Experiments in minimal cognition: An investigation into visual orientation and size discrimination.

Thomas A H Baker

Third Year Project
Artificial Intelligence
Department of Informatics
University of Sussex, UK
2006

Candidate Number: 83385

Under Supervision of Ezequiel Di Paolo

Statement of originality

This report is submitted as part requirement for the degree of Artificial Intelligence at the University of Sussex. It is the product of my own labour except where indicated in the text. The report may be freely copied and distributed provided the source is acknowledged.

Tom Baker

Abstract

The project attempts to explore the nature of size discrimination using techniques taken from the field of evolutionary robotics. Previous investigations in physiology and minimal cognition have shown that animals and artificial agents tend to use relational size discrimination (bigger than, smaller than, in-between) much more easily than absolute size discrimination (choosing objects of a specific size). To explore the validity of these results, simple idealised agents and environments are constructed with the aim being to follow the minimally cognitive framework. Artificial evolution is used to evolve biologically plausible neural networks, allowing the agents to form 'naturally' thus limiting the number of assumptions made about the required solutions to the tasks. Initially, a simple experiment in visual orientation is used as a trial (a necessary component for size discrimination behaviour in this case); the developed techniques and programming structure are then applied to evolving agents capable of size discrimination. Two agents are successfully evolved to perform relational tasks, one to select the larger of two circles and the other evolved to select the smaller of two circles. Testing confirmed that the strategies used were indeed based on relational discrimination thus supporting previous investigations. Mixed results were achieved when evolving an agent to perform absolute discrimination, evolution being significantly harder. High scoring agents were however evolved, although they surprisingly employed a relational approach to selecting the correct circle. Unfortunately an attempt to evolve an agent capable of selecting a circle not of absolute size failed, only producing an agent that relied heavily on relational discrimination. A further experiment was also conducted in which an attempt was made to evolve an agent capable of selecting the middle of three stimuli. Yet again mixed results were achieved, although the agent did evolve to use a relational technique closely related to the agent capable of absolute discrimination. The results strongly agree with the previous investigations and serve to emphasise many of evolutionary robotic concepts. In each case artificial selection was presented with the chance to use relational or absolute methods of discrimination, however in each case agents evolved to use relational discrimination. The poor absolute discrimination results therefore suggest that relational discrimination is indeed the easier discrimination to evolve and thus use. The evidence is therefore biased towards relational discrimination being the 'fundamental' cognitive primitive for size discrimination and even suggests that absolute discrimination might well be based on this. The project finishes with a discussion; this explores the findings and warns against drawing premature conclusions. Consequently an effort is made to view and understand the results from an evolutionary prospective. By doing so the discussion highlights many of the important issues needing to be addressed when performing similar investigations and provides a possible method for interpreting the data and applying it to real biological systems. Finally in light of the findings several extensions are proposed.

Acknowledgements

Many thanks to Ezequiel Di Paolo for consistent advice throughout the year and also to Eduardo Izquierdo Torres for clarifying certain concepts and techniques used within this investigation.

Contents

1. Introduction

1.1 Overview of Project (See appendix B for professional considerations made before starting the project).

The project uses the principles and techniques of evolutionary robotics to evolve and study a simulated agent[1] that exhibits minimally cognitive behaviour. The agent proposed will use a decentralised control system, in this case a continuous time recurrent neural network. A microbial genetic algorithm [11] will then be implemented to evolve this control system.

Although many experiments exist where the aim is to model simple dynamic control systems, there are only a few which are directed at exploring the concept of minimal cognition and that attempt to fully explain how such behaviour is achieved out of apparently simple agents (e.g. [3,18,31]).

The purpose of this project is to evolve an agent who's behaviour is complex enough to trigger interesting questions about its cognitive ability and cognition in general but simple enough to be understandable and thus can be explained through thorough analysis. In this case minimal cognition will be defined as in [3(page 2)] "the simplest behaviour that raises cognitively interesting issues".

The first goal of the project is to evolve an agent capable of visually orientating it self to 'falling circles/balls', a short analysis will follow however the main purpose is to make sure the experimental techniques and methodologies used are correct. The methodologies learnt are then implemented to achieve the main goal in which artificial evolution is used to evolve an agent capable of size discrimination. In this case the agent will be investigated through experimental procedure to explore its ability to discriminate between two circles both in relational terms (bigger than, small than) and absolute size (always choose a circle of a specific size). The investigation is concerned primarily with the resulting 'agent – environment' system, and examines how the internal dynamics of the control system and the dynamics of the agents' body when coupled with a specific environment result in the adaptive behaviour as selected for through artificial evolution. The 'evolveability' of each agent will also be noted as will the generalisation ability of the resulting agent (its ability to adjust to situations which it has not been evolved specifically for). Previous success in evolving artificial agents that display minimally cognitive behaviour such as in [3,18,22,31] support the likelihood that this investigation will succeed.

The first sections of the report attempt to motivate the reader to value the importance of such an investigation, to provide a background to the procedures mentioned above and to briefly introduce related experiments with simulated artificial agents; in addition an overview of relevant psychological and biological data is given.

The main section of the report deals with the methods used for the simulated experiments such as the evolutionary algorithm employed. The experiments are described in detail and the results presented and analysed. The final section of the report focuses on a discussion of the achieved results.

1.2 Motivation

For many the study of minimal cognition, specifically size discrimination, will seem to be a dull subject and some would even doubt the value of such experiments. These viewpoints are understandable, simple basic experiments implemented in simulation do at first glance appear dull and even unhelpful. They are a far cry from attempts such as [17] to model regions of the brain in detail or experiments in which robots are physically constructed e.g. [42]; all of which grab the imagination and seem to offer much hope for understanding intelligence. What then can simple experiments involving idealised agents tell us about anything? Surely the modelling is too abstract and the results impossible to apply to actual intelligence?

The above views are short sighted, if we hope to create agents that are capable of performing more advanced cognitive tasks then surely it makes sense to build up from simple systems? Even simple agents are notoriously hard to analyse, building highly complex models only involves huge amounts of time and investment and makes the situation more complex.

Instead of attempting to produce highly complex accurate simulations of agents with huge control systems, experiments in minimal cognition attempt to generate interesting behaviour out of the simplest of agents. Far from being simplistic/limited such idealised models often produce remarkable behaviour, and not only call into question previous approaches to artificial intelligence but also cast doubt over conventional reasoning about cognition. Of course such models do abstract away from the biological details, yet they also retain a degree of biological plausibility. Thus the abstractions do not hinder the usefulness of such experiments; indeed the actual functioning of biological mechanisms is still unclear and often debated therefore idealised models make a lot more sense, - hopefully capturing the aspects that are important (e.g. decentralisation of control system, internal dynamics,

[1]In this case 'an agent' can be understood as simple situated and embodied simulated robot, section 2 explains these terms in more detail.

naturally arising control systems, embodiment, situated-ness) while ignoring the fine details who's actual importance and role is still unclear.

It is in this way then that such experiments are useful, it must be noted though that the resulting agent and its internal dynamics can not be applied directly to living systems; It is not possible to take the firing rates of the artificial neurons and look for similar patterns in an actual living brain and thus say this is how animals solve this problem (a result of huge differences in physiology of the agent and organism as well as the differences in the complexity of the environment). However investigations into minimal cognition can inform biological studies and correct previous theories about living systems and their operation within the world (see Webb [38,39]). The evolution and analysis of minimally cognitive agents is thus a useful, and valid technique for contributing to the understanding of cognition; why then is size discrimination important?

The ability to discriminate between the size (or perhaps even another feature e.g. brightness) of two similar objects is an aspect of cognition that we tend to use unwittingly but it is non-the less an important ability present in many organisms. Unfortunately the simulation in this project is simple and abstract to aid analysis, it is therefore hard to see its application in the real world. However there are many potential applications for size discrimination in the natural world. For example on a low 'neural level', one could easily imagine individual neurons performing size discrimination between two inputs, i.e. firing according to the size of different inputs that are perhaps larger or smaller than some threshold (indeed it is even possible that 'size discrimination' of one sort or another is a fundamental operation in neurons, and can be used as a bases for other tasks, or alternatively perhaps it is itself based on more basic operations such as multiplication). Size discrimination may also be used for visual processing, many animals rely on optic flow to obtain distance information (e.g. locusts [5]), this is particularly evident in invertebrates, where if the image size of a object on their retina moves above a particular threshold (becomes larger) the animal will perform a particular behaviour (e.g. avoidance). Another application might be landmark recognition, for example evidence suggests that wood ants [10] attempt to position objects on their retina to fit the stored 'size' (memory) of the objects from a particular vantage point.

At a higher-level, size discrimination may be used in part for mate selection. There is research to suggest that in some animals a females or males choice of mate might be based on the size of an individual's physical feature. Evidence for this is seen in reindeer [1], wolf spiders [30] and crickets [2] where it is suggested that the size of a feature of an organism or of the organism itself may provide an indication as to the fitness of an animal (the quality of its genes). Other forms of assessment may well also rely on this, i.e. two males may analyse one another's features before confrontation, if one male's horn for example is much smaller than another then it is unlikely to win the fight to secure a female and so it might back down and avoid confrontation.

These observations are not completely dependant on size discrimination (for example mate selection often involves discrimination of different dimension e.g. length of a 'song' or 'dance'). However the above examples do give a good indication as to the variety of applications that size discrimination might well be used for or involved in, and shows that discrimination between two or more stimuli of one type or another seems to be fundamental to and tied up in many behaviours. It is thus an interesting and worthwhile subject to study; an idealised model of which may well help to inform and support biological theories.

2. Background Information

2.1 Evolutionary Robotics: Genetic Algorithms, Neural Networks and Simulated Environments

Evolutionary robotics (see [14,28] for an in depth review) consists of evolving robotic/agent control systems and/or morphology to perform specific behaviours through the use of genetic algorithms. The fitness (score) of these control systems increases over time (successive generations) and eventually a control system will emerge that produces the correct or 'wanted' behaviour.

Genetic algorithms represent a form of search. The idea is inspired from the process of natural evolution of a species. Each member of a species has a phenotype (its physical characteristics) that is encoded for by its genotype (its genetic information). When organisms reproduce this information is passed from both parents into their offspring; this is the process of recombination where the offspring inherits fifty percent of the fathers and fifty percent of their mothers genetic information. Consequently the resulting offspring has different characteristics from each of their parents, and it is the characteristics that determine the fitness of the offspring.

The environment is constantly selecting the 'fittest' members of a population, i.e. the best adapted members are more likely to survive and reproduce as a result of this the less fit will begin to die off being replaced by fitter members (for example which better at competing or avoiding predators). However if the population continues in this way, without 'fresh' genetic information entering into the cycle, it is likely to reach a peak fitness and not progress from there (after many generations there are still only certain genes to choose from). This would cause problems if the environment were to suddenly significantly change as there are likely to be very few, if any members of the population able to adapt (not enough variation). Consequently mutation plays an important role in the evolutionary process, acting against the selective process by introducing randomness into the population. Mutation may or may not end up with a positive result, in most cases this randomness will not improve the fitness of an individual, however every now and then it might 'get lucky' and give a particular individual a phenotype which is fitter then the rest of the population i.e. better adapted. Consequently mutation allows for the evolutionary process to continue finding the fittest members as selection pressure changes according to the environment. A modified form of the evolutionary process described above is used in this project (as described in [11]), although the underlying principles are the same; for a detailed explanation of the process see section three.

The experiments in this project involve an agent that is both situated and embodied as described by Brooks in [6] (in a simulation sense, the agent is not actually located in the real world). The control system is not seen as an isolated calculation device but as a system that interacts on sensible time scales dynamically with the environment around it; evolution can therefore use this, and the dynamics that arise due to embodiment to find novel solutions to problems. The environment and the fact that the agent is embodied thus play a large role in determining the scope of the agents' behaviour; the importance of both should be reiterated in the results.

The control systems employed in evolutionary robotics are usually of the form of neural networks, thus having a general level of biological plausibility. A continuous time recurrent neural network is implemented in this case, introducing temporal dynamics to the control system (neurons can fire out of phase/different times) where its behaviour is affected by recent previous experience. This adds further complexity to the dynamics of the control system thus allowing for the generation of additional interesting interactions between the agent and its environment (the agents behaviour can depend on recent previous interactions with its environment (a form of short-term memory). In addition because the evolutionary/genetic algorithm selects the network architecture, far fewer prior assumptions are made about the control system than with traditional robotic techniques. A highly important consequence of this and one that is vital to this investigation is that the correct behaviour has more chance of arising 'naturally' from the interaction between the agent and environment, perhaps exploiting unexpected invariants in the environment.

Further, the neural network is bilaterally symmetric, thus mimicking the symmetry found in nature (e.g. the left and right hemispheres of most mammalian brains), this reduces the number of parameters needing to be evolved (effectively evolving the connections for one side of the network only) and suits the symmetry of the tasks.

The experiments involved in this investigation are in themselves quite simple, however there are a variety of interesting and important questions that can be asked and investigated. For example how important is it to model any physics in the simulation and is this related to the agents' problem solving ability? An apparently complex problem that appears to be solved through 'calculation' might well be explained in terms of the simple dynamics between the embodied agent and its environment; simple environmental features might be exploited to solve the problem. For example, a robot controlled by a simple feed forward (reactive) neural network might be tasked to follow a ball. If the ball were to move behind another object the robot would immediately stop (having no inputs as the ball is occluded by another object). However if inertia was added to the simulation, then evolution might take advantage of this and use it to allow the robot to continue to move (due to its reluctance to slow down) thus when the ball became occluded the robot would continue to move past the object until the ball was visible again, consequently being seen to solve the problem. This poses interesting questions in terms of how much of our every day interaction exploits simple aspects of the environment in order to solve quite complex problems and warns us to be careful when assigning behavioural descriptions to agents (i.e. the behaviour might be much simpler than it appears).

The above is a prime example of the decisions that have to be made when experimenting via simulation and the interesting questions that can arise even from the simplest of simulations. Indeed our ability to solve problems might well depend on something that we forget to simulate, consequently causing difficulties to evolutionary process. If however the environment is too simple an agent may evolve to exploit unobvious simple patterns, appearing to solve the problem but actually just taking advantage of an unwanted or unrealistic regularity in the environment. For example, if the task was to evolve an agent capable of catching objects, it would be a mistake to start the agent off in the same position for each evaluation trial and have the object always fall from the left at a set distance. In such a situation, one way to solve the problem would be to simply go left for a certain amount of time and this is almost certainly what artificial evolution would exploit.

To combat the above, a degree of randomness has to be introduced to the simulation. Noise can be added to the environment (i.e. random starting positions, background noise – see Jakobi's minimal simulation approach in [19]) and to the agent itself (i.e. to the sensor inputs values, motor outputs and also to the internal neurons). However a balance has to be achieved, too much noise and it will make the analyses of the evolved agent-environment system difficult. As well as including noise, it is also important to include a 'representative' number of trials to judge the fitness of an agent, where the positions of objects or the sizes of objects are changed and an average score taken to remove 'fluke' results and to stimulate the evolutionary process. If applied, the resulting agent should be able to handle variations in its environment and thus generalise well to new situations.

In this case we want to keep the environment simple, and therefore only the essential aspects of the simulation are modelled to keep the analysis tractable. Internal neuron noise is also removed however 'environmental noise' in terms of the random placement of objects is kept. The specifics of the 'agent-environment' setup are described in section 3.

2.2 Previous Experiments in Minimal Cognition

Randall D. Beer has produced some of the most recent and relevant work in minimal cognition [3][31]. His research follows the evolutionary robotics methodology and concentrates on analysing the dynamics of idealised agent and environment systems; concentrating on agents capable of visual tasks relying on primitive vision systems. His initial experiments described in [3] evolve agents capable of visual orientation (agents had to adjust their horizontal position to catch falling objects with differing velocities, sizes and offsets), achieving fitness' of 99% over 100 random trials. Later agents capable of discriminating between different shapes are evolved. Populations of 300-400 agents are evolved for 100 –200 generations to catch circles but avoid diamonds and lines, the fittest achieving scores of 98.68% (for circles and diamonds) and 97.85% for (circles and lines) over one hundred random trials (Similar experiments are also performed by Eduardo Izquierdo-Torres achieving comparable results described in paper [18]).

Slocum and Beer's second paper [31] leads on from the above success, evolving agents capable of a number of cognitively hard tasks for example agents capable of judging their own body size relative to openings or switching their attention between multiple distal objects.

Two experiments stand out in this investigation as being relevant to this project. The first being the investigation in to 'perceived affordances', here agents are evolved that achieve a 96% over 1000 random trials, capable of judging the size of an aperture in a falling wall in relation to its own body size in enough time to prevent itself from becoming squashed.

Another experiment of relevance to this investigation is the experiments involving selective attention. In this case an agent has to focus on an object while ignoring another object in its environment and catch both of them (the objects moving at different horizontal and vertical velocities). Consequently the agent was confronted with many problems (this would be extremely hard even for a human to achieve!) such as to which object should it attend to first. The resulting agent, was evolved from a population of 100 individuals for over 9000 generations achieving 90 % fitness score over 1000 trials and is capable of "decoupling its behaviour from its immediate circumstances while remaining sensitive to them, make predictions about the future configurations of objects based on observations of their past motion" [31 (page 1)].

In all cases the evolved agents use dynamic strategies to solve the particular tasks, in the above example the strategy employed by the agent involved keeping both objects in its field of view, the dynamic behaviour involving the performance of large sweeps back and forth, with the agent gradually tightening its scan on fast moving objects.

Directly relevant to this investigation is the work achieved by Peter Law [22] whose work concentrated on the analysis on agents capably of size discrimination and to my knowledge this the only other investigation of a similar type. His work supports the idea of relational discriminations being the cognitive primitive, finding it much harder to evolve agents capable of absolute size discrimination.

The above experiments will be used to guide this project and it is expected that similar results will be obtained (size discrimination involving a degree visual orientation and selective attention). In addition it will be interesting to see if the results of this project support previous studies.

2.3 Previous Experiments in Psychology and Biology

There are numerous examples where animals might use size discrimination and as a consequence a large body of experiments concerned with its nature. Most research is concerned with identifying whether animals judge size through relationships between objects (relational discrimination) or whether animals use the absolute size of objects to make distinctions, perhaps one is even derived from the other. Therefore the psychology approach tends to concentrate on designing and conducting experiments to support either the relational discrimination or the absolute discrimination as the 'basic' cognitive primitive for size discrimination (i.e. which of the two do animals do more naturally?).

Although the above is not the fundamental aim of this investigation, (the goal purely being to evolve an agent capable of such behaviour and to gain some understanding of the dynamics of its control system when coupled with its body and environment) it is none the less important subject to explore whilst experimenting. Therefore it is useful to briefly analyse several of the main psychological experiments; through which several benefits can be had:

1) Identify the most probable cognitive primitive based on the psychological experiments.
2) Apply previous psychological experiments to simulated agent.
3) Compare results.

As Peter Law [22] describes in detail, Kohler's transposition test [21] is perhaps the most famous experiment in which chickens, chimpanzees and humans were tested to observe whether they make simple relationships between pairs of objects to make discriminations (although initially it was not directly concerned with size discrimination). In brief, the experiment consists of a training phase and a test phase. The training phase involves training a test subject to select one of a pair of stimuli (differing in one dimension e.g. in this investigation it would be size) e.g. stimuli A and stimuli B, where stimuli B is bigger than A. The subject was then tested (test phase), given a choice between the entrained stimuli, B and a new stimuli C, where C is bigger than B. If the subject chose stimuli B in the test phase then the subject could be seen as using absolute size to discriminate where as if the subject chose C then one could see this as evidence for the subject using relational discrimination. The results of his experiments identified that the majority of his subjects of all species tested tended to choose stimuli C, thus supporting the idea of relational discrimination as a cognitive primitive.

A selection of other relevant investigations is provided in the bibliography (section 8) [9,20,23,24,27,40,41], these experiments use similar techniques to test various animals from pigeons to chimpanzees. The experiments share many similarities and the overriding conclusion gained from each is that relational discrimination is much easier to perform than absolute discrimination. Although another interesting observation suggests that after many presentations of the same stimuli certain animals will start to use absolute discrimination, relative discrimination 'subsuming control' when again presented with novel stimuli.

Two representative examples of the experiments are firstly the investigation by Lawrenson and Bryant [23] which showed that performance of young children when carrying out tasks that are best solved by both techniques (absolute and relational) was much higher in relational discrimination tasks. Secondly, the investigation carried out by McGonigle and Jones [27], here squirrel monkeys were tested at size and brightness discrimination tasks, the monkeys' performance being much stronger on relational tasks, requiring fewer trials to complete/learn and producing fewer errors. In addition the relational discriminations were much more resilient to variations in the tests such as the addition of an extra object. Further evidence suggests that larger stimulus is learned with greater ease followed by smaller then middle. In terms of biological experimental evidence, a few examples are given in section 1.2, however a more specific example accompanied by strong evidence is the investigation into the detection of parasitic cuckoo eggs by the yellow-browed leaf warbler (host) in [26]; where the detection seems to based on the relative size of eggs in the clutch, rather than absolute sizes of individual eggs.

These results make sense, from ones own experience it tends to be much more difficult (considerable conscious effort often involved) to select an object of absolute size (when presented with two similar stimuli) with out a large degree of familiarity with the particular object (perhaps even requiring previous experience in measurement or physical involvement with a certain kind object). Relational discrimination however can be made instantly; one can immediately pick the larger or smaller or two similar objects.

In conclusion, there is a large body of evidence supporting relational discrimination as the primitive for size discrimination and it will be interesting to see if this is the case for the evolved agent.

3.Methods

The experiments are carried out in simulation only. No attempt is made to accurately model the physics of the robot, its sensors, or to make the environment completely realistic. The experiment is interested in the dynamics of the evolved control system within its simple environment and not in creating a detailed model of the real world. Keeping the simulation simple allows more attention to be given to the development of a working control system; secondly because the environment is simple the behaviour of the agent and the dynamics of the control system should be easier to explain.

The agent and environment is similar to the experiments described in papers [18,22]: although there are several modifications as described below. The parameters and values remain the same for each experiment and noise is not added to the simulation unless otherwise stated.

3.1 Environment Description

The environment consists of a 500x500 pixel ‘arena’, as described above the simulation does not attempt to model any real world physics. In both experiments circles fall from the top of the arena and the agent is positioned at the bottom of the arena. The vertical velocity of the circle dictates the amount of time the agent has to catch a circle. In the size discrimination experiments each circle is positioned just beyond the visual range of the agent with the base of each circle in the size discrimination experiments having the same vertical (Y) position (diagram 1 A). If this were not the case evolution would most likely take advantage of the first stimuli entering the environment being largest circle as reported in [22] (see Diagram 1 B).

