

Cognitive control in reactive agents: Surviving predators through the evolution of a circadian rhythm

Laurie Bax (0710504)

July 15, 2010

Bachelor Thesis
Author: Laurie Bax
Student number: 0710504
Email: lauriebax@student.ru.nl
Supervisors: Pim Haselager and Ida Sprinkhuizen-Kuyper
Radboud University Nijmegen

Cognitive control in reactive agents: Surviving predators through the evolution of a circadian rhythm

Laurie Bax (0710504)

July 15, 2010

Abstract

A theory of EEC proposes that the brain constitutes the role of a traffic regulator in simple behavior. In simulations was tested whether or not agents with a control structure, which could inhibit behavioral layers, lived longer than simple reactive agents. The world had a circadian day-night rhythm and a predator. By establishing a sleep-wake cycle the control structure agents improved their fitness. Especially when the control structure has a variable energy usage, depending on the amount of inhibition it exerts, the fitness is improved. The agents would also take siestas to save energy during the day. The behavior of the agents did not vary much between the different types of predators.

1 Introduction

For many species having a circadian rhythm is advantageous Ouyang et al. (1998); Berger and Phillips (1995); Toh (2008). It helps them regulate their energy consumption and synchronize it with the periodic factors in the environment. Even when the environment no longer displays the so called ‘Zeitgebers’ Kolb and Wishaw (2006), like the light and darkness of the day and night, organisms are still able to maintain a circadian rhythm Cambras and Diez-Noguera (1991). Such circadian clocks are found in species from single cellular organisms to higher species of plants and animals Ishida (1995); Toh (2008). Establishing of such a circadian clock therefore does not require a central control structure.

The fact that the sleep-waking cycles are often synchronized with the day-night cycles of the environment supports the theory of embodied embedded cognition, or EEC. According to van Dijk et al. van Dijk et al. (2008), EEC proposes that cognition and behavior emerge from the bodily interaction of an organism with its environment. In other words, the physical structure of both body, including the internal structure, and the world put constraints on the behavior of an organism. Although the brain plays a big role in behavior and cognition, it is not the origin of all behavior. It facilitates behavior by acting as a traffic regulator. It helps to select the appropriate responses or to block non-appropriate responses.

The brain acting as traffic regulator was tested by Lagarde Lagarde (2009). In his thesis he investigated in a simple world, whether or not agents with a

simple control system would perform better than purely reactive agents. The world had a day, during which everything was perfectly perceivable and a night, during which the agent had some trouble seeing. He found that with a simple control system, the agents survival rate improved by establishing a circadian rhythm.

In my research I intend to expand Lagare’s research, by adding a predator to the world. I will investigate the following questions:

- Will agents with an inhibitory control system still be more effective than purely reactive agents in an environment displaying a ‘day and night rhythm’ and which has a predator?
- If a circadian rhythm is developed here too, will the rhythms be different depending on the type of predator?
- Will these results still support the notion of the brain as a traffic regulator like EEC proposes?

The first two questions will be addressed by experiments; the third question is more philosophical in nature and will be discussed after the experimentation.

The remainder of this paper will be as follows. First I will describe the simulation environment followed by the parameters used, followed by the hypotheses about the results of the simulations. Then the results will be analyzed and the research questions will be answered. Finally there will be the conclusion and ideas for future research.

2 Methods

In this section the environment within the simulations is specified, and how the different agents and predators work. The base of the design will be the same as in the thesis of Lagarde Lagarde (2009). The differences are described in Section 2.4. The main difference is that a predator is added in the environment and the agents are adapted to be able react to it. The details are described below.

2.1 Environment

The world is divided into squares. A square can be occupied by either an obstacle, food or nothing. The obstacles are like quicksand, they are low to the ground and the agent can get through, but hitting it will result in loss of energy. By eating the food the agent will regain energy. Obstacles will remain on the same place during the simulation, the food will be randomly replaced in the world after the agent has eaten it.

The world will have a day and night rhythm. At daylight, the agent can see everything perfectly clear. At night the agent will have trouble seeing. This is what we all experience ourselves. When there is sufficient lighting we can see (almost) everything, but as soon as it gets dark we lose contrast and some objects blur into the background.

The obstacles can only be seen from close by, one square Manhattan distance. The food can be seen from two squares. If food is located behind an obstacle,

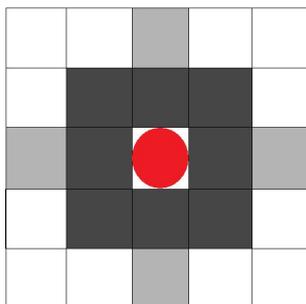


Figure 1: Sensory field of the agents. The circle is the agent, the dark gray squares are the locations where it can see obstacles and food, the light gray squares where it can see food.

the food can still be seen. Figure 1 shows the visual field of the agent. The circle is the agent, the dark gray squares are the locations where it can see obstacles and the light gray squares where it can see food.

The agent also has a sound detector. It can hear the predator from two squares away (Manhattan distance). Lighting conditions do not influence hearing, so the sound detector works the same during the night as during the day. It is less precise however. The agent can hear the predator from two squares away, but only knows its exact location when the predator is just one square away.

Both the predator and the agent can move in all eight directions to an adjacent square. Also they can move through the four boundaries of the environment and reappear at the opposite site of the world. Time will pass in discrete steps. In each step both the agent and the predator can make a move. The agent will always move before the predator does, which gives the agents a slight advantage.

The simulation will end when the agent dies. This can be either through starvation or being eaten.

2.2 Agent design

Agents have to survive in the environment for as long as possible. Survival is based on finding enough food in order not to run out of energy. The agent uses up energy when it moves in the environment and loses a lot of energy when it hits an obstacle. Also, the agents need to avoid being eaten by the predator.

There will be four types of agents but the basis of each agent is the same, namely the subsumption architecture Brooks (1991). The reactive and reactive-DN agents use solely this architecture, the control-1 and control-2 agents have an inhibitory neural network on top of it. This neural network is like the neocortex in mammals and (like the human brain) this mini brain will use extra energy Shulman et al. (2004); Oz et al. (2007).

The four agents that will be used in the simulations are:

- Reactive. A reactive agent with no higher control structures.
- Reactive-DN. A reactive agent with a built-in day and night rhythm. DN stands for Day and Night.

- Control-1. An agent with a higher control structure. The control structure may inhibit behavioral responses of the agents. The cost of having this control structure is some extra energy usage.
- Control-2. An agent with a higher control structure. The control-2 agent is the same as a control-1 agent with the only difference being the extra energy consumption of the control structure. The energy usage of the control-2 agent will be based on the number of behavioral responses inhibited by the control structure.

2.2.1 Reactive

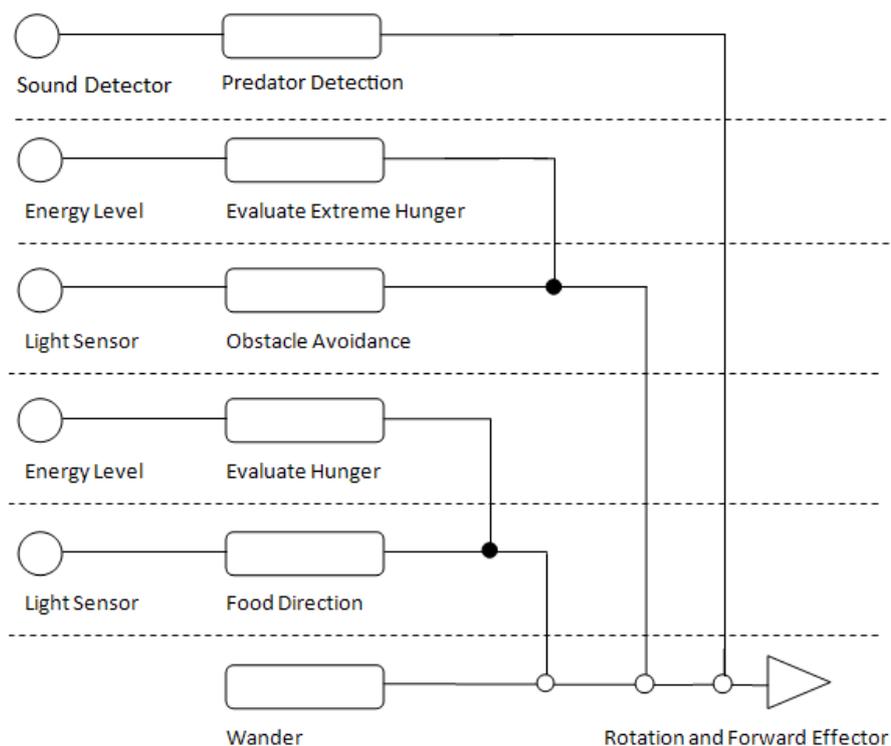


Figure 2: The behavioral layers of the agents

In Figure 2 is the ‘brain’ of the reactive agents displayed.

