J. H. and Page Jr, R. E. (2000). Colony-level selection effects on individual and colony foraging task performance in honeybees. *Behav. Ecol. Sociobiol.*, 48:173–181.
- [68] Ficici, S. G., Watson, R. A., and Pollack, J. B. (1999). Embodied evolution: A response to challenges in evolutionary robotics. In *Eighth European Workshop On Learning*, pages 14–22.
- [69] Flake, G. W. (1998). *The Computational Beauty of Nature: Computer Explorations of Fractals, Chaos, Complex Systems, and Adaptation*. "A Bradford Book" and MIT Press, Cambridge, Massachusetts.
- [70] Floreano, D. (1997). Reducing human design and increasing adaptivity in evolutionary robotics. In Gomi, T., editor, *Evolutionary Robotics*, pages 187–220.
- [71] Floreano, D., Epars, Y., Zufferey, J. C., and Mattiussi, C. (2005). Evolution of spiking neural circuits in autonomous mobile robots. *International Journal of Intelligent Systems*.
- [72] Floreano, D., Mitri, S., Magnenat, S., and Keller, L. (2007). Evolutionary Conditions for the Emergence of Communication in Robots. *Current Biology*, 17:514–519.
- [73] Floreano, D. and Mondada, F. (1996). Evolution of homing navigation in a real mobile robot. *IEEE Transactions On Systems, Man, and Cybernetics—Part B: Cybernetics*, 26(3):396–407.
- [74] Floreano, D. and Urzelai, J. (2000). Evolutionary robots with on-line self-organization. *Neural Networks*, 13:431–443.
- [75] Foster, K. R., Wenseleers, T., and Ratnieks, L. W. (2006). Kin selection is the key to altruism. *Trends Ecol. Evol.*, 21:57–60.
- [76] Frank, S. A. (1995). Mutual policing and repression of competition in the evolution of cooperative groups. *Nature*, 377:520–522.
- [77] Frank, S. A. (1998). *Foundations of Social Evolution*. Princeton University Press, Princeton, NJ.

- [78] Gautrais, J., Theraulaz, G., Deneubourg, J. L., and Anderson, C. (2002). Emergent polyethism as a consequence of increased colony size in insect societies. *J. Theor. Biol.*, 215:363–373.
- [79] Gerkey, B., Vaughan, R., Stoy, K., Howard, A., Sukhatme, G., and Mataric, M. (2001). Most valuable player: a robot device server for distributed control. In *Intelligent Robots and Systems, 2001. Proceedings. 2001 IEEE/RSJ International Conference on*, volume 3.
- [80] Goldberg, D. E. (1989). *Genetic Algorithms in Search Optimization & Machine Learning*. Addison-Wesley, Reading, MA.
- [81] Grafen, A. (1984). Natural selection, kin selection and group selection. *Behavioural Ecology: An Evolutionary Approach*, pages 62–84.
- [82] Grafen, A. (1985). A geometric view of relatedness. *Oxford Surveys in Evolutionary Biology*, 2:28–90.
- [83] Grefenstette, J. (1988). Credit assignment in rule discovery systems based on genetic algorithms. *Machine Learning*, 3(2):225–245.
- [84] Griffin, A. S. and West, S. A. (2002). Kin selection: Fact and fiction. *Trends Ecol. Evol.*, 17(1):15–21.
- [85] Griffin, A. S., West, S. A., and Buckling, A. (2004). Cooperation and competition in pathogenic bacteria. *Nature*, 430:1024–1027.
- [86] Grimm, V. (1999). Ten years of individual-based modelling in ecology: what have we learned and what could we learn in the future? *Ecological Modelling*, 115(2):129–148.
- [87] Gross, R. and Dorigo, M. (2004a). Cooperative transport of objects of different shapes and sizes. *Lecture Notes in Computer Science*, 3172:106–117.
- [88] Gross, R. and Dorigo, M. (2004b). Evolving a cooperative transport behavior for two simple robots. *Lecture Notes in Computer Science*, 2936:305–316.
- [89] Hamilton, W. D. (1964a). The genetical evolution of social behavior i. *J. Theor. Biol.*, 7:1–31.