A Visual representation of agent (red circle) showing its evenly spaced sensors rays (grey lines) and the environment displaying the starting positions of each circle (blue and green circles) is shown in diagram one:

Diagram 1 A:

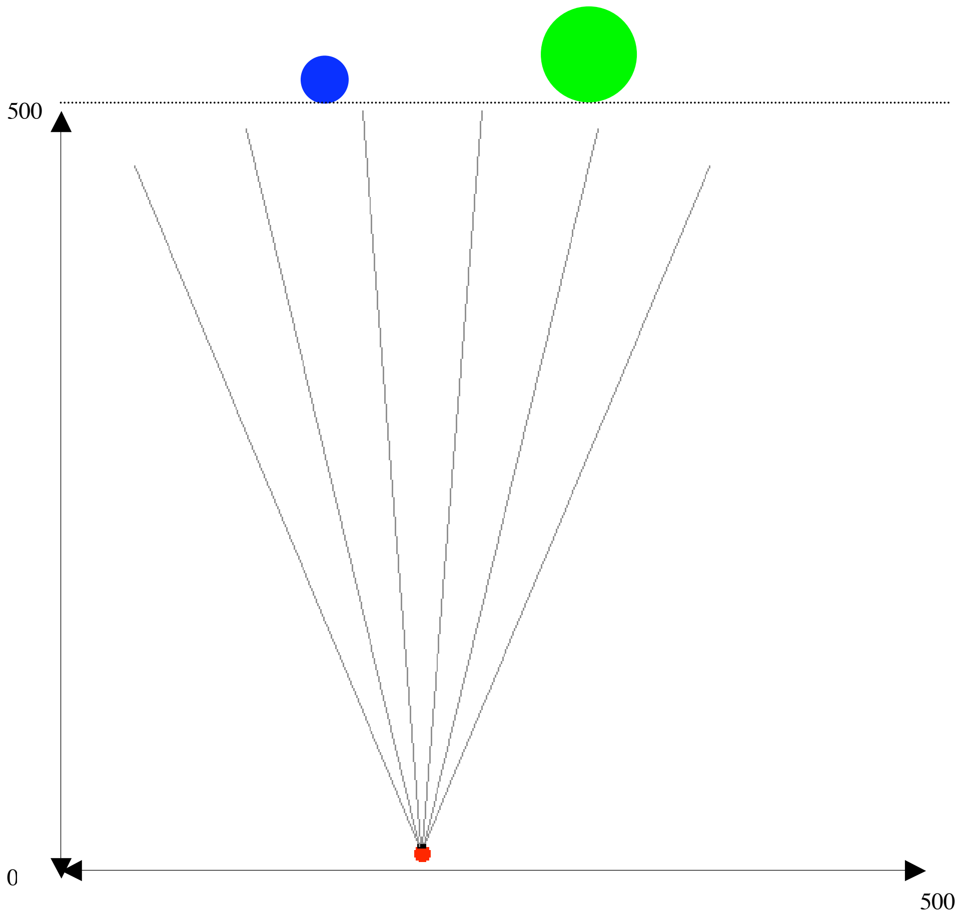

Diagram 1 B: (larger circle base ‘hanging’ below smaller circle base).

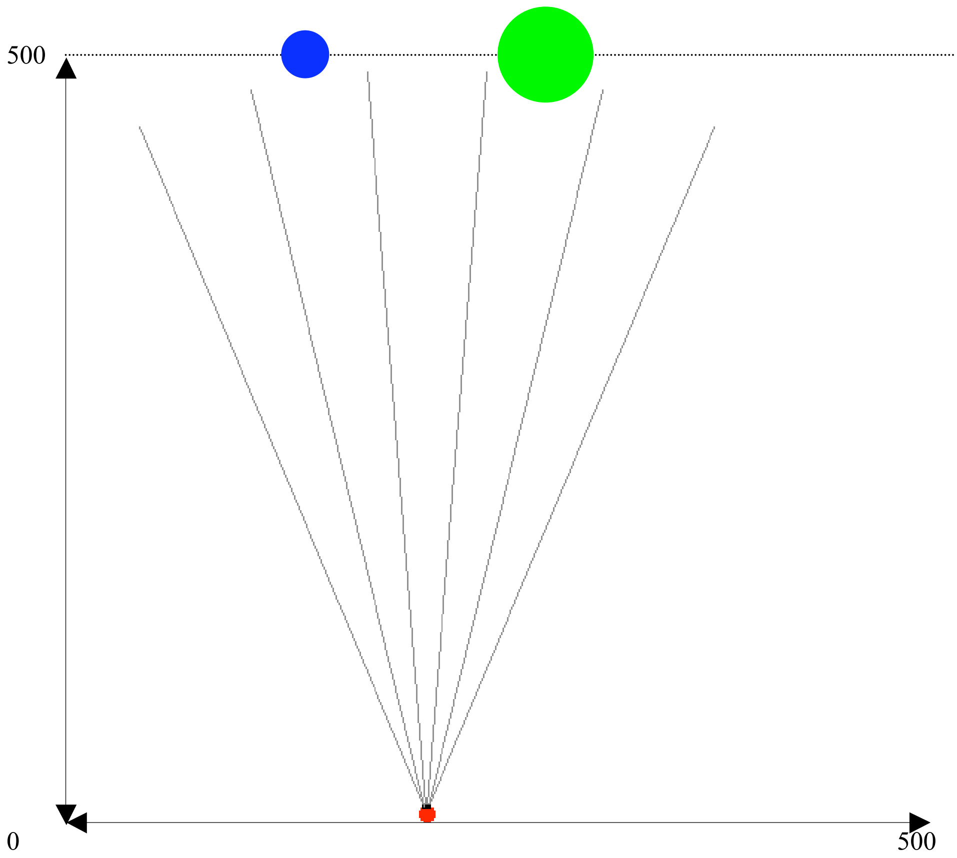

3.2 Agent Description

The agent consists of a circular shape (diameter 10 pixels) confined to a 1 dimensional plane. It is therefore restricted to moving horizontally, left or right only; the direction of movement dictated by the sum of two opposing motor outputs. In this experiment the maximum velocity is left uncapped.
The agent has six 'ray sensors' that allow it to perceive its environment evenly spaced over 45° facing 'upward' (diagram 2).

Diagram 2: 'Zoom in' on the agent showing evenly spaced sensors:

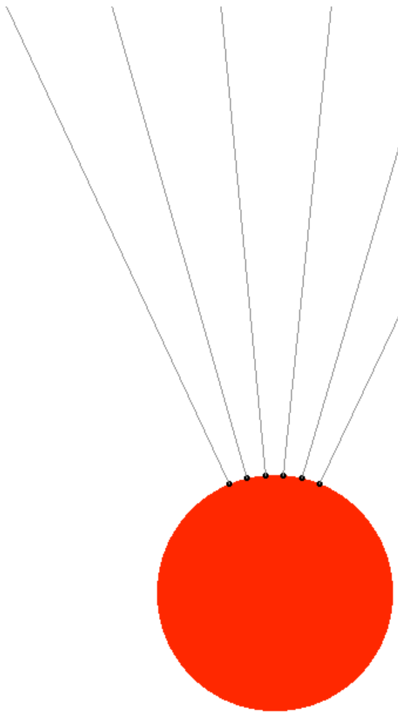

If a circle intersects (diagram 3) one or more of these rays an input value is calculated for those sensory nodes rays that are intersected. The value is inversely proportional to the distance of the intersecting object along a particular ray. Each ray has a length of 500 pixels, giving the agent a good visual range of the environment (diagram 1) (-needed for the particular task it will be evolved for); if the rays are at there maximum length no input is injected with a maximum input being injected when they are of zero length.
Therefore the resolution of the agent is dependent on the number of rays, the visual angle over which they are spread, how far away an object is (a more distant object will intersect fewer rays) and also the size of an object (a larger object intersecting more rays). In addition a sensor gain is evolved (shared by all sensors), large gains increasing the effect each sensory input has on the input neurons.

Diagram 3:

Example of the circle intersecting a ray:

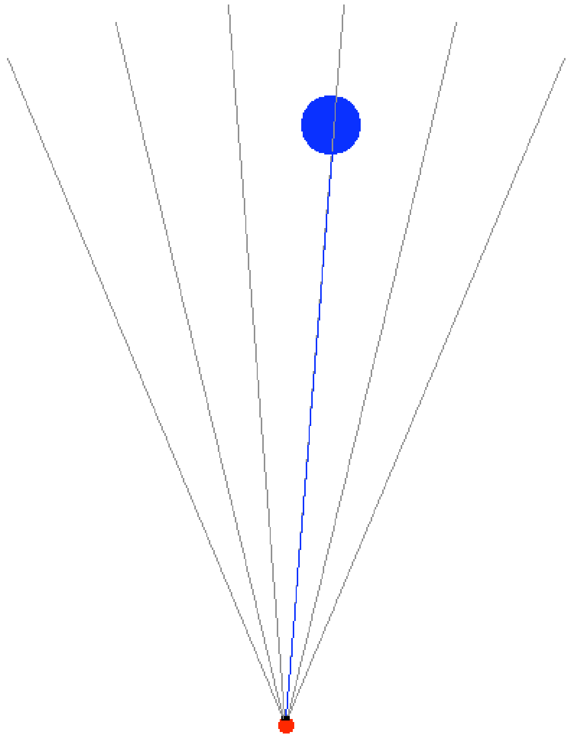

3.3 Details of Control System

A bilaterally symmetrical continuous time recurrent neural network (CTRNN) was evolved to control the agent. The network takes the following form:

$$\tau_i \, y_i = - y_i + \sum_{j=1}^{N} \omega_{ji} \, \sigma \, (y_j + \theta_j) + I_i$$

Where y is the activation function of each node, τ is its time constant, ω is the strength of the connection from the j^{th} to the i^{th} node, θ is a bias term, $\sigma(x) = 1/(1+ e^{-x})$ is the sigmoid non-linear activation function, I represents a constant external input and N is the number of nodes in the network.

As in [3,22,18] the node activations will be calculated forward through time by straightforward time slicing using Euler integration, in this case with a time step of 0.1. (Therefore a circle falling at a speed of five pixels per time step would now fall at (5 x 0.1) 0.5 pixels per time step. A 500 pixel ‘drop’/movement takes 1000 time steps (500/0.5)).

The neural network consists of six input nodes/neurons that receive inputs from each of the six sensors correspondingly. These project to six interconnected hidden neurons, each having a self-connection. The hidden neurons then project to two output neurons used to control the motors (see diagram 4).

Diagram 4:

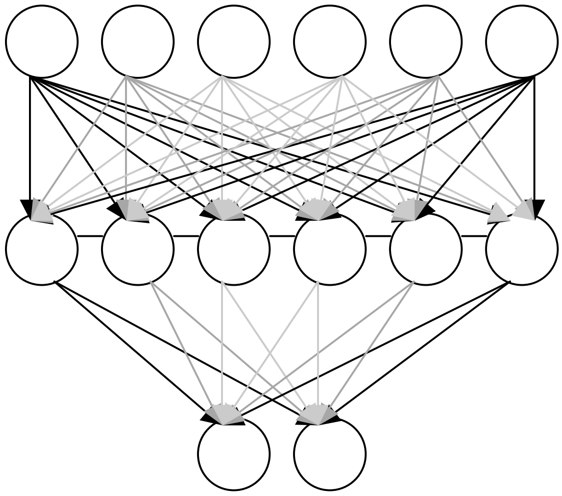

The network is bilaterally symmetrical in weights values, neuron time values, and biases however the input neurons share the same time constant and bias values (this is also true for the output layer neurons as there are own two neurons). The genotype was encoded as below: I =Input, H = hidden, O = output, wts = weights, T = time constant, B = Bias.

[I to H wts] + [H to H wts] + [H to O wts] + [T I]+ [T H] + [T O] + [B I] + [B H] + [B O] + [sensor gain])
[3*6] + [3*6] + [3*2] + [1] + [3] + [1] + [1] + [3] + [1] + [1] = 53 parameters.

3.4 Microbial Algorithm Details

A microbial genetic algorithm [11] is used to evolve the control systems. This type of reproduction involves a similar process of evolution as described in section one; however firstly, the reproduction is different (in terms of parental selection, recombination and mutation), secondly, the population size is also fixed and thirdly, the selection process remains constant.

The algorithm uses tournament selection, where two members of the population are randomly selected and evaluated. The genes of the winner of the tournament (the fitter of the two) are then copied to the losers' genotype with a particular probability (i.e. unconventional reproduction, the loser is almost 'infected' with the winner's genes); the losers' genotype is then mutated. This is therefore a steady state genetic algorithm meaning that instead of the whole population being placed each generation, only one member is; the fittest of each tournament remains the same (is kept in the population) therefore a form of elitism is used.

Initially in both experiments the population size was fixed to thirty individuals (unless otherwise stated). The mutation rate was set at 100% and the recombination rate was set at 90%; consequently at the end of a microbial tournament each gene on the losers/offspring's genotype had a 90% chance of being replaced and a 100% chance of being mutated.

The parameters for the neural network in phase two are initially encoded randomly as real numbers between 0 and 1 on the genotype. When an agent is evaluated its genes are linearly mapped to the neural network parameters using the following scales:

Time Constants = [1,10]
Biases = [-10,10]
Weight Values = [-10,10]
Sensor Gain = [1,20]

The genetic algorithm is as follows:

1. Construct a population of genotypes with random values. In this case the genes are encoded as real number values between 0 and 1 and the population size consists of thirty individuals.

2. Randomly select two members of the population (the beginning of a tournament).

3. Calculate the fitness of each of these individuals (see the evaluation function described below).

4. Compare both of the fitness scores.

5. The individual with the highest fitness score takes the place of the 'winner' and the one with a lower score takes the place of the 'loser'.

6. Perform recombination. With a 90% probability copy the winners genotype to the losers (each gene on the losers has a 90% probability of being changed).

7. Perform mutation. With a 100% probability mutate the losers genes. This is achieved by adding a small random number to each of the loser's genes. The value of which is drawn from a non-uniform Gaussian distribution with a mean of 0 and a standard deviation of 0.05. After each mutation, check to make sure the new value remains between 0 and 1 if not adjust (see appendix 1 source code).

8. Put both individuals back into the population into the same positions (This is the end of one tournament).

9. Repeat the procedure for a specified number of tournaments.

10. After thirty tournaments it is considered that roughly one generation has passed (depending on population size, if 100 individuals then after 100 generations = passing of one generation); at this point then the average fitness score and the highest fitness score are recorded. (The fittest fitness and the average fitness score are also sampled at the beginning and end of the genetic algorithm).

4. Visual Orientation Experiment

4.1 Evaluation Function

The evaluation function is used to assess the fitness of a particular member of the population based on its behaviour. The higher the returned fitness score (the closer it is to one) the better adapted the individual is; the algorithm is as follows:

- Transfer the agents' genotype to network parameters.
- Place the agent in the environment at a random horizontal position usually between the range of: [150,350].
- Reset its activation states (see Appendix C for details).
- Place a circle in the environment at a random horizontal usually between the ranges of +- [10,50] of the agent's horizontal position.
- Set the radius of the circle to a random size between 1 and 40 and set its vertical velocity to 5.
- Store the initial horizontal positions. Calculate the difference and store for use later.
- Run the simulation; allow the circle to fall 'downwards' and the agent to respond.
- Stop the simulation when the circles 'y' axis position is equal to the agents 'y' axis position (after 1000 time steps).

- Calculate the fitness of the agent represented by the equation below:

$$\text{Fitness} = 1 - \sum_{i=0}^{N} d_i / N$$

- Where 'd' is the horizontal distance between the centres of both objects when their vertical separation is zero.
- If 'd' is less then at the start of the evaluation trial then: divide the initial 'd' by the final 'd' and subtracted the result from one to give a fitness score for the particular trial.
- Add this to the total fitness score.
- Repeat this process twenty times, after which the total fitness score for the particular agent, is divided by the number of trials. This average fitness score is finally returned for the agents fitness.

4.2 Evolutionary Results

Figure 1: Graph of fitness v tournament number
(Solid blue line = fittest fitness score at each sample, solid red line = average fitness at each sample)

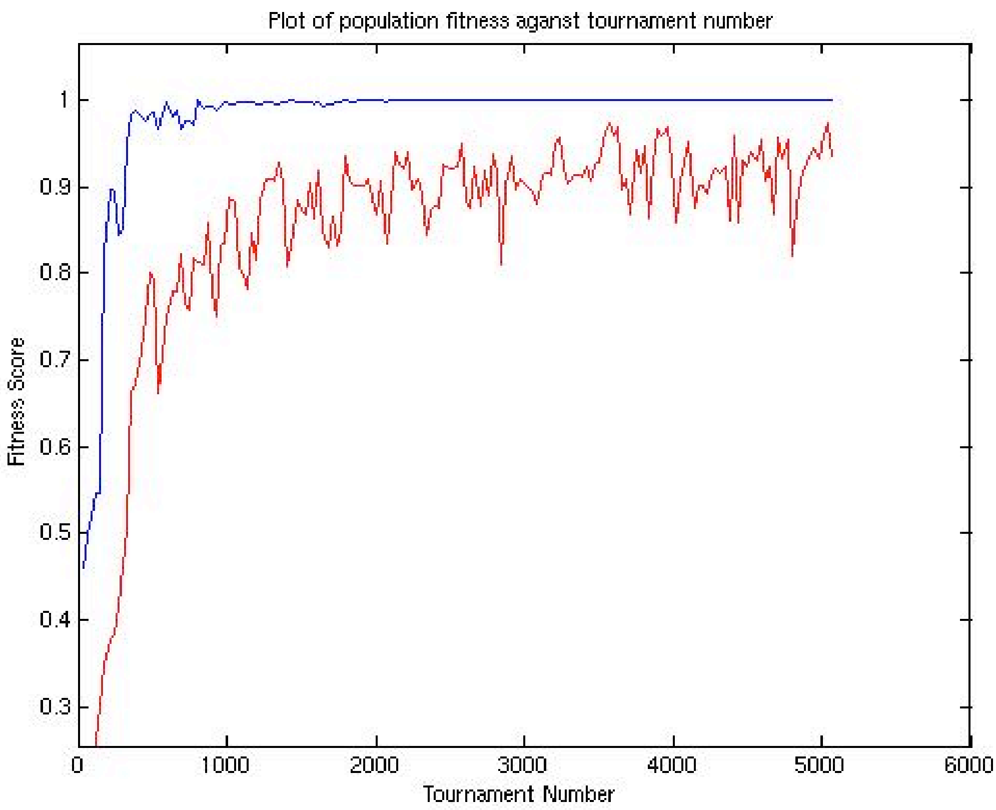

Agents with a high fitness were easy to evolve the fittest scored 99% over 20 evaluation trials (in evolution, figure 1) and 99.7% over 1000 random trials. The evolutionary algorithm worked well, producing adequate results within 1000 tournaments (approximately 33 generations).

4.3 Experiment Description

Two Experiments were carried out:

Experiment 1) An experiment to test the dependence of the agent on the offset value of the circle (how far away it is).
Experiment 2) An experiment to test the dependence of the agent on the size of the circle.

The following algorithm represents the main experimental procedure followed for each experiment (variations of this are used later in the size discrimination experiments). The details change according to the experiment, the changes being highlighted when appropriate.

- Transfer the agents' genotype to network parameters.
- Place the agent in the environment at a random horizontal position between [150,350].
- Reset its activation states (see source code for details).
- Place a circle in the environment (at the highest vertical position y = 0).
- Set the radius of the circle to equal a specific size and set its vertical velocity to 5.
- Store the initial horizontal positions. Calculate the difference and store for use later.
- Run the simulation; allow the circle to fall 'downwards' and the agent to respond.
- Stop the simulation when the circles 'y' axis position is equal to the agents 'y' axis position (after 1000 time steps).

- Calculate the fitness of the agent. This is given by the equation below:

$$\text{Fitness} = 1 - \sum_{i=0}^{N} d_i / N$$

- Where 'd' is the horizontal distance between the centres of both objects when their vertical separation is zero.
- If 'd' is less then at the start of the evaluation trial then: divide the initial 'd' by the final 'd' and subtracted the result from one to give a fitness score for the particular trial.
- Add this to the total fitness score.
- For the first fifty trials, the circle is positioned at a random position on the left hand side of the agent. For the final fifty trials, the circle is positioned on the right hand side of the agent.
- Repeat this procedure for specific number trials (varying form 100 to 1000 depending on specific experiment).

4.4 Experiment Results and Analysis

Continuous time recurrent neural networks with hidden neurons that are fully interconnected are notoriously hard to analyse. They have complex internal dynamics with neurons that fire at differing rates influenced by excitatory and inhibitory connection/synapse strengths and the external stimuli at one particular point in time. Therefore to gain a clear picture of the internal dynamics firstly the behaviour of the agent is observed, this is then correlated with the general pattern of neural activity. To make the referencing easier to follow the following diagram (diagram 5) represents each neuron within the network architecture.

Diagram 5:

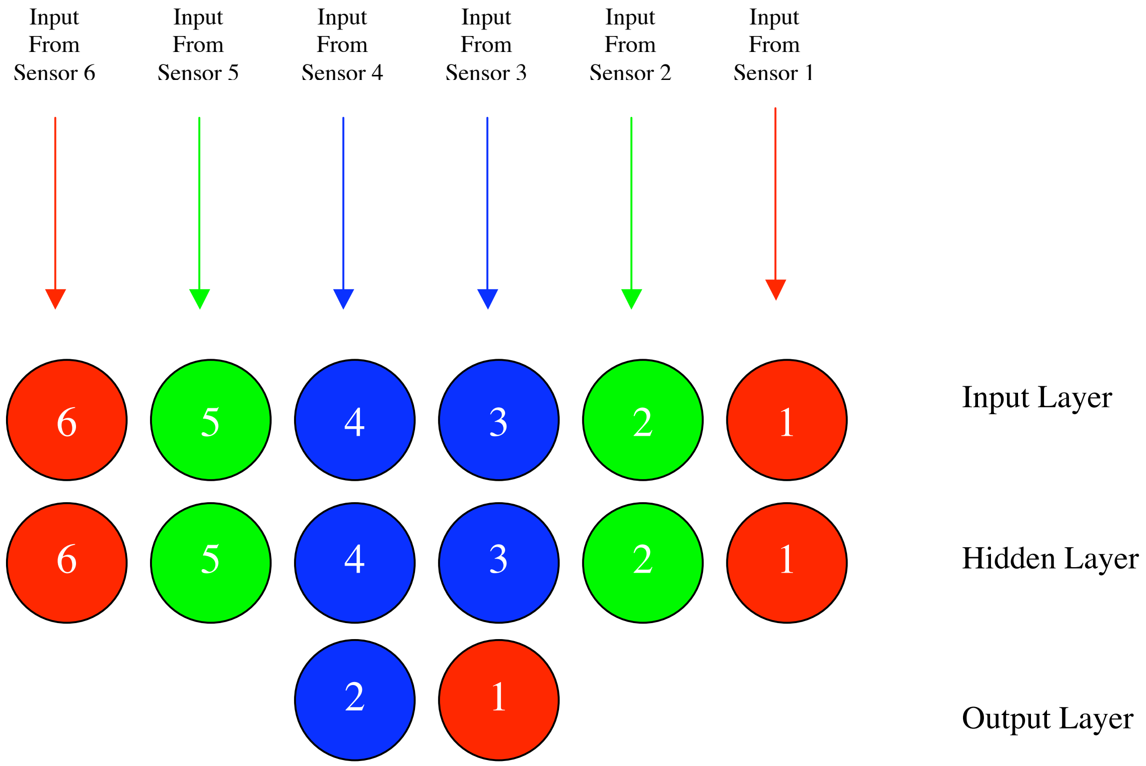

The sensors are set up as above, evenly spaced from right to left. Therefore as in the illustrated above, neuron one receives an input from the far right sensor, neuron two from the middle right sensor, neuron three from the inner right sensor, neuron 4 from the inner left sensor, neuron 5 from the middle left sensor and neuron six from the outer/far left sensor.