The subsumption architecture exists out of layers. Each layer individually has to connect sensing to action. It may not use other layers to carry out its action. If this is done properly the layers are not aware of other layers and the system can be built up incrementally. Also, because each layer takes care of its own perception-action coupling, “there is no single place where ‘perception’ delivers a representation of the world in the traditional sense” Brooks (1991).

The layers are connected with each other through inhibitory (black circles) connections. There are different levels within the model. The most basic layer, level 0, is the wander layer. This layer causes the agent to move in a random

direction. On top of this layer, there is the food location layer. It takes the input from the light sensor and, if the agent is hungry, moves towards a nearby food source. The next level is the obstacle avoidance layer. This layer checks if the direction the agent wants to go to is obstacle free. If not, it makes the agent go in an other direction. The same goes for the predator detection layer. If in the preferred direction a predator is observed, the agent will no go into that direction.

During the day this agent should have no problems functioning. However, at night the readings of the light sensor become highly unreliable and the agent will make mistakes.

2.2.2 Reactive-DN

The reactive-DN, or reactive-DayNight agent behaves during the day the same as the reactive agent. The only difference between the reactive-DN and the reactive agent is the behaviour during the night. Where the reactive agent will also move during the night, the reactive-DN agent is programmed to sleep. However, if the agent is extremely hungry it will move and search for food. When the agent is asleep the energy consumption will be reduced to a minimum. The energy usage when asleep will be the same for all agents. The energy costs are displayed in Table 1.

The reactive-DN agents serve as a control group to compare the control agents to if they develop a circadian rhythm. The reactive-DN agent always sleeps at night, even when there is a predator active.

2.2.3 Control-1

The control agents have, as the name suggests, a control structure on top of the reactive layers. It is a inhibitory neural network. How the network is wired with the layer structure, can be seen in Figure 3. The control agents can perform three types of actions: move, rest and sleep. The move action lets the agent move around in the world. It can go to every adjacent space. When the agent does not move, but is alert to the input of at least one of its sensors (sound or light or energy level), the agent is resting. When the agent does not attend to its sensors and does not move, the agent is asleep. The move action will cost the most energy, followed by resting. When the agent is asleep it will use the least energy. The energy costs are displayed in Table 1.

The neural network has 4 inputs. First there is the energy level to monitor the agent. Next are two light sensor inputs, one is to see where the agent is standing on this moment and one to see its surrounding. Finally, there is the sound detector input, for hearing the predator. The output consists of 7 inhibitory links. The network can inhibit each layer separately, or with the seventh connection inhibit the entire subsumption architecture at once, disabling the agent entirely. When the control structure uses this last link, the agent will go to sleep. When it is inhibiting the motors, but not the entire subsumption architecture, the agent is resting. The distinction between sleep and rest was not in the thesis of Lagarde (2009), where the world was static. If the agent did not move nothing else happened. Because there is a predator added to the world now, the world is more dynamic. Even when the agent does not move, the predator may do so (depending on the type of predator, see Section 2.3).

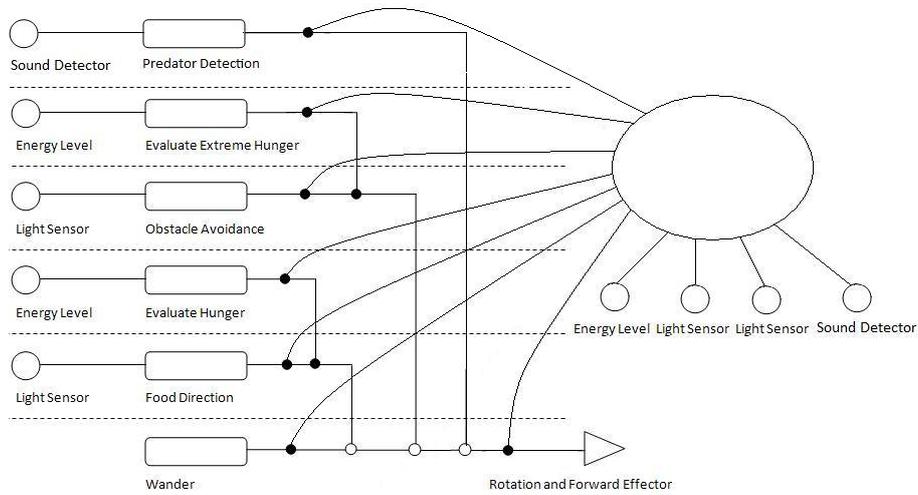


Figure 3: The structure of the control agents. The neural network has an inhibitory link to every layer of the subsumption architecture. The inputs of the control structure are the energy level of the agent, two light sensors (one for the surrounding and one for the current location) and the sound detector. The layers of the neural network are fully connected, resulting in 55 connections (4 input, 5 hidden and 7 output neurons).

Therefore, the agent is presented with the choice of either keeping an eye out for the predator or not.

Between the input and output layer is a hidden layer with 5 neurons. If the number of neurons is bigger, the number of links soon becomes very large and it would take a lot of time to optimize the weights. Also, the fitness did not improve very much if the number of hidden numbers became bigger (see Appendix B, Figure B5). The weights of the network will be optimized by the evolutionary algorithm described in Section 2.5. The activation functions will be the same as in the thesis of Lagarde(2009):

There are three activation functions used: the input layer uses the identity function, the hidden layers uses the hyperbolic tangent (tanh) and the output layer makes use of the logistic sigmoid function Bishop (2006). The logistic sigmoid is defined by

$$\sigma(a) = \frac{1}{1+\exp(-a)}$$

The logistic sigmoid has, as the name implies, an S-shaped curve and maps the whole real axis to the interval [0,1]. This property is then used to interpret the outputs as probabilities. The higher the output is, the higher the probability of activating the output link and the higher the probability of deactivating the behavioral layer to which the output link is connected.

The actual parameters that were used in the simulations can be found in Appendix A.

2.2.4 Control-2

The difference between the control-1 and control-2 agents is only in the amount of energy the brain uses. For the control-1 agents, the energy usage of the brain is always the same, despite the amount of activation. For the control-2 agents is the energy usage dependent on the amount of activations in the brain. If no links are activated, there is no extra energy usage. For each link that is activated, an amount of $\frac{1}{7}$ (one divided by the number of output links) is added to the energy consumption. When all links are activated, the control-2 agent will use the same amount of energy as the control-1 agent. In all other cases, the control-2 agent will use less. In Table 1 are the energy costs of all agents.

Agent	Sleep costs	Rest costs	Move costs
Reactive	-	-	3
Reactive-DN	2	-	4
Control-1	2	3	4
Control-2	1	2 (+ additional costs)	3 (+additional costs)

Table 1: Energy costs in units. A - means that this action is not available for that type of agent.

2.3 Predator design

The predator moves like an agent around in the world. The only task of the predator is to hunt the agent and eat it. It cannot starve to death, so it does not have to worry about its energy level. The predator catches the agent if they are one the same square during one time step. The predator does not have a strategy to catch the agent. It wanders around in the world. If it sees the agent, which will be from two squares Manhattan distance, it will chase its prey until the predator has either caught or lost the agent. During one time step the agent may move first, so the agent has a slight advantage over the predator. The time of day the predator will be active depends on the type of the predator. There are four types:

- stationary predator (Stat)
- diurnal predator (Day)
- nocturnal predator (Ngt)
- continuous or day-night predator (DN)

The stationary predator will remain on the same place the entire simulation. This is done to mimic the simulations of Lagarde Lagarde (2009). The day and night predators will be active on their respective part of the circadian cycle. The other half of the day they will stand still like the stationary predator. The day-night predator will be active during the entire simulation.

2.4 Previous work

As mentioned before, the majority of the simulations are ported from the thesis of Lagarde (2009). As much as possible is kept the same to be able to compare

Agent	Lagarde (2009)	Bax (2010)
Reactive	-/-/2	-/-/3
Reactive-DN	-/2/3	2/-/4
Control-1	-/2/3	2/3/4
Control-2	-/1/2 (+ additional costs)	1/2/3 (+ additional costs)

Table 2: Costs of sleeping/resting/moving. A - means that this action is not applicable by the agent. The costs are higher because of the sound detector the agents have now.

the data of his simulations to these simulations. Therefore, the only changes that are made are tied to the predator. The predator is added to the environment. The sound detector was added to enable the agents to perceive the predator even at night. Because of the predator, the world becomes more dynamic. As mentioned above, for the agents a choice was needed to either keep their sensors ‘open’ or ‘closed’. This is the difference between resting and sleeping.

