Structural coupling with environment and its modelling on neural driven agents

Master thesis

Matěj Hoffmann

Structural coupling with environment and its modelling on neural driven agents

Matěj Hoffmann

Over the past decades, H. Maturana and F. Varela have tried to characterize the interaction of organisms with environment, inventing the structural coupling concept. At the same time, many researchers in AI have appreciated the importance of embodiment for intelligence. In this work, I follow the idea that these two concepts have a lot in common and develop a relational model of this phenomenon. A simulated neural driven agent with plastic synapses is co-evolved with an environment, which contains an abiotic and a biotic part. The agent and environment mutually perturb each other's structure, on an evolutionary and individual time scale, and develop a structural fit. This congruence of their structures can be observed as emergent behaviour, visualised in a 3D scene. The model is interpreted, and partly formalized from a dynamical systems perspective. This allowed to quantify structural coupling and to draw interesting consequences concerning the symmetry of agent-environment interaction and the role of different time scales in it.

Dedication

I would like to thank the supervisor, Ivan M. Havel, for choosing this interesting topic and for help with its elaboration. I would also like to thank other people that have helped me with writing this thesis, especially Jan Romportl and Jan Burian.

Table of Contents

List of Figures

List of Tables

List of Equations

Preface

What is life? What is intelligence? How can we characterize them and what are the necessary conditions for them? How to approach evolution of organisms and how to describe their interaction with environment? What is the difference between natural and artificial systems?

It is very hard to answer these questions, but people cannot help themselves trying to answer them once again. These phenomena are addressed in everyday situations by ordinary people as well as by many scientists in various disciplines. Biology, philosophy, psychology, artificial intelligence are only a few among the disciplines that struggle with life or intelligence.

A Chilean biologist, Humberto Maturana, in the 1960's was unhappy with biological characterizations of life, which rested on enumerating features of living systems (e.g. organization, metabolism, growth, adaptation, response to stimuli, reproduction). He wanted to capture an invariant feature in the organization of living systems. With his student, Francesco Varela, they have found this invariant in a process that they named *autopoiesis* (self-creation). Life can be ascribed to autopoietic systems, systems that are able to maintain their organization over time. Later, the existence of such systems in their environment was captured by another construct: *structural coupling with the environment*.

This set of constructs, also referred to as *Autopoietic theory*, provides a framework on which the holy grails of life or intelligence may be characterized in terms of processes and relationships. It becomes irrelevant, whether a system is virtual or physical, whether it is from metal or cells. What matters are the relationships. Hence this theory may provide a common ground for scientists in various disciplines to understand themselves when dealing with these difficult phenomena.

These concepts are suitable for a branch of science that uses a synthetic methodology, or tries to understand things by building them. Making a model of a system, that becomes coupled to its environment may give us more insight into the nature of the phenomenon than observing organisms. Such a constructive approach may provide understanding of nature (when an organism is modelled) or, on the other hand, give clues for construction of artificial systems that will come closer to the magnificent abilities of organisms in their environments. However, in this work, I will stay in the middle – with the help of a computer model, I will try to understand the general principles of the interaction (and history of interaction) of a system with its environment. A virtual model in a computer is much more powerful than a flowchart on paper and the fact that it is only virtual is justified by the process-like nature of the phenomena modelled. We shall see whether it will help us to reveal some of the fundamentals of life or intelligence.

Chapter 1. Introduction

Let me start with an explicit formulation of my goals. I think that the concepts embodiment and structural coupling are fundamentally important for understanding intelligence or adaptive behaviour of animals and at the same time are of great importance when developing artificial systems. However, I find the definition and understanding of them insufficient. This is also because it is a transdisciplinary issue. In order to deepen this understanding I try to bring together many disciplines and use a synthetic methodology to develop a model. The road to this model (choice of level of modelling, architecture) and the interpretation of it reveal a lot about the nature of the modelled phenomena. With the help of a dynamical systems formalization, I will be able to show a way how to quantify the interaction of an organism (or agent) with its environment and to draw some interesting consequences regarding symmetry of this interaction and the time scales involved.

In this chapter I will first explain how has Artificial Intelligence arrived at the notion of embodiment, a key feature of intelligent systems. Then the work of Maturana and Varela will be introduced, arriving at structural coupling. I think that this context is relevant. Nevertheless, if the reader is familiar with the history of AI and its arrival at embodiment, or with the ideas of Maturana and Varela, these parts (Section 1.1 and Section 1.2) can be omitted. Then these concepts will be related and it will be argued that they have the same essence. Finally, it will be explained why it is worth modelling this phenomenon and how it can be done. The former section (Section 1.3) is a theoretical formulation of the phenomena that will be modelled. The latter, Section 1.4, answers the question why is such a model useful, and presents an overview of how it will be developed. The qualities we want from the model are followed by a note on methodology and a formulation of the level of modelling. The structure of the thesis will be presented in the last section of this Chapter.

1.1. AI on the road to understand intelligence

I will briefly outline, how difficult it is to define or even develop (artificial) intelligence. We will have a quick tour through the 50-year history of the field of Artificial Intelligence (AI) and see how it has arrived at the notion of embodiment and its importance.[1]

1.1.1. What is intelligence

In everyday situations, we often talk about intelligence. However, if we were asked to define it, we would hardly end up with same or even similar definitions. At least, we would be able to mention some capabilities an intelligent being should possess – these may include: think-

[1] Sections 1.1 and 1.2 benefit from Pfeifer, Scheier (2001), Chapter 1.

ing and problem solving, learning and memory, language, intuition and creativity, consciousness, emotions, surviving in a complex world, perceptual and motor abilities. In 1921 the Journal of Educational Psychology (Vol. 12, p. 123-147, 195-216) asked fourteen leading experts in the field to provide their definitions of intelligence. Not surprisingly, they got 14 different answers back, for instance:

- The ability to carry on abstract thinking. (L. M. Terman)
- The ability to adapt oneself adequately to relatively new situations in life. (R. Pintner)
- Having learned or ability to learn to adjust oneself to the environment. (S. S. Colvin)
- A biological mechanism by which the effects of a complexity of stimuli are brought together and given a somewhat unified effect in behaviour. (J. Peterson)

Although all experts were from one field, they have used very different levels for their characterization – Terman talking of abstract thinking, whereas Peterson of biological mechanisms. Apropos, are animals intelligent? Some experts also mentioned the environment in their definitions, some did not.

Nevertheless, Pfeifer and Scheier are able to find a common denominator. The underlying theme seems to be "coming up with something new." However, we would not consider something new, that is not advantageous, or even fatal, intelligent. The innovation has to be advantageous in the situation, where the system finds itself. “Both components, conforming and generating are always present. The key point is generation of diversity while complying with the givens. We call this the *diversity-compliance trade-off*.” (Pfeifer, Scheier 2001, p. 20-21) Using different words only, the mechanism enabling organisms to cope with environmental changes – adaptation – contains two components: complying with existing rules and generating new behaviour. In this sense, *intelligence corresponds to adaptive behaviour*. And intelligence in this sense will be investigated in this thesis.

1.1.2. Understanding by building

Empirical sciences, like biology, neuroscience or psychology, on their road to understand intelligence typically proceed by observing the behaviour of an animal, a human, a part of the brain etc. in normal situations or (more often) in carefully designed experiments, record the results and then analyse them. The *analytic approach* is applied. A model predicting the behaviour of an animal in a particular situation may be developed, but this would be a statistical one, as it is only the 'output' (e.g. number of correctly recognized objects or time for finding a way out of a maze) what matters.

Synthetic approach is different. In this approach, we are interested also in the internal mechanisms and try to develop an artificial system that reproduces certain aspects of a natural system. Similar outputs of the artificial and natural system are also important but the focus is shifted to how the results were brought about. Such an investigation of a natural system is called *synthetic modelling*. Interesting insights can also be gained by developing artificial systems that do not model a particular natural system; their design may, however, be inspired by mechanisms or processes that we suspect are present in interestingly behaving animals. This field of research may be described as *understanding* general principles of intelligence (life, adaptivity, behaviour) *by building*.[2] With the progress of computers, the synthetic methodology is becoming more and more powerful. If a remarkable system is developed, it may also have (industrial) applications (see Figure 1.1).

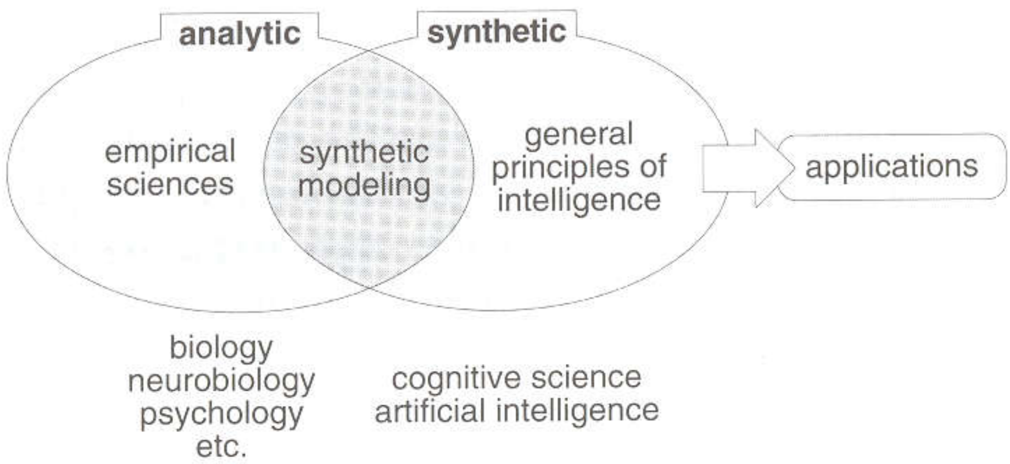

Figure 1.1. Overview of approaches to the study of intelligence (from Pfeifer, Scheier 2001, p. 22, with kind permission).

1.1.3. Cognitivism

The synthetic approach was adopted by the emerging field of Artificial Intelligence in 1950's. The aspects of intelligence that were modelled were the high-level cognitive functions, such as comprehension, inferencing, decision-making, planning, learning or generalization – similar to the commonsense notion of intelligent abilities. The *computer* was not only the *tool* on which these models could run, but also quickly became the leading *metaphor for mind*. The essence of this paradigm is that the key to intelligence is computation with symbols that rep-

[2] Havel (2001, p. 25) catches this difference as accurate (or authentic) vs. metaphorical modelling.

resent the world. The keywords are *algorithmic nature, symbolic computation* and *representation*.

Computation is independent of matter, of the physical substrate; what matters is only the *function* of individual components of a mechanism, not their physical instantiation. To think thus de facto means to run a correct program. In theory (that is what Strong AI claims), human mind can be realized on a silicon chip or with sticks and strings. It would suffice if we adequately understood how brain operates with symbols like "hunger" or "pain". Such a program could then be run on arbitrary hardware. The cognitivist paradigm is also known as *functionalism*, *computer functionalism* or *GOFAI* (Good Old-Fashioned Artificial Intelligence). (Hoffmann 2006a, p. 265)

Case study – chess

Artificial systems developed under this paradigm proved to be very efficient in fulfilling some tasks. They are very successful in limited domains with clear rules. Such a domain is for instance chess-playing – a touchstone for AI. Chess-playing at grandmaster level is considered a typical manifestation of intelligence. And GOFAI has achieved a major victory here, in 1997 IBM's Deep Blue has defeated the world's number one player – Garry Kasparov. How is this possible? This is because chess is a suitable domain for a computer. Chess has a comparatively simple discrete state space – the states being board positions at every move. Every state is specified by positions of the chessmen on the board. From every position, there is a finite set of transitions to a subsequent position – legal moves, which are defined by the rules of the game. The solution lies in calculating all the possible sequences of moves from the current to a final position and choosing the move leading to the best result (quickest victory). However, this leads to a combinatorial explosion. Human players thus investigate only a couple of sensible moves, count only a few moves ahead and rely on experience and generalization in their evaluation of moves. How about the computer? The search space (of all possible subsequent positions) is too large even for it and thus it has to limit it by some heuristics. Even though these are sophisticated, we can still say that when compared to the human player, the computer substitutes intelligence by huge computational power (calculating thousands of positions every second).[3]

But how did Deep Blue actually play against Kasparov? The man at the right on Figure 1.2 is not Deep Blue. We would not value moving the physical chessmen over the chessboard, or detecting the move of the opponent as anything intelligent. However, this would be more difficult for Deep Blue than playing the formal game.

[3] Although grand masters are able to calculate more moves ahead than hobbyists, it is not this ability what constitutes their strength.

Figure 1.2. Garry Kasparov playing against IBM's Deep Blue in 1997.

1.1.4. Problems of cognitivism

If we took Terman's definition of intelligence – ability to carry on abstract thinking – from the section What is intelligence, we could well say that GOFAI (or cognitivism) was quite successful in developing artificial intelligent systems. Deep Blue could be an example. It was this kind of problems and this kind of intelligence the traditional AI focused on.

However, let us examine the goal of *adaptive behaviour in the real world*, a task all organisms on Earth face. The end of the previous section has provided us with an introduction that acting in the *real world* is different than in the formal world of chess. Deep Blue has no body and copes with the abstract world of chess only; it needs a human who stands in for it in the real world.

Real world vs. formal world

Let me use the world of chess as a simple example of a formal world and confront it with the real world outside. What are the differences? The real one is much more complex as it is:

- *Inaccessible*: Contrary to the world of chess, a robot in the real world cannot obtain accurate up-to-date information about the whole environment. It is limited by its sensors and their operation.

- *Non-deterministic:* The outcome of some actions is not uniquely defined. For instance, a robot may decide to pick up an object. However, it cannot be sure whether it will succeed.
- *Dynamic*: The environment changes in ways beyond the robot's control.
- *Continuous in time/space.*
- *Real-time*: The robot does not have arbitrary time to make a decision.

Cognitivist robot

Once GOFAI has succeeded in formal worlds such as world of chess, let the real world be a new test bed for this way of creating intelligence. We are in the real world now and our artificial system thus needs a **body** to be able to act and perceive. This gives rise to a *cognitivist robot*.

However, cognitivist robots find it very hard to operate in such environments. Let us look at a typical cognitivist robot[4] and investigate why is it so. At the heart of the robot is a central procedure that operates on an internal representation of the world. On the basis of this representation, the central procedure decides which action to take, in order to fulfill a goal from a list. As we are in the real world, the robot gets a designed body with sensors and effectors (motors) to be able to act and perceive. These connect the symbolic core to the world. Then, this robot operates in a sense-think-act cycle: reads input from its sensors; updates the internal representation; based on this representation, the central procedure selects an action, which is performed by the effectors.

Problems of cognitivist robots[5]

Practical problems

- *Robustness and generalization*: The problem stems from the preprogrammed nature of the central procedure. This is tolerant neither to faults nor to noise (unless this was explicitly programmed, but noise that was not programmed will always exist in the real world). The robot is also unable to generalize, to cope with new situations, which the designer did not foresee.

[4] This is an 'idealized' model that enables to show the problems of this design.
[5] This section is adapted from Pfeifer, Scheier (2001), Chapter 3.

- *Central sequential processing vs. real-time world*: The world outside has its own pace, own dynamics and we have to keep up with it. This is very hard for a central sequential procedure if the data flow is large. On the contrary, the brain uses distributed parallel processing.

Fundamental problems

- *Frame problem*: This problem is concerned with keeping the internal representation of the world consistent with the real world outside. The internal world model will consist of a set of logical propositions. However, in the outside world, any of these can change at any point in time. If you imagine a large internal model and a dynamic environment, it becomes very hard to detect which changes in the outside are relevant for the internal model.

- *Symbol-grounding problem*: In classical AI systems (e.g. expert systems), the designer of the system defines symbols (in a purely syntactic way). Symbols relate to other symbols and are processed by a procedure. They do not relate to the outside world. It is the user giving meaning to the operation of such systems. A cognitivist robot can be viewed as such a symbolic system put into a body. The world model and control procedure are symbolic. This means that the programmer has introduced a set of symbols, such as `food` or `poison` and a set of rules, e.g. if `see(food) then eat(food)`. The programmer knows what is food, but the robot does not. The key to grounding these symbols from the robot's perspective is interaction with the real world (through sensors and motors): grounding them in own experience. However, it is very hard for the robot to map sensory stimulation to the symbols, unless this mapping is also provided by the programmer. It is then the programmer or observer, but not the robot himself, who is grounding the symbols in the world.

1.1.5. New AI

Let me repeat that we are aiming at the goal of creating an artificial system that will be able to survive in a complex world – we will then attribute intelligence to such a system. Animals do not have problems surviving in the wild but robots have. But why? It might be that the way GOFAI designs robots is not the best one. A new approach has emerged that is called *New AI*, *behaviour based AI* or *embodied cognitive science*.

Rodney Brooks in his lab at MIT in the 1980's was trying to develop autonomous mobile robots. He realized first, that the traditional AI's focus for more than 30 years was on abstract, i.e. disembodied systems, and second, that the methods developed under the cognitivist paradigm are not suitable for designing autonomous agents acting in the real world.

Let us start from the traditional AI point of view. There is one thing about GOFAI that was not mentioned explicitly here – it tries to construct intelligence using a *top-down method*. The symbolic representations and processes we have encountered in the previous sections very often correspond to higher cognitive functions (we may have modules like learner, planner, perception module etc.), which we know from psychology or introspection. However Brooks (1991) argued that higher level thought may be viewed rather as a tip of an iceberg. It rests on a very complex foundation that enables us to cope with the world outside. This base is also present in other animals and it has taken evolution a longer time to develop this base than our highest cognitive functions. And hence we should look at simple organisms for inspiration and start rather from *bottom-up*. (p. 1 and 16)

The key to the new approach lies in **situatedness** and **embodiment**. Let me introduce these ideas in a structure similar to the original paper, *Intelligence Without Reason* (Brooks 1991).

Situatedness

The problem with traditional AI systems is that they are problem-solvers, solving problems in a symbolic abstract domain – they are not situated in the real world. Even the cognitivist robots do not deal with the real world – they 'live' in their world model and update it, struggling with the Frame problem.

To illustrate that this might not be our normal way of dealing with the world, let me quote Margaret Boden, who has noticed this point already in 1977:

> In everyday life you usually remember your "place" largely because the external world is there to remind you what you have or haven't done. For instance, you can check up on whether you have already added the vanilla essence by sniffing or tasting the mixture, or perhaps by referring to the pencil and paper representation of the culinary task that you have drawn up for this mnemonic purpose. A computational system that solves its problems "in its head" rather than perceiving and acting in the real world, or pencil and paper models of it, has to have all its memory aids in the form of internal representations. (Boden, M. (1977) in *Artificial Intelligence and Natural Man*, 1977, p. 373; cited from Pfeifer, Scheier 2001, p. 71)

The same applies for mobile robots. Brooks showed that an agent can avoid obstacles in a much simpler way if it continuously refers to its sensors and uses simple reactive rules (such as `if (left front sensor active) then turn right`) than by updating a model of the world and planning paths around the objects in this model. This essence is in the famous idea: “*The world is its own best model.*” (Brooks 1991, p. 15)

Embodiment

The fundamental consideration is whether we can have a *disembodied mind*. Quoting from Brooks' (1991) passage on this: “Many believe that what is human about us is very directly related to our physical experiences. For instance Johnson[6] argues that a large amount of our language is actually metaphorically related to our physical connections to the world. Our mental 'concepts' are based on physically experienced exemplars. (p. 15)” The consequences of being a brain in a vat are also analysed by Wiedermann (2006).

Such a disembodied system simply lacks grounding: the real world that would give meaning to its processing. We have already met the symbol-grounding problem, but note that a non-symbolic system, such as neural network, will lack this grounding in a similar way. We can conclude that *intelligence* (at least the intelligence we are seeking) *requires a body*. Once we have en embodied system, we have the advantage that *it can be really tested* whether it is able to deal with a complex world outside: by putting it in the world (Brooks 1991, p. 15).

You may wonder what is the relationship of situatedness and embodiment. Embodiment certainly affects robot's situatedness: if we change the body, for instance some sensors, this will change what the robot is dealing with (Wiedermann 2006, p. 426). Also Pfeifer, Scheier (2001, p. 72) note that embodiment does not necessarily imply situatedness. This would be the case of a cognitivist robot, which has a body but relies on a world model that was given to it by the designer. On the other hand, it is not possible to be situated in a world without having a body. Thus situatedness implies embodiment.

At last it is necessary to emphasize that by embodiment we mean a *body interacting with the environment*. An embodied agent is part of the world dynamics, affecting the environment and being affected by it. This will be analysed further in the next section.

Intelligence and emergence

The importance of the environment can be illustrated by Simon's example. An ant on a beach scattered with pebbles follows a crooked path. The complicated path does not reflect the ant's internal complexity, but rather is a result of simple reactions of the ant and a complex terrain – a case of obstacle avoidance we have already discussed in the section called “Situatedness”. Brooks (1991, p.16) concludes: “*Intelligence is determined by the dynamics of interaction with the world*.”[7] The (intelligent) behaviour we see on the beach is a result of an interplay of the ant and the terrain, the environment. We may say, that this behaviour *emerges* from the agent-environment interaction and so does intelligence, if we define it as intelligent or adapt-

[6] Johnson, M. (1987): The Body in the Mind. University of Chicago Press, Chicago, IL.

[7] Brooks quotes Simon, Herbert A. (1969): The Sciences of the Artificial. MIT Press, Cambridge, MA.

ive behaviour.[8] Note that with preprogrammed robots, this is a different case. The behaviour is already programmed and there is no (or little) room for emergence.

Autonomous agents research and remedies to cognitivist robots' problems

The New AI thus stays with the synthetic methodology but uses a different approach. It works from bottom-up to create robots (or agents) that are situated and embodied in their environments and whose intelligent behaviour emerges from the interaction with the world. The research concentrates on *artificial autonomous agents*. Let us quote a definition by Beer (1995a):

> Autonomous agent: any embodied system designed to satisfy internal or external goals by its own actions while in continuous long-term interaction with the environment in which it is situated... encompassing at the very least all animals and autonomous robots. (p. 173)

In this thesis, it also encompasses virtual agents that have a simulated virtual body and interact with a virtual environment. I use robot when the agent is physical and deals with the real world outside. Let us look, how the agents built under New AI solve the problems of cognitivist robots.

Practical problems

- *Robustness and generalization*: A solution to these problems was already offered by connectionism. Neural networks as controllers of an agent have such properties that favour tolerance to noise and generalization. The New AI approach, working from bottom-up and employing artificial evolution, creates robust solutions of agents in their environments.
- *Central sequential processing vs. real-time world*: This problem is solved by favouring distributed parallel processing instead.

Fundamental problems

- *Frame problem*: This problem is greatly reduced by situatedness – relying on the world

[8] First, behaviour and intelligence in this sense, as emerging from agent-environment interaction, is used for instance in by Pfeifer, Scheier (2001). Brooks is talking about behaviour and intelligence in the agent. Second, this is only one possible use of the term emergence. For three ways of using this term see e.g., Pfeifer, Scheier (2001, p.125).

outside rather than on a large internal world model.

- *Symbol-grounding problem*: This is the most difficult problem. However, the New AI systems, if designed carefully, do not have to suffer from it. Let us quote an example:

> ... the Distributed Adaptive Control architecture is embedded in a physical robot. The sensory stimulation, and thus the input to the system, are the results of physical processes occurring in the robot's interaction with its environment. The robot's own movements have a strong effect on the sensory stimulation to which it is exposed. The outputs correspond to the robot's movements. In this sense, the outputs have meaning directly to the robot: They affect its behaviour. The robot selects the meaningful patterns from its own perspective. Meaningful patterns are the ones that cause the robot to change its behaviour; examples are activation patterns that get the robot to turn. Note that in order to learn something new the robot has to move – it has to interact physically with its environment. Physical processes are required, not only informational ones... In other words, if we want to make progress in resolving the fundamental issues, we have to move beyond pure information processing. (Pfeifer, Scheier 2001, p. 174)

Does embodiment have to be physical?

From what we have met so far, does an agent have to be really *physically* embodied to allow for the study of intelligence? Brooks or Pfeifer would suggest: yes, referring to the fact that "intelligence cannot merely exist in the form of an abstract algorithm but requires a physical instantiation, a body (Pfeifer, Scheier 2001, Glossary)." This is certainly safer, but sometimes not feasible, such as in artificial evolution. However, Pfeifer also agrees that simulated physical worlds, like for instance the one used by Karl Sims (1994) and the virtual creatures, are relevant for our study.

1.2. Maturana, Varela: life and interaction of organisms with environment[9]

In the 1960's Chilean biologist Humberto Maturana was unhappy with the characterization of living systems and their cognition. In psychology of that time, after a period of behaviourist dominance, cognitive psychology was experiencing a rebirth. Cognitive psychology is con-

[9]This section was developed mainly with the help of Quick (2006), Whitaker (2001) and Varela et al. (1991). Citations from original works that are not from Varela et al. are cited from Whitaker.

cerned with high-level mental processes such as memory, attention, perception, action, problem solving or mental imagery. However, Maturana thought that cognition is a biological phenomenon.[10]

Maturana did not think that we can characterize life or cognition from outside, by inventing some functional descriptions and abstractions. Instead, he has founded, and later developed together with Francesco Varela, a systemic framework that would characterize living systems, their cognition and also their phenomenology in terms of the organism itself. This body of work was originally labeled biology of cognition[11], but later became better known as *Autopoietic theory*. The core of this theory is the process of *autopoiesis*.

1.2.1. Autopoiesis

Autopoiesis is a process that should answer the question of what is the characteristic organization of living (or cognitive) systems.[12] Let us have look at some definitions of autopoiesis.

> In cybernetics, a term coined by Humberto Maturana for a special case of homeostasis in which the critical variable of the system that is held constant is that system's own organisation. (Fontana Dictionary of Modern Thought)

> An autopoietic machine is a machine organised (defined as a unity) as a network of processes of production (transformation and destruction) of components that produces the components which: (i) through their interactions and transformations continuously regenerate and realise the network of processes (relations) that produced them; and (ii) constitute it (the machine) as a concrete unity in the space in which they (the components) exist by specifying the topological domain of its realisation as such a network. (Maturana H. R., Varela F. J. (1980): Autopoiesis and Cognition. D. Reidel, Dordrecht, Holland, p. 78-79.)

With inspiration from cybernetics living systems are characterized as living machines. The word autopoiesis comes from Greek (*auto* (self) and *poiesis* (creation)). The system keeps on producing or creating itself. The key example is a biological cell. To rephrase the definition

[10]To make things clear, the work of Maturana and Varela is presented here separately from the developments in AI. In fact, the two fields may have influenced each other. As we see, both were influenced by cognitivism and also by cybernetics.

[11]In 1970 Maturana has published two papers: 1. The neurophysiology of cognition. In Garvin, P. (ed.): *Cognition: A Multiple View* (Spartan Books, New York, 1970), 3-24. 2. Biology of Cognition. Biological Computer Laboratory Research Report BCL 9.0., Urbana IL: Univ. of Illinois, 1970.

