A Biologically Inspired Solution for an Evolved Simulated Agent

Edgar Bermudez Contreras
School of Informatics
University of Sussex
Falmer, Brighton. UK.
ejb21@sussex.ac.uk

ABSTRACT

Biologically inspired designs can improve the design of artificial agents. In this paper we explain and explore the role of directional light sensors from an Evolutionary Robotics perspective using a dynamical systems approach. It was found that by using directionally specific sensors in the agent, there was a simplification of the neural controller employed. This simplification helped not only with the analysis of this type of controller but also improved the behavioural performance of the agents, thereby showing a good example of the ecological balance principle.

Categories and Subject Descriptors

I.2.9 [**Artificial Intelligence**]: Robotics Autonomous Vehicles; I.2.9 [**Artificial Intelligence**]: Robotics Sensors

General Terms

Design

Keywords

Evolutionary Robotics, Adaptive Behaviour, Evolution of Sensors

1. INTRODUCTION

Eyes have always been as important in animals as they are today. The appearance and improvement of eyes coincided with the increase in size and speed in animals around 530 million years ago [14]. With the development of visual systems, predation became a way of living.

Visual processes constitute a large part of the processes in the brain. Therefore, the comprehension of how vision works can help us to better understand neural processes. Studying the visual processes in the brain can also help with the design of artificial visually guided agents.

In order to understand visual systems in animals, it is useful to start with the study of very simple visual systems, like photoreceptors and primitive structures (i.e. those that came before the existence of eyes as we now know them).

In this paper, we employ an evolutionary robotics (ER) approach to study a very simple primitive visual system in an agent performing phototaxis. This approach has proven to be useful not only in the study and design of robot controllers [9], [7] but also in shedding some light on the understanding of cognitive phenomena (see [1], [2], [11], [16]) or to explore vision morphologies [4] and visual properties of sensors [12], [15].

In this work, we show some of the advantages of following an evolutionary process in the design of visual sensors for artificial agents. While simple hand-designed controllers (for example, Braitenberg vehicles or controllers that feed the motors with the difference of the sensor activity) can provide solutions for simple tasks, the purpose of this work is not to find a quick solution to the task but to understand the interaction between agent and environment. In order to analyse the role of sensors, we used a dynamical systems approach. The dynamical systems approach is a young but promising tool that can help us to understand cognition [3], [10], [17].

In particular, it was shown how directional light sensors (the most primitive precursors to eyes) can not only simplify the neural controller required for phototactic behaviour in an agent, but also improve the navigational strategies exploited. These advantages can help not only with the understanding of the dynamics of robot controllers but also can shed some light on the design of better controllers for visually guided robots or agents. This work can stand also as an example of the principle of ecological balance [18], showing that sensors play an important role in the nature of the control required to solve a simple task.

In the next section we describe the details about the arena, the agent architecture, the genetic algorithm (GA) and the task for the simulated agents in this work. This is followed by a detailed explanation of the experiments and scenarios in section 2. Next we present the results found in this work for the different configurations of sensors and neural controllers employed. Following that, an analysis and discussion about the role of biologically inspired sensors and the neural dynamics of the controllers are presented in section 4. Finally, the conclusions and subsequent possible lines of research for this work are outlined in section 5 and 6 respectively.

GECCO'07, July 7–11, 2007, London, England, United Kingdom.

2. METHODS

An object was placed in the center of an unlimited extension arena. This object emitted a signal dispersed according to the inverse of the distance from the object. For the purposes of this work, we are going to assume that this signal is light. An agent was placed in a random position within an area of 10 × 10 units around the object. At the beginning of every run, the agent started with a random orientation. Agents were evolved to carry out phototaxis, that is, to approach a source of light (object) guided by the intensity of the light emitted by the object.

2.1 The agent

The agent had a circular body with radius of 0.5 units and two wheels on each side driven by independent motors. The agent was able to sense the signal emitted by the object through two sensors placed at $\pm\pi/4$ radians from the line of orientation of the body (see figure 1). The body of the agent was symmetrical with respect to the axis of orientation.

