School of Humanities and Informatics
Masters Dissertation in Computer Science, 25 Credits
D-level
Spring 2006

How do multiple behaviours affect the process of competitive co-evolution?

An experimental study

Anders Roxell

How do multiple behaviours affect the process of competitive co-evolution?

An experimental study

Submitted by Anders Roxell to the University of Skövde, as a dissertation for the degree of Master of Science (M.Sc.) by examination and dissertation at the School of Humanities and Informatics.

August 30, 2006

I certify that all material in this dissertation which is not my own work has been clearly identified and that no material is included for which a degree has previously been conferred on me.

Signed: ______________________________

Supervisor: Nicklas Bergfeldt
Examiner: Tom Ziemke

How do multiple behaviours affect the process of competitive co-evolution?

An experimental study

Anders Roxell

Abstract

In evolutionary robotics there has been research about the pursuit problem with different numbers of predators and prey: *(i)* one predator and one prey, *(ii)* many predators against one prey, and *(iii)* many predators against many prey. However, these different experiments are only involving food chains with two populations (two trophic levels). This dissertation uses three trophic levels to investigate if individuals in the middle trophic level perform equally or better than those that are been evolved in a two trophic level environment.

The investigation was done in a simulator called YAKS. A statistical analysis was conducted to evalutate the results. The result indicated that a robot with two tasks gets better at hunting and evading than robots with one task (either hunt or evade). Robots from the middle trophic level that are moving in the same direction as the camera is facing, were the best predators and prey. This dissertation is a step towards more complex and animal-like behaviours of robots.

Keywords: Evolutionary Robotics (ER), the pursuit problem, trophic level, behaviour

Acknowledgements

I would like to thank my friends that have reviewed my ongoing work and those I could discuss my dissertation with. Then I want to thank my family and specially my girlfriend that they put up with me and for their support, last but not least I want to thank my supervisor for all his comments during the entire dissertation.

Contents

List of Figures

List of Tables

1 Introduction

Today, research in robotics is striving towards more and more complex behaviour while having as little human intervention as possible. This is partly why researchers are using evolutionary techniques to produce a useful behaviour for artificial neural network robot controllers. When using evolutionary techniques it is easier for engineers to let neural networks produce a useful behaviour without explicit guidance. However, this comes at a cost, thus the time it takes for the neural network to establish with a behaviour that actually solves the problem that the engineer wants to solve (Nolfi & Floreano 2000).

To evolve a network controller that solves the task can be very time consuming. A way to speed up the evolutionary process is to use competitive co-evolution, which means using two (or more) competing populations with coupled fitness (i.e., the fitness of one population depends on the fitness of the other population(s)). This speeds up the evolution because the different populations can incrementally drive each other forward; a process called an arms-race (Nolfi & Floreano 2000).

Competitive co-evolution where an individual from one of the populations is supposed to catch an individual from a competing population is often referred to as the *pursuit problem* (Nolfi & Floreano 2000). Different variations of the pursuit problem have been investigated throughout the literature (Haynes, Wainwright & Sen 1995, Nishimura & Ikegami 1997, Nolfi & Floreano 1998). Some of the pursuit scenarios that have been investigated can be classified into three different groups: *(i)* a single predator against a single prey (Nolfi & Floreano 1998), *(ii)* multiple predators against a single prey (Haynes et al. 1995), or *(iii)* multiple predators against multiple prey (Nishimura & Ikegami 1997).

Evolution in robotics is based upon the concept that the fitness landscape depends on an individual's ability to reach a given goal (e.g., for a predator to catch its prey). All of the above given examples of the pursuit problems are based upon two trophic levels (i.e., when one population tries to catch the other, and the other tries to escape). However, in biology there can be more than two trophic levels; actually, most food chains have between four or five links (Abrams & Roth 1994), which means that the individuals from the populations in the middle trophic levels have to act as both predator and prey. This in turn leads to the orientation of this dissertation, which will focus on trophic levels which have three levels. The robots in the middle trophic level will have more than one task, that is, they should both stay alive (by evading their predator) and catch their prey. Thus, one can argue that these robots will have to have two different behaviours. Other researchers who have conducted experiments with robots having more than one behaviour including Nolfi (1997) and Floreano & Mondada (1996).

In this dissertation we will investigate how good the robots from the middle trophic level become compared to the robots that are just hunters or prey. The work conducted in this dissertation was inspired by Búason (2002) who outlined this as potential future work and by contemporary research that is striving to evolve more and more complex behaviour.

2 Background

This section provides a background of the concepts that are used throughout this dissertation. First, Section 2.1 presents the concept of a food chain and the basics of the eye placement of an animal. This is followed by an overview of evolutionary robotics in a wide perspective in Section 2.2. Finally, Section 2.3 gives an introduction to work that is related to this dissertation.

2.1 Biology

In Section 2.1.1 an introduction to trophic levels is be given, followed by a description of the eye placement on animals in Section 2.1.2.

2.1.1 Trophic level

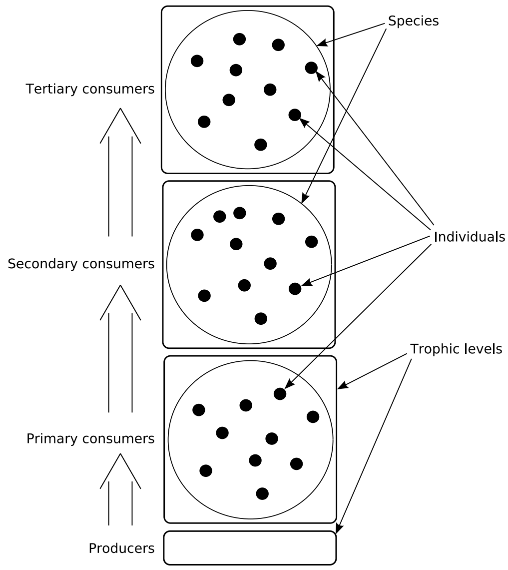

Figure 1: Trophic levels, species and individuals.

A trophic level contains one or many species, and each species contains one or many individuals. Figure 1 shows three different species and within every species there are a number of individuals.

Most of the food chains in nature are not longer than four or five trophic levels (see, for instance, Abrams & Roth 1994, Campbell & Reece 2002, Stauffer, Kunwar & Chowdhury 2005). Normally there is not one food chain, there are several of them. When food chains are connected to each other they form a food web (Campbell & Reece 2002) (depicted in Figure 2). One example of a food chain is when you have a plant on the bottom level (primary producer), next level are the herbivores (primary consumer) that eat the plant, such as a cricket. The next level in this food chain are the carnivores (secondary consumer, e.g., rodents) that consume the primary consumers (e.g., the cricket). On the fourth level there are carnivores as well (tertiary consumers). These could for instance be a snake that eats the rodent (Campbell & Reece 2002).

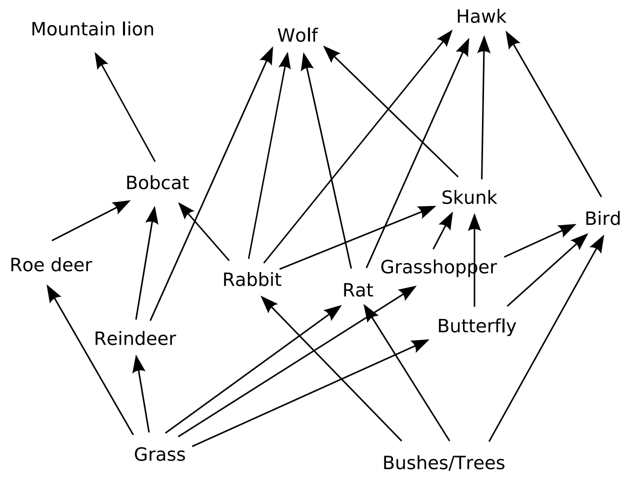

Figure 2: A food web (adapted from Campbell & Reece (2002)).

Campbell & Reece (2002) define a food chain as:

> "The pathway along which food is transferred from trophic level to trophic level, beginning with producers." (Campbell & Reece 2002, glossary)

Further they define the trophic level as:

> "Any of the several levels of a food chain whose species are based on their main nutritional source. The trophic level that ultimately supports all others consists of autotrophs, or primary producers." (Campbell & Reece 2002, glossary)

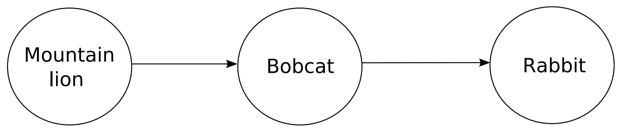

Figure 3: A pursuit scenario.

In Figure 3, the arrows between the mountain lion circle and the bobcat circle, and between the bobcat circle and the rabbit circle are called a food chain. The different circles stand for different trophic levels.

2.1.2 The animal eye

As vision is important for most animals one could assume that eyes of the animal in the middle trophic level must be able to support the behaviour of both predator and prey. Animals utilize their eyes for many different purposes, detecting a prey or a predator is just one. Thus, an animal's eye should be efficient in many different tasks (compared with the "eyes" of a robot which only needs to be good at fewer tasks). Animals often have a binocular viewing field created by two eyes. The placement of the eyes and the viewing field are dependent on the environment (Wimborne, Mark & Ibbotson 1999). Cats for example have frontal eyes with a binocular field of vision of approximately 98° and monocular field of vision of 143°. Rabbits instead have lateral eyes with a binocular field of vision of 24° and monocular field of vision of 192°. Wimborne et al. (1999) argues that a wide view angle and for example an upper field of vision (e.g., to watch out for birds if you are a frog) is more important to survive in the environment to protect against predators. Hence, requirements on the predator's eyes should also be dependent on what kind of pursuit method the animal uses (e.g., if the predator pursues its prey or ambushes it). Moreover, animals can move their heads. This increases the actual field of vision. The animal also has other senses whose information can guide the behaviour directly or indirectly through eye movements (e.g., hearing, electrical sense, lateral-line sense (for a detailed explanation, see Lythgoe 1979)).