[12]"Living systems are cognitive systems, and living as a process is a process of cognition." (Maturana 1970, p. 13)

once again: an autopoietic system is composed of parts and these have relationships between them (in the form of processes). This interaction is such that it keeps maintaining and regenerating the components, together with the relationships. The physical components may be replaced in the course of time, but the *organization* remains constant.

Autopoietic systems are living systems; they are self-producing their components. There is also a more general class of systems in Autopoietic theory: autonomous systems. These maintain their organization, but do not have to produce their constituent components. This more general class of systems can be attained by real artificial systems and it suffices for structural coupling. Varela defines these systems by the attribute of organizational closure, a term from cybernetics.[13] Since I already talk about autonomous agents and will talk about autonomous dynamical systems, for this notion, I will use *organizationally closed systems*.

Organization and structure

Structure is the physical instantiation of a system – physical components engaged in processes. The components may be replaced by new ones (as are components of a cell), but the network of processes of production keeps running – *organization* remains constant. The organization is such that the structure is constantly being regenerated. In other words, organization is an abstract network of processes, and structure its instantiation. In the following I will talk about structure only, but organization is also involved, like two sides of the same coin.

1.2.2. Structural coupling

In the previous section on autopoiesis, we were concerned with the organization and structure of a living system alone. However, living systems live in interaction with their environments. Not only their internal dynamics, but also their contacts with the environment and their course of change in interaction with it are determined by their structure. We talk about *structural determination*.

A central explanatory construct that addresses interaction of two systems (or unities) is **structural coupling**. Let us have a definition by Maturana first:

> In general, when two or more plastic dynamic systems interact recursively under conditions in which their identities are maintained, the process of structural coupling takes place as a process of reciprocal selection of congruent paths of structural changes in the interacting systems which result in the continuous selection in them of congruent dynamics of state. (Maturana H. R., Guiloff G. D. (1980): The quest for the intelligence of intelligence. *Journal of Social and Biological Structures* **3 (1980)**, p. 139.)

[13] For a more thorough explanation see Varela F. J. (1979): Principles of Biological Autonomy. Elsevier (North Holland), New York.

A plastic system is one that can be affected by outside events; it can be perturbed by another system. This structural coupling definition characterizes the interaction of two systems in general. In this thesis, we will be interested in a special case where one system will be an organism (or agent) and the other will be its environment: **structural coupling with the environment**. Maturana[14]offers this definition: "If one of the plastic systems is an organism and the other its medium, the result is ontogenic adaptation of the organism to its medium: the changes of state of the organism correspond to the change of state of the medium."(p. 326) Thus only the organism has a plastic structure. However, I think that this is a simplification and in my model, I will work with a plastic environment. Thus, I will approach the more general and more interesting case of two structurally coupled systems.

We have seen what is a structurally determined system. Its structure determines the interactions it can engage in. Looking at an organism as a plastic system, engagement in interactions with environment changes its structure; it is perturbed by the environment, and hence the possible future interactions with the environment also change. When the environment is a plastic system, its structure also continually changes, during the interaction with the organism. Over the course of time, the system (organism) and environment are engaged in mutual non-destructive perturbations and shape each other; a congruence between the two systems arises. They get structurally coupled to each other. In the following, we will sometimes use structural coupling only, implicitly meaning with environment. If this was not the case, it will be explicitly noted.

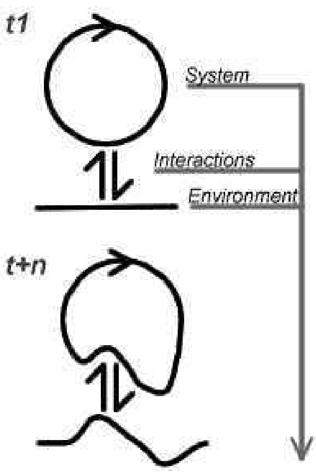

Figure 1.3. Structural coupling of system with environment (from Quick 2006).

The result of such a process, or *history of structural coupling*, are two organizationally closed systems, organism and environment, that are however shaped by the mutual interaction, as

[14]Maturana H. R. (1975): The organization of the living: A theory of the living organization. *International Journal of Man-Machine Studies* **7 (1975)**, 313-332.

Figure 1.3 shows. The word organizationally closed was used deliberately, as the two systems engaged in coupling do not have to be autopoietic. Such a congruence can occur between any two structurally determined plastic systems living (which would be also autopoietic) or not living, for instance some artificial systems, which may be only organizationally closed.

The instantaneous interaction of such systems will seem like "'*coordination*' in the sense that the coupled systems will (to an observer) exhibit some measure of correspondence or correlation in the manner and the course of their behaviours during coupling." (Whitaker 2001, Tutorial 2, section Structural coupling; emphasis added) "To an observer" refers to the fact that Maturana and Varela are careful to characterize the system within its own bounds. The behaviour is not in the system, but appears only in the eye of an *observer*. On a longer time scale, the coupling will (again, to an observer) seem like '*co-evolution*'. The systems get tuned to each other. On the part of an organism, this occurs on the time scale of an individual and also (indirectly) on the time scale of a species.

1.2.3. Conclusion

The aim of Maturana and Varela was to answer the question of cognition in a different way than cognitivism. The answer was already suggested by Maturana (1970) in the *Biology of Cognition*:

> A cognitive system is a system whose organization defines a domain of interactions in which it can act with relevance to the maintenance of itself, and the process of cognition is the actual (inductive) acting or behaving in this domain. Living systems are cognitive systems, and living as a process is a process of cognition. This statement is valid for all organisms, with and without a nervous system. (p. 13)

Varela nevertheless later develops another concept: *enaction*. This is a more phenomenological concept – addressing an organism in the first person. *Enactive cognitive science* has even become a new paradigm. Varela et al. (1991) explain the difference between cognitivism, connectionism and the enactive viewpoint. An organism is viewed not as an input/output machine, but rather as a network consisting of multiple levels of interconnected, sensorimotor subnetworks. A history of sensorimotor patterns gives rise to cognitive structures of an organism. And cognition is enaction: A history of structural coupling that brings forth a world. (p. 207) Note that this is not an objective world – every organism enacts a world of its own.

The whole autopoietic theory is a very inspirational way of looking at living systems. The constructs and language used are new and therefore it is not easy to immerse oneself into this theory. On the other hand, this can also be an advantage, as this theory does not (or not as much) carry the burden of existing fields and the constructs of this theory are concerned with systems, their dynamics and relationships only and thus it can be used in a number of disciplines. In this thesis, I will concentrate on the phenomenon of structural coupling of a system with its environment.

1.3. The common ground: embodiment = structural coupling?

In the previous sections I have briefly introduced two fields and their efforts. The first was artificial intelligence; more precisely I concentrated on its 'core' that tries to understand the principles of intelligence. After the cognitivist era, many researchers have arrived at the idea that bodies and interaction with environment are crucial for intelligence: in animals as well as in robots. Embodiment became a buzzword.

H. Maturana and F. Varela were concerned with living systems and their cognition. They were also unhappy with the cognitivist notion of this phenomenon. They have invented autopoiesis to characterize living systems and structural coupling to catch interactivity among them and with the environment. It is their structure what determines the possible interactions of a system (its cognitive domain) and over time, the structures of two plastic systems are changing, influenced by each other, entering a congruence, or getting structurally coupled.

Let me identify a number of parallels between the work in New AI and the ideas of Maturana and Varela:

- Researchers in AI as well as Maturana and Varela wanted to understand life and intelligence and thought that traditional empirical science is insufficient. AI has come with a synthetic methodology; Maturana and Varela analysed living systems as unities, defining them “systemically in terms of their processual configuration” (Whitaker 2001, Tutorial 1).
- Behaviour is not in a system. Maturana and Varela claim that behaviour is in the eyes of an observer watching system's interaction with environment. Researchers in AI also recognize behaviour as emerging from the interaction.
- Necessity to understand life and intelligence in a different manner than as information processing.
- Importance of the body of living or intelligent systems and their intimate relationship with environment.

It seems that both approaches have a lot in common and hence it might be fruitful to try to 'merge' them. In New AI, embodiment has become a key concept. However, it lacks a clear definition. The last parallel (and the previous text) suggest that embodiment and structural coupling with the environment are in fact very similar concepts. Hence, in AI we could use structural coupling to define and characterize embodiment. I think that this effort, as proposed in Quick et al. (1999), is very fruitful and I will demonstrate in this thesis.

1.3.1. On the notion of embodiment

In the section called "Embodiment" and the section called "Does embodiment have to be physical?" we have seen that New AI has arrived at the fundamental importance of embodiment and at the same time at a somewhat intuitive definition of it, which more or less rests on physical embodiment.

However, one has to ask what is so special on the physical, or material, embodiment? Quick et al. (1999) ask this question and note that in behavioural robotics, the fact that robots are physically embodied is exploited, however, the question of what it means to be embodied is not answered. Moreover, in evolutionary robotics, a similar setup is used first in simulation and then transferred on a real robot. Philosophers (such as Dreyfus) also think that there is nothing inherently material on embodiment (Quick et al. 1999, p.2-5).

Ziemke (2003) reviews the notions of embodiment and mentions six of them. He mentions *physical embodiment* and a restricted version of it, which he calls *organismoid embodiment*. *Social embodiment* addresses social interactions. Nevertheless, the other three notions stem from the work of Maturana and Varela! *Organismic embodiment* is based on autopoiesis and limits embodiment to living systems. The last two are the most interesting to us – they are relational definitions of embodiment based on structural coupling: *Embodiment as structural coupling* and *historical embodiment*, reflecting the history of structural coupling.

1.3.2. Relational embodiment definition

We are seeking a definition of embodiment that would catch what is going on and that would not be limited to a particular material instantiation. Structural coupling will be our relational definition of embodiment.

The definition of embodiment by Quick et al. (1999) "describes a minimal state of affairs whereby structural coupling is made possible" (p. 6). Let me quote their characterization of structural coupling:

> Structural coupling is a process that occurs when two structurally plastic systems (an organism and its environment, for example) repeatedly perturb one another's structure (their constituent components and the relationships between them) in a non-destructive fashion over a period of time. This leads to the development of structural 'fit' between the systems. There is an intimate relation-

> ship between this process and the emergence of 'appropriate' behaviour from the interplay between interacting systems, because the structure of a system determines its responses to perturbatory environmental events. (p. 6)

And a *minimal definition of embodiment* derived from that:

> A system X is embodied in an environment E if perturbatory channels exist between the two. That is, X is embodied in E for every time t at which both X and E exist, some subset of E's possible states have the capacity to perturb X's state, and some subset of X's possible states have the capacity to perturb E's state. (Quick et al., p. 7)

Note first that this definition involves a mutual perturbation and second that it is really minimalist: Quick et al. (1999) also show that an abandoned car rusting in a desert would also satisfy it (p. 8). I offer a different *example*, the notorious thermostat with a heater in a room. The room's states are different temperature levels and these are perturbed by the action of the heater (and heat losses to the outside). The thermostat measures the temperature and falls into one of the states heater on / heater off. It is thus also perturbed. This trivial situation thus fits both, embodiment and structural coupling. The structural fit is a certain level of temperature in the room. This minimalism is a problem, since we are not able to differentiate between 'interesting' and 'uninteresting' structural coupling. However, a quantification of structural coupling (as discussed further and applied on my model) will allow us to assess its degree.

A big positive of the relational approach is that it would fit physical agents (robots), where there was the physical notion of embodiment, as well as software agents, where there was a lack of embodiment definition.

> Software systems with no body in the usual physical sense can be intelligent. But they must be embodied in the situated sense of being autonomous agents structurally coupled with their environment. (Franklin, S. (1997): Autonomous Agents as Embodied AI. *Cybernetics and Systems* **25(8)**, 499-520; cited from Ziemke 2003, p. 1306)

Moreover, this definition can also be applied in many other disciplines, not just AI, as already structural coupling is, as Whitaker (2001) documents.

1.4. Modelling structural coupling

The concept of structural coupling has already been introduced. In this section I will deal with its modelling. First I will try to argue that this process is worth modelling and second, the details of how should it be modelled will be presented.

1.4.1. Why modelling it?

As already stated in Section 1.1.2, AI uses a synthetic methodology: trying to understand things by building or modelling them. We have also seen that embodiment has been identified by New AI researchers as a crucial phenomenon underlying intelligence. In the previous section, we have seen that embodiment can be defined in terms of structural coupling with the environment, a term of Maturana and Varela, which was identified by them as a construct explaining the interactivity of (living) systems with environment. And hence modelling structural coupling can be very productive on the road to understand intelligence and also a contribution to the effort of creating artificial systems.

There is also another, pragmatic, reason, why to model it. We have discussed that it is a relational construct, not tied to a particular field or physical instantiation. Hence, it can be modelled, even in software, and the lessons learned should be helpful to a wider interdisciplinary audience.

The synthetic approach may follow three aims: a) modelling natural systems, b) modelling general principles, and c) developing applications. This work concentrates on b), modelling structural coupling with the environment in general.

1.4.2. How to model?

We have already justified the choice of modelling in software, hence a *virtual model* is a good choice. Then, to model structural coupling, we need a system and its environment in interaction. Both, system and environment, should be plastic, i.e. their structure (constituents and corresponding relationships between them) should have the possibility to change over time. And the source of change (perturbation) of a system has to, at least partly, be in the environment. Similarly, the environment should also be perturbed by the system. This last point is not present in the original definition of structural coupling with the environment by Maturana, and typically also not present in experiments with autonomous agents. In my model, I will allow plasticity to the environment, so that the environment will actively add to the development of structural congruence with the agent.

Qualities

What qualities should our model have?

- *Compliance with principles*: It should comply with the principles of the phenomenon modelled.

- *Minimalism*: No unnecessary constructs should be introduced into the model.
- *Heuristic value*: The model should further our understanding. It should reveal some new, interesting consequences.
- *Quantifiability*: It would be nice to quantify the results (see next section).
- *Performance*: It should reliably produce some remarkable, visible and preferably also quantifiable results, which would increase the heuristic value.

There will be trade-offs between these, for instance between minimalism and performance.

Quantification

The features that could be quantified were proposed by Quick et al. (1999, p. 8), however, their application has lead me to an extension (italics):

- The size of system's and environment's structure. How complex are the interacting systems, or if viewed as dynamical systems, how large is their state space *and how complex are the maps or dynamical laws and how many system parameters are there* (corresponding to their structure).
- Plasticity of system and environment. To what extent (and how quickly) can they change.
- Degree of mutual effect of system and environment on each other. To what extent they cause changes in the other's structure. Or bandwidth of perturbatory channels.

I will also make a distinction between time scales, especially individual and evolutionary one.

Assumptions and methodology

I will not make a distinction between natural and artificial systems and assume that modelling of structural coupling is relevant to both. The systems modelled will not be autopoietic, but organizationally closed (see end of Section 1.2.1). I will not concentrate on phenomenological aspects. I will take care of the relation between observer, designer (or modeler), the artifact, the environment, and the observed agent – what is known as the Frame-of-reference problem (Pfeifer, Scheier 2001, p. 112). Note that this is something different than the Frame problem.

Level of modelling and platform

The notion of structural coupling is very general and many interacting systems, natural or artificial, simple or complex, could be described as engaged in structural coupling. Even trivial examples would fit the definition. In nature, structural coupling with the environment takes place in all organisms: from the simplest, unicellular, like bacteria, to the most complex, vertebrates. As analysed further, the structure of biological organisms is very complex and comprises many levels (e.g. processes within cells, among cells, in the nervous system; possibly all these interacting with the environment). For my model, I have to choose a level[15], the basic level of analogy. Nevertheless, recall that natural organisms are our inspiration but we will not model a particular one. Hence we are more free in our choice, not limited by the relations existing in a particular organism on a particular level.

In the choice of the level in this thesis, I was lead by the qualities I have set forth. We need a model that would reveal some interesting consequences and allow for a fruitful interpretation. Thus, we do not want a banal model, such as a thermostat with heater in a room. The examples from nature are much more interesting. Hence, our model, drawing inspiration from nature, could be best characterised as *co-evolution* what I will identify as a special case of structural coupling.

In my model, I will have an agent with a simulated physical body and a neural controller, interacting with a simulated physical 3D world, under simplified laws of physics. Modelled parts of the body will be sensory and motor surfaces, not metabolism. The neural controller will connect the sensors and effectors. Environment will consist of an abiotic part, a landscape, and a biotic part – food items. The body and 'brain' of the agent as well as the biotic part of the environment are co-evolved and will get coupled to each other in the course of time. The neural controller of the agent will also acquire plasticity in an individual agent.

The level of modelling also encompasses *time scales* that we will trace. An individual organism gets coupled to environment during growth (ontogenesis) and also in maturity: structural changes in interaction with the environment occur which affect the body and the nervous system (if it has one). However, there are also very important structural changes on a longer time scale: evolutionary. These are indirect, as they do not directly perturb a particular organism. Indirectly, evolution selects individuals with better genotypes ('starting structure') and thus a species gets coupled to the environment (phylogenesis). On this time scale, the environment also evolves and gets coupled to the agents living in it. The model will help us to clarify the issue of time scales. I will address structural changes in an individual and in a species. However, the genetic encoding used will not be a developmental one, as in nature. I will use a direct mapping from genotype to a mature phenotype. This is a simplification, for a model in-

[15]This is a simplification. Havel (2001, p. 46) accents that in living organisms there is a hierarchy of levels and these may not be distinct and it may also be difficult to say which level is above which.

cluding this see Quick et al. (2003), for example. On the environmental side, plasticity will involve mainly the evolutionary time scale.

The choice of level determines also choice of a modelling platform. Cellular automata (Bittorio) were used by Varela et al. (1991) to explain basic principles of structural coupling. Quick et al. (2003) modelled structural coupling on Genetic Regulatory Networks, a computational model based on a biological mechanism. These models and the corresponding platforms are more low-level. I have chosen a higher level: neural driven agents with plastic morphology. Therefore, I need a different platform: a software environment simulating agents with physical bodies and neural controllers, interacting with a (simulated) 3D physical world. This allows for modelling of sensorimotor patterns and the performance should be more visually attractive, hence enabling for easier explanations of the phenomenon. The software environment chosen is Framsticks.

1.5. Structure of the thesis

In the following chapter, we will look at the problem from a biological point of view. The terms adaptation and co-evolution will be related to structural coupling. Then, structural coupling of some organisms, bacterium E. Coli and a cnidarian Hydra, with their environment will be analysed. Finally we will look at what constitutes the structure of organisms. In Chapter 3 I will prepare ground for the model by choosing a level of modelling. Then I will analyse the suitability of neural architectures and how to model a body, environment, how to define a task and which time scales to consider. Chapter 4 is the model itself. The software environment and model architecture are presented, followed by experiments and their results. In Chapter 5, the experiments are interpreted from a structural coupling perspective and some interesting consequence are drawn. The experiments are evaluated. This chapter is the core of the thesis and contains most of the contribution of my work. Chapter 6 contains analysis of related work, Chapter 7 future development. In the last chapter, Conclusion, the contribution of this thesis is summarized.

Chapter 2. Biological considerations

Before I start with my model of an agent that is structurally coupled with its environment, let us look at natural agents for inspiration. It was primarily the principle of existence and operation of cells and organisms what Maturana and Varela were trying to explain. In the beginning, I will relate the term structural coupling to adaptation and co-evolution, terms used in biology. Then, we will look at structural coupling in organisms. First, we will observe the behaviour of a bacterium (E. coli) and analyse, how the behaviour emerges from its structure interacting with the environment. The second organism subject to analysis will be Hydra, one of the simplest organisms with a nervous system. Finally, we will reveal what actually constitutes the structure of these organisms.

2.1. Adaptation, co-evolution and structural coupling

A key term in our area of interest is *adaptation*. Let us look what is usually meant by adaptation in biology:

> A biological adaptation is an anatomical structure, physiological process or behavioral trait of an organism that has evolved over a period of time by the process of natural selection such that it increases the expected long-term reproductive success of the organism.[1]

This differs from what we mean by structural coupling in the following ways. First, structural coupling is a process, an ongoing history of coupling; the agent and environment have been interacting with one another and changing each other and will keep on doing this. This contrasts with adaptation as a result of this process. *Continual adaptation* is a better term depict the nature of structural coupling.

Second, adaptation concentrates on the agent (animal). The definition we quoted does not even mention environment. On the other hand, structural coupling is more symmetrical: both sides are engaged in this reciprocal interaction. Hence I prefer the term *mutual adaptation*.

Third, while adaptation is usually connected with evolution and natural selection, i.e. long-term changes on the time scale of species (phylogenetic time scale), structural coupling may comprise many time scales. Non-destructive structural perturbations can occur between an individual organism in an environment or, indirectly, also on the species' time scale. Those in-

[1] Adaptation (2006, July 20). In: *Wikipedia, The Free Encyclopedia*. Retrieved 16:57, August 4, 2006, from http://en.wikipedia.org/w/index.php?title=Adaptation&oldid=64875278.

dividuals achieving a 'better' congruence with the environment may live longer and produce more offspring and their genotypes expand (natural selection).[2]

In this thesis, I will differentiate between structural coupling on two basic time scales: First, the individual time scale. I will use the terms *structural coupling on individual time scale*, structural changes in an individual, or sometimes individual adaptation. Adaptation on individual time scale may include ontogenesis (origin and development of an organism from the fertilized egg to its mature form) and any further adaptation in a mature individual. In biological language it may include for instance acclimatization, acclimation. The term learning will be reserved for learning in the cognitive sense – acquiring new competences and accumulating knowledge. This is, however, exhibited by very complex organisms whose structural coupling is too complex and beyond the scope of this paper. And second, the time scale of species, or evolutionary – *structural coupling on evolutionary time scale*.

Hence, if we intended to use adaptation to describe structural coupling, we could speak of *continual mutual adaptation of agent and environment on evolutionary and individual time scales*. However, this does not capture the essence of the phenomenon – how the adaptation comes about.

A different term, which I want to relate to structural coupling, is *co-evolution*. This is a special type of evolution. 'Basic' evolution involves a species that over millions of years adapts to its *abiotic environment*, which stays more or less constant on this time scale. However, in co-evolution, there is also a *biotic environment*. Flegr (2005, p. 326) identifies two aspects in which it differs. First, the biotic part of the environment changes (or changes much more quickly than the abiotic part) and hence the organism has to deal with a changing environment. Second, the biotic environment actively reacts, by means of its own evolution, on the evolutionary changes in the organism. It thus may lead to an arms race or to a symbiosis. Therefore, co-evolution is a process involving mutual structural changes in organisms and their environment. Thus, I would say, that *co-evolution* is a natural *example of structural coupling* and I will use it as an inspiration for my model. However, note that co-evolution does not necessarily involve plasticity on individual time scale: the individuals do not have to be very adaptive. It may rest mainly on *indirect evolutionary perturbations*.

2.2. Structural coupling in nature

Before we start modelling, let us look at some natural organisms structurally coupled with

[2] Varela et al. (1991) offer rather *natural drift* than natural selection. This means that evolution does not select optimal solutions or is not directed to a particular goal, but rather any viable solutions randomly spread.

their environment. This is relevant to our understanding of the phenomenon. The analysis of a bacterium, E. coli, and a cnidarian, Hydra, from the structural coupling perspective will show us that structural coupling is extremely complex, even in the simplest organisms. We will also get a better picture of what constitutes the structure of organisms. This will be analysed in the last section and will help us to create an artificial structure.

2.2.1. Bacteria

Bacteria are unicellular organisms and (thus) do not have a nervous system. But still they are attracted and repelled by certain stimuli, for instance light (phototaxis) or chemicals in their environment (chemotaxis). They move in their environment, accepting and excreting substances by diffusion and try to get to the place where good substances prevail. How do they achieve this adaptive behaviour?

Case study of E. coli[3]

Structural dynamics of the bacterium *Escherichia Coli* is analysed by Quick et al. (1999). The behaviour observed when we watch this bacterium in its environment is *chemotaxis*. E. coli is sensitive to chemoattractants; propelled by flagella it is able to swim smoothly up a concentration gradient of attraction. When it finds itself in a uniform concentration, it switches to 'tumbling'.

How does this observed behaviour emerge from the dynamics of its structure, in interaction with environment? There are sensory (receptors) and motor (mainly flagella) surfaces in the bacterium, connected via a signalling pathway – a complex dynamical system. Basic elements interact mainly by means of the transfer of phosphoryl groups. These are produced within the cell and their presence at flagella supports tumbling. Meanwhile when attractants bind to receptors, the production of phosphoryl is inhibited, encouraging smooth swimming. From this the basic behaviour emerges. However, the dynamics is more complicated. First, there are also methyl groups involved, which act as *feedback loops*, producing a homeostatic effect – so that the balance of tumbling and swimming returns to a base level, whatever the concentrations of attractants may be. Second, there is also receptor clustering involved (see Quick et al. 1999 for details). All these processes are involved in the mutual perturbation of bacterium and environment, giving rise to the observed behaviour.

Evolutionary time scale

What I would like to add to this analysis, is the evolutionary time scale. Evolution of bacteria takes place via genes on a single chromosome and their mutation. Bacteria are very old and thus a long history of structural coupling with the environment on this time scale has oc-

[3]This section is adapted from Quick et al. (1999), p. 9-11.

curred. This is (indirect) structural coupling as co-evolution. A remarkable demonstration of the effect of this is the diversity of this group of organisms. For instance, the genetic distance between two bacteria species, E. coli and Thermus aquaticus, is greater than the distance between humans and oak trees! Every pair, bacterium species and its environment, has achieved a very different structural congruence.

2.2.2. Organisms with simple nervous systems

We have seen in the previous section that sensorimotor activity in bacteria was achieved by a dynamic coupling of sensory and motoric surfaces of the cell membrane. This activity is present in multi-cellular organism too. However, such organism can afford a cellular specialization. Neural systems have evolved in multicellular organisms as an efficient method of communication between cells, to co-ordinate sensorimotor activity as well as inner processes in the organism. Maturana and Varela[4] have characterized the role of the nervous system in structural coupling as follows:

> ...the neurons, the organism they integrate, and the environment in which they interact operate reciprocally as selectors of their corresponding structural changes and are coupled with each other structurally: the functioning organism, including its nervous system, selects the structural changes that permit it to continue operating, or it disintegrates. (p. 170-171)

Let us look at the phylum cnidarians containing the simplest organisms with nervous systems; we will examine one of the simplest members: hydra. A brief look at its nervous system and behaviour may give us inspiration for the artificial model and also give us a clue what can we possibly expect.