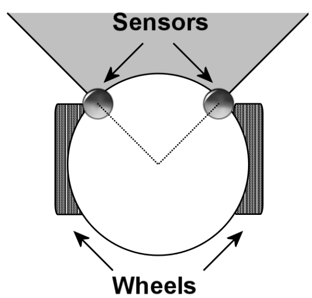

Figure 1: Agent body: two wheels on each side driven by independent motors. Two sensors placed at $\pm\pi/4$ radians from the line of orientation of the body. The body has a radius r of 0.5.

2.2 Controller

The controllers for the agent are Continuous Time Recurrent Neural Networks (CTRNN). These kinds of artificial neural networks show desirable properties as robot controllers for several reasons. First, a CTRNN shows rich complex dynamics as a universal approximator (any smooth dynamical system can be approximated by a CTRNN with any degree of accuracy) [10], [8]. Second, CTRNNs are biologically inspired so that the analysis of this type of model can help to understand the dynamics of real brains (see [1] for a detailed study of their characteristics as neural controllers).

The state y of neuron i changes in time according to the differential equation:

$$\tau_i \dot{y}_i = -y_i \sum_j w_{ij} \phi(y_j + \beta_j) + g \cdot I$$

That is, the state of each neuron is the integration of the weighted sum of all incoming connections (plus a gained input $g \cdot I$ for input neurons). The time constant ϕ is the sigmoid activation function, $\tau \in [0.2, 2.0]$ and the bias $\beta \in [-10, 10]$ and all the weights $w_{ij} \in [-5, 5]$ are shaped by the GA.

Initially, the controller consisted of eight neurons, specifically, two sensor neurons, four fully connected interneurons and two motor neurons. Another set of experiments was carried out using a neural controller with six neurons, two sensor neurons, two interneurons and two motor neurons (see figure 2).

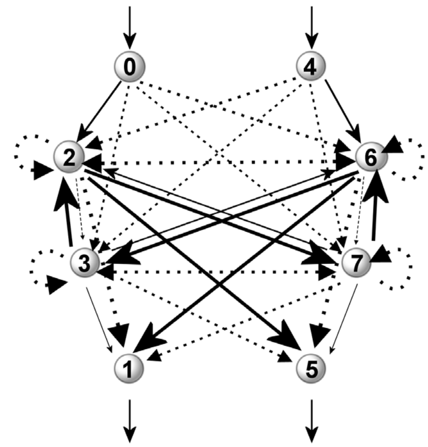

Figure 2: Neural controller: a CTRNN with 8 nodes. Neurons 0 and 4 are the sensor nodes, neurons 2, 3, 6 and 7 are fully connected interneurons and neurons 1 and 5 are the motor neurons. The width of each arrow represents the strength of the connection (weight). The solid lines in the arrows represent excitatory connections and the dotted lines in the arrows represent inhibitory connections.

The sensor neurons were activated by light (sensed as the inverse of the distance between each sensor in the agent and the object).

The motor neurons received activation from every interneuron only. The output of the motor neurons was connected to the motors of the wheels of the agent with a gain of 2.

Due to the nature of the body of the agent, a bilateral symmetry was imposed on the neural controller of the agent. This characteristic is very important for the dynamics of the neural activity (this is explained in the next section).

2.3 Genetic Algorithm

A distributed GA was employed to evolve the neural controllers for a phototactic task. A population of 400 individuals was evolved with mutation probability of 80% and 20% for mutation for each component. There was also a 5% probability of crossover and an elitism probability of 80%.

The genome of each individual was coded in a real vector of 25 elements, 4 for the time constants of each neuron, 4 for the bias of each neuron, one for the sensor gain and 16 for the weights. Each element was coded as a real number in $[0, 1]$ and linearly scaled according to the parameters previously described.

The fitness function F for this task was defined as:

$$F = \frac{1}{d_f}$$

where d_f is the distance from the agent to the object at the end of the trial. In this way the evolutionary pressure was towards individuals finishing as close as possible to the light source.

3. RESULTS

The controllers were evolved to perform phototaxis in order to explore the role of different types of light sensors and different configurations in the neural controllers. It is important to mention that controllers were also evolved using light sensed as $1/d^2$ under the same scenarios showing qualitatively the same results as the ones presented in this section.