2.2 Evolutionary robotics

Evolutionary robotics is a technique used to train robots to perform one or more tasks, with as little human intervention as possible. The technique is based on the Darwinian principle of selective reproduction of the fittest individuals (Nolfi & Floreano 2000). It is a very useful technique for training robots when little is known about the environment in which the robots are to operate (Nolfi & Floreano 2000).

2.2.1 Evolution in evolutionary robotics

Campbell & Reece (2002) define the term evolution as:

> "All the changes that have transformed life on Earth from its earliest beginning to the diversity that characterizes it today." (Campbell & Reece 2002, glossary)

The concept of evolutionary robotics is inspired by the real world (i.e., the principle of evolution). The pursuit problem is specifically aimed at exploring the concept of food chains, and contemporary research is mostly focused on food chains that have two trophic levels (e.g., Nolfi & Floreano 1998, Haynes et al. 1995, Nishimura & Ikegami 1997).

2.2.2 Co-evolution

Co-evolution often involves two separate species that compete in some way with each other, such as a predator and a prey species. The selection criteria for the two species are different, the predators' selection criteria is for pursuit behaviours and the prey's selection criteria is for evasion (Cliff & Miller 1995).

In co-evolution the fitness landscape can vary over time. When one of the species adapts its behaviour, the other species' fitness landscape will be deformed (Cliff & Miller 1995, Nolfi & Floreano 1998). This is called *the Red Queen effect* (Ridley 1993, Cliff & Miller 1995, Nolfi & Floreano 1998). The red queen effect makes it hard to monitor progress because the individuals' fitness landscape can vary dynamically during co-evolution. Therefore, one wants to avoid the red queen effect. This can be accomplished by using the CIAO (Cliff & Miller 1995) and/or master tournament (Nolfi & Floreano 1998) techniques. To detect that the red queen is present is difficult, because of the open-ended nature of co-evolutionary simulations (Cliff & Miller 1995, Nolfi & Floreano 1998).

In competitive co-evolution something one wants to achieve is an arms race, which occurs when two or more competing species drive each others evolution forward (Nolfi & Floreano 1998). If it takes more time for the predator to counter-adapt to the prey adaption than it take for the prey to counter-adapt to the predator, the prey is ahead in the race.

An arms race can be classified as symmetric or asymmetric:

- Symmetric means that the species compete to achieve the same goal (e.g., two predators tries to catch the same prey).
- Asymmetric is when two or more species have different goals (e.g., a predator and prey scenario) (Dawkins & Krebs 1979).

An arms race can be distinguished to be interspecific or intraspecific:

- Interspecific arms races describe the evolution between two or more species (e.g., predator and prey).
- Intraspecific arms races are between robots within the same species (e.g., only predators).

The experiments carried out in this dissertation would be classified as asymmetric and interspecific because the prey and predator populations (which are not the same species) have different (but coupled) goals.

There are several ways in which an arms race can end. One species can win if the other is driven to extinction; one species may have reached its optima; both species reach a local optima that is equally good; or the race will get caught in a cycle that is theoretically endless (Dawkins & Krebs 1979).

2.3 Related work

This section describes other work that is related to this dissertation. Related work can be divided into two groups: *(i)* Work describing similar behaviours to the behaviours in this dissertation (Section 2.3.1) and *(ii)* work describing control architectures which are able to evolve similar behaviours to the ones needed in this dissertation (Section 2.3.2).

2.3.1 Behaviours

The pursuit problem: According to Haynes et al. (1995), Haynes & Sen (1997) and Jim & Giles (2000) the original version of the pursuit problem was done by Brenda, Jagannathan & Dodhiawalla (1985). The experiment consisted of one red agent, whose goal was to not get caught by the four blue agents while it was running around in a grid-world, and their goal was to capture the red agent. For the four blue agents to capture the red agent they had to surround him. The agents could only move horizontally or vertically in the grid-world, and two agents could not be at the same spot at the same time. The red agent's movement was random, it could choose one of the four neighboring locations that were not occupied by a blue agent.

Since then there has been a number of experiments done with one predator that hunts one prey which is trying to avoid being caught (e.g., Floreano & Nolfi (1997*b*) and Nolfi & Floreano (1998)). Búason (2002) and Búason & Ziemke (2003*a*, 2003*b*, 2003*c*) tried to improve Nolfi & Floreano (1998). Búason tried to minimise the human bias by using evolutionary techniques and maximise the self-organisation. The different parameters that were particularly interesting was how the robot could evolve the camera direction, camera view angle and the camera view distance. The predators and prey preferred different camera directions, the prey preferred wider camera view angle and the predator preferred longer camera view range.

One pursuit problem experiment with multiple predators and prey in the environment has been done by Nishimura & Ikegami (1997), in which there were 50 individuals of each species in the beginning of the experiment in the environment. The environment was dynamic:

- The predators died if they did not catch any prey but if they caught a prey they could reproduce there genes.
- The prey could reproduce if they were not caught and they died if they got caught.
- The prey could fight back if the predator did not catch the prey from behind, which results in the death of both individuals.

All the aforementioned experiments involve two trophic levels. No other experiments with a larger number of trophic levels have been carried out to the knowledge of the author of this dissertation knows. One big difference between these experiments and the experiments that will be conducted in this dissertation is that an extra trophic level will be added (i.e., so there are three trophic levels).

ALVINN: Pomerleau (1993, 1995) and Batavia, Pomerleau & Thorpe (1996) trained a neural-net guided vehicle to use vision-based autonomous driving. This learning system was called ALVINN (Autonomous Land Vehicle In a Neural Network), where the aim was to automatically keep a self-controlled car driving on a road. The car had a camera that took images of the road. Every pixel of the image became one input-node to the neural network. The network was trained with back-propagation, where the target was to learn from a human driving the car. If the output from the neural network differed from the training output (when the human drove the car) the difference was propagated back through the network. The network was over-trained when it only drove straight on a straight road, so they drove curvy on a straight road so that the network could recover if it were to drive some degree out of line. They added another task to the network: to classify the road. It should for example be able to differentiate between roads that had a center line or if there were leaves on the road. ALVINN performed better if it had to solve *both of these tasks at the same time*.

One could therefore argue that ALVINN generalises better by learning two different behaviours at the same time (i.e., drive on the road and classify the surface). In this dissertation the individuals from the middle trophic level have to evolve two behaviours to solve their task; the tasks being to hunt their prey and to evade their predator. This is of interest for this dissertation since the aim is to investigate if robots learn better or faster if trained on more than one task at the time.

2.3.2 Control architectures

Modular network: Nolfi (1997) used a robot that was supposed to move around in an environment, find boxes, grab them and put them outside the environment. He argued that the robot had different behaviours: one behaviour for finding the boxes and one behaviour for putting the boxes outside the environment.

In these experiments it was argued that the robot had to evolve several behaviours (e.g., *(i)* locate and pick up garbage, and *(ii)* drop garbage outside of the arena). In this dissertation, the robot in the highest trophic level is hunting the robot from the middle trophic level, consequently the robot in the middle trophic level must be able to evade this robot. This is one of the behaviours the robot from the middle trophic level needs to evolve. As the task for the robot from the middle trophic level also is to hunt the robot in the lowest trophic level, it must also evolve this behaviour. Overall the robot from the middle tropic level must evolve at least these two different behaviours to become an optimal robot. As the robots have to evolve several behaviours, in the garbage collection experiment and in the experiments in this dissertation, the networks used in these experiments can potentially be use in the experiments conducted in this dissertation. Thus, since the cleaning robot was able to solve the task with the used control architecture (modular network), the same control architecture will be used for some of the comparable experiments in this dissertation.

Extended Sequential Cascaded Network (ESCN): Ziemke & Thieme (2002) conducted a series of so called road-sign experiments using an Extended Sequential Cascaded Network (ESCN), a particular type of recurrent network. The robot got some input through a light sensor and should act upon that at a later point in time.

Ziemke & Thieme (2002) argued that the ESCN could solve delayed response tasks (i.e., if something happens the robot can act upon that later in time). The robots from the middle trophic level in the experiments in this dissertation have to hunt their prey and at the same time try to stay away from their predator. So if a robot from the middle trophic level notices that its hunter is near it might have to begin running away from its predator instead of hunting its prey, and if escaping successfully start hunting its prey again. For the robot to quickly find its prey again, and resumes the hunt, some kind of memory of where it last saw its prey could be useful. This behaviour could be seen as a delayed response, because when the robot has succeeded in escaping its hunter it could go back to the place where it last saw its prey. This is similar to the task in the experiments conducted by Ziemke & Thieme (2002), and thus this type of neural controller (the ESCN) will also be used in the experiments conducted in this dissertation.

A two layer recurrent network: Floreano & Mondada (1996) conducted experiments to see if a robot could find out on its own when to recharge its battery so it could continue moving. The robot did get punished if it went back to the recharging area to early and it could not come back too late either of course. The environment was a rectangle and in one corner the floor was painted black, which was representing the recharging area. One could argue that this robot had two different behaviours: *(i)* control how much energy there is left and *(ii)* locate the recharging area. The control architecture that was used in this experiment is a network with one hidden layer of recurrent nodes. As the experiments conducted by Floreano & Mondada (1996) is (as the experiments in this dissertation) based upon multiple behaviours, it could be interesting to use this control architecture on the robot from the middle trophic level in this dissertation as well.

3 Problem

This section states the problem that is targeted by this dissertation. Further, the aim will be stated and the objectives that will be used to reach the aim are explained.

3.1 Motivation

In competitive co-evolution most of the research that has been done concerning the pursuit problem limited to two trophic levels (Haynes et al. 1995, Nishimura & Ikegami 1997, Floreano & Nolfi 1997*a*, Floreano & Nolfi 1997*b*, Nolfi & Floreano 1998, Floreano, Nolfi & Mondada 1998, Búason & Ziemke 2003*a*, Búason & Ziemke 2003*b*). However, in nature it is not uncommon to have as many as four to five trophic levels (see Section 2.1.1). If there are more than two trophic levels, the middle species must be both predator and prey at the same time. Thus, the middle species needs two different behaviours to survive.