Cnidarians and diffuse nerve net

This phylum contains aquatic multicellular organisms with radial symmetry. Hydra, jellyfish or anemones belong to this phylum. These are the simplest animals with a nervous system. They possess a *diffuse nerve net* (or more nerve nets); neurons are distributed (and connected) throughout their body without a centre (no cephalic ganglion or 'archaic brain'). These organisms exhibit internally generated rhythmic behaviour and complex co-ordinated sensorimotor behaviours (for instance the feeding behaviour of Hydra) in their environments.

[4]Maturana H. R., Varela F. J. (1987): The Tree of Knowledge: The Biological Roots of Human Understanding. Shambhala, Boston. Cited from Iglowitz J.: http://www.foothill.net/%7Ejerryi/ChapterOne.pdf

Hydra

Hydra belongs to the more primitive members of this phylum. It is a freshwater predatory animal with a tubular body. There is a foot which provides attachment to a place on one side of the body, and a mouth, surrounded by 5 to 12 tentacles, on the other (Figure 2.1). Its body consists of two layers: an outer one (epidermis) and an inner one (gastrodermis) with cells secreting digestive fluids. Hydra contains simple muscles, and photoreceptors and touch sensitive nerve cells (in body wall and tentacles). They are connected by specialised nerve cells.

Figure 2.1. Hydra. One of the simplest organisms with neurons, in the form of a diffuse nerve net.

Hydra's feeding behaviour is complex:

> When feeding, Hydras extend their body to maximum length and then slowly extend their tentacles. Despite their simple construction, the tentacles of hydra are extraordinarily extensible and can be 4 – 5 times the length of the body. Once fully extended, the tentacles are slowly manoeuvred around waiting for a suitable prey animal to touch a tentacle. Once contact has been made, nematocysts on the tentacle fire into the prey and the tentacle itself coils around the prey. Within 30 seconds, most of the remaining tentacles have already joined in the attack to subdue the struggling prey. Within 2 minutes, the tentacles will surround the prey and move it into the opened mouth aperture. Within 10 minutes, the prey will be enclosed within the gastrovascular cavity and digestion will have started. The hydra is able to stretch its body wall considerably in order to digest prey more than twice its size. After two or three days, the indigestible remains of the prey will be discharged by muscular contraction through the mouth aperture.[5]

As in the case of E. coli we shall ask: how does this behaviour emerge? Simplifying, we can say that there is a sensory surface of the touch sensitive cells and a motor surface, consisting of muscle and secretory cells. The co-ordination, which was in E. coli facilitated by a signalling pathway within one cell, is here brought about by the nerve net. Johnson and Rohrer (2006, p.7) say that neurons with elongated membranes “can extend over the length of the entire organism before terminating in the muscle cells. These tail-like cellular projections are the axons, and evolutionarily speaking they are the flagella of the multi-cellular organism.” The structure has become more complex: apart from the dynamics within one cell, we have many cells and these perturb each other directly (the neighbouring cells) and also via the neural communication (or perturbatory) channel. Moreover, the nerve net becomes a new player – a new dynamical system with its own structure, coupled with the other cells of the organism and the environment. The behaviour observed emerges from the interaction of all these systems within the organism among themselves and with the environment. Again, the structural coupling comprises both individual and evolutionary time scale.

More complex organisms

We have seen in the previous sections that the structure (and organization) of the simplest organisms is already very complicated and so is their coupling with the environment. For instance Prescott (2005) provides an insight on the evolution of substrate for action selection. He illustrates, how the nervous systems get more complicated, from jellyfish like *Aurelia aurita*, a more complex cnidarian with two nerve nets influencing each other in a natural example of subsumption architecture, over flatworms with archaic brains, to vertebrate nervous systems with 'central switching devices' in the form of basal ganglia.

2.2.3. What constitutes the 'structure'?

Structure, the substrate for intelligence or adaptive behaviour, encompasses all the constituent components of an organism and relationships between them. The structural dynamics of an organism consists in internal processes and interaction with the environment via sensory and motor surfaces. A unicellular organism, such as bacteria, can engage in complex structural coupling with its environment, through the structure of its single cell. Moreover, this is a living system and it is not only organizationally closed, but also autopoietic. The autopoiesis of individual cells, which are referred to by Maturana and Varela as 'first-order autopoietic units', lies also in the core of autopoiesis of multicellular organisms. An organism is as 'second-order autopoietic unit', emerging from the solidarity of individual cells.

Neural systems facilitate efficient perturbatory channels between cells, helping in coordina-

[5] Hydra (genus) (2006, June 7). In: *Wikipedia, The Free Encyclopedia*. Retrieved 12:29, July 4, 2006, from `http://en.wikipedia.org/ w/index.php?title= Hydra_% 28genus% 29&oldid= 57273874`.

tion of internal processes as well as in sensorimotor coordination. Nevertheless, Maturana (1970) reveals that the working of a nervous system is different in the following respect:

> The nervous system enlarges the domain of interactions of the organism by making its internal states also modifiable in a relevant manner by 'pure relations', not only by physical events. (p. 13)

The neural system may be viewed as an organizationally closed system of its own, the rest referring to as body. Mingers (1990) talking about humans (but this applies generally) says:

> As humans, we are autopoietic systems with a plastic nervous system that is organizationally closed. This is structurally coupled to our body and, through this, to the environment. Both the body and the nervous system are structure-determined systems – the changes which they undergo depend on their own prior structure and can be only *triggered*, not determined, by interactions with other systems. (p. 570)

Nevertheless, in living systems "...the function of the nervous system is subservient to the necessary circularity of the living organization (Maturana 1970, p. 13)." This is because it could not exist on its own; it is not autopoietic and does not regenerate its components.

Chapter 3. Level of modelling and design choices

I have already uncovered the nature of structural coupling and explained why this phenomenon is worth modelling. In this chapter, I will choose which aspects will be modelled in this work and also analyse possible architectures of the model. First, the level of modelling will be chosen and reasons for this choice will be presented. On the chosen level of analogy, the modelled systems will be an agent's neural controller, body, and the environment, and the time scales individual and evolutionary. Second, I will comment on how these can be modelled and discuss the suitability of selected neural architectures, artificial bodies and environments. The last sections analyse what is the task or desired behaviour in this model and how can we acquire a fitness function.

3.1. Level of modelling

The previous section has showed how complex is the structure of natural organisms. In physical terms, the structure and its plasticity comprise a number of levels, from microscopic to macroscopic. However, we are interested in processes and relations, not in the physical instantiation. In the relational domain, we have seen a complex network of processes in a bacterium. Moreover, in multi-cellular organisms with nervous systems, this comprises networks of processes within cells, among cells, plus interactions with the nervous system – all that interacting with the environment. In principle, any of these levels can be modelled. Recall that in this thesis we are interested in modelling general principles of structural coupling and we are thus not bound to a particular organism. The following section describes the level of analogy chosen here.

3.1.1. Neural driven agent with a plastic body

In this work I will use the discrimination of a body, which will encompass all the processes in the agent (switching over from organism to agent when talking about the model), and the neural network (nervous system in organisms). Together, viewed as one system, this will be the structure of the agent.

This model of structural coupling is comparatively high-level. It concentrates on sensorimotor activities of an agent, interacting with environment. The modelled component of the body are hence sensory and motor surfaces, which interact through a neural network. The body will be modelled in a physical world, but I will abstract from all metabolic processes (and from cells living and interacting in the body), apart from simple energy calculation. Both the body and the neural network will have plasticity, to allow for structural coupling.

I would also like to mention the difference between natural and artificial neural networks. Already in the nerve nets of cnidarians, there are multi functional neurons, action potentials, synapses, chemical transmission. The operation of a biological neural network is very complex. Artificial neural networks (ANNs) are very simplified models that model only some of the features. Similarly to the discussion on the body physiology, most of the biological processes will not be modelled and the choice will be made that would enable a good structural coupling demonstration. In the following, by neural networks we will understand artificial neural networks.

3.1.2. Individual and evolutionary time scale

It has been already discussed that structural coupling of an animal with its environment involves the coupling of the individual animal, and, indirectly, of a whole species. We would like to model both, in a close interaction. Structural coupling on *individual time scale*, structural changes in an individual, or sometimes individual adaptation may include (on the agent's side) ontogenesis (by this I will understand origin and development of an organism from the fertilized egg to its mature form) and any other adaptation in a mature individual. Structural changes in an individual from birth to maturity, in interaction with the environment, constitute an important part of the phenomenon modelled. This requires a developmental genotype encoding to be used, which is rare in agent experiments. Unfortunately, neither this model explores it (for reasons see Section 4.1.4 in the next Chapter). For a use of a developmental encoding, even on structural coupling modelling, see Quick et al. (2003). We will thus concentrate on structural changes in a mature agent – individual, while it interacts with the environment.

To model 'agent phylogenesis', we can use a genetic algorithm. The agent's neural controller as well as body will be subject to artificial evolution. Note, that the changes on the evolutionary time scale do not affect the structure of a particular individual: they are indirect. Rather, the genetic algorithm will, in an imitation of natural selection, prefer some genotypes – recipes for the initial structure of phenotypes. The 'starting structure' of a whole agent population is plastic and gets coupled to an environment. In our model, we want the coupling to be mutual. Therefore, the environment should evolve together with the agent. On this time scale we will evolve part of the environment that we call biotic environment.

3.1.3. Support for chosen level

I present the following comments on the choice of modelling level:

- *Compliance with structural coupling principles*: In living systems, Maturana reminds us that the nervous system is subservient to the body. This will not be our case and it is justified by the fact that for structural coupling, we do not need autopoietic systems. Hence we can have two organizationally closed systems: body and neural network, structurally coupled to each other.

- *Biological inspiration*: First, nervous systems play a key part in action and perception of animals. In modelling sensorimotor activity in artificial agents, artificial neural networks are a natural candidate for modelling the role of nervous systems. Second, co-evolution is typical for nature and it is a good example of structural coupling. Therefore our model will be inspired by co-evolution.

- *Relational domain*: In nature, nervous systems add the domain of pure relations to organisms (see Section 2.2.3). As I want to model a relational definition of embodiment, this domain is suitable for this purpose, as it is not fixed to particular 'living matter'.

- *Performance*: I have declared that the model should produce some visible results. The coupling modelled on a comparatively high level, concentrating on the sensorimotor interaction with the environment co-ordinated by the nervous system, may appear more prominent than modelling physiological processes of an agent on a lower level. Hence a higher-level model may be more understandable and have a greater heuristic value.

3.2. Architecture – discussion of possibilities

I have already decided to model structural coupling in software and on a neural driven agent interacting with an environment. Thus I have to specify the agent (its neural controller and body) and the environment, their interaction and how they will change over time. This chapter presents a discussion of architectural possibilities, with respect to the qualities we want from the model.

Before we start with the individual components, let me emphasize, that we must always try to see them in context of the other components. Together, body, brain and environment constitute one system. The complexity of all the subsystems involved has to be balanced – this is sometimes called the *ecological balance principle*. An example from nature: complex nervous systems are present in organisms with large bodies that live with an active lifestyle in a complex environment.

3.2.1. Agent neural controller

What do we want from the neural controller? First, it has to be able to coordinate the sensorimotor activity of the agent, interacting with the environment. Second, it should be plastic – the engagement in interaction with the environment should change its structure, allowing for structural coupling. Possible neural network architectures that may constitute a neural controller of an agent are discussed from this perspective.

Feed-forward neural networks (FFNNs)

This is the most widely known architecture. The main strength of these networks is in effective training by gradient-descent methods such as the *backpropagation* algorithm. However, from our perspective, it is an *unsuitable* architecture. This is because while backpropagation is a *supervised learning* algorithm, in structural coupling, there is *no* room for a *supervisor*. There is only an agent and environment, with plastic structures and perturbatory channels between them. The environment gives only indirect clues what is good.

It is still possible to use this architecture and evolve the connection weights (by using genetic algorithms), rather than use backpropagation. This gives plasticity on the evolutionary time scale. However, the neural controller of an individual agent would be a FFNN with fixed weights: a mapping from sensors to motors. Such an agent (provided other parts of the agent do not facilitate it) is then a *reactive* one – reacting always in the same way on the same environmental situation; it is a mere puppet of the environmental puppeteer. As directly follows from the architecture, the neural controller will have no plasticity in the life of an individual agent.

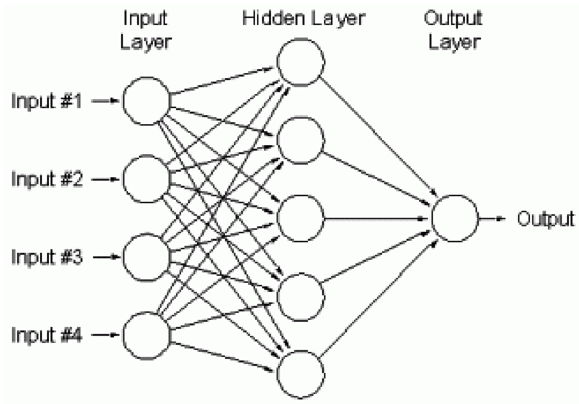

Figure 3.1. Feed-forward neural network (FFNN) with one hidden layer.

Simple recurrent neural network (SRNN)

Recurrent connections give a neural network temporal dynamics and enable a neural controller to deal with time dependent events. However, SRNNs are basically feed-forward networks with a partial feedback only. This is achieved through a context layer: the activation values from a previous time step are fed back into the network at a later time step. This is referred to as *first-order feedback* (Ziemke 2001, p. 6). Thus these networks do not provide rich temporal dynamics, but rather access to states of some units at a previous time step only. The connection weights to the context layer are typically fixed, whereas the other can be trained by a modified backpropagation algorithm or evolved. The former case is supervised learning, which we cannot use, the latter case means fixed weights during the operation of such a network. Such a neural controller can also support a simple sequence of actions by taking advantage of the context unit, thus being better than a FFNN. However, the use of previous states of the context unit is the only adaptivity to the environment or plasticity. Thus the potential for reaching a structural congruence with the body and the environment on individual time scale is very limited.

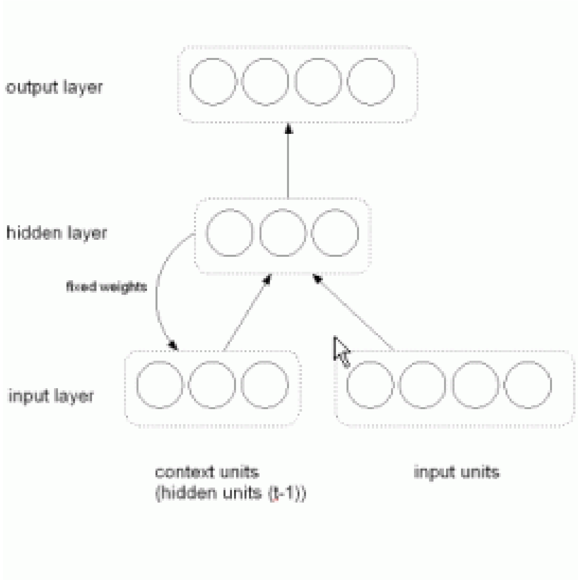

Figure 3.2. Simple recurrent neural network (SRNN). Activations from the hidden layer are fed back at a later time step through the context unit.

Continuous-time recurrent neural network (CTRNN)

These networks have usually a general recurrent architecture where every neuron is connected to every other and self-connections are allowed. Input from sensors can be connected to some neurons and activations of some neurons can be used as outputs. Time is continuous and the behaviour of neurons is governed by differential equations. These networks have a rich temporal dynamics and even networks of very few neurons may perform complex mappings (e.g., Beer 1995b).

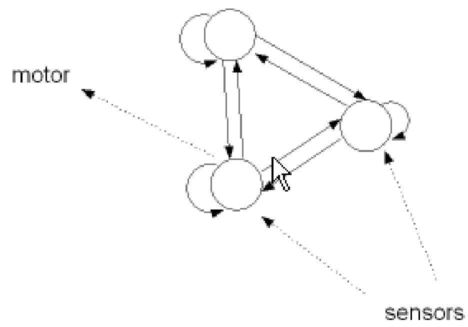

Figure 3.3. Continuous-time recurrent neural network (CTRNN). A general recurrent architecture.

Sequential behaviour and reinforcement learning

Yamauchi, Beer (1994) successfully evolved such networks[1]

1. to generate an output sequence (such as (01)*) in response to triggers from the environment – an equivalent of *sequential behaviour* of an agent.

2. to choose a correct sequence on reinforcement from the environment – '*reinforcement-like learning*'. ('Reinforcement-like' denotes that no explicit reinforcement-learning algorithm was used. However, from a behavioural point of view, there is no difference.)

These behaviours were achieved with no a priori discretization of states or time and no pre-designed learning algorithm. The genetic algorithm has prepared such a dynamics capable of accomplishing these tasks.

CTRNN and structural coupling perspective

These neural controllers in the example of the previous section were evolved: they are plastic on evolutionary time scale and can engage in structural coupling on this time scale. However, how is it with the time scale of individuals? The reinforcement-like capabilities demonstrate that such a controller can support adaptive behaviour. Hence, there must be some structural plasticity in the neural controller, even though the connection weights are fixed.

The plasticity is in the network dynamics and can be best explained by the language of dynamical systems theory (see Section A.1). The neural network is a dynamical system with

[1] The evolved parameters were connection weights, thresholds and time constants.

continuous dynamics. The state of this system consists in activations of neurons. The subsequent state is given by a dynamical law – in this case equations governing the activation of individual neurons. Some neurons may receive inputs from sensors, the activation of some may be sent to motors, taking part in the emergence of behaviour. The long-term behaviour of such a system is governed by *attractors* – subsets of the state space which attract nearby trajectories and once entered, the system remains there, unless perturbed. This may correspond for instance to a constant output to one motor. The distribution of attractors and of their basins of attraction is determined by the evolution of parameters (weights, thresholds), which shape the phase portrait. However, when some neurons are in contact with the environment (probably through the body), these inputs constantly perturb the system. A particular perturbation may eventually cause the system to leave an attractor and enter a different one. And this is the plasticity of this neural controller. It is perturbed by the environment and changes its structure as a result of this interaction.

The structural change may not be very intuitive here. If we imagine a physical instantiation of such a neural controller, then by entering a different attractor, there is no obvious physical change. However, the network is in a different state (and this will also be reflected materially), the relationships were perturbed and this determines further changes it can engage in. Hence, it is a structural change and such a neural controller *can be used as a substrate for structural coupling*.

Synaptic plasticity

Up to know only neuron activations were allowed to change in the active mode of the network (corresponding to life of an agent). We can introduce one more 'degree of freedom', if we allow also the connection weights to change. These are plastic synapses.

FFNN or SRNN with second order feedback

These are typically some forms of SRNN, i.e. a feed-forward network with a partial feedback. However, the feedback is not used to modify the output directly, but to adapt the connection weights and/or biases. We call this *second-order feedback*.

In Ziemke (2001) a sequential cascaded network (SCN) was used. This network consists of a feed-forward part (function network) – mapping from inputs to outputs – and a context network – state units connected to the input, which on the basis of their state change the weights from input to themselves and to output. The principal of operation in the experiment was basically this: By using an evolutionary algorithm the network (weights and biases) was evolved such that in the final configuration the context network was able to 'switch' between different weight settings on the basis of current state (current situation). And hence a different behaviour could have been produced. In other words, the state units do not affect the outputs (effectors) directly, but rather by changing the weights in the feed-forward part they change the mapping from sensors to effectors.

A similar architecture was used by Nolfi, Parisi (1996). A simple feed-forward network with no hidden units was used. There were only two extra teaching outputs, which were connected through fixed connections to the standard outputs and used for learning by backpropagation. The controller was evolved as a whole (weights on all connections) and finally a simulated robot was able to adapt itself to one of two possible versions of the environment. The network contained a predisposition to learn – an initial pattern of behaviour which enabled the agent to adapt to the environment. The action of the agent in the environment and resulting teaching outputs and backpropagation resulted in a gradual change of weights and adaptation on environment.

In principle, these neural architectures can be used as neural controllers for agents involved in structural coupling. Their structure is perturbed while interacting with the environment. However, what I object to these is that there is a considerable influence of the designer. The architecture consists of a feed-forward module and a teaching module ensuring the plasticity. Although the co-operation of these was evolved, I would still say that this plasticity, or the mechanism of change, is to a large extent *pre-designed*. Thus the designer influences to a large extent the possible structural changes in the neural controller. This is very different from the previously discussed CTRNNs, where a general recurrent topology was used and plasticity emerged from evolution of network parameters and interaction with environment. The separation of modules for sensory-motor mapping and adaptation, as used in these second-order feedback architectures, is also not biologically plausible.

Plastic Neural Network (PNN)

In the previous section, we saw architectures designed in such a way that one unit was responsible for agent's contribution to an emerging behaviour (by providing a mapping from sensors to motors) and another unit ensured changes of connection weights in the first unit, thus modifying the agent's activity. The cooperation of these units (weight setting of the context or teaching units) was evolved.

A different architecture was introduced by Floreano and Mondada (1996), further developed by Floreano, Urzelai (2000) and Blynel, Floreano (2002) and given the name Plastic Neural Networks (PNNs). There is no separate 'teacher'; the synaptic change is inspired by biological nervous systems and *Hebbian learning* is used. Thus there is only one network forming an agent neural controller, and its synapses dynamically change their strengths (in other words connection weights change) on the basis of neuron activations spreading through the network while the agent interacts with the environment. The architecture used in the experiments was a discrete-time neural network with recurrent connections. Individual synapses change their strength (denoted as w) as a result of activations of a *presynaptic neuron* x and *postsynaptic neuron* y (with range <0,1>), according to the following properties of the synapses:

- *learning rule (Hebbian)*
  - Plain Hebb: Strengthens the synapse proportionally to the correlated activity of x and y.

    $$\Delta w = (1 - w)\, xy$$

  - Postsynaptic rule: As plain Hebb, but includes also weakening of synaptic strength when y is active and x is not.

    $$\Delta w = w\,(-1 + x)\, y + (1 - w)\, xy$$

  - Presynaptic rule: Weakening occurs when x is active and y is not.

    $$\Delta w = wx\,(-1 + y) + (1 - w)\, xy$$

  - Covariance: Strengthening/weakening based on correlation of simultaneous activity. When the difference between activations of x and y is less than half their maximum activity, the synapse is strengthened, otherwise it is weakened.

    Equation 3.1. Hebbian learning rule, covariance 1.

    $$\Delta w = \begin{cases} (1 - w)\mathcal{F}(x, y) & \text{if } \mathcal{F}(x, y) > 0 \\ (w)\mathcal{F}(x, y) & \text{otherwise} \end{cases}$$

    Equation 3.2. Hebbian learning rule, covariance 2.

    $$\mathcal{F}(x, y) = \tanh(4(1 - |x - y|) - 2)$$

    where Equation 3.2 measures the difference between presynaptic and postsynaptic activity. It is positive for difference in activation greater than 0.5 (half the maximum activation) and negative otherwise.

- *learning rate*: magnitude of change.

The new weight (synaptic strength) is then the old weight (from previous time step), plus the weight change multiplied by the learning rate:

$$w_{ij}(n+1) = w_{ij}(n) + \eta \Delta w_{ij}(n)$$

The learning properties of individual synapses were genetically encoded and evolved. The evolved properties of synapses contained also *sign of a synapse* (whether it is excitatory or inhibitory) and in Floreano, Mondada (1996) there was one more property – whether a synapse has a driving or modulatory effect. In Floreano, Urzelai (2000) and Blynel, Floreano (2002) a more compact encoding was used, where all incoming synapses to a neuron shared the same learning properties and all outgoing synapses the same sign.

In individual agents, 'at birth', the synaptic strengths are initialized to small random values! Then these strengths dynamically change, while the agent interacts with the environment. In the experiments, it was shown that agents with such controllers are able to

- navigate in a simple arena (Floreano, Mondada 1996)
- learn on reinforcement (Blynel, Floreano 2002)
- perform a sequence of actions (Blynel, Floreano 2002)
- adapt to changes in sensor input (Blynel, Floreano 2002)

Thus, the behavioural repertoire supported by these networks is a rich one. Are they also a *good neural substrate for structural coupling*? Yes, they are. Basic structure is shaped by the evolutionary development of synaptic properties. This determines the relations between neurons and synapses. On the individual time scale, the structure is readily perturbed as a result of acting in the environment. Every neuron activation has an effect on synaptic strength, thus changes the structure and modifies the responses to perturbatory environmental events in the future. Moreover, the recurrent connections ensure also internal dynamics and perturbations coming from inside of the network.

Comparison of CTRNNs and PNNs

In the previous sections, I have identified two neural architectures as suitable candidates for the neural controller of an agent engaged in structural coupling with the environment: continuous-time recurrent neural networks (CTRNNs) and plastic neural networks (PNNs).

Which one is better? The abilities exhibited by agents driven by such networks were compared by Blynel, Floreano (2002) and were found to be similar (though there were differences identified). I would say that qualitatively, they can support the same 'class of behaviour'. A CTRNN can compensate for the synaptic plasticity in PNNs by continuous dynamics or more neurons.[2] As used in the experiments, there was an interaction of structural changes on evolutionary and individual time scale, as we would expect in structural coupling. Nevertheless, I would favour the PNNs for the reasons that follow.

First, *PNNs are more biologically plausible*. This is a complicated question, but let me highlight some of the major points. Although the dynamics of nervous systems plays a crucial role in brains, the synaptic strengths are changing, due to different mechanisms and on different time scales. The synaptic strengths are not encoded in the genotype. On the other hand, Hebbian learning and dynamic changes of synapses are based on neurophysiological evidence (for more details see Floreano, Mondada 1996). Second, the change of structure in PNNs in the form of changing synaptic strength is a more *intuitive* view of a structural change than a bifurcation causing transition to a different attractor, as in CTRNNs.

Other types

Let me only remark that there are many other neural architectures that were not mentioned. For instance, one could work with sophisticated neuron types (such as spiking neurons) or with modular architectures that are inspired by biological neural networks.

Summary – suitability of different neural architectures

Table 3.1. Neural architectures for structural coupling.

Neural architecture	**Dynamics**	**Pre-designed**	**Structural coupling**		
			Evolutionary time scale	**Individual time scale**	**Overall suitability**
FFNN[a]	discrete	highly	yes	no	very low
SRNN[a]	discrete	highly	yes	very low	very low
CTRNN	continuous	not much	yes	yes	high
FFNN/SRNN with second-order feedback	discrete	highly	yes	yes	medium
PNN	discrete	not much	yes	yes	very high

[a]We assume evolution of weights, not backpropagation.

[2]I have also analysed this from a dynamical systems perspective in Hoffmann 2006b.

3.2.2. Agent's body

In some cases the researchers are concerned with the neural controller only, providing it inputs from the 'environment', such as triggers and reinforcement, but there is no body (and perhaps no agent; e.g., Yamauchi, Beer 1996). However, we want to model embodiment and thus we certainly need a body. A very popular choice in experiments with autonomous agents are commercially available robots (such as Khepera or Koala, physical or simulated). These have a designed body, which cannot be changed. Therefore, such a body is not suitable, since it has no plasticity – it cannot be perturbed by the environment. Moreover, it is pre-designed, parachuted into an environment. On the other hand, we want a body that will be a result of a history of structural coupling.