Also, because of the extra sensor and the choice to keep their sensors opened or closed the agents had a different energy usage per time step than in the thesis of Lagarde. In Table 2 the energy costs are displayed for both this paper and the thesis of Lagarde.

2.5 Evolution

The evolution of the control system agents will go in the same way as in the thesis of Lagarde Lagarde (2009). In the evolutionary algorithm the weights of the neural network are represented in a vector. This is the genome of one agent. Each layer of the neural network is fully connected with the next layer, so there are 55 connections (4 inputs, 5 hidden nodes and 7 output nodes). A population of 125 individuals is created in which each individual has 55 genes. There are no bias weights. The weights of the neural network are constrained to the interval $[-3, 3]$, and are rounded to two decimal places. The individuals of the first generation have random values for the weights. The fitness of a genome is calculated as follows. The genome is converted to a neural network which is given to an agent. In a random environment (with a fixed number of obstacles and food sources) one simulation is ran. This is done for 5000 times in 5000 different environments. The number of time steps the agents survives is counted and averaged over the simulations. There is a maximum of 750 time steps during one simulation, so the fitness is bounded between zero and 750.

For the next generation, first parents are selected based on their fitness. By means of mutation and recombination Eiben and Smith (2007) new offspring is created. The Evolutionary Algorithm is implemented with the use of JGAP Rotstan (2009). The default configurations for mutation rate and selection method provided by this package were used. Details can be found in Appendix A2.

2.6 Simulation

The agents are tested in an environment of 10 by 10 squares, which has 10 food sources and 30 obstacles. The thesis of Lagarde Lagarde (2009) showed

that the differences in fitness of the different agents was clearly distinguishable with these values of the parameters. The obstacles and food sources are placed randomly. Because of the randomness, in one simulation the environment can be very friendly or hostile to the agent. To get reliable results, 5000 simulations are run and thus the agents get tested in 5000 different environments. The results are then averaged. The agents start with an energy level of 250. A full day night cycle consists of 30 time steps: 15 time steps day and 15 night. At the start of the simulation, the predator is placed in the middle of the world and the agent on the first free space from the top left corner. This ensures that the agent and predator always start on about the same distance from each other and that they cannot see or hear each other. A summary of these parameters can be found in Table A1.

2.7 Hypotheses

There are four types of agents and four types of predators. In total, all 16 possible combinations are tested. To compare the different conditions, the results will be averaged according to which aspect will be investigated. The results are about either the differences between the predator types or the agent types. So in each section only four conditions will be reviewed at each time.

2.7.1 Difference between predators

There will be four kinds of predators: a stationary predator, a diurnal predator, a nocturnal predator and a continuously hunting predator. The precise workings of the predators are explained in Section 2.3. The agents will survive the longest when they share their environment with a stationary predator. The reason is obvious, agents can only get eaten by this predator if they walk into it. They can hear the predator very clearly all of the time, so this should not happen at all. If the agents avoid the one spot the predator is they can live out their 'natural' lives.

The fitness will be the lowest when there is a continuously hunting predator. The agents can never really sleep, because they have to stay alert to avoid getting eaten. Therefore the agents cannot conserve energy very well and they die sooner of starvation. Also, the chances of getting eaten are of course bigger.

Finally, the survival rate will be slightly higher when the predator is active during the day than during the night. When the predator is active during the night, the agents will go to sleep during the day. However, this means that they will have to search for food during the night, increasing the chances of bumping into an obstacle. This will cost a lot of energy and therefore their energy supply will deplete quicker than that of the agents that can search for food during the day.

2.7.2 Difference between agents

In this thesis four different kind of agents are tested. One purely reactive agent, a reactive agent that was programmed to sleep during the night (reactive-DN), and two reactive agents with an inhibitory control structure on top of the reactive layers (control-1 and control-2). The difference between the control agents is in

the energy consumption of the control structure. The agents are described in more detail in Section 2.2.

Because the reactive agents have no means of energy conservation, they will probably have the lowest fitness. It will not really matter which predator is active, since they always are alert during both the day and the night.

The reactive-DN agents will do better than the reactive agents, if they do not get eaten during the night. They are able to conserve energy, but they cannot move at night. So when there is a nocturnal predator active, they will probably get eaten very soon.

Because this world is more complex than the one of Lagarde (2009), the differences between the reactive and the control agents can become either bigger or smaller. Nevertheless, the control agents will perform always as least as good as the reactive-DN agent. Because if sleeping at night is the best survival strategy, evolution will find it.

When the differences get bigger the simple inhibitory brain the control agents have can improve the agents behavior better in a more complex and dynamic world. Van Dartel et al. van Dartel et al. (2005) found in their experiments that agents with a multi-layered perceptron, like the control agents in this thesis, “can organize their behavior according to past sensory stimuli”. So if they can ‘remember’ where the predator was just one moment ago, they might be able to escape better.

When the differences get smaller, this brain structure might not be complex enough to cope with an extra dynamic component, namely the predator. If the latter is the case, future research can investigate what structures might be able to do so.

3 Results

Results are gathered for each condition (agent vs. predator) by averaging over 5000 simulations. The other settings of the simulation can be found in Table A1.

3.1 Visual inspection

Here are the most apparent behavior of the agents described. The apparent behavior of the agents does not vary much over the different predators. Therefore, it is only described per agent. If there is behavior that only occurs within a certain condition, it is mentioned separately.

The reactive agent will keep moving all the time. During the day it manages to keep clear of the obstacles. However, during the night it runs into obstacles often. Also, although it runs away from the predator when it sees it, the predator manages to intercept the agents from time to time. This is mostly a coincidence. The reactive agents always act according to their behavior structure, so there is no difference in behavior of the agent when there are different predators.

The reactive-DN agent elicits the same behavior, except for that it sleeps at night. It does not run into obstacles much and therefore it survives longer, if it does not get eaten.

For the control agents are able to stand still and save energy whenever they choose, they use that ability too. When the predator is stationary or hunting

during the day, the control-1 agent looks a lot like the reactive-DN agent. It wanders around during the day and sleeps during the night. When the predator is active during the night or continuously the agents tend to have siesta's, or little day-time naps. At night they stay on the same place too, but sometimes move one or two steps. The control-2 agent shows the same behavior but with slightly bigger siesta's.

Also, there was one curious tactic that the control agents used. The agent would not move a bit. They kept standing still on the same spot the entire simulation. Instead of searching for food, the agents saved energy to stretch out their life as long as possible. This strategy often ended with the agent getting eaten by the predator.

We have to keep in mind that from visual inspection we cannot see the difference between sleep (eyes and ears 'closed') or rest (eyes and ears 'open'). So when the agent does not move, it can either be resting or sleeping.

In the rest of the result section, the data collected by the simulations is used.

3.2 Fitness

To really investigate what the effects of the predators are on the different kinds of agents, the data will be divided into two categories: the agents and the predator. In the first category the data of each type of agent is averaged over the different predators, in the second category the data of the agents is averaged over the types of agents. This results in a bar plot of the average fitness of the agents by type of agent (Figure 4) and by type of predator (Figure 5). In Appendix B (see Figure B1) are the fitnesses shown of every type of agent for every type of predator.

Because over 90 % agents get eaten (see Appendix B, Figure B4) by every predator except for the stationary one, the question arose how well the non-eaten agents performed in comparison to the eaten agents. That is why the data is split up into three groups as well. One group for the non-eaten agents, one for the eaten agents and one with all the agents. Now the eaten agents can be compared to the non-eaten agents as well.

3.2.1 Effectiveness of agents

Remember there are four types of agents in the simulation, as described in Section 2.2. The reactive (R), reactive-DayNight (reactive-DN or RDN), control-1 (C1) and control-2 (C2). Also, as shown in Figure 4, the average fitnesses of the agents are split into three groups. One group represents the average fitness of the non-eaten agents, i.e. the agents that died of starvation, the average fitness of the eaten agents and the average fitness of all agents, regardless of their cause of death. The figures shows that the fitness of the eaten agents is lower than that of their non-eaten brethren. This is due to the fact that the non-eaten agents lived out their 'natural lives'. When an agent gets eaten this will always happen before an agents dies of starvation.