In the robotics area one could argue that ALVINN (Pomerleau 1993, Pomerleau 1995, Batavia et al. 1996) gets better when it is trained with two tasks (e.g., learning to drive down the road and learning to classify the surface of the road), see Section 2.3.1. It is therefore a reasonable hypothesis that the individuals in the middle trophic level could learn better if two goals were to be learned at the same time (i.e., to evade its predator and catch its prey).

Since the robots from the middle trophic level have two different tasks (hunt and evade) it is reasonable to assume that their vision is affected by having two behaviours (hunting and fleeing). As described in Section 2.1.2, eyes are important for most animals. According to Wimborne et al. (1999) the placement of the eyes depends on the environment (e.g., new objects can appear in the environment).

In evolutionary robotics, robots are normally only equipped with one camera (Floreano & Nolfi 1997*b*, Búason 2002) and consequently have only a monocular field of vision. Búason evaluated in what direction the predator and prey preferred their camera to be (in the pursuit problem experiments described in section 2.3.1), how far they could see, and how wide their field of vision was (Búason 2002, Búason & Ziemke 2003*a*, Búason & Ziemke 2003*b*, Búason & Ziemke 2003*c*). However, it has not been investigated in the case of in a three trophic levels environment, how the species in the middle trophic level prefers its vision to be (e.g., direction, length and width of the field of vision).

3.2 Aim of project

The aim of this dissertation is to investigate the behavioural differences between species that are both predator *and* prey with the species that are either predator *or* prey.

3.3 Objectives

The objectives for this dissertation are:

1. A foundation based on earlier research will be established.
   (a) The biology background of the food chain and animal evolution will be investigated.
   (b) The area of evolutionary robotics, and key concepts (such as the pursuit problem and the multiple robotic behaviours) will be described.

2. Experiments will be planned and carried out in order to collect data.

    (a) Plan and execute experiments based on Búason (2002) to get a baseline for the rest of the experiments.

    (b) Plan experiments to investigate the camera direction, angle and view range of the robot from the middle trophic level. The experimental setup is based upon the work of Búason (2002).

    (c) Conduct experiments in order to collect data for how the middle robot performs in comparison with the robots that are either predators or prey.

3. The extracted data will be analysed and conclusions for how learning multiple behaviours at the same time affects an individuals overall performance will be drawn.

4 Method

This section describes the research methods that have been chosen for the different objectives. Section 4.1 will discuss the research method that will be used and Section 4.2 will describe the planning. Finally, an experimental overview (Section 4.3) is illustrated.

4.1 Research method

In this thesis a literature analysis is carried out in order to establish the background. The literature analysis is necessary because of two things: *(i)* to investigate what has been done within the area of evolutionary robotics with focus on the pursuit problem, and *(ii)* to investigate how real animals from the middle trophic level prefer to have the eyes (frontal or lateral) and how wide they prefer their field of vision to be. It would be interesting to see if the robots from the middle trophic level would develop to have its camera view angle, camera view range and camera direction the same way as the field of vision of for instance, cats (a real life middle trophic level species). If such correlations could be found, this would strengthen the hypothesis that robots, like animals, adapt not only their behaviour but also their morphology to their environment.

Berndtsson, Hansson, Olsson & Lundell (2002, p.59) defines a literature analysis as: “a systematic examination of a problem, by means of an analysis of published sources, undertaken with a specific purpose in mind”. In this treatise, the literature analysis is used in order to gather background information about the area of evolutionary robotics and the connection to biology (objective 1).

Berndtsson et al. (2002, p.66) state that: “experiments are used to verify or falsify a previously formulated hypothesis”. In order to evaluate if a robot that is trained to do two tasks (hunting prey and evading predators) becomes a better hunter and a prey than a robot that is trained for only one task (either hunting or fleeing) a statistical analysis is conducted to evaluate how good the robot from the middle trophic level competes to the top predator and the bottom prey (objective 2) from the baseline experiment. Additionally, some time will be spent on looking at how the robots from the middle trophic level behave in the environment from a more subjective point of view (i.e., a visual analysis of the behavioural patterns of the robots). The experiments could be conducted in different ways. Two alternatives are: *(i)* either use real robots in a real environment, or *(ii)* use a simulator. The first alternative, to evolve a real robot in a real environment, is not feasible (in this dissertation) since it is time consuming and resource demanding.

Based on the experiments, an argumentative approach will be used to draw conclusions about the effect of learning multiple behaviours at the same time (objective 3). This methodology is suitable for this objective since an outside observation of the environment and the individuals is carried out.

4.2 Experimental setup

This section describes the parameter setup for the simulator (Section 4.2.1), followed by what parameter constraints there are (Section 4.2.2) and the fitness function (Section 4.2.3).

4.2.1 Parameter setup

The simulator Yet Another Khepera Simulator (YAKS) is used in this dissertation, since this work is based on the work by Búason (2002) which used this simulator. The YAKS simulator uses connectionist networks and is based on the widely used Khepera robot. For more information visit `http://his.se/iki/yaks`. YAKS is used by a number of researchers (e.g., Carlsson & Ziemke 2001, Búason 2002, Bergfeldt 2005) due to its open architecture that is fairly easy to adopt for special needs. The Khepera robot has a number of modules (infrared sensors and a simple vision module, just to mention a couple of them) (see Figure 4). The sensor values are based on values that has been recorded from a real Khepera robot's infrared sensors (Carlsson & Ziemke 2001). Normally, the robot is not equipped with a camera, but it will be for this experiment. The robot is designed with six proximity sensors in the front, and two in the back (see Figure 4).

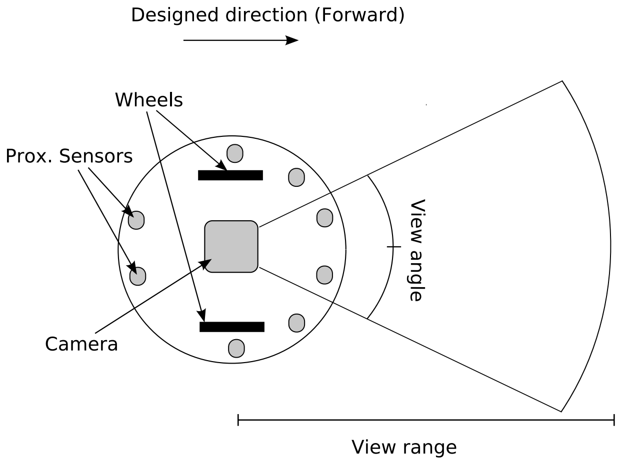

Figure 4: Khepera robot (adapted from Búason & Ziemke (2003*a*)).

Table 1: Parameter setup for YAKS during training

Parameters	**Quantity**
Species (populations)	3
Generations	1000
Individuals	100
Parents	20
Offsprings	5
Epochs	10
Time steps	500
Selection method	elitism
Speed	0 - 1 (0.5 means stop, 1 means full speed forward and 0 means backward full speed)
Camera resolution	5

The parameter setup for the YAKS simulator are in many ways similar to the setup used by Búason (2002). What differs is the network and the generations, the rest of the parameters are the same (see Table 1). Búason (2002) used eight input nodes from infrared sensors and five input nodes from the camera. In this dissertation two extra input nodes are added (see Figure 5). These input nodes represents the robots' "hearing", one node is activated if the robot hears a predator and the other node is active if the robot hears a prey, this for the mid trophic level robots to be able to distinguish its predator from its prey. There are different ways to achieve this (e.g., colour coding or hearing) but in this dissertation the hearing approach is chosen due to simplicity. The activation given by these nodes is normalized to be real values between $0 - 1$.

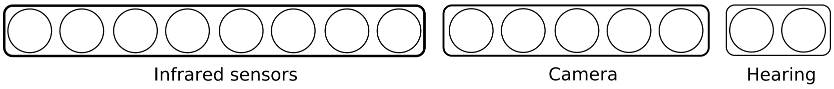

Figure 5: Input nodes from the infrared sensors, the camera and the hearing (adapted from Búason & Ziemke (2003*a*)).

In the experiment based on Búason (2002) the robots are evolved for 250 generations. This experiment works as a baseline for the rest of the experiments and to evaluate if it is possible to replicate the work of Búason (2002) even though two extra input nodes are added to simulate the robots hearing. The network used for these robots are depicted in Figure 6. This network is the same network that Búason (2002) used except for the extra two input nodes for hearing. The camera direction, camera view angle and camera view range that is evolved for the baseline robots (i.e., predator and prey) are used in the opposing robots to train the robot from the middle trophic level.

For the experiments with three robots, they are evolved for 1000 generations. The reason for evolving the robots for 1000 generations and not only for 250, is that the robot from the middle trophic level has a more complex task (i.e., hunt and evade) than the robots that either hunt or evade. One could argue that it is not fair to evolve the robots from the middle trophic level up to 1000 generations when the robots from the top and bottom trophic level are evolved for only 250 generations. One argument against this is that the robots from the middle trophic level have to evolve a more complex behaviour and this could take a longer time to achieve.

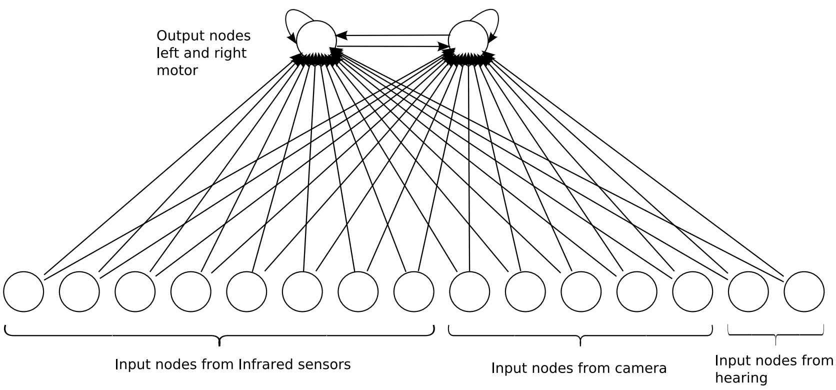

Figure 6: Neural network that is used for the robots that is only predator or prey (adapted from Búason & Ziemke (2003*a*)).

