

Virtual life

eco-evolutionary experiments with situated agents

Walter de Back
Department of Philosophy
Utrecht University
walter.deback.net
walter@deback.net

Supervisors

Marco Wiering
Intelligent Systems group
Department for Information and Computing Sciences

Edwin de Jong
Large Distributed Databases group
Department for Information and Computing Sciences

Daniel van der Post
Theoretical Biology / Bioinformatics group
Department of Biology

March 2006

Abstract

The growing interest in the biological roots of cognition leads to the cross-fertilization between the fields of autonomous robotics and artificial life. This requires new tools that facilitate research on the interface between embodied cognitive science and theoretical biology. In this thesis, a model is presented that enables the simulation of evolving ecosystems of situated agents. This model enables studies to the interplay between situated interaction, self-organised collective behaviour and evolution by natural selection.

The use of complex computer simulations as scientific tools requires a theoretical embedding. This is established by analysing and interpreting the results of ecological simulations (of bitrophic and tritrophic food chains) in terms of analytical models of population dynamics. These ordinary differential equation (ODE) models allow us to understand and control the population dynamics that emerge from simulations. Moreover, the evolutionary dynamics observed in eco-evolutionary simulations can be interpreted as changes in the ODE model, which enables us to understand evolvability of certain traits in terms of ecological viability.

In this thesis, several eco-evolutionary experiments are replicated that were previously conducted by more formal models: predator-prey systems, enrichment, tragedy of the commons, evolutionary arms races, red queen effect, evolution of reproductive restraint. The simulation model allows us to gain new insights by relaxing some of the assumptions of these formal models (infinite population sizes, global interactions, spatial homogeneity) and comparing the results. An indirect explanatory framework is used in which one gives explanations of an emergent pattern (e.g. evolutionary dynamics) by relating it to other emergent patterns (population and spatial dynamics), without the need to refer to the specification of the simulation model itself.

The combination of the complex simulation model, its embedding in theoretical ecology and the use of an indirect explanatory framework, provides a valuable new tool for use in research on the edge of artificial life and autonomous robotics, or theoretical biology and embodied cognitive science in general.

Contents

1 Introduction	3
1.1 Problems	4
1.2 Robotics and artificial life	4
1.3 Methodology	6
1.4 Contributions	9
1.5 Outline	11
2 Background	12
2.1 Embodied situated interaction	12
2.2 Collective behavior	14
2.3 Evolution	18
2.4 Conclusion	22
3 Simulation model	23
3.1 Virtual life model	24
3.2 Emergent patterns	32
3.3 Conclusions	40
4 Ecological analysis	41
4.1 Resource and consumer dynamics	42
4.2 Ecological models	45
4.3 Controlled ecosystem	53
4.4 Conclusions	55
5 Evolutionary experiments	56
5.1 Evolution	57
5.2 Coevolution	62
5.3 Multi-level selection	71
5.4 Conclusions	81
6 Conclusions	84
6.1 Summary	84
6.2 Discussion	86

Chapter 1

Introduction

The idea that there is a continuity between life and cognition causes a growing interest into the biological origins of cognition. The common denominator of living and cognitive processes is that they involve self-organization arising from interaction between underlying entities. This perspective has resulted in changes in cognitive science and artificial intelligence, as well as in theoretical biology and ecology, which has caused these fields to grow closer to each other.

This is most clearly demonstrated by cross-fertilization between autonomous robotics and artificial life. Traditionally, the field of autonomous robotics is concerned with the control of an individual embodied situated agent that interacts with its environment. The field of artificial life, instead, traditionally focuses on self-organization of collectives or on population-level evolutionary dynamics. Recently, attention is shifting to research in which the approaches of robotics and artificial life are combined.

In this thesis, we present the virtual life simulation model that aims to facilitate such research. The simulation implements spatially explicit individual-based models in which individuals are situated agents. The interactions between agents and their environment and among each other give rise to multiple self-organised spatial and temporal patterns. Theoretical ecological modelling is used to understand the patterns, in ecological and evolutionary contexts.

After an overview of the scientific embedding of this thesis, in Chapter 1, and its central background concepts, in Chapter 2, the simulation model is presented in Chapter 3. The population dynamics that arise from these simulations are analysed using models from theoretical ecology in Chapter 4. The theoretical understanding that results from these models is employed to gain insight into the results from a series of eco-evolutionary experiments in Chapter 5. Finally, conclusions are drawn and opportunities for future research are discussed in Chapter 6.

1.1 Problems

The last two decades have brought the rise of new modelling approaches in cognitive science, artificial intelligence and theoretical biology. Instead of modelling cognitive behavior in terms of complex internal mechanisms, new embodied approaches in cognitive science focus on the situated interaction between relatively simple robotic agents and their environment (Pfeifer and Scheier, 1999). Theoretical biology has shown a shift from mathematical mini-models of ecological and evolutionary processes towards individual-based and spatially explicit simulation models in which local interactions among individuals result in global phenomena (Hutson et al., 1988; Hogeweg and Hesper, 1990; Grimm, 1999). These developments have resulted in the rise of autonomous robotics and artificial life simulations as common modelling tools for these fields.

Whereas autonomous robotics concentrates on the control and interaction of an individual robot with its environment, artificial life simulations typically focus on group- and population-level processes. The field of artificial life can be divided into studies in either self-organised collective behaviors or evolutionary processes. Recently, the interests of robotics and artificial life have a growing overlap in combinations and interplay of the processes involved in embodied situated interaction, collective behaviors and evolutionary dynamics.

In these converging research fields, new simulation models are needed that facilitate this research. Moreover, methods of theoretical analysis are necessary to enable us to interpret the results of these simulations, and embed them in an established theoretical framework. Both these problems are addressed in this thesis. A simulation model is presented that allows the spatially explicit simulation of evolving ecosystems of populations of situated individuals. And a corresponding theoretical model is constructed that is based on well-studied models in ecology.

1.2 Robotics and artificial life

In new approaches to cognitive science and theoretical biology, the concept of interaction plays a vital role. In embodied approaches to cognitive science, behavior is re-conceptualised as the result of agent-environment interaction, which has moved the focus from cognition as logical reasoning towards situated interaction. Complexity of behavior is largely a reflection of the richness of the environment (Simon, 1969; Brooks, 1985) and can result from embodied agents (figure 1.1a) with simple control systems (Braitenberg, 1984). In a similar vein, a-life attempts to explain collective behavioral phenomena such as bird flocks (figure 1.1b) or coordination in social insects on the basis of self-organization through interactions between groups of relatively simple individuals (Reynolds, 1987; Melhuish and Holland, 1999). Evolutionary approaches to behavior have turned from a concept of evolution as an optimization process towards evolution as adaptation to the problems posed by the ecology.

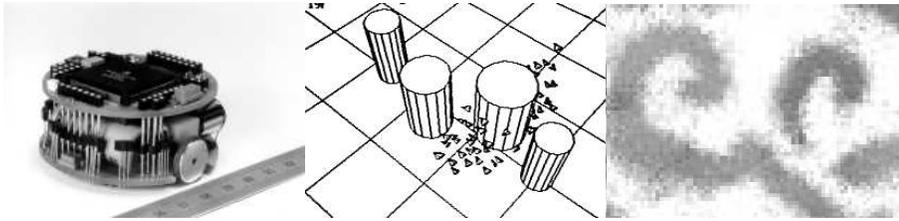


Figure 1.1: (Left) Khepera robot: used in evolutionary robotics (Nolfi and Floreano, 2000). (Middle) BOIDS model: example of collective a-life (Reynolds, 1987). (Right) Self-structured spiral waves in cellular automata model of evolution (Savill and Hogeweg, 1997)

This is reflected in a shift from exogenous fitness models in which selection is imposed externally by user-defined fitness functions towards endogenous fitness models (figure 1.1c) (Forrest and Mitchell, 1994; Ray, 1991; Menczer and Belew, 1996). In such models the fitness of individuals, defined as the rate of reproduction, results from interactions within the system (e.g. survival and reproduction).

The processes of embodied situated interaction, collective behavior and evolution by natural selection models have already been studied extensively as isolated mechanisms (see chapter 2). More recently, interest in combinations of these processes is increasing. In autonomous robotics, the potential of self-organization in collective behaviors and evolutionary design is acknowledged. And in theoretical biology, modern modelling approaches use individual-based and spatially explicit simulation that yield radically different dynamics than predicted by classical population-level models.

Swarm robotics, for example, combines the first two topics in studying collective phenomena that emerge from physically situated robots (Bonabeau et al., 1999). Evolutionary robotics combines robotics with evolutionary optimization, and recently attention has shifted towards less explicit fitness criteria (e.g. through applying coevolution) (Harvey et al., 1997; Nolfi and Floreano, 2000). In theoretical biology renewed attention is given to the controversial theme of group selection which combines self-organised collectives with evolutionary mechanisms by holding that such collectives can serve as a unit of selection and thus influence the course of evolution (Wilson, 1975; Johnson and Boerlijst, 2002; Savill and Hogeweg, 1997). Other studies combine embodied situated interaction with endogenous fitness models by evolving (simulated) robots through natural selection (Yaeger, 1994; Channon, 2000).

Attempts to study the combination and interplay of situated interaction with ecological and evolutionary processes have been rare, however. This thesis constitutes a modest first step towards this end. The present study is focused on simulation and theoretical understanding of ecological and evolutionary phenomena that arise from interacting situated agents.

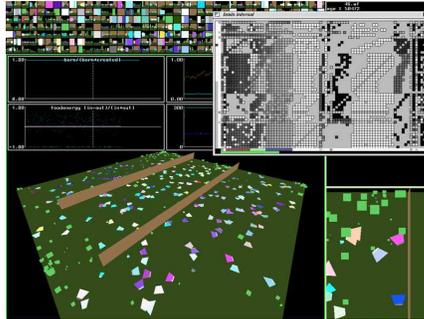


Figure 1.2: PolyWorld (Yaeger, 1994)

1.3 Methodology

A rare example of an artificial life model that simulates situated agents to study ecological and evolutionary processes is Yaeger's PolyWorld (see fig. 1.2). This simulator models an environment that is inhabited by an evolving population of complex situated individuals. As an early simulation model of astonishing complexity, it has been fruitful both in showing the potential of this modelling approach, as in showing the difficulties in interpreting the results of such simulations and thereby their use as tools in biological sciences. This is perhaps best described by Yaeger himself:

PolyWorld brings together biologically motivated genetics, simple simulated physiologies and metabolisms, Hebbian learning in arbitrary neural network architectures, a visual perceptive mechanism, and a suite of primitive behaviors in artificial organisms grounded in an ecology just complex enough to foster speciation and inter-species competition. Predation, mimicry, sexual reproduction, and even communication are all supported in a straightforward fashion. The resulting survival strategies, both individual and group, are purely emergent, as are the functionalities embodied in their neural network brains. Complex behaviors resulting from the simulated neural activity are unpredictable, and change as natural selection acts over multiple generations.

...

[The interpretation of these results] requires some well designed and implemented graphical analysis tools. For now, ethological level behaviors may be the best way to begin developing some understandings and intuitions about the evolutionary dynamics possible in such a system. (Yaeger, 1994)

Despite the wealth of interesting emergent behavioural, ecological and evolutionary processes observed from the simulations, the analysis of them entails

little more than anthropomorphic description of observed behaviours. It is unclear how the labelling of observed behaviours with exotic names like 'frenetic joggers' and 'indolent cannibals' contributes to a deeper understanding of the simulation results or biological science in general. Indeed, the benefits of simulation models are lost when the results are almost as hard to analyse as the real world.

This is not necessarily the case for complex artificial life simulations, as this thesis attempts to point out. There is, however, a need for a theoretical embedding in the initial construction of the model to be able to interpret the results of simulations in a later stage. A theoretical understanding has been an integral part of the development of the virtual life simulation model. This starts from a general perspective on the use of computer simulations in scientific inquiry.

Simulation as scientific tools¹

Computer simulations are among the most flexible and powerful new tools for theoretical development in biology. This does not imply, however, that they are necessarily easier to understand or more useful than other tools such as purely mathematical models. Although it is relatively easy to construct computer models that simulate complex situations that go beyond mathematical tractability, this can be a disadvantage with respect to scientific inquiry aimed at understanding biological processes. When simulations are too complex or too different from existing (mathematical) models, their scientific value is hard to assess since they defy comparison to existing biological theory.

A fruitful way of incorporating simulations into scientific activity is to consider them as ordinary tools (like hammers or microscopes) that are constructed to overcome our bodily shortcomings. Scientific activity can be seen as a cognitive form of skillful activity aimed at gaining a maximal grip on the environment in which we find ourselves (Merleau-Ponty, 1943; Dreyfus). A tool, or rather the *use* of a tool, serves as a corporeal enhancement by elaborating on the range of skillful activities of the user. The use of a microscope, for example, extends our visual abilities to discriminate between very small objects. Its use only becomes meaningful to the user when he can embed this particular form of visual experiences in the frameworks of embodied skills he is already familiar with through a process of skill acquisition.

Similarly, the use of computer simulations becomes meaningful and useful only relative to the degree of integration into existing theoretical frameworks. The value of a new scientific tool can only be assessed if it allows comparison to common norms. For the introduction of new simulation tools to be successful, the new tool should be accompanied by a body of theory that validates its use. This does not mean that simulations do not present genuine *new* tools merely because they are validated by existing theory. New tools can open new

¹The following two paragraphs are roughly based on the excellent overview of methodological issues by di Paolo (1999, ch 4).

grounds for research by (re)moving parts of the old web of constraints (as is illustrated by the development of cell biology following the invention of the microscope by Antonius van Leeuwenhoek).

The advantages of computer simulations lie mainly in the fact that they enable exploration of the emergent properties from local individual interactions. This renders simulation models new and potentially very useful. Within the domain of biology, the study of the emergence of new levels of organizations (e.g. cells, individuals, societies) and the interactions between the dynamics of the various levels is essential. New theoretical modelling shows that many traditional distinctions are not as clear-cut as they were supposed to be. In contrast to the habit of separating ecological and evolutionary timescale, for example, new eco-evolutionary models in which this simplification was lifted show that the interplay between these dynamics have strong influences on each other.

Direct and indirect explanation

A large part of research in artificial life is conducted to show and understand the emergence of complex global dynamics from simple local interactions. The power of self-organization has been demonstrated in early work such as the simulation of bird flocking by Reynolds (1987). Reynold's model (see figure 1.1) acted as a proof of concept by showing that a complex phenomenon such as coordinated collective behaviour can be reproduced by invoking three simple local rules between interacting individuals. This way of using computer simulations has been fruitful to gain recognition for self-organization as an important property of biological systems, which has since become widely accepted. What is also increasingly recognized, however, is the limitation of this approach, because it proceeds on the conjecture that the ability to replicate a certain phenomenon does not imply understanding of how the pattern arises from the model. In many cases, it may be difficult or even impossible to precisely know what aspects of the model are involved and how they relate. In any case, the explanation of the observed pattern is done by relating the phenomena to the simulation model directly (see left panel in figure 1.3). This is only one, rather limited, approach to the use of simulation models in biology.

A much more interesting and powerful way of using computer simulations is in applying an indirect way of explanation. Artificial life simulation models often give rise to multiple emergent patterns simultaneously². Some of these observations can be explained by the basic model itself, but others may require to be explained in terms of other emergent patterns (see right panel in figure 1.3).

The present study attempts to provide insight into (1) ecological dynamics (e.g. population dynamics and spatial self-organization) by relating them to the local interacting agents and (2) evolutionary dynamics (e.g. coevolutionary arms races and group selection) by relating them to ecological dynamics.

²This depends, of course, on the number of free parameters in the model and the choice of observables that show the patterns.

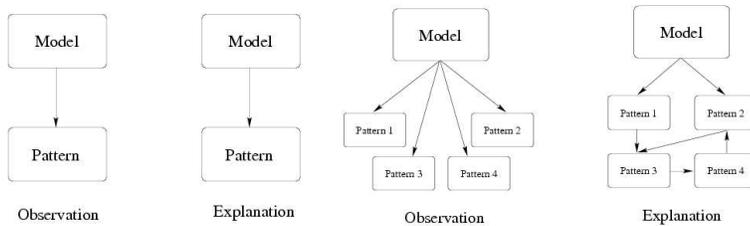


Figure 1.3: Direct explanation (left) and indirect explanation (right) (di Paolo, 1999)

As already noted, the difficulty of interpreting results from simulation models is often due to a lack of embedding in existing theoretical frameworks. Two approaches are taken as countermeasures to this problem. First, the simulation model is kept simple enough to allow meaningful analysis. Second, its results are embedded in established mathematical frameworks that form points of reference in our effort to explore the consequences of lifting the implicit assumption in these mathematical frameworks.

Apart from the methodological considerations raised above, simulations should be grounded in biological theory with respect to the subject under investigation to allow meaningful interpretation of its results. Existing ODE models in theoretical biology can be taken as a starting point. Replication of the results of such models by simulation models followed by an exploration of the consequences of relaxing the explicit or implicit assumptions within the ODE models can provide new insight and help theory development (Miller, 1996). Following this line of thought, most of the experiments described in this thesis are replications of ecological and eco-evolutionary experiments that were previously conducted in (more) formal models.

1.4 Contributions

This thesis aims to contribute to the fields of embodied cognitive science and theoretical biology by (further) blurring the distinction between these fields and thereby establishing modern theoretical biology as a fertile ground for studies in cognitive science, based on the continuity of life and cognition. The contribution to theoretical biology consists of the development of the simulation model and methods of analysis of simulation results. The content of the various experiments conducted in this simulation-analysis framework aims to contribute to a cognitive science that is embedded in theoretical biology.

1.4.1 Contributions to Artificial Life

Development of simulation platform

A simulation model is developed that facilitates the study of the interplay between situated interaction, self-organization and evolution by natural selection. The result is a spatially-explicit individual-based model in which ecological and evolutionary processes are simulated. Like most artificial life simulations, the model is defined on the level of (situated) individuals, while our main interest is at the level of (evolving) populations. Although individuals are modelled as simple agents throughout this thesis to keep results analytically tractable, the model is easily extended to more complex morphologies, more elaborate (neural) control systems and more complex agent-agent interactions (e.g. signalling).

Development of methodology for theoretical analysis and interpretation

To render the results of these simulations useful for scientific enquiry, the results are interpreted in the well-established framework of theoretical ecology. Analytical ODE models enable us to predict, manipulate and control the population dynamics that emerge from the simulations. Moreover, they enable us to understand the ecological constraints on evolvability.

Most work in artificial life is restricted to a 'simple to complex' paradigm and shows that simple local interactions lead to the complex global patterns by employing a direct explanation model. The simulations presented in this thesis are an attempt to transcend this paradigm towards the 'complex to complex' by employing indirect explanations, in which the emergence of one pattern is explained in terms of its relations to other emergent patterns. This explanatory strategy renders explanations more generic because they depend less on implementation details of the simulation model. It can therefore be considered a contribution to the use of computer simulations in scientific research in general.

1.4.2 Contributions to Cognitive Artificial Intelligence (CAI)

Although (the evolution of) cognition is not explicitly studied in this thesis, the eco-evolutionary simulations presented here are relevant to cognitive science because they examine the biological processes in ecology and evolution that underly the emergence of cognitive behaviour.

Cognition is an adaptation to ecological problems

From an evolutionary perspective, cognition is considered as an adaptation to problems posed by the demands of the ecology. The complexity of behaviour is related to the complexity of the environment, as Rössler (1974) pointed out. Even a random walk is sufficient for survival and reproduction if food is abundant in a given environment. If food are distributed more sparsely, more

coherent movement of individuals is required (e.g. chemotaxis). In even more scarce environments, there is a growing need for cognitive abilities such as the use of landmarks or cognitive maps in order to survive.

The selection pressures leading towards the evolution of cognitive behaviour arise from ecological interactions. Such eco-evolutionary pressures are examined in the first two simulation experiments in chapter 5.

Major transition in evolution

The evolution of cognition is a major transition in evolution, which is related to multi-level selection theory. Major transitions, such as the origin of life and the evolution of multicellularity, deal with the emergence of higher levels of organization through evolution. In each transition, a number of smaller entities that were originally capable of surviving and reproducing on their own, became aggregated into a larger entity, thus generating a new level of biological organization. To accomplish this, the smaller entities must have had some selective advantage stemming from aggregation and cooperation. The challenge is to understand these transitions in Darwinian terms. Multi-level selection theory, which states that selection operates on more than one level simultaneously, is proposed as a way to integrate the emergence of new levels of organization.

Although the evolution of cognition is not modelled explicitly, we focus on the emergence of new level of organization that serves as a new level of selection. The last simulation experiments in chapter 5 presents in which a new level of selection emerges that causes the evolution and maintainance of an altruistic trait.

1.5 Outline

The remainder of the thesis is structured as follows. Chapter 2 provides an introduction to the conceptual background of this thesis and introduces several themes that are central to the simulation model and the experiments described in later chapters. In chapter 3, the simulation model is presented together with the results of ecological experiments of a two- and a three-species food chain. Chapter 4 discussed the theoretical ecological framework which is developed to understand and manipulate the population dynamics that emerges from the simulated ecosystems. In chapter 5, these ecological model are used in the indirect explanation of the observed evolutionary dynamics in terms of its relations to population dynamics and spatial self-structuring. Finally, conclusions are drawn in chapter 6.

Chapter 2

Background

This chapter provides a brief overview of the basic concepts and themes that are central to this thesis: embodied situated interaction, self-organization, and evolution by natural selection. The design of the individuals in the Virtual life model is based on the principles of embodied cognitive science, discussed in Section 2.1. The global spatial and temporal patterns that are observed from simulations of the model emerge from the interactions among these individuals through a process of self-organization, examined in Section 2.2. In Section 2.3, after discussing endogenous fitness model, several evolutionary mechanisms are examined which are modelled and simulated in the experimental part of this thesis.

2.1 Embodied situated interaction

The field of artificial intelligence and cognitive science have long been dominated by a computational approach in which cognition is modelled as abstract disembodied reasoning, and the body is reduced to a rather uninteresting input-output interface. In this paradigm, the term 'behavior' refers to cognitive processes such as decision making. Cognitive reasoning is understood in isolation, without providing a grounding in the bodily behavioral fundamentals on which it is based. From the 1980's, an alternative paradigm emerged that, in contrast, stresses the importance of embodiment and interaction.

Embodied cognitive science is the field that studies how complete agents cope with challenges in their environment (Pfeifer and Scheier, 1999). These agents are embodied, i.e. they possess a physical body with sensors and effectors, and are situated in their environment, i.e. their actions influence their perceptions mediated through the environment.

In embodied cognitive science, the term 'behavior' refers to the result of the dynamical coupling of an agent to its environment. The ultimate aim for embodied cognitive science is to understand how high-level cognitive phenomena arise from low-level interactions with the environment. This study is obscured



Figure 2.1: Frame of reference problem. Adopted from Pfeifer and Scheier (1999).

by several problems that together make up the frame-of-reference problem in the design of autonomous agents (Clancey, 1991). A distinction has to be made between the observer's perspective and the perspective of the agent itself. In particular, descriptions from our observer's point of view should not be taken as the internal mechanisms underlying the behavior that is described (see figure 2.1a). Since behavior always results from its dynamic interaction with its environment, complex behavior cannot be understood on the basis of internal mechanisms alone (see figure 2.1b). The environment must be part of the explanation. Therefore, the complexity of behavior cannot simply be attributed to the complexity of the underlying control mechanism. Indeed, agents with extremely simple control mechanisms often exhibit complex behavior that must be attributed to the complexity of the environment (see figure 2.1c).

2.1.1 Braitenberg vehicles

This is illustrated by the Braitenberg's "vehicles", a sequence of imaginary robots that differ in their internal wiring between their sensors and motors which demonstrate that often even extremely simple brains can show behaviors that look remarkably sophisticated to outside observers (Braitenberg, 1984). Figure 2.2 shows two types of vehicles with light-sensitive sensors and two motors/wheels. The vehicles are controlled by two wires that are connected either lateral (figure 2.2a) or contra-lateral (figure 2.2b).