- 
- [90] Hamilton, W. D. (1964b). The genetical evolution of social behavior ii. *J. Theor. Biol.*, 7:17–52.
- [91] Hamilton, W. D. (1964c). The Genetical Evolution of Social Behavior I+II. *Journal of Theoretical Biology*, 7:1–52.
- [92] Hamilton, W. D. (1972). Altruism and related phenomena, mainly in social insects. *Annual Review of Ecology and Systematics*, 3:193–232.
- [93] Hammerstein, P. and Leimar, O. (2006). Cooperating for direct fitness benefits. *Journal of Evolutionary Biology*, 19(5):1400–1402.
- [94] Hara, A. and Nagao, T. (1999). Emergence of cooperative behavior using ADG; Automatically Defined Groups. In Banzhaf, W., Daida, J., Eiben, A. E., Garzon, M. H., Honavar, V., Jakiela, M., and Smith, R. E., editors, *Proceedings of the 1999 Genetic and Evolutionary Computation Conference (GECCO-99)*, pages 1038–1046. Morgan Kaufmann.
- [95] Harvey, I., Husbands, P., Cliff, D., Thompson, A., and Jakobi, N. (1997). Evolutionary robotics: the Sussex approach.
- [96] Hauert, C., De Monte, S., Hofbauer, J., and Sigmund, K. (2002). Replicator dynamics for optional public good games. *Journal of Theoretical Biology*, 218:187–194.
- [97] Hauert, C. and Doebeli, M. (2004). Spatial structure often inhibits the evolution of cooperation in the snowdrift game. *Nature*, 428:643–646.
- [98] Hauert, C. and Stenull, O. (2002). Simple adaptive strategy wins the prisoner’s dilemma. *J Theor Biol*, 218:261–272.
- [99] Haynes, T. and Sen, S. (1997). Crossover operators for evolving a team. In Koza, J. R., Deb, K., Dorigo, M., Fogel, D. B., M. Garzon, H. I., and Riolo, R. L., editors, *Genetic Programming 1997: Proceedings of the Second Annual Conference*, pages 162–167. Morgan Kaufmann: Stanford University, CA, USA.
- [100] Haynes, T., Wainwright, R., Sen, S., and Schoenefeld, D. (1995). Strongly typed genetic programming in evolving cooperation strategies. In Forrest, S., editor, *Proceedings of the Sixth International Conference on Genetic Algorithms*, pages 271–278. Morgan Kaufman.

- [101] Holland, O. (1997). Grey Walter: the pioneer of real artificial life. In *Proceedings of the 5th International Workshop on Artificial Life*.
- [102] Holland, O. and Melhuish, C. (1999). Stigmergy, self-organisation, and sorting in collective robotics. *Artificial Life*, 5(2):173–202.
- [103] Holland, O., Woods, J., De Nardi, R., and Clark, A. (2005). Beyond swarm intelligence : The ultraswarm. *IEEE Swarm Intelligence Symposium SIS2005*.
- [104] Hölldobler, B. and Wilson, E. (1994). *Journey to the ants. A story of scientific exploration*. Harvard University Press.
- [105] Huang, Z. Y. (1992). Honeybee colony integration: Worker-worker interactions mediate hormonally regulated plasticity in division of labor. *Proc. Natl. Acad. Sci. U.S.A.*, 89:11726–11729.
- [106] Huang, Z. Y. and Robinson, G. E. (1996). Regulation of honey bee division of labor by colony age demography. *Behav. Ecol. Sociobiol.*, 39:147–158.
- [107] Hunt, G. J., Page, R. E., Fondrk, M. K., and Dullum, C. J. (1995). Major quantitative trait loci affecting honey bee foraging behavior. *Genetics*, 141:1537–1545.
- [108] Husbands, P., Harvey, I., Cliff, D., and Miller, G. F. (1997). Artificial evolution: A new path for artificial intelligence. *Brain and Cognition*, 34:130–159.
- [109] Jakobi, N. (1997a). Evolutionary robotics and the radical envelope of noise. *Adaptive Behavior*, 6(2):325.
- [110] Jakobi, N. (1997b). Half-baked, ad-hoc, and noisy: Minimal simulations for evolutionary robotics. In *Proc. of the Fourth European Conf. On Artificial Life*, pages 348–357. MIT Press.
- [111] Jakobi, N. (1998). Running across the reality gap: Octopod locomotion evolved in a minimal simulation. *EvoRobots*, pages 39–58.
- [112] Jakobi, N., Husbands, P., and Harvey, I. (1995). Noise and the reality gap: The use of simulation in evolutionary robotics. *Lecture Notes in Computer Science*, 929:704–720.