With respect to our aim and the level of modelling chosen, a good choice is a simulated physical body with sensors and effectors connected to the neural controller, interacting with a simulated physical environment, but with no internal metabolism. A good example are the morphologies of Karl Sims' (1994) virtual creatures. Sims' creatures are also visually attractive and thus emphasize the concepts involved and help to bring them to a wider audience.

Plasticity of the bodies in such experiments is exclusively indirect – evolutionary. This is certainly a limitation. On the level on which the body is simulated, in animals there is bodily adaptation in individuals, such as long term changes in muscle strength (on exercise). Adding such plasticity would thus enhance the model of structural coupling. Nevertheless, we may say, that this kind of plasticity does not play a major part in (mature) individuals and as it makes simulations more complicated, it will be omitted here as well.

3.2.3. Environment

As a natural consequence of the analysis in previous sections (we want a neural driven agent with a simulated physical body), a suitable environment would be a simulated three-dimensional physical world, such as the one used in the work of Sims (1994). However, the problem from our perspective is the very plasticity of the environment. In the experiments similar to Sims', the method is usually such that the parameters of the environment (e.g., land or water, gravity, friction) are changed by the designer and then the agent is evolved to fit the environment. The environment is then subject to no perturbation, neither from the agent, nor any other. This is not suitable for our purposes, since such a structural coupling will be one-way only. I will introduce some plasticity into the simulated 3D world.

In our experiments, we will also use different environments: different terrain with different physical laws, and see how the agents couple to them. We call this part of the environment abiotic (see Section 2.1), and it will stay fixed during a simulation. However, to acquire also environmental plasticity, we will also introduce a biotic part that will be evolved together with the agent. This is very important and rarely present in other experiments.

For instance, if we look at most experiments with autonomous agents, the environment is typically a static one, such as an arena where an agent is supposed to learn a particular task, to avoid obstacles for instance. As we have said, this is not suitable for mutual structural coupling. Thus, as we claim that structural coupling underlies intelligent behaviour, we cannot expect a very sophisticated behaviour to emerge. This is because if one of the systems is not plastic, evolution will reach a limit soon, thus preventing also further structural changes in the other system: an agent can be only as intelligent as is its environment. Having a dynamically changing environment and perhaps trying different environments allows for further development of the agent.

3.2.4. Task or desired behaviour

According to Pfeifer, Scheier (2001) agent experiments always involve a *triad: agent – task/ desired behaviour – ecological niche* (I use environment). These have to be specified. However, the order in which they are specified depends on the goal of the research. The most common situation is this: it is decided on the task and environment (what Pfeifer, Scheier refer to as 'task environment') – this may be modelling of particular animal behaviour (e.g. cricket phonotaxis) or a general task, which may shed light on general principles of intelligence (such as whether a neural driven agent is capable of navigation, performing a sequence of actions or reinforce-learn) – and then the crux of the problem lies in developing an agent which will cope with the task in a given environment.

The difference between task and behaviour is approximately this: A task is more goal-oriented, for instance we may want a robot to cross an area with obstacles from point A to B. However, this task can be accomplished by different behaviours; the robot may detect the obstacles and avoid them, but it can also fly over the area or use brute force like a tank and go over the obstacles. The task can be even more general, such as survival in the world, and animals use very different behavioural repertoires to accomplish it. In AI, we are usually more concerned with (intelligent) behaviour than with tasks.

However, the goal followed in this thesis is special and thus requires special treatment. We are not modelling a particular behaviour but rather a concept underlying all behaviour: embodiment or structural coupling. Therefore, we can play with all members of the triad simultaneously to develop a model that would uncover something about this general principle. Moreover, structural coupling arises from the mutual interaction of agent and environment. Therefore, the environment should receive more attention than in usual agent experiments, which are to a large extent agent-centered. Do we then have a task? We have a goal: to develop a good model of structural coupling, but we do not have a specific task for our agent. This should be fine, as the task we talk about in agent experiments, are only in the eyes of the observer, not the agent. Nevertheless, we will need a fitness criterion, to guide our agent-environment interaction. We will deal with this in the following section.

3.2.5. Fitness

We have decided to evolve the structure of the agent's neural controller and body as well as the environment's structure. We thus need a selection criterion: a fitness function. However, in the previous section, we have found, that there is no task or desired behaviour. What should be the *fitness*? One way out would be to base fitness on quantifying structural coupling (see Section 4.2.2, “Quantification”). However, I do not consider this correct. Biological organisms are also not selected for their structural coupling. I will prefer leaving the attempts to quantify the phenomenon to evaluation. Instead, I prefer an *implicit fitness function*, such as lifespan of agents, or number of offspring, if they are able to reproduce. A spontaneous evolution model is ideal. Such a model will be more biologically realistic; structural coupling is the nature's answer to the selection criterion – number of offspring. As usual, we should not forget the environment. The biotic part of it will be co-evolved together with the agent and needs a fitness as well.

Chapter 4. Model

In this Chapter, I will first describe the software environment that was chosen, Framsticks. I will explain why it was chosen and then I will present an overview of those features that are relevant to my experiments. Then, the experiments conducted will be described. They share a common experimental setup: a population of agents is simulated in three different environments (flat, water and mountains), together with two populations of food, which can be eaten by the agents. Experiment 0 involves only testing of agents, which were developed by other users of Framsticks, in this new setup. The best agent is chosen and it is used as the starting agent in further experiments. Experiment 1 is a co-evolution of agents with food. In the three different environments, the agents as well as food populations undergo different structural changes, resulting in a structural fit. Experiment 2 is an extension of the first: the agent neural controller contains plastic synapses (see the section called “Plastic Neural Network (PNN)”), allowing for better structural coupling of an individual agent as well.

4.1. Architecture

In the previous chapter, we have analysed the design principles and parameters our model should have. Now, these have to be implemented (and perhaps modified) on a particular platform. The software environment chosen for my experiments is Framsticks.

4.1.1. Software environment – Framsticks

Framsticks[1] is a realistic, three-dimensional simulation of agents and their interactions. Simulated agents with a body and neural network act under simplified physical laws in the simulated world. This is also shown in a 3D OpenGL View. The agents have also their genotypes, on which genetic algorithms operate. After giving reasons for the choice of this software, I will go through the features relevant for my experiments.

Why Framsticks?

Framsticks was a good choice for my experiments for the following reasons:

- It simulates neural driven agents in a physical world. The agents' metabolism (or chemistry) is not simulated. This is exactly the level of modelling I have set forth in Section 3.1.

[1]You can find more information, download the software and manual etc. from the Framsticks website http://www.frams.alife.pl/. Description of the simulator in this thesis draws from there and from Komosinski, Rotaru-Varga (2001).

- It offers artificial evolution of agents, thus allowing for the evolutionary time scale of structural coupling.
- Simulator parameters can be defined by the user and there is a scripting language that enables the user to create own experiment definitions, fitness functions, or own neurons.
- This versatility allows me to introduce also environmental plasticity.
- 3D visualization is very important for understanding the model and for presenting the results.

4.1.2. Agents and environment

The *bodies* of agents are composed of sticks and are exposed to simplified physical laws. The physical interactions that will play a part in my experiments include action and reaction forces, static and dynamic friction, gravity, damping and uplift pressure in water (some of the forces depicts Figure 4.1). The step by step simulation uses a finite element method. There are articulations between the sticks. The relative position of the sticks may be controlled by muscles – bending and rotating – placed in these articulations. The bodies fulfil the criteria I have set forth in the previous Chapter.

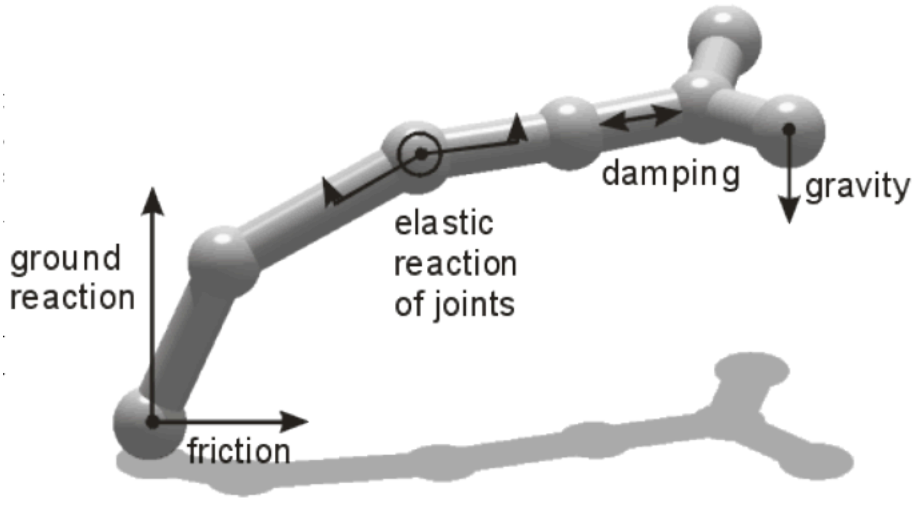

Figure 4.1. Forces on agent bodies (picture from Komosinski, Ulatowski 2006, with kind permission).

The *neural controllers* (in Framsticks they are called brains) may be composed of many neuron types. They share discrete time (simulation steps) and synchronous updating of activations. Initial activations are set stochastically to small values. Muscles (bending, rotating) have their respective neurons and so do the sensors. The major part will play touch sensors in my experiments; a gyroscope (orientation sensor) was used in the prologue. The internal nodes standardly used in Framsticks are N neurons. These are sigmoid neurons but their activation involves some inertia. Thus, apart from a parameter determining the steepness of the sigmoid activation function, there are two additional parameters, force and inertia. The former determines how quickly the neuron changes its activation on new inputs, the latter how strongly it keeps its trend. Some agents evolved by other Framsticks users contain for instance sinus generator neurons, or constant output neurons, but these will not be important in my experiments.

From the analysis of suitable neural architectures I have made in Section 3.2.1, it seems that the recommended ones are not implemented in this software environment. Since the simulation operates in discrete time, we cannot use a continuous-time recurrent neural network (CTRNN). The plastic neural networks (PNNs) were assessed as the most suitable substrate. However, in the experiments in Framsticks, neuron connection weights are evolved and stay fixed in individual agents. Nevertheless, it is possible to create custom neurons with plastic synapses and evolve synaptic properties rather than their strength. The strength can then change dynamically while the agent interacts with the environment. This will be done in Experiment 2.

The agent metabolism is substituted by an energy value. An agent has a particular energy at birth (e.g. a constant per stick) and loses it at a constant rate during its life, dying when energy reaches 0. It can get energy by finding food items in the environment. The energy consumption can also increase with muscle activity and it is also possible to turn on aging – the idle metabolism takes more energy after a certain period of life. This latter setting prevents successful creatures from living forever and is used in my experiments.

The *environment* is a landscape on which the agents live. Different environments (flat vs. mountains) provide different conditions for movement; some environments, like water, provide even new forces: the uplift pressure (buoyancy). Environment may also contain *food*. Food items may be placed into the environment at a constant rate or there may be a constant number of them. A new one is added every time a food is either eaten by an agent or it dies (food may also have energy that decreases at a certain rate). In my experiments, the food will play a little more sophisticated part.

4.1.3. Simulation

The 'life' of the agents in the world is made possible thanks to the following modules:

- *Physical simulation module*: Computes interaction of agents with the world, in 3D space. Forces on sticks are assessed and new positions calculated.

- *Neural module*: Manages the agents' neural controllers, provides them with sensory input and sends effector output to muscles.
- *Energetic module*: Takes care of energy in agents and food.
- *Creation module*: Creates new agents (e.g. mutation, crossover) and puts them into the world.

4.1.4. Evolution and genotype encoding

The idea of Framsticks is not to allow for simulation of agents designed by hand (this can be done as well, though), but rather to explore artificial evolutionary techniques. In order to be able to use a genetic algorithm, we need the agents' genotypes and hence a genotype encoding. In Framsticks, there are three full-fledged genotype encodings, but the user can also define his own. These encodings are *simul*, *recur* and *devel*.[2] *Simul* is a *low-level direct 'encoding'* which exactly corresponds to the way the agents are represented in the simulator – a list of objects the agent is composed of: body sticks and neurons, their attributes and how they are connected. Encoding and decoding is easy and there are no restrictions on agent phenotypes. However, this encoding is neither easy to use by a human, nor suitable for genetic algorithms – it is not compact and robust when the recombination operators (mutation, crossover) are applied.

Recur is a *direct recurrent encoding*. It is *higher-level* and can be better understood by humans. The decoding to phenotypes uses the *simul* encoding as an intermediary. A reverse procedure, encoding, is not possible. It is possible to represent only a subset of all possible phenotypes with this encoding. However, it is much more compact and robust in face of genetic operators and thus allows for better performance of a genetic algorithm.

Devel is an *indirect high-level developmental encoding*. It does not directly specify a phenotype but rather a developmental process leading to a mature individual. Otherwise it has similar properties as *recur*. Such an encoding is certainly more biologically plausible and developmental encodings were found to produce more structured and modular phenotypes than direct encodings (Komosinski, Rotaru-Varga 2001, p. 10-11).

Komosinski, Rotaru-Varga (2001) have compared the three encodings experimentally and showed that the *simul* encoding performs poorly in face of genetic algorithms. That is why it will not be used in my experiments. From our structural coupling perspective, a developmental encoding would be desirable. In nature, animals are not born mature, directly decoded from

[2] These names are used in Komosinski, Rotaru-Varga (2001). In the software and its manual, they are referred to as f0, f1 and f4 respectively.

their genotypes, but rather develop from a single cell containing genetic information. Cells divide and a multicellular organism emerges from the interaction of the cells and the genetic code in their nuclei. This mechanism of structural changes in an organism is greatly simplified and imitated by the *devel* encoding. However, in nature, this development occurs in an environment and is readily perturbed by it. *Devel* does not develop the organism in an environment, but rather before putting it into the world. As our main focus is on structural coupling with the environment, this does not help us a lot. This and also the fact that it is not as easy to use made me *choose* the *recur* genotype encoding. But the use of a developmental encoding, ideally one interacting with environment, would improve the model and is a theme for future work. Note that the encoding codes both morphology and neural controller and both are thus subject to evolution.

4.2. Experiments and results

4.2.1. Experimental setup common to all experiments

Environment

I use a set of three different environments in all experiments. They provide very different conditions for movement of agents. The reason for experimenting with different environments is to see structural changes of agents in different environments. Adaptive structural changes of agents in their environments can then be verified by testing them in the other environments.

Table 4.1. Environments.

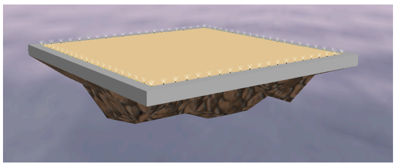 Flat.	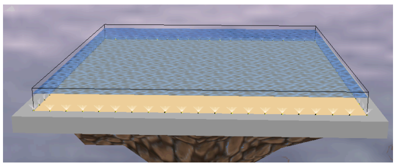 Water.	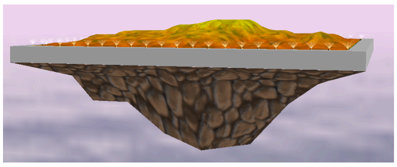 Mountains.

There will be a constant supply of food items in the world, namely 50 items, and a constant number of agents, 9. Agents as well as food will have a starting energy and it will be decreasing at a constant rate. Once it reaches 0, the agent/food 'dies' and a new one is put into the world, selected by the genetic algorithm. When an agent collides with food, it gains all its energy and the food thus disappears.

Genetic algorithm parameters

The parameters of the genetic algorithm are a result of experimentation. They are listed in Table 4.2. Some are default in Framsticks, some were changed in the course of experiments. In particular, originally I worked with roulette selection and deleting genotypes with inverse proportion to their fitness. However, this has lead to early homogenization of the gene pool. A tournament selection and deletion of worst genotypes produced better results.

The simulation is nondeterministic (random initialization of weights, random placements of agents and food into the world) and therefore more runs are needed to produce reliable results. More instances of a genotype may exist in the gene pool (its capacity limits the total number of instances of all genotypes). Sometimes a genotype is cloned and thus tested more than once, giving more reliable results within a single run.

Table 4.2. Parameters of genetic algorithm.

Type	Steady-state.
Gene pool size	200
Selection	Tournament, tournament size: 2.
Deleting genotypes	The worst.
Cloning probability	0.2
Crossing-over probability	0.16
Mutation probability	0.64
Fitness	Lifespan.
Genotype encoding	Recur.

4.2.2. Prologue and Experiment 0

For my experiments, I need a moving agent for all three environments. Rather than starting from scratch, I will take advantage of agents developed by other users of the simulator. I will describe this 'prologue' to my experiments and then, in Experiment 0, I will test the candidates in the three environments, and select the most suitable one for use in further experiments.

Prologue

Version 1.0 of Framsticks simulator has been released in 1998 and from that time a lot of agents (or creatures as the simulator users call them) have been developed. Very often they were evolved in a fixed environment for speed. The best ones are available in the Framsticks experimentation centre. In my experiments, I will also need walking (moving) agents and thus I find it useful to start from these, rather than from scratch.

A starting group of genotypes includes 9 genotypes. I have chosen these without any rigorous method, by looking at available agents. My criteria were these: I did not want the agents to have other sensors than touch and gyroscope (there are for instance creatures finding food with smell sensors but these I did not want this time). Then I wanted the agents to move well and to have a simple neural controller and body. Some of them are shown in Table 4.3.

Table 4.3. Moving agent candidates. Connection weights in neural controllers were evolved and stay fixed in individual agents. Neurons are distributed in the agent bodies.

Name and picture	Neural controller	Description
Amphibian jumper 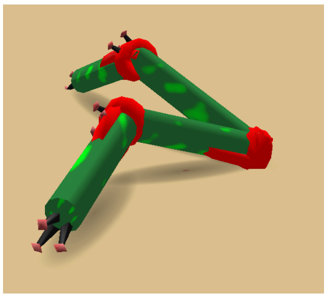	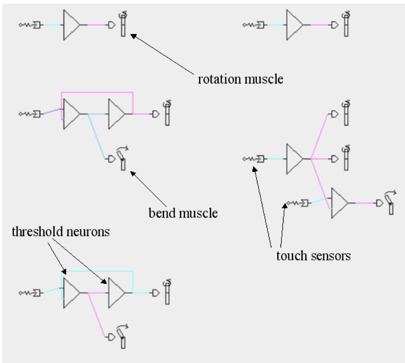 	Jumps in a random but efficient way. It falls on the ground and either throws itself directly back in the air, or bounces along first, till it reaches a position for a big jump. There are touch sensors in the articulations and these are connected to bending and rotating muscles in the same and other (mostly adjacent) articulations. They are co-ordinated by five small neural networks; two of them also contain simple recurrent connections.
Champ 	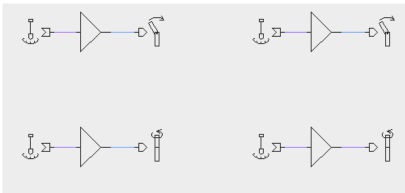	Hops on two legs. The neural controller is again distributed one, but this time gyroscopes are used instead of touch sensors.
Cheeta 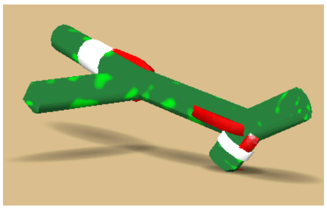	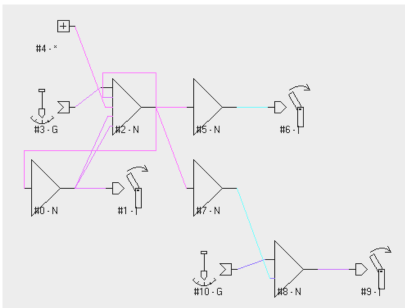	Jumps with the use of its hind leg. The coordinated movement is achieved by a single neural controller with gyroscopes.

Experimental setup (Exp. 0)

The agents were evolved for speed and often in a flat environment. I will however use a different experimental setup – different environments, food items and lifespan fitness (see Section 4.2.1). Therefore this experiment involves testing the agents in the new setup.

The agents can extend their lifespan by eating as much food as possible. Nevertheless, as the agents are not able to detect the food, they do best if they move efficiently in the environment. This fitness is thus positively correlated with velocity, but not the same. For instance the size of the agent has also an effect: a bigger agent has a greater chance of colliding with food. The food items used in this experiment are objects consisting of two sticks.

In every environment, all the 9 genotypes are tested together, each one at least ten times. When an agent dies, the simulator selects a new one (random selection of the genotype) and puts a new agent at a random place in the world. Food supply is kept constant. An instant from the testing in the Mountains is shown in Figure 4.2. The lifespans of all agent tests are recorded.

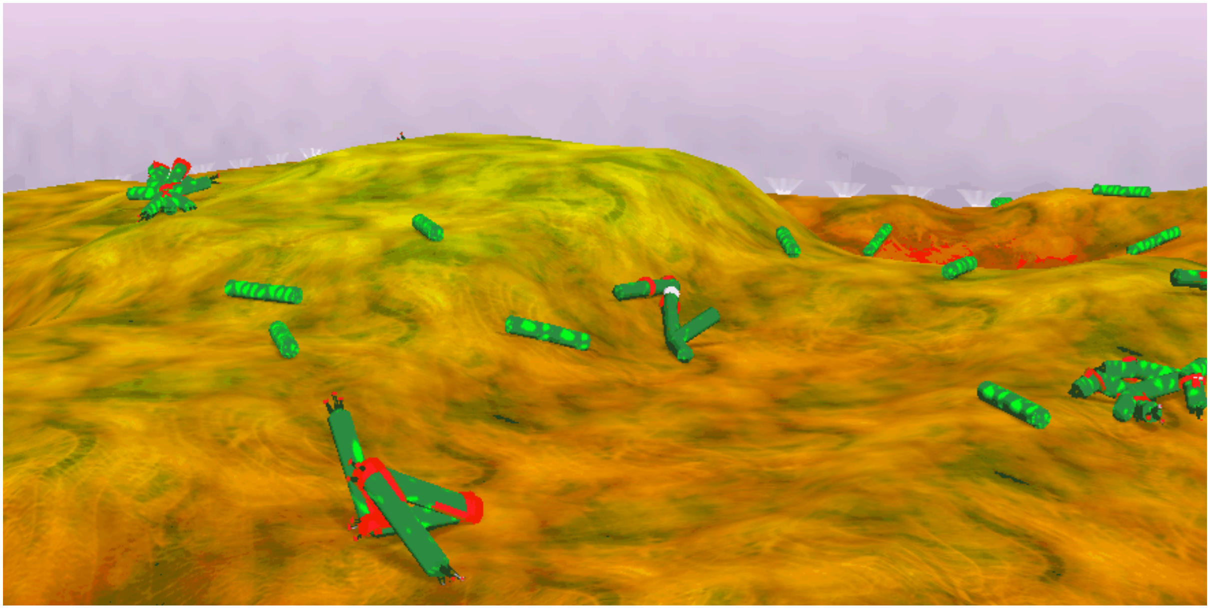

Figure 4.2. Experiment 0. Test in Mountains. The agent at the front is an amphibian jumper. The red parts are muscles, the white ring on another agent is a gyroscope. Food items are the simple sticks.

Results (Exp. 0)

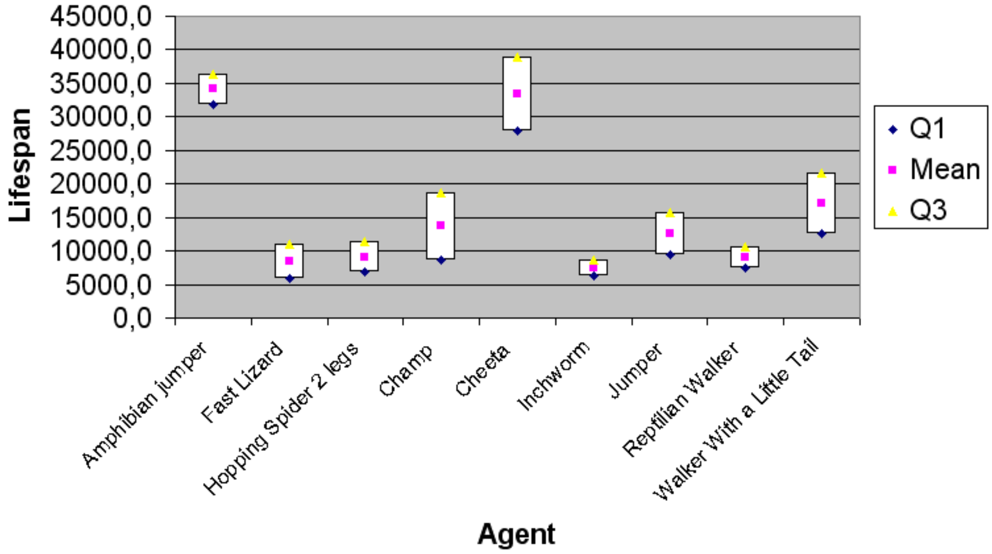

Figure 4.3. Experiment 0. Graph of average agent lifespan in Flat env. At least 10 tests per agent.[3]

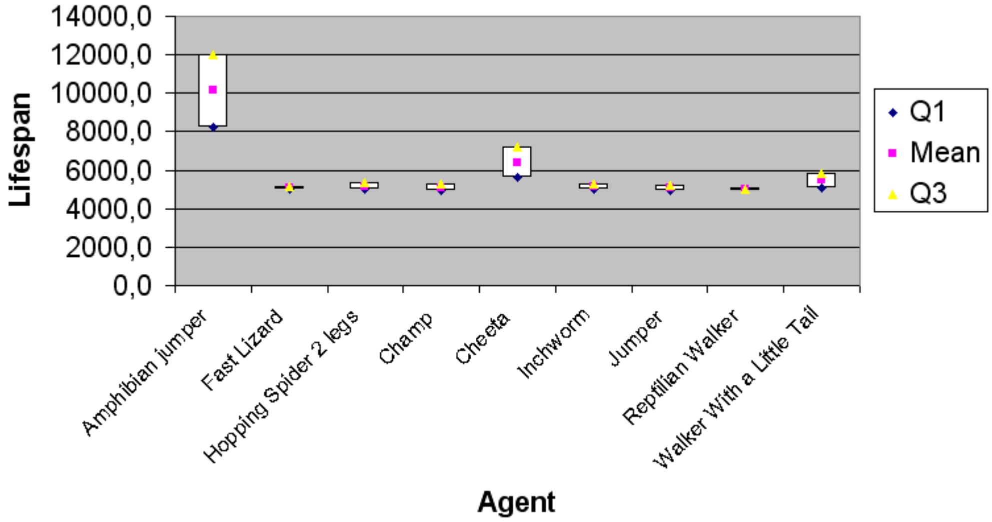

Figure 4.4. Experiment 0. Graph of average agent lifespan in Water env. At least 10 tests per agent.