The colours correspond with the line colouring later used on graphs of neuron activity. To make a distinction between the left (neurons 4-6) and right side (neurons 1-3) of the network and sensors, the right sensors and neuron activities are drawn with a 'dash-dot' / ' - .' coloured line. The left sensors and neurons are drawn with solid coloured lines accordingly. As there are only two neurons in the output layer, the right motor neuron (1) is drawn with a solid red line and the left motor neuron (2) is drawn with a blue solid line.

Experiment one and two are an attempt to understand the neural activity of the agent in response to different stimuli.

Experiment one:

This experiment adjusted the offset of the circle at the beginning of each trial.

Constants: Circle radius = 15.

Random Variables: Starting position of the agent, set randomly between [150,350].

Controlled Variables: The value of circles offset, - for the first fifty trials the circle is positioned on the right of the agent, at the beginning of each trial the offset value is increased by two until it reaches a maximum of one hundred pixels on the fiftieth trial. The offset value is then reset to two on the fiftieth and the process is repeated with the circle now positioned on the left of the agent (for another fifty trials).

Results:

The agent successfully completed each trial, catching the ball with precision, the offset of the circle apparently having no effect on the effectiveness of the agent, the final score for the one hundred trials being 99.9%.

The typical behaviour/movement shown by the agent when caching balls on the right far (figure 2) and far left (figure 3) is shown below:

Figure 2: Horizontal Position (Movement towards the top of the Y axis = movement right, movement towards the bottom of the Y axis = movement left).

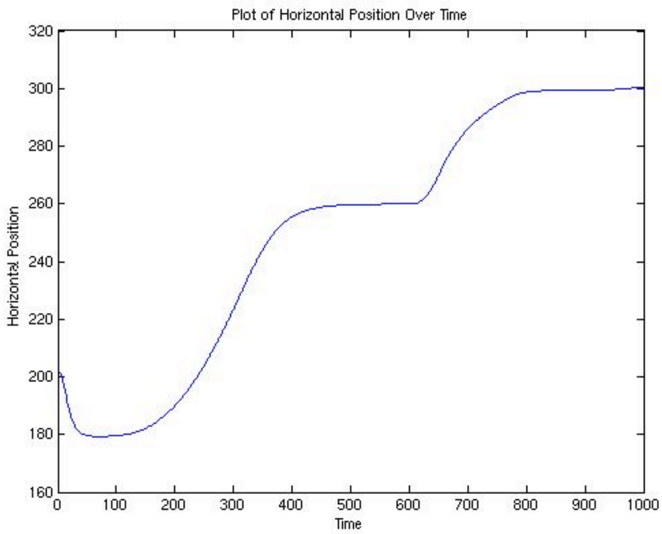

Figure 3: Horizontal Position:

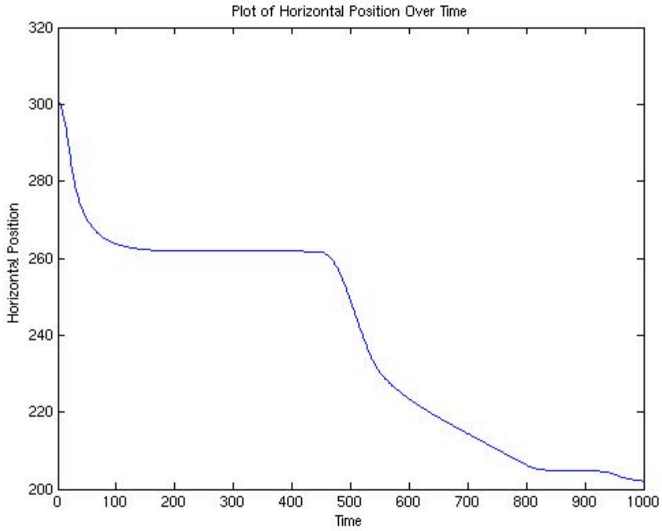

Despite not having an effect on the performance of the agent, the starting position/horizontal offset of the circle does have an effect on the behaviour of the agent. The agent has six evenly spread sensors, although the ray length is long enough to extend the full height of the environment, the spatial acuity of the agent is relatively poor. Some horizontal offsets will intersect the rays in the first few time steps, giving an indication to the agent as to where the circle is falling, however in other cases the horizontal offset might be outside or in between sensor rays, in which case the circle will not intersect the rays.

Both figure two and figure three show the agent performing similar behaviour, although in figure one the circle is to the right of the agent and in figure two the circle is to the left of the agent. The agent does not actively seek out a falling circle but instead appears to 'drift' either left or right until it receives a sensory input unless it receives an immediate input at the start of a trial. Once a sensory input has been received the agent proceeds to move in the direction of the input (much like a reactive control system). As more sensory inputs are received (more rays being intersected) the agent continues its movement towards the centre of the circle. Towards the end of the trial the circle intersects most, if not all of the sensor rays (the distance between each ray being less, nearer the agent thus the agent has a greater 'visual resolution' of close objects), the agent then remains stationary (having equal inputs left and right).

The neural activity over each trial supports the behavioural description, the input neuron activations and firings closely mirroring the sensor inputs and the output neuron activations and firings mirroring the input neurons. Thus the firing of the output neurons dictate the horizontal movement of the agent, i.e. the right neurons becoming more active when the agent receives sensory information from the right and the opposite occurring when a left sensor is active (left neurons becoming more active), consequently the agent moves left and right accordingly.

The following graphs (figure 4,5,6,7) relate to the behaviour in figure one above displaying the neuron activations in trial 49 where the circle is furthest from the agent on the right, the accompanying descriptions refer to diagram 5.

Figure 4: Sensory Inputs

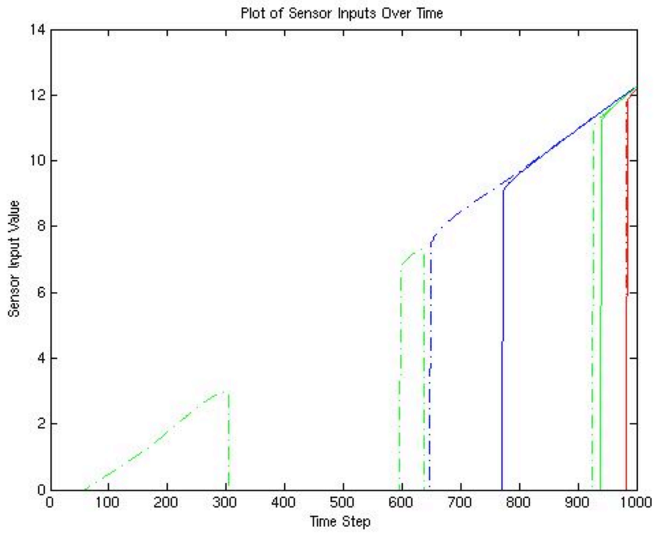

Although initially drifting left, at approximately the ninetieth time step the second right sensor is activated (green -. line in the figure four above), the value is small as the circle is still relatively far away, yet the activation is enough to cause the agent to take action. The input neurons show a mirror image of the sensory inputs, for example between 100 and 400 time steps only the right middle input neuron (neuron 2 in the network illustration (diagram 5)) is activated (figure 5).

Figure 5: Input neuron activations

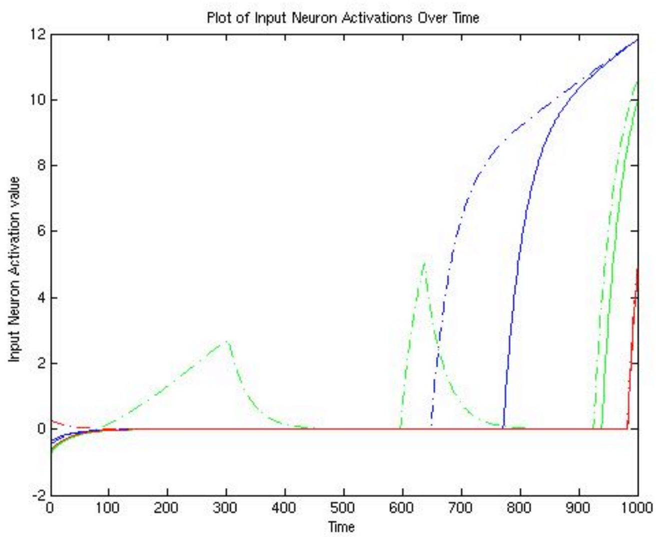

Figure 6: Hidden Neuron Activations

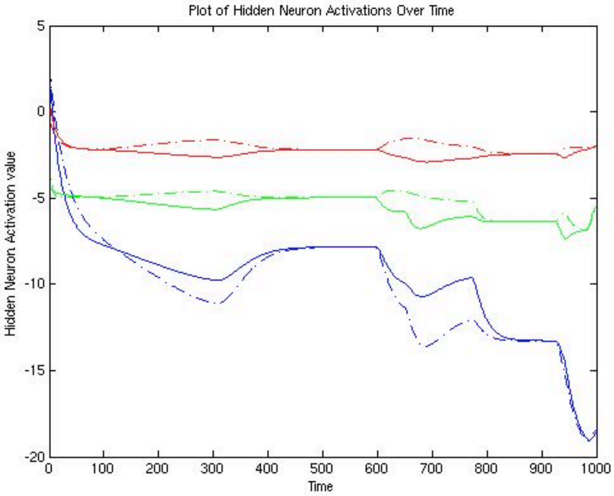

Figure 7: Hidden Neuron Firings

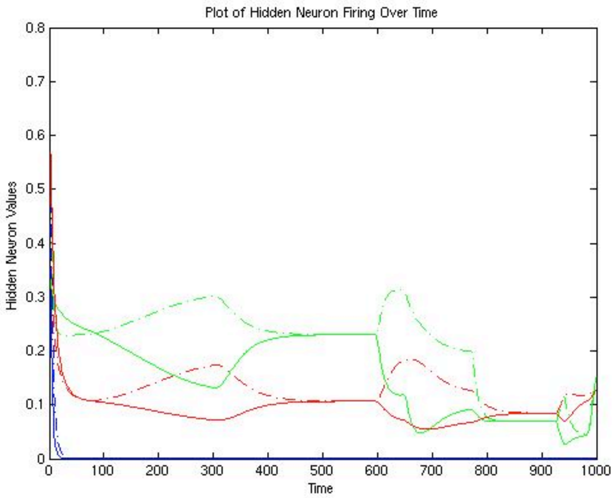

The hidden layer contains the most interesting activation and firing patterns (figures 6 and 7), with neuron 1 and 2 (red and green 'slash dot'; lines respectively) on the right hand side of the network oscillating in an opposite way to neuron 5 and 6 (red and green solid lines) on the left hand side of the network. The firing patterns for neurons 1,2,5,6 mirror the input neuron activations. These patterns again correspond to the behaviour, when there is an input to the left sensors the left neurons are activated whilst the right neurons become inhibitory, on the other hand an input to the right causes the opposite result. The firing of these neurons therefore has a direct effect on the behaviour of the agent.

Unexpected behaviour is seen by the inner hidden neurons (neurons 3 and 4 (blue lines)), which seem to have a inhibitory effect on the network, their activations oscillating in an opposite direction to the other neurons and the neurons themselves failing to fire throughout the trial (perhaps because they have relatively large negative bias' associated with them). The data suggests that these neurons are used to control the stimulation of the other neurons; the inputs to the network are highest towards the end of a trial, yet the amplitude of the activation and firing oscillations of the other hidden neurons (1,2,5,6) and especially the motor neurons, decrease in size (figure 8), at the same time however the activation states of the inner hidden neurons become increasingly negative (especially when the two middle sensors are strongly activated (sensors 3 and 4)) thus apparently offsetting/compromising for the strong sensory inputs. Therefore although not firing, this 'non action' in it self affects the whole network; their negative activations performing an inhibitory action on the other hidden neuron activations (being fully interconnected) and thus dictating the behaviour of the neurons in the output layer (figure 8). The end result of which is that as the sensory inputs become stronger the size of the agents horizontal movements decrease. The agent therefore remains stationery when the circle is very close and thus remains in the centre of the circle until the end of the trial.

Figure 8: Output Neuron Activations

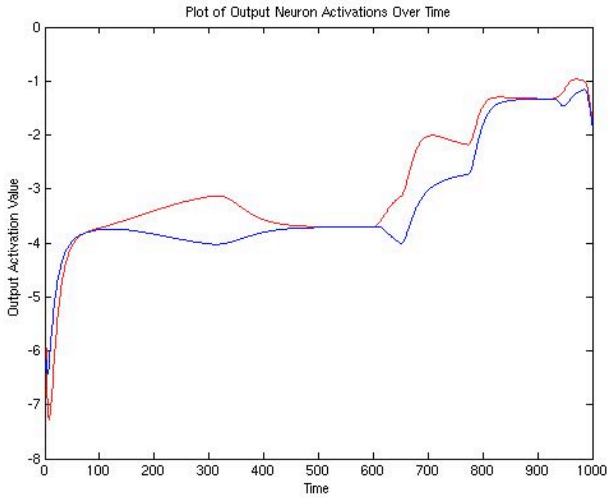

The above motor activations and firing patterns (figure 8) coincide with the recorded behaviour with the left neuron becoming more activated in response to stimuli to the left, the opposite applying for a right stimuli. When there are no or equal input activations the activations cancel each other out (remain level).

<u>Experiment two:</u>

This experiment was conducted to judge the effect of different circle sizes on the agent.

Constants: Horizontal offset of the circle, set to 45 for each trial.

Random Variables: Horizontal starting position of the agent, - set to a random value between [150,350].

Controlled Variables: For the first fifty trails the circle is positioned to the right of the agent, initially the radius of the circle is set to as value of: 1. The radius value is then incremented by 1 until it reaches a maximum value of 50 on the fiftieth trial. The radius value is then reset as in experiment two and the process repeats it self with the circle now appearing to the left of the agent.

<u>Results:</u>

Again the agent had no trouble adjusting to catching circles of different sizes achieving a score of 99.9% over one hundred random trials. Exactly the same behaviour was recorded as in experiment one, bigger circles intersecting more rays for longer thus producing larger sustained inputs (figure 9), whilst small circles result in single inputs every now and then (figure 11).

Whilst the agent could see the larger circle for almost the whole trial, and therefore gradually adjust its position (figure 10) smaller objects required the agent to make a quick decision, however despite seeming to be a more complex form of behaviour, the behaviour is generated by the same mechanisms as described in experiment one. Often the agent would wait for one ray to be intersected then move rapidly in the direction of the intersection (i.e. around the 700 time step), having only a few time steps to centre itself on the circle (figure 12).

Figure 9: Sensory inputs for a large circle:

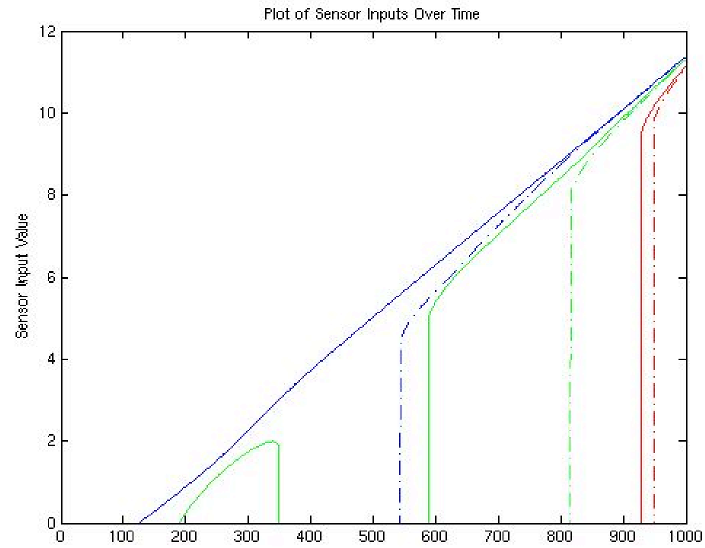

Figure 10: Corresponding horizontal movement in response to a large circle (gradual)

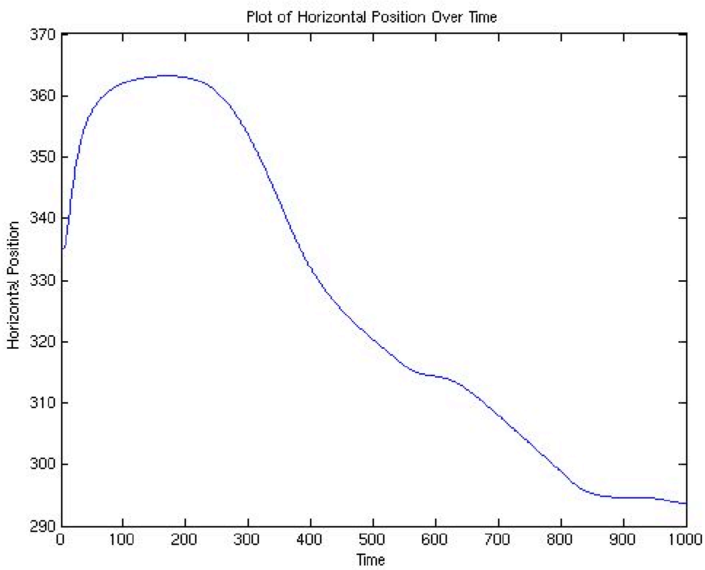

Figure 11: Sensory inputs for a small circle:

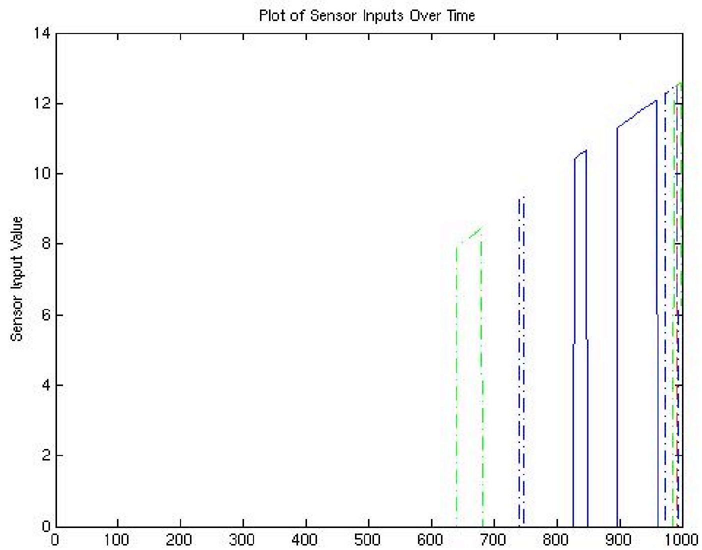

Figure 12: Corresponding horizontal movement in response to the small circle (rapid movement towards the end of the trial).

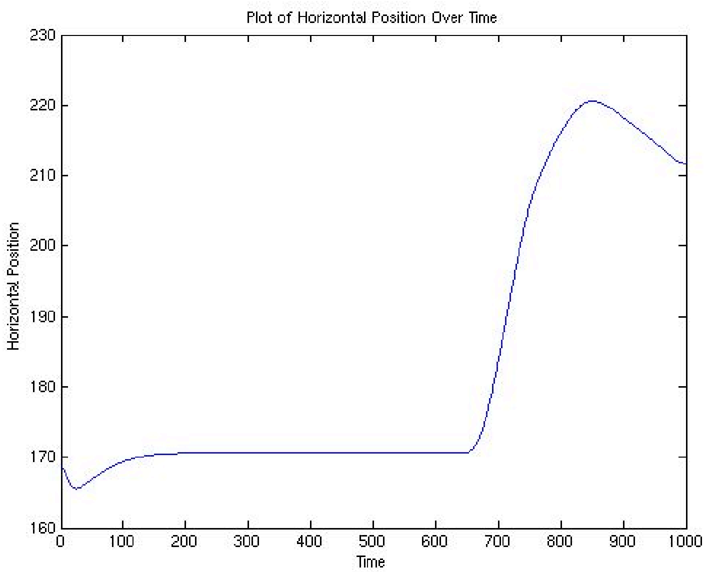

5.Experiements in size discrimination:

5.1 Experimentation Overview:

The aim of the following experiments was to analyse the physical and neural behaviour of different size discriminatory agents. In performing such experiments, it was hoped that a deeper understanding of the complexity of the task could be grasped. In addition the results should provide evidence of either relational discrimination or absolute discrimination as being the 'cognitive primitive' and when combined with the recorded data will hopefully go some way to explaining the results.

This section is split into four experiments:

5.3 Experiment One A: Agent evolved to choose the larger of two stimuli.

5.4 Experiment One B: Agent evolved to choose the smaller of two stimuli.

5.5 Experiment Two A: Agent evolved to choose stimuli of an absolute size.

5.6 Experiment Two B: Agent evolved to choose stimuli not of absolute size.

5.7 Experiment Three: Agent evolved to pick the 'middle size' of three stimuli.

If successful, in each experiment the evolved agent was usually subjected to three tests (although not necessarily in order), the general methodology behind each test is given below:

Test 1: Agent subjected to 1000 random trials, where the circle positions were set using the same horizontal position selector as described below in 5.2.

Test 2: Agent subjected to circles where the separation between the two increased by two on each new trial until a maximum of 100 was reached on the 50th trial. The circles positions were then swapped, so that the circle on the right would now be on the left and so forth. Therefore each circle was kept on the opposite side to the other circle in this test. The process is repeated for another 50 trials for a total maximum of 100 trials.

Test3: Agent subjected to two circles, at least one of which would have its radius size increased incrementally by one at the beginning of each new trial (according to the specific experiment). The same position swapping occurred after 50 trials as in test 2. Thus both circles were kept on separate sides (right and left) but in this case well within the agent's visual field (± 60 of the agents horizontal position).

The specific details of each experiment are mentioned when modified, but in each experiment the agent was scored in each trial using the same procedure as described in the original evolutionary evaluation function. The details of each circle and agent were recorded at each time step and stored in separate files; these were later imported into Matlab for processing and analyses (appendix C). The relevant and important results are reported in this section with appendix A used to display additional data (referenced appropriately within this section to emphasis specific points).

5.2 Evaluation Function.

The same evaluation and fitness function are used for these experiments (as described in section 4.1), however there are a few modifications to the evaluation function:

- Firstly, the experiments involve discrimination between two similar stimuli (distinct in size only). Two circles are dropped within the environment with their bases starting at the same height to avoid the problems reported in paper [22].