The learning of multiple behaviours for the robot from the middle trophic level could affect the number of necessary generations as follows:

- 500 generations could be enough, i.e., twice as many as Búason (2002) used in his experiment, since two behaviours are evolved.
- It is possible that it will take less time to train two different behaviours than twice the amount since the robot is learning these two behaviours at the same time (cf. Section 2.1.1).
- It could take more than twice the amount of time since the robot has to evolve the two behaviours that may take 250 generations times two and then the robot has to learn when to use a specific behaviour. This could occur because multiple behaviours maybe more complex to evolve than a single behaviour.

Thus, the robot is evolved for 1000 generations but it will be evaluated after 250, 500, 750 and 1000 generations. For each of these "steps" the middle trophic level robots will be evaluated against the top/bottom trophic level robots that have been evolved for 250 generations. This is when (and if) the middle trophic level robots will be as good as (or better) than the robots only having been exposed to an environment with two trophic levels. There are 100 individuals, the selection method is *elitism* and the 20 best individuals out of these 100 are selected to reproduce. Thus, 20 individuals are reproducing 5 new individuals each in every generation. Each individual in a generation runs for 500 time steps and 10 epochs, to let the robots start from 10 different starting positions. The robots' starting positions and starting directions are randomised. Table 1 summarises the evolved parameters.

Table 2: Parameter setup for the robot from the middle trophic level during training

Parameters	**Range**
Camera direction	0 - 360°
Camera angle	0 - 360°
Camera view range	5 - 500 mm

The initial camera view angle, camera direction and camera view range in each experiment are generated for the robots from the middle trophic level using a uniform distribution function, the same that Búason (2002) used. The two robots that are either predator or prey have their camera view angle, camera direction and camera view range fixed to the values that evolved in the baseline experiment based on Búason (2002).

The network weights and the camera (i.e., camera direction, camera view angle and camera view range) are initialised using a Gaussian distribution with a standard deviation of 2.0. The camera direction and the camera view angle is within the interval 0 - 360 degrees, and for the camera view range the interval is 5 - 500 millimeters. During evolution, the camera view angle, camera direction and camera view range is evolved using Gaussian distribution with a standard deviation of 5.0. The network weights use the Gaussian distribution as well but with a standard deviation of 2.0. If the camera direction gets higher than 360° the value is wrapped (e.g., if the value is 373°then it is wrapped to 13°). However, if the camera view angle gets over 360° or below 0°, the angle will be set to 360° or 0° respectively (shown in Table 2). The parameters above are the same as used in the setup used by Búason (2002).

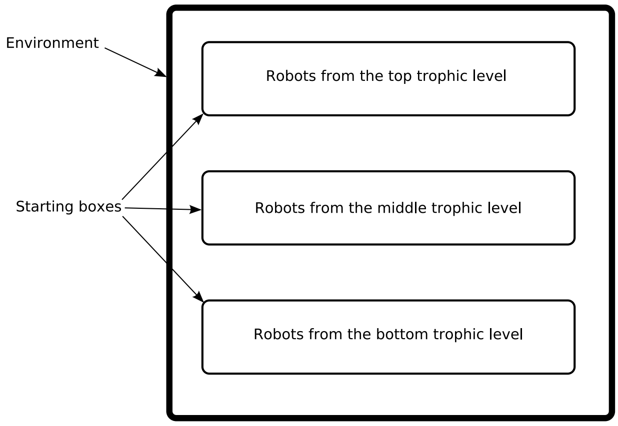

Figure 7: The starting boxes for the robots in the environment.

The environment has the same dimensions as was used by Búason (2002), that is, a square with the dimension 470 x 470 mm and the starting positions of the robots are randomised. Robots from the top trophic level start somewhere at the top, the robots from the middle trophic level starts in the middle of the environment and the robots from the bottom trophic level start at the bottom of the square. The starting positions of the robots are depicted in Figure 7.

In initial experiments the starting positions of the robots were three small squares within the rectangles, however, the robots caught their prey and evaded predators irrespective of the different starting positions. Thus, the starting position had no noticeable effect, and the robots were therefore allowed to start from a random position within the entire rectangles (depicted in Figure 7).

The neural network that is used for the two robots that are either predator or prey is almost the same as the network that was used by Búason (2002), with a small change for the hearing that has been explained earlier in this section (see Figure 6).

Table 3: The camera view angle in contrast to the speed and camera view range

Degree of camera view angle	**% of maximum speed**	**Maximum camera view range in mm**
$0 \leq x < 36$	100	500
$36 \leq x < 72$	90	450
$72 \leq x < 108$	80	400
$\vdots$	$\vdots$	$\vdots$
$324 \leq x \leq 360$	10	50

4.2.2 Constraints

There are some constraints added to the robots. The camera view angle was divided into 10 equally sized intervals, as done by Búason (2002), this means that one interval is 36°. The camera view range and speed is also divided into an interval of 10. This means, if the robot has a camera view angle between 36 - 72° then the maximum speed is decreased by 10% and the camera view range cannot be greater than 450 millimeters (see Table 3).

The robots hearing range is limited to 2,5 times the robot diameter (i.e., $2,5 * diameter$) because the robot should not hear across the whole arena. The limited hearing is 2.5 times the robots diameter because that is 25% of the arena (470 x 470 mm) and the robots diameter is 50 mm.That means that the hearing range is $50mm * 2.5 = 125mm$. In this dissertation the hearing range is fixed as described above, it would therefore be interesting to let the robots themselves evolve an appropriate hearing range. This is however left for future work.

4.2.3 Fitness function

The fitness function for the predators is a simple time-to-contact measurement. This means that the predator should catch the prey as early as possible to get high fitness and the prey should avoid the predator to achieve fitness. The calculation of the fitness for the robots from the middle trophic level is basically a combination of the top and bottom trophic level robots' fitness functions. The fitness functions from the top and bottom tropic levels are used for the robots in the mid trophic level and these two fitness values and then summed up and divided the result by 2 (to normalise the fitness given to the robots in the top and bottom trophic level).

4.3 Experimental overview

This section gives an overview of the competition (Section 4.3.1) and in Section 4.3.2 gives an overview of the selected control architectures.

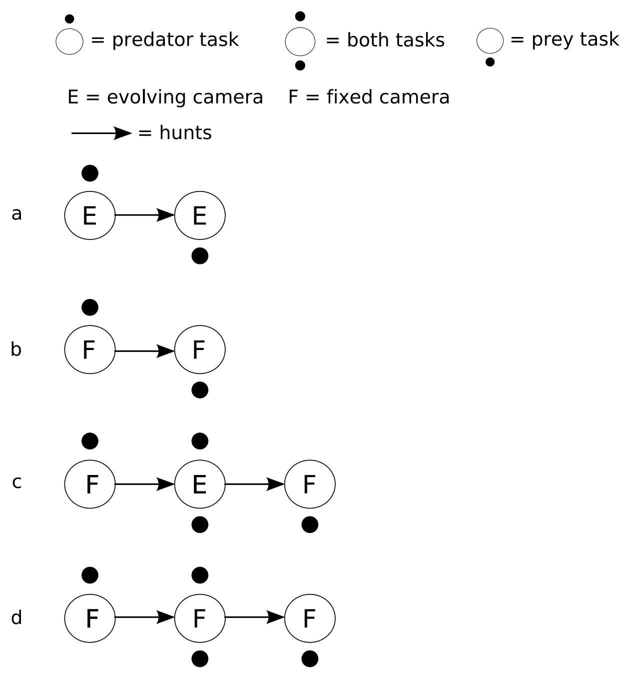

Figure 8: Overview of the experiments.

The experiments are carried out in these steps:

1. Evolve the different baseline robots cameras (see Figure 8a).

2. Evolve the different baseline robots with fixed cameras (using the camera values from step 1) (see Figure 8b).

3. Evolve the robots cameras with different control architectures. The robots in the top and bottom trophic level have fixed cameras (using values from step 1) (see Figure 8c).

4. Evolve the robots with fixed cameras using the values from step 3 for the mid trophic level robots and values from step 1 for the robots from the top and bottom trophic level (see Figure 8d).

5. Evaluate the best baseline robots from step 2, in a two trophic level environment.

6. Evaluate the best robots from the mid trophic level (step 4) against the robots from the top trophic level (step 2), in a two trophic level environment.

7. Evaluate the best robots from the mid trophic level (step 4) against the robots from the bottom trophic level (step 2), in a two trophic level environment.

When designing the experiments a number of test experiments were carried out in order to verify the experimental setup. One experiment was conducted in order to extend the experiments of Búason (2002) with the two extra input nodes to see how that influenced the robots. Further, those two robots will be evolved with the same parameter setup as when there are three robots (see Table 1), except that the number of species is 2, and the number of generations is 250 because that was used by Búason (2002). This is used as a baseline, that is to evolve the camera view angle, camera view range, and the camera direction for one predator and one prey. After the baseline experiment, different experiments are conducted with three robots (three trophic levels), where the predator and the prey had fixed camera direction, camera view angle and camera view range, while the robot that is both predator and prey has to evolve those parameters. After the baseline experiment there were four different experiments (one for each control architecture) with three trophic levels conducted. Every experiment is executed 30 times to get it statistically verified, because the network weights are randomised at initiation.

When all the different robots have evolved their own camera direction, camera view angle and camera view range, the robots run again from the start with a fixed camera, as was concluded in the earlier experiments.