[3]Please note the quartiles shown are an approximation. The data were recorded in the simulator and then their mean and standard deviation for every agent extracted. The quartiles shown assume a Gaussian distribution and that the quartiles then lie (2/3)*standard deviation either side of the mean. This will be used in the consecutive graphs as well.

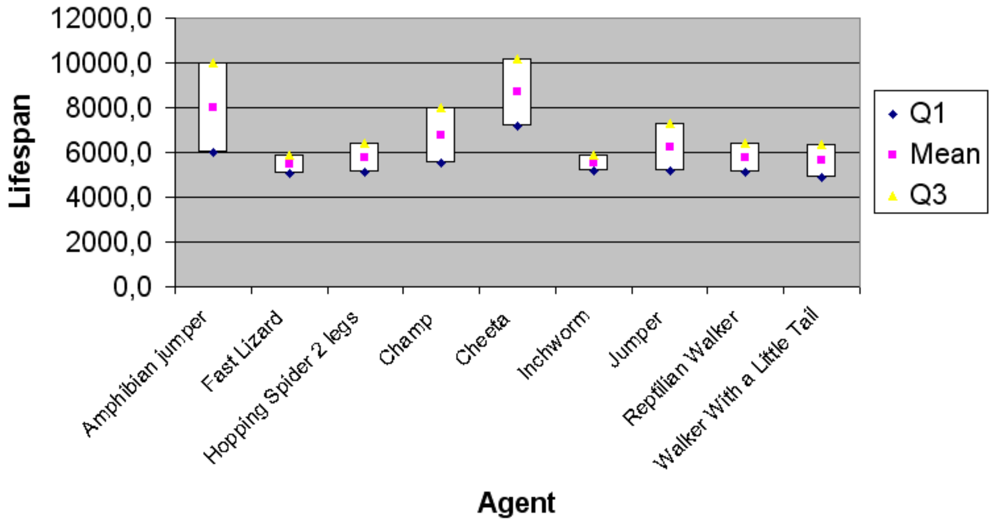

Figure 4.5. Experiment 0. Graph of average agent lifespan in Mountains. At least 10 tests per agent.

We see that the best agents in Flat and Mountains environments are Amphibian jumper and Cheeta. However, Cheeta does not cope very well with the Water environment and thus *I will use the Amphibian jumper*[4] in further experiments.

4.2.3. Experiment 1

This experiment involves evolution and then testing. Agents will be evolved in the three environments, starting from the Amphibian jumper genotype. Following the objectives set forth in the previous chapters, I intend some plasticity on the side of the environment as well. I have added the plasticity to food items. There are two food populations – Food1 and Food2 – and their morphology is evolved. Food1 evolves to get eaten and Food2 not to get eaten. This corresponds to the biological notion of co-evolution (Section 2.1). The food populations are the biotic part of environment. Food1 is in an arms race with the agents (e.g. a plant trying to hide from herbivores), whereas Food2 is in symbiosis (e.g. a plant or organism that needs to get eaten in order to reproduce). Note that the agents have no chance of detecting the food, as they do not have sensors for that.

Recur genotype encoding

To understand the agents' bodies and neural networks and how they are subject to evolution, we need to understand the basics of the genotype encoding used. The *body* is composed of

[4]The author of this agent is Victor Myagkov (victormnet@mtu-net.ru). The agent's morphology and neural controller are a result of an incremental combination of design and evolution.

sticks, denoted 'X'. The body may branch; branches are enclosed in parentheses, '()'. New sticks are joined with ends of the previous ones, forming a tree structure. The angle between branching sticks is divided by commas. The basic morphology of Amphibian jumper is simply this: `(XX,(XX,))`. Two two-stick parts branched at a certain angle.

To allow more possibilities, the sticks may be modified by special characters, which modify the stick's features and position. The *modifiers* used in my experiments are shown in Table 4.4. If the modifier is a small letter, the property is decreased; a capital letter has the opposite effect. In the Amphibian jumper's genotype, there are many L, M, and F modifiers used – making it larger, with stronger muscles and greater friction – contributing to better movement.

Table 4.4. Recur encoding – modifiers used in my experiments.

Sticks' joints properties	
R	Rotation.
Q	Twist.
C	Curvature.
Physical properties	
L	Length.
W	Weight (light sticks float in water).
F	Friction.
Biological properties	
M	Muscle strength.

Neurons, if there are any, are also written in the genotype, in square brackets. They are written after an X and then they reside on this stick. Inside the brackets, there is the neuron type specified, neuron properties and list of inputs, i.e. which neurons send their activations to this one.

Experimental setup (Exp. 1)

The setup is same as in the previous experiment. There are 9 agents simulated at a time and 50 food items, 25 of each food population. The reason for these counts is experimentation: for the world of chosen size, these numbers of agents and food items in the world at a time are balanced. The agents use the space available, collide rarely, and find enough food to have an effect on their lifespan. Food is often eaten by the agents, but has a chance of surviving till natural death. In other words, the *mutual perturbatory channels have a good bandwidth*. All populations are evolved: the agents' morphology as well as neural controllers are evolved, in the case of food only morphology is. Table 4.5 shows the parameters in detail. The initial genotype of the agents is Amphibian jumper; the food initial genotype is a single stick, 'X'. The 'strategy' of Food1 is to grow large, hence the food populations have a limit imposed on their maximum number of sticks. In all the three environments, all the three populations were evolved for millions of time steps, corresponding to approximately hundreds of generations.

Table 4.5. Detailed parameters of genetic algorithm – Experiment 1.

	Agents	**Food1**	**Food2**
Fitness	Lifespan	-Lifespan	Lifespan
Maximum number of sticks		7	7
Active modifiers for evolution	RrQqCcLlWw-FfMm	LlWwFf	LlWwFf
Mutation probabilities[a] – morphology:			
Add/remove a stick 'X'	0.05	0.05	0.05
Add/remove a junction '()'	0.02	0.02	0.02
Add/remove a comma ','	0.02	0.02	0.02
Add/remove a modifier	0.1	0.1	0.1
Mutation probabilities – neuron net:			
Add/remove a neuron	0.05	0	0
Add/remove neural connection	0.1		
Add/remove neuron property setting	0.1		
Change connection weight	1		
Change property value	0.05		

[a]These can be set by the user. They are relative probabilities. Real ones can be obtained by dividing by their sum.

Results (Exp. 1)

The experiments were conducted in all three environments. However, for reasons of brevity, let me describe the evolved populations for the water environment only. The evolved genotypes are the most interesting, as food can either float, or stay at the ground. Otherwise the interpretation of evolved structures would be similar in the other environments. That does not mean, however, that the evolved agents and food are similar. The difference will be demonstrated in the verifying experiment (the section called “Verification”)

Water environment

The simulation was run for 15 million time steps. A trend of average lifespan of evolved populations is shown in Figure 4.6 and snapshots from the start and end of experiment in Table 4.6. At the beginning the agents, Amphibian jumpers, are able to move, but their jumps are not very efficient. The food items are all identical, short sticks floating in the middle of the water slope. Over time, a *structural fit develops*. The agents are able to jump better, almost like dolphins. Whereas the initial agents' average velocities and distances traveled in their life were around 0.016 and 150 respectively, the agents after evolution achieve velocities around 0.03, thus twice as big, and the distances are around 900 (6x more). Thus they cover a larger area and are more likely to get more food: this is reflected in their rising lifespan. The two food populations can be distinguished easily. Food1, which 'wants' to get eaten, grows very large (attacking the maximum) and floats either on the water level, or, partly on the level and partly below. This gives a greater chance of being captured by an agent jumping from the ground. On the other hand, Food2 becomes very small and heavy, 'hiding' at the bottom.

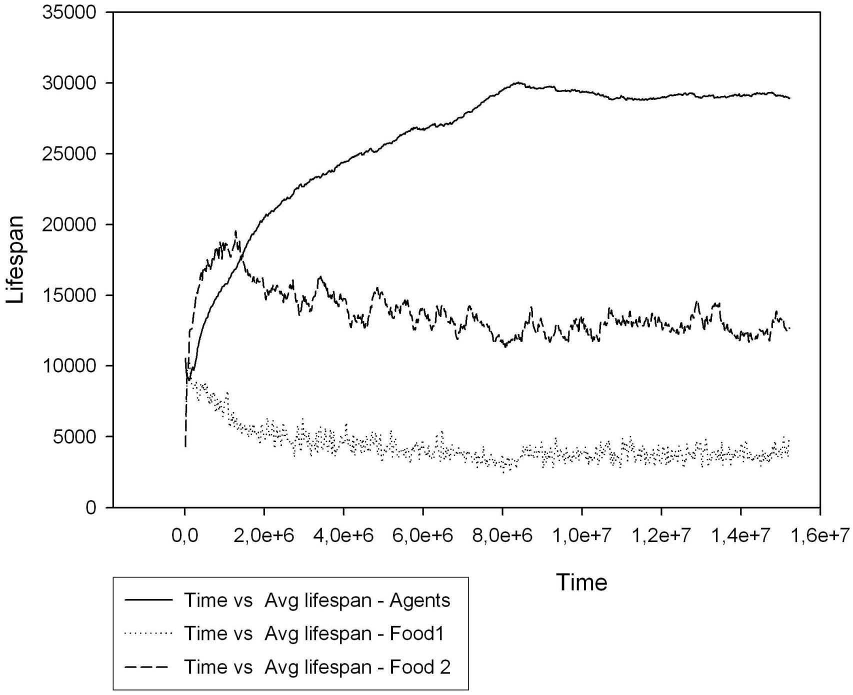

Figure 4.6. Experiment 1. Graph of agent and food population average lifespan during evolution in Water environment. The lifespan of agents rises steadily up to a maximum lifespan, around 30 000 time steps. The lifespan of Food 1 decreases, in an opposite fashion to the agents, showing symbiosis. The lifespan of Food2 ('prey') increases initially, but then falls a bit, due to better performance of agents.[5]

[5]Note that the graph shows lifespan, not fitness. That is important in case of Food1 population – its fitness was negative of lifespan.

Table 4.6. Experiment 1. Water environment – situation at start and end of experiment.

Start	**After 15 million time steps**
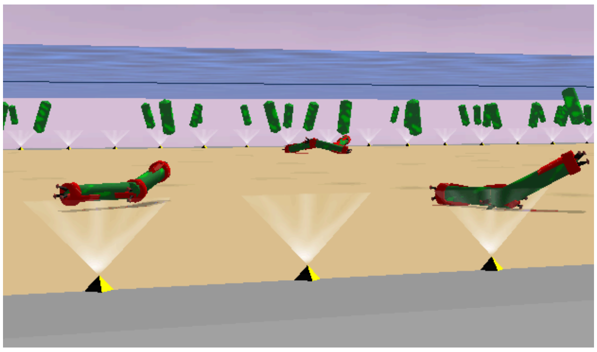	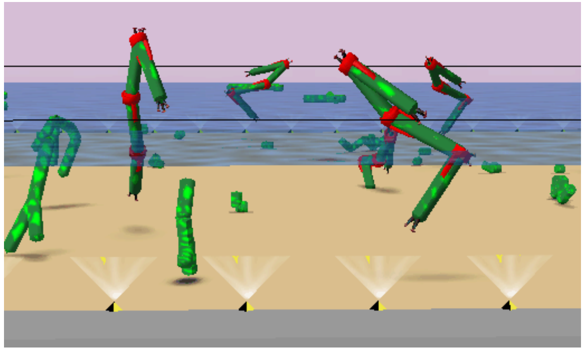
The agents near the bottom are not able to jump very efficiently. All the sticks standing vertically in the water slope are food items – both food populations look the same.	At the front, you see two agents jumping out of the water. The two objects standing from the bottom belong to Food1; the tiny ones scattered around the bottom belong to Food2.

What are the *changes in structure* of the genotypes in the populations? A look at the *agent genotypes* shows that the change in the emerging behaviour is due to changes in the connection weights in their neural controller. Although there was enough plasticity for evolution of morphology, it was not used, apart from minor changes in modifiers. Also apart from the weight changes, the other plasticity in the neural network, such as number of neurons and connections was also not used. This is because we have started with a body and neural controller already coupled to each other and to a flat environment, from the prologue. The adaptation to new environments was achieved through the weight changes, which were quite dramatic. However, it is not my point here to analyse, how they bring about different movement of the agents.

The changes in morphology of food, on the other hand, are big. Let us first look at *Food1*, which 'wants' to be eaten by the agents. The dominant change in the body is growing bigger. Whereas the original food consisted of a single stick, the evolved food attacks the limit of 7 sticks. The increase in size is also due to 'L' modifiers. Then there is also a use of 'w' modifiers, making the food light and float. This gives a better chance of colliding with the jumping creatures. Such a big food at the level, compared to the original food items around, is shown in Table 4.7, A, with its genotype underneath (the original genotype was just 'X'). Some food items even combine 'W' and 'w' modifiers in such a way that they have a vertical position from the level (B) or from the ground (C). Such a vertical position also increases the probability of being hit by an agent.

Table 4.7. Experiment 1. Evolved Food1 (the big ones), with genotypes shown below.

A	B	C
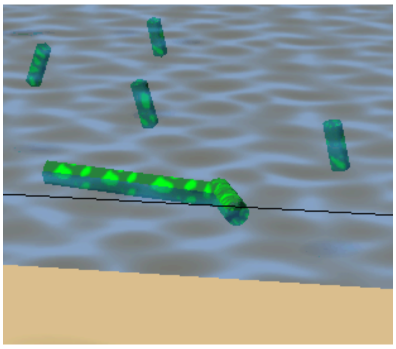	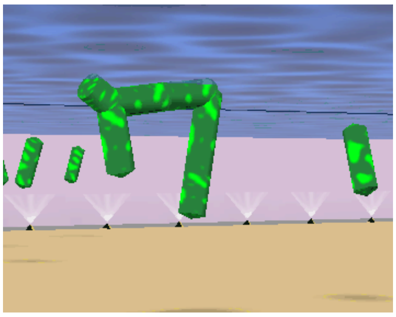	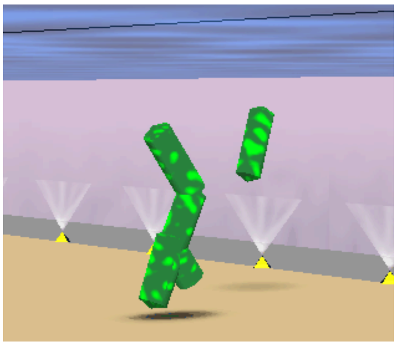
Lf (lllfWwwwflllFwW (w ((llwFllffffffffw (ffWX, fffwX LXlfffXwwwwX,, wwf (lllfff wwwwwwwXffX) ,) ,),) ,))	L (LFFw (lff (WXlllX, fffwlX XlllllFw (, Lw (fffffff WWWWX,) ,X,) , ,),,,) ,)	(f (lX (, lWw (llX) , lffffwX (llllw (, LFwf (ffffffff WWWW X,) ,,) ,,) ,,) ,wX) ,)

On the other hand, Food2 has gone through different changes. As it 'wants' to 'hide' from the agents, evolution has made its body small. Evolved genotypes consist of one to three sticks most often. There are many 'l' modifiers which make all the sticks very small. Also, there is a use of the 'W' modifiers, making the food heavy and stay at the ground. A single stick (A) and two-stick (B) Food2 are shown in Table 4.8, confronted again with the original colleagues floating nearby.

Table 4.8. Exp.1. Evolved Food2 (at the bottom), with genotypes shown below.

A	B
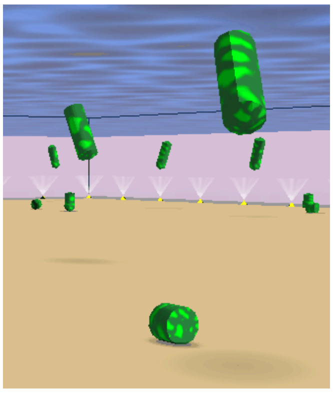	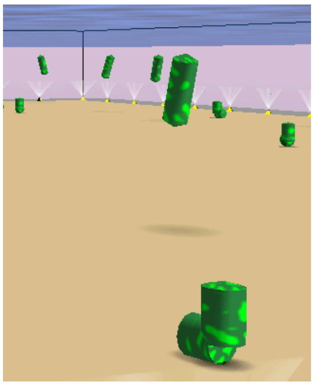
lllFWW(lfwwF((,llllFFFFFWX)),)	llWWwll((lw(, lllFFFWWX,), FwX,))

The evolved Food1 genotypes also use 'f' modifiers (smaller friction), whereas their counterparts, Food2, utilize 'F'. The effect of these on lifespan would have to be further investigated.

Verification

I will now show experimentally that, on the evolutionary time scale, the structure of the agents (their constituent components and relations between them) got coupled to the particular environment. I will test a collection of genotypes, created as follows: original Amphibian jumper and then there are two genotypes per environment, each from an independent evolutionary run in that environment, selected as the best one after 6.5 million time steps.[6]This collection was then tested in the three different environments. In every environment, the food items were also the best food genotypes evolved after 6.5 million time steps in that environment, in two independent runs. Graphs showing agent performance follow. Values of lifespan are averages of three independent testing runs in every environment; every agent was tested at least ten times in each test.

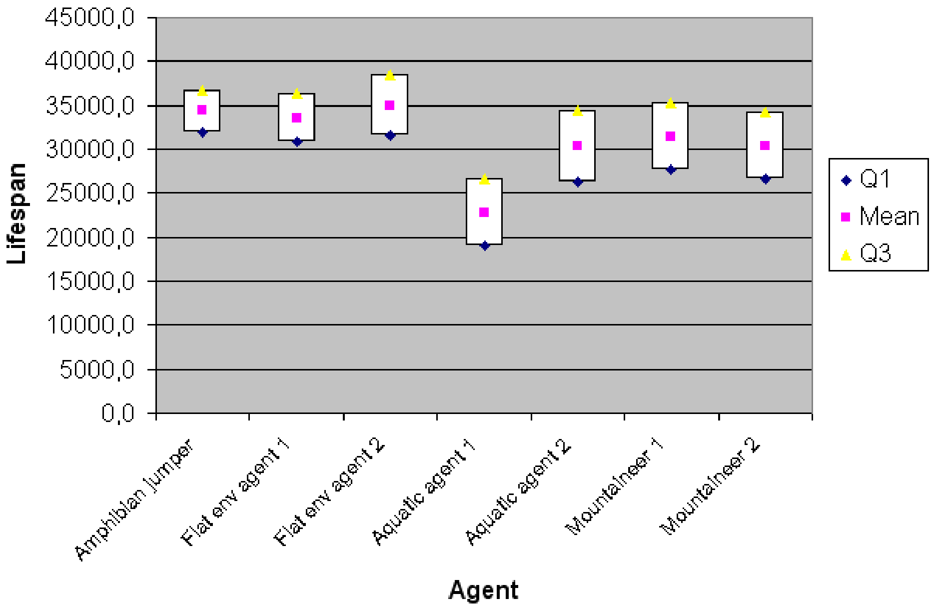

Figure 4.7. Experiment 1, Testing of evolved agents from different environments in the Flat one. Agents evolved in this environment have greatest lifespans, together with the original Amphibian jumper. This is because it was evolved in a flat env. for speed. It is thus a near optimal solution in my experiment as well.

[6]The criterion for selection was lifespan (fitness), however, only from genotypes with a certain number of instances. This is needed due to the nondeterministic nature of the simulation. For agent genotypes, this minimum number of instances was 5. From the genotypes with 5 or more instances, the one with greatest lifespan was selected.

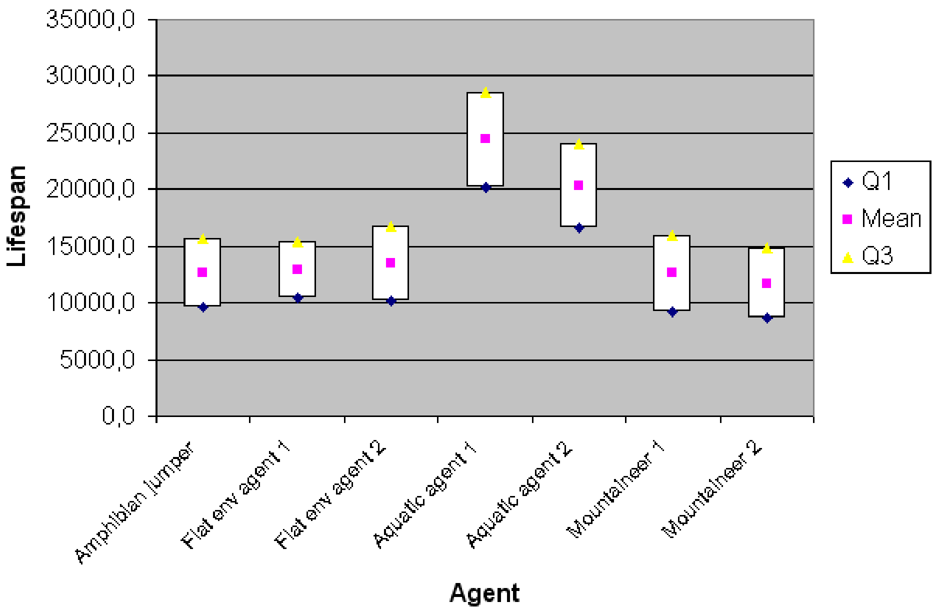

Figure 4.8. Experiment 1, Testing of evolved agents from different environments in Water environment.. Agents evolved in this environment have significantly longer lifespans.

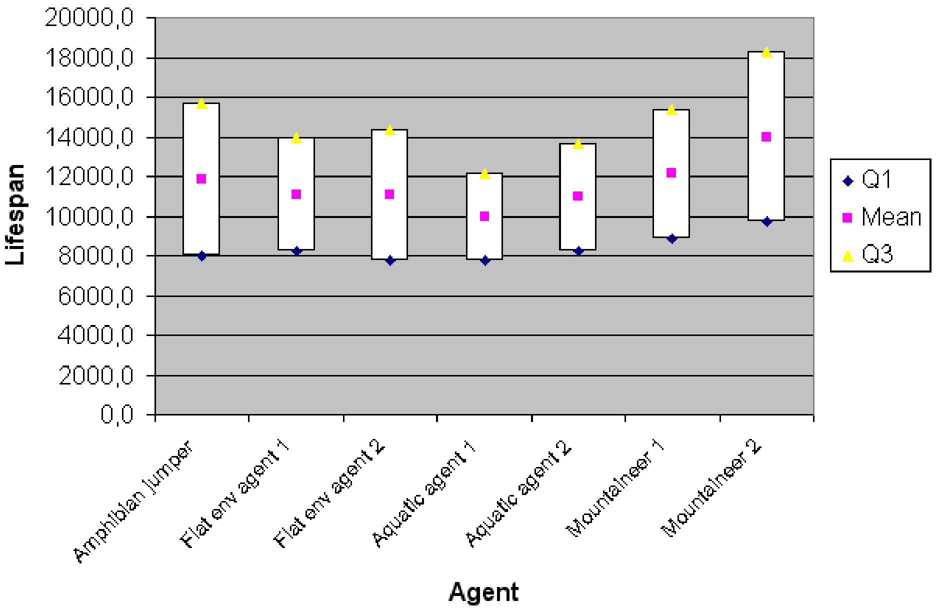

Figure 4.9. Experiment 1, Testing of evolved agents from different environments in Mountains. Agents evolved in this environment perform better. In this environment an agent's lifespan is highly dependent on its initial position. If it is for instance born in a deep valley, it is hard for it to escape it and thus it encounters less food. This causes a greater interquartile range of average lifespan.

The results show that the agents got coupled to their environments already after 6,5 million time steps. In every environment, those agents evolved in it perform best. After a longer evolution, the difference will grow. The greatest difference is in the water environment: agents evolved in different environments perform much worse.

4.2.4. Experiment 2

In Experiment 1, the agent's structure was changing on an evolutionary time scale, due to evolution of morphology and connection weights. However, we are interested in structural coupling on an individual time scale as well. Therefore, we have to introduce some plasticity into the individuals. According to the analysis from Section 3.2.1 we will use a neural architecture that allows this, namely a Plastic neural network (the section called “Plastic Neural Network (PNN)”).

For the initial agent genotypes, I have started from the Amphibian jumper again. I have preserved the morphology, but changed the neural controller. The neural controller of the Amphibian jumper consists of 5 feed-forward neural networks, mappings from touch sensors to nearby muscles. However, I intended to introduce neurons with plastic synapses. Therefore, I have created my custom 'plastic neurons', i.e. neurons with dynamically changing synapses. I used the version used in Blynel, Floreano (2002). All synapses leading to a neuron shared the same properties. So I replaced the threshold neurons leading to muscles in the original network with the plastic ones. I have manually created a collection of initial agents with different neural controllers. I tried different connections of neurons, sometimes also connections between separate networks (see the original Amphibian jumper in Table 4.3) to give more variability to the starting population. I have also added recurrent connections and self-connections.

The initial synapse parameters were random and were subject to evolution. The mutation probabilities differ from Experiment 1 and are shown in Table 4.9. Again, they were a result of experimentation. There are different settings concerning the agent neural controller than in Experiment 1. We want to concentrate on the synaptic properties. In the genotype, they are properties of neurons. Therefore the 'change property value' has by far the largest value. Apart from that, new connections are added or removed. Changing connection weights does not make sense here, as they are initialised to small random values and change dynamically in individuals, according to the synapses' properties. Neurons are also not added, however, there is a variability in the initial population and different neural architectures with different neuron numbers can be achieved by crossover. On the food's part, the plasticity of its morphology was increased.

Table 4.9. Experiment 2. Detailed parameters of genetic algorithm.

	Agents	Food1	Food2
Fitness	Lifespan	-Lifespan	Lifespan
Maximum number of sticks		7	7
Active modifiers for evolution	RrQqCcLlWw-FfMm	LlWwFf	LlWwFf
Mutation probabilities[a] – morphology:			
Add/remove a stick 'X'	0.05	0.2	0.2
Add/remove a junction '()'	0.02	0.04	0.04
Add/remove a comma ','	0.02	0.04	0.04
Add/remove a modifier	0.1	0.4	0.4
Mutation probabilities – neuron net:			
Add/remove a neuron	0.0	0	0
Add/remove neural connection	0.1		
Add/remove neuron property setting	0.0		
Change connection weight	0.0		
Change property value	7.0		

[a]These can be set by the user. They are relative probabilities. Real ones can be obtained by dividing by their sum.

Results (Exp. 2)

The experiments were conducted in all three environments. We will look very briefly at water environment and more closely on the mountains.