- Secondly, the radius size of each circle is set according to the discrimination task; if the goal is to always pick the larger of two stimuli then one circle is set to be bigger than other, the opposite applies if the task is to pick the smaller of two stimuli. If however the task is set to pick a stimulus of a specific size, then one circle remains a constant size throughout the evolutionary trials and the other changes - either having a radius bigger than, smaller than, or very similar to that of the constant stimuli (details are provided for each experiment). Therefore before evolving the agent the radius assignment was checked to ensure that the fitness function rewarded the agent according to the specified task, if not then either the radius was adjusted or the fitness function changed using position data from another circle.

The horizontal position of both the circles will have an impact on the capability of the agent to generalise and will also affect how the agent solves the problem. It was decided to always keep the two stimuli separate with at least a small gap between them in the evolutionary process. The circles horizontal position was set randomly and could either be both positioned on the right or on the left (see below) of the agent (one being between [1,50] from the agents horizontal position and the other being between [100,150] pixels (within the agents visual range)), or the circles could be positioned individually on opposite sides of the agent (one on the right and one on the left of the agent between [20,100]).

To present the agent with a representative number of examples, in terms of the placement of each circle, the positioning of each circle was also controlled by a variable 'num'. If the num value equalled '0' the agent would be presented with the 'wanted circle' on the left'; a random variable would then dictated whether both circles would be placed on the left and if so which circle would be closest to the agent. At the end of the trial the 'num' variable was set to 1. The next trial would then see the opposite happen (circles placed to the right of the agent) and at the end of the trial the variable would then be set back to zero. This placement mechanism alternated the positions of the circles back a forth for the set number of trials, thus making sure there were an equal number of trials right to left.

Therefore each agent was evaluated over 20 trials with the average score being taken as it fitness. The mutation rate was left at 100% and recombination rate left at 90% unless otherwise stated.

5.3 Experiment One A: Selecting the larger of two stimuli

Details

In the evolutionary run the circle radius values were kept constant at:

Circle one: 20
Circle two: 10

Evolutionary results

Figure 12 Plot of fitness against tournament number:

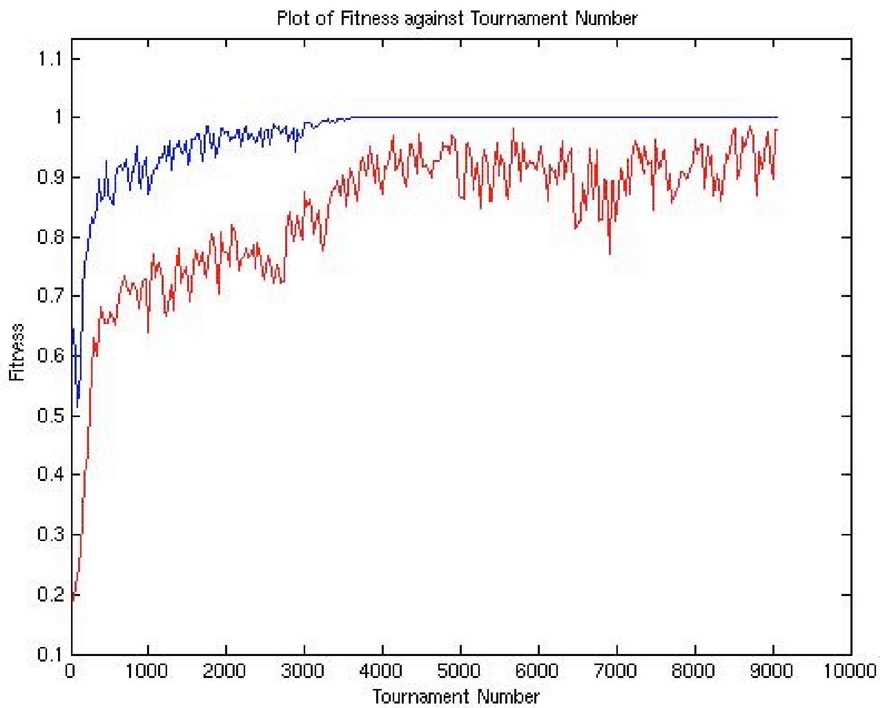

Agents with very high fitness were relatively easy to evolve (figure 12), achieved in fewer than 2000 tournaments (approx 66 generations), the maximum fitness at the end of all the tournaments being 99.9%.

The agent was tested over 1000 random trials (test1), in each trial the radius of circle two was set to 20, and the radius of circle one varied randomly between the range of [30,40]. The agent was rewarded as in the evolutionary run, - the closer it moved to circle one (the larger stimuli) the higher the score. The resulting test score of the agent was 98.2%, the agent mostly catching the larger of two circles. Likewise, if the agent is scored on approaching circle 2 then a very low score is achieved 0.08% over 1000 trials. Thus evolution seems to be exploiting relational discrimination rather than absolute size (despite having the opportunity to use either, i.e. in the evolutionary run the agent had been trained to catch the circle of radius: 20, this circle happened to be the larger of the two stimuli, but as both circles remained a constant size the agent could have easily relied on judging the absolute diameter of the circle.)

The neural activity is now briefly analysed to try and explain this behaviour (recorded in test 2 and 3), in test 3 both the circles were kept in constant positions, one circle positioned plus 60 pixels from the centre of the agent (on the right) and the other – 60 pixels (on the left of the agent)). However the radius of the circle two was incremented by one each new trial and circle one assigned to be double circle twos' radius. Please use the referenced appendix to aid understanding.

The agent achieved an average score of 75.3% in test 2 and 77.2% over the 100 trials in test 3. The neural activity of the agent is very similar to the agent in section 4 (capable of visual orientation). The input neuron activations and firings mirroring the sensory inputs at each time step and the left and right motor firings clearly corresponding to the behaviour of the agent (its horizontal position). Unsurprisingly the activations and firings oscillate one direction if the bigger stimulus is on the right and in the opposite direction if the large circle is on the left (a demonstration of the networks bisymmetrical nature). At a higher level the behaviour is slightly more complex, with the agent appearing to engage in a form of active scanning, moving to a from both stimuli then eventually making a decision. Typically the agent performs an initial relatively large active scan between 200^{th} and 400^{th} time steps and then gradually moves in the correct direction, performing another significant active scan typically between the 600^{th} and 800^{th}time steps; the amplitude of active scans/oscillatory movement decreasing with time (see appendix A 1.1).

The physical and neurological behaviour of the agent over all the trials remains relatively constant. The agent appears to rely on the largest stimuli to dictate its decision. For example in most trials the agent can see both circles, as both circles fall they intersect the rays at different points, the larger circle tending to intersect more rays and to a greater extent. On the trials that achieve high fitness scores the agent makes a decision to travel in one direction very quickly, after which the active scanning is a result of the agent briefly encountering a strong stimulus from another direction causing the network to reverse its oscillations temporarily. Slowly however the inputs from the bigger circle 'override' these interruptions and although in most cases this is only by a very small amount (inputs similar), this input becomes amplified through the network and is enough to change the agents' direction (see appendix A 1.2).

Where the circles are very small the agent suffers from the same visual hindrances ad described in section four, smaller circles rarely intersecting the rays due to the agents poor visual resolution. Initially the agent moves towards whichever stimuli it sees first. Often when the circle sizes are small there is a 50:50 chance that the larger of the two stimuli will be chosen by the agent. If however the smaller is chosen, the agent fails to correct its decision. The reason for this is that the difference between the two stimuli is very small (although the relative relationship of one being double the size of the other is still held); therefore if the wrong choice is initially made, the agent begins to move in the direction of the smaller circle, by the time the larger circle intersects a ray sensor the agent is already approaching the smaller circle, if the size of both circles only differs by a very small amount then the bigger circle is nothing more than a 'glitch' in the sensory inputs, the size of the stimuli being comparable to the smaller circle. Again because this glitch happens late in the trial the agent is often too far away to catch it even if it does decide to review its choice (see appendix A 1.1).

In addition when there is a very large relative difference between circle sizes (both being large) the agent also fails to make the correct choice, the performance of the agent in this case strongly connected to he positioning of the circles and is reiterated in the results of test 2. The decrease in performance is a consequence of 'circle overlapping'; in such a case the agent is influenced strongly by both stimuli, becoming pulled back and forth for one or two oscillations and then settling in the middle of both stimuli (see appendix A 1.3). There is no clear distinction that the agent can make as both circles form one large stimulus, the agents performance only increasing when the separation between the circles is larger than 34 pixels (although achieving a high score at trial one in test 2 where the both circles are 'on top of each other', the larger covering the smaller completely).

5.4 Experiment One B: Selecting the smaller stimuli.

Details

In the evolutionary run the circle radius values were kept constant at:

Circle one: 20.
Circle two: 40.

Evolutionary results

Figure 13: Plot of fitness against tournament number

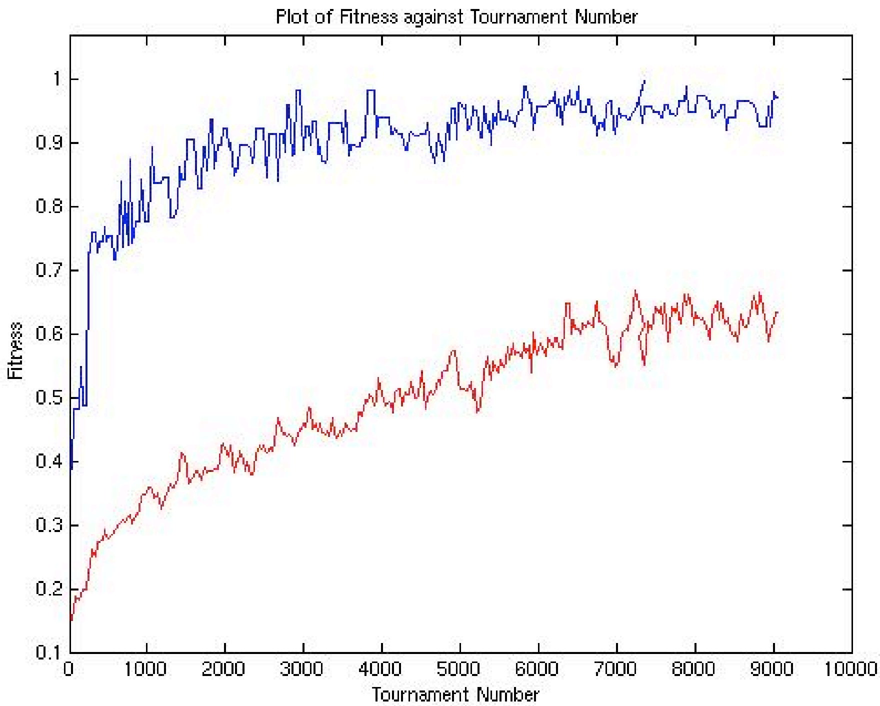

Evolving an agent to choose the smaller of two stimuli was considerably harder, requiring large relative circle radius size differences to be chosen and a population size of one hundred individuals.

After 9000 tournaments an agent of fitness: 97.3% was evolved, though the average fitness of the population only being 63.4% in the last tournament (figure 13). The agent was subjected to exactly the same trials as in experiment 1A (above) and achieved a score of 85.4% over 1000 random trials in which the agent was scored on how close it moved to the smaller of the two circles (test 1). The results indicate that the agent evolved to use a form of relational discrimination rather than always choosing the circle of the same size. To support this when a circle of radius 10 and a circle of radius 20 are presented to the agent it repeatedly selects the smaller of the two stimuli.

When subjected to test 3, the average score over the 100 trials was 72.4%, the agent capable of choosing the smaller of two objects reasonably accurately even when the two circles overlap. The performance does start to decrease slowly as the relative size difference between the circles increases (i.e. both the circles become very big and overlap completely). However the main cause of poor performance is the agents' inability to reliably judge the difference between two small circles (radiuses under 10).

The agent appears to engage in similar behaviour to that described in experiment 1A, however in this case it clearly cannot just be following the larger stimuli. On closer analysis one will see the agents behaviour depending strongly on the intersection of specific sensory rays. For example if the outer right hand side sensor (sensor 1 –refer to diagram 5) was stimulated, the agent moves in the opposite direction to the stimuli (in this case to the left). The opposite is true for sensor 6 (outer left sensor) where the agent moves to right following an intersection. The inner sensors behave in the same way, although appearing to have a slightly smaller effect (if sensor 3 (right) is intersected the agent moves left, if sensor 4 (left), it moves right). However if sensor 2 or 5 (middle sensors) is intersected, the agent will move towards the stimuli, moving right if sensor 2 (right) receives an input and left if sensor 5 (left) receives an input.

In this way then the agent is both attracted and repelled from circles, the decision happens when the ‘repulsion sensor’ inputs exceeds the ‘attraction sensor inputs. Often in the 100 trials, the bigger of the two circles will intersect a ray first (usually one of the middle sensors 2 or 5), whilst the smaller circle falls ‘in-between the rays’. This sensory input initially causes the agent to move towards the larger stimuli however as it does so another ray is likely to be intersected either by the same circle or by the smaller circle, depending on which ray it is the agents movement will be hindered or helped. As the trial continues the agent is pulled/moved back and forth accordingly, the bigger stimuli having the most control and in addition an intersection of the ‘attractive sensors’ (sensors 2 and 5) appearing to have a much stronger influence on the agents choice (However two ‘inhibitory stimuli’ or another ‘attractor stimuli’ from another circle can work together to combat this strong pull). Towards the end of the trial the agent makes its decision, on most occasions being dictated by whichever circle breaks the outer rays first. The outer rays having a relatively strong effect that overcomes the ‘attractive force’ of the larger circle and combines with the ‘attractive force’ of the correct circle consequently dictating which direction the agent will move in.

This strategy works because the larger circle will tend to intersect the outer ray first (larger diameter). The reason why the agent mostly fails when the circles are small is that by the time the bigger circle intersects the outer ray the trial is finished. Having also been subjected to test 2 (score achieved = 78%) another weakness in the agent’s technique is highlighted, the performance of the agent decreasing as the circles separation falls below 30 pixels, at this point the circles overlap considerably and thus the agent is presented with one large stimuli/circle.

Although the tests on the agents did not represent all the possible variations it did however give a vital clue as to how the agent solves the problem see Appendix A Experiment One B Details for an example.

5.5 Experiment Two A: Selecting a circle of absolute size:

Details

During evolution the circle radius values were:

Circle one: 20.

Circle two: Random value between [1,40] drawn from a uniform distribution.

Evolutionary results

Figure 14: Plot of fitness against tournament number:

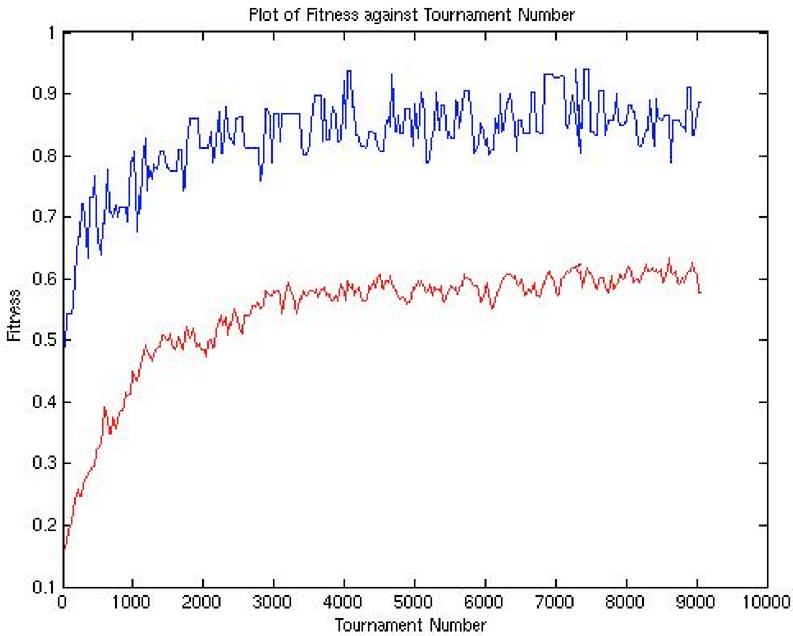

A population of 100 individuals was evolved for 9000 tournaments (roughly 90 generations), the resulting top scoring genotype was 88.7% fit and the average fitness of the whole population was 57.6% fit. Figure 14 displays the evolutionary progression the red line representing the average fitness and the blue line representing the fitness score of the fittest member of the population every thirty tournaments. Large fluctuations are clearly pronounced in the fittest fitness scores, this instability continuing to the last tournament (unlike figure 13, where the amplitude of the fluctuations decrease in size over time). In addition the average fitness remains low, with a value oscillating between 0.5 (50%) and 0.6 (60%). Despite these average results the agent was tested over 1000 random trials (test 1, in this case the radius of circle 2 was between 1:40) and interestingly achieved a score of 65%. The agent was also subjected to test 3, in this case the size of circle 1's radius remained constant at 20 throughout each trial, whilst the radius of circle 2 was incremented by 1 each new trial, reaching a maximum of 50). The resulting score was low, and further analysis revealed that the agent was merely picking the larger of the two stimuli. A selection of scores from variations of test 1 demonstrates this:

	Circle 1 radius	Circle 2 radius	Average score over three trials:
1)	20	40	48%
2)	40	20	62%
3)	10	20	30%
4)	20	10	84%
5)	10	5	82%
6)	20	20	58%
7)	25	20	5%;
8)	20	16	61%

The agent managed a score of around 60% in 1000 random trials, this above average result is most probably due to the fact that the radius of circle two was chosen randomly, there was an equal chance of the random number being less than 20 or greater than 20. If the agent was picking circles of absolute size then one would expect that if presented with two circles of a similar size there would be a 50:50 chance of selecting either. Indeed this does reflect in the score (20:20 –58%), however if relying on purely absolute size discrimination the agent should ignore two circles that it has not been evolved to select (I.e. number 5 above). Yet the data above demonstrates that given the choice, the agent will select the bigger of the two stimuli. Consequently the agent achieves low scores when circle 1 is smaller than circle 2 regardless of the absolute size of circle one. The neural behaviour was also briefly analysed and contained similar traits to that observed in part 5.3 Experiment One A.

The agent was indeed selecting the larger of the two stimuli, the evolutionary process seemed to have exploited the fact that on average over the twenty trials, in at least half the cases (in most cases more) choosing the larger of the two stimuli would be the correct action to take (i.e. circle 2 would be smaller than circle one). The extra points accumulated being facilitated by the relatively random positioning of the circles. Further investigation was not pursued in this case.

To combat this and to hopefully increase the likelihood of evolving a robust, fit agent the evaluation function was modified. Rather than placing the circles (individually or both) randomly left or right, the agent was presented with four different circle position configurations.

1) The first trial placed the circle with the constant radius size on the left of the agent and the circle with the varying radius on the right.
2) The second trial placed both circles on the left of the agent (with a sufficient gap between them), the circle closest to the agent chosen randomly.
3) The third trial placed the constant circle on the right and the varying circle on the left.
4) The final fourth trial placed both circles on the right, again with the circle nearest to the agent chosen randomly.

(This modification is now employed in all new experiments).

Therefore the agent was always subjected to four circle configurations for five times over the twenty trials. The results of implementing the above proved to be significantly better:
(Agent 1)
Fittest fitness: 94.7%.
Average fitness: 72%. (Figure 15)

Test 1 score: 77%.
Test 2 score: 74.8%
Test3 score: 82%.

Figure 15: Plot of fitness against tournament number:

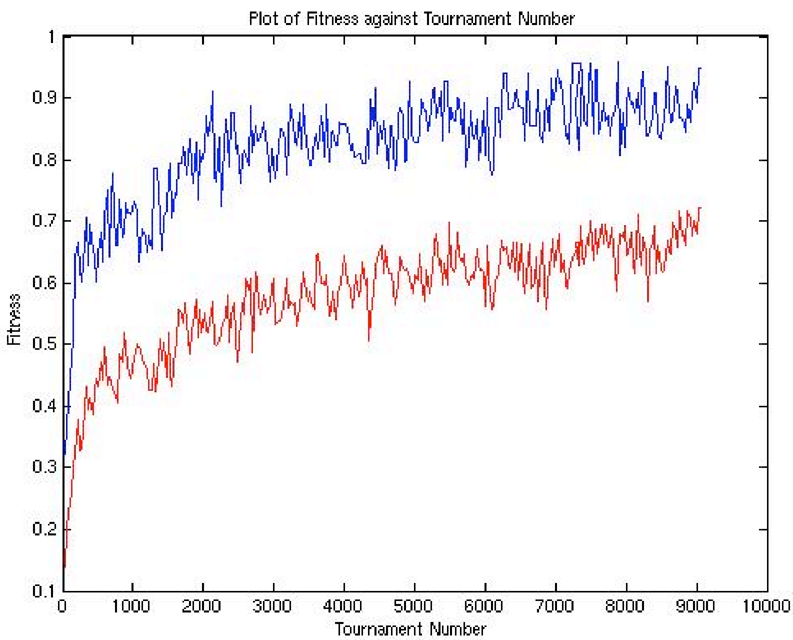

A further agent (agent 2) was evolved in attempt to evolve an agent with smother fitness curve/progression. To achieve this the number of trials in the evaluation function was increased from 20 to 40. Having implemented this change the resulting top fitness and average fitness were 82.9% and 64.8% respectively. Despite having a lower fitness compared with agent 1, agent 2 achieved a reasonable result in test 1 (73%), and in test 3 scored 78.7%. More importantly the extra trials seemed to have produced a much more decisive agent with its evolutionary process being considerably smoother (Figure 16) (fewer large amplitude fluctuations).

Figure 16: Plot of fitness against tournament number:

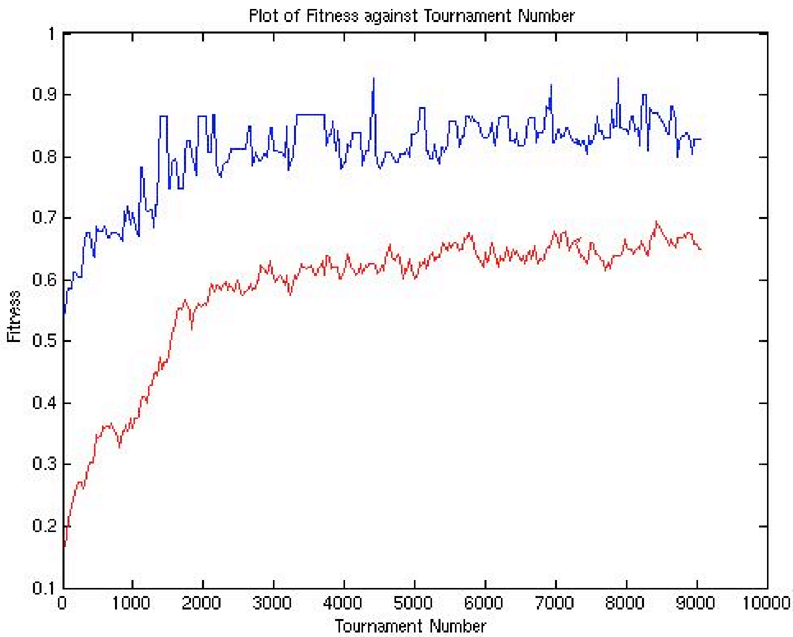

Having evolved two adequate agents, an attempt was made to understand how the agents achieved their scores. Interestingly both agents used an almost identical technique, both of which seem to be heavily based on the strategy identified in 5.4 where an agent is evolved to select the smaller of two stimuli. As in 5.4 the decision to select a circle is based on the response of the neural network to the intersection of certain sensors. An intersection of the inner sensors (3 and 4 (diagram 5)) will have a weak repulsive effect, causing the agent to gradually move in the opposite direction to the intersection (if the input is from the left sensor the agent will respond by slowly moving right). An intersection of the middle sensors (2 and 5, see diagram 5) will have a strong attractive effect, often 'overriding' sensors 3 and 4 causing the agent to move towards the intersection and therefore approach a circle/stimuli. Finally the outer sensors (1 and 6) also result in a 'repulsive' response from the agent, although seeming to have a much stronger effect than the inner sensors pushing the agent away from an intersection.