Starting from a distance, vehicle 2a initially moves towards the light because the difference between sensor activity is very low. During the ever faster approach, due to the increasing sensor activity, the vehicle eventually turns away from the light with the difference between sensors. The rotational speed of the vehicle also grows during the turning, since the difference between the

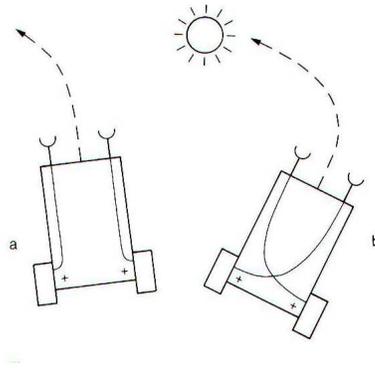


Figure 2.2: Braitenberg vehicles. Adopted From(Braitenberg, 1984).

sensors increases through turning. Vehicle 2b also approaches the light at an increasing speed. This vehicle does however not avoid the light, but instead it steers towards the source and hit it at top speed, possibly destroying the light bulb.

External observers, lacking knowledge of the internal mechanisms, might characterize these behavior as either timid or aggressive, and give explanations based on mental states such as beliefs, desires and intentions. Such observer-side ascriptions of behavior are characteristic of traditional cognitive science and artificial intelligence. Embodied cognitive science and artificial life, instead, attempt to understand behavior in terms of situated sensorimotor activity involved in agent-environment interaction.

2.2 Collective behavior

The environment with which animals interact typically includes other animals. Local interactions between agents, as a subset of agent-environment interactions, can result in the spontaneous formation of novel behavioral patterns on a global scale, as a result of a process of self-organization.

2.2.1 Self-organization

Self-organization is ubiquitous in nature. In many scientific domains, varying from physics and chemistry, to biology and economics, models have been developed that show that simple interaction between many system components at the local level may lead to (often unexpected) complex phenomena at the global level. This spontaneous emergence of global phenomena has led to intriguing new insights, since it shows that complex patterns need not be guided by central control, nor predesigned in the behavioural rules of its components.

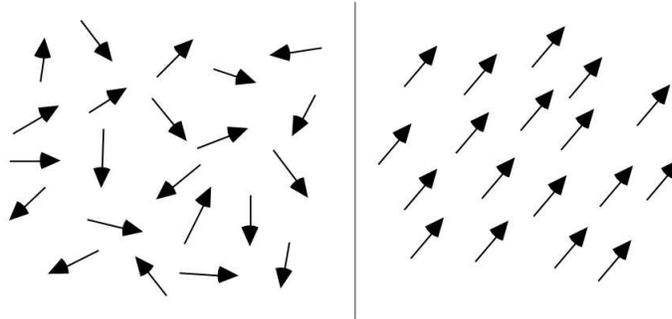


Figure 2.3: Self-organization of magnetic spins (Heylighen, 1999)

Instead, they emerge from the interactions among the components and their environment.

The process of self-organization can be illustrated by the simple case of magnetism. Consider a potentially magnetic piece of metal consisting of a multitude of tiny weak magnets called spins. All spins have a particular orientation that corresponds to the direction of its magnetic field. When the metal is heated, the spins point in all directions due to the kinetic energy (figure 2.3a). In this case, the magnetic fields of the spins cancel each other out, and the total metal piece will not be magnetic at all. Since the magnetic field of the metal looks the same from every perspective we look at it, it is symmetric.

When the temperature drops, however, the movement of the spins slowly decreases. Since the North poles of two individual spins repel each other, while the opposing poles attract each other, the spins will spontaneously align themselves. When two neighbouring spins align (due to a random movement) they will exert an increased influence on all their neighbours, causing them to align through local interaction. In this way, all spins end up pointing in the same direction (figure 2.3b). After this process, the magnetic field does not look the same from all perspectives anymore. Self-organization has broken the initial symmetry. Now, the fields of the spins do not cancel each other out, but add up. This results in a strong overall magnetic field of the metal piece.

A number of ingredients for self-organizing systems can be extracted from this example. First, self-organizing systems consists of many entities that interact with each other on a local scale. Such systems, being dynamical systems, always tend to evolve towards a state of equilibrium (end up in an attractor). Second, random events (noise) is necessary for self-organization to occur. Noise makes the dynamical system make undirectional movements through the state space (possible configurations of the system), and thereby end up in places where it is attracted to an equilibrium state, or do so more quickly. Third, self-organization requires regulation by positive and negative feedback. Positive feedback causes a self-reinforcement of a phenomenon (e.g. alignment of

spins) which ends when all components are absorbed by a new system configuration, leaving the system in a stable negative feedback state.

It has been argued that actually most dynamical systems can be said to be self-organizing in one way or the other, depending on the way and level of description. Self-organization is thus a way of *modelling* systems, rather than a *class* of systems. A system becomes describable as a self-organizing system when the level of description is being pitched on a lower level (Gershenson and Heylighen, 2003). If, for example, we describe the brain in relation to the bodily functions, it serves as a central control unit. When the level of description is moved to the neuronal level, however, the function of the brain is describable as a self-organised system.

The perspective of self-organizing system nevertheless remains essential in the field of artificial life modelling since explanations of emergence, evolution, and development of the life and behavior of living systems cannot be given by restricting models to a single level of organization or abstraction (Gershenson and Heylighen, 2003).

2.2.2 Biological systems

Self-organized processes are very common in biological systems. The mechanisms of self-organization in biological systems differ from physical systems in an important way: The interacting entities are usually much more complex. Whereas physical systems are composed of magnetic spins or grains of sand, the components of biological systems are neurons, ants or birds (Camazine et al., 2001).

As a result of this, the interaction between entities is typically far more elaborate than in physical systems. The interaction is not limited to physical laws as magnetism or gravity, but often show rich patterns of behavioral interactions. Moreover, insofar as these behavioral rules are genetically specified, an interesting dimension to collective behavior is added, because natural selection can adapt the rules of interaction, and thereby shaping the collective behaviors that can be formed (Camazine et al., 2001).

Many collective behaviors in biological systems can be understood as resulting from self-organization. Social insect behaviors are the primary source of inspiration for theories of collective behavior, but they are also applied to birds, primates and human beings (e.g. crowding in emergency situations) (Couzin and Krause, 2003). Not all collective behavior is due to self-organization, however. This is only the case where this group behavior is constituted by local interaction between individuals, independent of global or external control. Some counterexamples are regulation by queens in colonies of social insects or collective migration of birds in the direction of the sun.

Stigmergy

An interesting mechanism by which biological systems are capable of collective behavior is stigmergy, an indirect form of communication. By interacting

with the environment, an individual can change this environment. This can influence the behavior of other individuals to form collective phenomena.

Grasse (1959) found that the construction work done by a termite in a particular location changes the sensory input of the termite which in turn alters its behavior, and that of other individuals visiting this location. In this way, the building of termite hills is (spatially) coordinated by the environmental changes the work itself induces. Later, it was found that the emission of pheromones also played a significant role in the (temporal coordination of the) construction process. Both coordinational processes are completely distributed.

Melhuish and Holland (1999) distinguish between active and passive stigmergy. Active stigmergy changes the behavior of other individuals in a qualitative (do something else) or quantitative (do something more frequently) way. Passive stigmergy, on the other hand, does not influence the individual activity in any way (not qualitatively nor quantitatively), but *does* affect the outcome of a behavior. This source of self-organization in biological systems is probably the simplest and most basic form of coordinated action, since it requires the least behavioral complexity of its participants.

Didabots

Maris and te Boekhorst (1996) have used passive stigmergy, which they call a 'strategy of errors' (Deneubourg et al., 1983), with simple robotic models. They used robots, Didabots, with Braitenberg-like control systems¹. They have proximity sensors making them turn away from close obstacles and walls. Their environment consists of a rectangular area in which boxes are randomly distributed (figure 2.4a). When one of the robot's proximity sensors is activated because it detects an obstacle, the robot turns away from it. However, the morphology of the robots is chosen such that the sensors are in front but are directed outward. Therefore, a robot does not detect a box when it collides with it head-on (e.g. see left robot in figure 2.4a). This behavioral error causes a robot to push the box until it detects another box and turns away while leaving the box it was pushing. The result of this is that there are now two boxes where previously there was only one.

This change in the environment does not influence the behavior of other individuals in a qualitative or quantitative way. The robots keep doing the only thing they can. However, due to the physical changes that were made the chances increase that a third box is left at this location, since two boxes have more chance than one to be actually detected by the robots. The result of this auto-catalytic process is the formation of a cluster of boxes. For emergence to a single cluster to occur, it is necessary to use more than one robot to generate the crucial 'mistakes' in the obstacle avoidance behavior of the robots. Here and indeed in many collective behavioral phenomena the frame-of-reference problem arises again, since this behavior may be described from an observer's

¹A difference with Braitenberg's vehicles is that the Didabots also move in the absence of sensor activation.

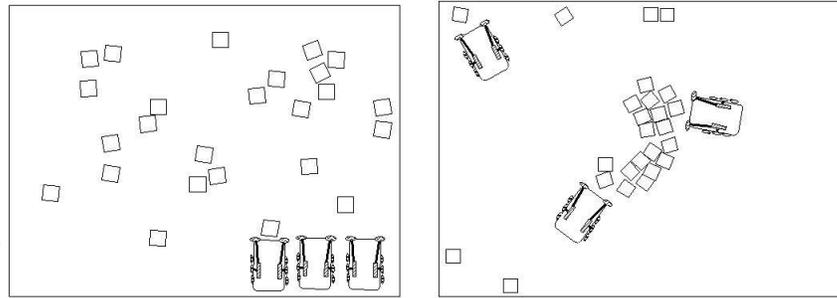


Figure 2.4: Clustering by passive stigmergy (Maris and te Boekhorst, 1996)

perspective as active ‘clustering’ or ‘cleaning’. The robots are however only avoiding obstacles ‘erroneously’.

2.3 Evolution

The self-organization of systems of biological organisms is perhaps most successfully illustrated by the concept of Darwinian evolution. The processes of replication, inheritance variation and interaction are sufficient to enable individuals to adapt to their environment by means of natural selection. Although the basic evolutionary principles are easy to understand, much is unclear about evolutionary mechanisms and selective forces by which natural species evolve. The Virtual life simulation model offers a way to investigate the influences of individual or ecological dynamics on the evolutionary dynamics.

2.3.1 Natural selection

As an evolutionary model, the Virtual life model is part of the class of *endogenous* fitness models (Forrest and Mitchell, 1994). This is a class of systems that is contrasted with traditional genetic algorithms (GAs) that employ *exogenous* fitness. The key difference between these classes is that in GAs the fitness of an individual is evaluated according to an explicit fitness function defined *a priori* by the experimenter. The selection mechanisms in GAs select individuals with relatively high fitness values for reproduction.

In endogenous fitness model, by contrast, reproduction is determined by the interactions between individuals and environment. Selective reproduction is not based on the relative fitness of the individuals in the population. Rather, endogenous fitness is based on the success of an individual in selective reproduction. It can be said that in exogenous fitness model ‘selection is based on fitness’, while in endogenous models ‘fitness is based on selection’. Or, traditional evolutionary algorithms use artificial selection, while endoge-

nous fitness models employ natural selection. The former model evolutionary *optimization* and the latter model evolutionary *adaptation* (Menczer and Belew, 1996).

Evolutionary adaptation by differential reproduction can only operate when there is variation among individuals. Without variation in the population (i.e. when all individuals are identical), natural selection does not change trait frequencies in populations, but it does result in changes in population sizes over time through reproduction and mortality processes. Endogenous fitness models without variation among individuals define non-evolving ecosystems (discussed in chapters 3 and 4). When variation among individuals is included in a trait that affects reproduction, we obtain eco-evolutionary systems (discussed in chapter 5).

In general, evolutionary change of a population can be attributed to either natural selection and genetic drift. The latter occurs when the evolvable trait does not affect differential reproduction. In this case, the genetic changes are basically a stochastic process that arises from the fact that mutations occur in traits that affect fitness. Genetic drift is observed most strongly in small evolving populations. Since the population in the simulation model presented in the next chapter are rather small, genetic drift must be ruled out in order to identify genuine selection. In the experiments presented in this thesis, this is done by choosing the evolving trait such that they directly affect fitness.

Selection cannot be based on explicit fitness criteria since these are absent in endogenous fitness models. Instead, it arises through competition among individuals for limited resources. This is true for trophic interactions (in food chains) between populations, but also between individuals of the same population. Thus, the fitness of an individual always strongly depends on other individuals.

Darwin was well aware of this interdependency: “the structure of every organic being is related, in the most essential yet often hidden manner, to that of all the other organic beings, with which it comes into competition for food and residence, or from which it has to escape, or on which it preys” (Darwin, 1859). In the struggle for energy and reproductive resource, all individuals are enemies of each other, even within a single species. This drives adaptation of the population as a whole through the evolutionary changes that arise in response to changes of other members of the population competing for the same limited resources.

2.3.2 Coevolution

The same evolutionary principle, both sides of interactions driving adaptation, holds for interactions between different populations or species, called coevolution. Coevolution is the mutual evolutionary influence between populations in which all populations exert selective pressures on others. The classic example is coevolution in predators and prey species which can lead to evolutionary arms races that can exhibit the red queen effect.



Figure 2.5: Red queen: *“In this place, it takes all the running you can do to stay in the same place.”*

Arms races

Foxes and rabbits are involved in a competitive race in two senses. The individual fox and its rabbit prey are competing on a behavioural time scale in the same sense as a submarine tries to sink a ship. But on a historical time scale, the designers of the submarine may learn from earlier mistakes, and as technology progresses, become better at sinking ships. Likewise, the fox population may evolve improvements for catching rabbits over evolutionary time. In response to this, the rabbit population may evolve adaptations to outwit the foxes. Biologists refer to such ongoing evolutionary mutual counter-adaptations as an evolutionary ‘arms race’ (Dawkins and Krebs, 1979).

Red queen effect

The adaptations obtained in an arms race do, however, not necessarily lead to an improved fitness in the sense of being better adapted to the environment, since this environment evolves as well. van Valen (1973), who studied the extinction rates within and between taxa, noticed that species with longer evolutionary histories need not be better adapted to their environment than ‘young’ species. He proposed the red queen hypothesis to point out that in evolution, populations must continuously adapt to maintain the same level of fitness (see figure 2.5). In this situation, both species adapt to the other, without either one becoming more efficient.

2.3.3 Multi-level selection

Darwinian evolutionary theory seems to predict that individuals will always act to increase their own fitness. For this reason, the evolution of altruism (acting to increase the fitness of another individual at the expense of its own) has long been a benchmark problem in evolutionary biology. Not only is it an important issue for understanding altruistic behavior of biological systems alive today, but also to explain the emergence of new levels of organization, since these major transitions in evolution often require the (partially) release of self-interest for the sake of the group (Maynard-Smith and Szathmary, 1995).

A first proposal to resolve this problem was the idea that traits can spread through a population because of the benefits they have for groups of individuals, regardless of the fitness of individuals within that group (Wynne-Edwards, 1962, 1963), was heavily criticized on the theoretical grounds that large population of altruist individuals would be very susceptible to invasion by selfish individual and therefore not constitute a stable evolutionary mechanism (Williams, 1966). The latter author developed a concept of the gene as the fundamental unit of selection, which has become a grand theme in biology, especially following the popularization of the idea of the 'selfish gene' by Dawkins (1976). The unit of selection was moved down (to a level lower than the individual) instead of up (to a level higher than the individual).

Group selection was even more pressured by two influential theories in the 60's and 70's. Inclusive fitness theory (a.k.a. kin selection) developed by Hamilton (1964) explained how altruism can evolve among genetic relatives, since it pays off in terms of gene propagation to behave altruistically towards kin since relatives carry copies of the genes. Even if one sacrifices himself for close relatives, the genetic material is preserved for reproduction in another individual. This has greatly contributed to the understanding of altruism in species using haplodiploidy (e.g. eusocial communities such as ant colonies). Game-theoretical accounts of the evolution of altruism showed that such behavior can also evolve among non-relatives if they participate in reciprocal altruism in Tit-for-Tat situations (Axelrod and Hamilton, 1981).

More recently, however, the concept of group or multi-level selection was reintroduced by D.S. Wilson et al. (Wilson, 1975; Wilson and Sober, 1994). They state that groups of individuals can have functional organisations in the same way as individuals do and can thereby act as a 'vehicle' for selection. This higher level of selection influences the course of evolution since selection *between* groups can be in another direction than selection *within* groups.

Spatially explicit individual-based models are used to emphasise the influence of *spatial* instead of the *genetic* relatedness between members of groups (as spiral waves, patches or clusters)². It is argued that Hamilton's relatedness is then not merely the genetical correlation, but depends of local interactions and population dynamics of viscous populations³ (van Baalen and Rand, 1998). In spatial models of evolution self-organised spatial structures can emerge that form a new level of selection (Johnson and Boerlijst, 2002). Boerlijst and Hogeweg (1991) provide a striking example of the differences of individual- and group-level selection. In a model of prebiotic evolution, they showed that competition between spiral waves of molecular species favours high mortality rates of individuals, which is clearly not beneficial to the individual. Spiral waves with short-living individuals have a competitive advantage over other spirals because they rotate faster and can annihilate waves consisting of longer living individuals.

²Hamilton has also noted that his inclusive fitness theory is more general than kin selection.

³Viscous populations are populations without imposed subdivision but with limited dispersal and migration. Because offspring tend to remain close to their relatives, individuals are likely to have relatives in their neighbourhood.

The study of the influence of spatial self-organization on the evolutionary dynamics of evolving populations is an exciting task that has only recently become feasible by spatially explicit individual-based simulation models, such as the virtual life model. However, it is often easier to construct and run such a simulation model, than it is to interpret its results. Especially because this involves studying, not merely the emergence of a single, but the interplay between various emergent patterns: spatial self-structuring, population dynamics, evolutionary dynamics. It is therefore important to relate emergent patterns to theoretical understanding where possible, both conceptual and mathematical.

This chapter provided a brief overview of some several theoretical concepts underlying the construction of the simulation model. The next chapter introduces the model and shows its ability to simulate emerging population dynamics and spatial clusters. In chapter 4, the emergent population dynamics are analysed and understood in terms of classic mathematical models from theoretical ecology. The evolutionary experiments in chapter 5 first show the interplay between evolution and population dynamics, and after allowing spatial self-structuring, the interplay between all these processes.

2.4 Conclusion

In this chapter, we have introduced several concepts and themes that are central to modern approaches in embodied cognitive science and theoretical biology in general, and to this thesis in particular.

Interaction between simple situated and embodied agents with their environment gives rise to complex behavioural patterns. In accordance with this view, complex (and cognitive) behaviour of natural agents is reinterpreted by embodied cognitive science as an emergent property of such agent-environment interactions. When the environment of an individual consists of many other individuals, the same process can result in self-organised collective behaviours. In contrast to the concept of evolution as an optimization process, new modelling approaches view evolution as an adaptation process in which the selection criteria are endogenous (produced from within), because the selective pressures are not predefined, but emerge from the interactions between individuals. The simulation models described in the rest of this thesis combine these processes.

Chapter 3

Simulation model

In this chapter, a simulation model is presented that combines the processes of situated interaction, spatial self-organization and endogenous fitness that were discussed in the previous chapter. This model aims to facilitate the study of self-organising ecological and evolutionary dynamics that emerge through interactions between situated agents.

The virtual life model is an individual-based, spatially explicit model in which individuals are specified as either resources or consumers. Trophic interactions result in an energy flow through the food chain from resources to consumers (and topconsumers). Whereas resources are immobile individuals, consumers are modelled as simple situated agents that perform taxic behaviour (towards their resource) by Braitenberg-like sensorimotor control. Reproduction and mortality depends on the age (in resources) or energy level (in consumers) of individuals.

The model enables the controlled simulation of ecological processes. Although population dynamics are not explicitly specified in the model, temporal patterns in relative population sizes emerge through the local trophic interactions between individuals. Spatial self-structuring can be introduced in the model by specifying the growth of the resource population locally. Non-homogenous spatial structure (emergent clusters/patches) can influence both ecological and evolutionary dynamics.

This controlled ecosystem is already an endogenous fitness model, since reproduction is not based on an explicit fitness function, but on individual interaction. Evolution is therefore easily incorporated in the model by introducing some inheritable variation between individuals of a population. Differential reproduction of individuals with different traits causes populations to (co)evolve. In this case, the simulations are eco-evolutionary models in which the timescales between these processes are not separated (van der Laan and Hogeweg, 1995). It therefore enables the study of the interplay between (spatial and temporal) ecological and evolutionary dynamics based on the conjecture that evolution does not occur as a clean universal process, but is always embedded in a web of ecological and historical dynamics (di Paolo, 1999).

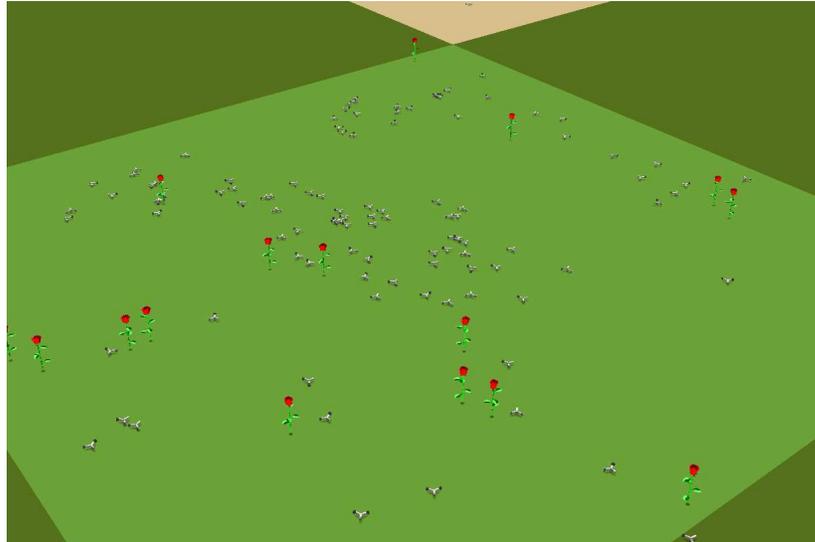


Figure 3.1: Screenshot of plants-herbivore simulation

This chapter overviews the most important aspects of the simulation model. The following section discusses the simulator in which the virtual life model is implemented. Section 3.1 examines the main part of the model by describing the specifications of individuals and populations, the trophic interactions and evolutionary process, and providing an overview of the most important simulation parameters. The results of ecological experiments with emergent population dynamics and spatial self-organization of resources are presented in section 3.2. In section 3.3, conclusions are drawn from this chapter.

3.1 Virtual life model

Framsticks

The Virtual life simulation model is implemented in the a-life simulator which allows customization for nearly all processes involved in the simulation by using the scripting language FramScript (Komosinski and Ulatowski, 1999; Komosinski and Rotaru-Varga, 2000; Komosinski, 2003; Adamatzky and Komosinski, 2005). The models presented here are simulations that are fully customized for ecological and evolutionary experiments. Among the features that are adjusted to the Virtual life model are the physical simulation of the environment, the sensors, control, and actuators of individuals, and most importantly the experiment definition.

Algorithm 1 Basic structure of experiment definition

Initialization:

- generate populations
- fill population with individuals

Step:

- process sensors, control and update locations
- if resource > reproductive age, reproduce (and mutate)
- if consumer > reproductive energy, reproduce (and mutate)
- subtract metabolic energy of consumers

Interaction (collision):

- transfer energy from resource to consumer
- if energy ≤ 0 , kill individual

Interval:

- write population and genotype data to logfiles
-

Event-based

The experiment definition is an event-based program that tells the simulation what to do when certain events occur. This includes user events (e.g. loading and initialization), interaction events (e.g. collisions between individuals), and population events (e.g. spawning and removing individuals). This definition is customized to implement an endogenous fitness model. The core of the Virtual life model is implemented as a set of experiment definitions that share the following basic structure.