Lets focus on the non-eaten agents first. The control-2 agents clearly live the longest, this was also the hypothesis. The control-2 agents found a way to save energy and therefore keep away the starvation for a longer period of time. The reactive agents were expected to perform the worst and so they did. Since they do not have any means for energy conservation and every other type

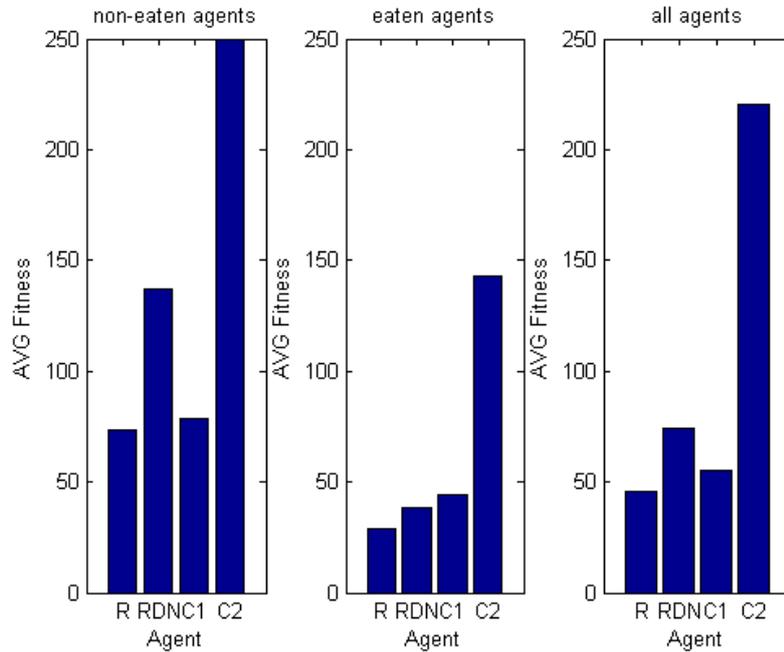


Figure 4: Average fitness of the agents per agent type

of agent does, they have the shortest natural lifespan. The hypothesis states that the control-1 (and control-2) agents should perform at least as good as the reactive-DN agents. This is clearly not the case for the control-1 agents. They almost perform on the level of the reactive agents. This is probably due to the fact that the evolution was cut off too early. Figure B6 shows that the average fitness of the control-1 agents increases a little when the number of evolutions is drastically heightend.

For the eaten agents, the pattern is exactly as the hypothesis predicted. Reactive agents perform the worst, followed by respectively the reactive-DN, control-1 and control-2 agents. Apparently the control-1 and control-2 agents developed a way to keep ahead of the predator for a longer period of time. Especially the fitness of the control-2 agents is relatively high compared to the other agents.

The results of the stationary predator can be compared to the results of the thesis of Lagarde (2009). However, in this experiment the total energy costs are higher. The agents are in these simulations equipped with a sound detector, the agents in the thesis of Lagarde (2009) did not have this. The sound detector has an additional energy usage. Also, the agents in this paper have an extra option: to go to sleep. This allows them to save an extra energy unit. In Table 3 are the average fitnesses shown.

The results of the reactive and reactive-DN agents are very similar to the ones of the agents of Lagarde (see Table 3). The average fitness is in this paper lower because of the extra energy usage. The control-1 agent performs here dramatically bad in comparison to its counterpart in the thesis of Lagarde

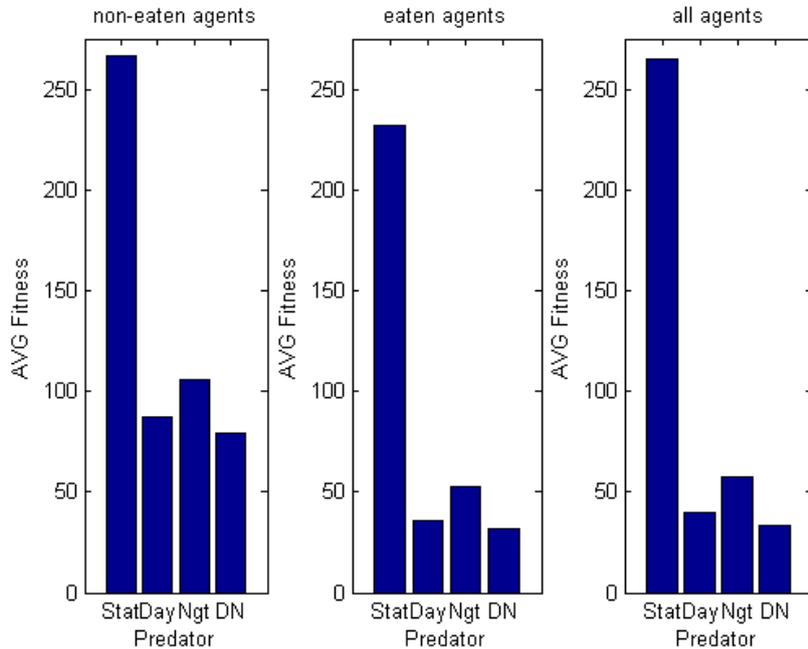


Figure 5: Average fitness of the agents per predator type

(2009). This is probably because of the premature cut-off in the evolution of the control structure. The control-2 agent is not bothered by this and performs even better despite the extra energy usage. The differences here can be explained by the difference in behavior. In the thesis of Lagarde the control-2 agents did have siestas, but they were only about 1.6% of the total behavior. When the stationary predator is added to the environment, the agent will sleep for about 10% during the day (see Figure 8). The agents save more energy and thereby live longer.

Agent	Lagarde (2009)	Stat predator
Reactive	111	82
Reactive-DN	200	175
Control-1	278	99
Control-2	544	705

Table 3: A comparison of the average fitnesses of the agents

3.2.2 Effectiveness of predator

Now we will look what the general influence is of the type of predator. Remember that there are also four types of predators. Stationary (Stat) which remains on the same place all the time, Day which hunts (moves) only during the day, Night (Ngt) which hunts only during the night and the DayNight (DN) which hunts continuously.

In Figure 5 can be seen that the average fitness of the agents is (relatively) large in the stationary predator condition, in comparison to the other predators. The predator does not move so the only way the agents get eaten is when they literally walk into its mouth. The predator avoidance layer can prevent this. Although an agent lives longer when it does not get eaten, some control agents still do manage to get eaten by the stationary predator. This is a very weird phenomenon. Inhibiting the predator avoidance layer only costs energy. It would be cheaper to leave it on. However, because of some reason the control-2 agents inhibit it and get eaten by the stationary predator, like it wants to commit suicide. Besides the fact that the control agents do not want to live anymore, there is one other explanation why this could be happening. It may be the case that the neural network just inhibits some layers when there is enough energy. For the survival of the agent there is no immediate danger of depleting its energy, since the inhibition can be annulled every time step.

In all bar plots of Figure 5 the same pattern is emerging. Besides the stationary predator, the fitness of the agents is the highest for the nocturnal predator, then for the diurnal predator and the lowest for the continuous predator. In the next section is the behavior of the agents analysed, and it shows that all agents sleep at night, despite of the type of predator. So even when the predator is active at night, the agents go to sleep. This means that they have to search for food during the day. If the predator is active during the day, the search for food can be interrupted from time to time by the predator. This may cost the agent lots of energy and reduces the energy income. Preisser et al. (2005) found this effect too and call it the trait-mediated interaction (TMI). TMI causes changes in the phenotype, behavior, etc. of the prey (the traits of the prey), in response of the presence of a predator. Because of the predator, the agents have to change their behavior. For instance, the agents can get disturbed when they search for food. This results in a lower energy income and the energy level of the agents will be depleted sooner, resulting in the death of the agent. A similar effect was found by Schmitz and others (1997); Beckerman et al. (1997); Schmitz (1998). In their experiments they found that grasshoppers switched to a less energetic food source because of the presence of spiders. The spiders mouthparts were glued shut, so the grasshoppers could not get eaten. Like in these experiments, the prey had a decreased energy income.

3.2.3 Conclusion

As expected perform the control-2 agents best of all. A simple inhibitory neural network on top of a reactive subsumption structure can improve the fitness of the agent even in a more dynamic world. However, the energy consumption of the network is also important. This is shown by the fact that the control-1 agents do not perform as well as expected. If they can improve by evolving further, they almost certainly will not come to the level of the control-2 agents, the difference is simply too big.

The fitness distribution over the different predators is a little surprising. The nocturnal predator was not expected to be less effective than the diurnal predator. The secondary effects, i.e. the interruptions while searching for food, were not taken into account in the hypothesis.

3.3 Behavior

Just as in the previous section, the data is split up into the same two categories: division by agent type (Figure 6) and division by predator (Figure 7). The behavior is separated into daytime behavior and nighttime behavior. During every time step the agent can do one of three things: it can move, rest or sleep. Each of these actions are counted for both the night and the day for every simulation. At the end, the ratio of every action is determined. These ratios are displayed in the figures.