The general pattern of neural activity is also similar to that previously recorded, although the hidden neuron behaviour is distinct for each agent, in the first agent (20 trials per evaluation) only the middle hidden layer neurons (3,4) fire and have direct control over the output neural firing, whilst in agent 2 (40 trials per evaluation) the outer hidden neurons control the network outputs unaccompanied. In both cases the agent moves according to the sum of the attractive and repulsive forces (refer to Appendix A: Experiment Two A for specific examples).

The evolved agent is therefore capable of choosing between the two circles when circle 2 is larger than circle 1 (radius = 20). In this case the agent sees the circles early on in the trial, being attracted to both equally as both tend to intersect the middle sensory rays. The agent is pulled back and forth in an identical fashion to that described in 5.4. Eventually however as in 5.4 the larger circle will inevitably break the outer ray before the smaller circle, consequently moving the bias of movement towards circle one (see Appendix A: Experiment Two A: Example 2). At this point then the strategy does not appear to be relying on absolute size. This is supported by further tests, including a modification of test 2, that show that the agents' strategy is biased towards selecting the smaller of the two circles. For example when test 2 is undertaken with the radius of circle 1 set to 20 and radius of circle 2 set to 40 high scores are achieved whilst if circle 2 is set to 10 much lower scores are achieved, - indicating that the agent is selecting the smaller circle. Despite this the agent still scores reasonably, and as described below if the circles are positioned favourably and have the right diameter the agent will succeed at choosing the correct circle (radius 20).

If circle 2 is smaller than circle 1 the above technique will not work. Despite this, the results of test 3 indicate that the agent is successfully choosing the correct stimuli. When analysing the agents behaviour in such a situation, the common solution employed by the agent takes the following form: assuming both circles are within the agent's field of view, if circle 2 is smaller than circle 1 which in test 3 is set to be 20, the agent will tend to detect circle 1 (the correct circle) first (see Appendix A: Experiment Two A: Example 1). Initially then the agent will become attracted to the larger stimuli and may start to move slowly towards it. As it does so the smaller circle (circle 2) is likely to intersect another ray and for a short period of time the attractive and/or repulsive 'forces' are balanced. However this period of equilibrium is usually brief, as the smaller circle will pass through the sensory ray rapidly, whilst the larger circle will intersect the ray for much longer. Therefore when the smaller circle falls past the 'middle sensory rays' the agent begins to move towards the larger stimuli (circle 1). As it does so however the smaller circle intersects the agent's outer ray, this then causes the agent to move rapidly towards the larger stimuli. The final decision is again dictated when the outer ray is intersected by the smaller circle; here to 'tip the balance' the repulsive force of the outer sensor (intersected by circle 2) combines with the attractive force of the middle sensor (intersected by circle 1) on the opposite side of the agent (see Appendix A: Experiment Two A: Example 1).

In this way then both agents appear to select the circle of absolute size, however it is more by luck than judgement; the results depending heavily on the initial positions of each circle and the circle sizes themselves (as reiterated when subjected to test 2). The results of test 3 are misleading, on the one hand they clearly show that the performance of the agent decreases when discriminating between similar sized circles (For both agents ± 10 of the constant radius appeared to be the limit, after which (any smaller differences) the results become unreliable, with generally a 50:50 chance of the agent making the right choice). On the other hand the agent achieved high scores (although less reliable than with larger circles) when circle 2 is small. With the evidence given above one would expect lower scores. However the horizontal position of each circle in test 3 is kept constant at +- 60, and it just so happens that this particular set up is favourable to the above strategy (the smaller circle tending to fall between the rays for long periods of time). The agent's ability to select the larger (circle 1) in a situation where circle 2 is smaller than circle 1 (circle of specific radius) is strongly dependant on the diameter and the horizontal positions of both circles.

Finally to gather further evidence of relational discrimination being used, the agent was tested with one circle only. The results show that the agent will tend to orientate towards any stimuli regardless of the size of the circle. Which ever breaks the outer sensors first the agent will avoid; if another circle is present then it will choose that. The third experiment has only had partial success then, yet raises several interesting questions commented on in section 6 (discussion).

5.6 Experiment Two B: Selecting stimuli not of absolute size

Details

During evolution the radius values of each circle were:

Circle one: A random value between [1,40] drawn from a uniform distribution.

Circle two: 20.

Evolutionary results

Figure 17: Plot of fitness against tournament number:

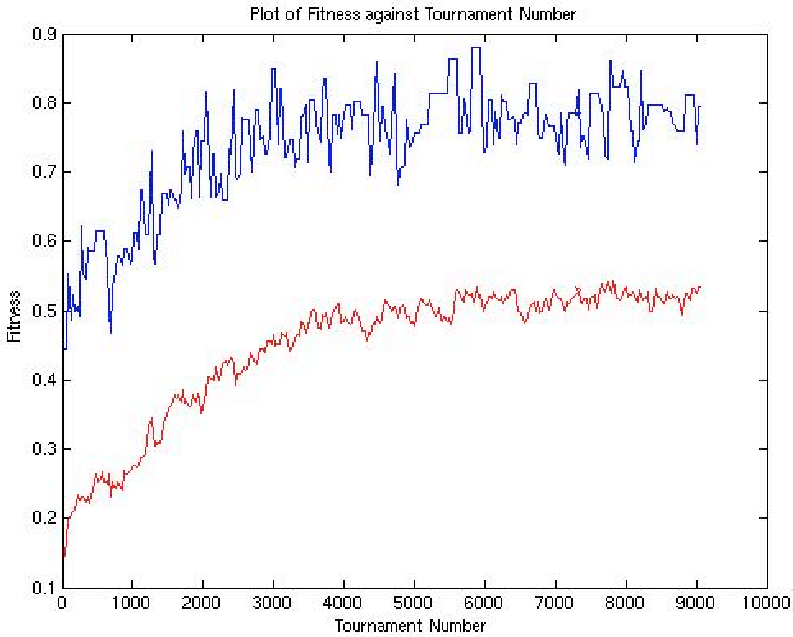

Poor evolutionary results were achieved in this case (figure 17). All attempted improvements/modifications to the evolutionary process such as varying the size of the population; the number of trial numbers and/or varying the mutation rate had little effect. The fitness score at the end of the evolutionary period was 79.4% for the fittest fitness and 53.4% for the average fitness (using a population of 100 with 40 trials in each evaluation).

Unsurprisingly the agent performed unsatisfactorily when subjected to test 1, where in this case the radius size of circle one was set to be a random number between one and forty whilst circle two remained constant at twenty. It was expected that the agent should score slightly higher than average and should for the majority of the time select circle one over circle two despite it's average fitness. However the agent only achieved a score of 57.2% over the 1000 random trials. If the agent were accurately and deliberately choosing not to select the circle of a specific size (circle 2) one would expect a higher score. Further testing revealed that the agent was in fact selecting the bigger of two stimuli (having similar patterns of neural behaviour to the low scoring agent in 5.5). The figures below display this, in each trial the agent marked on how far it moved towards circle one (modifications of test 1).

Test	Circle 1 radius	Circle 2 radius	Average score over three 1000 long trials:
1)	20	40	14.6%
2)	40	20	92.8%
3)	10	20	26.5%
4)	20	10	85.3%
5)	10	5	73.4%
6)	20	20	49%

Test 1 and 2 above return positive results, on both occasions the agent seemingly avoiding the circle of radius size 20. However tests 3 and 4 soon show that the agent is apparently selecting the bigger of two stimuli regardless of the circles absolute size. Test 5 and 6 confirm this, the agent apparently achieving better results when the circles are of a larger size and therefore a larger relative difference. Test 2 and 3 of the main tests were also completed confirming the above results. In experiment three the agent was scored on how far it moved in the direction of circle 1 (circles 1 radius size increasing by 1 each trial, whilst the radius of circle 2 was kept constant at 20). Test 3 surprisingly returned a score of 65%, the agent however achieving high scores when circle 1 is bigger than circle 2 whilst poor unreliable scores when the radius of circle one was below 20.

One slightly strange observation was that the hidden neurons fire opposite to the input stimuli and then the output neurons then fire in opposite way to the hidden neurons. For example if left sensor stimulated, right hidden neurons would become more stimulated and the activations of the left neurons become more negative, this causes the left output neuron to become stimulated, and the agent would consequently move left.

In conclusion then evolving an agent to select a stimuli not of a specific size was particularly hard, the product of which had evolved to always choose the larger of each stimuli.

5.7 Experiment Three: Selecting the 'middle size' of three stimuli (Extension).

The source code and thus experimental set up was modified (diagram 6) to accommodate three circles (also involving updating the agents sensory system, so that it could detect three circles). Several pilot tests were conducted the results of which proved that evolving an agent to pick the circle of middle size between three stimuli was harder than anticipated. Although seeming to be a relational task, the agent once again had the chance to evolve an absolute strategy; therefore the three circles were kept of constant size throughout each evaluation trial. To make the task easier for the agent, each circle was kept separate with one to the left of the agent between 100 – 150 pixels from the centre of the agent, another being roughly central to the agent between 20 pixels either side, and finally one positioned to the right of the agent again between 100 – 150 pixels form the agents initial horizontal position. Each circle was assigned a given radius size, circle 1 (middle circle -on which the agent is scored on approaching) = 20, circle 2 (the smaller circle) = 5 and circle 3 (the larger circle) = 40. The large differences were needed to achieve agents of high fitness; several previous attempts to evolve an agent to distinguish between circles of size 5,10 and 15 had failed.

There are six combinations of initial circle positions: (S = small circle, M = middle sized circle, L = large circle).

Right	Middle	Left
L	M	S
L	S	M
S	M	L
S	L	M
M	L	S
M	S	L

For evolution the population size was increased to 200 and in order to reliably evolve robust agents the number of evaluation trials was set to 24 (4 presentations of each combination). Initially the evaluation function presented each of the above combinations in tern, one after the other. However the resulting agents proved to be poor at generalising, the agents only scoring highly if the circles were presented in the correct order or otherwise only favouring specific combinations. Despite this intriguing result no further experimentation was undertaken. Instead another agent was evolved in an attempt to generate more robust behaviour. This included subjecting each agent to 48 random trials per evaluation (the circle sizes were kept the same although the order in which they were presented was randomly chosen).

To obtain agents of adequate fitness the evolutionary regime required a considerable number of tournaments (10000 – 20000 - most probably due to the additional trials). After 20000 trials (figure 18) an agent with a fitness of 83.2% was produced although the average fitness of the population was only 48.7%. In this case 20000 tournaments is approximately equal to 666 generations, - the population was sampled every thirty tournaments to gain a detailed picture of fitness change over time (despite the population size being 200).

The evolved agent returned an average score of 67% over three 1000 random trials (variations of test 1). The circle sizes were also changed and the agent re-subjected to the same three trials. Unfortunately poor results were achieved with the agent often returning a score of approximately 30%. The generalisation ability of the agent is considerably better than that of previous agents, yet still not particularly good, - the agents technique highly reliant on the specific size of each circle.

Next the agent was tested with each combination of circle positions (given above), each was presented individually and an average taken over three trials. The table below displays the resulting scores (calculated for approaching the middle circle, - reordered from above according to score).

Combination:	Average Score:
LSM	53%
MSL	56%
MLS	61%
SLM	63%
SML	84%
LMS	85%

The agent scores particularly well when the 'middle sized' circle is positioned centrally between the other two circles. If the largest circle is in positioned centrally, the score drops, with the worst score being achieved when the small circle is located in the centre.

Interestingly the agent employs exactly the same strategy as seen in section 5.5 (agent evolved to select a circle of absolute size) and section 5.4 (small size discriminator). Although the connections were not initially realised, on closer inspection the similarities between the tasks become apparent. Firstly, in section 5.5 the agent is tasked to select a circle of absolute size, the specific size of which was set to be in the middle of the presented random range (from which the size of the other circle was drawn). At any one time the agent could be presented with either a circle that was larger or smaller than the circle of specific size. In this case exactly the same applies, here the agent was again evolved to select a circle with a size almost exactly central between the larger and smaller circle. The same techniques used to solve the problem in 5.5 therefore work in this situation although in this case the agent has to perform its choice when both the larger and smaller circles are present (deal with three stimuli at one time). A close study of the neural behaviour confirms that the agent is indeed following the same environmental features as in 5.5, at a behavioural level we again see that an intersection of the outer sensor rays (1 and 6) will cause the agent to move away from the direction of the intersection; whilst an intersection of the middle sensory rays will cause the agent to move towards the direction of the intersection (as in 5.5 the inner rays also appear to have less of an effect). At a neural level, the hidden layer outer neurons (1 and 6) firing in response to the strongest inputs, whilst the interconnections cause the motor neurons oscillate in an opposite direction. (See appendix A: Experiment three for specific example).

Unfortunately the agent is vulnerable to the same mistakes, in some cases the trial can be favourable, for example if the 'middle sized circle' is positioned above the agent (between the other two circles), then firstly the agent is already in a good position (being directly below the middle circle). Secondly, the larger circle is likely to intersect the outer ray first (being to the left or right of the middle circle) and the smaller circle is unlikely to influence the agent for long if at all (intersecting few rays very quickly or the agent being pulled towards the larger circles initially, the smaller circle thus passing out of visual range). Consequently, although the agent will be pulled back and forth it will end up selecting the correct circle. On other occasions the larger circle is centrally positioned above the agent, and the middle circle will tend to intersect an outer ray first, the performance drops in such a case. If the smaller circle is positioned centrally the agent obtains the worse score, only catching the correct circle 50% of the time. In such a situation the agent can be attracted in either direction according to the positioning of the circles. Ultimately the larger circles will intersect an outer ray, the agent will consequently move towards the smaller circle (thus obtaining a low score).

In conclusion the ability of the agent to select the correct circle is once again highly dependent on the size and horizontal position of each circle. Surprisingly although a very similar strategy is used as described in section 5.5, evolution was considerably more difficult. This may be a reflection of the increased complexity of the task (more stimuli to respond to) or purely a flawed evaluation function and/or a poor choice of 'evolutionary parameters'. Pilot tests have revealed that similar scoring agents can be evolved with smaller populations of around 100 individuals, fewer trials and fewer tournaments; this is therefore 'on par' with the evolutionary parameters and performance in section 5.5. The experiment has been semi successful, although time and report limitations have prevented further work. More research is needed to understand why the agent performed quite so badly when exposed to circles of different sizes.

Diagram 6: two examples of the agent interacting with three circles.

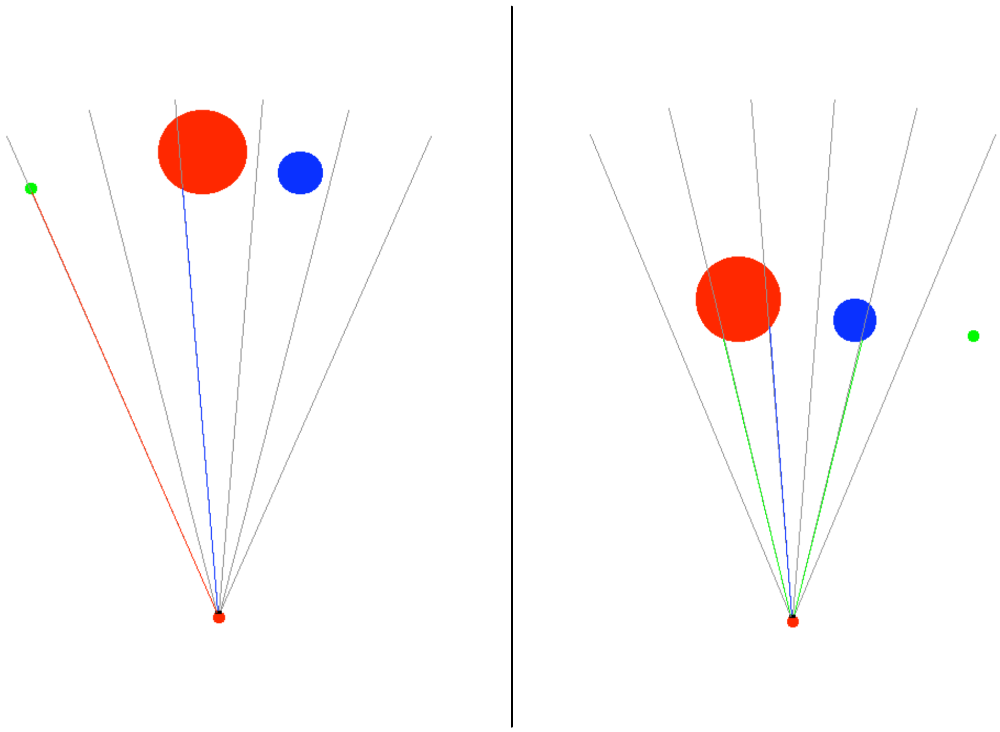

Figure 18: Plot of fitness against generation number

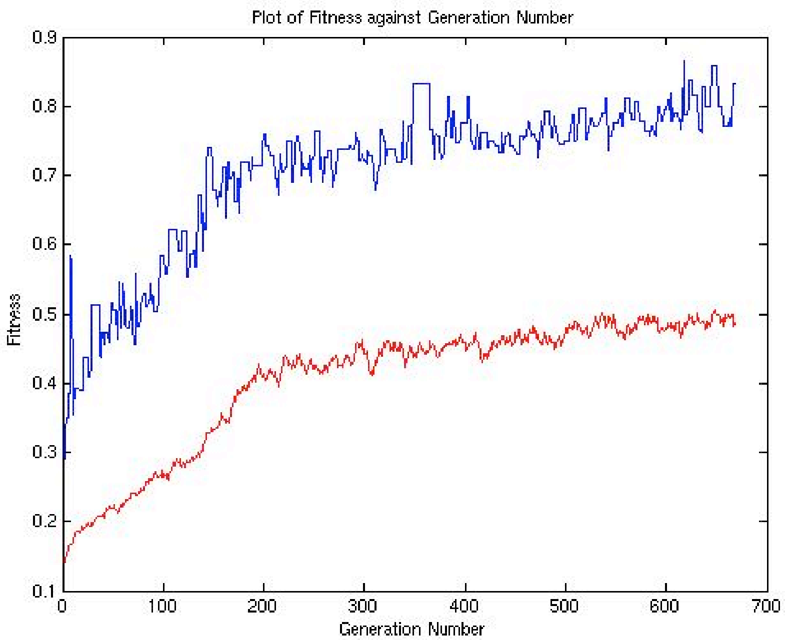

6. Discussion and Extensions

6.1 Discussion

The first observation that can be made is that although being endowed with very simple control and visual systems, in some experiments agents of high fitness were successfully evolved. One would expect a task such as size discrimination (especially absolute discrimination) to require visual and control systems of high precision and complexity, however in some cases the results suggest that the control systems used might well have been overly complex for the task. The agent in 5.3 (selecting the bigger of two stimuli) evolved neuron time constants very close to the value of one, this together with the very basic strategy employed by the agent (moving to the largest stimuli) suggests that the same high fitness results could have been achieved with a much simpler feed forward neural network, perhaps even with out a hidden layer (i.e. with direct connections from the sensors to the motors). Other agents did however involve more complex hidden neuron timings and connections although a general understanding could be gained.

Therefore the results show that in each case evolution has exploited computationally cheap techniques in order to solve the different problems. Unfortunately a 'Non-absolute size discriminator' could not be evolved (5.6) and mixed results were obtained when evolving a 'middle size discriminator' (5.7). However it was possible to produce agents capable of performing relational discriminations and there was also a certain degree of success when evolving an agent able to select a circle of absolute size. Surprisingly although evolved to perform different behaviours, the strategies employed by each apparently distinct agent are remarkably similar. A clear example of this can be seen in section 5.5 (Agents capable of selecting a circle of a an absolute size) where the two agents evolved different neural architectures (neural connections, time constants and biases) but exploited exactly the same environmental pattern (strategy) in order to solve the same specific problem. This in it self is unremarkable, as one would expect the evolutionary algorithm (being an optimisation process) to select the easiest strategy; in addition the neural architecture/control system is highly configurable, thus there are likely to many combinations of parameter values that produce similar behaviour. More interestingly however is that the strategy used by both agents is heavily based on, and almost identical to that used by the agent in section 5.4, where the agent was evolved to pick the smaller of two circles (relational discrimination) (although not completely successful section 5.7 also reported near identical behaviour emphasising the similarities between the two tasks (see 5.7)). This seems to suggest that absolute discrimination is strongly related to, if not fundamentally based on relational discrimination. Before committing to any wider claims about the nature of size discrimination it is important to understand why this behaviour has arisen and thus view the strategy in context.

Predictably, the strategies employed by each agent are tightly linked with the specific evaluation function used in the evolutionary algorithm and therefore have a strong connection to the environment in which the agent has developed. From the experiments involving the small size discriminator we know that the agent relied on the larger circle intersecting its outer ray first (larger diameter), thus pushing the sway of attraction towards the smaller circle. Within the evolutionary trials the circle diameters were kept constant, the large circle always intersecting the outer rays first due its larger diameter. This is clearly a useful and predictable environmental feature/invariant and is in this case exploited by evolution. Because this is such a common feature in the agents environment the strategy is robust, therefore an emergent feature of the agents' dynamics' is that it can reliably choose the smaller circle out of two, regardless of both of their sizes, as long as one is bigger than the other, the bigger will always intersect the outer ray first.

With this in mind, we see a similar situation with the 'absolute' discriminating agent. One would imagine that an agent capable of choosing a circle of a specific size might use some form of 'neural imprinting' where the size of a circle is some how stored within the neural network, the agent perhaps succeeding through some form of active pattern matching. However in this case the agent succeeds at an absolute size task through using much simpler relational techniques as described above. Once again there is a tight relationship between the strategy employed and the evolutionary process/training. The evolutionary task was to ignore circles that are bigger and smaller (range 1 – 40) than the circle of set size (radius 20) and therefore select the constant sized circle. In this case then the circle of specific size is exactly in the middle/centre of the random range presented to the agent and evolution has exploited the same relationship as with the small size discriminator, where if the circle presented to the agent were larger than that of the specific sized circle then it would intersect an outer ray first. In fifty percent of trials this strategy is likely to work, however to achieve a higher score a further hidden environmental invariant is exploited. When circle 2 is smaller than circle 1 (radius 20) and with both the circles being placed in favourable positions, the small circle will intersect the sensor rays less often, and in many cases falls 'in between' the sensor rays due to its small diameter. In such a case the agent moves towards circle 1 initially (circle 1 being the only circle it sees), as it does this, the outer ray is likely to be intersected by circle 2 thus adding the final encouragement to move towards circle 1. Low scores only occur when the circles are of a similar size or are placed in unfavourable positions (i.e. the smaller circle intersecting an middle (attractive) sensory ray initially). Section 5.7 also reports similar behaviour, evolution has successfully applied the same solution to multiple problems.