At initialization, populations of individuals are generated and certain amounts of individuals for each population are spawned in some locations in the environment. In evolutionary experiments, initial variation among individuals is induced by stochastic mutations in a certain trait. The physical simulation, as well as the processing of sensors, control and actuation (i.e. updating locations) are left to the Framsticks simulator itself. However, most other aspect of individuals are updated every simulation step in this definition, such as age and energy level. At a certain interval, data about observables are written to logfiles, such as population sizes and values of evolvable traits.

When collisions between individuals occur, the experiment definition determines whether these are ignored or handled. Collisions between con-specifics (individuals belonging to the same population) are ignored in the simulations reported in this thesis (but can include e.g. sexual reproduction or dominance interactions). Inter-specific collisions are handled as trophic interactions in terms of consumption and predation when certain conditions are met.

3.1.1 Environment

The environment is a continuous space in the form of a square (see fig. 3.1). Although the physical simulation in Framsticks simulates many physical forces on the mechanical bodies of creatures, these are largely ignored in these simulations. In the scope of this thesis, forces such as friction, inertia and gravity are

simulated but irrelevant for present purposes. Mechanical collisions between individuals are ignored¹.

Boundary conditions

The boundaries of the environment can be either *fixed*, *wrapped* around, or *absent*. Different boundary condition can also be used for different populations². Usually, resources are distributed only within the boundaries (fixed boundaries). Consumers can, in principle, travel away from the environment (i.e. absent boundaries). However, they can only sustain themselves for a limited time and distance, since no resources are found outside the environment. Individuals will die beyond a certain distance. Moreover, consumer individuals are controlled such that they turn in the direction of the highest sensory activation gradient, such that consumers are likely to stay relatively close to resources.

3.1.2 Populations

The environment is inhabited by individuals that belongs to a certain population. A population, or species, is a group of individuals that are equal in the following respects. They consume (and are consumed by) the same species, they have the same (im)mobility, the same (metabolic) energy costs, reproductive thresholds and evolvable traits. Populations are predefined by the model in the sense that the interactions between individuals are defined for the populations to which the individuals belong. These interactions determine the food chain in which the populations are structured.

The implementation of the Virtual life model allows for many populations to interact in intricate food webs, but the simulations in this thesis are restricted to simple linear food chains of three species: a resource, consumer and top-consumer species.

Energy and trophic interactions

The flow of energy is a common theme in all ecosystems and it is a means of understanding how all ecosystems having many properties in common irrespective of their apparent differences. Therefore, it is often used as the key feature of individual-based ecological models and endogenous fitness evolutionary models.

Likewise, in the virtual life model, individuals are specified by an energy level. The ecosystem is structured as a food chain, such that individuals of different populations transfer their energy when they are consumed (see figure 3.2). Energy is passed up the food chain every time an individual consumes

¹The physical simulation may be important in future research that includes embodiment more explicitly, e.g. in studies in the coevolution of morphology and control.

²This can be convenient in simulations with consumers that require absent boundaries in combination with self-clustering resources that require fixed boundaries.

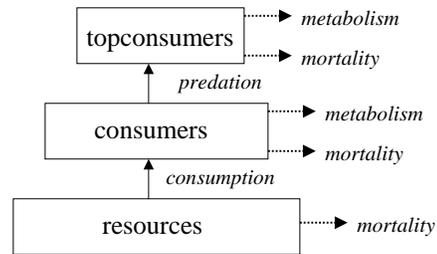


Figure 3.2: Energy flow and dissipation through food chain

either a resource or another consumer. Each successive transfer of energy involves dissipation of energy in the form of loss of energy by metabolism and natural mortality (but not by starvation). The modelled ecosystem can therefore be understood as a dissipative system, obeying the second law of thermodynamics.

The ecosystem model involves a structure of autotrophs and heterotrophs. Resources are modelled as autotrophic (i.e. self-nourishing) individuals, or primary producers, that are able to fix energy themselves (e.g. photosynthesis). Consumers are heterotrophs that cannot produce their own energy. They need to make use of energy stored in individuals of lower trophic levels. This class of individuals can be structured into primary consumers that feed on the autotrophs, and secondary consumers that feed on the primary consumers, etc.

3.1.3 Individuals

There are two main types of individuals: resources and (top)consumers. Resources are stationary individuals, and consumers are simple situated agents with Braitenberg-like sensorimotor control.

Resources

Resources are stationary individuals, such as plants or trees, that are specified by an internal level of *resource energy*. This energy level remains constant, except when the resource is consumed by a consumer individual. In this case, the energy of the resource is transferred to the consumer, and the resource dies and is removed. Apart from consumption, resources have a small probability, a natural *mortality*, of being removed randomly. A resource can reproduce once every *reproductive period*. The probability of reproducing successfully depends on the size of the resource population. The number of resources is limited to a maximum: as the population size grows to the maximum, the probability of successful reproduction of a resource decrease proportionally. The maximum number of resources that a certain environment can sustain is called the *carrying capacity*.



Figure 3.3: Herbivore body and control system

The carrying capacity of resource population can be specified *globally*, over the whole environment, or *locally*, over the local neighbourhood of the individual that attempts to reproduce. Offspring can also be placed *globally*, using a uniform random distribution over the environment, or *locally*, using a Gaussian random distribution with the parent location as its mean.

With these options the spatiality of the ecological simulation can be altered. When both the carrying capacity and the placement of offspring are defined globally, the resource population is homogeneously distributed and presents a well-mixed non-structured environment (to consumers). Under these settings, the spatiality of the simulation model is explicitly made homogeneous to match the implicit assumptions in the use of mean-field approximations in classic differential equation models. Such classic ecological models can then be used to model the population dynamics that emerge in such ecosystems. The ecological analysis of population dynamics in chapter 4 is largely restricted to the analysis of such 'homogenous space' models.

When the placement of offspring and carrying capacity are defined locally, however, the resource distribution is no longer homogeneous. The resource population organizes itself in clusters or patches that can move (over generations of resources), divide and compete with other patches. In the next section, the ecological and evolutionary consequences of this self-structuring is examined in more detail.

Consumers

Consumers are modelled as mobile situated agents with simple sensors, a control system and actuation devices (see fig. 3.3). Consumers move forward at a constant velocity, and (simultaneously) turn at a constant rotational velocity in the direction of the sensor with the highest activation. Individuals have a body in the form of a three-stick structure. Two of the sticks are equipped with

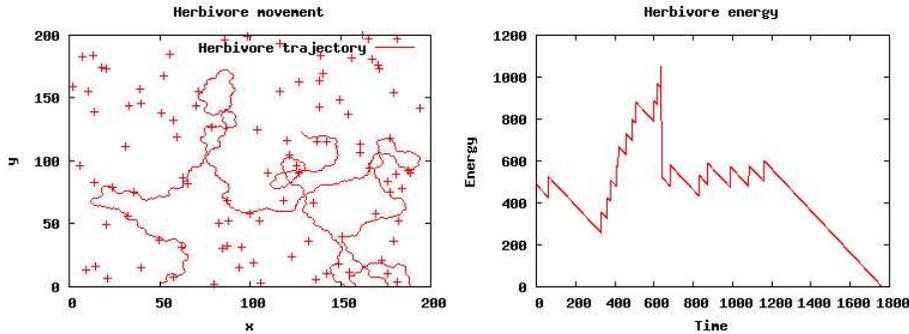


Figure 3.4: Consumer movement and energy

sensors whose activation is proportional to the energy it detects from individuals of the population it feeds on. The sensors of primary consumers detect the energy of resources, the sensors of secondary consumers detect energy of primary consumers, etc. The activation of a sensor is calculated by adding the energy divided by the distance from the sensor to the source plus noise:

$$sensor\ activation = noise + \sum_{i \in P} \frac{resource\ energy_i}{distance_i} \quad (3.1)$$

where P is the size of the population the consumer feeds on. Noise is added to the sensors, which makes the sensors unreliable when sources are distant (since the noise is equal or greater than the smell of sources). The sensory noise ensures that individuals can only sense their local environment. Since the sensors of a single individual are in slightly different locations and thus detect different energy levels, the difference between the sensors can be used to steer in the direction with the highest gradient. The control architecture of consumer individuals is similar to the phototactic Braitenberg vehicle in figure 2.2 on page 14. Consumers turn in the direction of the sensor with the highest activity³, and thus (if not disturbed by noise) move towards some individuals of the population they feed on. A typical trajectory of a consumer over its lifetime is depicted in the left panel in figure 3.4. After its birth close to location (50,0) it travels around and consumes several plants (crosses). This individual leaves the environment twice, but returns to forage, and eventually dies at location (125,125).

Reproduction and mortality Birth and death processes are determined by the consumer's energy level. At birth, consumers receive a certain amount of energy, and this level is decreased every step as a *metabolism cost*. When the

³They differ from Braitenberg vehicles in the sense that the velocities are not proportional to the sensor activity, among other things.

energy is below zero, the individual dies and is removed from the environment. Energy levels can be increased by consumption during which energy is transferred from the resource to consumer. When the energy level exceeds the *reproduction energy*, the individual reproduces. Offspring receives half of the energy of its parent.

The right panel in figure 3.4 shows the energy level over the lifetime of a consumer. It decreases linearly every step and increasing instantly on consumption. When the energy level rises above the reproductive threshold (of 1000), it reproduces and half of its energy is inherited by (or invested in) the offspring. After a period of sustaining a constant energy level, the consumer dies of starvation.

Placement of offspring The spatial placement of offspring can be done in several ways. They can be placed (1) in a random location in the environment, (2) close to its parent, and (3) close to a random parent. The first can be used to approach the homogeneity that is assumed by classic differential equation models, and thus can be understood better in the terms of such models. The second introduces spatial heterogeneity that potentially influences ecological and evolutionary dynamics, and the third placement option is used to show this influence by distorting the correlation between spatial and genetic relatedness.

3.1.4 Evolution

Evolution is easily incorporated in the model by having some inheritable variation among individuals in the populations. Within the scope of this thesis, the evolution dynamics are extremely simple. Individuals are specified by only a single trait, restricting the evolution of the population to a single dimension. Moreover, simple one-to-one genotype-to-phenotype mapping are used, and learning or development processes are excluded.

When an individual reproduces its offspring inherits its genotype (=phenotype), while mutations occur with a certain *mutation probability*. Mutations on real-valued genotypes are usually modelled as Gaussian diffusion, and by mutating natural-valued genotypes is done by altering the parent genotype by adding or subtracting one. In coevolutionary experiments, the mutation probabilities are equal for the populations involved. Phenotypic traits are chosen such that there is a clear relationship to parameters in the theoretical model and the evolvable trait.

Evolution is not included from the ecological experiments presented in this and the next chapter, in which all individuals are identical. Evolutionary dynamics that emerge from the simulation model are explored in more detail in chapter 5.

Class	Simulation parameter	Description	Default (unit)
Environment	<i>size</i>	size of the environment	200-400 (length)
	<i>boundary</i>	boundary conditions can be fixed, wrap-around, or none	fixed, none
Resource	<i>energy level</i>	fixed energy level, transferred on consumption	250-500 (energy)
	<i>reproductive period</i>	resources can reproduce after this time (amount of steps)	15-25 (timesteps)
	<i>carrying capacity</i>	probability of successful reproduction is $prob(reproduction) = 1 - (R/K)$; can be defined globally or locally	200-300 (# of ind.), global
	<i>carrying area</i>	if carrying capacity is local, this determines the area over which it is defined	25-100 (length)
	<i>placement</i>	offspring is placed either randomly or close to parent	random
	<i>natural mortality</i>	resources have a small probability of mortality	0.0-0.01 (per step)
	<i>metabolistic cost</i>	consumers lose this amount of energy every step	5.0-10.0 (energy)
Consumers	<i>reproductive energy</i>	consumers reproduce above this energy level	1000 (energy)
	<i>placement</i>	offspring is placed randomly, close to parent or close to random parent	close to parent
	<i>natural mortality</i>	(top)consumers have small probability of mortality	0.0-0.001 (per step)
	<i>metabolistic cost multiplier</i>	topconsumers lose this amount of energy every step: metabolistic cost * cost multiplier	2.0-10.0 (*energy)
	<i>handling time</i>	amount of steps a topconsumer cannot consume after a consumption	20-30 (timesteps)
Evolution	<i>mutational operator</i>	real-valued: Gaussian random distribution (specify st.dev.), natural-valued: add or subtract 1	
	<i>mutation probability</i>	probability of mutation per offspring	0.001-0.1 (per offspring)

Table 3.1: Simulation parameters

3.1.5 Simulation parameters

In the above description of the model, various simulation parameters have been discussed that influences spatial, energetic, reproductive and mutational processes in the simulated ecosystem. An overview of the most important simulation parameters is given in table 3.1.

These parameters enable us to influence simulated (evolving) ecosystems in many ways. The results of various experiments studying the influence of parameter settings are reported in the next section. In the next chapter, these simulation parameters are related to parameters in theoretical ecological models. Correspondences between these two enable us to manipulate the ecological and evolutionary dynamics.

3.2 Emergent patterns

The simulated ecosystem, consisting of many interacting individuals, can be understood as a self-organizing dissipative system. The energy flow through the food chain results in temporal patterns in population sizes, or in spatial pattern formation of resources. In this section, several emergent patterns resulting from ecological simulations are reported.

Two sets of simulation experiments are conducted. First, simulations with the virtual life model are conducted with a resource-consumer and a resource-consumer-topconsumer system. These experiments focus on the population dynamics that emerge through the trophic interactions in a homogeneous environment. Second, the influence of spatial self-structuring on ecological and evolutionary scales is studied by a simulation of only a resource population.

3.2.1 Population dynamics

To study population dynamics in the virtual life model, a simple ecosystem with a resource and consumer population is simulated. Trophic interactions between resources and consumers result in changes in population sizes over time. The emergent population dynamics show nonlinear behaviors such as fixed points attractors, limit cycle oscillations and strange attractors. The results of simulation of a resource-consumer and a resource-consumer-topconsumer ecosystem are presented below. In these experiments, the spatial distribution of resources was defined globally to obtain spatial homogeneity.

Resource-Consumer ecosystem The ecosystem model is initialised with a resource and a consumer population. The simulation parameters are shown in figure 3.2.

Figure 3.5 shows the resulting population dynamics, transients and phase plots, for three different values of the *carrying capacity*. The panels in the top row show population dynamics that approach a stable fixed point equilibrium, with random fluctuation. The population dynamics in the second row exhibit

	Simulation parameter	Value - bitrophic	Value - tritrophic
Environment	<i>size</i>	250	250
	<i>boundary</i>	R=fixed, C=none	R=fixed, C=T=none
Resource	<i>energy level</i>	250	250
	<i>reproductive period</i>	20	20
	<i>carrying capacity</i>	varied (100,200,300)	300
Consumers	<i>metabolistic cost</i>	10.0	10.0
	<i>reproductive energy</i>	1000.0	1000.0
Topconsumers	<i>metabolistic cost multiplier</i>		2.0
	<i>handling time</i>		20.0

Table 3.2: Simulation parameters setting for bitrophic and tritrophic systems

ongoing oscillatory behaviour, known as a limit cycle. Finally, the bottom plots show the oscillatory behavior with large amplitude of an unstable system, which eventually leads to the extinction of the consumer population.

Resource-Consumer-Topconsumer ecosystem When a topconsumer population is added to the ecosystem to obtain a tritrophic food chain, the behaviour looks more interesting. The transient of the population sizes in the top panel of figure 3.6 shows chaotic (aperiodic) behavior, until in the end the consumer population, and subsequently the topconsumer populations go extinct. Although the behaviour can be called chaotic, it still uses only a small portion of the phase space (possible configurations). In the 3D phase plot, the population dynamics can be seen to be governed by a strange attractor that involves oscillations in two distinct directions. By plotting two 2D phase space of the same data, one can see that there is oscillatory behaviour in both the resource-consumer plane (bottom right), and the consumer-topconsumer plane (bottom left).

Conclusions

The population dynamics of these ecosystems are not defined in the simulation model. Reproduction and mortality events are specified for individuals, and not for populations. The regularities observed in the population sizes over time emerge as a result of interactions between individuals of various populations. The temporal patterns in population dynamics come about due to self-organization.

It is an example of the artificial life paradigm that local interactions between simple entities can result in complex global patterns. In this case, the simple entities are the immobile resource individuals and the situated (top)consumer agents. The local interactions are twofold: between individuals of various population, and between individuals of the same population. The former is specified by the Braitenberg-like control mechanism of consumers which causes

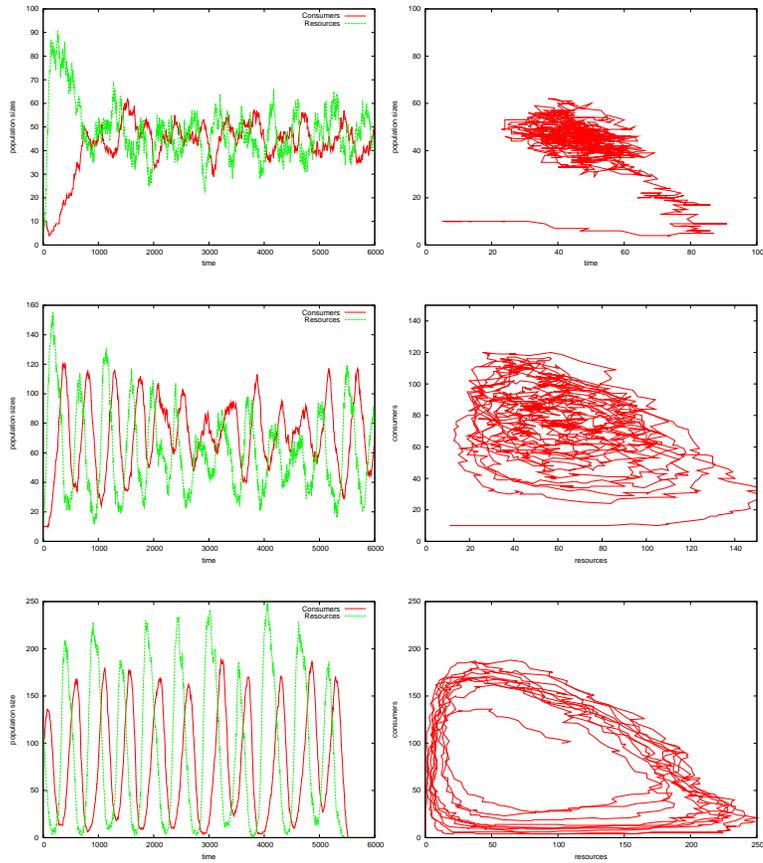


Figure 3.5: Population dynamics of the resource consumer, for various values of *carrying capacity*. (Left) Population sizes over time. (Right) Phase plots. Note the increasing scale of the plots. (Top) Fixed point attractor ($K=100$), (b) oscillations, limit cycles ($K=200$), (c) oscillations with large amplitude leading to extinction of consumer population ($K=300$).

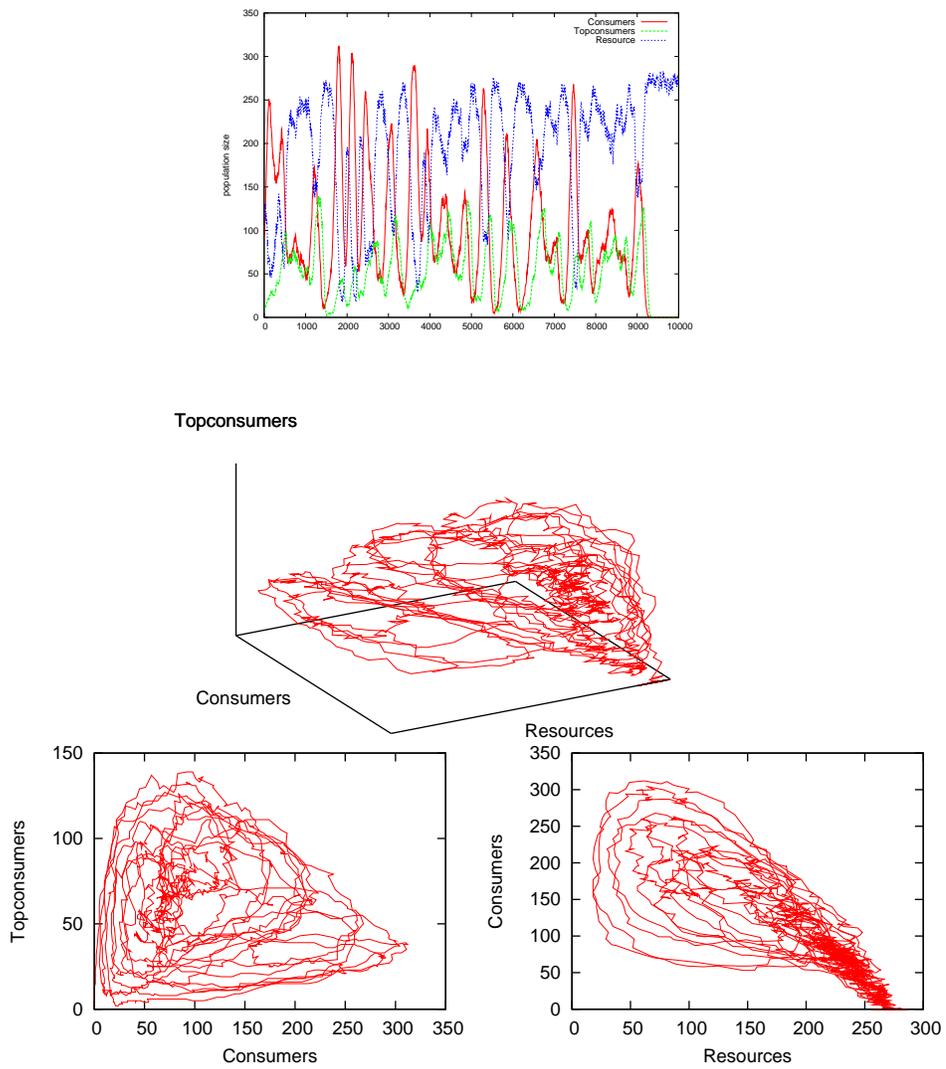


Figure 3.6: Population dynamics of tritrophic foodchain

	Simulation parameter	Ecological time	Evolutionary time
Environment	<i>size</i>	200	200
	<i>boundary</i>	wrap-around	fixed
Resource	<i>reproductive period</i>	15	20
	<i>carrying capacity</i>	25	25
	<i>carrying area</i>	50	50
	<i>placement</i>	local, $\sigma = 5.0$	local, $\sigma = 5.0$
	<i>natural mortality</i>	0.005	0.005
	<i>mutation probability</i>	-	0.1
	<i>mutational operator</i>	-	± 1

Table 3.3: Simulation parameters setting for spatial self-structuring

chemotactic behaviour towards resources (see section 2.1). When individuals interact directly (i.e. collide), energy is transferred which influences reproduction and mortality events.

No direct interaction rules between individuals of the same population are specified by the simulation model. Individuals are not even able to detect the presence of conspecifics. However, consumers do interact with each other indirectly, through their local environments. When a consumer consumes a resource, the behavior of other consumers is affected, since they do not move towards this resource anymore. Since the behaviour of individuals is based on their local environment, the activity of consumers can alter the outcome of the behavior of other consumers. In other words, consumers are engaged in indirect, stigmergic interactions (see section 2.2). Combined with the process of natural selection (see section 2.3), this results in the observed population dynamics.

3.2.2 Spatial clustering

To study spatial self-organization, an ecosystem with only a resource population is simulated. The carrying capacity is defined locally and offspring is placed close to parents, and the emergent spatial distribution of the resource population is heterogeneous. Although resources are stationary individuals, the resource population structures itself spatially by mortality and reproduction. The reproduction rates differ between individuals, according to the density of others in the local environment of individuals. In combination with offspring being placed close to parents, this can cause patches to emerge that remain stable over sustained periods. The self-structuring of resources is first studied on a short term ecological timescale, and later at a long term evolutionary timescale.