Water environment

Evolution was run for 15 million time steps (cca 1000 generations). I will pick up the situation at two points and demonstrate the effect of agents' movement on food. (1) At around 1.2 million time steps, agents crawl at the bottom. Food genotypes of the two populations do not differ significantly in size, but they differ in 'W'/'w' modifiers. Whereas Food1 contains 'W' modifiers, which make it heavy, and hence it stays at the bottom, Food2 'hides' at the top thanks to 'w' modifiers. Thus the two food populations have developed exactly opposite structural changes than in Exp. 1 where agents were jumping above water. (2) At 15 million time steps, the agents move more efficiently, making little jumps from the bottom with one end stretched to the top. Food1 has lost the advantage of being at the bottom, and has moved to the middle. There are not many 'W' modifiers anymore.

Mountains environment

The results from an evolution of 5,5 million time steps show interesting results when compared to the results from Experiment 1. The evolved populations in the world are shown in Figure 4.10. The agent moves not by jumping, but rather crawls in the terrain and is smaller – its body has three sticks only. The two food populations can be distinguished easily, and are similar to those from Experiment 1. Food1, which 'wants' to get eaten, grows very large (attacking the maximum). On the other hand, Food2 becomes very small.

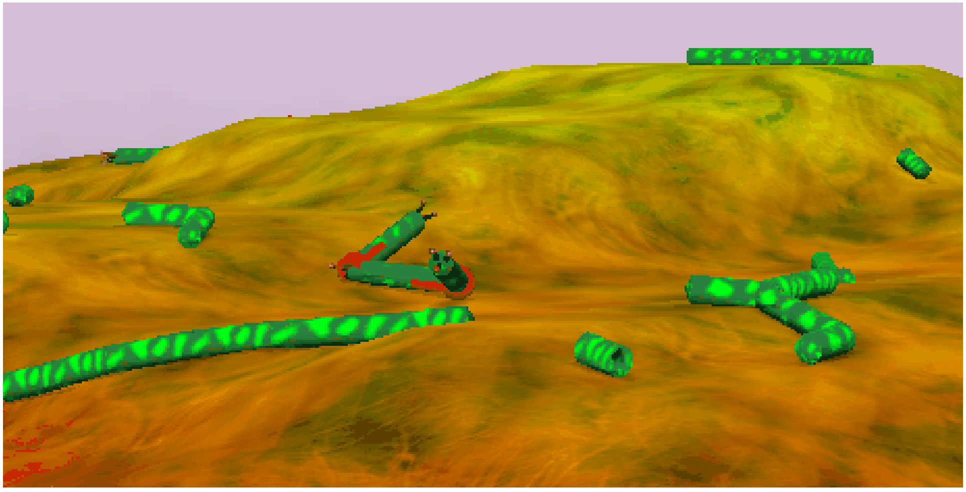

Figure 4.10. Experiment 2, Evolved populations in the Mountains environment. The giant objects belong to Food1, the tiny to Food2. In the middle is a crawling agent.

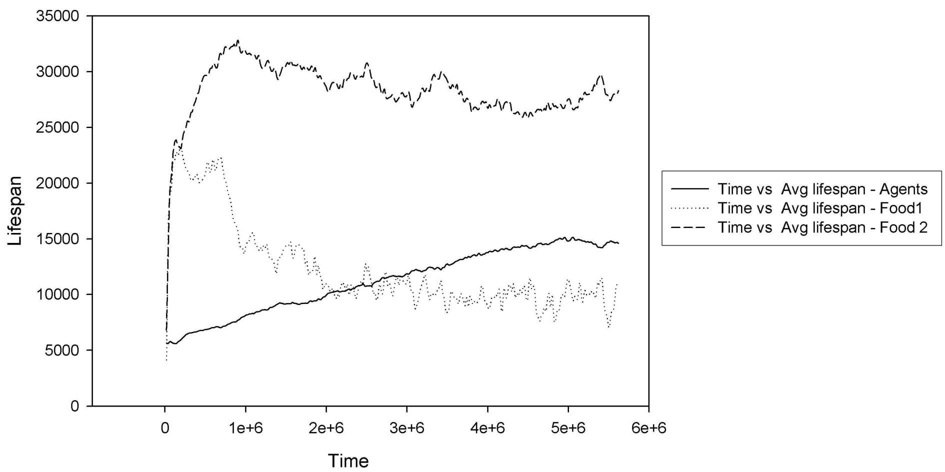

Figure 4.11. Experiment 2, Graph of agent and food population average lifespan during evolution.[7]

[7]The graph shows lifespan, not fitness. In case of Food1 population, fitness was negative of lifespan.

Let us examine an evolved agent after 5.5 million steps in more detail. Its crawling behaviour emerges from its body and neural controller (Table 4.10) interacting with environment. On the body there is above all one stick missing (and corresponding sensors and muscles) compared to the original morphology. Why has the morphology significantly changed, contrary to Exp. 1? This is because the initial neural controller did not support a co-ordinated activity. Evolution has used the plasticity of both, neural controller as well as body to find a moving agent. (In Exp. 1 we have started with a genotype of higher fitness – an already coupled body and neural controller. Removing one stick would result in dramatic fall in fitness.)

Table 4.10. Experiment 2. Evolved agent. Its genotype shows the body morphology as well as neural controller. Let me decode the neural network: In brackets, there are neurons. [T] is a touch sensor, [@] a rotation and [I] a bending muscle. [DS] are the 'plastic neurons' – with incoming dynamical synapses. [DS,0:1,-1:1...] means that this neuron is the plastic one, with a self-connection a and an input from its predecessor in the genotype. Then, these neurons contain evolved synaptic properties. Hebb_rule 1-4 is plain Hebb, postsynaptic, presynaptic, and covariance rule respectively. Learning_rate codes the rate of learning of the synapses and excit_inhib 0 means an excitatory synapse, 1 inhibitory. All incoming synapses to a neuron share the properties. However, they do not share the strength. In this agent, we see that most neurons have evolved synapses that are excitatory and with either a postsynaptic or covariance Hebbian rule.

<table>
<tr><th>Picture & Neural controller</th><th>Genotype</th></tr>
<tr><td>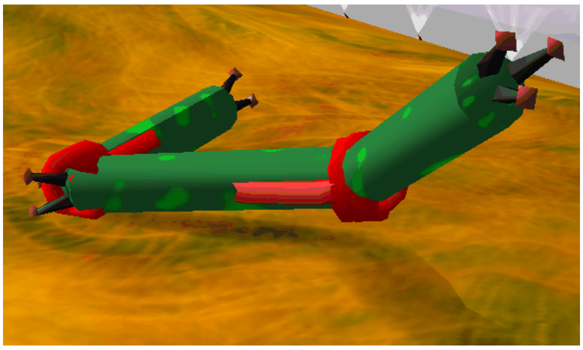
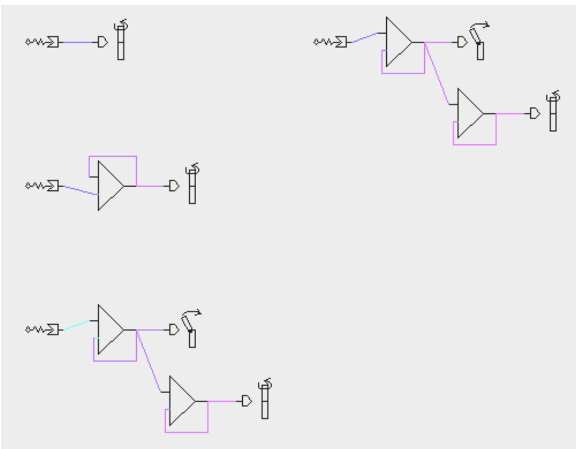</td><td>(I(rrLLLFW(LaFMMMM
X[T][@,-1:1][T][DS,0:1,-1:1,hebb_rule:2,
learning_rate:0.63, excit_inhib:1][@,-1:1]
,MMLFMMMMFFLLLF LMW
((rrLLLMMMMFFLMFMMMLFMLLLLLL
LLMMLLFM LLLLLLLMMLFMMLL-
MMLFMMMMFFLLLFLFMM
X[T][DS,-1:1,0:1,hebb_rule:2, learn-
ing_rate:0.5, excit_inhib:0]
[|,-1:1][DS,-2:1,0:1,hebb_rule:4, learn-
ing_rate:0.5, excit_inhib:0]
[@,-1:1]LLLLLFFFMMMM
X[T][DS,-1:1,0:1,hebb_rule:4, learn-
ing_rate:0.62, excit_inhib:0] [|,-1:1]
[DS,-2:1,0:1,hebb_rule:4, learning_rate:0.49,
excit_inhib:0] [@,-1:1],)))))</td></tr>
</table>

The neural controller is also simpler. It has nevertheless preserved the idea of separate modules; the connections between these were not supported by evolution. Self-connections (present in most of the initial genotypes), however, remained. Above all, in contrast to Exp. 1, the neurons shown contain dynamic synapses. Their strengths are initialized to small random values and then change according to Hebbian rules, while the agent interacts with the environment.

How the synapses of one neuron of an individual agent change during a small part of agent's life is shown in Figure 4.12. Whereas the strength of the second synapse quickly develops into its maximum strength, during few tens of sensory-motor cycles, the second one keeps dynamically changing while the agent interacts with the environment, possibly responding to different terrain. This means that the agent is much less reactive than before. The synaptic strength is key to its sensorimotor co-ordination and as it is dynamically changing, the reaction of muscles to same sensory input is different.

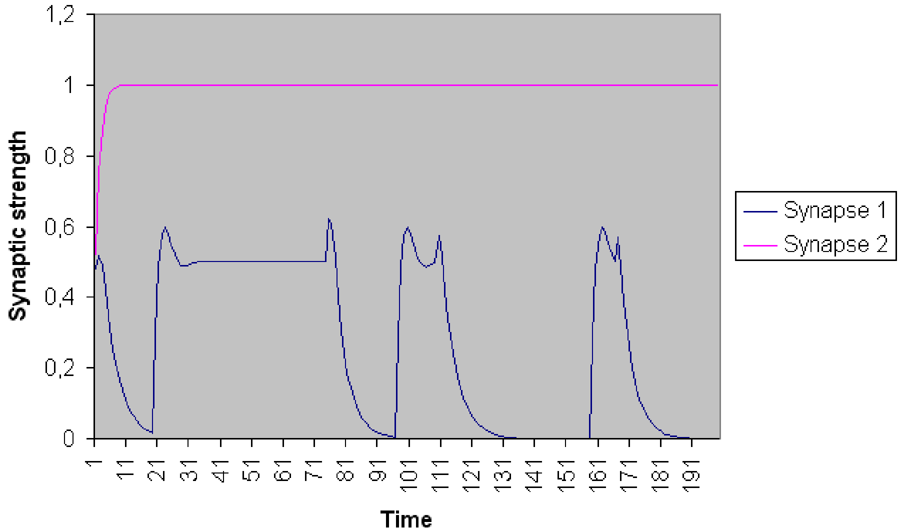

Figure 4.12. Experiment 2, Graph of changing synaptic strength in two incoming synapses of one neuron while the agent moves in a mountain environment. The evolved parameters of this synapse are the following: the Hebbian learning rule of the synapse is covariance, the learning rate 0.5 and the synapse is excitatory.

A thorough analysis of the neural controller dynamics will be relevant to answer how structural changes in it add to the emerging agent behaviour. However, such an analysis will be difficult and long. In this thesis, I will settle for the above demonstration of structural plasticity in the form of dynamic changes of synaptic strength. For an analysis of agents driven by PNNs, including neural network dynamics, see Floreano, Mondada (1996), or Blynel, Floreano (2002).

Chapter 5. Analysis and evaluation

This chapter is divided into three parts. First, I will interpret my experiments from a structural coupling perspective. I will apply R. Beer's (1995a) dynamical systems perspective on agent-environment interaction, modified for my case, in an attempt to formalize and quantify structural coupling. Second, I will present some consequences that can be drawn from the model. At last, the experiments together with the analysis will be evaluated.

5.1. Structural coupling – analysis of model

At the beginning of a structural coupling interpretation of the experiments, I will first identify the agent and environment and what constitutes their structure. Then I will observe their interaction and mutual perturbation, attempting to *quantify* it, according to the criteria already mentioned in the first chapter (Section 1.4.2), based on Quick et al. (1999, p.8), but with an extended notion of structure in dynamical systems perspective:

Quantifying structural coupling

1. The size of system's and environment's structure. How complex are the interacting systems, or if viewed as dynamical systems, how large is their state space and how complex are the maps or dynamical laws and how many system parameters are there (corresponding to their structure).

2. Plasticity of system and environment. To what extent (and how quickly) can they change.

3. Degree of mutual effect of system and environment on each other. To what extent they cause changes in the other's structure. Or bandwidth of perturbatory channels.

I will add a *distinction between individual and evolutionary time scale* to this and I will also not treat the agent as a single structure, but rather as a *body and neural controller*, these *coupled to each other*. In order to be able to talk more clearly about the systems involved, I will use a *dynamical systems perspective on agent-environment interaction*. The basics of dynamical systems theory and its application on agent and environment interacting, as introduced by Beer (1995a), is summarized in Appendix A, but I have adjusted it for discrete dynamical systems, as this is our case. Also, Beer concentrates on the agent's neural controller. The body, which is simple and not plastic, is made part of the environment. However, I am interested in the body as well as the neural controller, and thus I will have one more dynamical system. The dynamical systems perspective fits very nicely with our (structural coupling perspective) and provides a more formal mathematical structure.

In this section, I will go through my experiments. The agent and environment will be formalised as discrete dynamical systems; I will address their state space and dynamical law. On this basis, I will quantify structural coupling according to the criteria set above. I will analyse Experiment 1 where environment is co-evolved with the agent, and Experiment 2 where agents driven by plastic neural networks are introduced. I will show that this last one is a good structural coupling model[1].

5.1.1. Experiment 1

Let us look at the agent genotype, Amphibian jumper, to first informally uncover its structure and how it has come about. The basic structure of the body is simple: two sticks, a branch, and another two sticks. However, every articulation contains a bending and rotating muscle (one articulation contains only a rotating one). Hence the body has many degrees of freedom. Then, there are touch sensors at the end and at the articulations. The sensors send their signal to the neural controller which in turn sends signals to the effectors – muscles (for a picture of the body and neural controller see Table 4.3). There are 5 separate small neural networks, containing a touch sensor, and one or two effectors. They are thus like feed-forward neural networks, with no hidden layer, mapping signal from sensor onto the effector. However, the situation is a little more complicated. In some networks, the activation from sensor goes through one muscle neuron to another one and in two of the networks, there is also a recurrent connection. Also, in Framsticks, neurons associated with muscles have the same properties as the standard threshold neurons, thus having inertia. Therefore, even with no recurrent connections, this inertia provides a short-term memory, or internal state. How the agents move is a result of all the networks and a demonstration of a *principle of parallel loosely coupled processes*, as Pfeifer, Scheier (2001) call it. There are 9 agents in the environment at a time and they are evolved for lifespan. The agents can extend their lifespan by eating food. There are two food populations (25 of each in the world at a time) and their morphology is evolved.

Agent and environment as dynamical systems (Exp. 1)

Let us have a look at evolved agents interacting with environment. Let me introduce three discrete dynamical systems, `N` for the agent's neural controller, `B` for the body and `E` for the environment. Every individual agent (phenotype) will be described by the two dynamical systems, and `N` and `B` together make up a single dynamical system `A` (agent). The environment with all food items will constitute a single dynamical system, `E`, interacting with all agent systems that are in the environment. All the systems are nonautonomous coupled dynamical

[1] Please note that Experiment 2 is an extension of Experiment 1 and its analysis is also not self-contained, but rather based on its predecessor's interpretation.

systems (see Appendix A, and guard the difference between autonomous agent and autonomous dynamical system). The description will help us in the analysis of all the experiments.[2]

A. `N` (Neural controller, Figure 5.1)

- *State space*: The state space is a set of all possible values of neuron activations. The activations have an effect on activations at a later time step due to their inertia and also if there are recurrent connections. There are 8 neurons in the Amphibian jumper, with activations in the range [-1,1]. These determine the size of the state space.

  Note, that if the neurons did not have inertia and recurrent connections, we would have an ordinary FFNN and hence a trivial state space – only a mapping from the input coming from the body sensors to the output going to body effectors. Such a mapping could be included into the body's map and we would not even need a system `N`.

- *Map*:

  Equation 5.1. Map – neural controller.

  $$\mathbf{X}_{n+1}^{N} = \mathrm{F}^{N}(\mathbf{X}_{n}^{N}; \mathrm{I}(\mathbf{X}_{n}^{B}); \mathbf{u}^{N})$$

  where

  - $\mathbf{X}_{n}^{N}$ is a vector of values of state variables at time step n.
  - $\mathrm{I}(\mathbf{X}_{n}^{B})$ are variable parameters of the system. The function `I` is a mapping from the values of a subset of state variables of the body, `B`.
  - $\mathbf{u}^{N}$ are fixed parameters of the system.

  The map, `F`, maps the values of state variables from time step n to their values in the next time step, n+1. It depends on many parameters. Some of them are fixed, `u`. These are connection weights, thresholds and other neuron parameters. However, some of them are not fixed, but they are a function of the state variables of another system, `B`. These are inputs to the neurons from body sensors, denoted as function `I`. These constantly change the phase portrait of system `N`. Moreover, they are the key

[2] Note that the division between the systems is arbitrary, as already analysed by Beer (1995a, p. 181). Also, this formalization is only schematic. For instance, the food items that I make all part of the environmental dynamical system could also be treated as agents with passive bodies, subject to evolution. Every food item could be described as a simple dynamical system. Nevertheless, the division I have chosen allows for the most intuitive structural coupling interpretation.

variables influencing the dynamics of this system, the neuron inertia and recurrent connections being only a complement.

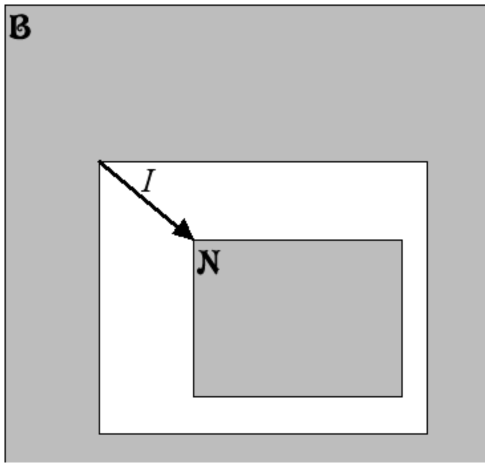

Figure 5.1. Agent neural controller (N) as a dynamical system. It is enclosed in the body (B). N acquires the values of its variable parameters through the function I, which is a mapping from a subset of B's states.

B. `B` (Body)

- *State space*: The state variables of this system are properties of the points – ends of sticks – as these are the simulated parts by the finite element method used by the simulator. They are angle between the sticks in 3D space and some other properties like damping of body parts. Current state of sensors, as acquired from the environment, will also be among state variables.

- *Map*:

Equation 5.2. Map – body.

$$\mathbf{X}_{n+1}^{B} = \mathrm{F}^{B}(\mathbf{X}_{n}^{B}; S(\mathbf{X}_{n}^{E}), M(\mathbf{X}_{n}^{N}); \mathbf{u}^{B})$$

The map `F` maps the relative position of joints and states of sensors and effectors from time step n onto their values at time step n+1. The fixed parameters, `u`, are the properties of sticks (size, weight, curvature...) and constants in the laws guiding the physics of the body (e.g. elasticity). The variable parameters are a function of a subset of the neural controller's states (motor function `M`), and of the environmental states (sensory function `S`), which include interaction with food and also other agents. The new angles at joints are a result of combining the activation going to muscles from the motor function `M` and the forces exerted by the environment, coming through the sensory function `S`. The states of sensors are determined by the signals coming from the environment, also through the function `S`. The body acts as an intermediary only, bringing the sensations from the environment to the neural controller.

Note that the system has no dynamics of its own. Rather it combines the input from muscles with forces exerted on it by the environment to produce new angles of joints. The body is thus reactive only.

C. `E` (Environment)

- *State space*: The state space is a set of all possible values of state variables. These are positions of agents' and food items' parts in the environment, their energy content, and other values such as speed and direction of movement of the parts.

- *Map*:

Equation 5.3. Map – environment.

$$\mathbf{X}_{n+1}^{E} = \mathrm{F}^{E}(\mathbf{X}_{n}^{E}; P(\mathbf{X}_{n}^{B1}, \ldots, \mathbf{X}_{n}^{B9}); \mathbf{u}^{E})$$

where $\mathbf{X}_{n}^{B1}, \ldots, \mathbf{X}_{n}^{B9}$ are subsets of state variables of bodies of the 9 agents in simulated world.

The map transforms the position of agents' parts at time step n to the positions at the next time step. It does this on the basis of previous positions, inertia of the bodies in the environment, and muscle forces from the body parts, acquired as variable parameters through the function `P`. The map will also include dying of agents and food, either from natural death or on colliding with agent, and birth of new agents and food items.

The fixed parameters, `u`, contain the environmental landscape and various constants, which are needed for the physical laws of the simulated 3D world. The gene pool of food populations (determining their morphology) can also be in the `E`'s fixed parameters, i.e. fixed on the individual time scale.

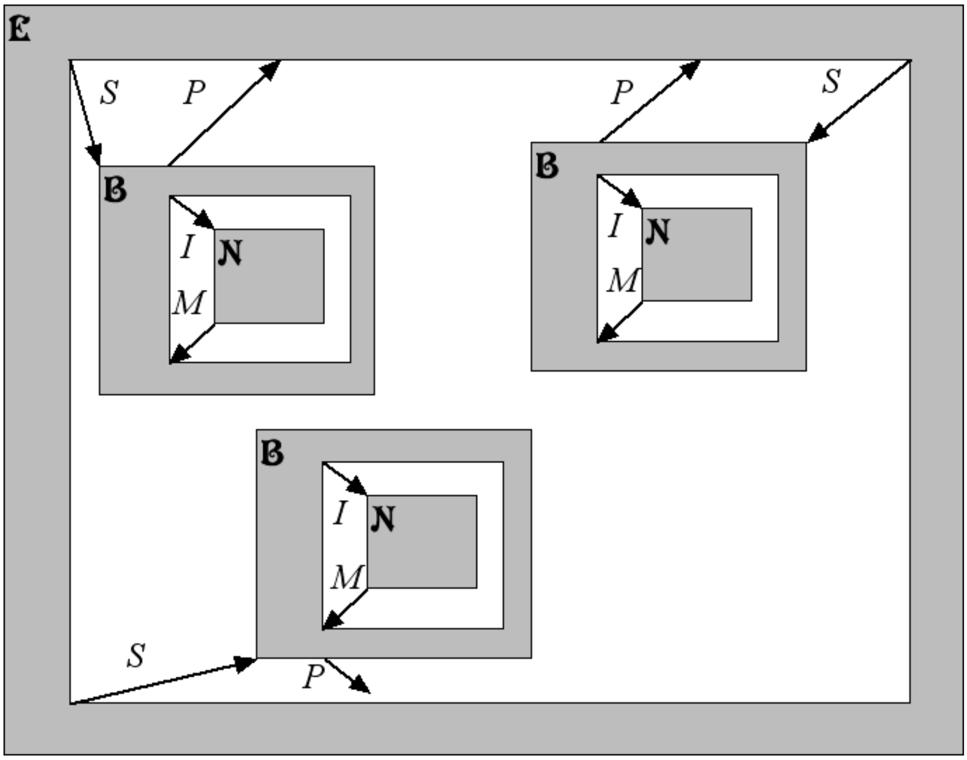

Figure 5.2. Interaction of agent (neural controller and body) with environment. There are three agents shown. In my experiments, there is 9 of them.

Individual time scale (Exp. 1)

Let me propose a quantification of structural coupling of all the systems altogether, on individual time scale, and according to the criteria set forth in Quantifying structural coupling.

1. *Size of structure*: `N`'s state space consists of all possible activations of 8 neurons; `B`'s of angles at articulations and states of sensors, and `E`'s of positions and velocities of agents' and food parts. The maps are complicated and have many parameters.

2. *Plasticity*: The values of state variables of the systems can change quickly. Neuron activations as well as body angles and agent positions in the environment change every time step. The maps and parameters do not change.

3. *Degree of mutual effect of systems on each other*: The activations of neurons are highly dependent on inputs from body. The body's state variables are completely dependent on the neural controller's and environmental states. The positions of agents in the environment are perturbed by the agent activities.

Although this separation into three components has helped us to uncover the nature of structural coupling, the interpretation is still far from straightforward. A lot of complexity comes from the simulated physics and 3D world. The positions of agents' parts in the environment have to be in some state space, however, they are not what we will intuitively understand as structure. Another problem is with very fast changes of neuron activations or angles at bodily articulations. These changes correspond to active dynamics of the systems. On the other hand, what we will intuitively perceive as adaptive (structural) changes in an agent, needs a greater time scale of changes.

A possible solution to this problem is to *further divide the time scale*. We could define an *active dynamics time scale*. This would comprise changes from time step to time step, such as change of agent's position and neuron activations while the agent moves in the environment. And we could say that changes on this time scale are not structural changes. A greater time scale could be called *learning time scale*. This would comprise slower changes that *persist* and affect future behaviour, causing for instance different future response to the same perturbation. This is a structural change. Note also that the structure are the components and relationships between them. Hence, the map and parameters are an equally important part of the structure as the state space. And these do not change on an individual time scale.

Looking at the situation from this perspective, we see a reactive-like agent jumping in an unchanging environment. The agent is not completely reactive thanks to the inertia of the neuron activations and some recurrent connections, thus having a short-term 'memory'. These structural changes have an effect in a few next steps only. On the side of the environment, the only interesting structural change caused by the agents is an accumulation of food items in some regions. This occurs in the mountains environment, where the mountain peaks are difficult to reach by the agents and hence the food survives longer there. The individual food items have no plasticity and hence there are otherwise no interesting structural changes on this time scale. Therefore, I would argue that *structural coupling* on individual time scale is actually *very limited*.

What lesson can we draw from that? The perspective of dynamical systems and a quantification cannot be applied mechanically. *Not all changes in the dynamical systems can be perceived as structural changes*. I have further divided the individual time scale into a smaller one, active dynamics, and a greater one, learning time scale, which better suites the notion of structural coupling. I have also evaluated the structural coupling on this time scale as weak. Nevertheless we see a structural fit: co-ordinated jumping of agents, in symbiosis with one food population and arms race with the other. I propose that this congruence is mainly thanks to the time scale discussed in the following section.

Evolutionary time scale (Exp. 1)

Structural congruence of agents and food populations is mainly a result of structural coupling on a different time scale – evolutionary. On the *agent's part*, the most structural changes took place already in the prologue – evolution of Amphibian jumper for speed. Then, the agents got coupled to the new environment and task mainly through changes in connection weights. Let me analyse this two-step agent evolution together.