3.1 Panoramic light sensors

The first set of experiments was carried out using unrestricted sensors (panoramic sensors), meaning that the sensors could detect light coming from any direction (including light coming from behind the agent).

After several thousands of generations, the evolved agents performed phototaxis successfully. The agents exhibited different approaching behaviours (i.e. searching and staying around the object). Most of the successful agents approached the object in a straight line and then remained close and continuously circled or "patrolled" it (see figure 3).

After finding controllers that could solve the task, we performed tests to understand the behaviour of the agent and its interaction within the environment. During the test runs, the neural activity of the controller and the positions of the agent were stored.

In order to understand the behaviour of the agent we examined the internal dynamics of the best evolved controllers during a test run. The neural dynamics corresponding to the trajectory in figure 3 (presented in the inset 3.A.) can explain the behaviour of the agent during that test run. Around timestep 190 (see figure 3.B), the agent gets very close to the object. However, the sensor neurons were saturated well before this and remained so even after the agent passed the object. If we observe what happens after timestep 190, when the agent passes the object (so the distance increases), we can see that just before timestep 300 (see the output of the neurons 1 and 5 in figure 3.A) the agent starts to change direction and returns to the object. This time corresponds to the time when the sensor neurons start to be deactivated again (so the searching behaviour is once again triggered).

How can we explain the behaviour of the agent? How are these decisions taken? The numerical analysis of the 8 differential equations that describe the system seems too complex to be able to explain the behaviour of the agent in general. Also, it seems too difficult to analyse the network structure of the controller (figure 2) as a way to explain the interactions and roles of the neurons. However, we can have a better idea about how the evolved controller works if we study the different situations that the agent can be in.

Apart from describing the dynamics of the neural activity during a test run, we created a long term steady state map. This map consists of the neural activation of each neuron after a period of time for every position of the object when the agent remained fixed in the center of the arena.

If we place the agent in a fixed position and we translate the object around it we can see how the neural activity responds to that location of the object. That is, we place the agent in a fixed position (not allowing the agent to move) and place the object a certain distance away and, after a period of time, examine the state of the neurons. In this way, we move the object around the agent with steps of one degree, describing a circumference. We store the neural activation for each step and then repeat this process. Once we cover all the positions around the object, we increase

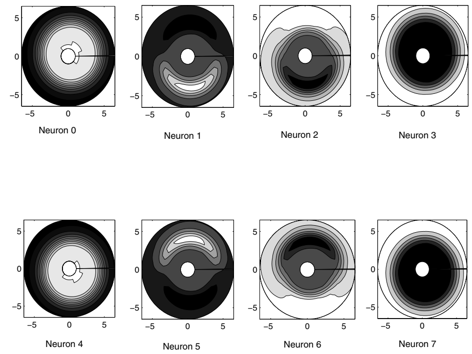

Figure 4: Long term steady state of the neural controller: the agent was fixed in a position facing right (indicated by a line) and the object was moved around it. After 50 timesteps the activation of each neuron is stored. White regions represent 1 in the output of the neuron when the object is in that position and black regions represent 0 in the output of the neuron when the object is in that position.

the radius of the circumference and start placing the object around the agent again and so on. Finally we plot the activation of each neuron for each of these positions in a steady state map (shown in figure 4). The white regions represent high activation and black regions represent no activation.

With this map we can observe the activation of each neuron for all the situations that the agent can find the object (within a limited range). For instance, when the object is to the left of the agent (see figure 4), the left sensor neuron is very active (white) and, as a consequence, the left motor neuron is inhibited (black) and the right motor neuron is excited (white), so the agent turns to the left. Following that movement, once the object is in front of the agent, the activity of the neurons is uniform and the agent goes in the direction of the object (the analogous situation happens when the object is located to the right of the agent). This map gives us a good idea about the way the motors and sensors interact, but still we can not explain exactly how.

Even when we can describe and explain the situations that the agent faces in its environment, the explanation of how the agent decides to turn and navigate, remains unclear. Can we reduce the dimensionality of the controller so we can fully analyse it and explain its neural dynamics?