The results emphasise many of the concepts involved in evolutionary robotics (situatedness and embodiment) and it is obvious then, that the agent's behaviour or 'cognitive ability' is grounded in its environment. The agent can be seen as exploiting an ecological niche, its control system set up to respond to predictable, common environmental stimuli, and in doing so fulfil its evolutionary roll. The agent itself does not actively make a decision or seek to solve the particular problem but works more like a reactive system, where the environment controls its decisions. The environment and sensory systems are simple and in many ways so are the resulting control systems (though one should not get the impression that control system is overly basic, neural interaction can become extremely complex!). Evolution has thus solved each problem using relatively simple, but none the less interesting techniques (perhaps the complexity of the environment dictates the complexity of the control system and therefore strategy used). It is more appropriate then to view the agent as a form of 'signal processor', where input stimuli from different regions of its visual space are processed internally according to various inhibitory and excitatory connections. Typically the processed inputs from both directions in space tend to be summed together, the larger and more persistent inputs having more of an effect and thus dictate which motor will activate (according to excitatory and inhibitory connections) and therefore the motion of the agent (one could imagine and it is plausible that similar signal processing could be involved at a 'low' 'neuron level' in biological systems). When thought through this is a perfectly reasonable method given the visual acuity of the agent and may well provide an explanation as to why agents capable of absolute discrimination were difficult to evolve and therefore the reasons why evolution resulted to using relational methods to 'fulfil' the fitness function.

Relational discrimination can therefore occur at a very simple/basic level, requiring extremely small amounts of computing, one would image however that absolute discrimination would require a much more computationally intensive process. For example perhaps a memory would need to be formed to remember a particular circle size (some kind of distributed neural representation). To form this memory the agent would almost certainly require a higher visual resolution, as to fully appreciate the diameter of a particular circle (although conceivably this could be achieved with just two sensors) and consequently would most probably involve a more substantial control system. Perhaps the agent would be more mobile and perform some kind of active scanning or active pattern matching procedure, where it would attempt to position (and thus match) a circle in a particular region of its visual field. This is certainly an appealing topic to explore in the future and it would be interesting to see if the agent resorted to similar techniques as seen in nature. For example insects appear to identify local land marks using a process known as ' image matching', in which the insect moves as to maximise the fit between the stored view of the landmark (perhaps the circle diameter in this case) and the current retinal image (see [10] and [32] for examples of image matching in ants and wasps accordingly).

In any case it is clear that absolute size discrimination most probably requires an agent of slightly greater complexity. In these experiments the environment and sensory and/or control system endowment of the agent favoured the development of relational discrimination techniques. If the main goal of the project were to evolve an agent purely to select a circle of one specific size it would also be likely that the environment and even the fitness function might need to be improved. However this was not the purpose of the project; it was however important that there was an equal chance for both methods to arise naturally, with as few prior assumptions and 'specific selections' made as possible. In any case the employed fitness function is very general (only scoring the agent on moving towards the correct circle) and the environment un-biased.

The results of section 5.5 (the absolute size discriminator) present a problem; clearly the agent is using a relational strategy, yet at a higher level it is performing the correct behaviour - reliably choosing the circle of a specific size. Whether this suggests that relational discrimination is the easiest, more robust and therefore the most fundamental discrimination ability, on which absolute discrimination is built, is not entirely clear. What can be concluded however is that even simple control systems can be difficult and time consuming to analyse, yet an understanding albeit at a relatively 'high' level can be gained through a combination of procedural testing and behaviour correlation. More importantly the project has shown that agents capable of relational discriminations are much easier to evolve, and suggests that given the chance evolution will attempt to use relational cues rather than resorting to absolute measures of discrimination.

6.2 Extensions and improvements

There are a large number of questions that need to be answered and therefore a large amount of experimental work needing to be conducted before any firm conclusions can be drawn about the nature of size discrimination. Yet the results achieved in this experiment and in previous studies provide a starting point on which to base further investigations. A variety of suggestions for future extensions that I feel are worthwhile pursuing are listed below:

1) The sensory system used within this project was kept basic to aid analysis; a consequence of which was the agent had poor visual acuity. It might be that this was the limiting factor in the investigation, preventing the evolution of agents capable of judging absolute size. It would be useful to repeat the experiments described above using an agent that has a much higher visual resolution (adding more rays could be easily implemented). Perhaps a sharper visual system would allow the agent to fully appreciate the size of circles; it would be extremely interesting to observe the environmental features and therefore strategies exploited by evolution in this case (perhaps the agent would use image matching?). It is likely that more sensors would need at least an equal number of input neurons, therefore increasing the number of parameters that need to be evolved.

2) The simulation contained no noise; again this was left out in order to make the analysis of the experimental data easier. Pilot tests with noise suggest that the fit agents are reasonably robust, able to handle between 1 and 10% noise on its sensors inputs and motor outputs with ease. However a formal investigation is required to observe the effect of noise on the different strategies evolved. Too much noise however and the tractability of analysis might be disrupted. It might also be interesting to model inter neural noise (noise within the neural network). Therefore two investigations can be conducted, one testing the evolved agents robustness to noise and the other evolving fresh agents to perform the same tasks but including noise within the simulation and evolution. This might result in evolution finding alternative, more robust solutions to the specific problems.

3) Only one type of control system was implemented in this investigation. It would be of use to compare alternative control systems, both simpler feed forward networks and more complex networks capable of self-organisation (for example the use of Hebbian learning rules to change the strength of neural connections at each time step/slice). A more interesting extension would be to enable evolution to adjust the architecture of the agent (i.e. number of sensors, sensor ranges, number of internal neurons, layers), thus enabling a more natural construction of the agent (making even fewer prior assumptions, this could however make analysis extremely complex!)

4) Psychological evidence suggests absolute sizes require much more time to learn requiring multiple presentations of the same stimuli. Perhaps an experiment could be conducted to test this. Theoretically, the evolutionary process should present the specific sized circle to the agent many times, however perhaps more trials are needed. A useful question to answer would be whether the agent resorts to relational discrimination in novel circumstances and then gradually over time learns the absolute size of a circle?

5) Less dramatic, a further extension might be to test the reliance of agent on the vertical movement of the circle. Perhaps the agent could be evolved to perform the same behaviour if the circles are kept in static positions. Would the agent rely on a more active strategy or maybe just resort to static pattern matching?

6) It would be of interest to see if similar behaviours are observed when evolving agents to discriminate between other shapes.

7) Even when noise is incorporated into the experiments the simulations are still abstract, involving huge simplifications. For realism then it would be useful to evolve the agents in hardware. This would mean that the agents are in contact with the real world and it is therefore less likely that unrealistic strategies will be evolved. Although useful, implementing artificial evolution in hardware can be an expensive and time-consuming process. To save time a robust control system might be evolved using a 'minimal simulation' [9] noise technique, and then transferred to a physically robot for experimental trials.

8) Variations of the evaluation function might be implemented to make the evaluation procedure harder. Modifications such as allowing the agent less time to make a decision, presenting more trials to the agent and adding larger amounts of variation in terms of agent and circle positioning and/or circle size might enable more successful behaviours to be evolved. Perhaps the agent could be presented with one circle at a time within a trial and then have to make a decision at the end of the trial when no circles are present (i.e. by moving left or right, according to where the correct circle was previously positioned). Such an investigation may also provide insights into short-term memory.

9) Finally, It might be interesting to explore absolute discrimination by intensively training/evolving the agent on one circle and then presenting it to an environment with two circles (rather than training it with both circles). Perhaps incremental evolution could be used, the agent could be initially evolved to orientate to the single circle. The evolved agent (genotype) would then be re-exposed to evolution; this time the agent would be subjected individually both to the same circle and to circles whose size is dictated randomly (the fitness judged on its approach to the correct circle and avoidance of random circles). If successful evolution could then continue, this time perhaps using a new evaluation function where both the circles are shown to the agent at the same time (scored for selecting the correct circle of specific size). Such a procedure may result in a much more robust absolute size discriminator agent.

There are potentially a huge number of extensions that could be explored (separately or in combination), however with any extension undertaken it is important to bear in mind the concepts behind minimally cognitive models. Any extensions should seek to build on previous results in a step-by-step fashion perhaps including a greater degree of complexity but retaining the key features and benefits of idealised models.

7. Conclusion

The project has contributed to the study of minimally cognitive behaviour by supporting and extending related work. Previous studies both in evolutionary robotics and psychology have provided evidence that relational discrimination is the ‘cognitive primitive’ naturally employed when discriminating between two similar stimuli and the results of this project firmly support this. Each experiment has not only raised several interesting questions about the nature of size discrimination and the necessary equipment needed to perform it but has also provided a clear demonstration of the fundamental principles involved in evolutionary robotics research, emphasising the strong link between an evolved organisms/agents environment and cognition. In addition the project has provided an insight into the problems and benefits of the minimally cognitive framework, showing that although unable to directly apply the results to biological systems the results do however provide evidence capable of informing and supporting biological investigations.

(See Appendix B: Project log, for personal conclusions).

8. References + Bibliography

1) Barton, L. and Bahbouh R. (2006). Antler size and fluctuating asymmetry in red deer (Cervus elaphus) stags and probability of becoming a harem holder in rut. Biological Journal of the Linnean Society: Volume 87 Page 59.

2) Bateman P. W., Gilson, L. N. and Fergus, J. W. (2000). Male size and sequential mate preference in the cricket Gryllus bimaculatus, Department of Zoology and Entomology. University of Pretoria.

3) Beer, R.D.: (1996) Toward the evolution of dynamical neural networks for minimally cognitive behavior. In Proc. of Simulation of Adaptive Behavior. MIT Press.

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

5) Blanchard, Rind and Verschure. (2000). “Collision avoidance using a model of the locust LGMD neuron”, Robotics and Autonomous Systems.30, 17-38.

6) Brooks, R. A., (April 1991)."Intelligence Without Reason". MIT AI Lab Memo No. 1293.

7) Clark, A., Thornton, C. (1997). Trading spaces: computation, representation and the limits of uninformed learning. Behavioral and Brain Sciences 20, 57–90.

8) Chiel, H.J., Beer, R.D. (1997). The brain has a body: Adaptive behavior emerges from interactions of nervous system, body and environment. Trends in Neurosciences 20, 553–557.

9) Gonzalez, R. C., Gentry G. V., and Bitterman, M. E. (1954). Relational discrimination of intermediate size in the chimpanzee. J Comp Physiol Psychol: 47(5): 385–388.

10) Graham, P. Fauria, K. and Collett, T.S. (2003). The influence of beacon-aiming on the routes of wood ants. J. Exp. Biol. 206, 535-541.

11) Harvey, I. (1996). The microbial genetic algorithm. Unpublished manuscript: http://www.cogs.susx.ac.uk/users/inmanh/Microbial.pdf

12) Harvey, I., Husbands, P., Cliff, D., Thompson, A. & Jakobi, N. (1997). Evolutionary Robotics: the Sussex Approach In Robotics and Autonomous Systems, v. 20 pp. 205--224.

13) Harvey, I., Vaughan, E., and Di Paolo, E. (2004). Time and Motion Studies: The Dynamics of Cognition, Computation and Humanoid Walking. Presented at HART 2004, Fourth Intl. Symp. On Human and Artificial Intelligence Systems: From Control to Autonomy. December 5-6, 2004, University of Fukui, Fukui, Japan.

14) Harvey, I., Di Paolo, E., Wood, R., Quinn, M. and Tuci, E. A. (2005). Evolutionary Robotics: A new scientific tool for studying cognition. Artificial Life, 11(1-2): 79-98.

15) Hawkins, J. (1986). An Investigation of Adaptive Behavior: Towards a theory of Neocortical Function.

16) Hawkins, J. (2004). On intelligence, Times books.

17) ‘IBM blue brain project’. http://domino.watson.ibm.com/comm/pr.nsf/pages/rsc.bluegene_cognitive.html

18) Izquierdo-Torres, E. and Di Paolo, E. (2005) Is an Embodied System Ever Purely Reactive? In M. Capcarrere et al (Eds.) Proceedings of the 8th European Conference on Artificial Life. Pages 252-261. Springer-Verlag. LNAI 3630.

19) Jakobi, N. (1998). The minimal simulation approach to evolutionary robotics. In Gomi, T.m editor, Evolutionary Robotics – From intelligent Robots to Artificial Life (ER’98). AAI Books, Ontario, Canada.

20) Kendler, T. S. (1950). An experimental investigation of transposition as a function of the difference between training and test stimuli. J Exp Psychol. 40(5): 552–562.

21) Kohler, W. (1918). Simple structural functions in the chimpanzee’s and in the chicken’s representation. In Ellis, W.D., Editor, A sourcebook of Gestalt Psychology, pages 217-227. London: Routledge.

22) Law, P. (2002). A Minimally Cognitive Behavior Model of the Discrimination of Simple Size Relation, M.Sc. Thesis, 2002.

23) Lawrenson, W. and Bryant, P. (1972). Absolute and relative codes in young children. Journal of child psychology and psychiatry, 13:25-35.

24) Lazareva, F. O., Wasserman, E. A. & Young, M. E. (2005). Transposition in pigeons: Reassessing Spence (1937) with multiple discrimination training. Learning & Behavior, Volume 33, Number 1, February, pp. 22-46(25) Psychonomic Society Publications.

25) Macinnes, I. and Di Paolo, E. (2005). From the Inside Looking Out: Self Extinguishing Perceptual Cues and the Constructed Worlds of Animats. In Capcarrere, M., Freitas, A.A., Bentley, P.J., Johnson, C.G. and Timmis, J., editors, Advances in Artificial Life: 8th European Conference, ECAL 2005, Canterbury, UK, September 5-9, 2005, Proceedings, pages 11-20. Springer.

26) Marchetti, K. (2000). Egg rejection in a passerine bird: size does matter. Animal Behaviour 59, 877-883.

27) McGonigle, B. O. & Jones, B. T. (1978). Levels of stimulus processing by the squirrel monkey: relative and absolute judgments compared. Perception, 7:635-659.

28) Nolfi,S. and Floreano,D. (2000). Evolutionary Robotics: The Biology, Intelligence, and Technology of Self-Organizing Machines.

29) Nolfi, S. (2002). Power and the Limits of Reactive Agents. Neurocomputing 49, 119–145.

30) Persons, M.H. & Lietz, G. W. (2005). Sexual cannibalism and mate choice decisions in wolf spiders: influence of male size and secondary sexual characters. Animal Behaviour. 69:83-94.

31) Slocum, A.C., Downey, D.C. and Beer, R.D. (2000). Further experiments in the evolution of minimally cognitive behavior: From perceiving affordances to selective attention. In J. Meyer, A. Berthoz, D. Floreano, H. Roitblat and S. Wilson (Eds.), From Animals to Animats 6: Proceedings of the Sixth International Conference on Simulation of Adaptive Behavior (pp. 430-439). MIT Press

32) Tinbergen, N. (1932). Über die Orientierung des Bienenwolfes (Philanthus triangulumFabr.). Z. Vergl. Physiol. 16, 305–334.

33) Vaughan, E. (2003). Evolution of 3-Dimensional Bipedal Walking with Hips and Ankles. MSc thesis, Dept. of Informatics, Sussex University.

34) Vaughan, E. (2003). Bilaterally Symmetric Segmented Neural Networks for Multi-Jointed Arm Articulation. Adaptive Systems, Dept. of Informatics, Sussex University.

35) Vaughan, E. (2003). Flight of the Biped. Simulation of Adaptive Behavior, Dept. of Informatics, Sussex University.

36) Vaughan, E., Di Paolo, E. A. and Harvey, I. (2004). The Evolution of Control and Adaptation in a 3D Powered Passive Dynamic Walker. Proceedings of the Ninth International Conference on the Simulation and Synthesis of Living Systems, ALIFE'9 Boston, September 12th-15th, 2004. MIT Press.

37) Vaughan, E., Di Paolo, E. A. and Harvey, I. (2004). The Tango of a Load Balancing Biped. Proceedings of the Seventh International Conference on Climbing and Walking Robots, CLAWAR Madrid, September 22th-24th, 2004.

38) Webb, B. (2000). What does robotics offer animal behaviour?. Animal Behaviour, 60, 545-558.

39) Webb, B. (2001). Can robots make good models of biological behaviour? Target article for Behavioural and Brain Sciences 24 (6).

40) Wills S., Mackintosh N. J. (1999). Relational Learning in Pigeons? The Quarterly Journal of Experimental Psychology B, Volume 52, Number 1, pp. 31-52(22) Psychology Press, part of the Taylor & Francis Group.

41) Wright, A., Cook, R. G., Kendrick, D. F. (1989). Journal of the experimental analysis of behavior relational and absolute stimulus learning by monkeys in a memory task 529, 237 - 248.

42) Zufferey, J.C. and Floreano, D. (2005) Toward 30-gram Autonomous Indoor Aircraft: Vision-based Obstacle Avoidance and Altitude Control. IEEE International Conference on Robotics and Automation (ICRA'2005).

Appendix A

Experiment data and analysis (to be used with the descriptions in section 5)

Appendix A:

Experiment One A details:

1.1 Small relative difference in radius size:

Plot of the horizontal movement of agent over time: (vertical axis representing horizontal position with larger values corresponding to the right of the agent (top of the y axis) and smaller values corresponding to the left (bottom of the y axis).

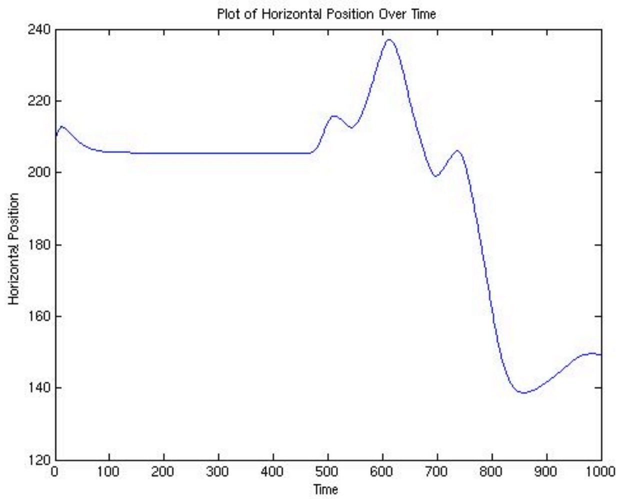

Plot of the sensory inputs over time:

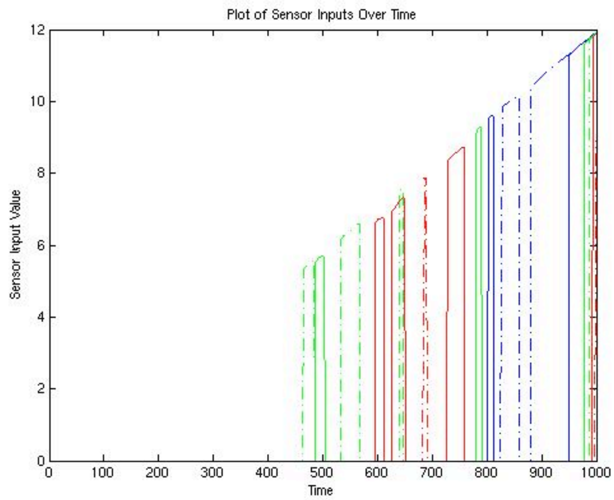

The above correlates to the recorded horizontal movement of the agent, the solid coloured line representing sensor inputs from the left, the dotted dashed lines representing sensor inputs from the right. There are no sensory inputs until just under half way into a trial, after which the both the left and right sensor values are very similar. On this trial (52) the larger circle was positioned to the right, in this case the agent achieved a low score of 0.01. One can clearly see the agent move in the direction of the bigger input stimuli at each time step (see sensory input graph above) and ultimately moves to the left. The agent appears to be working more as a feed forward neural network, this idea being supported by the time constants of most of the neurons being assigned to values very close to one.

1.2 Medium relative difference in radius size:

Plot of the horizontal movement of the agent over time:

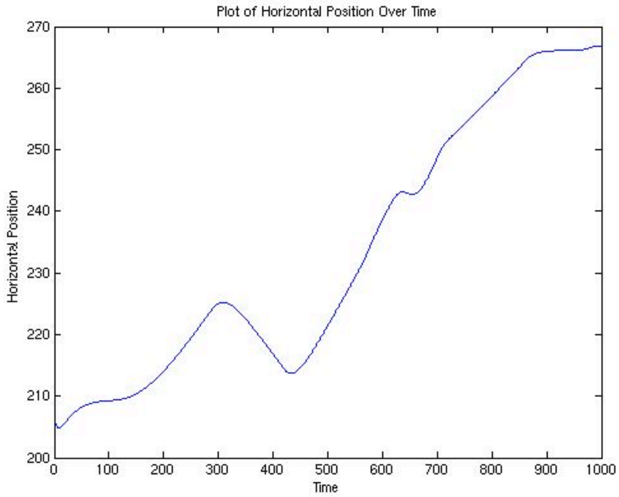

The above graph shows the horizontal movement of the agent in trial 17, the larger circle is to the right of the agent, in this case the agent achieves a high score of 1. The above graph displays the observed ‘active scanning’ although as the text in section 5.3 describes this is not actually the case.

As referred to in the text the input activations and firings mimic the sensory inputs:

Plot of the sensory inputs over time:

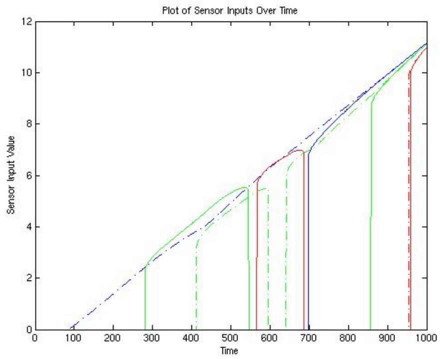

Plot of the input neuron activations over time:

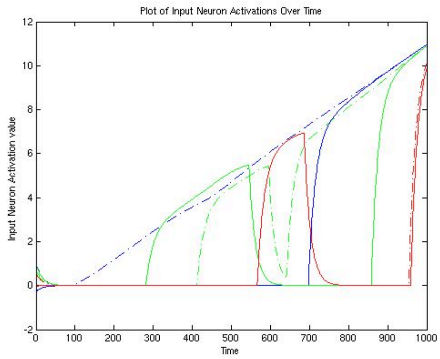

Likewise the hidden layer activations and output layer activations oscillate accordingly, the biases and weights appearing to ‘filter out’ and amplify the more important stimuli.