Structural coupling is *indirect*, operating on the genotypes which code the structure of individual agents. These changes do not occur in an individual agent (which is described by `N` and `B`), but gradually, in the agent species. The genetic algorithm selects genotypes on the basis of their phenotypes' (individual agents') performance, applies genetic operators, and again tests the resulting phenotypes. The genotypes slowly change and thus the initial structures of individual agents (phenotypes) change. Plasticity is much greater than on the individual time scale. On the neural controller's part, the changes are in the state space's size (adding/removing neurons) and in the whole map. Changing neural connections, changing neuron properties, and changing weights – fixed parameters on individual time scale – perturbates the relationships and thus changes the structure. In Framsticks, adding of muscles and sensors is part of the neural controller evolution. In my division, this constitutes the body. In the body there is also a great deal of plasticity. Apart from the muscles and sensors, the whole morphology is evolved – number, size, shape and other properties of sticks, and how they are joined together. The sensory function `S` and the motor one, `M`, are also evolved (through the changes in body, sensors and muscles). Thus the whole state space as well as the map with the functions `S` and `M` and the fixed parameters is evolved. The plasticity of the agent's structure is large. Note also that evolution ensures a tight coupling between the agent's neural controller and body. The agent's movement emerges from their interaction with each other and the environment.

To what extent are the perturbations in agent's structure caused by the environment? The environment has a large, albeit indirect, effect. The genotype that codes the structure is directly perturbed by the genetic algorithm and these perturbations are random. However, the genotypes of agents performing well in the particular environment expand. In the dynamical systems view, fitness may be scored for staying in a particular volume of the state space of a single autonomous dynamical system, a union of `N`, `B` and `E` in our case (as I propose in Hoffmann 2006b). Thus the environment is involved in shaping the state space and map of this single autonomous dynamical system and consequently in the selection of genotypes by the genetic algorithm.

Let us look at the *environment*. It follows from our dynamical systems formalization that the system E, contrary to N or B, never ceases to exist. Therefore, its *plasticity* is always *direct*, albeit on different time scales. Plasticity on the evolutionary time scale consists in changes of food morphology. Similarly to the agent bodies, the number, connection, and properties of sticks are evolved. The structural changes of both food populations are big. What remains to be explained is whether this perturbatory channel is from the agent. This was best shown in the water environment where the food adapts on the way the agents move. (At the same time, however, the agents adapt to where the food is. Thus we have an illustration of mutual structural coupling on this time scale.)

Summary

In this experiment, a structural fit has developed between the agents and environment, thanks to the changes on evolutionary time scale. In all the three environments, we can observe the results of structural coupling, a sense of co-ordination. The agents move in an efficient way, often colliding with Food1, which is typically big and in water also in an appropriate part of the slope, less often colliding with Food2, which is typically small, in water also hiding at the bottom. This co-ordination is, however, a result of co-evolution, structural coupling on the evolutionary time scale. On the individual time scale, there are no further big or interesting structural changes caused by mutual perturbations. The environment is passive and the agents move in a reactive fashion.

We have performed this experiment in three different environments to show how this affects the structural changes. In the different environments, all the systems were perturbed in a different way, and have developed different structures. Hence there is a different structural fit and observed behaviour. I suggest a biological parallel here: the different genotypes of the agents and food which evolved in the different environments are very different and this resembles the case of vast genetic distances in organisms, like bacteria species, as discussed in the last part of Section 2.2.1. Both are a result of a history of structural coupling.

The definition of structural coupling is very general and minimalist. I propose treating structural coupling as a continuum, rather than an all or nothing property. The formalization and quantification I have suggested helps us in deciding on the *degree of structural coupling*. If it was extended then it could be used to compare different experiments on a quantitative basis.

5.1.2. Experiment 2

The neural controller of the agent is a plastic neural network (PNN). In my view, this experiment constitutes a good model of structural coupling. Let us first see how the dynamical systems have changed when compared to the previous experiment. Then I will interpret structural coupling.

Agent and environment as dynamical systems (Exp. 2)

Please confront this part with the section called “Agent and environment as dynamical systems (Exp. 1)”. In this section only the differences, which apply only to the neural controller, will be addressed.

A. `N` (Neural controller)

- *State space*: The state variables of this system are neuron activations (Exp. 1), plus *synaptic strengths* at every time step. The evolved agent has 5 neurons with two synapses on each. In my neurons with dynamic synapses, there is no inertia anymore and the evolved controllers also do not have recurrent connections. Nevertheless, the activations still have an effect on the next time step due to self-connections.

- *Map*:

  In the fixed parameters, `u`, we do not need connection weights anymore, as these are now in the state space (are changing dynamically). The neuron parameters added, however, are the *properties of* every neuron's *synapses*: Hebbian rule, learning rate of incoming synapses and whether the synapses leading from this neuron are excitatory or inhibitory.

  The variable parameters – inputs to the neurons from the body sensors, as represented by the output of the function `I` – remain the key variables influencing the dynamics of this system, whose internal dynamics is limited.

B. `B` (Body) – same as in Experiment 1.

C. `E` (Environment) – same as in Experiment 1.

Individual time scale (Exp. 2)

Let us look at the quantitative analysis of structural coupling, on individual time scale, again.

1. *Size of structure:* `N`'s state space consists of a couple of neuron activations plus synaptic strengths of every synapse, `B`'s of angles at articulations and states of sensors and `E`'s of positions and velocities of agents' and bodies' parts. The maps and parameters are complicated.

2. *Plasticity:* These values of state variables of the systems can change quickly: neuron activations, synaptic strengths, body angles and agent positions in the environment. The food positions change more slowly – only if food is eaten and a new is born. The maps and fixed parameters do not change. When compared to Experiment 1, the synaptic strengths dynamically change.

3. *Degree of mutual effect of systems on each other:* The synaptic strengths change on the basis of neuron activations which in turn are highly dependent on inputs from body. The

body's state variables are completely dependent on the neural controller's and environmental states. The position of food as well as that of the agents themselves in the environment is perturbed by the agent activities.

The contribution of Exp. 2 is an increase of the neural controller's structure and its plasticity on the individual time scale. The structure of the neural controller consists in the neurons' activations and strengths of incoming synapses (state space), and in the connectivity of the network, activation functions etc., as represented by the map and its parameters. Some of the parameters are variable, namely the inputs from sensors. The neuron activations and synaptic strengths dynamically change and these changes are to a large extent caused by the inputs coming from the body (which in turn come from the environment). The structure of the neural controller is therefore large and plastic, readily perturbed by the environment. Note also that the changes in structure affect the emerging behaviour of the agent. Nevertheless, note that the synapses change so quickly that this time scale is very close to the active dynamics, thus making a separation of the active dynamics and learning time scales difficult.

Evolutionary time scale (Exp. 2)

What are the structural changes in the agent? On the neural controller's part, there are changes in the size of the state space – neurons may be added or removed (thanks to crossover in Exp. 2) and the same is true for connections, thus changing the number of neuron activations and synaptic strengths, which constitute the state space. There are also important changes in `N`'s parameters, namely the synaptic properties, which are subject to evolution. These determine the synaptic plasticity on the individual time scale. The body is also evolved. The fact that evolution has removed one stick from the agent's body and corresponding sensors and muscles (as we saw in the section called “Results (Exp. 2)”) is a dramatic change in the structure of the body. Not only does this affect `B`'s state space but also the sensory function, `S`, and motor function, `M`. The parameters of body parts (like size, shape) are also evolved. The whole map is thus changed. The structure of the agent is thus large and plastic on this time scale. Although the changes by genetic operators are essentially random, over time only the useful changes expand. By useful we mean those that increase the fitness of the agent. However, as I have already suggested, the fitness involves the whole system that includes `N`, `B` and `E`. Therefore, for instance useful changes in `N`'s structure are those that produce a more efficient behaviour in interaction with `B` and `E`. And hence, in co-evolution, all the systems mutually perturbate the structure of each other.

Summary (Exp. 2) and overall interpretation

I have started with three different landscapes, put the same starting agent and food genotypes in each, and started co-evolution. After millions of time steps (hundreds of generations), we see a sensible behaviour of all the systems, different in every terrain though. The agents move in an efficient way to find food and the first food population has developed such morphologies that it is easily reached by agents, whereas the second food population successfully 'hides' from them. This allows for a biological interpretation – the first food population is in a symbiosis with the agents, whereas the second one is like prey (and the agent a predator).

This behaviour emerges from the interaction of the systems. The systems themselves have structures (constituent component and relationships between them) and these determine how they interact, how they perceive and act. The sensible behaviour corresponds to structural fit of the systems. And this structural congruence is a result of structural coupling. The systems' structures are plastic on evolutionary as well as individual time scales and perturb each other. The interplay of these has produced the fit of the structures. A behaviour observed over a short period is a result of a history of structural coupling of all the systems and of instantaneous structural coupling of the individuals observed at this very moment.

More specifically, we have an agent and environment. The agent can be further divided into neural controller and body. We have prepared the experiments, such that these are plastic on an evolutionary time scale and in Experiment 2, the neural controller has gained considerable plasticity in individual agents as well. The body and 'brain' get tightly coupled to each other, and to the environment. The environmental plasticity is in the food morphology and is thus significant on the evolutionary time scale only. However, the structural changes are large and are a result of the agents' activities.

We thus have a structural coupling model. When looking at it in the perspective of continual mutual adaptation of agent and environment on evolutionary and individual time scale, we are only missing a mutual adaptation on individual time scale, as the environment's structure is rather passive. However, this is in accordance with our view of a biotic part of environment.

5.2. Interesting consequences

This section will point out some of the interesting consequences that were not yet (explicitly) mentioned.

5.2.1. (A)symmetry of agent-environment interaction

We have seen in Section 1.2.2 that Maturana differentiates between structural coupling of two systems and structural coupling of an organism with its environment. Whereas the former was symmetrical, the latter involved only adaptation of the organism to its environment. However, I have treated structural coupling with the environment in a more symmetrical manner; the agent and environment mutually perturb each other. Quick et al. (1999) treat structural coupling with the environment in a way that is even more similar to structural coupling of two general systems. However, the symmetry has its limits.

The agent is in the environment. This has also showed up in the dynamical systems formalization. We had to somehow include the agent's position into the environmental state space. This asymmetry is also hidden in the embodiment definition of Quick et al. (1999), which I quoted in Section 1.3.2. “A system X is embodied in an environment E if perturbatory channels exist between the two.” However, what about the environment? Is it also embodied?

Thus the structural coupling of agent and environment is special, different than structural coupling of arbitrary systems, for instance of two agents. The agent and environment both perturb each other's structure, however, the changes in each differ. We do not expect the environment to be as adaptive or intelligent as the agents (although we still want to maintain that intelligence emerges from their interaction). This is a justification for the fact that in my model the environment is not very plastic on individual time scale. It is in accordance with biological reality.

5.2.2. Time scales

When analysing structural coupling, it is not enough to define the systems involved as some entities in space, for instance, and their interactions. In the previous section, we have used a neural controller, body and environment. However, it is inevitable to define also a *time scale of interest*. Plasticity of the systems is always with respect to a time scale. We can observe changes and perturbations on this time scale; other time scales are either too small (changes are too fast) or too large (changes too slow), so that no change is observed – it has moved behind the horizon. Variables are replaced by constants or average values. We can differentiate between time scales in the 'life' of a single agent. In connectionism, there is usually a time scale of changes in neuron activations and then a greater time scale on which connection weights change, what corresponds to learning. We track changes of the same dynamical system, but on different time scales. Nevertheless, in the PNN we have used the connection weights (or synaptic strengths) change almost as quickly as the activations, thus making such a separation difficult.

However, *evolution* of agents is a different case. *Changes* are not only on a much greater time scale, but they also *involve a different system*. They are gradual changes of the species genome. We can view this as gradual changes of a representative of the species. The traits of this representative will be random variables, with a probability distribution derived from the traits of all the individuals in the population. Now, we can observe also 'vertical' coupling, across time scales (Figure 5.3). The representative of a species determines, on a probabilistic basis, the traits (initial structure) of a newborn individual. On the other hand, the performance of the individuals determines the changes in the representative – through the performance of individuals and natural selection. The agent species thus gets coupled to environment *indirectly*.

The same holds for the biotic part of the environment, the food populations. Coupling of species does not occur directly, but rather through changes in individuals that are in direct contact.

Note that we usually do not think of embodiment on evolutionary time scale. We think of the body of an individual agent. On the other hand, in this thesis, the indirect, evolutionary structural coupling plays a major part. We can thus either restrict our notion of embodiment to structural coupling of individual organisms, or, we can think of plastic bodies shaped by evolution. I would suggest the latter, and propose a term '*as-if embodiment*' (or phylogenetic embodiment), for embodiment of a species (representative).

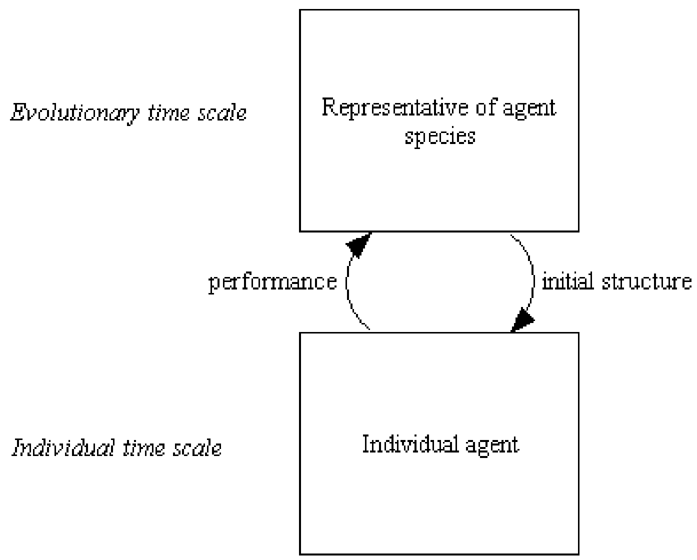

Figure 5.3. An agent on an evolutionary and individual time scale. On the evolutionary one, it is a representative of the species.

5.2.3. Dominance of evolutionary time scale

Structural coupling is best observed on an individual agent interacting with its environment and this interaction should be plastic, i.e. it should perturbate the structures of the systems involved. However, in this model, most of the plasticity is not on this time scale, but rather on the evolutionary one. This holds even more for many experiments with autonomous agents, when there is often a fixed body and an evolved controller, a form of a feed-forward network with evolved connection weights.

Does this mean that our model does not concentrate on the most important 'level' of coupling? I would not say so. In nature, plasticity involves many time scales: from evolution, through ontogenesis, individual adaptation (or 'learning') to active dynamics from moment to moment.

However, I argue that in simpler organisms, the 'biggest' structural coupling goes on in evolution. Every organism with its environment exhibit a structural fit. Nevertheless the structural changes in individuals are limited – up to insects, with some exceptions, animals are not able to learn. Although we have seen the adaptive behaviour in E. coli, involving complex structural changes, I would say, that most of the structural complexity has evolved and is not plastic in the lifetime of an individual bacterium.

Note, that our model agent is even simpler than E. coli or Hydra. Hence, it is justified that its structural plasticity is mainly evolutionary. Sophisticated adaptivity of individuals is usually connected with learning and genuine learning is present in very sophisticated animals, vertebrates. The road to them in artificial systems from bottom-up is, however, extremely long.

5.3. Evaluation

In this section, I will evaluate the model according to the criteria I have set forth in Section 1.4.2. It is an evolution of agents in environments consisting of abiotic and biotic parts. Thus, it has the essential features of co-evolution. This is a special case of structural coupling. The agent with a plastic neural controller also undergoes structural changes in an individual. The outcome or performance of the model is not very special, when compared to other work on artificial co-evolution, however, it is interpreted in a structural coupling perspective and reveals interesting consequences.

- *Compliance with principles*: This model complies with the principles of structural coupling. The agent and environment are both structurally plastic systems and they mutually perturb their structure, over two different time scales, and develop a structural fit. The principles are very general though, and can be fulfilled by trivial models too. However, this model is not a trivial one, but rather tries to catch the important features of structural coupling of organisms in their environments. It necessarily contains simplifications: metabolism is not modelled, body is not plastic in individual agents etc. This, however, does not reduce the compliance with principles.

- *Minimalism*: Minimalism was to a large extent sacrificed to increase the performance and heuristic value of the model. There are three populations co-evolved in different environments in a 3D world, the morphologies of the populations have many parameters and degrees of freedom. These are unnecessary constructs from the point of view of complying only with the basic principles, but were needed to produce some interesting and visible results.

- *Heuristic value*: I think that the model has furthered our understanding and revealed some interesting consequences. These are:

  - Co-evolution is a good example of structural coupling with the environment.
  - The model has shown how different abiotic parts of the environment affect the co-evolution of agents and a biotic part of it.
  - Structural coupling of agent and environment is not entirely symmetrical.
  - Structural coupling occurs on different time scales, directly or indirectly. There are many time scales involved on which the systems may change. Some of them may not be interesting from the structural coupling perspective, such as the shortest one, corresponding to active dynamics of the systems – their instantaneous behaviour.
  - Formalization using dynamical systems is fruitful. It has shown that the structure cannot be reduced to a system's state space, but rather has to include all parameters of the system. It has also helped to assess the degree of structural coupling and derive some of the interesting consequences.
  - Structural coupling of organisms with their environment is very complex. When building artificial systems, we have to simplify a lot. We have also discovered that in simpler organisms, most of adaptation goes on the evolutionary time scale. This is in accordance with artificial models created in the bottom-up way.

- *Quantifiability*: We have shown that dynamical systems theory can be used to formalize structural coupling of systems. This helped us in quantifying it. However, we have gone only halfway: although deriving some equations, these are only semi-formal and have only helped in our interpretation. We have not, for instance, measured the bandwidths of perturbatory channels. Nevertheless, this is justified in this heuristic model and we have shown a basis that can be further extended. Moreover, I have also found out that the structural coupling definitions are really general and it is necessary to find a way to assess not whether it is structural coupling, but rather what is its degree.

- *Performance*: The performance of the model is good – it has produced remarkable, visible structural fits of agents with their environments. It is a good platform, on which we can further our understanding of the phenomena modelled.

  However, the performance we have reached, when compared with experiments on co-evolution or autonomous agents experiments is in no sense outstanding. The agents are in fact very simple; although those with plastic synapses do allow an interpretation we were seeking, they do not perform better than those with fixed connection weights. I have also not analysed properly how the emerging behaviour comes about and did not perform enough runs to produce statistically more valid results. From this perspective, this model is really a prototype and can be extended in future work.

Chapter 6. Related work

This chapter presents a tour through related work to this thesis. However, the aim is not to present a comprehensive review, but rather to select some representatives of particular type of work and relate them to this thesis. In the first section, I will analyse models of structural coupling. In the second, I will relate my thesis to experiments with neural driven autonomous agents and to artificial creatures of Karl Sims (1994).

6.1. Models of structural coupling

The most relevant models to my thesis are those of structural coupling. However, there are very few of them in the literature. There are models of autopoiesis (e.g. McMullin and Varela[1], or Beer 2004). Varela et al. (1991) have presented a simple cellular automaton, Bittorio, as a model of structural coupling. Recently, Quick et al. (1999) and Quick et al. (2003) have defined embodiment in terms of structural coupling and modelled it.

6.1.1. Bittorio

Varela et al. (1991, p. 151-156) demonstrate structural coupling and enaction on Bittorio, a cellular automaton. It is put in an environment, a random milieu of 0s and 1s. When one of the cells encounters one of the two alternatives, the state of the cell is replaced by the perturbation it encountered (a 0 or 1). An interesting case occurs, if Bittorio finds itself in a quasiperiodic attractor (see Section A.1). Varela et al. describe Bittorio with a particular rule that is responsive to odd sequences of perturbations, by going from one to another spatiotemporal configuration (or attractor). Their interpretation is the following:

> ... given its rule and given its form of structural coupling, this Bittorio becomes an "odd sequence recognizer." (p. 152)
>
> ... over time this coupling selects or enacts from a world of randomness a domain of distinctions ("odd sequences" or "two successive perturbations") that has relevance for the structure of the system. In other words, on the basis of its autonomy the system selects or enacts a domain of significance. (p. 155-156)

In what ways does Bittorio differ from our model? Bittorio is really a minimalist model. However, it would not fit the notion of structural coupling with the environment, as was presented in this thesis. It is true that Bittorio is structurally determined: its structure determ-

[1] McMullin F. V., Varela F. J. (1997): Rediscovering computational autopoiesis. In: *Proceedings of the Fourth European Conference on Artificial Life* (MIT Press, Cambridge, MA), 38-47.

ines its response to perturbations. It can be perturbed by the environment, changing states and even the attractor. This is plasticity on a certain time scale. However, I would not say that it "becomes" a recognizer or that it "enacts" this significance, since the basic structure was given from the beginning and did not change any further. This would need plasticity on a greater time scale, such as evolving the rule of the automaton.

What is missing is a perturbatory channel from Bittorio to the environment. The environment is just a random milieu and thus we may say that it even lacks a structure and thus cannot go through structural changes to develop a fit with Bittorio. It is also questionable whether Bittorio has a body. Thus, Bittorio is a minimalist model showing structural determination of a system, but not a full-fledged model of structural coupling with the environment.

6.1.2. Tom Quick et al.

Model of E. coli

Quick et al. (1999) have analysed structural dynamics of the bacterium E. coli (see Section 2.2.1), and also developed a model based on its dynamics, called Phenomorph. This is a cellular automaton (CA) implementing a highly simplified version of the signalling pathway of the bacterium. Then, three different artificial sensory and motor surfaces were added, so that Phenomorph can be put in three different environments: World Wide Web, Abstract 2D parameter space, and Real world by means of putting the CA in a Khepera robot. The implementation was, however, only partially completed at the time of writing the paper and is not described in detail.

Nevertheless, it differs from the model developed in this thesis in the following aspects: (1) Phenomorph is a different type of model. It models, to a large extent, a natural system, a bacterium, while I have decided to develop a heuristic model, modelling only the general principles of structural coupling (see Section 1.4.1). (2) A different platform was chosen, a CA, as opposed to a neural driven agent. This has to do with the chosen level of modelling. (3) E. coli exhibits chemotactic behaviour and this is also expected from Phenomorph. This is a more realistic behaviour than only random jumps, as exhibited by the agents developed in my model. (4) Phenomorph resembles Bittorio in the way it was designed: there is a lot of significance from the beginning, coming from the bacterium. The coupling is then modelled on individual CAs. On the other hand, in my model, I focus on evolution and structural perturbations on evolutionary time scale. As the environment used by Quick et al. is not described in detail, I will not relate it to my model.

Embodied Genetic Regulatory Network-Driven Control Systems

This is a different platform, used by Quick et al. (2003) to investigate embodiment. Genetic

Regulatory Networks (GRNs) play a major part in the development of living organism and also in cellular metabolism of individual organisms in their lifetimes. There is a continual interaction of genes, protein machinery in cells, and environment. In this paper, Quick et al. use an experimental system, Biosys, to model artificial organisms controlled by such networks, coupling to an environment. A cell consists of genome and proteins. Proteins are produced by genes and act as 'messengers', regulating other genes (through binding to regulatory sites) and producing effector output. Sensory input is also facilitated by them – via proteins introduced into the cell (Figure 6.1).

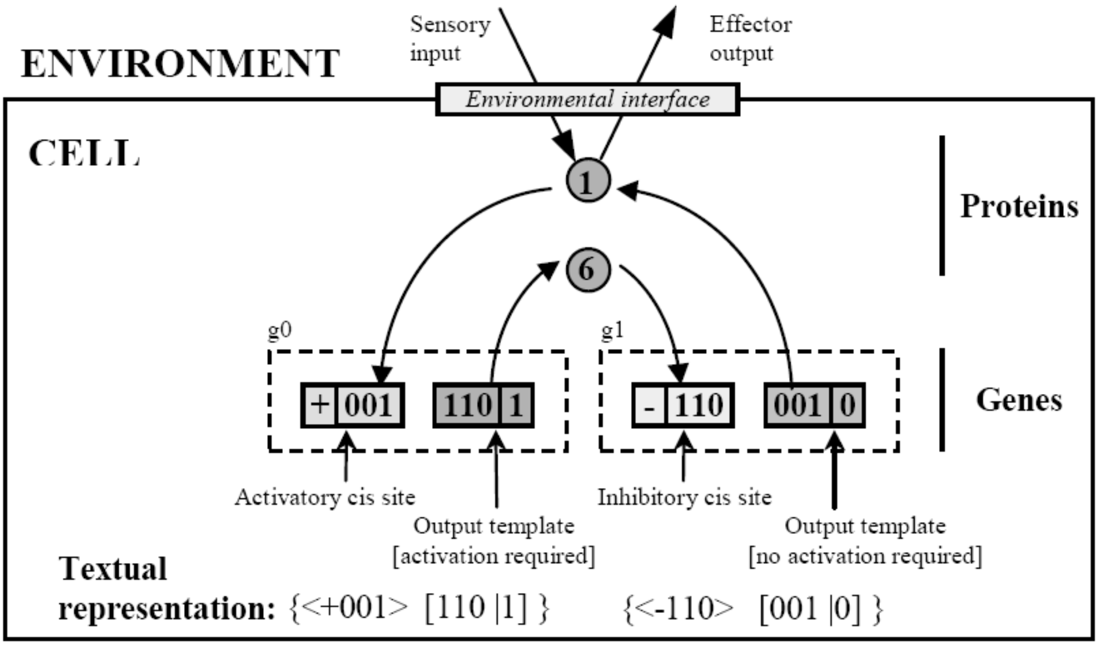

Figure 6.1. Biosys model. Schematic gene-protein-environment interaction, with a textual representation of genes illustrated (from Quick et al. 2003).

This model involves two time scales. Some parameters are evolved using a genetic algorithm, namely properties of regulatory sites (whether they are activatory or inhibitory, their matching protein type, their output protein type) and global parameters (protein decay rates, binding proportion, saturation value). The evolved networks are used as controllers in two different experiments:

a. in an abstract software environment where the controller acts as a thermostat with a heating unit

b. in real world, controlling a Khepera robot that should maximize the amount of encountered light

Let me relate this to the model developed in my work. The greatest difference is in the level of modelling chosen. Quick et al. model a cell and its structure, whereas I have chosen a higher level. On the level chosen by Quick et al., the emerging homeostatic behaviour in experiment a) involves complex structural coupling. This is in contrast to thermostat and heater as a trivial example of structural coupling on a higher level of abstraction.

The model of Quick et al. is certainly much more biologically realistic. Mapping between genotype and phenotype is at least partially realistic; moreover, the genes play a part in a mature individual as well. This contrasts with the direct genotype encoding in my experiments which is very far from biological reality and also prevents this coupling of genome with other parts that constitute a phenotype.