3.2 Directional light sensors

Visual systems in animals evolved from primitive simple light detectors into directional and spatial light sensors when motility represented a great advantage with the appearance of predators [6], [13], [14]. Mimicking this evolutionary visual sensor development in animals, we restricted the light sensors to be directional. That is, they can sense an object if it is within an angle range. As seen in figure 1, the agent can sense an object if the object is within the grey area. Considering this restriction, we evolved controllers in analogous conditions to the last experiments.

After restricting the sensory activity to a particular angle,

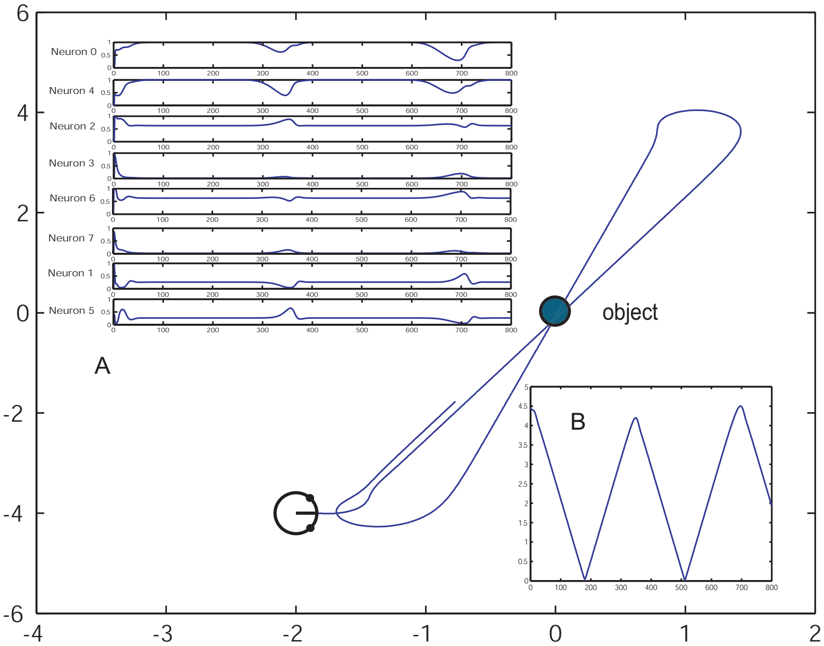

Figure 3: Dynamics of an evolved controller of 8 neurons using panoramic sensors. A. Positions of an evolved agent during a test run. The object (target) is placed in the center of the arena (0,0). [A] Neural dynamics during a test run. Neurons 0 and 4 are sensor neurons, neurons 1 and 5 are motor neurons and neurons 2, 3,6 and 7 are interneurons (hidden layer). [B] Distance between the agent and the object during the test run (timestep vs distance).

and finding successful agents, the best evolved controllers were tested again systematically. In figure 5 we can observe that, for these evolved controllers, the neural activation was different.

We observe that in this case, the agent has to be more active (exploratory, wandering around to locate the object) in the environment. In this situation, the agent cannot always sense the object (these sensors are more limited) so the agent has to sweep and find the object. The neural dynamics show this behaviour, as they oscillate much more. However, in this particular experiment, some of the neurons are saturated for the majority of the time (see neurons 3 and 7 in the figure 5). This saturation suggests that these neurons might be redundant.

To try to find a simpler controller that could solve the task, a 6 neuron controller was evolved to perform phototaxis in the same conditions described above but using restricted (directional) sensors. The successful evolved controllers had a simpler neural structure, making it easier to analyse (see figure 6).

The behaviour for the evolved controllers was similar to the ones found in the previous experiments. In general, it consisted of approaching the object and then describing a "flower" navigational pattern (see figure 7).

Now, with a simple neural controller, can we analyse its neural dynamics and finally fully explain the phototactic behaviour of the evolved agents?

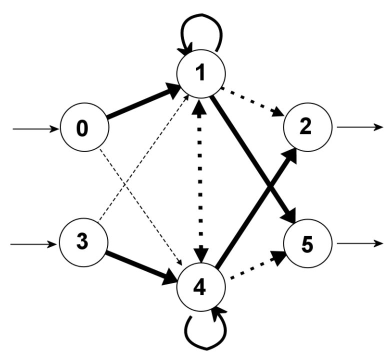

Figure 6: Simplified neural network: 6 nodes. Only two interneurons. The dotted line arrows represent inhibitory connections and solid arrows represent excitatory connections. The width of the arrows is proportional to the strength of the connection.