Plot of the hidden neuron activations over time:

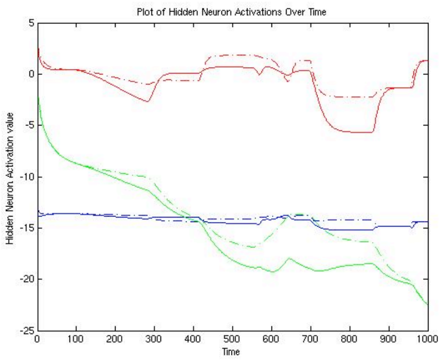

Plot of the output neuron activations over time:

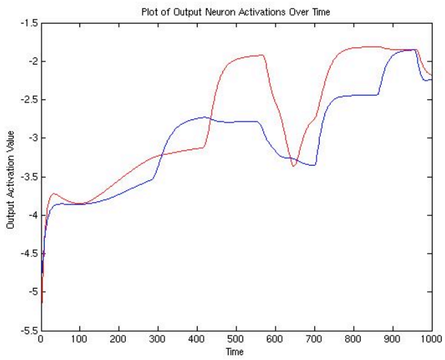

1.3 Large relative difference in radius size:

Plot of the horizontal movement of the agent over time:

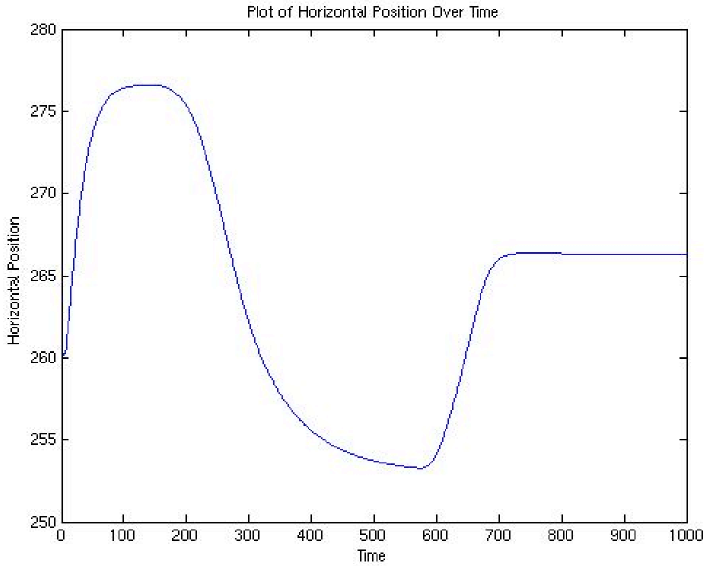

In this case both circles overlap, the agent moves far less and much more slowly, performing one large oscillation over 700 time steps and then finally settling in the centre of both stimuli. The agent in this case scored low around 0.12%.

The sensory input data shows strong consistent inputs from most of the sensors over the course of the trial.

Plot of the sensory inputs over time:

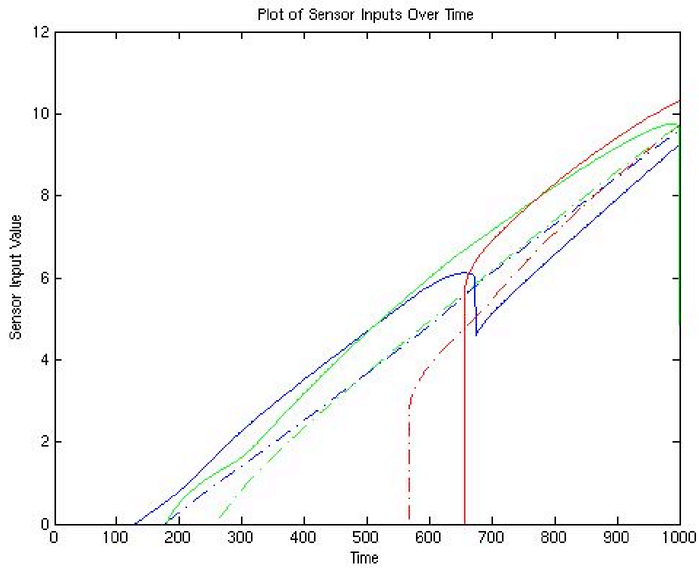

The hidden neuron activations become smoother (more constant) with the strong sensory inputs as inputs from both the left and right sensors are similar and consequently have the effect of cancelling each other out, this is clearly visible from time step 700 onwards.

Plot of the hidden neuron activations/states over time:

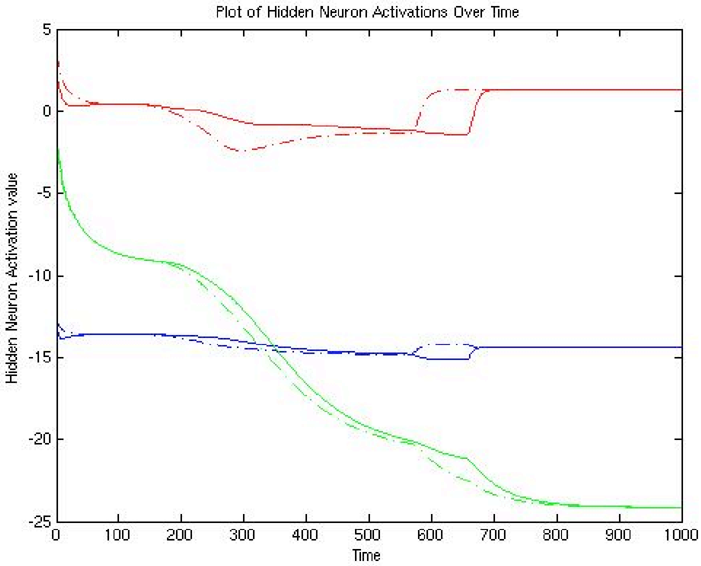

1.4 A further example: when both circles are positioned to the right of the agent:

Left input sensors are initially stimulated (as no circles are positioned on the right); a similar pattern of activity then commences as seen above (the agent eventually moving right towards the largest stimuli).

Plot of the sensory inputs over time:

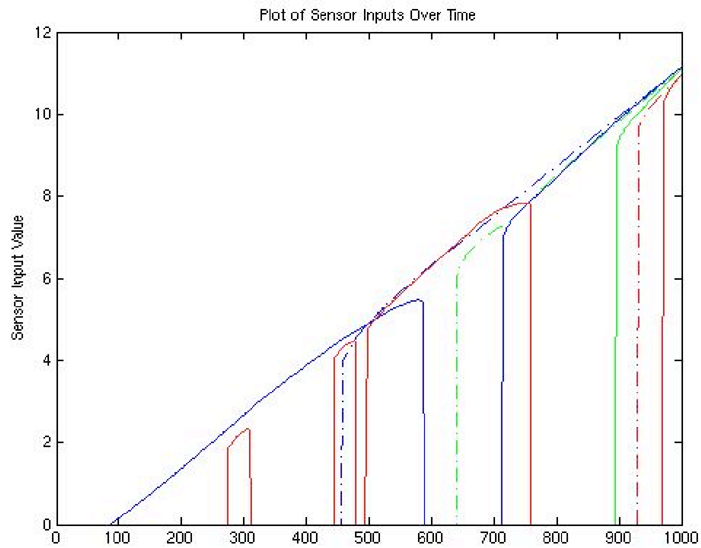

Plot of the horizontal movement of the agent over time:

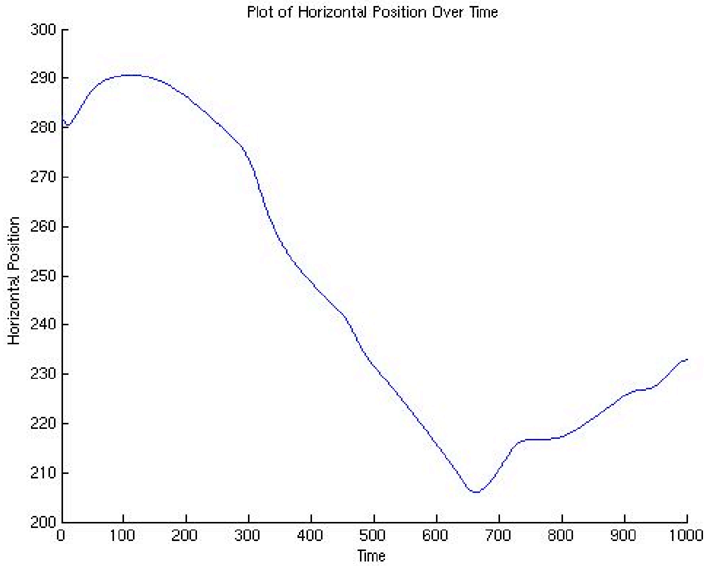

Experiment One B details:

Plot of the input neuron firing over time:

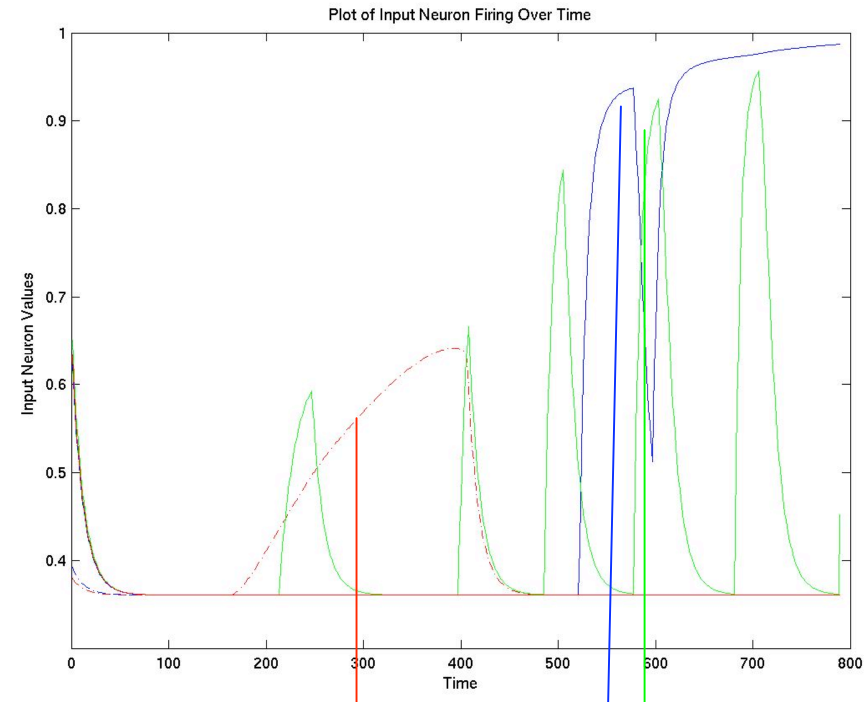

Plot of the horizontal agent movement over time:

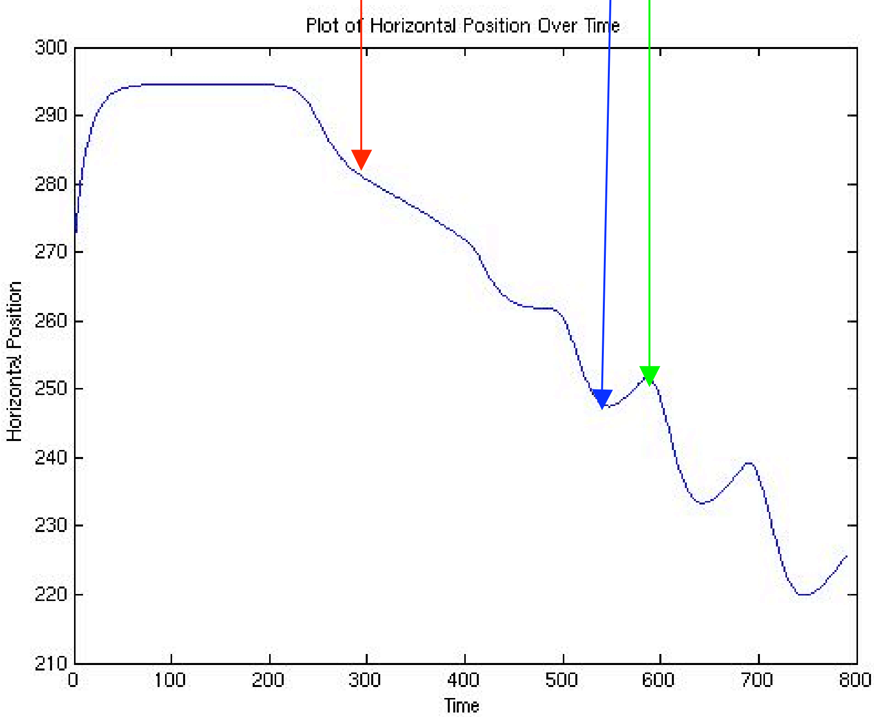

In this example a small circle was 'dropped' to the left of the agent and a larger circle to the far right. The effect of different sensor intersections can be clearly seen, with the left middle sensor (5) having a strong attractive effect on the agent (solid green line, the resulting behaviour correlated by green vertical arrow above) and the outer right sensor having a strong repulsive effect on the agent, as soon as the larger circle touches the right outer sensory ray

(1) (slash dotted red line) the agent is committed to moving left (exemplified by red vertical arrow above). The inner left sensor (4) (solid blue line) also clearly shows its repulsive effect, inhibiting the 'attractive force' of sensor (5) (exemplified by blue vertical arrow above). The agent managing to centre itself near the end of the trial when the 'forces' are equal.

Plot of the hidden neuron firings over time:

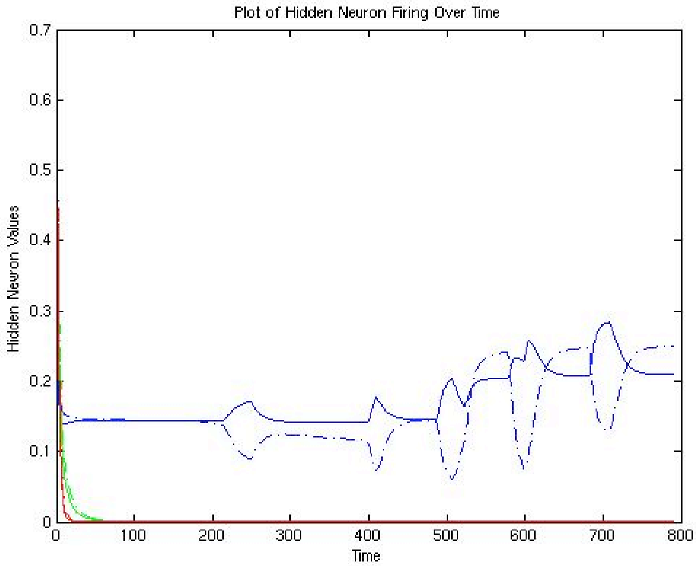

The firings of the neurons dictate the oscillating of the output/motor neurons and thus determine the movement of the agent. The firings seem to correspond to the sum of the attractive and repulsive forces (Which ever is bigger at a particular time step).

Plot of the hidden neuron activations over time:

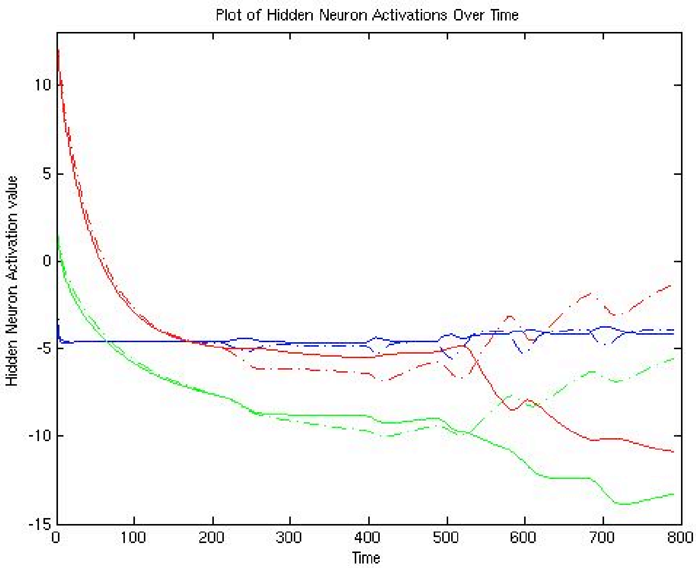

The Hidden Neuron activations are considerably more complex than in experiment 1A, neurons oscillate out of sync with each other, Experiment Two A details:

Example 1) Example of Agent 1 selecting circle 1 (radius= 20) when also presented with a smaller stimuli (radius = 3). – Example taken from Test 3 data.

Plot of the input neuron firings over the trial:

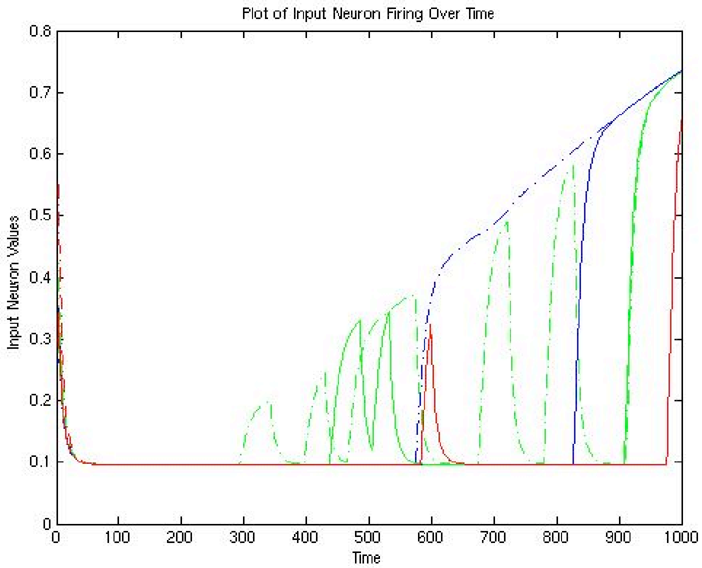

Plot of the horizontal position of the agent over the trial:

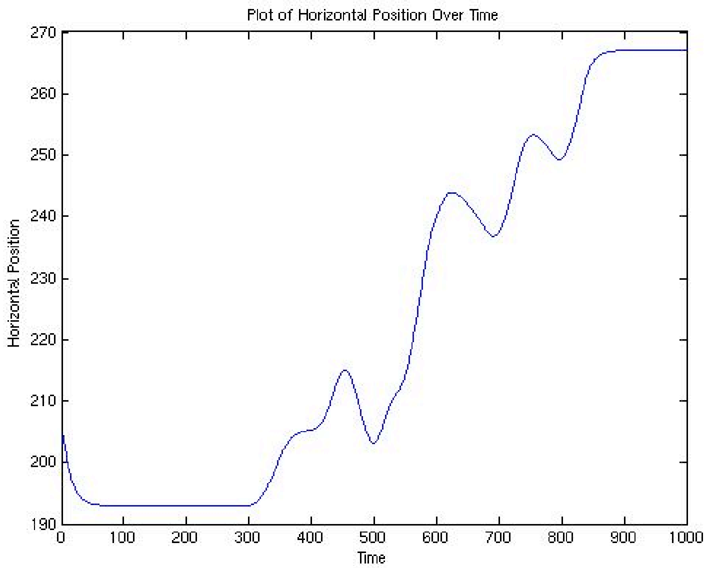

As with the small size discriminator the agents control system processes the current sensory inputs into repulsive and attractive forces/movements. At the 300th time step the right middle ray sensor (green dash-dot line, sensor 2 on diagram 5) is intersected by circle one, the agent responds by moving right (an attractive response). At approximately the 450th time step however the left middle sensor is briefly intersected (solid green line, sensor 5 on diagram 5), this time attracting the agent left. Therefore the two 'attractor rays' are pulling the agent in both directions, the strongest input winning. Circle 2 (on the left hand side/bottom of screen) intersects the outer left ray (solid red line, sensor 6 on diagram 5) of the agent first (approx the 600th time step), and despite receiving an inhibitory response from the right inner sensor, the agent continues moving right slowing down and gradually moving left as the inhibitory response grows in strength (blue dash-dot line, sensor 3 on diagram 5). As it begins to move away the right middle sensor (sensor 2) ray is intersected (approx 700th time step), thus the agent is strongly attracted back to the right. This interaction continues to the end of the trial, the important 'threshold point' is when the outer left ray is intersected at around the 600th trial. Although the hidden activations do not show a significant change, the only real change being the activations of the outer hidden neurons (neurons 1 and 6 on diagram 5) becoming increasingly negative at around the same time (Highlighted by the red circle below), however the full effect of this is still unclear.

Plot of the hidden neuron activations over the trial:

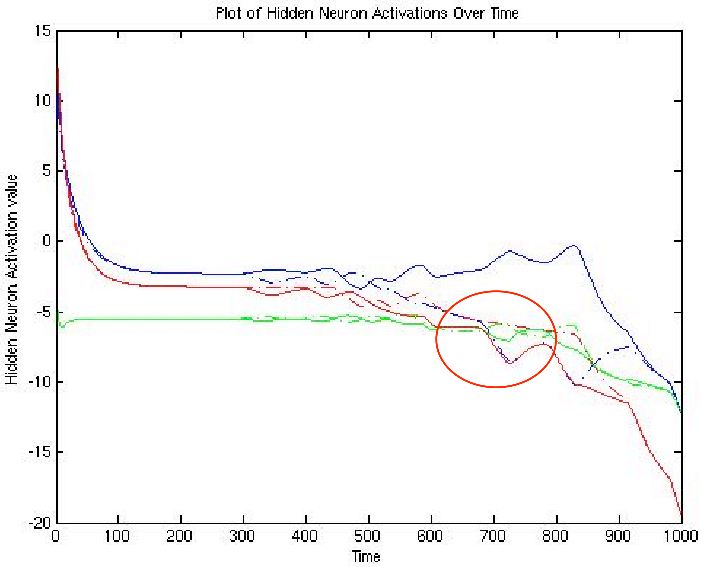

Interestingly only the middle hidden neurons actually produce an output/fire, the output neurons mimicking more or less the same pattern.

Plot of the hidden neuron firings over the trial.

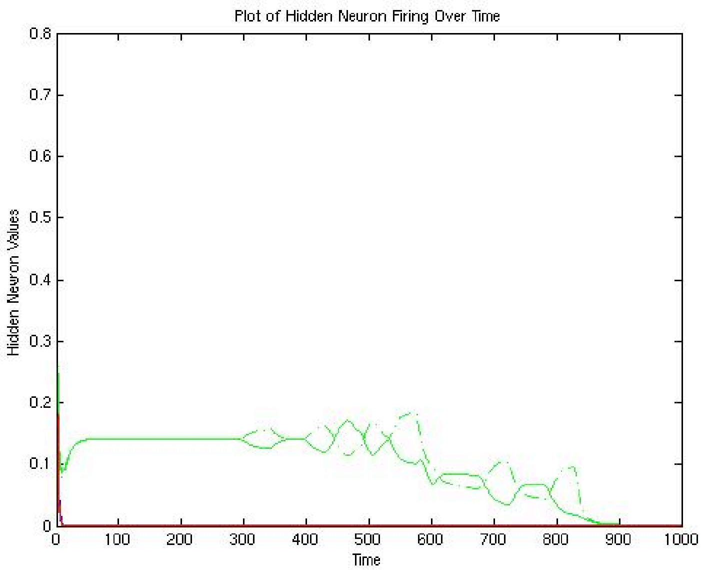

Example 2) Example of Agent 2 selecting circle 1 (radius= 20) when also presented with a larger stimuli (radius = 33).