- 
- [113] Jim, K.-C. and Giles, C. L. (2001). Talking helps: Evolving communicating agents for the predator-prey pursuit problem. *Artificial Life*, 6:237–254.
- [114] Johnson, M. and Gaines, M. (1990). Evolution of Dispersal: Theoretical Models and Empirical Tests Using Birds and Mammals. *Annual Review of Ecology and Systematics*, 21:449–480.
- [115] Kauffman, S. A. (1969). Metabolic stability and epigenesis in randomly constructed genetic nets. *J. Theor. Biol.*, 22:437–467.
- [116] Keijzer, F. (2001). *Representation and Behavior*. The MIT Press, Cambridge, Massachusetts.
- [117] Keller, L. (1997). Indiscriminate altruism: Unduly nice parents and siblings. *Trends Ecol. Evol.*, 12:99–103.
- [118] Keller, L., editor (1999). *Levels of Selection in Evolution*. Princeton University Press Princeton, NJ.
- [119] Keller, L. and Chapuisat, M. (2002). *Eusociality and Cooperation*, chapter 1, pages 1–8. Nature Publishing Group, London.
- [120] Keller, L. and Ross, K. G. (1998). Selfish genes: a green beard in the red fire ant. *Nature*, 394:573–575.
- [121] Kennedy, J. and Eberhart, R. (1995). Particle swarm optimization. In *Neural Networks, 1995. Proceedings., IEEE International Conference on*, volume 4, pages 1942–1948.
- [122] Kessin, R. H. (2000). Cooperation can be dangerous. *Nature*, 408:917–919.
- [123] Koza, J. (1992). *Genetic Programming: on the programming of computers by means of natural selection*. Bradford Book.
- [124] Krieger, M. J. B., Billeter, J. B., and Keller, L. (2000). Ant-like task allocation and recruitment in cooperative robots. *Nature*, 406:992–995.
- [125] Lehmann, L. and Keller, L. (2006). The evolution of cooperation and altruism: A general framework and a classification of models. *Journal of Evolutionary Biology*, 19:1365–1376.

- [126] Lehmann, L., Keller, L., West, S., and Roze, D. (2007). Group selection and kin selection: Two concepts but one process. *Proceedings of the National Academy of Sciences*, 104(16):6736.
- [127] Lohn, J. and Reggia, J. (1995). Discovery of self-replicating structures using a genetic algorithm. In *Evolutionary Computation, 1995., IEEE International Conference on*, volume 2, pages 678–683.
- [128] Luke, S. (1998). Genetic programming produced competitive soccer softbot teams for robocup97. In Banzhaf, W., Chellapilla, K., Deb, K., Dorigo, M., Fogel, D., Garzon, M., Goldberg, D., H., I., Koza, J., and Riolo, R., editors, *Genetic Programming 1998: Proceedings of the Third Annual Conference*, pages 214–222. Morgan Kaufmann.
- [129] Luke, S., Hohn, C., Farris, J., Jackson, G., and Hendler, J. (1997). Co-evolving soccer softbot team coordination with genetic programming. In *Proceedings of The First International Workshop On RoboCup, IJCAI-97, Nagoya*, pages 214–222. Springer-Verlag.
- [130] Luke, S. and Spector, L. (1996). Evolving teamwork and coordination with genetic programming. In *Proceedings of Genetic Programming 1996 (GP96), Stanford*, pages 150–156.
- [131] MacColl, A. and Hatchwell, B. (2004). Determinants of lifetime fitness in a cooperative breeder, the long-tailed tit *Aegithalos caudatus*. *Journal of Animal Ecology*, 73(6):1137–1148.
- [132] Magnenat, S. and Waibel, M. (2003-2007). Enki, a fast physics-based 2D simulator. <http://teem.epfl.ch>.
- [133] Magnenat, S., Waibel, M., and Beyeler, A. (2003-2007). Teem, an open evolutionary framework. <http://teem.epfl.ch>.
- [134] Marocco, D. and Nolfi, S. (2006). Self-organization of communication in evolving robots. In Rocha, L. M., Yaeger, L. S., Bedau, M. A., Floreano, D., Goldstone, R. L., and Vespignani, A., editors, *Proceeding of the Tenth International Conference On Artificial Life: AlifeX*, pages 199–205. Boomington: MIT Press.