4. DISCUSSION

First we must examine how the agent makes its decisions.

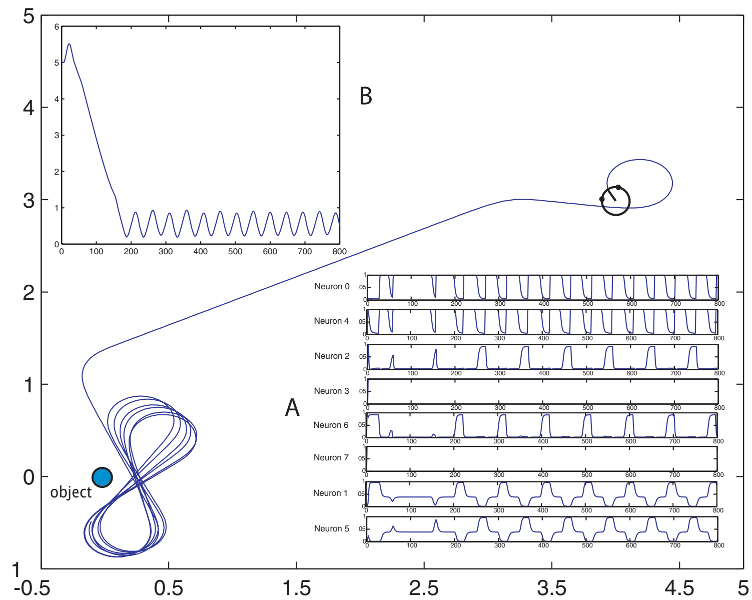

Figure 5: Dynamics of an evolved controller of 8 neurons using directional sensors. Positions of an evolved agent during a test run. The object (target) is placed in the center of the arena (0,0). [A] Neural activity during a test run. Neurons 0 and 4 are sensor neurons, neurons 1 and 5 are motor neurons and neurons 2, 3,6 and 7 are interneurons. [B] Distance between the agent and the object during the test run.

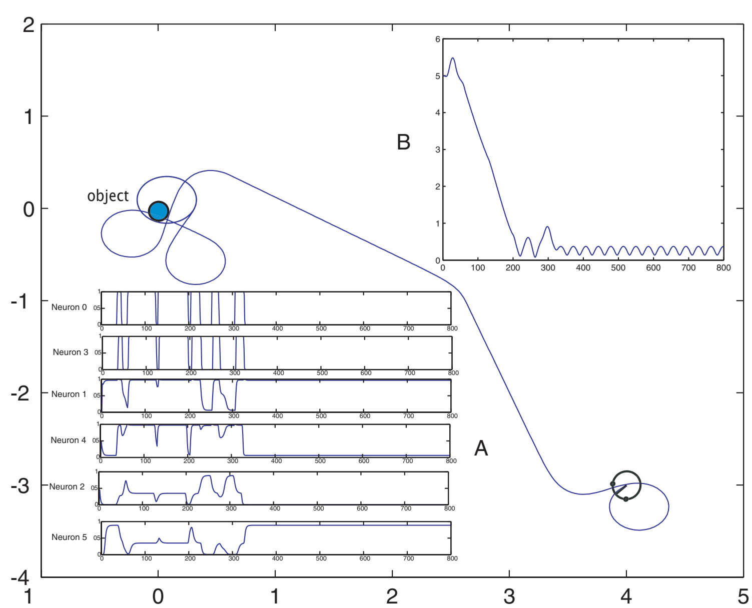

Figure 7: Dynamics of an evolved controller of 6 neurons using directional sensors. Positions of an evolved agent during a test run. The object (target) is placed in the center of the arena (0,0). [A] Neural activity during a test run. Neurons 0 and 3 are sensor neurons, neurons 1 and 4 are interneuron and neurons 2 and 5 are motor neurons . [B] Distance between the agentt and the object during the test run.