In Quick's model a), the behaviour that is observed on individual time scale is a good model of structural coupling with the environment. The organism and environment perturb the structure of each other via proteins. The environment is perturbed in such a way that its temperature changes. Conversely, proteins from the environment perturb the organism's structure and change its outputs in the future. This is a more intuitive model than the one developed in my thesis – an agent jumping in environment with food items.

However, I would see also strengths in my model. The structural fit in a 3D world can be better seen and explained. I have also evolved the environment and reached what may be best described as co-evolution; the environment also has a complex structure evolved by a genetic algorithm. On the evolutionary time scale, there is thus mutual perturbation of the systems and that is also visible in the 3D scene – on the environment's part, this is the morphology of food items. In my experiments, the agents are also not just control systems but do have a body, engaged in structural coupling, as well. I would say that models on various levels have different heuristic values.

6.2. Other models or experiments

6.2.1. Experiments with neural driven autonomous agents

There is a large body of work on autonomous agents and robots. The aim is usually to achieve some adaptive behaviour of the agents that may be interesting from a scientific as well as engineering perspective. The agents are often equipped with artificial neural controllers and these usually receive the most attention. If we look at the experiments from a structural coupling perspective, then we realize that from the triad brain-body-environment, body and environment are very often fixed and pre-designed (such as a Khepera robot in an arena) and their structures are not plastic. The 'brain', which may have more or less plasticity and be more or less pre-designed (see Section 3.2.1), gets coupled to the other two. Thus they are not what I would describe as structural coupling with environment.

Over time, it was found out that this imbalance may in fact hinder the development even from an engineering perspective. Thus in the last ten or fifteen years there were also experiments involving evolution of robot morphology. Some researchers have also used co-evolution of agents, typically predator and prey experiments, to stimulate evolution.

These efforts are summarized and further explored by Buason et al. (2005). They conduct a series of predator and prey experiments with robots, evolving their neural controllers and morphology. Moreover, they recognize that the environment is usually left out from evolution and introduce plasticity in the environment as well, by letting either the predator or the prey evolve the size of the arena, or position of some objects in it. Such experiments can be interpreted as structural coupling with environment.

Let us compare the experiment with mine. The fact that in the work of Buason et al. the agents evolve the environment to suit their needs is not realistic from a biological point of view. Also, the plasticity of bodies is very limited: Khepera robots are used and only camera view angles, range and robot speed are evolved. On the other hand, in the experiments of mine, the body has many degrees of freedom, subject to evolution. Also, the biotic part of environment has its own fitness and this results in a more realistic co-evolution. The plastic neural network also allows for more interesting structural plasticity than the neural controllers used by Buason et al.

At last, please note that in Buason et al. there are two agent populations – predator and prey. They have neural controllers and bodies and then there is an environment, an arena. On the other hand, I have one agent population and an environment consisting of a biotic and abiotic part. This division is arbitrary and my Food 2 population can easily be viewed as prey, 'trying to hide' from the agents. The abiotic part would then correspond to the arena. Nevertheless, my division is biologically motivated, as discussed in Section 2.1. In fact, in the experiments of Buason et al., there is a more interesting structural coupling between the predator and the prey, than between the agents and environment.

6.2.2. Sims' artificial creatures

Karl Sims (1994) has also used a simulated 3D physical world to evolve entire creatures: their control systems and morphologies. This was tried in different environments, with different physical laws. The 'brains' and bodies were evolved together and thus got structurally coupled to each other and to the environments. This is similar to my experiments. The importance of the visual performance of the result is also similar. Sims' evolved creatures are similar to those developed in Framsticks or to mine, but I must admit that they are better. They beautifully show how the creatures have developed a structural fit in very different environments.

Nevertheless, apart from some technical differences (genotype encoding, generational evolution, Sims' 'higher neurons'), there are also differences related to a possible structural coupling interpretation. Most importantly, in Sims' models, the environment is not plastic. He manually changes the parameters of the environment (as I change the abiotic part of the environment – the landscape) and then lets the creatures get coupled to it. On the other hand, my model contains also a biotic part of environment that gets co-evolved with the agents. Whereas Sims uses a more specific fitness, optimizing the creatures for a specific behaviour (sometimes even choosing the most aesthetically desirable ones), I use lifespan, a more general and biologically plausible one. On the other hand, Sims' starting creature genotypes are not hand-designed at all, which is better: there is more room left for self-organization. To conclude, Sims' work is more directed to the goal of evolving aesthetically impressive creatures than to creating a good model for scientific interpretations.

Chapter 7. Future development

The model developed in this thesis is a prototype. It does have a good heuristic value, but it can be extended in many ways. Some of them would, however, not be feasible in the current version of Framsticks, the software environment.

The agent developed in Experiment 2 has a neural controller with plastic synapses. However, it does not perform better than the one with fixed weights. Thus this plasticity of the agent has a value for our interpretation but not for the agent. If a better performance could not be reached by a longer evolution or adjustment of genetic algorithm parameters, then it may be that such plasticity is not needed. Then it would be good to modify the experiment in such a way that plastic synapses would be an advantage. It would also be much better if the agent had a way to perceive the food. Then it could move to it, in a form of chemotactic behaviour. This could be achieved by means of smell sensors, which are available in Framsticks. The agent could also get a feedback from the food eaten – its energy.

The direct genotype encoding used is neither biologically plausible nor ideal for structural coupling. A developmental encoding that would enable ontogenesis of the agent, in interaction with the environment, would be ideal. Framsticks allows for a developmental encoding which may be tried, but the decoding does not take environment into account. Models where genes would play an active part also in the agents' life are not supported by the software environment. However, if we were not bound to Framsticks, then it would be interesting to use a more realistic mapping between genotype and phenotype, making structural coupling more complex. A good one is used in Quick et al. (2003, see Section 6.1.2) and it can be also utilized to grow neural networks.

The method of evolution used is a directed one. The simulator maintains a constant number of agents and food items in the world. If they die, new ones are created from genotypes and put into the world. This allows for a stable simulation and evolution. However, it would more realistic to use a spontaneous evolution, which can also be done in Framsticks. The agents could reproduce on acquiring a certain level of energy. Food 2 population could reproduce at a certain age. The 'symbiosis' of agents with Food 1 could also be modelled. This food may reproduce when eaten by the agent; the agent may for instance drop the 'seeds' after a certain time. This would also increase the structural fit; the number and position of agents and food will be a result of their interaction. For instance, Food 2 will accumulate in those regions that are hard to access by agents.

As the last extension I would offer an experimental setup that I intended originally and that I also partially implemented. It was more ambitious. However, it could not be realised because of Framsticks' limitations. I intended the agents to use feelers with touch sensors to feel the shape of the food. The problem was that the simulation of the world is not accurate enough to model it. Nevertheless, should the simulator be more precise in the future, or should a different software environment be used, my idea was the following: The agent will have for instance two feelers that would not touch the ground and will not be used for movement, but rather would detect food items. The agent will not automatically eat food on collision, but rather would eat it only by means of a certain part – its 'mouth'. The feelers will be moving and able to put some food into the mouth and some other food draw aside. Then, similarly to our experiment, two food populations can be co-evolved with the agent's body and 'brain'. Food 1 population will again 'want' to get eaten, Food 2 will not. Moreover, this may be also 'poisonous' and the agent would lose energy on eating it. The agent could be able to sense it by a special neuron with activation proportional to energy of food just eaten. The emerging behaviour would be the agent moving in the environment, putting Food 1 items into his mouth and Food 2 items drawing aside. This would be a more intimate structural congruence – one food fitting into the agent's mouth, one not. It could also be extended to a more difficult version, demanding an ability to eat both types of food and reinforcement-learning from the agent. The energies of food can be switched and the agent should learn in his lifetime from the reinforcement neuron (energy of food eaten), which food is good and which is poisonous.

Chapter 8. Conclusion

I have brought together the work in AI with the work of Maturana and Varela and studied a unifying concept that underlies interaction of organisms with environment and their intelligence or adaptive behaviour. The concept resides in this very interaction. It is embodiment or structural coupling. In this thesis, I do not see the essence of embodiment in the material domain, but rather in the relationships between an organism and its environment, thus allowing me to equate it with structural coupling, a relational construct. The relationships are valid for natural as well as for artificial systems, and for systems in the physical space as well as for virtual ones. This justifies the development of a nonmaterial model in software. The idea of understanding by building is the answer to the question why such a model is useful: I want to uncover the general principles of this phenomenon by building its model. The goal is thus not the final model as such (as it is not an application), but the road to it and the interpretation of the model. The analysis and interpretation should help us to discover new, interesting consequences. This is also the case of this model. The strength or contribution of this work lies above all in theoretical analysis, with the structural coupling model at its centre.

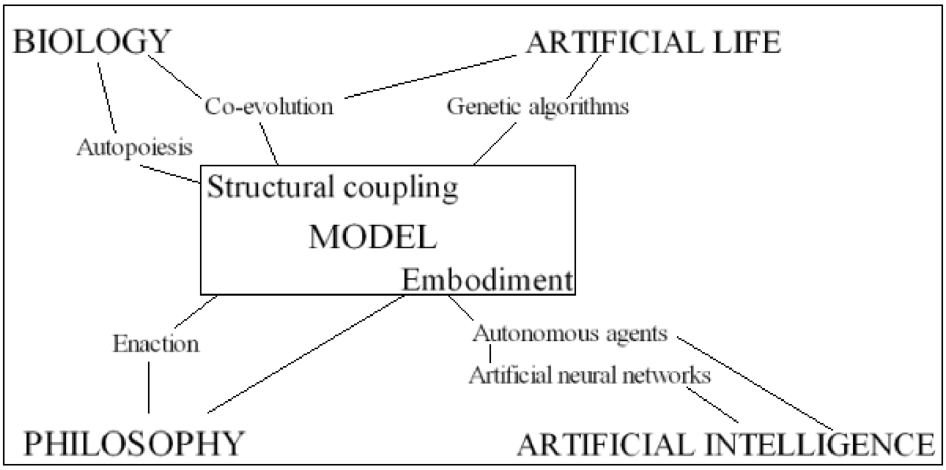

Figure 8.1. Transdisciplinary nature of structural coupling modelling. This is only a very schematic diagram of the fields and disciplines involved. All the consequences cannot fit into a 2D diagram – it is thus greatly simplified. For instance, the dynamical systems perspective that also plays a part is not depicted in the figure. Also, I rather did not draw any circles around the disciplines, because there are no distinct boundaries.

This work presents an extensive *analysis of a transdisciplinary area* (Figure 8.1). Putting the terms and concepts from diverse fields together is one of the contributions of this work. Starting from the idea of Quick et al. (1999) by bringing together embodiment and structural coupling, this work also relates structural coupling to biological terms adaptation and co-evolution. In Section 2.1, I have proposed that in biological terms structural coupling could be described as continual mutual adaptation of organism and environment, on individual and evolutionary time scale, and that co-evolution is an interesting instantiation of structural coupling. The meaning of the term structure is analysed on examples from nature. In Section 6.2 I have also analysed the work from the research field of intelligent autonomous agents in the context of my work.

This model differs from the few other structural coupling models that are in literature mainly in two aspects. First, the *level of modelling* chosen in this model is higher: a *neural driven agent with a body*, as opposed to modelling on cellular automata or genetic regulatory networks. I concentrate on modelling sensory and motor surfaces of the body, not bodily metabolism. Section 3.2 contains a detailed analysis how the neural controller, body and environment should be modelled at this level of analogy. Recall that we claim that the phenomenon modelled is central to intelligence and hence such an analysis may be exploited in other, also application domains. From this perspective, the level of modelling chosen is an advantage, as in the research on autonomous agents and robots we often use neural controllers and are not interested in bodily metabolism. The review and analysis of artificial neural controllers is one contribution to the field of neural driven agents.[1] The other is that it is fruitful to concentrate not only on the agent neural controller, but on the body and environment as well. Another advantage of the higher level of modelling is that the results have a better visualisation, adding to the heuristic value of the model.

The other major difference when compared to other structural coupling models (Section 6.1) is that the *environment plays a more important role* in my model. It is divided into an abiotic and biotic part where the latter one is evolved together with the agents. Thus, the model can be best described as *co-evolution*. Structural coupling is more symmetrical in such a model. The structure of the environment is perturbed in such a way that, over time, it adds to a structural fit with the agents.

Contribution of this work is also in a step forward toward *formalizing and quantifying structural coupling*. The structural coupling concept is very inspirational but it is not formalized enough to describe interacting systems in a more exact way, depicting the relationships more precisely. The definitions of structural coupling are also very general and minimalist. I have outlined a solution to these problems by supplementing the language of Maturana and Varela by dynamical systems theory, which provides a more formal mathematical structure. The dy-

[1] Analysis of possible neural substrates is sketched by Nolfi, Marocco (2001), however, in the end it is concluded that "unfortunately, a systematic comparison of the advantages and disadvantages of each neural network type has not been carried out yet." (p. 8) A step forward is the analysis of CTRNNs and PNNs in Blynel, Floreano (2002).

namical systems perspective on agent-environment interaction of R. Beer (1995a) was utilised and the triad neural controller-body-environment was expressed as three nonautonomous discrete dynamical systems. This allowed me to propose a structural coupling quantification on my model. A method suggested by Quick et al. (1999) was utilised and the notion of structure in this domain expanded to include not just the state space of a dynamical system, but also the dynamical law and system parameters, which are equally important for the relationships between the system's constituent components. These parameters can be also plastic, on an evolutionary time scale. Studying the size of structure, plasticity and mutual effect of the systems on each other enables us to assess the degree of structural coupling.

The model enabled me to draw a number of *interesting consequences*. First, the *(a)symmetry of the agent-environment relationship* was analysed. It was concluded that there is always asymmetry involved. Roughly speaking, the agent is in the environment and hence the two systems cannot be at the same level. Second, the formalization of my model enabled me to shed more light on *different time scales* of structural coupling. On an individual time scale, the structure of agents is directly perturbed by the environment, while on an evolutionary time scale, different parameters of the systems are perturbed, in an indirect manner. However, the situation is more complicated: we can identify more time scales of changes and assess their relevance for structural coupling. I have also proposed a term *as-if embodiment* for embodiment of a species representative. The evolutionary time scale is dominant in my model. Third, analysis of simple organisms as well as of models involving autonomous agents created in the bottom-up way has lead to an observation that most of the adaptive behaviour of the organisms or agents, arising from a structural congruence with their environment, is shaped by evolution. Individual adaptation is a complement only, which becomes more important in complex organisms.

This model is a model of structural coupling on the level of sensorimotor co-ordination; a neural driven agent co-evolves with its environment. Some of the qualities of minimalism and biological plausibility were sacrificed in favour of a visible performance in 3D visualisation, in order to increase the heuristic value of the model. Hence, video sequences from this model, accompanied by the interpretations suggested here, could be used to show, teach, investigate or study structural coupling, embodiment, or intelligence in the sense of surviving in a complex world.

Bibliography

Beer R. D. (1995a): A Dynamical Systems Perspective on Agent-Environment Interaction. *Artificial Intelligence* **72**, 173-215.

Beer R. D. (1995b): On the Dynamics of Small Continuous-Time Recurrent Neural Networks. *Adaptive Behavior* **3 (4)**, 469-509.

Beer R. D. (2004): Autopoiesis and Cognition in the Game of Life. *Artificial Life Journal* **10 (3)**, 309-326.

Blynel J., Floreano D. (2002): Levels of Dynamics and Adaptive Behavior in Evolutionary Neural Controllers. In: *From Animals to Animats: Proceedings of the Seventh International Conference on Simulation of Adaptive Behavior* (MIT Press, Cambridge, MA, 2002), 272-281.

Brooks R. A. (1991): Intelligence Without Reason. In: *Proceedings of the 12th International Joint Conference on Artificial Intelligence, Sydney, Australia*, 569–595.

Buason G., Bergfeldt N., Ziemke T. (2005): Brains, Bodies and Beyond: Competitive Co-Evolution of Robot Controllers, Morphologies and Environments. *Genetic Programming and Evolvable Machines* **6 (1)**, 25-51.

Flegr J. (2005): Evoluční biologie. Academia, Praha.

Floreano D., Mondada F. (1996): Evolution of Plastic Neurocontrollers for Situated Agents. In: *From Animals to Animats: Proceedings of the Fourth International Conference on Simulation of Adaptive Behavior* (MIT Press, Cambridge, MA, 1996).

Floreano D., Urzelai J. (2000): Evolutionary Robots with On-line Self-Organization and Behavioral Fitness. *Neural Networks* **13**, 431-443.

Havel I. M. (2001): Přirozené a umělé myšlení jako filozofický problem. In: *Mařík V., Štěpánková O., Lažanský J. et al.: Umělá inteligence (3)* (Academia, Praha, 2001), 17-75.

Hoffmann M. (2006a): Patří Čapkovi Roboti stále do říše sci-fi? *Vesmír* **5 (2006)**, 264-267.

Hoffmann M. (2006b): Pohled na agenty řízené neuronovými sítěmi skrze teorii dynamických systémů. In: *Kognice a umělý život VI, Proceedings of 6th Czech-Slovak colloquium* (Slezská univerzita v Opavě, Opava, 2006), 181-188.

Johnson M., Rohrer T. (forthcoming 2006): We Are Live Creatures: Embodiment, American

Pragmatism, and the Cognitive Organism. In: *Body, Language, and Mind, vol.1* (Mouton de Gruyter, Berlin).

Komosinski M., Ulatowski S. (2006): Framsticks website and Framsticks Manual. URL: http://www.frams.alife.pl

Komosinski M., Rotaru-Varga A. (2001): Comparison of Different Genotype Encodings for Simulated 3D Agents. *Artificial Life Journal* **7 (4)**, 395-418.

Maturana H. R. (1970): Biology of Cognition. In: *Maturana H. R., Varela F. J.: Autopoiesis and Cognition: The Realization of the Living* (D. Reidel Publishing Co., Dordecht, 1980), 5-58.

Mingers J. (1990): The Philosophical Implications of Maturana's Cognitive Theories. *Systems Practice* **3 (6)**, 569-584.

Nolfi S., Marocco D. (2001): Evolving robots able to integrate sensory-motor information over time. *Theory in Biosciences* **120**, 287-310.

Nolfi S., Parisi D. (1995): Learning to Adapt to Changing Environments in Evolving Neural Networks. Technical Report 95-15, Institute of Psychology, National Research Council, Rome, Italy.

Pfeifer R., Scheier C. (2001): Understanding Intelligence. MIT Press, Cambridge, MA.

Prescott T. J. (2005): Forced Moves or Good Tricks in Design Space? Great Moments in the Evolution of the Neural Substrate for Action Selection. In: *Modeling Natural Action Selection: Proceedings of an International Workshop Edinburgh, UK* (AISB Press, 2005), 36-44.

Quick T., Dautenhahn K., Nehaniv C. L., Roberts G. (1999): The Essence of Embodiment: A Framework for Understanding and Exploiting Structural Coupling Between System and Environment. In: *Proceedings of the Third International Conference on Computing Anticipatory Systems, Liège, Belgium, 1999*, 649-660.

Quick T., Dautenhahn K., Nehaniv C. L., Roberts G. (2003): Evolving Embodied Genetic Regulatory Network-driven Control Systems. In: *Proceedings of the Seventh European Conference on Artificial Life, Dortmund, Germany, 2003*, 266-277.

Quick T. (2006): Tom Quick's PhD page, sections Research and Autopoiesis. URL: http://www.cs.ucl.ac.uk/staff/t.quick/

Sims K. (1994): Evolving Virtual Creatures. *Computer Graphics* **28**, 15-34.

Varela F. J., Thompson E., Rosch E. (1991): The Embodied Mind. MIT Press, Cambridge, MA, London, England.

Whitaker R. (2001): The Observer Web: Autopoiesis and Enaction. URL: http://www.enolagaia.com/AT.html

Wiedermann J. (2006): Chtěli byste být mozkem v baňce, anebo o důležitosti vtělenosti a situovanosti při rozvoji mentálních schopností kognitivního agenta. In: *Kognice a umělý život VI, Proceedings of 6th Czech-Slovak colloquium* (Slezská univerzita v Opavě, Opava, 2006), 421-428.

Wikipedia contributors: Wikipedia, The Free Encyclopedia. URL: http://en.wikipedia.org

Yamauchi B. M., Beer R. D. (1994): Sequential Behavior and Learning in Evolved Dynamical Neural Networks. *Adaptive Behavior* **2 (3)**, 219-246.

Ziemke T. (2001): Remembering How to Behave: Recurrent Neural Networks for Adaptive Robot Behavior. In: *Medsker L., Jain L.: Recurrent Neural Networks: Design and Applications* (CRC Press, Boca Raton, 2001).

Ziemke T. (2003): What's That Thing Called Embodiment? In: *Proceedings of the 25th Annual Meeting of the Cognitive Science Society* (Lawrence Erlbaum, 2003), 1305-1310.

Appendix A. Dynamical systems perspective

This appendix presents the very basics of dynamical systems theory and the framework of R. Beer (1995a), who uses this perspective to characterize agent-environment interaction. This framework stems from the perspective of Maturana and Varela, autopoiesis and structural coupling, but employs the language of dynamical systems theory, since it provides a more mature mathematical structure. Thus it can be viewed as an extension of the work of Maturana and Varela (which is a more philosophical framework), providing a formal mathematical structure.

A.1. Dynamical systems theory

This section presents a very brief summary of relevant concepts and terminology of dynamical systems theory, according to Beer (1995a) but modified to concentrate on the discrete case.

A *dynamical system* is characterized by:

- a set of *state variables* **x** (discrete or continuous)
- a *dynamical law* `F` that governs how the values of state variables change with time

Set of all possible combinations of values of the state variables constitute the system's *state space* (also called the phase space). A dynamical system is *autonomous* if `F` only depends upon values of state variables and a on set of fixed *parameters* **u**. The dynamical system can be either *discrete-time* or *continuous-time*. In the former case the dynamical law is a map from current state to the next state. A simple example is an iterated map (Equation A.1). In the latter case, we have a set of differential equations and the dynamical law defines a *vector field* on the state space. A dynamical system is *linear* or *nonlinear*, depending on the linearity of `F` in the state variables. While Beer concentrates on continuous dynamics, I have chosen discrete-time dynamics for my model and I will further describe the theory for the discrete case only.

Equation A.1. Example of a discrete dynamical system: iterated map.

$$x_{n+1} = \mu x_n (1 - x_n)$$

From every initial state of the system, the dynamical law generates a sequence of states, a *trajectory* of the system. If the values of state variables do not diverge to infinity, the system has convergent dynamics and eventually converges to a *limit set*. This is a set of points invariant with respect to the dynamical law (the system will remain there). We will be most interested in *stable* limit sets. These are known as *attractors* and have an additional property of 'attracting' nearby trajectories, which converge to the attractor. Attractors govern the long-term behaviour of a system. The open set of initial states converging to a given attractor is called its *basin of attraction*. The simplest attractors are *equilibrium points* – single points invariant with respect to the dynamical law – and *periodic attractors* or *limit cycles* – closed trajectories in the state space which are stable limit sets. More complex are *quasiperiodic attractors*, formed from the sum of periodic solutions with incommensurate periods. *Chaotic attractors* have an apparently random behaviour. A dynamical system may contain many attractors, each with a basin of attraction around it.

The 'boundaries' between the basins are called separatrices. They create a cellular structure – the system's *phase portrait*. For minor changes in the parameters, the phase portrait ('the cellular diagram') changes only slightly. Such a system is said to be structurally stable. However, at some parameter values, a small change may produce a qualitative change in the phase portrait, a *bifurcation*. A *nonautonomous dynamical system* is one in which some parameters are allowed to vary in time. These may correspond to inputs of the system and cause the phase portrait to change in time, perhaps even drastically.

A.2. Beer's framework

Beer (1995a, p. 180-186) models the agent and its environment in constant interaction as two coupled nonautonomous dynamical systems. Some parameters of the system `A` (agent) are a function of the environment's state variables and some parameters of the system `E` (environment) are a function of the agent's state variables. We obtain a sensory function `S` from environmental state variables to agent parameters and a motor function `M` from agent state variables to environmental parameters, giving the following equations (Equation A.2) and a schematic diagram (Figure A.1).

Equation A.2. Agent and environment as coupled discrete dynamical systems.

$$\mathbf{X}_{n+1}^{A} = \mathrm{F}^{A}(\mathbf{X}_{n}^{A}; S(\mathbf{X}_{n}^{E}); \mathbf{u}^{A})$$
$$\mathbf{X}_{n+1}^{E} = \mathrm{F}^{E}(\mathbf{X}_{n}^{E}; M(\mathbf{X}_{n}^{A}); \mathbf{u}^{E})$$

where

- $\mathbf{X}_n^A$, $\mathbf{X}_n^E$ are vectors of values of state variables of systems A and E, respectively, at time step n.
- F^A, F^E are the maps or dynamical laws.
- $S(\mathbf{X}_n^E)$, $M(\mathbf{X}_n^A)$, are the sensory and motor function respectively, variable parameters of the systems.
- $\mathbf{u}^A$, $\mathbf{u}^E$ are any remaining system parameters. They are fixed, do not participate in coupling.

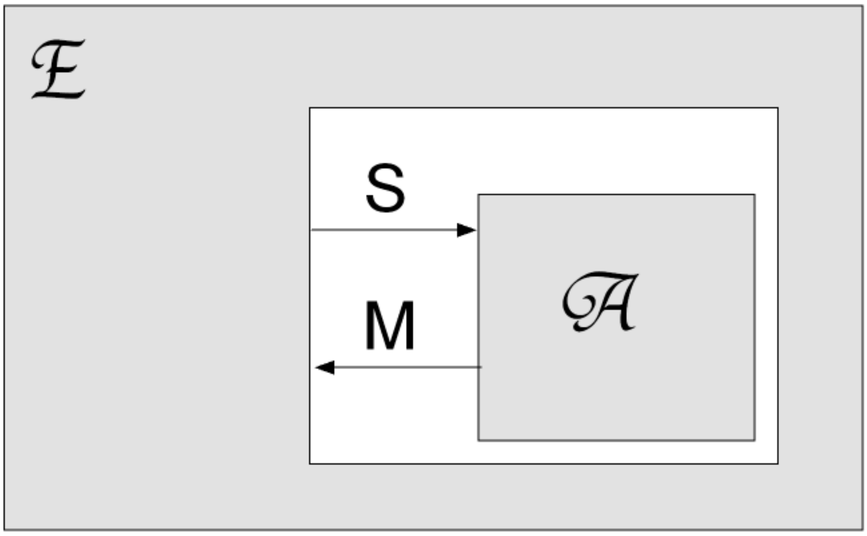

Figure A.1. Agent and its environment as coupled dynamical systems (redrawn from Beer 1995a).

Complementary to this view, the two coupled nonautonomous dynamical systems can be viewed as a single autonomous dynamical system U, whose state variables are the union of state variables of A and E and whose dynamical laws are all of the internal relations (now including S and M) among the new set of state variables.