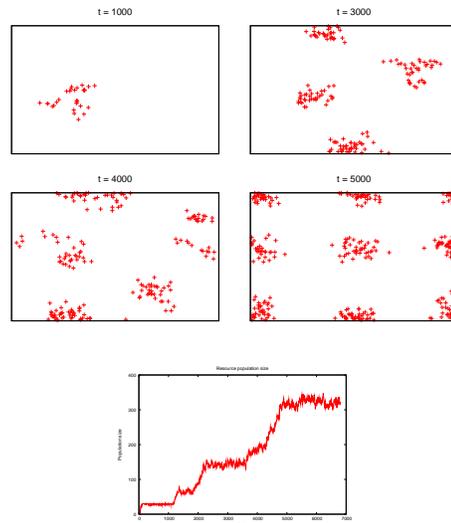


Figure 3.7: Spatial clustering of resources by local carrying capacity

Ecological timescale

Figure 3.7 shows spatial clustering of a resource population that reproduces locally in an environment with wrap-around boundary and a locally defined carrying capacity. The simulator parameter settings are shown in table 3.3. The simulation is initialised with a single individual in the middle of the environment. A patch of resources surrounding this individual soon grows to its local carrying capacity. Over time, the patch slowly moves and changes shape by natural mortality and reproduction. A patch can split into multiple patches. When a new patch arises, the total population size grows accordingly. The size of the resource population over time is given in the bottom panel in figure 3.7.

The top left plate in figure 3.7 shows a single patch that is almost dividing. The top right plate show the population in three patches. The clusters move through the environment and are dividing in bottom left panel. Finally, in the bottom right plate, the population has settled in a stable configuration of four patches arranged in a regular grid.

The corresponding population dynamics shows that the spatial structure of the resources strongly influences the population size. When a cluster splits up or divides, the new clusters soon establishes themselves near its local carrying capacity. As more clusters arise, clusters move through the environment (over generations) to maximize their local carrying capacity. Through an implicit competition for space between clusters, the clusters organise themselves as to optimize the global population size.

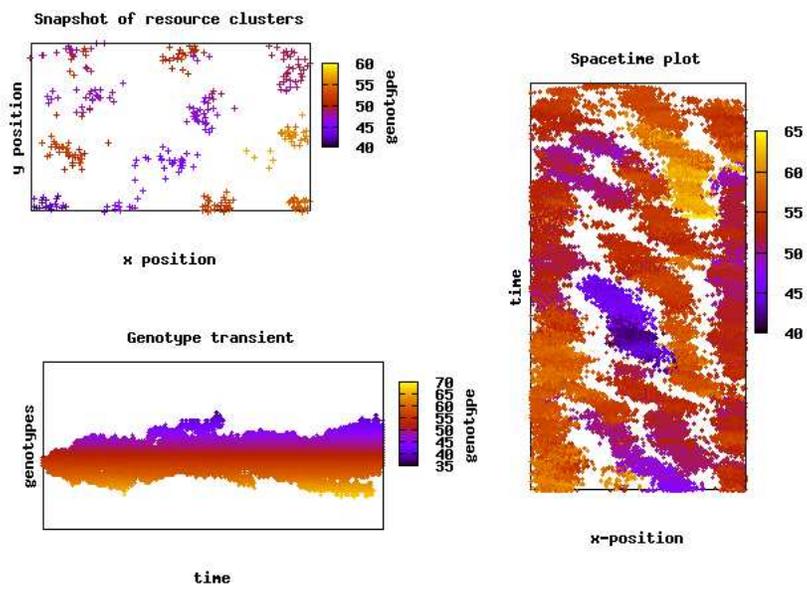


Figure 3.8: Clustering in evolutionary time

Evolutionary timescale

In another experiment, the influence of self-structured clusters is simulated over a longer evolutionary timescale. On this timescale, the population size is (after the initialization period) always near its global carrying capacity (of approx. 325 resources), and is therefore not plotted. In contrast to the previous experiment, the environment has fixed boundary condition which means offspring cannot be placed outside the environment (see table 3.3).

The results are shown in figure 3.8. First, the colours are ignored. The top left panel is a snapshot of the environment which shows the spatial distribution of resources into clusters. Note that, due to the fixed boundaries, the population does not organize itself in a rigid regular grid as in the previous experiment.

The panel on the right is a so-called spacetime plot which shows the clustered individuals in a certain section of the environment over time. At every 50 time steps, the x-position of the individuals in the y-region between 100 and 150 are plotted in the space-time plot. It thus gives an idea of the development of the clusters over time.

The space-time plot shows that clusters are long-term structures. Clusters have a lifetime that is far greater than the lifespan of individuals. Clusters can last for hundreds or thousands of generations. The plot also show that there is an ongoing movement of clusters. Its spatial configuration does not, in contrast to the previous experiment, reach a stable equilibrium. Clusters (at least in the recorded section) slowly move to the right. The empty space is filled because other spatial clusters divide. There is an ongoing (implicit) competition for space among clusters.

It can be argued that clusters have a life cycle of their own. Clusters can get born, mature (up to carrying capacity), they have interaction and competition with other clusters, they can divide or reproduce, and can die (when all its members go extinct). In this sense, the emergent spatial clusters form a new level of organization. In order to see whether this new level can serve as a new level of selection, we need to consider evolution.

To study the evolutionary consequences of this spatial self-organization, a mutation probability is introduced. Resources are specified with a natural-valued genotype that is mutated by adding or subtracting one. Genotype values are depicted in all panels as different colours. Individuals are completely identical, irrespective of their genotypes. There is no genuine evolution, since there is no selection. All evolutionary change is attributed to genetic drift. Changes in genotypes over time is plotted in the bottom left panel in figure 3.8.

In the genotype transient (the bottom left panel), one can see the emergence of clusters, not in real space, but in genotype space. The emergence of such 'quasi-species', small ensembles with similar genotypes, introduces a new meso-scale level into the model, between micro-scale (individuals of single genotype) and macro-scale (all possible genotypes). The pattern formation

in genotype space can lead to an interlocking of ecological and evolutionary timescales (van der Laan and Hogeweg, 1995).

The genotypic clusters can be seen in the space-time plot as well. By looking at the colours of the various clusters (in real space), the genetic similarity within the spatial clusters is evidently higher than between clusters. An example of a correlated spatial/genotypic cluster can be seen in the middle of the space-time plot (blue/purple/black cluster), which corresponds to the cluster in the middle of the genotype transient. The snapshot also shows a strong correlation between spatial relatedness and genetic relatedness, since the colours within a cluster are more similar than the colours between clusters.

This experiment does not include evolution in the sense of a mutation-selection mechanism, since there is no selection (acting upon geno- or phenotypes). It does show, however, that the formation of a new level or organization (i.e. the spatial/genotypic clusters) introduces a new level of historical dependence to the dynamics of the system over evolutionary timescale. This historical dependence causes a diversity of genotypes in spatial clusters, which introduces a new level on which selection can act.

Conclusions

The spatial structuring is not specified by the simulation model. Therefore, the simulation model only specifies the conditions that allow spatial self-organization to occur. The clusters of resource that appear over time are not predefined, but emerge through interactions between resources that stem from the local definitions of carrying capacity and placement of offspring.

3.3 Conclusions

This chapter described the key features of the virtual life simulation model: the simulated environment, the resource and consumer individuals that inhabit this environment and how they are structured in a trophic food chain. Nonlinear population dynamics emerge from the (trophic and behavioural) interactions between stationary resources and situated (top)consumers, due to a combination of stigmergic and energetic processes. Spatial self-structuring of the resource population occurs under local conditions for placement of offspring and carrying capacity. The emergent spatial clusters or patches introduce a new level of organisation (between individual and population) that can influence ecological dynamics and evolutionary dynamics.

The population dynamics are analysed in terms of differential equation models by mean-field approximation in the next chapter. This enables us to gain theoretical understanding of the dynamics that emerge from the simulation models. The correspondence and relations between the simulation parameters and the parameters in the theoretical model enable us to predict and manipulate the population dynamics in ecological and evolutionary contexts.

Chapter 4

Ecological analysis

The simulation model described in the previous chapter showed the emergence of complex bi- and tritrophic population dynamics. In order to understand, predict and control these ecological dynamics, a corresponding theoretical framework is developed in this chapter. The density dependencies of the various ecological processes (e.g. population growth, consumption, mortality) that arise from the model are experimentally determined and put together in a classical ecological model that consist of a set of ordinary differential equations. The models obtained in this fashion allow us to gain insight into the emergent dynamics in two ways: (1) it enables the control and prediction of the ecological dynamical consequences by of changes in simulation parameters or evolutionary changes (discussed in next chapter), (2) it provides a theoretical basis to study the consequences of lifting some of the assumptions implicit in the classical model.

The goal of the simulation model, and in fact of all individual-based models in theoretical biology and ecology, is to allow modelers to investigate questions that have been difficult or impossible to address using a classical state-variable approach (Hutson et al., 1988). The state variables most often used in ecology are population numbers or densities, and the ecosystem is described in terms of changes thereof in space and/or time. In contrast, individual-based models use individuals as their basic unit. The homogeneity assumed in classical models, caused by 'lumping' individuals in population averages, makes them inappropriate to study the self-organization of spatiotemporal patterns that arise from local interaction between individuals (Hogeweg and Hesper, 1990).

The self-organised spatial clustering and temporal patterns in population dynamics observed in the simulation model in the previous chapter were not part of the model specifications itself. Rather, they emerged from the behavioural and trophic interactions between individuals. This, however, does not necessarily imply that classical models are inappropriate to analyse these dynamics, as the regularities observed in the population dynamics (in figs. 3.5 and 3.6) already suggested. In these cases, the simulation model was constructed in co-

herence with some classical assumptions. The populations are monomorphic since all individuals in a population are equal. And the spatial distribution is approximately homogeneous, because all offspring is placed at a random location in the environment.¹

The most interesting phenomena are of course observed in situations where these assumptions are lifted, and the spatial and individual-based nature of the simulation model gives rise to qualitatively different outcomes. This will be shown in the evolutionary and spatially heterogeneous experiments in chapter 5. The classical model that is developed in this chapter serves as the theoretical framework in which the consequences of these changes can be studied.

In section 4.1, the various ecological processes in and between the resource and consumer population are experimentally determined and theoretically modelled. These processes are combined to form theoretical models for bitrophic and tritrophic ecosystems, for which a brief qualitative analysis is provided in section 4.2. Section 4.3 discussed the way in which these models allows us to manipulate and control the emergent population dynamics, and conclusions are drawn in section 4.4.

4.1 Resource and consumer dynamics

The increase and decrease of population sizes of resources and consumers depends on various factors. In this section, the various density dependent processes are modelled by examining the specification of the simulation model itself or the emergent processes it gives rise to.

4.1.1 Resource dynamics

To determine the change of the population size of the resource population over time, we first consider the simple situation of an ecosystem consisting of a resource population without consumers ($C = 0$). Resources can reproduce after a certain *reproductive period*, causing a fast growth at low values, and slow population growth at high values. The density dependence of the reproduction rate is explicitly defined by the simulation model. The probability of successful reproduction is determined by the difference between current population size and a user-defined carrying capacity. The resulting population dynamics shows a logistic growth. Such population growth can be described by the Verhulst or logistic model:

$$\frac{dR}{dt} = rR\left(1 - \frac{R}{K}\right) \quad (4.1)$$

where $r = 1/\text{reproductive period}$ and $K = \text{carrying capacity}$. A negative term dR can be included if natural mortality of resource is nonzero.

¹The simulation does not comply with the assumption of infinite populations, however. The populations in the simulation model are discrete and small.

The logistic model is applicable to populations whose growth depends on the population density. The birth and/or death rates are not fixed but are density dependent. At low population sizes, resources are abundant and birth rates are high. At high densities, however, birth rates decrease and death rates increase because of competition for reproduction. At population sizes above the carrying capacity, death rates are greater than birth rates and the population size decreases.

4.1.2 Consumer dynamics

Now we introduce consumers to the ecosystem that have interaction with resources. When a resource and a consumer interact, the resource dies, and energy is transferred to the consumer that reproduces when its energy level exceeds the *reproductive energy* level. Translated to the differential equation model, consumption constitutes a negative term to the resource dynamics and a positive term to the consumer dynamics. The relation between consumption and reproduction, and the mortality is analysed.

Consumption

Understanding the relationship between resources and consumers, or predators and prey, is a central topic in ecology. Several models of consumption (or predator feeding) rate exist. In the famous Lotka-Volterra model (Lotka, 1925; Volterra, 1926) the *per capita* consumption rate is a linear relationship which means that the number of consumed resources is directly proportional to the number of resources in the environment. This assumption has been shown to be wrong in two ways. The *functional response* indicates that the *per capita* consumption rate decreases with increased resource density (Holling, 1959, 1965). And *consumer interference* indicates that this also decreases with increased consumer density (Beddington, 1975; DeAngelis et al., 1975).

Functional response Several standard functions exist for describing the total consumption of an individual consumer depending on the resource density that are grouped under the name *functional response*. Holling (1959) defined three types of non-linear functional responses: maximized linear response, Monod saturated and sigmoid saturated response. All three types have a maximum number of resources that a consumer can consume within a certain time. Holling's motivation for this maximum was the *handling time*. Although the time to search resources decreases proportionally with the resource density, the time needed to handle the resource (e.g. chasing, eating, digesting) does not decrease with this density. Therefore, the consumption cannot grow infinitely fast, even in the presence of infinite resources. The functional response of the *per capita* consumption rate is described by:

$$f(R) = \frac{aR}{h + R} \quad (4.2)$$

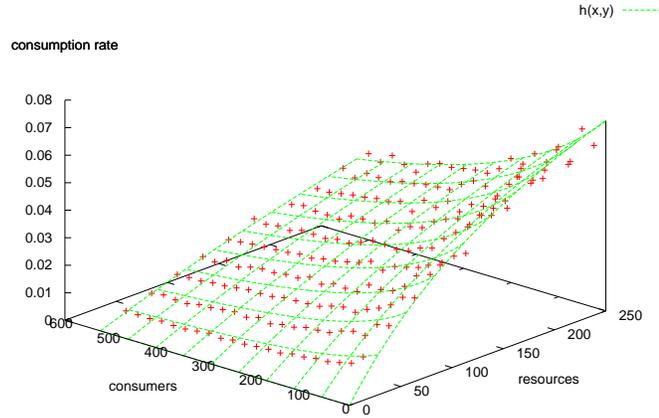


Figure 4.1: Beddington model: Functional response and consumer interference. Consumption rates saturate with increasing resource density, and decrease with increasing consumer density.

where a is maximum consumption rate, and h is the handling time.

Figure 4.1 shows the results of a series of experiments in which the population sizes of resources and consumers was controlled and varied across simulations. The average consumption rate of the consumers over the course of the simulation was measured and is depicted in red points. The consumption rates for one consumer ($C = 1$) over various resource population densities is saturated according to type II functional response.

Beddington model However, figure 4.1 also shows that the *per capita* consumption rates also dependent on the density of the consumers themselves. The consumption rates decreases as the consumer density grows. A reason for this that the searching efficiency decreases at high consumer densities. Beddington (1975) and DeAngelis et al. (1975) separately proposed an extension to the functional response that incorporates competition or interference between consumers:

$$f(R, C) = \frac{aR}{h + R + eC} \quad (4.3)$$

where e is the strength of the interference between consumers (see figure 4.3). Without interference, i.e. when $e = 0$, this function reduces to 4.2. With consumer interference, i.e. when $e > 0$, the consumption decreases as the herbivore population grows.

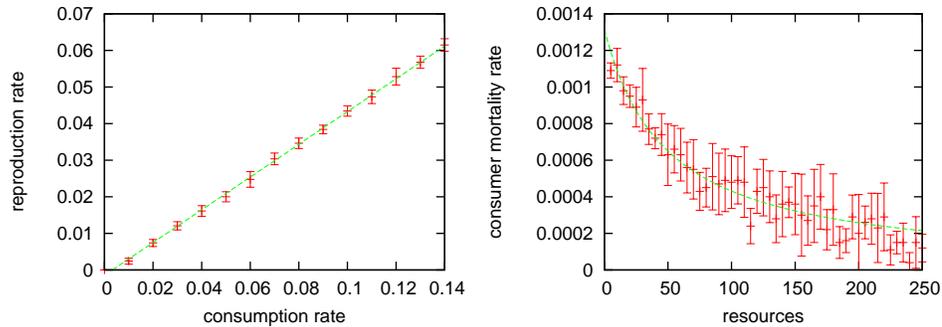


Figure 4.2: (Left) Linear conversion of consumed resources to consumer offspring. (Right) Mortality rate .

Reproduction

Most ecological resource-consumer models abide the conversion rule which states that the consumer reproduction rate is a function of its consumption rate (Ginzburg, 1998). In the simulation model, reproduction is based on individual energy levels, which can only be increased by consumption. It can thus be expected that the reproduction rate and consumption rate will show a strong correlation and this is evident from figure 4.2 which plot the linear relationship between them. The conversion efficiency coefficient c is the slope of this line.

Mortality

Eventually, consumers die of starvation. The right panel in figure 4.2 shows that the *per capita* mortality rate depends on the resource density, according to the function:

$$f(R) = d - \frac{dR}{i + R} \quad (4.4)$$

where d is the death or mortality rate, which is maximal when $R = 0$ (no resources), and half its maximum when $R = i$.

The conducted experiments show no dependence on consumer density, however. In the ecological models that are used in the following, the *per capita* mortality is simplified to be constant.

4.2 Ecological models

In the previous paragraphs, the various functional dependencies of the resource and consumer populations were examined. These relations are the components of the ecological model that is constructed in this section. By putting

the functional relations together in a set of differential equations, a theoretical model is obtained that captures the dynamics of the population dynamics that are observed in the simulations described in chapter 3. A mathematical and graphical analysis of such a simplified 'mini-model' allows us to gain insight into qualitative and quantitative aspects of the emergent population dynamics. Moreover, the relation between the simulation model to some well-studied theoretical models enables us to make use of the existing ecological literature that deals with these models.

In this section, theoretical models are constructed from the components laid down in the previous section that model the population dynamics of the bi-trophic (two-species) and tritrophic (three-species) food chains, as observed in section 3.2. In the construction and analysis of these models, a graphical approach is adopted (in which we follow de Boer (2006)).

4.2.1 Bitrophic model

To establish an ecological model describing the population dynamics of a bi-trophic ecosystem, as shown in figure 3.5, the functions describing logistic resource population growth (eq. 4.1), Beddington consumption (eq. 4.3), linear consumer conversion and constant *per capita* mortality are combined²:

$$\begin{aligned}\frac{dR}{dt} &= rR\left(1 - \frac{R}{K}\right) - \frac{aRC}{h + R + eC} \\ \frac{dC}{dt} &= \frac{caRC}{h + R + eC} - dC\end{aligned}\quad (4.5)$$

This ecological model enables us to study the behaviour of this system by graphical phase space analysis. This does not give us analytical or numerical solutions, but it does give us insight into the behaviour of the system over time. To do a graphical analysis, the nullclines of the system and its steady states are calculated, and sketched in a phase space diagram with a vector field. Nullclines are graphs of the set of points in the phase space for which the change of a population is zero, and steady states are located at intersections of nullclines.

The nullcline of the consumer population in the model 4.5 is obtained by solving the system when setting $dC/dt = 0$:

$$\frac{dC}{dt} = 0 \rightarrow C = 0 \vee C = \frac{ca - d}{de}R - \frac{h}{e} = \frac{R_0 - 1}{e}R - \frac{h}{e} \quad (4.6)$$

where $R_0 = ca/d$ represent the *per capita* consumption-dependent birth rate ca of individuals over their expected life span of $1/d$ time units. This gives the

²In the construction of the model in eq. 4.5 some simplifications are made: natural mortality of resources (apart from consumption-induced mortality) is not included, and the *per capita* mortality of consumers is simplified to a constant mortality.

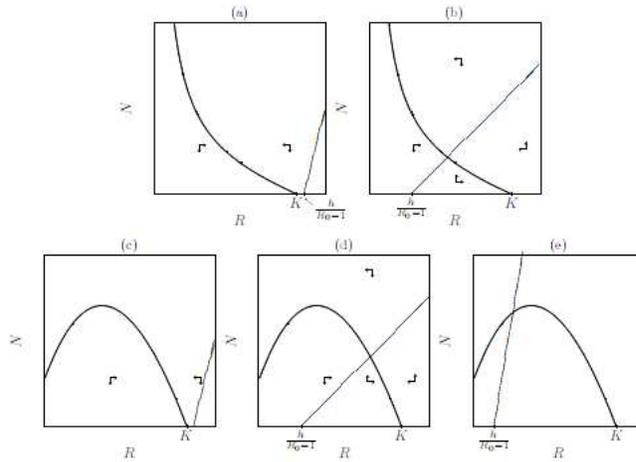


Figure 4.3: Five different qualitative states of bitrophic model with beddington functional response de Boer (2006)

expected number of offspring of an individual R_0 , or its the expected “fitness”³. The consumer nullcline in equation 4.6 defines a slanted line with slope $(R_0 - 1)/e$ that intersects the horizontal axis at $R = h/(R_0 - 1)$.

Graphical construction reveals that the resource nullcline can either have a vertical asymptote, or be a parabola, depending on the level of consumption. If consumption is limited, that is when the maximum consumption (at infinite resource density) is smaller than the rate of reproduction (i.e. $a/e < r$), one obtains a resource nullcline with a vertical asymptote. If $a/e > r$, one obtains a parabola for the resource nullcline that is truncated at the left (as a result of consumption limiting handling time). Figure 4.3 sketches these two situations in the top and bottom row respectively.

Enrichment If we look at the situations depicted in the bottom row, we see that the model has a stable equilibrium when the consumer nullcline intersects at the right side of the top of the parabola (situation (c) and (d)). Although one cannot determine the stability of the equilibrium analytically, numerical analysis has shown that the steady state can be unstable and that the behaviour of the model approaches a stable limit cycle. Enriching the system with nutrients, i.e. increasing the carrying capacity, destabilizes the system by moving from the qualitative situation (d) to (e). This is exactly what was done in the bitrophic food chain simulations reported in section 3.2.1.

³The notation of fitness as R_0 is adopted from epidemiology, where this concept originates. This should not be confused with our notation of R which represents the density of the resource population.