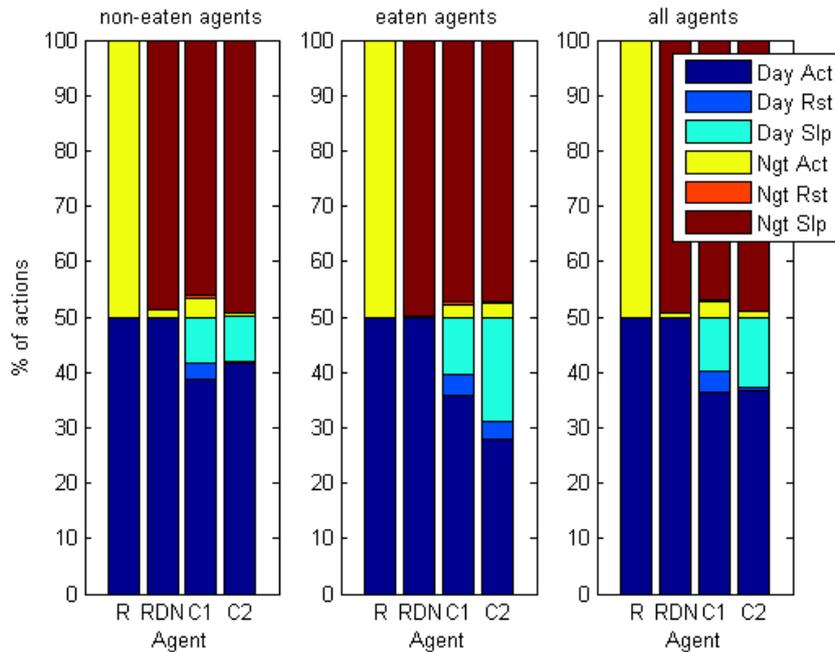


Figure 6: Mean average behavior per agent type

3.3.1 Effectiveness of agent

Again, the non-eaten agents will be looked at first. The behavior of the reactive agent consists only out of actions (movements). The reactive-DN agent moves during the day and sleeps in the night, like programmed. The control agents display a very different behavior than the reactive agents. The control agents sleep, and sometimes rest during the day. This is the same energy saving technique Lagarde Lagarde (2009) found in his thesis. This means that this technique is robust in a more dynamic environment. Also the agents elicit almost no actions at night, they sleep most of the time during the night even when a predator is active. Another remarkable observation is that the agents rest very little, this suggests that resting is not saving enough energy to be profitable and the risk of getting eaten when sleeping is acceptable. With a more effective predator this might be different.

Another striking observation is that the agents that were eaten slept a lot

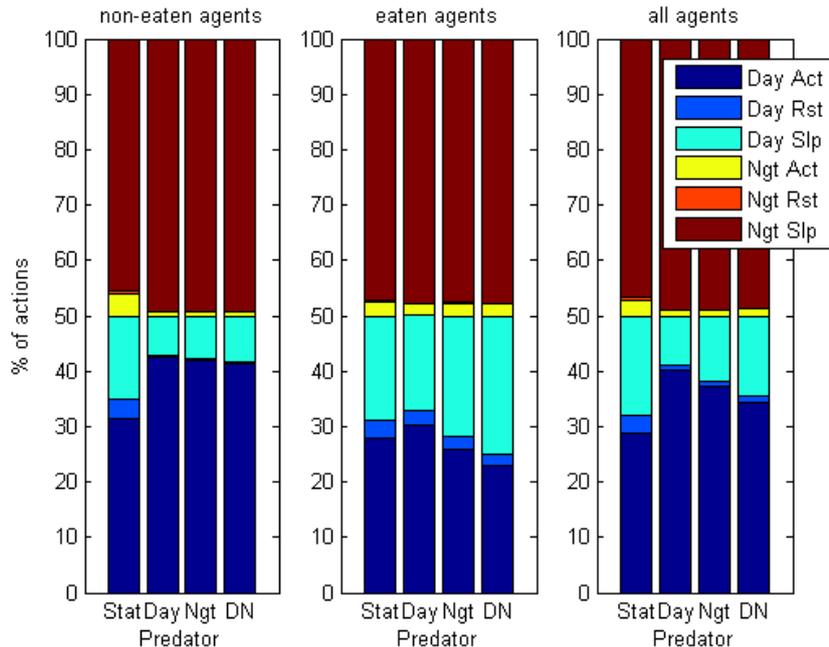


Figure 7: Mean average behavior per predator type

more than the non-eaten agents. The fitness of the eaten agents is lower, and since a great percentage of all agents gets eaten (see Figure 9), this might not be a usable survival strategy after all in this dangerous environment. There were some agents which deployed the curious technique of standing still (and probably sleep) the entire simulation (see Section 3.1). The agents that were visually inspected and used this technique were all eaten by the predator. If enough agents have done this, it could have influenced the behavioral results.

Again, we can make a comparison between the results obtained by the simulations of Lagarde(2009) and the simulations done here. In Figure 8 are the relative behaviors shown.

What immediately pops out is the fact that the control structure agents sleep a lot more in the new environment. Remember the costs to sleep/move were respectively $2/3$ for the control-1 agents and $1/2$ (+additional costs) for the control-2 agents in the thesis of Lagarde. The costs in these experiments are $2/4$ for the control-1 and to $1/3$ (+ additional costs) for the control-2 agents (see Table 2). Sleeping is relatively more energy saving in these simulations. Because it improves the fitness more the sleeping behavior is selected more often.

3.3.2 Effectiveness of predator

As expected the agents sleep the most when they share the environment with the stationary predator. There is no danger of getting eaten so it is safe to close your eyes and ears for a moment. When the predator is only active at night it makes sense to sleep during the day. However, as we saw earlier, all agents sleep at night, despite of whether or not a predator is roaming around. Sleeping

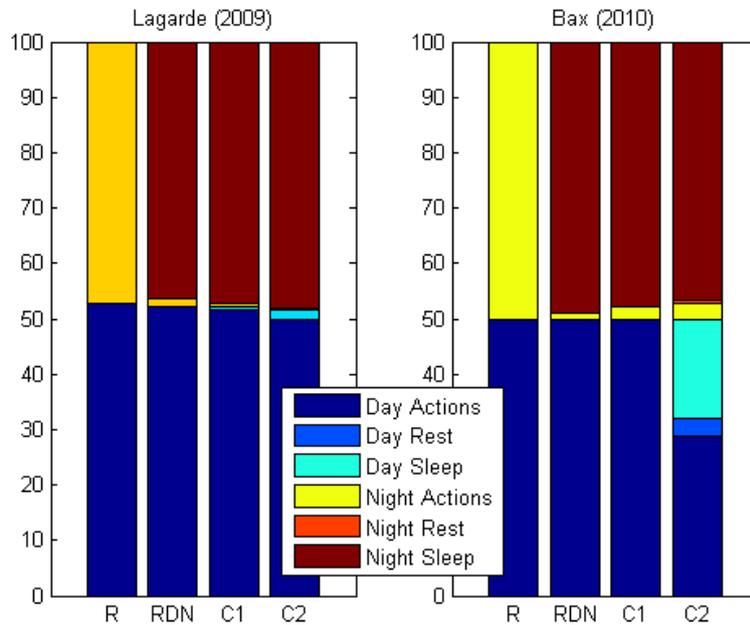


Figure 8: A comparison of the average behavior of agents

during the day acts purely as extra energy conservation, not as a substitute for sleeping at night. This is probably the reason why the agents also sleep during the day when the predator is active both during the day and night. Saving energy by sleeping is worth the risk of getting eaten. However, the agents take the least amount of siestas when the predator is active during the day only. This can have two main reasons: a direct reason and an indirect reason. The direct reason is obvious. If the agents sleep when the predator is active the risk of getting eaten is fairly big, so the agent might try to sleep as little as possible. The indirect reason, trait-mediated interaction (TMI, Preisser et al. (2005)), is less obvious. Like mentioned before, TMI causes the behavior of the agents to change. All agents search for food during the day. If there is a predator roaming around then, the agents might get interrupted in their search. This will result in less energy income and the agents cannot afford much sleeping during the day.

3.3.3 Conclusion

The complex behavior of the control agents is about the same as found in the thesis of Lagarde Lagarde (2009). This means that a circadian rhythm is advantageous even when there is a predator. The rhythm does not adapt to the type of predator (diurnal or nocturnal), however. The agents will always sleep at night and search for food during the day. The daytime naps remain, whether or not there is a predator, so this means the siestas are a fairly robust phenomenon.