- 
- [135] Martinoli, A. (1999a). *Chapter 5: Evolution of Distributed Control Strategies*. PhD thesis, EPFL.
- [136] Martinoli, A. (1999b). *Swarm Intelligence in Autonomous Collective Robotics: From Tools to the Analysis and Synthesis of Distributed Control Strategies*. PhD thesis, Ecole polytechnique fédérale de Lausanne (EPFL).
- [137] Mataric, M. and Cliff, D. (1996). Challenges in evolving controllers for physical robots. *Robotics and Autonomous Systems*, 19(1):67–83.
- [138] Mataric, M. J., Sukhatme, G. S., and Ostergaard, E. (2003). Multi-robot task allocation in uncertain environments. *Autonomous Robots*, 14(02.03.2004):255–263.
- [139] Mattiussi, C. (2005). *Evolutionary synthesis of analog networks*. PhD thesis, Ecole polytechnique fédérale de Lausanne (EPFL), Lausanne.
- [140] Mattiussi, C., Dürr, P., and Floreano, D. (2007a). Center of Mass Encoding: A self-adaptive representation with adjustable redundancy for real-valued parameters. In *GECCO 2007*.
- [141] Mattiussi, C. and Floreano, D. (2007). Analog Genetic Encoding for the Evolution of Circuits and Networks. *IEEE Transactions on Evolutionary Computation*.
- [142] Mattiussi, C., Marbach, D., Dürr, P., and Floreano, D. (2007b). The Age of Analog Networks. *AI Magazine*.
- [143] Mattiussi, C., Waibel, M., and Floreano, D. (2004). Measures of diversity for populations and distances between individuals with highly reorganizable genomes. *Evolutionary Computation*, 12(4):495–515.
- [144] May, R. M. (1976). Simple mathematical models with very complicated dynamics. *Nature*, 261:459.
- [145] Maynard Smith, J. (1982). *Evolution and the Theory of Games*. Cambridge University Press, Cambridge, UK.
- [146] Maynard Smith, J. and Szathmáry, E. (1995). *The Major Transitions in Evolution*. Oxford University Press, New York, NY.



- [147] Maynard Smith, J. and Szathmáry, E. (1999). *The Origins of Life*. Oxford University Press, New York, NY.
- [148] McGill, R., Tukey, J., and Larsen, W. (1978). Variations of box plots. *The American Statistician*, 32(1):12–16.
- [149] Michel, O. (2004). Webots™: Professional Mobile Robot Simulation. *Arxiv preprint cs.RO/0412052*.
- [150] Michod, R. (2000). *Darwinian Dynamics: Evolutionary Transitions in Fitness and Individuality*. Princeton University Press.
- [151] Miconi, T. (2003). When evolving populations is better than coevolving individuals: The blind mice problem. In Gottlob, G. and Walsh, T., editors, *Proceedings of the Eighteenth International Joint Conference on Artificial Intelligence*, pages 647–652. Morgan Kaufmann.
- [152] Miglino, O., Lund, H., and Nolfi, S. (1996). Evolving mobile robots in simulated and real environments. *Artificial Life*, 2(4):417–434.
- [153] Miglino, O., Nafasi, K., and Taylor, C. (1994). Selection for Wandering Behavior in a Small Robot. *Artificial Life*, 2(1):101–116.
- [154] Minar, N., Burkhart, R., Langton, C., and Askenazi, M. (1996). The Swarm Simulation System: A Toolkit for Building Multi-Agent Systems. Technical report, Santa Fe NM: Santa Fe Institute Working Paper.
- [155] Minsky, M. (1961). Steps toward artificial intelligence. *Proceedings of the Institute of Radio Engineers*, 49:8–30. Reprinted in *Computers and Thought*, McGraw-Hill, 1963.
- [156] 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.
- [157] Mondada, F., Bonani, M., Guignard, A., Magnenat, S., Studer, C., and Floreano, D. (2005). Superlinear Physical Performances in a SWARM-BOT. In *Proceedings of the 8th European Conference in Artificial Life (ECAL05)*, pages 282–291. Springer.



- 
- [158] Mondada, F., Pettinaro, G. C., Guignard, A., Kwee, I. W., Floreano, D., De-neubourg, J. L., Nolfi, S., Gambardella, L. M., and Dorigo, M. (2004). Swarm-bot: A new distributed robotic concept. *Autonomous Robots*, 17:193–221.
- [159] Murphy, M. and O’Neill, L., editors (1995). *What is Life? The Next Fifty Years: Speculations on the Future of Biology*. Cambridge University Press, Cambridge, UK.
- [160] Noda, I. (1995). Soccer server: a simulator of robocup. Proceedings of the JSAI AI-Symposium.
- [161] Nolfi, S. and Floreano, D. (2000). *Evolutionary Robotics. The Biology, Intelligence, and Technology of Self-organizing Machines*. MIT Press, Cambridge, Massachusetts. 2001 (2nd print), 2000 (1st print).
- [162] Nolfi, S. and Parisi, D. (1997). Learning to adapt to changing environments in evolving neural. *Adaptive Behavior*, 5(1):75–98.
- [163] Nowak, M.; Sigmund, K. (1993). A strategy of win-stay, lose-shift that outperforms tit-for-tat in the prisoner’s dilemma game. *Nature*, 364:12–13.
- [164] Nowak, M. and Sigmund, K. (1998). Evolution of indirect reciprocity by image scoring. *Nature*, 393:573–575.
- [165] Nowak, M. and Sigmund, K. (2004). Evolutionary dynamics of biological games. *Science*, 303:793–799.
- [166] O’Donnell, S. (1998). Genetic effects on task performance, but not on age polyethism, in a swarm-founding eusocial wasp. *Anim. Behav.*, 55:417–426.
- [167] Oli, M. (2003). Hamilton goes empirical: estimation of inclusive fitness from life-history data. *Proceedings: Biological Sciences*, 270(1512):307–311.
- [168] Page, R. E., Erber, J., and Fondrk, M. K. (1998). The effect of genotype on response thresholds to sucrose and foraging behavior of honey bees (*apis mellifera* l.). *J. Comp. Physiol. A*, 182:489–500.
- [169] Page, R. E., Fondrk, M. K., Hunt, G. J., Guzmán-Novoa, E., Humphries, M. A., Nguyen, K., and Green, A. S. (2000). Genetic dissection of honeybee (*apis mellifera* l.) foraging behavior. *J. Hered.*, 91:474–479.