4.3.1 Evaluation by competition

This section presents how the evaluation of the robots in the middle trophic level is done. The best individuals from the middle trophic level in the 30 experiments are competing against the best individuals from the 30 experiments that were conducted where only two trophic levels were competing against each other. This evaluation is conducted in a master tournament when the mid trophic level robots are prey and one where they are predators (in an environment with two trophic levels) with 100 epochs. To calculate how many times one predator catches its prey a counter is used (i.e., the maximum amount of catches for one predator is $30 * 30 * 100 = 90000$ times). The value will be divided with $30 * 30 = 900$, that means that the maximum value that the predator can achieve is 100.

4.3.2 Control architectures

The four different control architectures (networks) (depicted in Figure 9) will be tested on the robot from the middle trophic level to evaluate which control architecture that

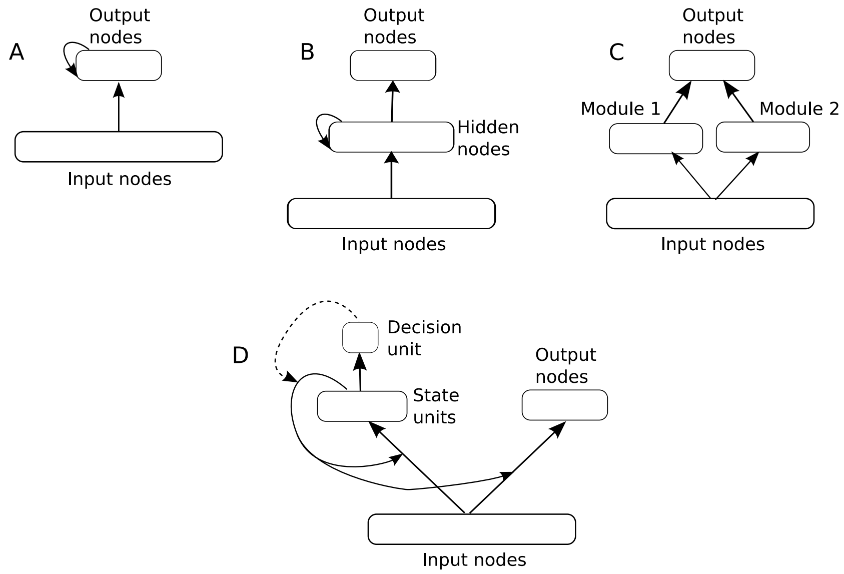

Figure 9: Control architecture (adapted from Búason (2002), Floreano & Mondada (1996), Nolfi (1997) and Ziemke & Thieme (2002)).

is most appropriate. The initial weights for the middle trophic level robots' networks are always randomised (see Section 4.2.1). However, there are two different options regarding the initial network weights of the top and bottom trophic level: *(i)* always the best individual (i.e., already evolved robots), *(ii)* also randomised individuals (network weights). If option *(i)* is used the robot from the middle trophic level might suffer from what is called *the Bootstrap Problem* (Nolfi & Floreano 2000) which is when a population has a hard time accumulating fitness because its opponents are too good at their task. Finally, if option *(ii)* is used this means that the robots are equally "bad" at the beginning and this could lead to an arms race (see Section 2.2.2). In this dissertation option *(ii)* is used because this can avoid the bootstrap problem and it can lead to an arms race.

The robot from the middle trophic level will be tested with different control architectures as shown in Figure 9 and all control architectures have the same amount of input nodes. There are eight nodes of infrared sensors, five nodes for the camera and two nodes for hearing. Further, there are two output nodes, one for the right motor and one for the left motor (depicted in Figure 6).

According to Callan (1999) is it easier to achieve a control architecture that can generalise with few nodes and few hidden layers, these are the control architectures that are used in this dissertation:

- A one layer neural network with recurrent connections on the output nodes (Figure 9A)
- A two layer neural network with recurrent connections on the hidden nodes (Figure 9B)
- Modular network (Figure 9C)
- Extended sequential cascaded network (Figure 9D)

The control architecture in Figure 9A is a one layered network and the output layer is self-recurrent, that is, the same network as Búason (2002) used for the robots that are only predator or prey. However, this control architecture may be too simple to solve the tasks that the robot from the middle trophic level has (hunting and being hunted).

The control architecture illustrated in Figure 9B has two layers that are fully connected and the middle layer of nodes is self-recurrent. Floreano & Mondada (1996) uses 5 hidden nodes and 12 input nodes (Section 2.3.2), and according to Callan (1999) between $30 - 50\%$ of the input nodes should be used as hidden nodes. Floreano & Mondada (1996) has $5/12 \approx 40\%$, which is according to the suggestion made by Callan (1999). In this dissertation there will be five nodes in the hidden layer and that corresponds to approximately 33% of the input nodes. This control architecture, as described in Section 2.3.2, did solve the task to recharge the battery when that was needed. The control architecture could fit to solve the complex task of being a hunter and a prey (i.e., having two behaviours) at the same time because, as argued in Section 2.3.2, one could argue that the robot has different behaviours to solve the task in the experiments conducted by Floreano & Mondada (1996).

Figure 9C, illustrates a modular network with two modules. This type of control architecture was used by Nolfi (1997) as described in Section 2.3.2. Each module has one node that makes the decision if the module is activated or not. This control architecture could be good for solving this task because it solved another task where the robot needed a number of different behaviours, those experiments were conducted by Nolfi (1997).

Finally, Figure 9D, illustrates an ESCN. This control architecture was used by Ziemke & Thieme (2002), as described in Section 2.3.2. The experiments that were conducted by Ziemke & Thieme (2002) showed that the network solved delayed response tasks. A control architecture that can be able to handle delayed response tasks could be good to solve this task as well. For example if the robot from the middle trophic level is hunting its prey (can see the prey) and sees its hunter that runs towards him, then the robot from the middle trophic level can remember where it last saw its prey and can go back there when it has moved away from its predator.

5 Results and analysis

This section presents the results of the experiments in Section 5.1 and the analysis in Section 5.2.

5.1 Results

In this section the results of the experiments are presented. First, Section 5.1.1 presents how the camera were evolved for the robots that was evolved in an environment with two trophic levels (i.e., either predator or prey). The same section presents how the robot from the middle trophic level that has two behaviours (i.e., hunt and evade) preferred to have the camera direction, camera view angle and camera view range. These were the experiments that were based on Búason (2002). Secondly, Section 5.1.2 presents how good the robot from the middle trophic level performed with its different control architectures compared to the robots (i.e., the predator and the prey) that has only one task to solve (i.e., one behaviour).

5.1.1 Vision evolution

In this section the results from objective 2c are presented. The results for the evolved camera view angle, camera view range and camera direction for the predator is presented in Table 4 and the results for the prey is presented in Table 5. The last table in this section presents how the robot from the middle trophic level wants to have its camera view angle, camera view range and camera direction for the different control architectures (for details, see Table 6).

Table 4: The evolved camera of the **predator** robot (standard deviation in parenthesis)

Parameters	**Range**
Camera direction	186° (96°)
Camera angle	160° (102°)
Camera view range	200 mm (111 mm)

Table 5: The evolved camera of the **prey** robot (standard deviation in parenthesis)

Parameters	**Range**
Camera direction	181° (99°)
Camera angle	200° (86°)
Camera view range	160 mm (76 mm)

5.1.2 Competition evaluation

This section will present the results that show how well the robot from the middle trophic level performs in comparison to the predator in trophic level three and the prey in trophic level one. The design of the robots implies a direction, however, the moving direction can be be different (i.e., the robots can move backwards). Due to reasons

Table 6: The evolved camera of the robot from the middle trophic level with different control architectures (standard deviation in parenthesis)

Network in figure	**Camera direction**	**Camera view angle**	**Camera view range**
9A	191° (108°)	156° (100°)	133 mm (116 mm)
9B	136° (92°)	147° (102°)	140 mm (81 mm)
9C	177° (82°)	175° (110°)	179 mm (113 mm)
9D	146° (102°)	146° (97°)	213 mm (103 mm)

given in Section 5.1.1 the camera direction is backwards (see Figure 10). Figure 12 presents how well the robots performed as predators. The performance of the robots as prey's is shown in Figure 13. In Figure 15 the predator performance and in Figure 14 the prey performance is shown, when the moving direction is restricted to move forward (in the designed direction of the robot). When the moving direction is backwards (in the opposite direction to the designed direction of the robot), the performance of the predator robots as shown in Figure 17 and those for the prey robots in Figure 16.

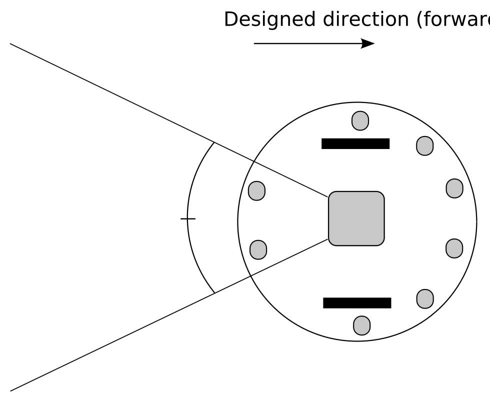

Figure 10: The design direction of the Khepera robot, forward and backward

5.2 Analysis

This section compares the different results, in Section 5.2.1 and 5.2.2, and presents the analysis of these results.