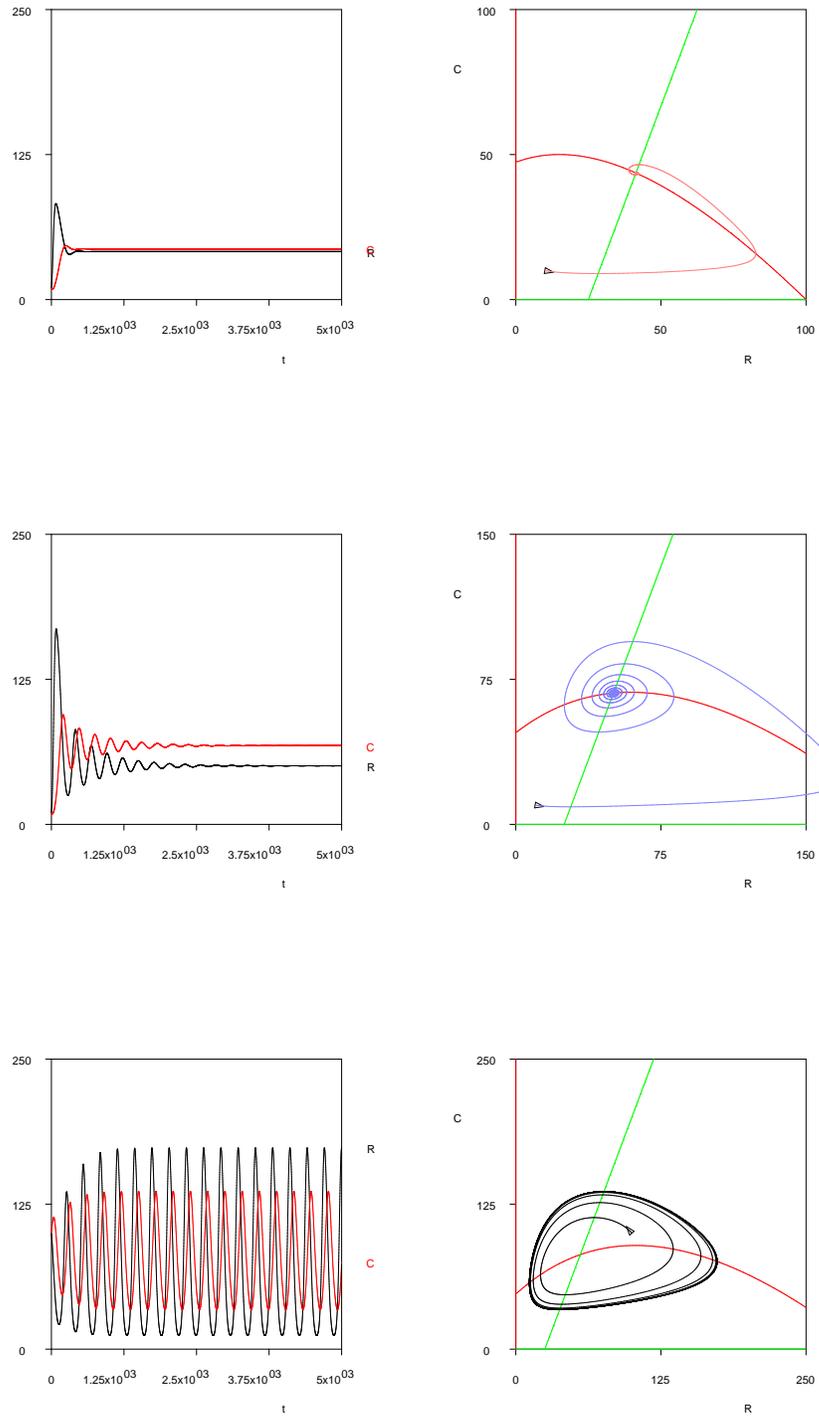


Figure 4.4: Time plots (left) and 2D phase plots (right) of bitrophic food chain model. Compare with figure3.5

The dynamics observed in these simulations can now be understood on theoretical grounds. To do this, we take theoretical model in 4.5 and fill it in it with reasonable estimates of the parameters. These estimates are the results of curve fitting of the functional dependencies. We can now change the parameter in the theoretical model that corresponds to the changes in simulation parameters across the various runs, in this case the carrying capacity K , and observe what happens to the behaviour of the model.

Figure 4.4 shows three numerical runs for different values of K . As can be expected, increasing the carrying capacity does not affect the consumer nullcline, since this parameter does not occur in equation 4.6. The truncated parabola describing the resource nullcline does change and expands to the right, since the logistic self-limitation of this population is decreased.

Comparing these results with the emergent population dynamics from the simulation model in figure 3.2.1 on page 32. This confirms the applicability of this theoretical model to our simulation model, and shows why the simulated bitrophic food chain is destabilized through enrichment.

Although the simulations were carried out in a way as to allow meaningful theoretical analysis, by tuning down spatial heterogeneity (by defining placement and carrying capacity globally), some notable differences in the behaviour of both models remain. Whereas the numerical solution shows a fixed *point* equilibrium for $K = 100$ (top row panels), the simulation model shows something that is more appropriately called a fixed *cloud* equilibrium. The simulated system continues to move around the equilibrium, because it is pushed out of the steady state by random fluctuations. Likewise, the situation for $K = 200$ (middle row panels), shows damping oscillations in the numerical case. The simulation model also shows damping oscillations towards a stable fixed point (or cloud), but when the system moves toward this equilibrium, it starts to show limit cycle behaviour with small amplitude, indicating the presence of an unstable equilibrium.

These differences can be attributed to the fact that the theoretical model presented here is 'ideal' in the sense that it assumes (1) infinite population numbers and (2) deterministic interactions and events. The simulation model, by contrast, deals with rather small populations and involves many stochastic or random events. Moreover, the theoretical models averages over differences between individuals by using state variables pitched at the population level. Although this seems justified by the fact that all simulated individuals are in fact identical, differences in consumption and reproduction rates between individuals do occur in the simulations. These fitness-affecting differences are due to differences in spatial location of individuals (e.g. by being born in an area with high or low resource density), and are only limited due to the fact spatial homogeneity was tuned down. When space is left free to organise itself, and the environment for consumers becomes more structured or heterogeneous, lumping individuals in population averages becomes increasingly dangerous, and the applicability of these theoretical mini-model is undermined.

Fortunately, many interesting phenomena remain to be investigated using the simulation model, even in the case of an approximately homogeneous

space. One of these things is extending the food chain to include a topconsumer.

4.2.2 Tritrophic model

The classic ecological models of interacting populations are typically restricted to two competing populations. It has long been recognized, however, that this limited caricature of ecosystems only applies to a small part of natural systems, where many species interact with each other. Even the simplest models in involving three interacting species, such as the tritrophic food chain discussed here, can show remarkably complex behaviour.

Topconsumers are identical to the consumer population discussed before, except that these secondary consumer feed on the primary consumer. Therefore, the dynamics of consumption, reproduction and mortality of topconsumers is qualitatively the same as examined for the consumer population. This allows us to extend our bitrophic model to a tritrophic model by simply adding a topconsumer equation dT/dt :

$$\frac{dR}{dt} = rR\left(1 - \frac{R}{K}\right) - \frac{aRC}{h + R + eC}$$

$$\frac{dC}{dt} = \frac{c_C aRC}{h + R + eC} - \frac{bCT}{i + C + fT} - d_C C \quad (4.7)$$

$$\frac{dT}{dt} = \frac{c_T bCT}{i + C + fT} - d_T T \quad (4.8)$$

in which b is the maximum predation rate⁴, i is the handling time, f is the predation interference, and c_T and d_T represent the conversion and mortality rates of topconsumers respectively.

The nullclines of this system are sketched in the three-dimensional phase space in figure 4.5. In the absence of topconsumers, when $T = 0$, the system reduces to the bitrophic system. Indeed, in the ground plane, we can recognize the bitrophic phase space (red = resource nullcline, and green = consumer nullcline). When $T > 0$, the consumer population is limited by the topconsumers along the vertical plane (blue = topconsumer nullcline). The consumer nullcline is curved in all three dimension as this population depends on the resource density by reproduction, on the topconsumer density by predation and on its own density by consumer interference.

Strange attractors Tritrophic food chain models can exhibit many interesting types of chaotic dynamics involving multiple attractors not observed in bitrophic systems. Systems of three interacting population models have been used to emphasize the importance of occurrence chaos in natural ecologies,

⁴The term 'predation' is chosen here to distinguish it from the consumption of resource by primary consumers, although these terms represent the same theoretical process.

where typically more than two species interact in trophic structures. Hasting and Powell (1991) used numerical simulation of a tritrophic model similar to the one above (except with Holling type II functional response) to show such systems exhibit chaotic behaviour for biologically reasonable parameters. They initialised the system such that in the absence of topconsumers the bitrophic system was involved in limit cycle behaviour. For different values of the handling time h of consumption, the system was observed to exhibit stability, limit cycle behaviour and chaos. A strange attractor found in a wide range of parameters within the chaotic regime was labelled a (up-side-down) tea cup attractor, because of its characteristic appearance in 3D phase space.

Although a range of beautiful complex patterns can arise from theoretical tritrophic models, these may be more interesting to mathematicians than to ecologists, since the latter deal with systems that are inherently discrete and stochastic. The fine structures in complex behaviours as Lorenz and teacup attractors are unlikely to be observed in noisy natural ecosystems. Field ecologists, therefore, have good reason to be suspicious towards *too* theoretical work in ecology (Ives and Jansen, 1998).

This does not imply, however, that theoretical studies of complex and chaotic behaviour are unimportant to ecology. Complex nonlinear behaviour is ubiquitous in biological systems. Moreover, stochasticity and discreteness not only reduce, but in fact play roles in the formation of such behaviour. Although many fine structures may indeed be obscured by stochasticity, these are replaced by other interesting complex behaviours that arise from stochasticity and discreteness. Theoretical models incorporating these features may prove more applicable to natural ecologies, and spatially explicit individual-based systems like the simulations presented in this thesis are part of this class of models.

Ives and Jansen (1998) used discrete and stochastic versions of a general theoretical tritrophic model and found the occurrence of a complex attractor in the shape of a torus in all cases. This kind of attractor is also observed in numerical simulation our theoretical model. The middle panel of figure 4.5 shows a trajectory in 3D phase plot of toroidal attractor obtained in model 4.8. The 2D phase plots (bottom row) show the same trajectory in the resource-consumer plane (left panel) and the consumer-topconsumer plane (right panel). This does not only shows similarity in shape to the torus found by Ives and Jansen (1998), but also to the population dynamics observed in the tritrophic simulation in figure 3.6 on page 35.

The torus attractor consist of short-term oscillations of resource and consumer and longer-term oscillations of consumer and topconsumers. Whereas in torus attractors these oscillations operate relatively independent of each other, systems involved in Lorenz-type attractors can switch back and forth between two oscillations in different planes. Large oscillations in the resource-consumer plane occur when T is small (but stable) and can switch to a situation of oscillations in the consumer-topconsumer plane when resource density is R is high (and stable near carrying capacity). For a small parameters domain, the

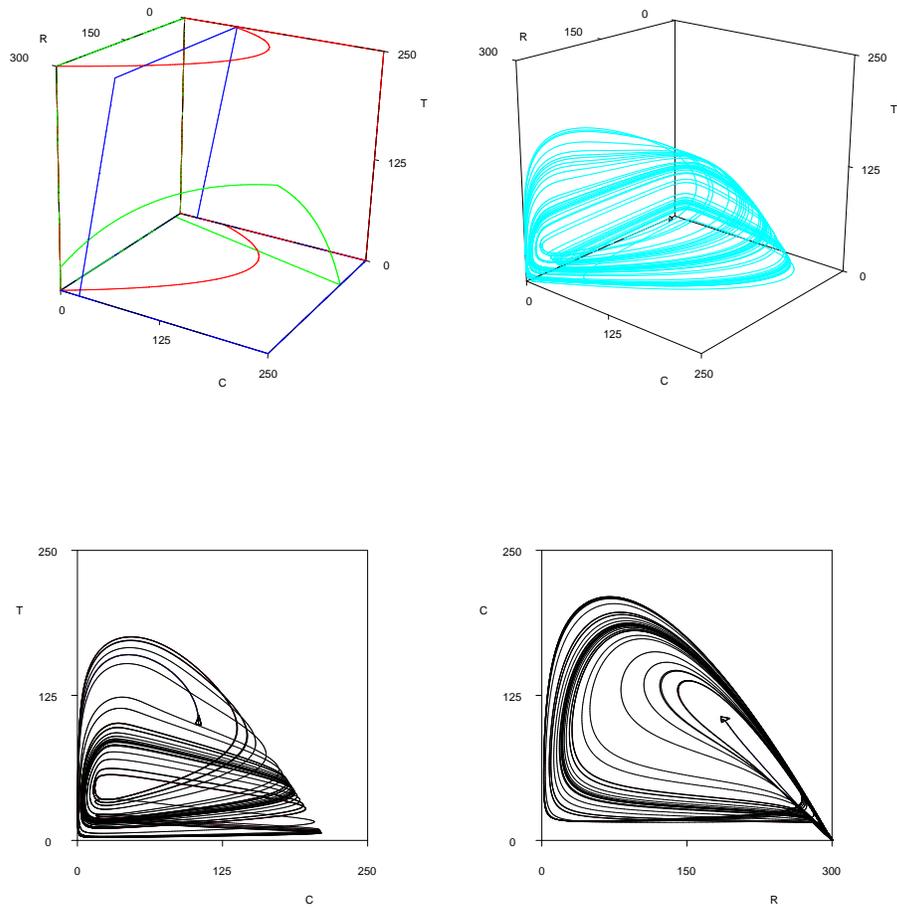


Figure 4.5: (Top left) Nullclines of tritrophic model in 3D phase space. (Top right) 3D trajectory plot of torus-shape attractor. (Bottom) 2D plots of the same trajectory. Compare with fig. 3.6 on page 35

torus is reduced to such a regime in the theoretical model. This type of complex behaviour was also found in simulation runs, but was unable to persist. The situation of oscillation in the resource-consumer plane, where T is very small, often leads to the extinction of the topconsumer population by random fluctuations in population size.

4.3 Controlled ecosystem

In many cases, and especially in evolutionary studies, we will be interested in simulations of complex trophic ecosystems in which all participating populations persist over long periods. The knowledge acquired of the simulation model in chapter 3 and the theoretical model in the previous sections of the present chapter provides us with some powerful tools to understand and manipulate the emergent population dynamics. In the last section, the emphasis was put on the knowledge we can gain about our simulation results by constructing a theoretical mini-model, through mathematical and graphical analysis, as well as from existing literature. This theoretical knowledge can, of course, also be put to use in predicting or manipulating the simulation model.

Despite the underlying complexity of the spatially explicit individual-based model, the virtual life simulation model can serve as a controlled ecosystem to allow ecological and eco-evolutionary studies (and comparison to ODE models). The need for control over the emergent population dynamics is important in the effort to understand the interplay between various patterns that emerge from spatial and/or eco-evolutionary systems.

Suppose, for example, we want to study chaos in tritrophic ecosystems, like Hasting and Powell (1991). For this we may need to initialise the simulation model in a situation where the resource-consumer population are involved in limit cycle behaviour in absence of topconsumers (i.e. $T=0$). The theoretical knowledge acquired over the simulated ecosystem can now help to achieve this, by the process of enrichment, for example. As described above, enriching the bitrophic system by increasing the carrying capacity causes the bitrophic system to destabilize. This knowledge can subsequently be used to obtain limit cycle behaviour from simulation.

The theoretical analysis has, however, also pointed out some limitations in our ability to manipulate the population dynamics. Although this is easy in the example of destabilization by enrichment, this control is hindered by emergence. The parameter K describing the carrying capacity in the theoretical model is directly specified in the simulation model by the parameter *carrying capacity*, which allows direct control over the nutritious richness of the environment. Other parameters that appear in the theoretical model, such as the maximum consumption rate a or handling time h , are harder to control. These parameters are not specified by the simulation model, but represent population-level approximations that describe the emergent interaction processes between population. The functional response in consumption, responsible for these parameters, emerges from the simulation due to temporal-

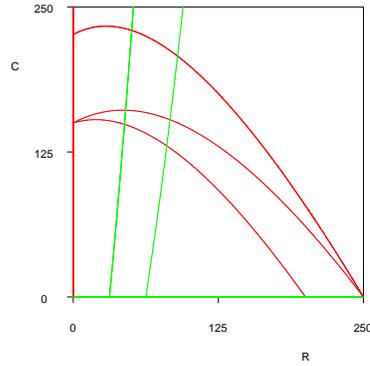


Figure 4.6: Controlling the emergent population dynamics

Param in ODE	Description	Influenced by simulation param	Default value
r	resource reproduction rate	<i>seeding rate</i>	20
K	carrying capacity	<i>carrying capacity</i>	250
a	maximum consumption rate	-not specified-	
h	handling time	-not specified-	
e	consumer interference	-not specified-	
c	conversion rate	<i>resource energy</i>	500
d	consumer mortality rate	<i>metabolism</i>	5.0

Table 4.1: Relations parameters in simulation model and ODE model

spatial limitations of consumer individual that must travel a distance towards their resources, even in abundance of the latter. Therefore, the maximum consumption and handling time can only be influenced indirectly. Controlling these processes requires one to change the interacting between individuals on the level of their situated activity. Increased consumption could, for example, be achieved by making consumer individuals faster or more efficient in their chemotactic behaviour.

Table 4.1 summarizes the correspondence between the parameters in the bitrophic theoretical model in eq. 4.5 and the simulation parameters in table 3.1. Figure shows an impression of the ways in which the simulated ecosystem can be manipulated by varying the simulation parameters in table 4.1.

4.4 Conclusions

In this chapter, the ecological interactions observed in resource-consumer interactions were analysed in terms of population-level variables to construct theoretical ecological mini-models. These models enable us to understand the population dynamics that emerge from interaction between simulated individuals. The simulation results of the previous chapter could be interpreted in terms of enrichment in the bitrophic case and a torus attractor in the tritrophic system. This theoretical modelling also enables us to control the emergent population dynamics, as was shown in the last section.

Chapter 5

Evolutionary experiments

So far, we have examined the virtual life simulation model (in chapter 3) and have developed analytical models to understand the ecological interactions that the simulations give rise to (in chapter 4). In this chapter, evolution is incorporated in the simulation model by introducing variation among individuals. Individuals of evolving populations are specified with inheritable and mutable traits that affect their fitness (henceforth called phenotypes) which causes differential reproduction. The evolutionary dynamics that emerge from these eco-evolutionary systems are analysed in terms of the ecological framework developed in the previous chapter.

A series of three eco-evolutionary experiments is presented that addresses several fundamental biological problems and processes that are considered to be important to explain major transitions and complexification in evolution. Moreover, these experiments aim to illustrate the methods of ecological analysis and exemplify the use of direct and indirect explanatory models.

The first experiment models the evolution of the reproductive efficiency of a consumer population. This causes over-exploitation of their common resources which results in the extinction of the consumer population. This process is known as the (evolutionary) “tragedy of the commons” (Hardin, 1968). The second experiment models coevolution in a predator-prey ecosystem. This results in the emergence of an evolutionary arms race (Dawkins and Krebs, 1979) that displays red queen dynamics (van Valen, 1973). The third experiment is similar to the first with the exception that here spatial self-structuring of resources and consumers is allowed. This causes the formation of spatio-phenotypic consumer groups causes reproductive restraint, which prevents the consumer population to fall victim to the “tragedy of the commons” through a process of group selection (Wilson, 1975; Wilson and Sober, 1994).

The applicability of ODE models to analyse these evolutionary processes is limited. The assumptions in ODE models that the populations consist of identical individuals and that their spatial distributions is homogeneous are both lifted in the experiments in this chapter. Therefore, the ODE models cannot

accurately model the ecological interaction of the simulation model. However, they can still serve as a theoretical framework in which to understand these ecological interactions. Moreover, they can serve to show the differences in evolutionary dynamics and outcome between the ODE model and the individual-based simulation model.

In order to allow the ODE model to serve as a meaningful framework, even when confronted with these differences, the 'genetics' involved in the evolutionary experiments is kept extremely simple. Individuals carry only a single gene that directly affects its behaviour (phenotype) without development or learning. Such phenotypes are single real values that directly control consumption rates or reproduction efficiencies. This means that there are clear relations between the evolving phenotypes in the simulation models and the parameters in the ODE models. Therefore, changes in the (population averages of) phenotypes can be interpreted as changes in one of the parameters in the ODE model.

Explanations of observed evolutionary dynamics differ from standard practice in artificial life. Instead of relating the observed emergent phenomenon back to specifications in the simulation model directly, the emergent evolutionary dynamics are explained indirectly by relating them to other emergent patterns (i.e. population dynamics and spatial self-structuring). The structure of the explanatory model grows more intricate as the interplay between these emergent patterns grows increasingly complex over the three experiments presented below. This results in explanations of evolutionary phenomena that are based on generic biological processes (evolution, population dynamics, and spatial structure) and are less dependent on the specific details of the simulation models.

5.1 Evolution

5.1.1 Introduction

The first simulation implements a simple eco-evolutionary system. It is similar to the ecological experiments in chapter 3, except that it includes variation among consumer individuals. This variation serves as a substrate for natural selection and alters the phenotypic configuration of the population by differential reproduction.

In this experiment, evolution can alter the energy that a consumer obtains from eating a resource. Natural selection favours consumers that gain more energy over the ones gaining less energy from resources, since the former reproduce more often. This leads to an overall increase of population size of the consumer population, ecological instability and eventually extinction of consumers. The evolutionary maximization of individual gain causes an over-exploitation of common resources as the "selfish" consumers fall victim to the "tragedy of the commons" (Hardin, 1968).

	Simulation parameter	Value
Environment	<i>size</i>	200
	<i>boundary</i>	R=fixed, C=none
Resource	<i>energy for consumers</i>	800 * <i>consumer phenotype</i>
	<i>reproductive period</i>	20
	<i>carrying capacity</i>	300
Consumers	<i>metabolic cost</i>	10.0
	<i>reproductive energy</i>	1000.0
	<i>mutation rate</i>	0.25
	<i>mutational operator</i>	$\sigma=0.015$

Table 5.1: Simulation parameters setting for evolution of conversion

5.1.2 Methods

A resource-consumer ecosystem is simulated in which only the consumer population evolves. Consumer individuals are specified with an inheritable and mutable trait or phenotype. The phenotype determines the energy that a consumer obtains from eating a resource. This phenotype is a real value between 0.0 and 1.0 which a child inherits from its parent directly or after mutation. Mutation is implemented as a Gaussian distribution with the parent trait value as mean and a small standard deviation. Resources and consumers are placed randomly over the environment at birth, resulting in a near-homogeneous spatial distribution. The simulation parameters are given in table 5.1.

5.1.3 Results

The main observables in this experiment are the population sizes and the phenotypes of consumers over time. The results are shown in figure 5.1. The evolution of phenotypes is shown in the top left panel. The top right panel shows a phase plot with nullclines that is obtained from the analytic resource-consumer model (see equation 4.5). The bottom row of figure 5.1 shows the population dynamics in a time plot (left) and phase plot (right).

Evolutionary dynamics

Each point in the phenotype time plot (top left panel) represents the phenotypic value of a consumer individual. The ancestors of consumer population were initialised with phenotype values around 0.2, which means they extract $0.2 * \text{resource energy}$ from the resources they consume. Mutations causes the phenotypes to diffuse which creates variation among consumers. Natural selection favours individuals with higher phenotypes, i.e. individuals with higher consumption efficiency. The phenotypes therefore shows an increase over time.

Although selection favours faster reproducing individuals, the process of

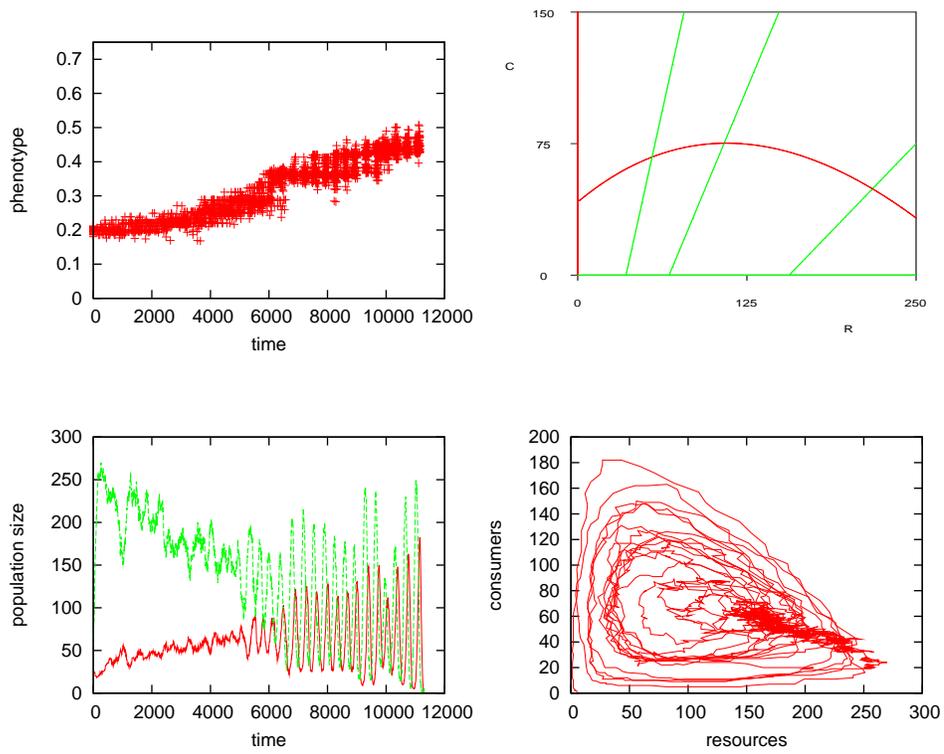


Figure 5.1: (Top left) Phenotypes of consumers over time. (Top right) Analytic phase plot of corresponding ODE model. (Bottom left) Population sizes over time. (Bottom right) Phase plot showing population dynamical trajectory.

removal of suboptimal phenotypes from the population is slower than the reproduction of the individuals carrying suboptimal phenotypes. Relatively slow reproducing individuals can therefore remain present in the population, and optimal phenotypes do not fix in the population (i.e. no complete dominance). Instead, the time plot of phenotypes shows the evolution of “clouds” of closely related phenotypes, known as quasispecies in molecular evolutionary biology (Eigen and Schuster, 1979; Eigen, 1971).