- [170] Pamilo, P. (1989). Estimating relatedness in social groups. *Trends in Ecology & Evolution*, 4(11):353–355.
- [171] Panait, L. and Luke, S. (2005). Cooperative multi-agent learning: The state of the art. *Autonomous Agents and Multi-Agent Systems*, 11:387–434.
- [172] Parker, L. E. (2000). Current state of the art in distributed autonomous mobile robotics. In Parker, L. E., Bekey, G., and Barhen, J., editors, *Distributed Autonomous Robotic Systems*, volume 4, pages 3–12. Springer-Verlag, Tokyo.
- [173] Pettinaro, G., Kwee, I., and Gambardella, L. (2003). Definition, Implementation, and Calibration of the Swarmbot3D Simulator. Technical report, Istituto Dalle Molle di studi sull' intelligenza artificiale (IDSIA/USI-SUPSI).
- [174] Pfeifer, R. (2002). On the role of embodiment in the emergence of cognition: Grey walter's turtles and beyond. In *Proceedings of the Workshop "The Legacy of Grey Walter"*.
- [175] Pfeifer, R. and Bongard, J. (2006). *How the Body Shapes the Way We Think: A New View of Intelligence (Bradford Books)*. The MIT Press, Cambridge, Massachusetts.
- [176] Potter, M. A. and De Jong, D. V. (1994). A cooperative coevolutionary approach to function optimization. *Lecture Notes in Computer Science*, 866:249–257.
- [177] Potter, M. A. and De Jong, D. V. (2000). Cooperative coevolution: An architecture for evolving coadapted subcomponents. *Evolutionary Computation*, 8:1–29.
- [178] Potter, M. A., Meeden, L. A., and Schultz, A. C. (2001). Heterogeneity in the coevolved behaviors of mobile robots: The emergence of specialists. In *Proceedings of The Seventeenth International Conference on Artificial Intelligence*, pages 1337–1343.
- [179] Pyke, G. H., Pulliam, H. R., and Charnov, E. L. (1977). Optimal foraging: A selective review of theory and tests. *Quarterly Review of Biology*, 52:137–154.
- [180] Queller, D. C. (1994). Genetic relatedness in viscous populations. *Evolutionary Ecology*, 8:70–73.