5.2.1 Vision evolution

The predators' and the prey's camera direction, camera view angle and camera view range differ from Búason (2002) (shown in Section 5.1.1). This is because the robot in that experiment had to rely on information from vision only, and in this dissertation the robot can rely on both visual information and the categorical information from hearing. In these experiments with only two trophic levels that are based on Búason (2002), the robot that was a predator used its hearing first to locate its opponent if the opponent was close enough but outside the predators' field of vision. For example, if the predator has the camera directed in the opposite direction away from the prey the predator uses its hearing to locate the prey's location (if the prey is behind him) and then turns around to use its vision to catch the prey. On the other hand if the predator does not hear the

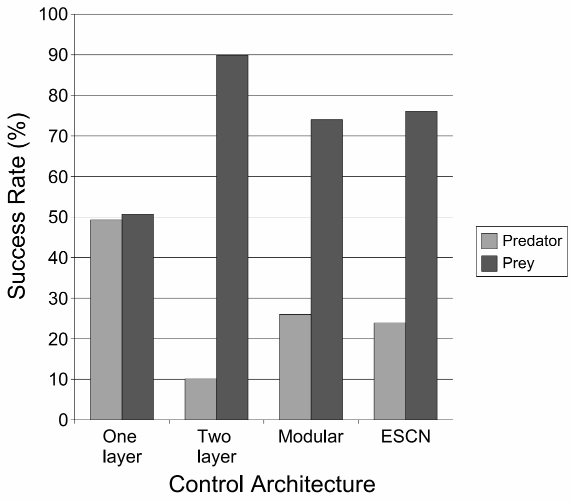

Figure 11: Percentage of **catches** for the **baseline top trophic level predator** robots and **evasions** for the **baseline bottom trophic level prey** robots (for more details see Appendix B, Table 9 and Table 10)

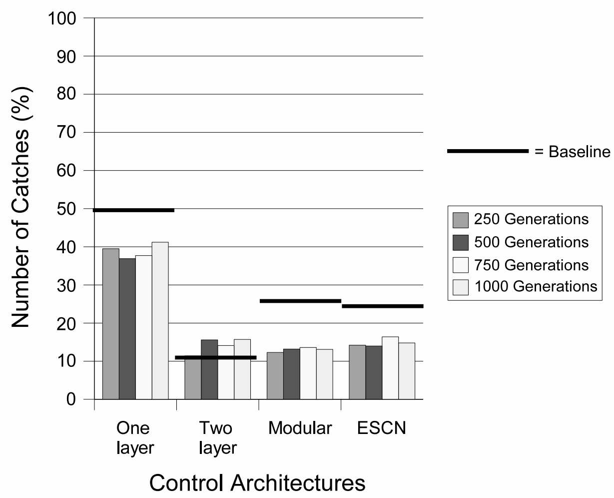

Figure 12: Percentage of **catches** for the **mid trophic level predator** robots (for more details see Appendix A.1, Table 7)

prey he moves around in the environment to search for the prey. This is probably why the evolved camera view angle is smaller and the evolved camera view range is bigger in Búason than on the predator robot that was evolved with hearing nodes. However, on the prey robot the camera view range is smaller and the view angle the same. The camera direction was the same as Búason (2002).

The robot in the middle trophic level evolves its camera to be more like the robot that was only a predator than the robot that was only a prey. This is true fore all the different control architectures (see Figure 9). The camera view angle, camera view range and camera direction is presented in Table 6. This is comparable to cats (e.g., the bobcats in Figure 3) that have a frontal field of vision (described in Section 2.1.2), because they are predators and prey’s at the same time.

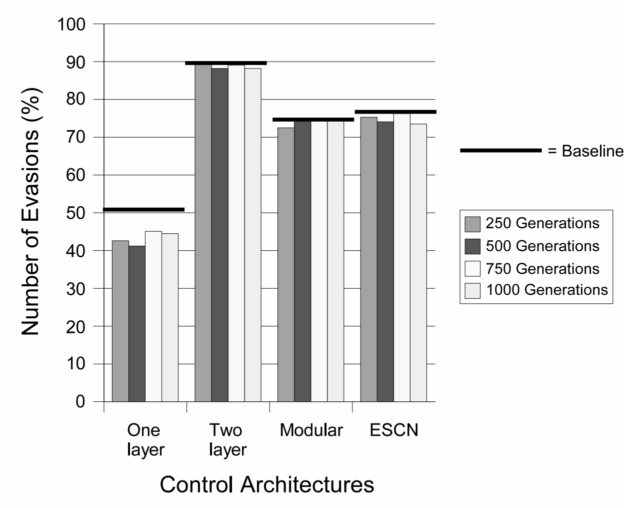

Figure 13: Percentage of **evasions** for the **mid trophic level prey** robots (for more details see Appendix A.1, Table 8)

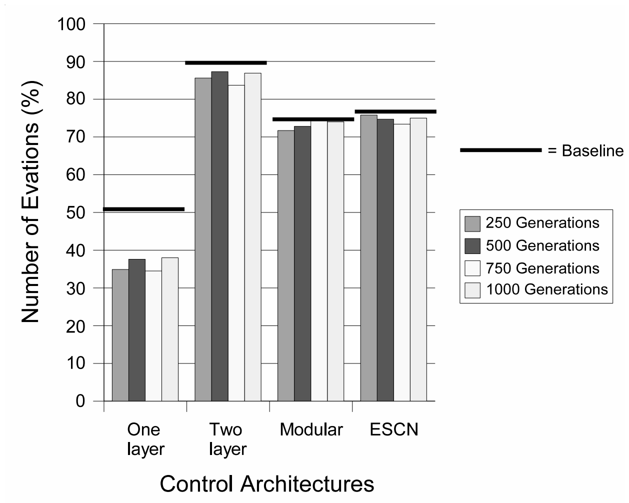

Figure 14: Percentage of **evasions** for the **mid trophic level prey** for the **forward** moving robots (for more details see Appendix C.1, Table 12)

5.2.2 Competition evaluation

The first paragraph shows the results for the mid trophic level robots when they are placed in an two trophic level environment and compete against the predators of the top trophic level (see Figure 12). Figure 13 shows the results for the mid trophic level robots compete against the robots of the bottom trophic level.

Two Trophic Levels: The results for the baseline robots from the top and bottom trophic level (that are only a predator *or* a prey) are presented in Figure 11. When the robots from the middle trophic level act as a predator in a two trophic level environment the robots with control architecture 9B is the only that is better than the respective baseline robots after 250 generations.

Regardless of the control architecture, some of the robots prefer to move in the

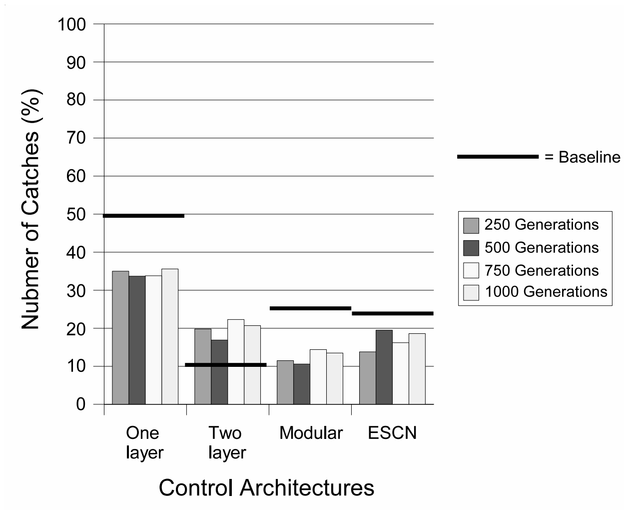

Figure 15: Percentage of **catches** for the **mid trophic level predator** for the **forward** moving robots (for more details see Appendix C.1, Table 11)

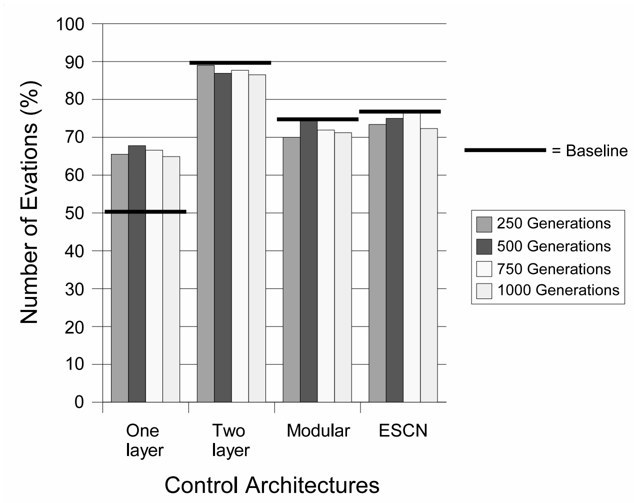

Figure 16: Percentage of **evasions** for the **mid trophic level prey** for the **backward** moving robots (for more details see Appendix D.1, Table 14)

camera direction, and some preferred the opposite. Figure 15 shows how the robots that are moving in the design direction and having the camera facing backwards perform some percent-unit worse or equal than the robots that are mixed up (i.e., all robots, some are moving in the design direction and the others are moving against it) as in Figure 12. However, in Figure 14 they are as good as the mixed up robots (see Figure 13).

The robots in the middle trophic level that are moving against the designed direction (see Figure 10) and that are hunters (see Figure 17) are as good as the respective baseline hunters, except from the robots with control architectures 9A and 9B. Thus, the robots with control architecture 9A are 9%-units better than the baseline hunters. The same can be observed when comparing the prey's that are moving backwards (see Figure 16) with the prey from the baseline. Only the robots with control architecture 9A are better than their respective baseline prey, almost 15%-units better.

These results show that these robots are as good at evading its predator as the robots

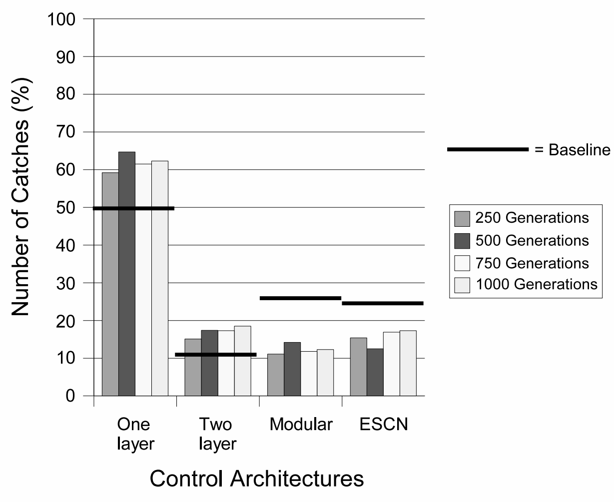

Figure 17: Percentage of **catches** for the **mid trophic level predator** for the **backward** moving robots (for more details see Appendix D.1, Table 13)

that use the same control architecture with only one behaviour (see Figure 11). It does not matter if the robots are moving forward or backwards when evading its predator. All robots with different control architectures, are equally good as the robots with the same control architecture except for control architecture 9A (see Figure 13, Figure 14 and Figure 16). However, when hunting its prey there are only two control architectures (9A and 9B) that perform better than the network in the baseline experiments, i.e., when the robots are moving backwards (see Figure 17). When the robots are moving forwards the only control architecture that performs better than the networks in the baseline experiments is 9B (see Figure 15).