The increase of phenotypes does not proceed entirely smoothly. Most notably, a sudden increase is observed after 6000 time steps. Just before this point, the number of different phenotypes in the population increases and the start of a speciation event is observed (single quasispecies cloud breaking up into two clouds). The abrupt phenotypic increase is due to selection of the upper quasispecies at the expense of the lower part. Interestingly, these events coincide with changes in the population dynamics in the bottom left plot.

After approximately 11500 time steps, evolution comes to a halt when the consumers have phenotypic values around 0.45. To see why evolution is interrupted at these values, instead of evolving towards the maximum of 1.0, the population dynamics must be taken into account.

Population dynamics

The population dynamics (bottom left panel) shows an increasing instability of the ecological dynamics. Initially, the population sizes settle at a relatively stable equilibrium (only perturbed by stochasticity). During the first phase, the consumer population slowly increases in size at the expense of the resource population. In the phase plot (bottom right panel), this is reflected by the dark region of the trajectory. After 5500 timesteps, regular oscillations start to emerge. The amplitude of these oscillations increases over time which is shown in the phase plot as limit cycles. Eventually, this instability leads to the extinction of the resource populations, followed by the consumers that depend on them.

The increasing instability of the population dynamics can only be attributed to the phenotypic character of the consumer population (because nothing else changes) and visual inspection of both time plots (left panels) shows strong correlation. As the phenotypes of the consumers increase, the ecosystem shows increasing instability. Consumers evolve towards their own extinction.

Eco-evolutionary analysis

The observed dynamics can be understood by relating the simulated eco-evolutionary system to the analytical resource-consumer model:

$$\begin{aligned}\frac{dR}{dt} &= rR\left(1 - \frac{R}{K}\right) - \frac{aRC}{h + R + eC} \\ \frac{dC}{dt} &= \frac{caRC}{h + R + eC} - dC\end{aligned}\tag{5.1}$$

which is equal to model 4.5 as examined in more detail on page 46.

From the simulation model, we know that consumer reproduction fully depends on energy levels, and that energy levels depends on consumption (energy gain) and metabolism (energy loss). The phenotype determines the amount of energy that a consumer gains from eating a resource. The energy required for reproduction and the energy lost by metabolism is equal for all consumers, independent of their phenotypes. Therefore, an increase in energy gain causes an increase in the conversion of consumed resource into new offspring.

In the analytical model, this conversion is captured in the biomass conversion coefficient c . Changes in the population average of the consumer phenotypes in simulation can be interpreted as a change of this conversion parameter in the analytical model. Since evolution of the consumers adapts their biomass conversion parameter, we can study the observed population dynamics by examining the effect of changes of c in the ODE model.

To do this, we sketch the nullclines of the theoretical model with reasonable estimates of the parameters, for various values of the conversion parameter c . This is shown in the phase plot (top right panel), where the green consumer nullcline is drawn for three value of c . The values roughly correspond to population average of phenotypes in the first, middle and last periods in the evolutionary run.

The rightmost consumer nullcline corresponds to the initial period in the simulation. Stability analysis shows that the equilibrium in this configuration is stable (i.e. see panel (d) in fig. 4.3). As the consumer nullcline moves to the left as a result of an increase in c , the equilibria become unstable. Trajectories in phase space of the situation with the middle nullcline show limit cycles. The amplitudes of these limit cycles grows larger as the conversion parameter increases. Eventually, the amplitude of the limit cycles becomes too large and both populations go extinct.

The ecological interactions put a constraint on the evolvability of the biomass conversion. High conversion parameter values (e.g. greater than 0.5) are not evolvable in this ecosystem, because such values are not ecologically viable.

Tragedy of the commons

Evolution, the process of mutation and selective reproduction, acts to increase the fitness of individuals. This causes the growth of the consumer population, at the expense of the resource population. All consumers share a common resource pool in which the resources are renewed at a limited rate. The consumer population evolves to increase its reproduction rate, but the reproduction of

resources remains constant. This leads to an overexploitation of resources, followed by an explosion of the consumer population. As this offspring has little or no resources left to eat, because the resource population do not have a chance to grow back, the consumer population goes extinct. The consumer population has fallen victim to the “tragedy of the commons” (Hardin, 1968).

5.1.4 Conclusion

The results of this eco-evolutionary experiment illustrate the relation between evolutionary and population dynamics. Natural selection favours consumer individuals with higher energy gain, because they have more offspring on average, and therefore out-compete their less efficient conspecifics. The energy gain of the consumer population also has major influence on the population dynamics. In the ODE model, the evolution of this trait could be interpreted as change in the biomass conversion parameter c . Varying this parameter enables us to understand the observed population dynamics as well as understand the ecological constraints on the evolvability of conversion values. The evolutionary dynamics are constrained by the ecological principle of the “tragedy of the commons”.

This tragedy seems to suggest that cooperation and altruism cannot evolve because natural selection always favours selfish individuals. And indeed, this is predicted from the analytical model, and was confirmed by the simulation model. This results does, however, depend on the homogeneous spatial distribution that is implicit in the ODE model and was explicitly modelled in the simulation model. We will return to this point in section 5.3, in which similar experiments are conducted in a system that allows spatial heterogeneity. First, however, a coevolutionary simulation is presented in which an evolutionary arms race is observed that exhibits red queen dynamics.

5.2 Coevolution¹

5.2.1 Introduction

Coevolution of predators and prey is often understood in terms of evolutionary arms races. These arms races occur because predators and prey do not only compete on a behavioral scale, but also on an evolutionary scale (Dawkins and Krebs, 1979). Arms races between coevolving populations are considered as an important source of evolutionary complexification. However, arms races do not necessarily cause populations to be better adapted to their environment. Van Valen (1973) noticed that populations with long histories go extinct with the same frequency as populations with relatively short evolutionary histories. This indicates that the coevolutionary coupling between populations do not

¹A modified version of this section has been submitted to GECCO 2006 conference as “Red queen dynamics in predator-prey ecosystem”, W. de Back, M. Wiering, E. de Jong.

cause them to be better adapted to their environment, because this environment is evolving with them. Van Valen (1973) adopted the term 'red queen' to denote the fact that in coevolution, populations evolve only to maintain instead of improve their fitness.

Coevolutionary systems can be contracting, expanding, or involved in red queen dynamics, depending on the lag load. The lag load describes the reduction of mean population fitness due to the average trait of a population differing from its optimal value. That is, it is the distance a population lags behind, relative to its evolving antagonist. When the lag load of a population increases over time, it is losing relative to its antagonist, which (by contraction) can lead to extinction. The antagonist, in this case, has a decreasing lag load and will grow (expand) in population size. Both coevolutionary modes are not sustainable as they lead to the extinction of one of the competitors, unless these modes alternate. Red queen dynamics is a steady state in the evolutionary dynamics which occurs when both populations have a positive but constant lag load. In this case, both populations continually adapt to each other, while remaining at the same fitness.

The simulation results show the simultaneous emergence of patterns in population and evolutionary dynamics. We do not merely explain these dynamics as the emergent outcome of interaction between individuals, but explain these two patterns by relating them to each other, using the population dynamical model. In this fashion, it is established that the coevolving ecosystem is involved in an evolutionary arms race. After a short asymmetric period, in which the relative fitnesses (and lag load) changes, the arms race settles in a relatively steady state in which both populations do not receive any fitness benefits. That is, the arms race settles in a red queen dynamics.

5.2.2 Methods

A resource-consumer (or predator-prey) ecosystem is used in which both populations evolve. In contrast to the previous experiment, here, both predators and prey individuals are specified with a phenotype. The coevolutionary model that is used in this experiment here is adopted from van der Laan and Hogeweg (1995) and its spatial extension in Savill and Hogeweg (1997).

Phenotypes

This phenotype is a natural value between 0 and 100. When a predator and a prey interact (collide), the difference between the two phenotypes determines the probability of predation (i.e. energy transfer) according to

$$p(\text{predation}) \propto a = \exp\left(-\left(\frac{\text{pheno}_R - \text{pheno}_C}{\sigma}\right)^2\right) \quad (5.2)$$

which defines a Gaussian distribution, where the standard deviation σ can be interpreted as the specialism or generalism of the predators (kept constant in these experiments). Predation probabilities are highest (1.0) for predator-prey

	Simulation parameter	Value
Environment	<i>size</i>	250
	<i>boundary</i>	R=fixed, C=none
Resource	<i>energy for consumers</i>	550
	<i>reproductive period</i>	20
	<i>carrying capacity</i>	300
Consumers	<i>metabolic cost</i>	10.0
	<i>reproductive energy</i>	1000.0
	<i>mutation rate</i>	0.05
	<i>mutational operator</i>	+1 or -1
	<i>generalist-specialist σ</i>	10.0

Table 5.2: Simulation parameters setting for coevolutionary experiment

pairs that have equal phenotypes, and drops off with an increasing phenotypic difference. The phenotype axis is wrapped around, which means there are two ways of calculating the difference between phenotypes. The minimal difference is used to determine the probability of predation upon interaction.

Mutation Offspring have a small probability of mutation, which is equal for prey and predators. The mutation operator is implemented as a shift by one unit ($phenotype\ mutant = phenotype\ parent \pm 1$). Phenotypes greater than 100 or smaller than 0 are wrapped around by subtract resp. adding 100 units.

The simulation parameters are shown in table 5.3.

5.2.3 Results

The observables are the population sizes and the phenotypes of predators and prey individuals over a period of 100000 steps, and are plotted in figure 5.2. The patterns that are observed along these dimensions are understood by relating them to each other with the aid of the analytical resource-consumer model.

Population dynamics

The population sizes observed during the simulation run are plotted in the top row of figure 5.2. The phase plot (top right) shows the trajectory in phase space which takes a horn-like shape. This plot indicates movements of the equilibrium during evolution, as well as its stability. The left of the horn-like shape consists of large oscillations that reduce in amplitude as the attractors moves to the middle. The dark curved basis of the horn represents situations in which population dynamics are relatively stable. It is thus clear that the eco-evolutionary system traverses through large parts of the phase space encountering stable and unstable ecological equilibria over evolutionary time.

In contrast to the previous experiment, this does not occur in a unidirectional manner (as in the previous evolutionary experiment), as is clear from the time plot (top left in fig. 5.2). This plot shows alternating periods with qualitatively different behaviours: periods of instability characterised by oscillations, relatively stable periods, and near-extinction events of consumers. The system, for example, shows oscillatory behaviour in the initial phase, and after a period of relative stability, oscillations are re-encountered after some 75000 time steps. The coevolving system is thus able to move back and forth through ecological phase space.

To understand why this is possible, the analytical resource-consumer (or predator-prey) model is useful again:

$$\begin{aligned}\frac{dR}{dt} &= rR\left(1 - \frac{R}{K}\right) - \frac{aRC}{h + R + eC} \\ \frac{dC}{dt} &= \frac{caRC}{h + R + eC} - dC\end{aligned}\quad (5.3)$$

The difference between phenotypes determine the probability of actual predation on interaction between a predator and a prey (eq. 5.2). The more similar the phenotypes of interacting antagonist, the higher the probability of successful predation, thus increasing the predation rate. The movements of the equilibrium through phase space can thus be interpreted as evolutionary changes in the maximum predation rate a .

The bottom right plot in figure 5.2 shows an graphical construction of the analytical model in eq. 5.3. Nullclines are sketched for various values of the maximum predation rate a , and the (in)stability of the various equilibria are indicated by open (unstable) and closed (stable) boxes. The equilibria move to the right as the maximum predation rate decreases. By comparison of both (experimental and analytical) phase plots, the parts of the phase space through which the system traverses can be understood as changes in maximum predation rate a . It does not, however, enable us to appreciate why the equilibria can move back and forth through phase space without extinction. To understand this, the evolutionary dynamics must be examined in more detail.

Evolutionary arms race

The evolution of phenotypes of predators (red) and prey (green) are depicted in the bottom left plot in figure 5.2. The double stripes show the predators chasing the prey (and prey evading the predator) phenotypes phenotypes through phenotype space (with wrap-around boundaries). The predators chase their prey in real space over ecological time, as well as in phenotype space over evolutionary time. They are bound in an evolutionary arms race.

From eq. 5.2 we know that predation rates are determined by the *differences* between the phenotypes of predators and prey. The solid line in the plot represents the population average phenotypic difference between the predators

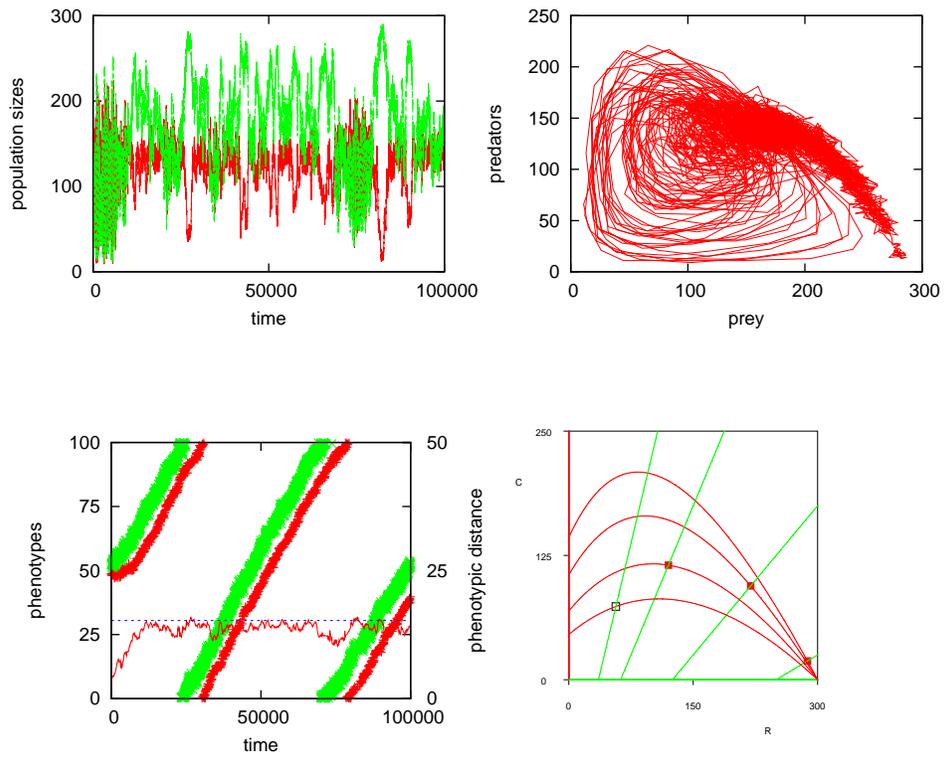


Figure 5.2: (Top) Population dynamics over time (left) and phase plot (right). (Bottom left) Phenotypes and phenotypic distance over time. (Bottom right)

and prey (read from right axis). The dotted line represents a critical value of phenotypic difference that is used later. This enables us to understand why the system moves back and forth through the phase space: the phenotypic difference grows larger and smaller during evolution.

In the first period, the phenotypic difference increases rapidly. In the background, we see that this is caused by the prey phenotypes (green) that evolve away from the predator phenotypes (red). During this period, the prey population expands (grows in size), because they are predated upon less.

A paradox of predation Surprisingly, the expansion of the prey population does not go at the expense of the predator population. The predator population is not contracting, but is expanding simultaneously with the prey population. This can be understood by looking at the two leftmost equilibria in the analytical phase plot (bottom right). At initialization, the phenotypic difference is low, and therefore the predation rates are high, which is represented by the leftmost (unstable) equilibrium. As the prey evolve away from the predators, the predation rates decrease, and the system moves towards the second left equilibrium. This equilibrium is not only more stable, it also harbours larger population sizes of both prey *and* predators. This is somewhat paradoxical, since it means that decreasing predation rates are beneficial for the size of the predator population (as well). At the least, it is contrary to the expectations based on the expansion-contraction modes of coevolution, which seems to be an unsuitable model, even for the simple eco-evolutionary system used here.

In this initial period, the prey evolves away from the predators, but the predators (representing the environment for the prey) do not follow suit. The prey does not evolve, however, away from the predators far enough to cause extinction of the predators. After the initial increase, the phenotypic differences start to decrease when the predators start to chase the prey phenotypes. To understand the evolutionary dynamics of the predators, the analytical model is employed in a more elaborate sense to define a concept of fitness.

Fitness We adopt the concept of expected fitness R_0 from epidemiology. The expected fitness is the number of offspring per time unit times the expected lifespan of an individual. A population with $R_0 < 1$ decreases in size (and eventually goes extinct), population sizes are stable for populations with $R_0 = 1$, and increase for $R_0 > 1$. To determine the expected fitness of a population, the reproduction and mortality of a population must be defined explicitly. Because this is not the case for the prey population (reproduction and natural mortality are collapsed into the intrinsic growth rate r), the R_0 is undefined for the prey population. Therefore, we only study the fitness R_0 for the predators. This is justified by the fact that the prey fitness is directly coupled to the predator fitness. Moreover, assessing the predator fitness is sufficient for determining coexistence.

Predators have an average mortality rate of d every step, and its expected

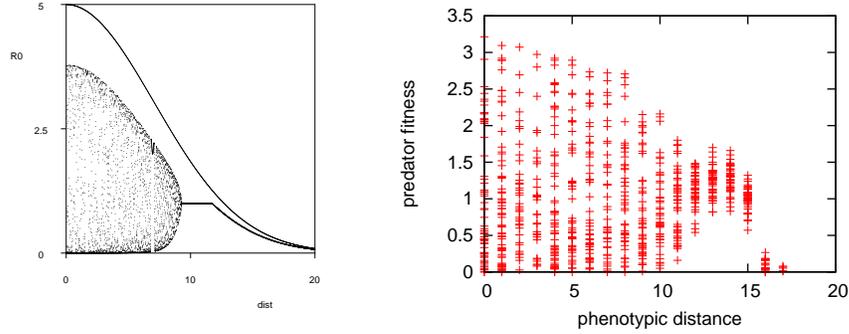


Figure 5.3: Analytical (left) and experimental (right) fitness R_0 over phenotypic difference.

lifespan is therefore $1/d$. The number of offspring per time unit, or predators' rate of reproduction, can be estimated in different ways. When ideal circumstances for the predators are assumed, the expected fitness of a predator is

$$R_0 = \frac{ca}{d} \quad (5.4)$$

where ca is the maximum predation rate multiplied by the number of offspring per consumed prey. When this concept of fitness is assumed, the fitness landscape can be represented as a smooth hill (upper curve in left plot of fig. 5.3). This is expected, as R_0 (eq. 5.4) is proportional to the maximum predation rate a , which decreases with the phenotypic distance (eq. 5.2).

Note that this fitness landscape is not static. The fitness is plotted against phenotypic difference, and this difference changes for any given phenotype due to evolution of the antagonist species. Therefore, the "hill" in the landscape moves along with prey phenotypes. To remain at a given fitness, the predator population continuously adapt and climb the hill, whereas the prey phenotypes adapt and move the hill to the left. Since in coevolving ecosystem these opposing forces are continually active, these have to be balanced in order to obtain sustainable coevolution.

The fitness landscape of the eco-evolutionary system is not as smooth as suggested by the smooth curve. We already showed that the evolution of predation rates influences the stability of the ecological system. This renders evolutionary fitness optimization difficult and puts ecological constraints on the evolutionary dynamics.

The fitness values along the smooth curve can only be under situations that are ideal for predator reproduction. The size of the prey population is assumed to be maximal ($R = K$), such that the predation is always saturated at its maximum. This maximum is only obtained when there are no other predators

present, such that there is no predator interference ($eC = 0$). These situations do, however, not occur in ecologically embedded predator-prey coevolution.

When these idealistic assumptions are lifted, and the various density dependences are taken into account, the expected fitness of predators becomes

$$R_1 = \frac{caR/(h + R + eC)}{d} \quad (5.5)$$

where the numerator again represents the predators' rate of reproduction. When this fitness is plotted against the phenotypic distance, a much more complex and dynamic landscape appears (bottom curve in left plot of fig. 5.3). When the phenotypic distance is large and decreases, the expected fitness slowly increases to $R_1 = 1$ and remains at this value for a relatively wide domain of phenotypic differences. When the phenotypic difference decreases below a critical value (at approx. *difference* = 9), the predator population can have many different fitness values.

In this region, left of the bifurcation, the predator population is involved in population dynamical limit cycles, which render fitness evaluation difficult. For a population with a given phenotypic difference, the expected fitness can be high ($R_1 \gg 1$) or low ($R_1 \ll 1$), since the birth and death processes are subject to oscillations. The fitness realised by a population in this region depends more on population dynamics than on its exact predation rate.

The right plot shows the results of experimental measurements of the predators' fitness over various values of phenotypic difference. Comparison with the analytical plot shows that the experimental data is in coherence with the concept of fitness as in equation 5.5.

Red Queen dynamics

This concept of fitness enables us to understand why the prey population can initially evolve away from the predator population, but is chased by the predators in a later stage when the difference grows. The fitness of the predators shows a steep drop when the phenotypic difference increases over 15. Above this difference, the predator fitness $R_1 < 1$, which means an average predator has less than one offspring. If all predators have fitness $R_1 < 1$, the population goes extinct. This critical value was plotted as the dotted line in the phenotype plot in figure 5.2.

This does not happen in the simulation, since there is phenotypic variation among individuals in the population. When part of the predator population comes to have a fitness $R_1 < 1$ by evolution of prey, selection favours other predators that have less phenotypic difference to the prey. When the critical phenotypic difference is reached, there are suddenly fitness differences causing selection pressures within the predator population. Indeed, in figure 5.2 we see that the predator phenotypes plotted in the background plot have increased their phenotypic adaptation rate to the level of the prey at the time that the difference encounters this critical value. This means that the predators (the

environment of the evolving prey) starts to change as fast as the prey is changing.

During the rest of the simulation, the phenotypic difference remains mainly in the region between 12 and 15, with some fluctuations. These fluctuations can be related to the predator fitnesses in this region. In figure 5.3, we see that in this region, the fitness landscape is fairly flat and therefore induce little selective pressure. For the predators to remain in this comfortable region, their phenotypes must evolve as quickly as the prey phenotypes. In absence of a strong selection pressure, however, the predators are likely to lag behind. When the difference has again grown to the critical value 15, strong selection cause the predators to catch up with the prey. This enables us to interpret the long evolutionary walk close to this critical difference value (fig. 5.2). Moreover, it explains the alternating behaviours and reversals of movements through phase space that were observed in the population dynamics. These reversals coincide with the encounters with this maximum.

Although the phenotypic difference is likely to be near the critical value, since the selection pressure in the predator population is weak, so that the lag load increases until strong selection reinstalls when the critical difference is reached. This is indeed most frequently observed in the simulation. There is, however, also one event (observed around 75000 steps) in which the differences suddenly decreases significantly. This sudden movement is caused by the fact that a small portion of the evolving predator population that has rapidly evolved towards the prey (possibly facilitated by the near-neutral evolution in region 12-15). As figure 5.3 shows, predators left of the bifurcation can gain high fitness $R_1 \gg 1$, and therefore quickly take over the predator population. This is only temporary, however, since the increased predation rate causes stronger selection pressure in the prey population which evolves away, as it did during the initial phase.

The two limiting mechanisms discussed above cause the evolutionary dynamics (of the observable of phenotypic difference) to be constrained to a certain parameter range in which predators and prey can coexist. The evolutionary adaptation rates of predators and prey balances themselves, such that the environment of an evolving population moves as fast as it is adapting to it. This is where the red queen reigns: "Now, here, you see, it takes all the running you can do, to keep in the same place" Carroll (1899).