- 
- [181] Queller, D. C. and Goodnight, K. F. (1989). Estimating relatedness using genetic markers. *Evolution*, 43(2):258–275.
- [182] Quinn, M. (2001). A comparison of approaches to the evolution of homogeneous multi-robot teams. In *Proceedings of the 2001 Congress On Evolutionary Computation: CEC2001*, pages 128–135, Seoul and Korea. IEEE.
- [183] Quinn, M., Smith, L., Mayley, G., and Husbands, P. (2002a). Evolving formation movement for a homogeneous multi-robot system: Teamwork and role-allocation with real robots. *Cognitive Science Research Papers*, 515.
- [184] Quinn, M., Smith, L., Mayley, G., and Husbands, P. (2002b). Evolving teamwork and role-allocation with real robots. In *Artificial Life VIII: Proceedings of the Eighth International Conference*, pages 302–311. MIT Press.
- [185] Quinn, M., Smith, L., Mayley, G., and Husbands, P. (2003). Evolving controllers for a homogeneous system of physical robots: Structured cooperation with minimal sensors. *Philosophical Transactions of the Royal Society of London and Series A: Mathematical and Physical and Engineering Sciences*, 361:2321–2344.
- [186] Raik, S. and Durnota, B. (1994). The evolution of sporting strategies. In Stonier, R. J. and Yu, X. H., editors, *Complex Systems: Mechanisms of Adaption*, pages 85–92, Amsterdam, Netherlands. IOS Press.
- [187] Rainey, P. and Rainey, K. (2003). Evolution of cooperation and conflict in experimental bacterial populations. *Nature*, 425:72–74.
- [188] Reyes Lopez, J. L. (1987). Optimal foraging in seed-harvester ants: Computer-aided simulation. *Ecology*, 68:1630–1633.
- [189] Reynolds, C. (1993). An evolved, vision-based behavioral model of coordinated group motion. In *From Animals to Animats 2*, pages 384–392. MIT Press.
- [190] Richards, M., Whitley, D., Beveridge, J., Mytkowicz, T., Nguyen, D., and Rome, D. (2005). Evolving cooperative strategies for UAV teams. In *Proceedings of the 2005 Conference on Genetic and Evolutionary Computation*, pages 1721–1728, New York, NY, USA. ACM Press.
- [191] Riolo, R. L., Cohen, M. D., and Axelrod, R. (2001). Evolution of cooperation without reciprocity. *Nature*, 414:441–443.

- [192] Robinson, A. and Spector, L. (2002). Using genetic programming with multiple data types and automatic modularization to evolve decentralized and coordinated navigation in multi-agent systems. In *Late-Breaking Papers of the Genetic and Evolutionary Computation Conference (GECCO-2002)*. The International Society for Genetic and Evolutionary Computation.
- [193] Robinson, G. E. (1992). Regulation of division of labor in insect societies. *Annu. Rev. Entomol.*, 37:637–65.
- [194] Robinson, G. E., Page, R. E., Strambi, C., and Strambi, A. (1989). Hormonal and genetic control of behavioral integration in honey bee colonies. *Science*, 246:109–112.
- [195] Ross, K. G. and Keller, L. (1995). Ecology and evolution of social organization: Insights from fire ants and other highly eusocial insects. *Annu. Rev. Ecol. Syst.*, 26:631–656.
- [196] Rowell, J., Ellner, S., and Reeve, H. (2006). Why Animals Lie: How Dishonesty and Belief Can Coexist in a Signaling System. *American Naturalist*, 168:180–204.
- [197] Rüppele, O., Pankiw, T., and Page Jr, R. E. (2004). Pleiotropy, Epistasis and New QTL: The Genetic Architecture of Honey Bee Foraging Behavior. *J. Hered.*, 95:481–491.
- [198] Sachs, J. L., Mueller, U. G., Wilcox, T. P., and Bull, J. J. (2004). The evolution of cooperation. *The Quarterly Review of Biology*, 79:135–160.
- [199] Saunders, G. and Pollack, J. (1996). The evolution of communication schemes over continuous channels. In *From Animals to Animats 4 - Proceedings of the Fourth International Conference on Adaptive Behaviour*.
- [200] Schrödinger, E. (2003). *Was ist Leben?* Piper Verlag GmbH, München, 6th edition.
- [201] Seeley, T. D. (1995). *The Wisdom of the Hive: The Social Physiology of Honey Bee Colonies*. Harvard University Press.
- [202] Seyfarth, R., Cheney, D., and Marler, P. (1980). Monkey responses to three different alarm calls: evidence of predator classification and semantic communication. *Science*, 210(4471):801.

- 
- [203] Sigmund, K. and Nowak, M. (2001). Tides of tolerance. *Nature*, 414(22):403–405.
- [204] Simoes, E. D. V. and Barone, D. A. C. (2002). Predation: An approach to improving the evolution of real robots with a distributed evolutionary controller. In *Proceedings of the 2002 IEEE International Conference On Robotics and Automation*, pages 664–669.
- [205] Smith, T. (1998). Blurred vision: Simulation-reality transfer of a visually guided robot. In *Evolutionary Robotics: First European Workshop, EvoRob'98*, pages 152–164. Springer.
- [206] Sober, E. (2001). *Thinking about Evolution: Historical, Philosophical, and Political Perspectives*, chapter The Two Faces of Fitness, pages 309–321. Cambridge University Press, Cambridge, UK.
- [207] Spector, L., Klein, J., Perry, C., and Feinstein, M. (2003). Emergence of collective behavior in evolving populations of flying agents. *Lecture Notes in Computer Science*, 2723:61–73.
- [208] Spector, L., Klein, J., Perry, C., and Feinstein, M. (2005). Emergence of collective behavior in evolving populations of flying agents. *Genetic Programming and Evolvable Machines*, 6:111–125.
- [209] Stanley, K., Bryant, B., and Miikkulainen, R. (2005). Real-time neuroevolution in the NERO video game. *Evolutionary Computation, IEEE Transactions on*, 9:653–668.
- [210] Stone, P. and Veloso, M. (2000). Multiagent systems: A survey from a machine learning perspective. *Autonomous Robots*, 8:345–383.
- [211] Strassmann, J. E., Zhu, Y., and Queller, D. C. (2000). Altruism and social cheating in the social amoeba *dictyostelium discoideum*. *Nature*, 408:965–967.
- [212] Sycara, K. (1998). Multiagent Systems. *AI Magazine*, 19(2):79–92.
- [213] Tarapore, D., Floreano, D., and Keller, L. (2006). Influence of the level of polyandry and genetic architecture on division of labour. In *The Tenth International Conference on the Simulation and Synthesis of Living Systems, Bloomington, IN*, pages 358–364.