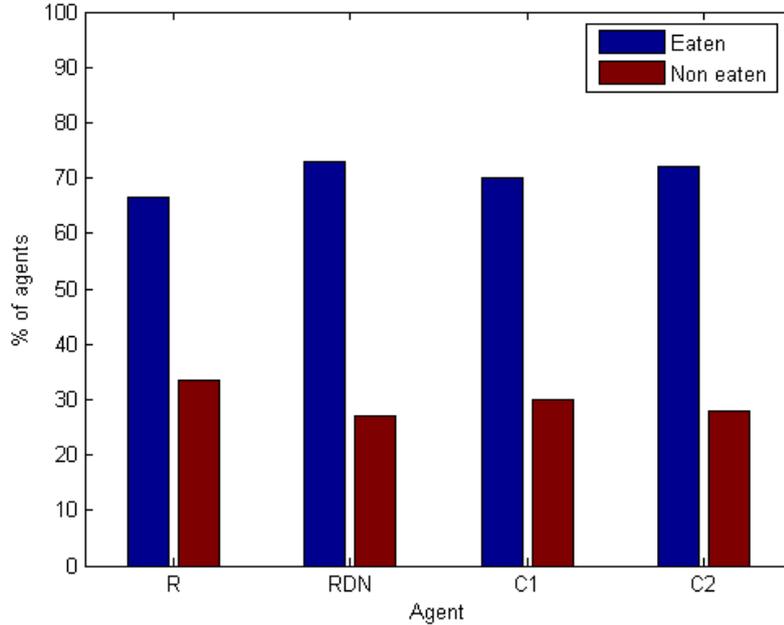


Figure 9: Percentage of eaten and not eaten agents per type of agent

3.4 Embodied Embedded Cognition

The results above show that the control agents do develop a circadian rhythm to increase their fitness. To what extent do these results support EEC? Or at least the notion of the brain as a traffic regulator? To answer these questions both the embodied and embedded part of EEC will be discussed separately. Finally the interaction between the two will be discussed.

3.4.1 Embodiment

The control agents are able to use their sensors and internal states to act in the environment. This supports the embodied part of EEC. When the light sensor does not work very well at night, the control agents adjust their behavior to the poor lighting conditions. They use the information from the sensors to select appropriate behavior. The motor system gets disabled at night to prevent bumping into obstacles they cannot see and thereby they prevent losing energy. Furthermore, the risk of getting eaten by the predator is acceptable in terms of energy costs. The sensors get shut off and this leaves the agent very vulnerable to a predator attack. When the agent does not move and shuts off its sensors it is practically asleep. This results in the day and night rhythm we see in many visually guided animals. The agents use their internal energy level to select their behavior too. A strong example is given by the control-2 agents. As explained in Section 3.3, they sleep during the day too when their energy level is high enough. So the internal state influences directly the higher control structure of the agent.

3.4.2 Embeddedness

The embeddedness of the agents is shown in the reaction of the agents to the different predators. The predator is a dynamic part of the environment and different predators therefore result in different environments. The behavior at night does not vary at all between the different predators, the behavior during the day does. When the predator is stationary, the control agents sleep and rest more. With the diurnal predator the agents sleep the least during the day. As explained in Section 3.3 this can have two main reasons, a direct and an indirect reason.

Also, the agents might not sleep during the day when their energy level is sufficient. The siesta time of the control agents goes up if the predator is active during both the day and night, but the fitness goes down. This may suggest that in this condition the agents sleep to prolong their lives as long as possible instead of searching for food. The same behavior therefore emerges in different environments for different reasons.

Although there are two different types of control structures (control-1 and control-2) the same behavior still emerges. Both control structure agents establish a circadian rhythm and have siestas. Even in different environments which share the most essential feature for this behavior: a light-dark cycle. This supports the embeddedness part of EEC.

3.4.3 EEC

The interaction between this embodiment and embeddedness causes the agents to develop the day and night rhythm. The environment changes and because of that the input of the sensors changes. At night, moving becomes too dangerous because of the unreliable input from the light sensors. Even when the predator is near, moving is not an option. The expected energy loss of getting eaten by the predator (i.e. the chance of getting eaten) is less than the expected energy loss by keeping the sensors on and run away from the predator when it is near. This results in a control system that puts the agent asleep at night.

The notion of the brain as a traffic regulator is supported by the results found in the experiments. By use of the inhibitory control structure and variable costs, the control-2 agents gained the most survival advantage. The control structure is used as a traffic regulator, by assisting in the selection of the appropriate behavior. This survival advantage could have helped the development of more complex control structures in animals.

4 Conclusion and Future Research

Koechlin et al. (2003) proposed a model for cognitive control in the human prefrontal cortex. They proposed that the control is a hierarchical process originating from the premotor cortex. Koechlin et al. state that “there are at least three nested levels of processing” that influence the cognitive control: sensory, contextual and episodic control. Each of these levels are embedded in distinction areas of the prefrontal cortex. This is pictured in Figure 10. This model can be applied partly to the control structure agents. The agents act and adapt their behavior upon the environment (stimulus) and the circadian rhythm (context). To investigate whether or not the simple control structure can handle

the episodic control too, there could be a home base or safe spot added to the environment. In this safe spot, the agent would be safe from the predator. It should be stored on an extra layer of memory. This way there are also three levels of control in the agents. The subsumption architecture would be the sensory control, the control structure the contextual control and the new extra layer would be the episodic control.

Although a great part of the agents will still be reactive, when episodic memory is added there will also representation involved. To resolve this the use of ‘sweet spots’ can be investigated. These so-called sweet spots are also safe from the predator, but the agent does not know this in contrast to the episodic control. It has to find out by evolution that they exist.

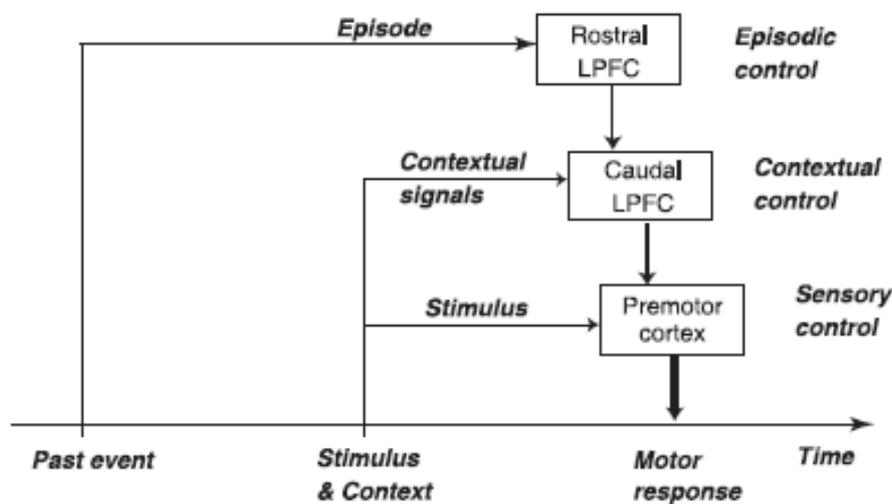


Figure 10: Levels of cognitive control Koechlin et al. (2003)

It was expected that the agents would switch their entire rhythm when they were put in the same environment as a nocturnal predator, since the only defence of the agents is to flee. This was not the case, however. The agents do not have night vision, nor they can develop it. This makes the agents not suitable to move at night, even when it is dangerous to stay put. If there is an option for the agents to evolve their vision in some way, the results may be different. It can also be interesting to investigate how dangerous the predator has to be to obtain a major change in behavior with different predators.

In their paper, van Dijk et al. (2008) talked about four different actions the brain could take to facilitate behavior: suppressing behavior, maintaining behavior, enhancing behavior or favoring actions. The control structure agents could only suppress, or inhibit, different behaviors. A logical expansion of this research is to allow the control structure to do more, like maintaining or favoring certain behavior. This way, the agents may behave differently when the predator is not in sight (or hearing range).

Finally, for future research the role of the agent could be switched from prey to predator. Now, the agent searches for static food sources. It might be interesting to see what the agent would do if it has to chase after its food.

However, as Nolfi and Floreano pointed out Nolfi and Floreano (1998), there is a possibility that the agent does not catch its prey in the first generations. If this is the case all agents are awarded with the same fitness and an optimal solution cannot be found. This can be solved by co-evolving the predator and prey.