To do so we have to analyse the situations the agent faces. In this instance, the interaction between the agent and the environment now can be described by two general situations. The first one being when the object is not within the visual field and the second one when it is within. What happens when the object is not present in the visual field? The agent spins around "searching" for the object. Once the object is within the visual field, the agent approaches it following a straight line.

To study the first case, we did not put the object in the arena. The agent was only spinning if the neural activity was randomly initialised, otherwise it would go in a straight line (due to the bilateral symmetry of the controller). The hypothesis is that this behaviour corresponds to an attractor of the dynamical system described by the differential equations in the controller. The sensory input would perturb this attractor until it reaches the point where it has similar activation in both sensors and it would then proceed in a straight line until it passes the object. Having no sensory input, which would follow in the attractor (first case), the agent would start spinning around until it was able to find the object again.

The most important aspect to explain is how the agent decides when to turn and navigate to find the object. This should be described by modulation of sensory inputs and motors. This modulatory process has to be done by the interneurons since there are no feedback connections in the controller.

Due to the simplification of the neural controller we first analyse how these processes take place in the controller with 6 neurons and directional sensors. The behaviour of the agent can be divided into the same two main situations: 1) when the object is not within the field of view and 2) when the object is within the field of view.

For the first situation with the simple controller (6 neurons and directional sensors), neurons 1 and 4 play the modulatory role. Due to the symmetry and mutual inhibition (see figure 6), if the sensor activity is different (due to random initialisation), the interneuron that is more active overcomes the other one until they reach an equilibrium point. This equilibrium point is then reflected in the motors and finally translates into a steady spinning of the agent. This is an attractor of the dynamical system described by the differential equations of the neurons. This spinning behaviour helps the agent to find the direction of the object.

At some point after the agent is spinning, the object will be within the field of view (situation 2) and the activation of the sensor neurons will be different. This difference will perturbate the equilibrium point that causes the aforementioned spinning behaviour. For example, if neuron 1 was more active than neuron 4, the spin would be towards the left until the object was within the field of view (see figure 7A). At this point the right sensor is increasingly more active and neuron 2 starts to dominate until both neurons reach an equilibrium point. This regulatory process also happens to the motors so the agent starts to go in straight line. In this way the agent changes direction towards the object.

Due to the symmetry of the controller, if the activation is the same in all neurons, the agent is going to go in a straight line indefinitely. However, any difference in the sensors can break that symmetry (due to the mutual inhibitory connection in the interneurons). So this repellor (see [19] for a good text about dynamical systems) occurs only when the line of direction of the agent intersects with the center of the object (so the activation of the sensors is exactly the same). In order to avoid this situation in the controller, it is important to add noise into the sensors.

How did the agent try to find the object once it had passed it? If the object is behind the agent, the agent can not sense it, so the agent has to have a mechanism to find the object again. In principle, this situation seems to be more difficult than the one in which the object evokes a sensor signal, regardless of the orientation of the agent (which was the case of the panoramic sensors), so it seems natural to think that this controller should be more complicated. However, since the dimensionality was reduced it is easier to analyse the dynamics of the evolved controllers in this case.

Why was the restriction of the sensors useful in the reduction of the dimensionality? As mentioned above, it was considerably more difficult to find a successful controller that had 6 neurons and used panoramic (unrestricted) sensors. However, restricting the sensors (intuitively more limited) resulted in the requirement of less neurons to solve the problem. In fact, the problem of localising the light source (object) is more complex for panoramic sensors, since the agent has to disambiguate the two possible positions of the object. That is, the activation of the sensors would be exactly the same for situations where the object is in front or behind but with the same distance from the agent (see figure 8).

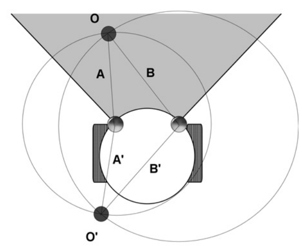

Figure 8: Ambiguous situation: the activation in the sensors when the object O is in front of the agent, is equivalent to the activation generated from the object O'. The distance from O to the sensors is the same as the distance from O'. That is A = A' and B = B'.