- [214] Taylor, P. D. (1992). Inclusive fitness in a homogeneous environment. *Proc. R. Soc. Lond., B, Biol. Sci.*, 249:299–302.
- [215] Theraulaz, G., Bonabeau, E., and Deneubourg, J. L. (1998). Response threshold reinforcement and division of labour in insect societies. *Proc. R. Soc. Lond., B, Biol. Sci.*, 265:327–332.
- [216] Thompson, A. (1995). Evolving Electronic Robot Controller that Exploit Hardware Resources. In *Proceedings of the European Conference on Artificial Life*, pages 640–656.
- [217] Traniello, J. F. A. (1989). Foraging strategies of ants. *Annual Review Entomology*, 34:191–210.
- [218] Traulsen, A. and Nowak, M. (2006). Evolution of cooperation by multilevel selection. *Proceedings of the National Academy of Sciences*, 103(29):10952.
- [219] Trianni, V., Gross, R., Labella, T. H., Sahin, E., and Dorigo, M. (2003). Evolving aggregation behaviors in a swarm of robots. *Lecture Notes in Computer Science*, 2801:865–874.
- [220] Trianni, V., Nolfi, S., and Dorigo, M. (2006). Cooperative hole-avoidance in a swarm-bot. *Robotics and Autonomous Systems*, 54:97–103.
- [221] Velicer, G. J. and Yu, Y. N. (2003). Evolution of novel cooperative swarming in the bacterium *myxococcus xanthus*. *Nature*, 425:75–78.
- [222] Waibel, M., Floreano, D., Magnenat, S., and Keller, L. (2006). Division of labour and colony efficiency in social insects: Effects of interactions between genetic architecture, colony kin structure and rate of perturbations. *Proceedings of the Royal Society B*, 273:1815–23.
- [223] Waibel, M., Keller, L., and Floreano, D. (2007). Genetic team composition and level of selection in the evolution of multi-agent systems. *IEEE Transactions on Evolutionary Computation*. submitted.
- [224] Ward, C. R., Gobet, F., and Kendall, G. (2001). Evolving collective behavior in an artificial ecology. *Artificial Life*, 7:191–209.

- 
- [225] Watson, R. A., Ficici, S. G., and Pollack, J. B. (1999). Embodied evolution: Embodying an evolutionary algorithm in a population of robots. In *Proceedings of the 1999 Congress on Evolutionary Computation, 1999. CEC 99.*, volume 1, pages 335–342.
- [226] Watson, R. A., Ficici, S. G., and Pollack, J. B. (2002). Embodied evolution: Distributing an evolutionary algorithm in a population of robots. *Robotics and Autonomous Systems*, 39:1–18.
- [227] Webb, B. (2000). What does robotics have to offer animal behaviour? *Anim. Behav.*, 60:545–558.
- [228] Webb, B. (2001). Can robots make good models of biological behaviour? *Behavioral and Brain Sciences*, 24(06):1033–1050.
- [229] Werner, G. M. and Dyer, M. G. (1992a). Evolution of communication in artificial organisms. In Langton, C., Taylor, C., Farmer, D., and Rasmussen, S., editors, *Artificial Life II*, pages 659–687, Redwood City, CA. Addison-Wesley Pub.
- [230] Werner, G. M. and Dyer, M. G. (1992b). Evolution of herding behavior in artificial animals. In *Proceedings of the Second International Conference On From Animals to Animats 2 : Simulation of Adaptive Behavior*, pages 393–399. MIT Press Cambridge, MA, USA.
- [231] West, S., Gardner, A., and Griffin, A. (2006). Altruism. *Current Biology*, 16(13):482–483.
- [232] West, S., Griffin, A., and Gardner, A. (2007a). Social semantics: altruism, cooperation, mutualism, strong reciprocity and group selection. *Journal of Evolutionary Biology*, 20(2):415–432.
- [233] West, S. A., Griffin, A. S., and Gardner, A. (2007b). Evolutionary explanations for cooperation. *Current Biology*, 17:661–672.
- [234] 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*, 409:510–513.