References

- Beckerman, A., Uriarte, M., and Schmitz, O. (1997). Experimental evidence for a behavior-mediated trophic cascade in a terrestrial food chain. *Proceedings of the National Academy of Sciences*, 94:10735–10738.
- Berger, R. J. and Phillips, N. H. (1995). Energy conservation and sleep. *Behavioural Brain Research*, 69(1):65–73.
- Bishop, C. (2006). *Pattern recognition and machine learning*. Springer.
- Brooks, R. A. (1991). Intelligence without representation. *Artificial Intelligence*, 47:139–159.
- Cambras, T. and Diez-Noguera, A. (1991). Evolution of rat motor activity circadian rhythm under three different light patterns. *Physiology & Behavior*, 49:63–68.
- Eiben, A. E. and Smith, J. E. (2007). *Introduction to evolutionary computing*. Springer.
- Ishida, N. (1995). Molecular biological approach to the circadian clock mechanism. *Neuroscience Research*, 23(3):231–240.
- Koechlin, E., Ody, C., and Kouneiher, F. (2003). The architecture of cognitive control in the human prefrontal cortex. *Science*, 302:1181–1185.
- Kolb, B. and Wishaw, I. Q. (2006). *An Introduction to Brain and Behavior*. Worth Publishers, 2 edition.
- Lagarde, S. (2009). Evolving a circadian rhythm (ba-thesis). Supervisor: Hase-lager, W.F.G and Sprinkhuizen-Kuyper, I.G.
- Murphy, R. R. (2000). *Introduction to AI Robotics*. The MIT Press.
- Nolfi, S. and Floreano, D. (1998). Co-evolving predator and prey robots: Do ‘arms races’ arise in artificial evolution? *Artificial Life*, 4:311–335.
- Ouyang, Y., Andersson, C. R., Kondo, T., Golden, S. S., and Johnson, C. H. (1998). Resonating circadian clocks enhance fitness in cyanobacteria. *Evolution*, 95(15):8660–8664.
- Oz, G., Seaquist, E. R., Kumar, A., Criego, A. B., Benedict, L. E., Rao, J. P., Henry, P.-G., Moortele, P.-F. V. D., and Gruetter, R. (2007). Human brain glycogen content and metabolism: implications on its role in brain energy metabolism. *American Journal of Physiology: Endocrinology and Metabolism*, 292:E946–E951.

- Preisser, E. L., Bolnick, D. I., and Bernard, M. F. (2005). Scared to death? the effects of intimidation and consumption in predator-prey interactions. *Ecology*, 86(2):501–509.
- Rotstan, N. (2009). Jgap: The java genetic algorithm package. Available from <http://jgap.sourceforge.net/> (Version 3.4.3).
- Schmitz, O. (1998). Direct and indirect effects of predation and predation risk in old-field interaction webs. *American Naturalist*, 151:327–342.
- Shulman, R. G., Rothman, D. L., Behar, K. L., and Hyder, F. (2004). Energetic basis of brain activity: implications for neuroimaging. *TRENDS in Neurosciences*, 27(8):489–495.
- Toh, K. L. (2008). Basic science review on circadian rhythm biology and circadian sleep disorders. *Annals academy of Medicine*, 37(8):663–668.
- van Dartel, M., Sprinkhuizen-Kuyper, I., Postma, E., and van den Herik, J. (2005). Reactive agents and perceptual ambiguity. *Adaptive Behavior*, 13(3):227–242.
- van Dijk, J., Kerkhofs, R., van Rooij, I., and Haselager, W. F. G. (2008). Can there be such a thing as embodied embedded cognitive neuroscience? *Theory & Psychology*, 18:297–316.

Appendix

A Parameters

A1 Simulation parameters

Here are all the parameters used to run the simulation. These values for the parameters were chosen on the basis of the data from the thesis of Lagarde Lagarde (2009). The number of hidden neurons were did determine by trial (see Figure B6). If the value is higher, the results did not differ a lot from the lower values, but the computing time did take a lot longer.

Map size	10 x 10
Number of obstacles	30
Number of feeding places	10
Agents initial and maximum energy	250
Maximum number of time steps	750
Number of day-/nighttime steps	15/15
Maximum number of simulations	5000
Number of input neurons	4
Number of hidden neurons	5
Number of output neurons	7

Table A1: Simulation parameters

A2 JGAP Settings

The settings of the JGAP algorithm are the same as in the thesis of Lagarde Lagarde (2009). JGAP is an open-source project which provides a robust Java framework for Genetic Algorithms. The default configuration provided by this package was used to run the evolutionary algorithm. The specific values and methods used in the algorithm are displayed in Table A2.

Population size	125
Representation	Integer Vector
Mutation	Addition of value in interval [-1,1]
Mutation probability	$\frac{1}{12}$
Recombination	One-point crossover
Recombination probability	0.5
Parent selection	Roulette wheel
Survivor selection	Generational
Termination condition	400 generations

Table A2: Simulation parameters

B Data

B1 Fitness

In Figure B1 are the fitnesses shown of every type of agent for every type of predator. In Section 3.2 this data is averaged over either the agents (Figure 4 or the predators (Figure 5). The figure is clipped to make it easier to see the differences between the agents. The control-2 agents (C2) had a fitness of almost 750 (full length of the simulation) when they shared their world with the stationary predator.

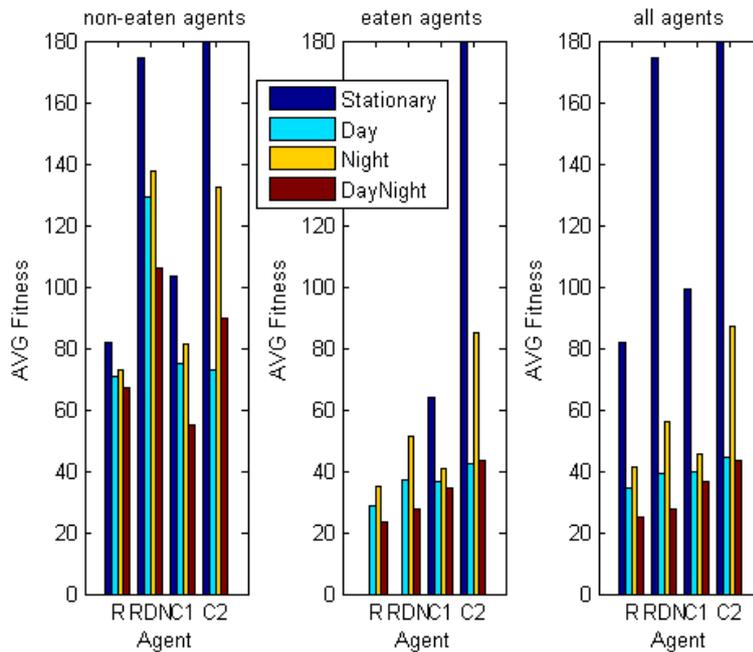


Figure B1: The fitnesses of every type of agent for every type of predator. The control-2 agents (C2) had a fitness of almost 750. To make the differences easier to observe the plots were clipped.

B2 Behavior

In Figure B2 is the average behavior shown of every type of agent for every type of predator. In Section 3.3 this data is averaged over either the agents (Figure 6) or the predators (Figure 7).

B3 Cause of death

In Figure B3 are the absolute numbers displayed of how many agents are eaten per predator. There were 5000 simulations per condition, so the maximum number of a type of agent a predator could eat is 5000. Not all agents got

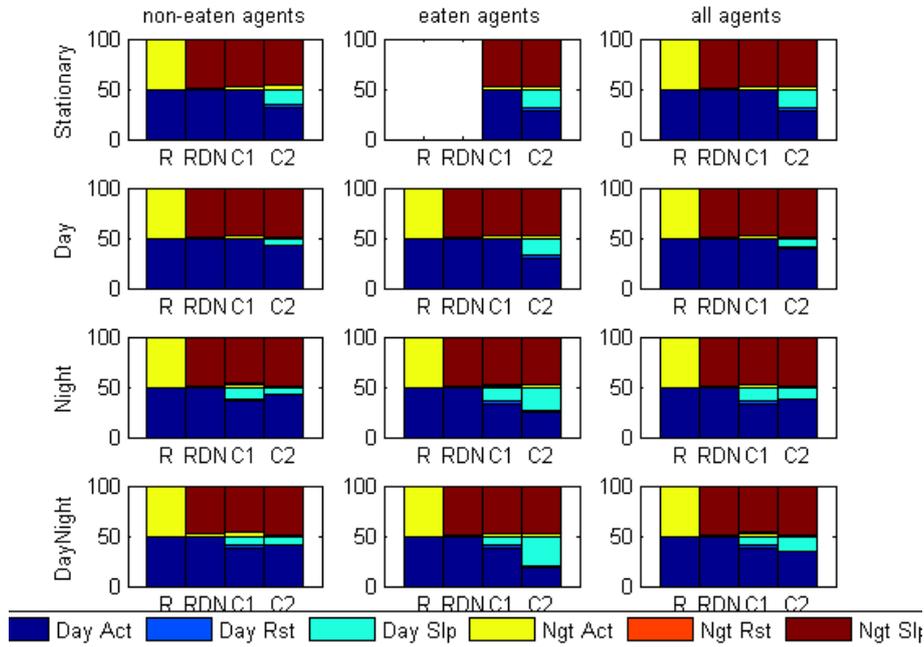


Figure B2: The average behavior of every type of agent for every type of predator.

eaten, the remainder died of starvation or, in a single case, survived for the full 750 time steps.