In order to discriminate these two possible positions of the object, the agent has to move continuously (since it is only by moving that the sensory activation differs for these points). And, since the sensory activation is changing smoothly for the panoramic sensors (because the object can be sensed all the time), the agent has to go far away from the object to unsaturate the sensor neurons. Once the neurons are unsaturated, the neural dynamics fall into the attractor of spinning, and the agent turns around until the equilibrium of the sensors is reached again. Therefore, the agent spends more time going away from the object (compare distances in figure 3.B and figure 7.B), making it harder for this type of controller to have a high fitness on average (see figure 9.A).

An important point to notice is the nature of the solutions found for the different kind of sensors and controllers. For example, on average, having panoramic sensors could help to find the object slightly more reliably than by having directional sensors (see figure 9.B). However, having directional sensors allowed the agents to be closer to the object in the most adapted cases (see figure 9.A).

This is because, by having panoramic sensors, in general it is easier to approach the object (since it can be always sensed), rather than requiring a search behaviour (spinning) for directional sensors. But as a trade off, by having panoramic sensors, it is necessary to go away from the object to unsaturate the sensor neurons (as explained above) and to stay further away most of the time due to the smooth continuous sensory activation. In contrast, the discrete activation of the sensory neurons in directional sensors resulted in more "reactive" phototactic behaviour allowing the agent to remain closer to the object.

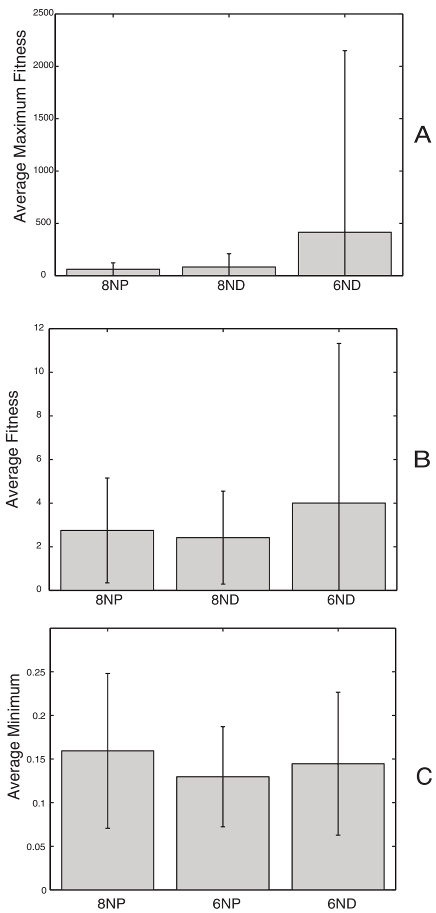

Figure 9: Results of 10 evolutionary runs for 1000 generations for different controllers and type of sensors. Controllers with 8 neurons and panoramic sensors (8NP), controllers with 8 neurons and directional sensors (8ND) and controllers with 6 neurons and directional sensors (6ND). Figure 9.A shows the average of the maximum fitnesses for the 1000 generations, figure 9.B the average of the average fitness for the 1000 generations and figure 9.C shows the average of the minimum fitness for the 1000 generations. In general, the controllers with 6 neurons and directional sensors had better fitness function than the controllers with 8 neurons.

It is also relevant to note that, in general, the solutions found by the evolutionary processes for the different scenarios were not very different in terms of maladaptation. The performance of the most poorly adapted agents was very similar regardless of the type of sensor employed (see figure 9.C).

5. CONCLUSION

In this work controllers were evolved to perform phototaxis. An analysis of the controllers through a dynamical systems approach showed that by restricting the sensory system, a simplification not only of the interaction between the environment and the agent, but also of the neural controller can be achieved. In particular, it was possible to evolve a controller with less neurons to perform the same task. By this reduction of the dimensionality of the problem, it was easier to analyse the neural dynamics of the controller and the behaviour of the evolved agent.

It is interesting that in this work it was found that such restriction of the sensors caused a better employment of the embodiment and situatedness by the evolved agents (see [18] for a good explanation of these concepts). Restrictions on the body and the visual sensors could actually improve the performance and simplify the design of neural controllers. These findings could be analogous to the evolutionary process of the visual systems in animals, which suggests that the appearance of directional sensors was one of the key issues in the exploitation of movement [14].