Evaluation of Behaviour: In the experiment based on Búason (2002) the predator moved around in the environment to find its prey, moving “backwards” (i.e., against the designed direction of the robot). The prey on the other hand stood in one place near walls or corners in the environment and roted around trying to spot the predator. When the predator was spotted the prey tried to outrun it, and that worked out approximately 50% of the time (see the two left bars (one layer) in Figure 11 and the baseline).

The robot’s different control architectures evolved different behaviours: When the robot uses control architecture 9A it moves around in the environment trying to catch the prey and avoiding its predator. When it catches its prey it stopped moving around and stayed close to the walls and corners, until it hears the predator. Then it tries to outrun the predator. If the robot sees the predator (but cannot hear its prey) after it has caught its prey, the robot moves towards the predator until it hears that it is the predator and then tries to outrun the predator.

The robots with control architecture 9B moved around in the environment clockwise near the walls trying to avoid getting caught, and in the same time trying to catch its prey, the behaviour continues after the prey has been caught. However, when it sees its predator it tries to outrun it and if the robot sees its predator but cannot hear it, it moves closer and tried to outrun it when it hears that it is its predator.

With control architecture 9C the robot continuously moves around in the environment trying to catch its prey while avoiding the predator. The robots with control architecture 9C, behave like the robots with control architecture 9A when they are able to see but cannot hear their predator.

Finally, when the robots use control architecture 9D they try to hide from their predator. When they spots their prey, they set off and try to catch it. This type of control architecture could possibly perform better when evolved for a longer time. An initial experiment was executed with 6000 generation and the result showed that the robot performed as good as its respective baseline robots.

6 Conclusion

This section discusses the results that have been presented in Section 6.1 and shows the contribution in Section 6.2. Finally, in Section 6.3 ideas for possible future work is outlined.

6.1 Discussion

Section 6.1.1 describes how the camera was evolved for different robots using different networks. The evaluation of the performance of the robot from the middle trophic level was compared with the robots that were only predator or prey is discussed in Section 6.1.2.

It is interesting that the robots from the middle trophic level, using the control architecture from Figure 9A, and moving around in the same direction as its field of vision perform better than the prey-only robot at evading the predator, and also performs better than the predator-only robot catching the prey. The interesting part about this is that in nature animals from a middle trophic level have their eyes placed similarity (cf. Section 2.1.2).

The robots with control architecture 9A and multiple behaviours solve the task better than these with only one behaviour. However, multiple behaviours do not have the same effect when using other control architectures (9B, 9C and 9D). The robots with control architecture 9B perform better than the respective baseline robots when they are acting as a predator (see Figure 17) and perform equally well when acting as a prey (see Figure 16). The predator robots in the baseline experiments perform poorly. However, they perform really well as prey (see Figure 11). The robots (all control architectures) that are evading are as good as the baseline prey robots (see Figure 13, Figure 14 and Figure 16); however, when the robots try to catch the prey, only robots with two out of the four control architectures perform better than the baseline robots (see Figure 17).

6.1.1 Vision evolution

The robots in an environment with two trophic levels (extending the work of Búason (2002)) evolved almost the same camera direction as in Búason (2002), despite having two extra hearing nodes. It was the predators' camera view angle and camera view range that was different from the experiment conducted by Búason (2002), possibly this is an effect due to the robots' hearing. It is better to have a wide field of vision when you have hearing, because the hearing helps the robot to locate if the prey is near, and when the robot has turned around and can see the prey with its vision (camera) it can move towards it. In nature, animals use both vision and hearing to locate their prey or predators and then tend to rely more on vision than hearing when trying to catch or evade it. However, hearing is often used for an initial detection, similar to the behaviour evolved in the experiments. Vision and hearing were not stereoscopic in the experiments conducted in this dissertation, however, this would be interesting to implement in order for the environment to resemble the real world as much as possible.

6.1.2 Competition evaluation

The robots from the middle trophic level that had two behaviours in general not as good hunters as the robots that only had one behaviour (see Figure 12). They are not as good at hunting if they move forward (in the design direction) with the camera aimed backwards (see Figure 15). However, they are not as good at hunting if they are moving in the same direction as camera vision direction, except for the robots with control architectures 9A and 9B. They outperformed the robots with one behaviour. The evaluation shows that robots with control architecture 9A are better hunters with 9%-units better success than the baseline hunters (see Figure 17) and robots with control architecture 9B are almost 5%-units better than their respective baseline robots. Robots with control architecture 9A are better than the baseline prey by almost 15%-units (see Figure 16). All other evaluations of the prey moving in the direction of the camera with the other control architectures shows similar success at evading as the respective baseline prey (see Figure 16). When using control architecture 9A in a robot it is moving in the field of vision direction. One could argue that this is the same behaviour as a bobcat has (i.e., eyes in a frontal position, so it can catch its prey). For the other control architectures (9B, 9C and 9D) they perform nearly as well at evading as the respective baseline prey (see Figure 16). Thus, it is possible that the robots with these control architectures need to evolve for a longer period before they get better than the baseline robots because their control architecture are more complex. An initial experiment on robots with control architecture 9D was allowed to evolved for 6000 generations and those robots performed equally good as a hunter baseline robot with that control architecture.

Using three trophic levels makes the experiments more realistic than just using two trophic levels, since not only pure predators and preys are considered. Instead predators that are not only top predators (e.g., bobcat which hunts rabbits and evades mountain lions) are represented as well.

6.2 Contribution

This dissertation extends the work by Búason (2002). In contrast with his work an extra trophic level has been added and the camera position of the robot in the middle trophic level has been evolved.

The main contribution is shown by the evolution of robots from a middle trophic level, that has to solve multiple tasks. The results that are presented in this dissertation indicate that robots that need to solve multiple tasks can perform as well as or better than robots that only need to solve one task. However, it shows that considering only two trophic levels is both unrealistic and less successful in some cases. This indicates that the study of two trophic levels in isolation as done so far might not resemble the complexities of reality as most animals have more than one role. Moreover, a robot that has to solve two tasks at the same time, can perform as well as or better than a robot that only needs to solve one of the tasks. This is not always the case, however, therefore more research is needed in this area, before any strong conclusions can be drawn. Section 6.3 gives suggestions for future work in this area.

6.3 Future work

The main aim of future work should be to make robots more animal-like in order to improve the level of autonomy towards a level of autonomy that animals have. There are two main directions for future work: *(i)* to enhance the robot and *(ii)* to use more advanced and realistic environments.

One interesting topic *(i)* would be to evolve two eyes, as animals normally have two eyes. This would allow the robots to evolve stereoscopic vision (cf. Section 6.1.1).

Maybe one could place the eyes like a typical predator or a typical prey. One could investigate what happens if they could use a stereoscopic vision which allows the robot, for instance, to estimate distance. Most animals can move their heads, therefore it might be worth investigating artificial agents whose cameras have several degrees of freedom (i.e., a more narrow field of vision might be sufficient then). One might also think of implementing stereoscopic hearing which would allow the agents to detect the precise direction towards its predator or prey. Another way to create more realistic agents might be the inclusion of energy resource systems, thus an agent might get exhausted after chasing for some time even die of starvation if it did not catch sufficient prey in a certain period of time.

An interesting topic for *(ii)* would be to use multiple robots in each of the three trophic levels (e.g., to investigate group behaviours). This could be specially interesting when investigating a robot food web. Already with only one robot of each species there might be interesting effects such as emergent cooperation amongst the predators in order to catch a common prey. The introduction of objects into the environment might also be worth investigating because it might make the environment more realistic and both predator and mrey might use them in their strategies (i.e., they would use them for hiding or trapping).

Finally, it might be attractive to extend some of the robot controll architectures to more complex ones or just to evolve the weights of the control architecture 9D for a longer time to see if the robot performance improves even more (cf. Section 6.1). These ideas could be regarded as initial steps on the way to the creation of robots that display animal-like behaviour in real environments.

References

Abrams, P. A. & Roth, J. (1994), 'The responses of unstable food chains to enrichment', *Evolution Ecology* **8**, 150–171.

Batavia, P., Pomerleau, D. & Thorpe, C. (1996), Applying advanced learning algorithms to alvinn, Technical Report CMU-RI-TR-96-31, Robotics Institute, Carnegie Mellon University, Pittsburgh, PA.

Bergfeldt, N. (2005), 'Towards detached communication for robot cooperation'. Licentiate Thesis No. 1192, University of Linköping.

Berndtsson, M., Hansson, J., Olsson, B. & Lundell, B. (2002), *Planning and Implementing your Final Year Project with Success!: A Guide for Students in Computer Science and Information Systems*, Springer - Verlag London Limited.

Brenda, M., Jagannathan, V. & Dodhiawalla, R. (1985), On optimal cooperation of knowledge sources, Technical Report BCS-G2010-28, Boeing AI Center, Boeing Computer Services, Bellevue, WA.

Búason, G. (2002), 'Competetive co-evolution of sensory-motor systems'. Master Thesis No. HS-IDA-MD-02-004, University of Skövde.