- [235] West, S. A., Pen, I., and Griffin, A. S. (2002). Cooperation and competition between relatives. *Science*, 296:72–75.
- [236] Wiegand, R. P., Liles, W. C., and DeJong, D. V. (2001). An empirical analysis of collaboration methods in cooperative coevolutionary algorithms. *Proceedings of the Genetic and Evolutionary Computation Conference 2001*.
- [237] Wilson, D. (1983). The Group Selection Controversy: History and Current Status. *Annual Review of Ecology and Systematics*, 14(1):159–187.
- [238] Wilson, E. O. (1971). *The Insect Societies*. Harvard University Press, Cambridge, Massachusetts.
- [239] Wilson, E. O. (1985). Between-caste aversion as a basis for division of labor in the ant *Pheidole pubiventris* (hymenoptera: Formicidae). *Behav. Ecol. Sociobiol.*, 17:35–37.
- [240] Wilson, E. O. (2005). Kin selection as the key to altruism: Its rise and fall. *Social Research*, 72:159–166.
- [241] Wilson, E. O. and Hölldobler, B. (2005). Eusociality: Origin and consequences. *PNAS*, 102:13367–13371.
- [242] Wu, A. S., Schultz, A. C., and Agah, A. (1999). Evolving control for distributed micro air vehicles. In *IEEE Conference on Computational Intelligence in Robotics*, pages 174–179.
- [243] Wynne-Edwards, V. (1962). *Animal Dispersion in Relation to Social Behaviour*. Oliver and Boyd. London, GB.
- [244] Yao, X. (1999). Evolving artificial neural networks. *Proceedings of the IEEE*, 87:1423–1447.
- [245] Zaera, N., Cliff, C., and Bruten, J. (1996). (Not) evolving collective behaviours in synthetic fish. In *Proceedings of the Fourth International Conference On Simulation of Adaptive Behavior (SAB'96)*, volume 4, pages 635–644.
- [246] Zufferey, J., Floreano, D., van Leeuwen, M., and Merenda, T. (2002). Evolving Vision-based Flying Robots. In *2nd International Workshop on Biologically Motivated Computer Vision (BMCV'2002)*, Lecture Notes in Computer Science, pages 592–600.



- [247] Zufferey, J., Guanella, A., Beyeler, A., and Floreano, D. (2006). Flying over the Reality Gap: From Simulated to Real Indoor Airships. *Autonomous Robots*, 21(3):243–254. The original publication is available at [www.springerlink.com](http://www.springerlink.com).



---

# Curriculum vitæ

---

On September 29th 1977 I was born in Vienna, Austria. I grew up in Wels, Austria where I received my secondary education from the Bundesgymnasium Dr. Schauerstrasse. In grade 11 I spent an exchange year at Cochrane high school in Alberta, Canada before graduating in Wels in 1996. The same year I took up my studies of Electrical Engineering at the Technical University of Vienna and switched to studying Technical Physics a year later. From 2001 to 2002 I completed an exchange year at the EPFL in Lausanne. In 2003 I graduated from the Technical University of Vienna with a M.Sc. in Technical Physics.

Since November 2003, I pursue my Ph.D. on the Evolution of Cooperation with Prof. Dario Floreano at the Laboratory of Intelligent Systems at the EPFL and Prof. Laurent Keller at the Department of Ecology and Evolution at the University of Lausanne. As part of this project I have co-developed a setup for the automatic evolution of robot teams in hardware, a physics based simulator (Magenat and Waibel [132]) as well as a software framework for evolutionary robotics (Magenat et al. [133]). In addition to the work presented in this thesis I have co-authored an article on measures of genetic diversity (Mattiussi et al. [143]). In parallel to my research activity, I am involved in educational tasks such as the preparation of practical assignments for undergraduate students. I have also proposed and supervised three Master theses and eight semester projects.

In collaboration with Prof. Dario Floreano I have created and coordinated a bi-weekly internet radio podcast on robotics and artificial intelligence. As part of this project I conducted 17 interviews with leading professionals which have been downloaded more than 55000 times. I have also co-organized a workshop on The Emergence of Social Behavior held at the 9th European Conference on Artificial Life in Lisbon, Portugal in September 2007.