6. FUTURE WORK

Although the tasks employed for this work are very simple, they provide an example of the importance of the ecological balance principle.

The properties of the sensors employed showed an important correlation to the complexity of the controller required to solve a simple task. This correlation shows the crucial role of defining the interaction between environment and agent [5].

A point that would be interesting to study is the possibility of exploiting the low costs of simplified visual sensors. For example, these types of sensors could provide a cheap and fast alternative for the exploration of controllers performing more complex tasks (such as discrimination and recognition) using richer visual information.

Acknowledgments

Thanks to Eduardo Izquierdo for useful discussions about the dynamical systems approach. I want to thank Jenna Bailey for enhancing the readability of this work and thanks as well to Emmet Spier and Anil Seth for useful comments on draft versions.

7. REFERENCES

[1] R. Beer. On the dynamics of small continuous-time recurrent neural networks. *Adaptive Behavior*, 3(4):469–509, 1995.

[2] R. D. Beer. Dynamical approaches to cognitive science. *Trends in Cognitive Sciences*, 4(3):91–99, March 2000.

[3] R. D. Beer. The dynamics of active categorical perception in an evolved model agent. *Adaptive Behavior*, 11(4):209–243, 2003.

[4] D. Cliff and G. F. Miller. Co-evolution of pursuit and evasion II: Simulation methods and results. In P. Maes, M. J. Mataric, J.-A. Meyer, J. B. Pollack, and S. W. Wilson, editors, *From animals to animats 4*, pages 506–515, Cambridge, MA, 1996. MIT Press.

[5] K. Dautenhahn, T. Uthmann, and P. D. Special issue on 'evolution of sensors in nature, hardware and simulation. *Artificial Life*, 7(2), 2001.

[6] R. Fernald. Evolving eyes. *Int. J. Dev. Biol.*, 48:701–705, 2004.

[7] D. Floreano, T. Kato, D. Marocco, and E. Sauser. Coevolution of active vision and feature selection. *Biological Cybernetics*, 90:218–228, 2004.

[8] K.-I. Funahashi and Y. Nakamura. Approximation of dynamical systems by continuous time recurrent neural networks. *Neural Netw.*, 6(6):801–806, 1993.

[9] I. Harvey, P. Husbands, and D. Cliff. Seeing the light: artificial evolution, real vision. In D. Cliff, P. Husbands, J. Meyer, and S. Wilson, editors, *From animals to animats III*, pages 392–401, 1994.

[10] I. Harvey, P. Husbands, D. Cliff, A. Thompson, and N. Jakobi. Evolutionary robotics: the sussex approach. *In Press*, 1996.

[11] I. Harvey, E. A. D. Paolo, E. Tuci, R. Wood, and M. Quinn. Evolutionary robotics: A new scientific tool for studying cognition. *Artificial Life*, 11:79–98, 2005.

[12] R. Kortmann, E. Postma, and J. Herik. Evolution of visual resolution constrained by a trade-off. *Special Issue on 'Evolution of Sensors in Nature, Hardware and Simulation. Artificial Life*, 7(2):125–145, 2001.

[13] M. Land and R. Fernald. The evolution of eyes. *Annu. Rev. Neuroscience*, 15:1–29, 1992.

[14] M. F. Land and D. E. Nilsson. *Animal Eyes*. Oxford University Press, 2002.

[15] A. Liese, D. Polani, and T. Uthmann. A study of the simulated evolution of the spectral sensitivity of visual agent receptors. *Special Issue on 'Evolution of Sensors in Nature, Hardware and Simulation. Artificial Life*, 7(2):99–124, 2001.

[16] S. Nolfi and D. Floreano. Synthesis of autonomous robots through evolution, 2002.

[17] E. A. D. Paolo and I. Harvey. Decisions and noise: the scope of evolutionary synthesis and dynamical analysis. *Adaptive Behavior*, 11(4):284–288, 2004.

[18] R. Pfeifer and C. Scheier. *Understanding Intelligence*. MIT Press, 1999.

[19] S. H. Strogatz. *Nonlinear Dynamics and Chaos: With Applications to Physics, Biology, Chemistry and Engineering*. Perseus Books Group, 2001.