5.2.4 Conclusions

In this experiment, an eco-evolutionary simulation of predator-prey coevolution was presented. The results show the emergence of an evolutionary arms race, a paradox of predation and red queen dynamics. These processes were explained by relating the population and evolutionary dynamics to each other with the aid of an analytical predator-prey model.

Initially, the evolutionary arms race was asymmetrical as the prey population evolves away from the predator phenotypes. During this period, a paradox of predation is observed, as both the prey and the predator populations expand when predation rates decrease. This paradox was resolved by analytical modelling of the evolutionary change in the maximum predation rate a . The coevolving system is able to move back and forth through phase space, dependent on the phenotypic differences. After the initial increase of this difference, a decrease is observed after an encounter with the critical phenotypic difference of 15. This reversal was explained by developing a concept of fitness R_1 and by showing that $R_1 < 1$ for predator individuals with a phenotypic difference greater than 15. As a result, the phenotypic difference stays in the comfortable region between approximately 12 and 15.

The sustainability of the evolutionary arms race between predators and prey, i.e. the ecological coexistence of both populations over evolutionary time, is due to Red Queen dynamics. A lower bound of phenotypic difference causes the prey population to evolve away from the predators, and an upper bound causes a strong selection in the predator population towards the prey phenotypes. The Red Queen dynamics emerge as ecological constraints on the evolvability of the phenotypic difference.

5.3 Multi-level selection

5.3.1 Introduction

There has been a much controversy and debate about the level upon which Darwinian selection operates. In the last decade, however, biologists have adopted spatially explicit and individual-oriented modelling approaches which show the emergence of new levels of selection through spatial self-structuring (Boerlijst and Hogeweg, 1991; Johnson and Boerlijst, 2002). When the phenotypic similarity of individuals within a self-structured patch, cluster or wave (henceforth: group) is higher than the similarity of individuals between groups, evolution can use these groups as a substrate for selection. Selection on the level of groups can be opposite to the direction of individual selection, and therefore influences evolutionary dynamics (discussed in section 2.3.3).

Although group selection is accepted among biologists as a theoretical possibility, many hold that its impact is too weak to be of much interest for the study of natural evolving ecosystems of small and viscous populations (Williams, 1966). To investigate the emergence of group selection under these conditions in our simulation model, we use a modified version of the simulation model of the evolution of the energy gain in section 5.1.

Consumers are again specified with an inheritable and mutable phenotype that determines the amount of energy they extract when eating a resource. In contrast to the previous experiment, however, the spatial structure of the model is no longer homogeneous (not well-mixed). The carrying capacity and resource placement are locally defined, which enables spatial self-structuring

of the resource population (see section 3.2.2). Moreover, the offspring of consumers are placed close to their parents, through which the spatial distribution of the consumer population becomes heterogenous as well.

This small change in the model causes qualitatively different ecological and evolutionary dynamics. The consumer population does not go extinct, and the phenotypes shows the evolution of reproductive restraint. The analysis of the emergent evolutionary dynamics is based on studying its relation to other emergent patterns in population dynamics and spatial structure. The emergence of spatio-phenotypic clusters results in a selection pressure between clusters that is opposite to the selection pressure acting upon individuals.

Although restraining reproduction is harmful to the fitness of the restraining individual, it confers an advantage to the fitness of the group it belongs to. This evolution of altruistic, in which individual-level fitness is sacrificed for the persistence of the population, is enabled by group selection. Groups of selfish individuals fall victim to the “tragedy of the commons” leaving a population of groups of more restrained consumers.

5.3.2 Methods

The model is similar to the one described in section 5.1, with the exception that the resource carrying capacity and the placement of resources and consumers is now defined locally². The spatial distribution of the consumer population structures itself by placing new consumers close to their parents. Since the consumers are reactive agents that act upon their local environment, consumers in the same environment behave in similar ways (although sensory noise prevents agents to travel exactly the same trajectories). Therefore, small areas with relatively high consumer populations are formed. The metabolic energy cost is increased (compared to previous experiments) such that the exploration of a consumer is limited in (life)time and space (ensuring local interaction).

Simulation parameters are shown in table 5.3.

Models

Spatial structuring of population has influence on both ecological and evolutionary dynamics. Our aim is to investigate the influence of spatial structuring on the evolutionary dynamics. Therefore, the ecological impact should be distinguished from the evolutionary. To this end, three different models are simulated that differ in the way individuals are placed in the environment:

Model 1 Random placement: uniform distribution over whole environment

²The local definition of resource carrying capacity and placement is here implemented as follows: The environment is divided into an array of ‘demes’ in which the resource population growth is logistic. The carrying capacity is determined locally for an individual over its neighbouring demes (Moore neighbourhood). Resource offspring is placed near its parent according to a Gaussian distribution, such that empty neighbouring demes can be seeded and can grow. In absence of consumers, the resource population structures itself in various patches (similar to fig. 3.8 on page 38)

	Simulation parameter	Value
Environment	<i>size</i>	400
	<i>boundary</i>	R=fixed, C=none
Resource	<i>energy for consumers</i>	1000 * <i>consumer phenotype</i>
	<i>reproductive period</i>	20
	<i>carrying capacity</i>	local: 16*16 demes of max. 25 res.
	<i>placement</i>	Gaussian around parent; $\sigma = 25$
Consumers	<i>metabolistic cost</i>	15.0
	<i>reproductive energy</i>	1000.0
	<i>mutation rate</i>	0.25
	<i>mutational operator</i>	$ph_{Mutant} = ph_{Parent} \pm 0.01$
	<i>placement</i>	Gaussian; $\sigma = 5$

Table 5.3: Simulation parameters

Model 2 Close to random conspecific: Gaussian distribution around (uniform) randomly chosen conspecific

Model 3 Close to parent: Gaussian distribution around parent

The first model corresponds closely to the experiment in section 5.1. This forms a baseline experiment in which resources and consumer are distributed homogenously. Model 2 results in spatial clusters (groups) whose members are not phenotypically related, because offspring is placed in a random group. In model 3, the spatial clusters mainly consist of phenotypic relatives because offspring is placed close to their parents.

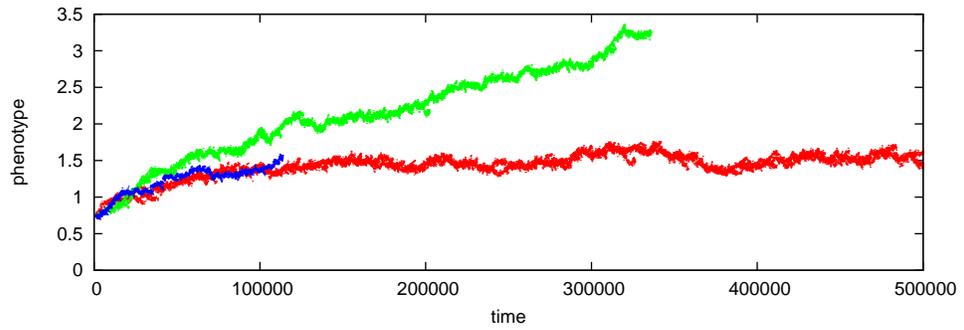
5.3.3 Results

Figure 5.4 shows the evolutionary and population dynamical results of a typical simulation run for three scenarios. The top panel shows the evolution of the phenotypes, and the bottom panel shows the population dynamics corresponding to these simulations.

The blue line shows a gradual evolution of phenotypes in model 1, which leads to a depletion of the resource population, followed by the extinction in the consumer population, comparable to the results in section 5.1. This is also observed in model 2, but only after evolving to much higher phenotypic values and much larger population dynamical oscillations. The eco-evolutionary system in 3rd model does not evolve to extinction and shows a stable moderation of phenotypes. Since the phenotype (energy extracted from a consumed resource) here almost directly determines the reproduction rate of individuals, this represents an example of the evolution of reproductive restraint.

First, we examine the eco-evolutionary difference between first and the second model which can be analysed in terms of population dynamics. But to understand the difference between the model 2 and 3, that are ecologically rather

Evolution of phenotypes for three models (blue=model 1, green=2, red=3)



Population dynamics of three models over time (left) and phase plot (right)

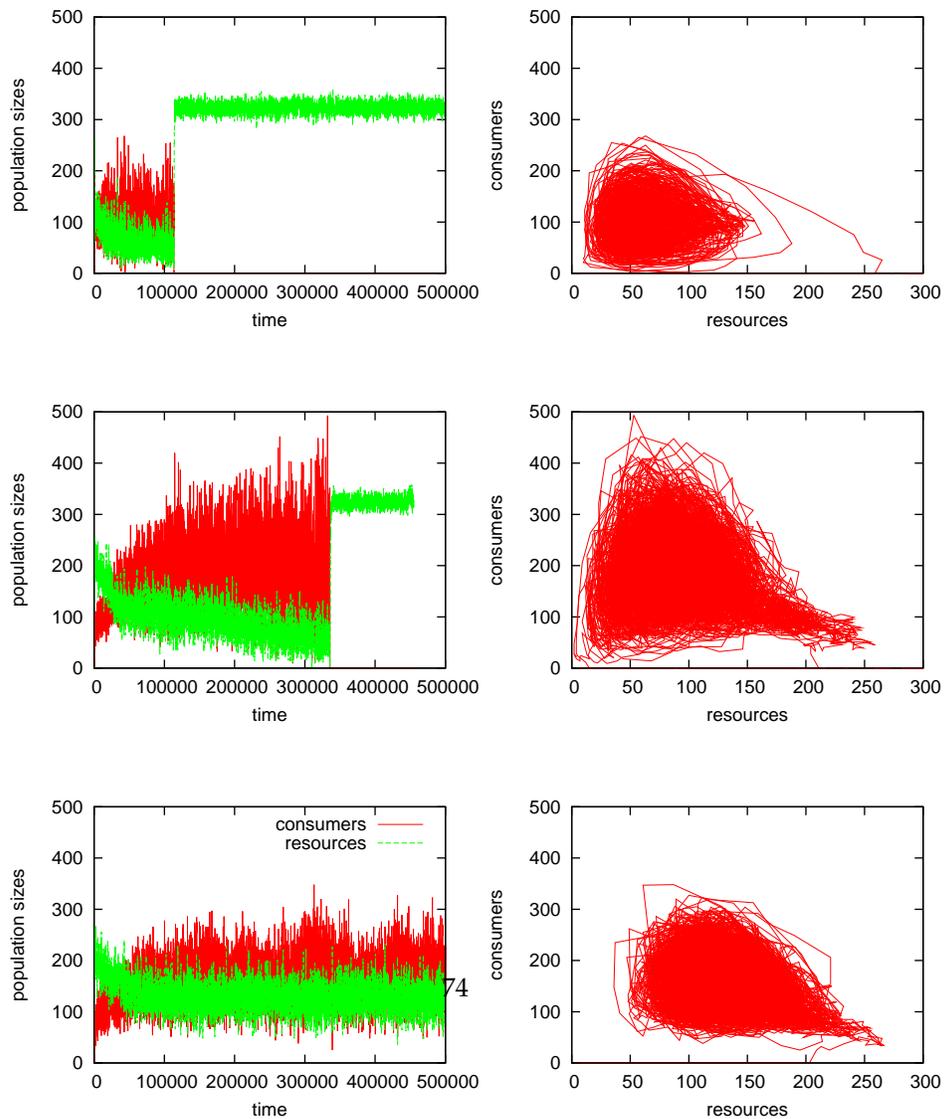


Figure 5.4: Evolution of reproductive restraint.

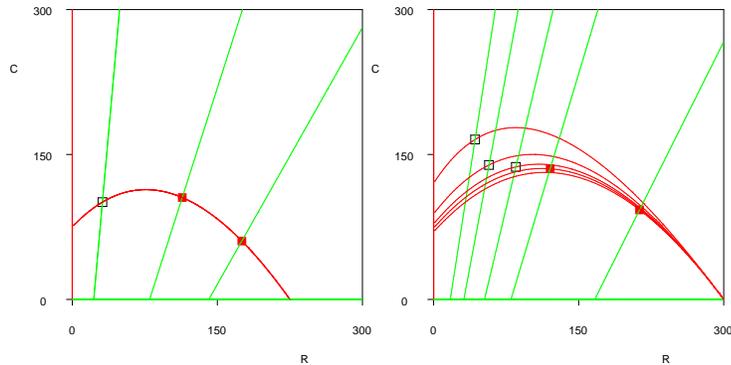


Figure 5.5: Analytical phase plots of model 1 (left) and models 2 (right) with nullclines

similar, but quite different from an evolutionary point of view, requires a different view. An inspection of the dynamics of the spatial structuring, shows that the formation of spatio-phenotypic clusters gives rise to a new higher level of selection.

Population dynamics

The results obtained in model 1 and 2 are qualitatively similar. Both evolutionary trajectories increase in phenotypic values, and therefore in reproduction per consumed resource. This increase results in a “tragedy of the commons” for both models. These situations are comparable to the results of the experiment in section 5.1.

The population and evolutionary dynamics observed from both models are, however, quite different quantitatively. Not only do the spatial heterogeneous consumers in model 2 survive three times as long as its homogeneous counterpart in model 1, moreover it evolves phenotypes that are twice as high as model 2. These differences are attributed to the emergence of spatial groups of consumers.

The analytical ecological models we have developed in chapter 4 assume a homogeneous distribution of an infinite population, and cannot be used to accurately model the global population dynamics of small stochastic populations of heterogeneously distributed individuals. They do, however, offer us a framework to get insight into the remarkable quantitative difference between models 1 and 2.

As we saw in section 5.1, the evolution of the phenotypes can be interpreted a change in the biomass conversion parameter c in the analytical model. The left panel in 5.5 shows the change of the conversion. In model 1, the evolutionary change could be interpreted a change in only this parameter. The spatial structuring of resources and consumer in model 2, however, requires us to take

several other parameters in the analytical model into account next to the conversion parameter c .

The heterogeneous spatial distribution of resources in model 2 causes the global carrying capacity K to be variable. The phase plots in figure 5.4 shows that the global carrying capacity realised in model 1, with approximate homogeneous consumer distribution, is lower than the others. This can be understood from the local growth of resources and consumers. Random placement of consumers in locations prevents the formation of resource clusters, because the resources are consumed before a cluster can be grown. When consumers are in groups, however, resources are depleted in some areas, but can flourish and form clusters where consumers are scarce. Therefore, the global carrying capacity is larger in the heterogeneous model 2 than in homogeneous model 1 (see fig. 5.5).

Another implication of the formation of consumer groups, is the increase of consumer interference e between consumers. The consumers in model 2 are expected to experience more interference from their neighbours. Moreover, the level of interference is indirectly related to the phenotype, since higher phenotypes causes more dense groups, in which the consumer interference increases even more. Since higher consumer interference e results increased ecological stability, the evolution towards instability (by increase of the conversion parameter c) is slightly inhibited by increased interference in situations with spatial heterogeneity (see fig. 5.5).

Although in both models, evolution causes an increase of the biomass conversion parameter c , the quantitative differences between model 1 and model 2 can be understood in terms of an increased global carrying capacity K due to resource heterogeneity, and increased consumer interference e due to consumers heterogeneity.

Groups

The evolutionary difference between model 2 and 3 as shown in the top panel in figure 5.4 cannot be attributed to such differences in population dynamics. In both experiments, offspring are placed close to other consumers which gives rise to (moving and viscous) groups of consumer individuals with the similar ecological consequences. To understand why evolutionary dynamics in one case leads to extinction while the other shows moderation and survival, a closer look at the group dynamics is required.

The consumer population organises itself into spatial clusters of consumer individuals, that we call groups. Groups are formed by reproduction of individuals, but persist over many generations of its members. Groups move around, divide and interact with each other, break down due to dispersion or go extinct. In a sense, they have a life cycle of their own.

The top panel of fig. 5.6 shows snapshots of the spatial distribution of the consumers population over the environment, with the colour representing the phenotypes relative to the others. In scenario 2 (left panel), the groups are

phenotypically mixed, while the individuals in the groups in scenario 3 (right panel) shows strong phenotypic similarity.

The short term spacetime plots in fig. 5.6 illustrates the movement of consumer groups over time. The spacetime plots presented here are created by recording the x-positions of consumers that inhabit a slice of the snapshot (with $25 < y < 75$) at a fixed interval.

Within-group selection

The difference between the two short term plots (taken over period 200000-210000) lies mainly in the phenotypic distribution. The groups in model 2 are phenotypically mixed. This is most clear in the leftmost region of the plot for model 2 in which groups are observed that contain red and blue phenotypes. These are, in fact, individuals of two distinct (quasi)species that have emerged just before (the only speciation event observed in model 2, see top panel in fig. 5.4). Since the individuals of these phenotypically distinct species compete with other directly within spatial groups, the blue (lower) phenotypes are quickly outcompeted by the red (higher) phenotypes, because the latter reproduce faster. The phenotypic differences within groups in the rest of the plot is smaller, i.e. there are no distinct species, because phenotypic variation is suppressed due to global mixing of phenotypes.

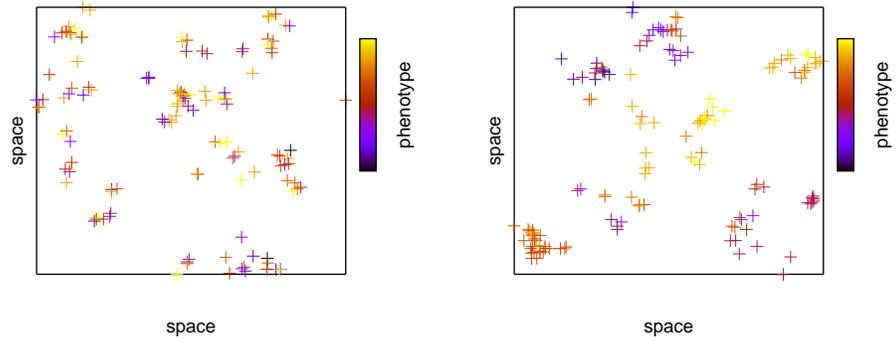
Differential reproduction in individual selection acts best if all individuals/phenotypes interact with each other, and are evaluated relative to all others in a well-mixed situation. Since individuals interact locally with others in their neighbourhood, the phenotypical composition of this neighbourhood influences the selective pressure. When the spatial distribution of phenotypes becomes structured into groups, individuals mostly interact with individuals that are phenotypic relatives, and therefore distinct phenotypes are sparsely evaluated against each other.

The short term spacetime plot for model 3 in fig. 5.6 shows groups with individuals that are phenotypically similar. Individuals in these groups mostly compete with similar individuals. Because the within-group variation is low, individual selection has little substrate and operates relatively slowly.

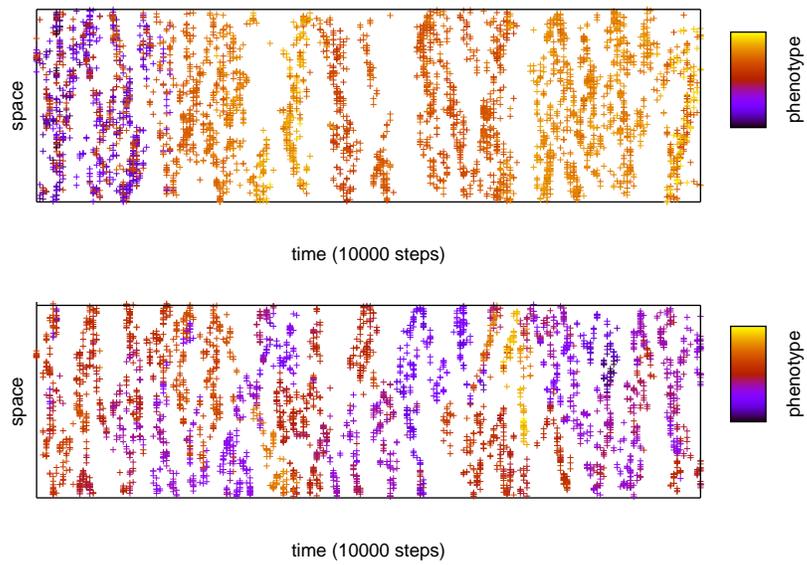
Because the individuals that make up a group interact and reproduce locally, phenotypes are not frequently evaluated against individuals of other groups. Therefore, a group with relatively low phenotypes is able to maintain itself, even when groups with much faster reproducing individuals are already present in the environment.

If individual selection pressures differ between homogeneous (model 1 and 2) and heterogeneous (model 3) spatial distribution of phenotypes, this is expected to be reflected in the observed population-level variation. Figure 5.7 shows the standard deviation of the phenotypes present in the population over evolutionary time for the three models (moving average of 1000 timesteps). Sudden increases in variation in model 2 and 3 indicate speciation events. The phenotypic variation in model 3 is indeed significantly higher than the other models over the entire course of evolution.

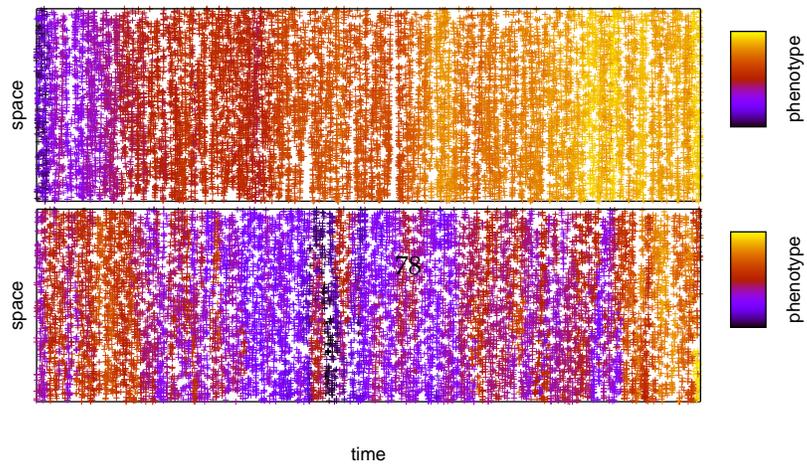
Snapshots of clusters in model 2 (left) and 3 (right)



Spacetime plot (short term); model 2 (top) and 3 (bottom)



Spacetime plots (long term); model 2 (top) and 3 (bottom)



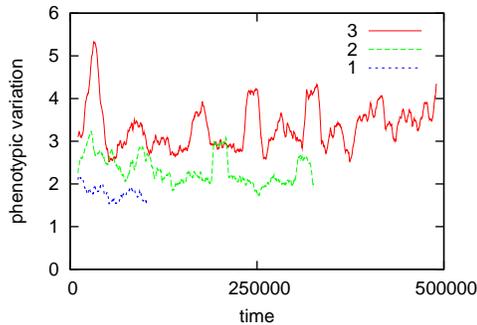


Figure 5.7: Phenotypic variation for three models

Between-group selection

The variation observed in model 3 in fig. 5.7 reflects the variation between groups, rather than the variation within groups. Selection within groups acts slowly due to the phenotypical similarity of their members. The relatively infrequent interaction between groups can cause large variation between groups as a result of the remaining within-group selection and genetic drift. This between-group variation is used as a substrate for selection.

We can regard a group as correlated emergent patterns in real space and phenotypic space that has properties not shared by the individuals they are made of (perhaps even their own life cycle). One such property of a group, seen as a subpopulation, is that it is involved in (local) population dynamics. The local population dynamics control the life cycle of consumer groups, in a way comparable as the energetic scheme controls the life cycle of individual consumers: group formation and growth, reproduction in abundance, and extinction when the group's population size declines.

The phenotypes of group members control the local population dynamics, since it controls the consumers rate of reproduction, in the same way as shown for global populations in section 5.1. The evolvability of phenotypes in populations and groups is constrained by their ecological viability.

Groups of individuals with low reproduction rates are generally more stable, but are outcompeted by faster expanding groups when they interact. When there is little interaction, however, such groups can persist. Groups of fast reproducing individuals run the risk of over-exploiting their local environment and fall victim to a local version of the "tragedy of the commons". Other, more stable groups that remain in the population occupy the vacant territory in real space and in phenotype space.