In Figure B4 is the number of eaten agents averaged over the agents. In Figure 9 is the number of eaten agents averaged over the predators.

B4 Simulation parameters

In Figure B5 is the development showed of the average fitness of the control agents. The fitness is averaged over every type of predator. The control-1 structure agents have almost no overall benefit from an increasing amount of hidden neurons. The control-2 agent have benefits from hidden neurons, but it stabilizes very soon.

Because the control-1 agents did not perform as well as expected, they did not get the same fitness as the reactive-DN agents (see Section 3.2), the control structures of the agents were evolved for more generations. The control-2 structure has almost no benefit of the prolonged evolution. After 400 evolution the control-2 structure has found the best survival strategy. The control-1 agents do have a benefit. This still does not compare to the average fitness of the reactive-DN agents (see Figure 6), but it shows that it might reach that level eventually if the evolution goes on even longer.

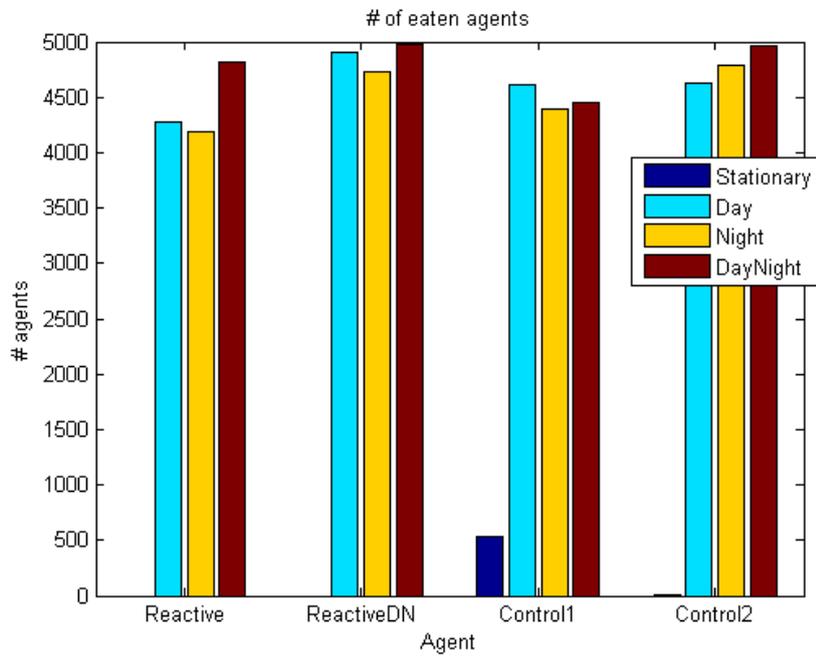


Figure B3: Number of eaten agents per agent type per predator. The number of agents in each condition was 5000.

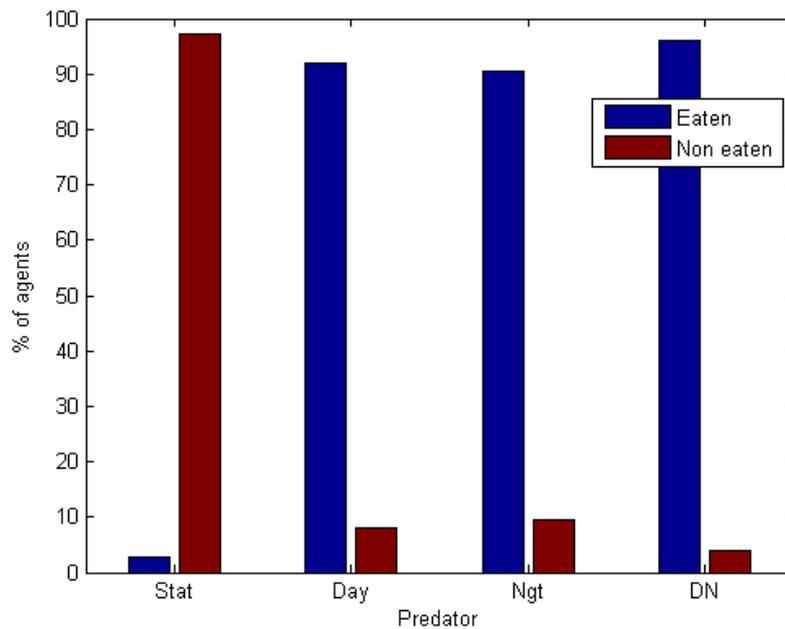


Figure B4: Percentage of eaten agents per predator type

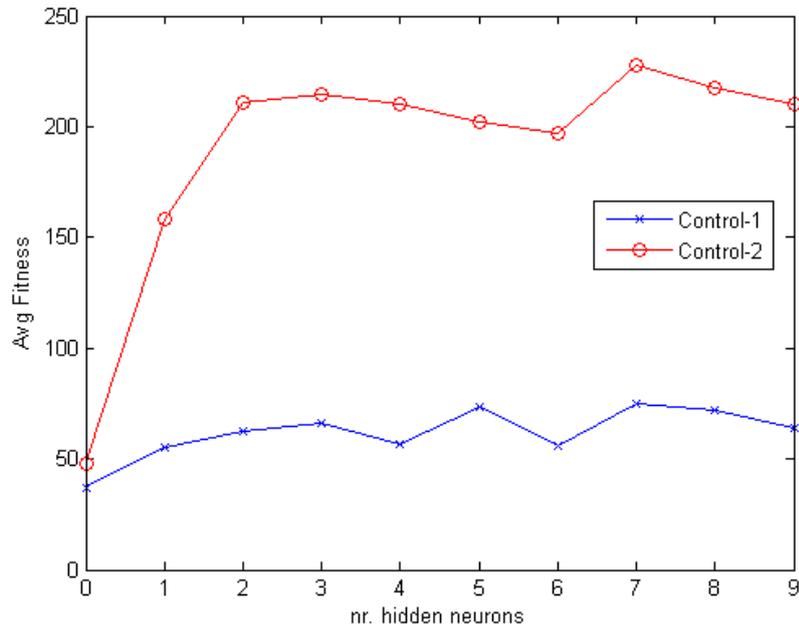


Figure B5: Average fitness of agents by number of hidden neurons. The fitness was averaged over the predators.

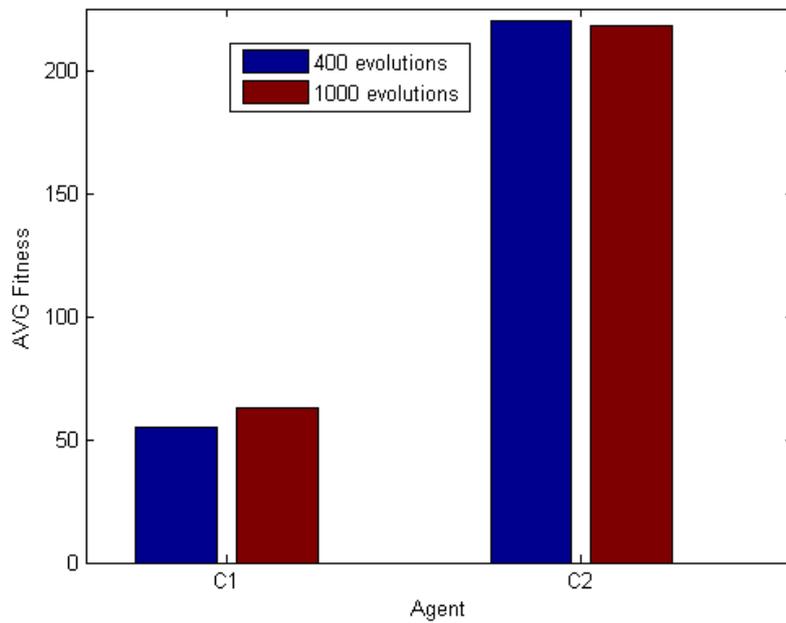


Figure B6: Fitness of agents with more evolutions. This data is also averaged over all the predators.