Búason, G. & Ziemke, T. (2003*a*), Co-evolving task-dependent visual morphologies in predator-prey experiments., *in* C.-P. et al., ed., '*Genetic and Evolutionary Computation Conference*', Berlin: Springer-Verlag, pp. 458–469.

Búason, G. & Ziemke, T. (2003*b*), Competitive co-evolution of predator and prey sensory-motor systems, *in* C. et al., ed., '*Applications of Evolutionary Computing - EvoWorkshops 2003*', Berlin: Springer-Verlag, pp. 605–615.

Búason, G. & Ziemke, T. (2003*c*), Explorations of task-dependent visual morphologies in competitive co-evolutionary experiments, *in* W. Banzhaf, T. Christaller, P. Dittrich, J. T. Kim & J. Ziegler, eds, '*Advances in Artifial Life - Proceedings of the seventh European Conference on Artificial Life*', Berlin: Springer-Verlag, pp. 763–770.

Callan, R. (1999), *The essence of Neural Networks*, Prentice Hall.

Campbell, N. A. & Reece, J. B. (2002), *Biology*, International Edition, sixth edn, Pearson Education, Inc., publishing as Benjamin Cummings.

Carlsson, J. & Ziemke, T. (2001), Yaks - yet another khepera simulator, *in* '*Proceedings of the 5th International Heinz Nixdorf Symposium*', HNI-Verlagsschriftenreihe, Paderborn, Germany, pp. 235–241.

Cliff, D. & Miller, G. F. (1995), Tracking the red queen: Measurements of adaptive progress in co-evolutionary simulations, *in* J. J. M. F. Moran, A. Moreano & P. Chacon, eds, '*Advances in Artificial Life: Proceedings of the Third European Conference on Artificial Life*', Berlin: Springer-Verlag.

Dawkins, R. & Krebs, J. R. (1979), Arms races between and within species, *in* '*Proceedings of the Royal Society of London, Series B, 205*', Department of Zoology, South Parks Road, Oxford OX1 3PS, U.K., pp. 489–511.

Floreano, D. & Mondada, F. (1996), 'Evolution of homing navigation in a real mobile robot', *Systems, Man and Cybernetics, Part B, IEEE Transactions on* **26**(3), 396–407.

Floreano, D. & Nolfi, S. (1997*a*), Adaptive behavior in competing co-evolving species, *in* P. Husbands & I. Harvey, eds, '*Fourth European Conference on Artificial Life*', The MIT Press, Cambridge, MA, pp. 378–387.

Floreano, D. & Nolfi, S. (1997*b*), God save the red queen! competition in co-evolutionary robotics, *in* J. R. Koza, K. Deb, M. Dorigo, D. B. Fogel, M. Garzon, H. Iba & R. L. Riolo, eds, '*Genetic Programming 1997: Proceedings of the Second Annual Conference*', Morgan Kaufmann, Stanford University, CA, USA, pp. 398–406.

Floreano, D., Nolfi, S. & Mondada, F. (1998), Competitive co-evolutionary robotics: From theory to practice, *in* R. Pfeifer, ed., '*Proceedings of the Fifth International Conference on Simulation of Adaptive Behavior*', MIT Press-Bradford Books, Cambridge, MA.

Haynes, T. & Sen, S. (1997), 'The evolution of multiagent coordination strategies', *Adaptive Behavior*.

Haynes, T., Wainwright, R. & Sen, S. (1995), Evolving cooperation strategies, *in* V. Lesser, ed., '*Proceedings of the First International Conference on Multi–Agent Systems*', MIT Press, San Francisco, CA, p. 450. citeseer.ist.psu.edu/haynes94evolving.html.

Jim, K.-C. & Giles, C. L. (2000), 'Talking helps: Evolving communicating agents for the predator-prey pursuit problem', *Artificial Life* **6**(3), 237–254.

Lythgoe, J. N. (1979), *The Ecoloty of Vision*, Clarendon Press - Oxford.

Nishimura, S. & Ikegami, T. (1997), 'Emergence of collective strategies in prey-predator game model'. citeseer.ist.psu.edu/article/nishimura97emergence.html.

Nolfi, S. (1997), 'Using emergent modularity to develop control system for mobile robots', *Adaptive Behavior* **3-4**(5), 343–363.

Nolfi, S. & Floreano, D. (1998), 'Coevolving predator and prey robots: Do "arms races" arise in artificial evolution?', *Artificial Life* **4**(4), 311–335. citeseer.ist.psu.edu/nolfi98coevolving.html.

Nolfi, S. & Floreano, D. (2000), *Evolutionary Robotics: The Biology, Intelligence, and Technology of Self-Organizing Machines*, MIT Press, Cambridge, Massachusetts.

Pomerleau, D. (1995), Neural network vision for robot driving, *in* M. Arbib, ed., '*The Handbook of Brain Theory and Neural Networks*', MIT Press.

Pomerleau, D. A. (1993), *Neural Network Perception for Mobile Robot Guidance*, Kluwer Academic Publishers.

Ridley, M. (1993), *The Red Queen: Sex and the evolution of human nature*, London : Penguin Books.

Stauffer, D., Kunwar, A. & Chowdhury, D. (2005), 'Evolutionary ecology in silico: evolving food webs, migrating population and speciation.', *Physica A* **352**(1), 202–216.

Wimborne, B., Mark, R. & Ibbotson, M. (1999), 'Distribution of retinogeniculate cells in the tammar wallaby in relation to decussation at the optic chiasm', *The Journal of Comparative Neurology* **405**(1), 128–140.

Ziemke, T. & Thieme, M. (2002), 'Neuromodulation of reactive sensorimotor mappings as a short-term memory mechanism in delayed response tasks', *Adaptive Behavior* **10**(3-4), 185–199.

Appendix A Forward and backwards moving robots

Section A.1 presents two tables, where the predator and the prey robots are evaluated against a predator and a prey with the same control architecture but they have only one behaviour (hunt or evade). The robots from the middle trophic level has been evolved for 250, 500, 750 and 1000 generations and the standard deviations are $\approx 3.0\%$.

A.1 Catches and evades with different baselines

This section shows in Table 7, the percentage of all catches that the predators that is trying to catch the baseline robots that uses the same control architectures does. Table 8, shows the percentage of all evasions that the preys that is evading the baseline robots that uses the same control architectures does.

Table 7: Percentage of catches for all the predator robot with different baselines

Number of Generations	**250**	**500**	**750**	**1000**
Network shown in Figure	**Number of Catches (%)**			
9A	39.5	36.9	37.7	41.2
9B	11.4	15.6	14.1	15.7
9C	12.3	13.2	13.6	13.1
9D	14.2	14.0	16.4	14.8

Table 8: Percentage of evasions for all the prey robot with different baselines

Number of Generations	**250**	**500**	**750**	**1000**
Network shown in Figure	**Number of Evasions (%)**			
9A	42.6	41.2	45.1	44.5
9B	89.6	88.2	89.1	88.2
9C	72.5	74.6	74.6	74.4
9D	75.3	74.1	76.3	73.5

Appendix B Baseline robots

Table 9 shows the percentage of catches for the baseline predator robots' with the different control architectures and in Table 10 shows the percentage of evasions for the baseline prey robots' with different control architectures. These different baseline robots has been evolved for 250 generations.

Table 9: Percentage of catches for the baseline predator robots

Number of Generations	**250**
Network shown in Figure	**Number of Catches (%)**
9A	49.3
9B	10.1
9C	26.0
9D	23.9

Table 10: Percentage of evasions for the baseline prey robots

Number of Generations	**250**
Network shown in Figure	**Number of Evasions (%)**
9A	50.7
9B	89.9
9C	74.0
9D	76.1

Appendix C Forward moving robots

This section will present the robots that is moving forward (designed direction), in Section C.1 the predator and prey robots are evaluated against robots with the same control architecture, however, only one behaviour. The robots with the different control architecture have been evolved for 250, 500, 750 and 1000 generations and had a standard deviations are $\approx 3.0\%$.

C.1 Catches and evades with different baselines

Table 11 and Table 12 presents the robots that has been evaluated against baseline robots with the same control architectures.

Table 11: Percentage of catches for forward moving predator robots with different baseline robots

Number of Generations	**250**	**500**	**750**	**1000**
Network shown in Figure	**Number of Catches (%)**			
9A	35.0	33.7	33.8	35.6
9B	19.8	15.9	22.3	20.7
9C	11.5	10.6	14.4	13.5
9D	13.8	19.5	16.2	18.6

Table 12: Percentage of evasions for forward moving prey robots with different baseline robots

Number of Generations	**250**	**500**	**750**	**1000**
Network shown in Figure	**Number of Evasions (%)**			
9A	34.9	37.6	34.5	38.0
9B	85.6	87.3	83.7	86.9
9C	71.7	72.8	74.9	74.0
9D	75.8	74.7	73.4	75.0

Appendix D Backwards moving robots

The robots has been evolved for 250, 500, 750 and 1000 generations and got a standard deviations of $\approx 3.0\%$. This section presents robots that are moving around backwards (opposite to the designed direction). Section D.1 presents robots that has competed with baseline robots that uses the same control architecture.

D.1 Catches and evades with different baselines

In this section the tables (Table 13 and Table 14) presents the results from when the robots that are moving backwards competes with robots that uses the same control architecture as they do.

Table 13: Percentage of catches for backward moving predator robots with different baseline robots

Number of Generations	**250**	**500**	**750**	**1000**
Network shown in Figure	**Number of Catches (%)**			
9A	59.2	64.7	61.5	62.3
9B	15.1	17.4	17.3	18.5
9C	11.1	14.2	11.8	12.3
9D	15.4	12.5	16.9	17.3

Table 14: Percentage of evasions for backward moving prey robots with different baseline robots

Number of Generations	**250**	**500**	**750**	**1000**
Network shown in Figure	**Number of Evasions (%)**			
9A	65.5	67.8	66.6	64.9
9B	89.0	86.9	87.7	86.5
9C	70.0	74.4	71.9	71.2
9D	73.4	75.0	76.4	72.3