Individual and group selection

Group selection in model 3 uses the variation between the local population dynamical properties of spatial self-structured groups as a substrate. It acts by truncating the evolution of phenotypes through local “tragedies of the commons”. The direction of the group selection pressure is opposite to the direction of individual selection. While individual selection favours “selfish” consumers, group selection favours “altruistic” consumer that restrain their reproduction.

5.3.4 Conclusion

In this experiment, the impact of spatial self-structuring on ecological and evolutionary dynamics on a resource-consumer system was studied.

The formation of spatial groups of resources and consumers, as a result of local definition of carrying capacity and placement of offspring, causes improved ecological stability. This ecological difference between spatial homogeneous and heterogeneous distribution of resources and consumers (i.e. model 1 and 2) was understood in terms of the analytical model as an increased global carrying capacity of resources, and increased consumer interference in consumer groups. The evolutionary dynamics in both models is qualitatively similar in that they eventually result in the extinction of the consumer population. Evolution in the spatially heterogeneous ecosystem however differs quantitatively from the homogeneous system as it results in the evolution of much higher rates of reproduction in the consumer population before going extinct. This evolutionary difference is due to the increased ecological stability in spatially heterogeneous systems, which causes high reproductive rates to be ecologically viable and thus put a weaker constraint on the evolutionary dynamics.

Although consumer individuals in model 2 are spatially clustered, their phenotypes are distributed over these clusters. In model 3, consumers and their phenotypes are self-structured in groups. A comparison between the results of these two models show that spatial self-structuring has a major impact on the outcome of the selection process.

Since the phenotypic variation in groups in model 2 is greater than in model 3, individual selection operates more effectively in the former. Since there is little within-group variation in model 3, selection on the basis of individual differences in reproduction has little substrate. The remaining individual selection pressures and random genetic drift causes the between-group variation to grow. Initially, faster growing groups outcompete slower growing groups. Above a critical phenotypic level, groups that over-exploit their local environment become involved in a local “tragedy of the commons”. The between-group variation ensures that more stable groups exist that occupy the vacancy in real space and phenotype space, once the resource population has grown back. At a population level, this results on the evolution of reproductive restraint. Since restraining reproduction does not increase the number of

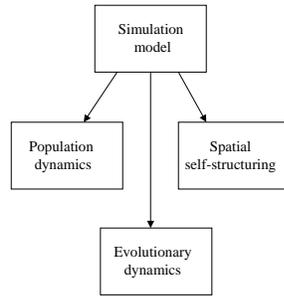


Figure 5.8: Observation of multiple emergent patterns

offspring of the restraining individual, but does increase the number of offspring of competing individuals, the evolutionary process in model 3 can be interpreted as the evolution of altruism.

5.4 Conclusions

The experiments in this chapter have demonstrated the use of eco-evolutionary computer simulations of a resource-consumer system as a means to investigate various biological evolutionary processes and their ecological constraints. The evolution of the amount of energy gained from eating resources resulted in the overexploitation of resources by increasingly “selfish” consumers which lead to their extinction in a process known as the “tragedy of the commons”. The coevolution experiment showed the emergence of an evolutionary arms race between predators and prey that settles in red queen dynamics by ecological constraints on both sides that prevented extinction of predators. The last experiments showed the ecological and evolutionary impact of spatial self-structuring formation of spatial subpopulations whose evolution of phenotypes is truncated through local “tragedy of the commons”. The emergent spatial groups serve as a substrate for group selection.

The explanatory methodology that was used in these eco-evolutionary experiments has remained implicit, and is briefly outlined here. In the eco-evolutionary experiments, the full explanatory framework developed in this thesis was employed. The emergent patterns that directly arise from the simulation (population dynamics and spatial structure) were analysed and explained by direct reference to the specifications of the model. In chapter 3 and 4, a direct explanatory model was used by exploring how the simple rules (behavioural, trophic, resource growth and placement) gave to the complex temporal and spatial patterns that were observed.

The observed evolutionary dynamics could, however, not be explained in a direct way. Since evolution in the simulated eco-evolutionary system is em-

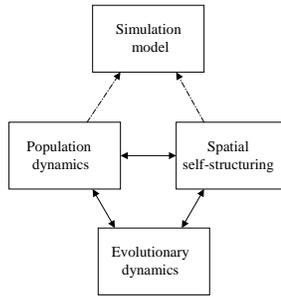


Figure 5.9: Indirect explanation of evolution dynamics in terms of population and spatial patterns

bedded in ecological interactions which constrain these dynamics, an indirect explanatory model was adopted. In section 5.1 and 5.2, evolution was understood in terms of the underlying population dynamics. In the explanation of group selection observed in the experiment in section 5.3 an even more elaborate indirect explanatory model has been used.

In the group selection experiment, multiple emergent patterns were observed (see fig. 5.8). In order to explain the evolutionary dynamics showing the evolution of reproductive restraint, an elaborate explanatory organization was constructed, in which the observation was explained on the basis of the interplay between the emergent patterns in evolutionary dynamics, local population dynamics and spatial structure (see fig. 5.9).

In short, spatial self-structuring causes the formation of groups that are subject to local population dynamics. Evolution uses these groups as substrate for selection by truncating groups that have unstable local population dynamics by overexploitation of their resources. The evolutionary dynamics are explained in terms of population dynamical and spatial self-structuring processes, without reference to the specifications of the simulation model itself.

Such indirect explanations are a powerful way of using computer simulations as tools for scientific enquiry. The basic specifications of the simulation model (local carrying capacity, local placement, chemotactic agents) are uncontroversial as biological processes and lead to the emergence of temporal and spatial self-organised patterns (population dynamics, patch formation) that are ubiquitous in nature. It is on the basis of these rather unsurprising and uncontroversial processes, that the explanation of the controversial process of group selection is based. It has been shown that these emergent patterns are involved in an interplay that arises 'for free'. In our experiment, this interplay was enough to explain the evolution of the altruistic act of restraining reproduction.

The explanation of group selection that is obtained in this way is relatively independent of the details of the simulation model, as it is completely based on

generic biological self-organised processes (albeit emerging from a simulation model). Therefore, such explanations are much less sensitive to the criticism that the outcome of the simulations is the product of the (often opaque) specific implementation of the simulation. Having overcome this major criticism towards artificial life models by indirect explanation, the use of complex simulation models as scientific tools for theoretical biological has become plausible.

Chapter 6

Conclusions

In this thesis, a versatile spatially explicit individual-based simulation model was presented that facilitates the study of the combination of situated interaction, self-organised collective behaviours and evolution by natural selection. It provides a platform for investigation of the interplay of processes at different levels of organization: behavioural interactions between individuals, spatial and temporal patterns on an ecological scale, and evolutionary dynamics of (co)evolving populations. The simulation model facilitates the construction of models that share the focus on individual embodiment and situatedness as in autonomous robotics, as well as the collective and evolutionary approaches in artificial life. The combination of the development of the virtual life simulation model, with providing it with a sound embedding in theoretical ecology and the use of direct and indirect explanations, constitute an important step towards closing the gap between models used in embodied cognitive science and theoretical biology.

This concluding chapter provides a brief discussion of the results of simulation experiments in terms of the main background concepts as presented in chapter 2 (section 6.1). For more elaborate conclusions about specific ecological and eco-evolutionary simulations, the reader is referred to the conclusion given above. After the summary, we proceed to draw general conclusions with respect to the methodological issues raised in section 1.3 about the scientific use of complex simulations (section 6.2).

6.1 Summary

We have seen that the virtual life simulation model incorporates the various concepts and themes that were examined in chapter 2.

Situated interaction

The simulated individuals can be simple static plant-like individuals, or mobile situated agents equipped with sensors, effectors, and a (neural) control system. The consumer individuals that are used in the experiments in this thesis, implement a simple control architecture similar to the Braitenberg vehicle (type 2b, see fig. 2.2). When these situated agents inhabit the same environment, and alter aspects of this environment, the behaviour of one individual influences the behaviour of others by a process of stigmergy from which spatiotemporal patterns can arise.

Self-organization

In the ecological simulations in chapters 3 and 4, the stigmergic and trophic interactions gave rise to self-organised temporal patterns in population dynamics and spatial self-structuring. The emergent patterns in population dynamics that were observed in simulations show that local interactions between simple individuals agents results in the formation of global patterns. This enabled us to capture the global dynamics in terms of an analytical model, which allowed us to predict and control emergent dynamics of the simulated ecosystem.

Evolution

By including variation in the population along inheritable and mutable traits, the experiments in chapter 5 have shown that eco-evolutionary simulations give rise to interesting patterns in population dynamics, spatial self-structuring and evolutionary dynamics. Evolution is not modelled as an optimization process (as in traditional genetic algorithms), but as a process of adaptation in which an evolving population adapts to the ecological challenges of survival and reproduction. This was sufficient to model natural selection of “selfish” consumer leading to a “tragedy of the commons”, evolutionary arms races and red queen effect in a coevolutionary setting, and group selection as resulting from the interplay between local population dynamics and spatial self-structuring.

It seems reasonable to conclude that the aims of the virtual life model, as stated in the introduction, has been satisfied. The virtual life model has been shown to be capable of simulating situated interaction, self-organized collective behaviour and evolution by natural selection. Moreover, the model has proven useful in the effort of the study of the interplay between emergent patterns that arise from the above mentioned processes.

6.2 Discussion

Theoretical embedding

The ability to construct complex simulation models does not, by itself, render such simulation useful as tools for scientific enquiry, however. To be able to interpret simulation results and use them in the scientific theory construction, the model needs to be embedded in, or related to, a well-understood theoretical framework. This is, of course, true for all simulation models, but is especially important for complex simulations such as the virtual life model, as been illustrated by the lack of such theoretical embedding of the PolyWorld simulator.

In the construction of the virtual life model, the theoretical embedding and tractability have been an integral part of the development. The theoretical framework has been provided by the well-studied models of theoretical ecology, as laid down in chapter 4. This has provided a useful framework from which simulations can be controlled and to which results can be compared and understood.

We have use these ecological models in a slightly odd way. The simulation model was used as a model for ecosystems, and the ODE models were used as analytical model for the simulation model. The latter is a more powerful formalism to study the interplay between generic biological processes than these classic state-variables models in many respects, due to its spatial and individual-based nature. It is more powerful in the sense that it allows the study of the interplay between many processes that have traditionally been studied separately and has previously been unavailable. Despite this fact, the simulation experiments and their analysis in this thesis has shown that the classical analytical models has not lost their value. They have served here to understand the emergent population dynamics in terms of density dependences, to control the population dynamics in the simulated ecosystems, and to understand evolutionary dynamics in terms of population dynamics. Moreover, analytical models have guided the design of the simulation model and the various experiments themselves. Most simulation experiments presented in this thesis replicated the results of existing theoretical work using (more) formal models (enrichment, tritrophic foodchain, evolution, coevolution, group selection), and is therefore itself part of this modelling approach.

But perhaps the most important way in which these theoretical model have been proven useful in analyzing simulation results is their use as baseline models to which the results of relaxing of their implicit assumptions (infinite population sizes, monomorphic and well-mixed populations) can be compared and understood. In a sense, simulation results only tend to get interesting where the applicability of ODE models breaks down and the simulation show qualitatively different behaviour. There is no way of telling what is interesting, however, without knowing what is not.

Indirect explanation

In chapter 1, we have identified the direct and indirect explanatory strategies as two ways to use artificial life models and interpret their results. Both have been applied in this thesis, but the focus was on the indirect way of interpreting the simulation results.

An example of direct explanation was provided for the regular patterns in the population dynamics that were observed in the simulations. These were interpreted as an emergent property of the stigmergic and trophic interactions between static resource and chemotactic consumer individuals. This statement relates the observed phenomenon directly to the specifications of the simulation model¹. This can be interesting from the point of view of behavioural ecology, since it enables the study of the relation between situated behaviour and the ecological processes it gives rise to. It is not unexpected, however, since it merely restates that simple local interactions can give rise to complex global patterns.

We have also seen that when phenotypic variation and spatial heterogeneity are allowed in the simulation model, the system gives rise to multiple emergent patterns: in population dynamics, in phenotype space, and in spatial distribution. In these cases, an explanation that directly relates the observed phenomenon to specifications of the model is often not sufficient, and sometimes not even necessary. The evolutionary arms race and red queen dynamics were related to population dynamics, and the evolution of reproductive restraint was related to spatial distribution. Here, the emergent patterns in evolutionary dynamics are understood in terms of patterns that are themselves emergent properties of the simulation model. Instead of focusing on the self-organizing processes themselves, we have studied the interplay between them, and used this as an explanatory model.

This strategy potentially renders theory formation based on simulation results more generic. Indirect explanations are based on an organizational structure between emergent higher level patterns and is therefore less dependent on the implementation details of the model. Explanations of processes obtained in this way are relatively independent of the details of the simulation model, as they are based on generic biological self-organised processes. Therefore, such explanations are much less sensitive to the criticism that the outcome of the simulations is the product of the (often opaque) specific implementation of the simulation. Indirect explanatory models enable the use of artificial life simulations as scientific tools in a fashion that transcends the 'simple to complex' paradigm by allowing the study of the 'complex to complex'.

The latter is essential in biological science to be able to gain insight in the interplay between processes at the different level of organization in biological systems (e.g. molecules, cells, organisms, ecosystems), and to explain the exist-

¹If one considers the stigmergic interactions itself as an emergent property of the spatial and chemotactic behaviour of consumers, one could argue that this is indirect explanation as well. Here, stigmergy is assumed to be the self-organizing process itself, instead of the property this process gives rise to.

tence of these multiple levels in the first place (e.g. the origin of life). We believe that it is essential for (embodied) cognitive science as well, since it studies a different part of the same multi-level spectrum, and also attempts to explain the emergence of major transitions (e.g. origin of sociality and cognition). Therefore, we feel that complex artificial life simulation, combined with a theoretical biological embedding and an indirect explanatory framework, is a new and valuable tool for cognitive science.

Bibliography

- A. Adamatzky and M. Komosinski. *Artificial life models in software*. Springer, 2005.
- R. Axelrod and W. Hamilton. The evolution of cooperation. *Science*, 211(4489): 1390–1396, 1981.
- J.R. Beddington. Mutual interference between parasites or predators and its effect on searching efficiency. *Journal of Animal Ecology*, 44:331–340, 1975.
- M.C. Boerlijst and P. Hogeweg. Self-structuring and selection: spiral waves as a substrate for evolution. In Langton, Taylor, Farmer, and Rasmussen, editors, *Artificial Life II*, pages 255–276. Addison-Wesley, 1991.
- E. Bonabeau, M. Dorigo, and G. Theraulaz. *Swarm intelligence: from natural to artificial systems*. Oxford University Press, 1999.
- V. Braitenberg. *Vehicles: experiments in synthetic psychology*. MIT Press, 1984.
- R. Brooks. Achieving artificial intelligence through building robots. *MIT Technical report*, 1985.
- S. Camazine, J. Deneubourg, N. Franks, J. Sneyd, G. Theraulaz, and E. Bonabeau. *Self-organization in biological systems*. Princeton University Press, 2001.
- L. Carroll. *Through the looking glass*. Mansfield and Wessels, New York, 1899.
- A. Channon. Towards the evolutionary emergence of increasingly complex advantageous behaviours. *International Journal of Systems Science, Special Issue*, 2000.
- W. Clancey. The frame of reference problem in the design of intelligent machines. In K. van Lehn, editor, *Architecture for Intelligence*, pages 357–424, 1991.
- I. Couzin and J. Krause. Self-organization and collective behavior in vertebrates. *Advances in the Study of Behavior*, 32:1–75, 2003.

- C. Darwin. *On the origins of species by means of natural selection, or the preservation of favoured species in the struggle for life*. J. Murray London, 1859.
- R. Dawkins. *The selfish gene*. Oxford University Press, 1976.
- R. Dawkins and D. Krebs. Arms races between and within species. *Proceedings of the Royal Society of London*, 205:489–511, 1979.
- R. de Boer. Modelling population dynamics: a graphical approach. Lecture notes, Theoretical Biology/Bioinformatics group, Utrecht University., 2006.
- D.L. DeAngelis, R.A. Goldstein, and R.V. O'Neill. A model for trophic interactions. *Ecology*, 56:881–892, 1975.
- J. Deneubourg, J. Pasteels, and J. Verhaege. Probabilistic behaviour in ants: a strategy of errors? *Journal of Theoretical Biology*, (105):259–271, 1983.
- E. di Paolo. *On the evolutionary and behavioral dynamics of social coordination: models and theoretical aspects*. PhD thesis, University of Sussex, 1999.
- H. Dreyfus. A phenomenology of skill acquisition as a basis for a Merleau-Pontian non-representationalist cognitive science.
- M. Eigen. Selforganization of Matter and the Evolution of Biological Macromolecules. *Naturwissenschaften*, 58:465–523, 1971.
- M. Eigen and P. Schuster. *The Hypercycle: A Principle of Natural Self-Organization*. Springer, Berlin, 1979.
- S. Forrest and M. Mitchell. Genetic algorithms and artificial life. *Artificial life*, 1 (3):267–289, 1994.
- C. Gershenson and F. Heylighen. When can we call a system self-organizing? In W. et al Banzhaf, editor, *ECAL-7*, pages 606–614. Springer, 2003.
- L.R. Ginzburg. Assuming reproduction to be a function of consumption raises doubts about some popular predator-prey models. *Journal of Animal Ecology*, 67:325–327, 1998.
- P. Grasse. La theorie de la stigmergie: essai d'interpretation du comportement des termites constructeurs. *Insectes Sociaux*, 6:41–81, 1959.
- V. Grimm. Ten years of individual-based modelling in ecology: what have we learned and what could we learn in the future? *Ecological modelling*, 115(2-3): 129–148, 1999.
- W. Hamilton. The genetical evolution of social behaviour (I and II). *Journal of Theoretical Biology*, 7:1–16, 17–52, 1964.
- G. Hardin. The Tragedy of the Commons. *Science*, 162(3859):1243–1248, 1968.

- I. Harvey, D. Husbands, A. Thompson, and N. Jakobi. Evolutionary robotics: the Sussex approach. *Robotics and Autonomous Systems*, 20:205–224, 1997.
- A. Hasting and T. Powell. Chaos in a three-species food chain. *Ecology*, 72: 896–903, 1991.
- F. Heylighen. The science of self-organization and adaptivity. The Encyclopedia of Life Support Systems, 1999.
- P. Hogeweg and B. Hesper. Individual-based modelling in ecology. *Mathematical and computer modelling*, 13(6):83–90, 1990.
- C.S. Holling. Some characteristics of simple types of predation and parasitism. *The Canadian Entomologist*, 91:385–398, 1959.
- C.S. Holling. Functional response of predators to prey density and its role in mimicry and population regulation. *Memoirs of Entomological Society of Canada*, 45:3–60, 1965.
- M. Hutson, D. DeAngelis, and W. Post. New computer models unify ecological theory. *BioScience*, 1988.
- A.R. Ives and V. Jansen. Complex dynamics in stochastic tritrophic models. *Ecology*, 79:1039–1052, 1998.
- C. Johnson and M. Boerlijst. Selection at the level of the community: the important of spatial structure. *Trends in Ecology and Evolution*, 17:83–90, 2002.
- M. Komosinski. The Framsticks system: versatile simulator of 3D agents and their evolution. *Kybernetes: International Journal of Systems & Cybernetics*, 32 (1/2):156–173, 2003.
- M. Komosinski and A. Rotaru-Varga. From directed to open-ended evolution in a complex simulation model. In *Artificial Life VII*, pages 293–299. MIT Press, 2000.
- M. Komosinski and S. Ulatowski. Framsticks: towards a simulation of a nature-like world, creature and evolution. In *ECAL-5*, pages 261–265, 1999.
- A. Lotka. *Elements of physical biology*. Williams & Wilkins Co., Baltimore, 1925.
- M. Maris and R. te Boekhorst. Exploiting physical constraints: heap formation through behavioural error in a group of robots. In *IROS96*, 1996.
- J. Maynard-Smith and E. Szathmary. *The major transitions in evolution*. Oxford University Press, 1995.
- C. Melhuish and O. Holland. Stigmergy, self-organisation, and sorting in collective robotics. *Artificial Life*, 5(2):173–202, 1999.
- F. Menczer and R. Belew. Latent energy environments. In *Adaptive individuals in evolving populations: models and algorithms*, pages 191–208, 1996.

- M. Merleau-Ponty. *Phenomenology of Perception*. 1943. translation by Colin Smith 1963.
- G. Miller. Artificial life as theoretical biology: How to do real science with computer simulations. In M. Boden, editor, *Philosophy of Artificial Life*. Oxford University Press, 1996.
- S. Nolfi and D. Floreano. *Evolutionary robotics: the biology, intelligence and technology of self-organizing machines*. MIT Press, 2000.
- R. Pfeifer and C. Scheier. *Understanding Intelligence*. MIT Press, 1999.
- T. Ray. An approach to the synthesis of life. In C. et al Langton, editor, *Artificial life II*, 1991.
- C. Reynolds. Flocks, herds, and schools: a distributed behavioral model. *Computer Graphics*, 21(4):25–34, 1987.
- O. Rössler. Adequate Locomotion Strategies for an Abstract Environment: A Relational Approach. In *Physics and Mathematics of the Nervous System*, pages 399–418. Springer, 1974.
- N. Savill and P. Hogeweg. Evolutionary stagnation due to pattern-pattern interactions in a co-evolutionary predator-prey model. *Artificial Life*, 3:81–100, 1997.
- H. Simon. *The sciences of the artificial*. MIT Press, 1969.
- M. van Baalen and D.A. Rand. The unit of selection in viscous population and the evolution of altruism. *Journal of Theoretical Biology*, 143:631–648, 1998.
- J.D. van der Laan and P. Hogeweg. Predator-prey coevolution: interactions across different timescales. *Proceedings of Royal Society London*, 259:35–42, 1995.
- L. van Valen. A new evolutionary law. *Evolutionary Theory*, 1:1–30, 1973.
- V. Volterra. Variazione e fluttuazione del numero d'individui in specie animali conviventi. *Memorie della R. Acc. Naz. dei Lincei*, 2:30–111, 1926.
- G.C. Williams. *Adaptation and natural selection: a critique of some current evolutionary thought*. Princeton University Press, 1966.
- D.S. Wilson. A theory of group selection. *PNAS*, 72(1):143–146, 1975.
- D.S. Wilson and E. Sober. Reintroducing group selection to the human behavioral sciences. *Behavioral and Brain Sciences*, 17(4):585–654, 1994.
- V. Wynne-Edwards. *Animal dispersion in relation to social behaviour*. Oliver & Boyd: Edinburgh, 1962.

- V. Wynne-Edwards. Intergroup selection in the evolution of social systems. *Nature*, 200:623–629, 1963.
- L. Yaeger. Computational genetics, physiology, metabolism, neural systems, learning, vision, and behavior or PolyWorld: life in a new context. In *Artificial life III*, 1994.

Software

The simulations, analysis, visualization of data, and text processing involved in the research and writing of this thesis has been facilitated by following software packages. All these programs are freely available.

Framsticks

Framsticks is a 3D artificial life simulator that models mechanical bodies that can be controlled by sensors, neural networks and actuators, and evolved by applying evolutionary algorithms. This simulation offers an intuitive visual interface and many tools for analysis. Framsticks is also available in command line interface, and a 'theater' application for demonstration purposes. Developed by Maciej Komosinski and Szymon Ulatowski, Poznan University of Technology, Poland. Website: www.frams.alife.pl

GRIND

GRIND is a powerful and simple system for analyzing systems of ordinary differential equations (ODEs) and maps. This program allows you to draw trajectories, timeplots, nullclines, compute eigen vectors, and do simple bifurcation analyses. It is driven by a simple command language. Developed by Rob de Boer, Theoretical Biology/Bioinformatics group, Utrecht University, the Netherlands. Website: theory.bio.uu.nl/~rdb/software.html

Gnuplot

Gnuplot is a command driven interactive data and function plotting utility. It allows scientists and students to visualize mathematical functions and data. Website: www.gnuplot.info

Lyx

Lyx is a WYSIWYM ("what you see is what you mean") word processor that allows you to use the lay-out and type-setting features of Latex, in a familiar word processing environment. Website: www.lyx.org