– Example taken from Test 3 data.

Plot of the input neuron firings over the trial:

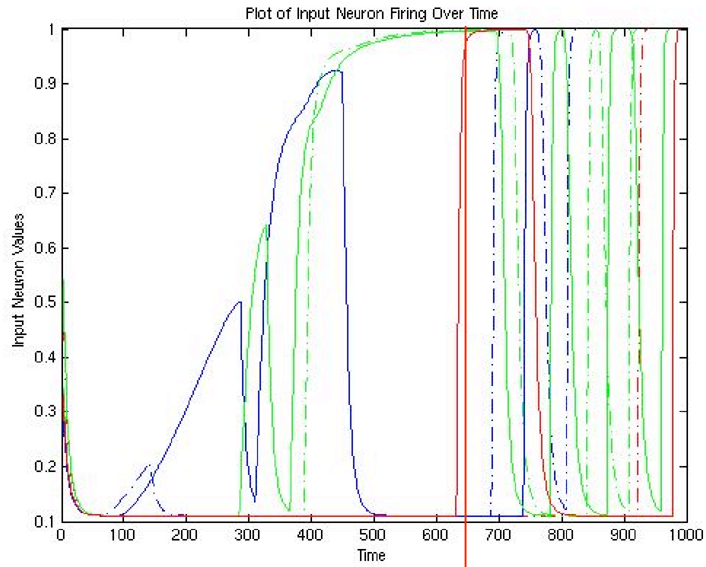

Plot of the horizontal position of the agent over the trial:

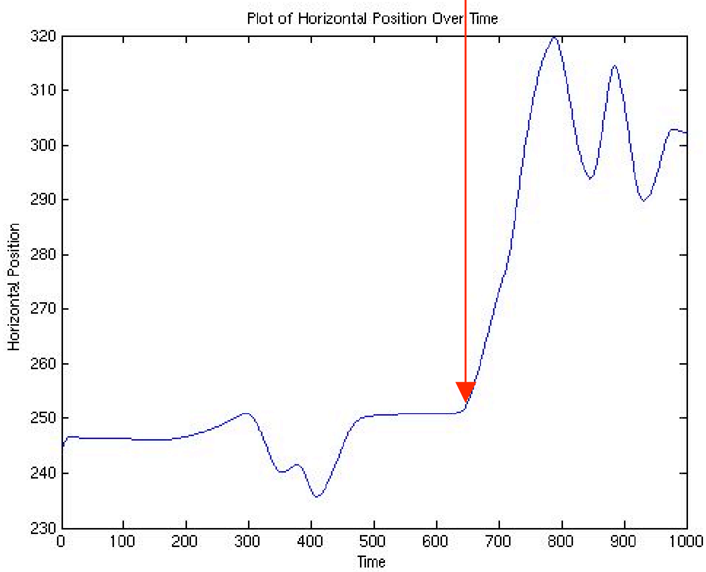

As with agent 1 (above) this agent uses the intersection of an outer ray to determine its final choice. At the 600 – 650th time step (red arrow) the outer left ray (solid red line on graph) is intersected (sensor 6 on diagram 5) by circle 2 (the larger of the two circles). This corresponds tightly with the horizontal position of the agent as the above graphs show. The effect is pronounced, in this case where both circles are clearly visible to the agent (both being reasonably large). From the 500th to 600th time step the agent is stationary with equal inputs from both middle sensors (sensors 3 and 5 on diagram 5), however just after the 600th time step the intersection of the left outer sensory ray 'tips the balance' of attractive and repulsive forces towards the right direction.

Plot of the hidden neuron activations over the trial:

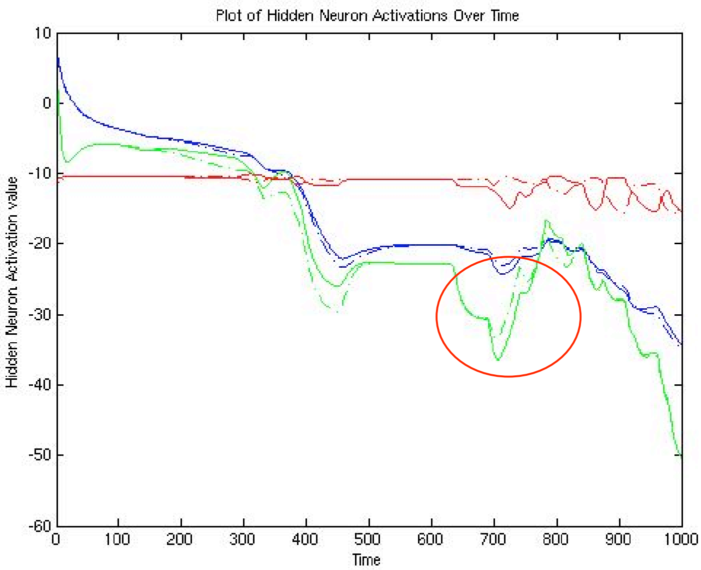

Again as in agent two there appears to be little significant change in the activation values, the activations of the middle hidden neurons becoming more negative (highlighted above) perhaps inhibiting the effect of the middle input firings. In this agent the outer hidden (neurons 1 and 6) neuron-firing pattern directly dictates the neural activity in the output neurons (The firings seeming to correspond to the sum of attractive and repulsive forces or maybe directly related to the sum of the hidden layers activation oscillations. More research is needed!).

Plot of the hidden neuron firings over the trial:

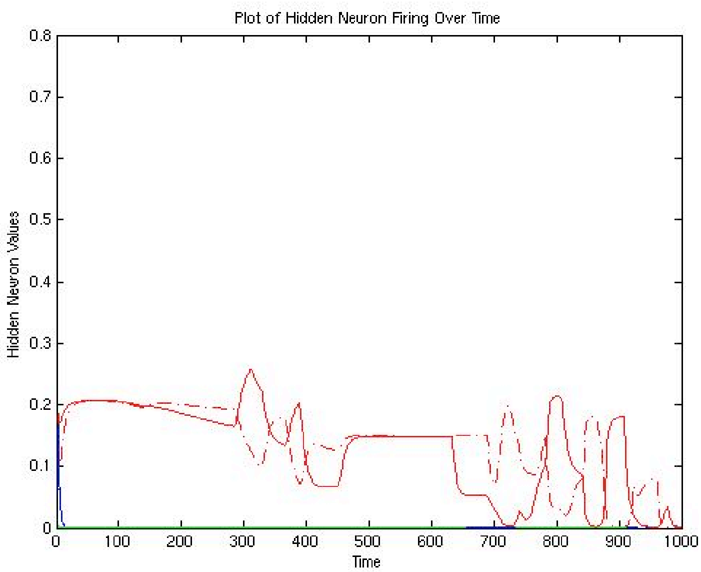

Experiment 3 Details:

In this example the middle circle was positioned on the right of the agent, the large circle central/above the agent and the smallest to the left of the agent. In such a situation the agent has about a 50:50 chance of choosing the correct circle (In this case the positioning was favourable). The decision to select a particular circle is less pronounced than in experiment 2A, in this case the agent is much more 'active' as a result of the extra stimuli. Here the smaller circle intersects the left outer ray on around the 200th time step encouraging the agent to move right (small red circle below). A further intersection between the 600th – 800th time steps by the larger circle (large red circle below highlights the point) encourages the agent to continue its path.

Plot of sensory inputs of the trial:

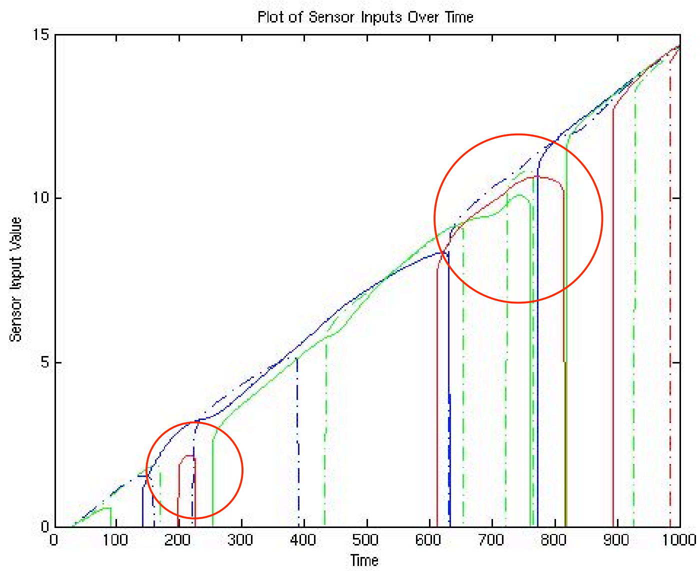

Plot of horizontal position over the trial:

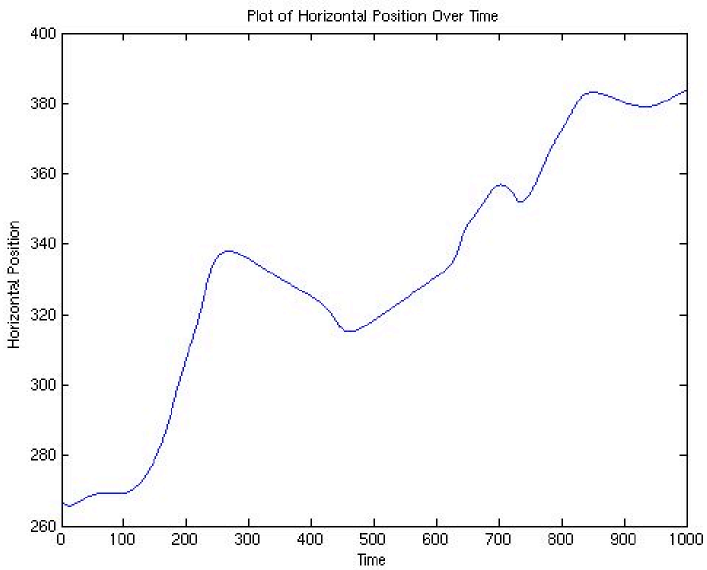

Plot of hidden neuron activations over the trial:

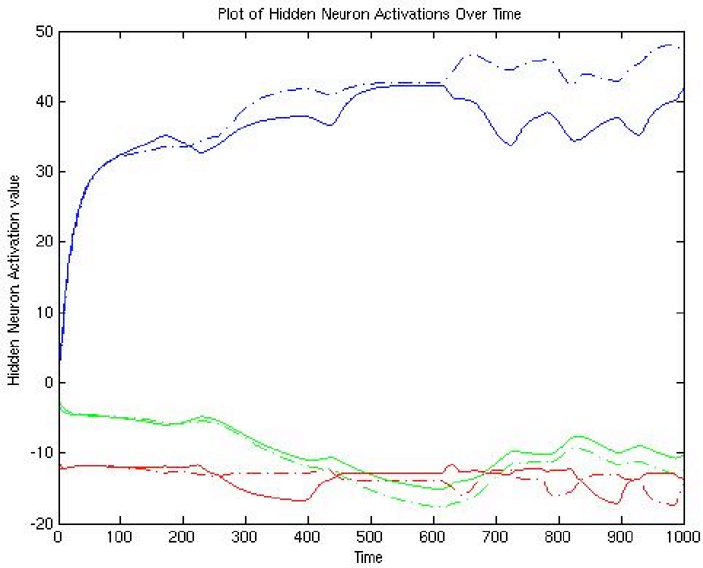

Plot of hidden neuron firing over the trial:

The outer neurons (1,6) (below) fire according to the outer neuron activations (red line above). One would expect the hidden neurons to fire according to the input neurons, however in this case they activate and fire in an opposite ‘direction’. If the inputs indicate that a right movement should occur, the left outer neuron activates and fires (solid red line below), if left then the right outer neuron (slash-dot red line below) becomes more active.

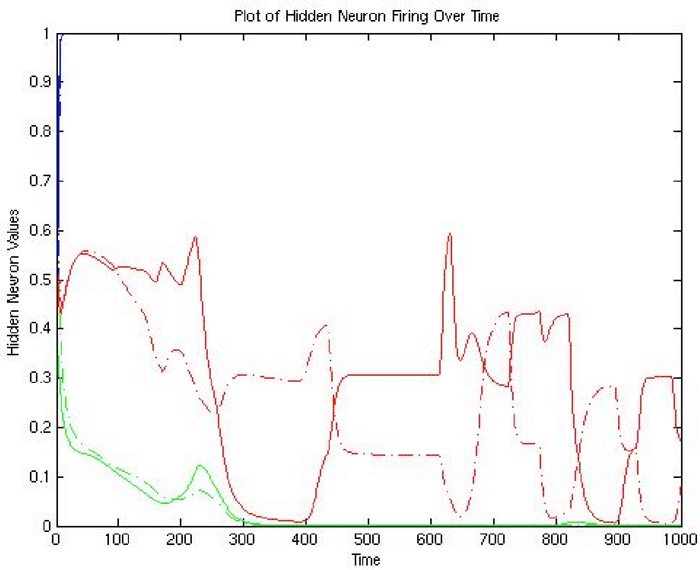

Plot of output activations over the trial;

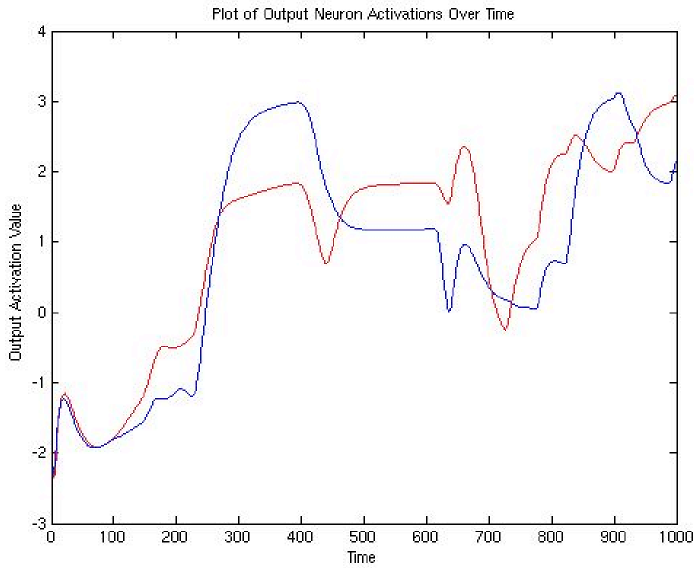

The output neurons are connected to the hidden neurons in such a way that they fire in an opposite direction to that of the hidden neurons. I.e. if the right hidden neurons fire more, the left output neuron (blue line above) fires to a greater extent, if the left hidden neurons fire more the right output neuron (red line above) fires to a greater extent. This is a good example of how simple control systems can produce complex behaviour. In addition it clearly displays that evolution will often use different architectures for the same problem. Although the neurons 'process' the information differently from example 2 A, the resulting behaviour is almost identical (different neural solutions for a particular strategy but based on similar environmental patterns).

Appendix B

(Program/Code Description including debugging procedures and professional considerations that were made before starting the project. As part of the requirements a project log is also provided at the end of this section.)

Appendix B

The project is produced and completed in accordance with the code of conduct and the code of practice as published by the British Computer Society. As the project is scientific in nature and the developed software will not be used by any end users, only a few of the ethical considerations are appropriate. The main considerations that I feel are applicable are listed in the attached interim report, compiled before work had commenced on the project. For convenience a copy of the professional considerations is given at the end of this section.

Although the project was started with the knowledge that the developed software would not be used by any end users it still had to be designed within certain guidelines. Firstly the object-orientated approach is used to some degree in the project, to separate out the functionality of different components, the different classes representing the source code being split up according to the distinct functions. Secondly, the code had to be flexible enough as to be easily edited and changed to fit the nature of the project. For example, for each main experiment where a new agent was evolved, many of the same classes were used as before, however one or two might have to modified, for example to change the fitness function. To easily allow such modifications variable names were labeled sensibly as to avoid confusion and the code clearly commented to aid the understanding. In addition variables were used as often as possible, the values of which either passed in as parameters or set at the beginning/top of each class thus making the changing of parameters a much quicker process. Overall the class structure is simple and remains reasonably elegant. Repetition is kept to a minimum and used only if it makes the code easier to understand and maintain.

As the project progressed it was important to debug each stage, the project was developed using the BlueJ compiler and thus the provided debugger was used. This is a relatively basic complier and chosen for two main reasons, firstly it provides a clear visual/object orientated representation of the class structure (see below) and secondly it allows users to easily run individual methods separately and work interactively with objects (viewing contents of objects and variables, while the program is running). The project is in the form of a scientific report, the software will not be used in an everyday situation or by novice users, consequently it does not have to be hugely robust (although it must function correctly and reliably!); thus the usability was tailored to my own specifications/requirements. Despite this fact, the source code was debugged for both compile and runtime errors at each major stage or addition. The results proved that the code is stable enough to repeatedly run numerous experiments. It should also be noted that each section of code, especially the parts that dealt with complex interactions using information from different data sources and involving potentially complex mathematics were thoroughly debugged. Short simulations were run step by step to assure the correct performance of each of the individual functions. A very simple graphical interface/visualization display was also constructed to aid the debugging of sensors and agent, this displayed the agent and circle positions at each time step and highlighted the intersection of specific sensory rays.

The resulting source code used for experiment consisted of six main classes.

1) Sim - The top level of the hierarchy, its main method sets up instances of the other classes, initiating and starting an evolutionary run and displaying the resulting agent after. The display window is constructed and tailored in this class, the code is also written so that it can be easily modified to display previously evolved agents from stored genotypes (through commenting four lines of code (at the top of the class).

2) Run - This class is very simple and is used to control the current thread that dictates how often the display is updated and the length of the simulation time step when displaying the evolved agent.

3) DrawCanvas - Handles the drawing of the visual display (agent, sensors, environment, canvas in general). Also contains method used to reset the simulation.

4) MicrobialGA - Contains the evolutionary algorithm used to evolve agent.

5) Agent – Deals with agent and circle movement, interactions, network calculations, sensor calculations.

6) ReadGenes - Reads information from a specified text file and stores it in an accessible array.

These classes are usually accompanied by three other classes that handle individual experiments:
Labeled exp1 exp2 exp3; the first of which is normally used to randomly test the agent over 1000 trials; the last two vary according to the investigation. The diagram below represents the above class structure and illustrates some of the class connections/dependences (Diagram taken from blue J compiler).

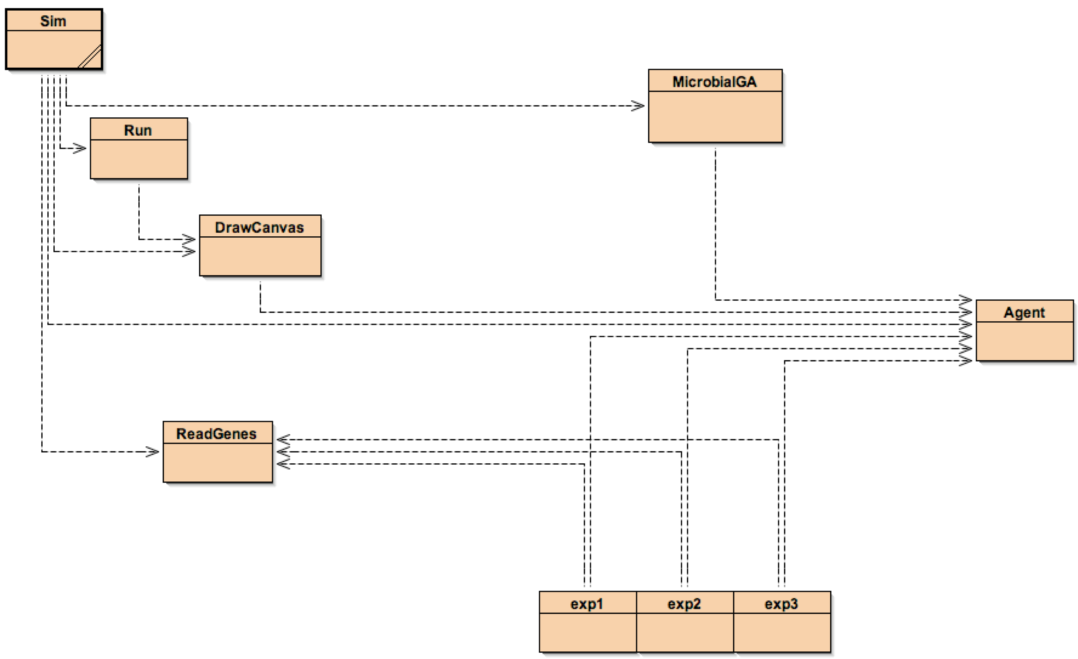

Copy of Professional Considerations (Copied from interim report - identified before starting the project):

"The project will be produced and completed in accordance with the code of conduct and the code of practice as published by the British Computer Society. As the project is scientific in nature and the software developed will not be used by any end users, only a few of the ethical considerations are appropriate. The main considerations that I feel are applicable are explained below.

- What is my level of competence at the subject, are the goals likely to be achieved? Although inexperienced in the implementation of certain aspects of the project, I have some background in the use of neural networks and genetic algorithms. I feel confidant that any novel concepts can be grasped and that the project goal is not beyond my capability.

- What are the goals that are realistically achievable in the given time? There are many variations of the described experiments that would provide valuable scientific data; however exploring the effect of every variable can be time consuming. Only a few major goals are given in section five these are the goals that I aim to achieve within the given time. If ahead of schedule there are a number of possible extensions that could be implemented (such as evolving the agent in a two dimensional environment).

- How will I deal with the knowledge that I lack and any inexperience? The tactic that will be used is to split the project up into separate phases as described above. My experience is limited in evolving continuous time recurrent neural networks and simulating agents. To combat this problem, the implementation of the project was started immediately so that I could gain experience with simulating and evolving very simple control systems for a simpler problem. Step by step the simulation, neural network and genetic algorithm will be increased in complexity and gradually more analysis tools are added, thus the project follows a form of incremental evolution. In this way I will always have achieved something rather than leaving the whole project to do in one go. Having completed the simple simulations, the programming already produced will enable me to easily 'plug in' the more complex neural network into the simulation with little change to the programming. In addition the knowledge obtained through working out the details of and investigating simpler control systems will aid the study of problem solving in more complex control systems.

- The results obtained by this project will be valid and repeatable, and the programming style will be following the code of practice guidelines.

- Will the report be understandable? The report will be written in a style that is easy for others to understand and will also clearly state the methodology employed so that the experiments can be replicated with ease."

Project Log

First Term:

Week 2, 2005:

Discussed proposal by email with supervisor, identified a subject matter to explore, and identified a main goal for the project (as described in the introduction).

Week 3 2005:

Met with supervisor for half an hour. Discussed project ideas, and created a rough plan for the term.
Week spent reading papers relevant to minimal cognition experiments (listed in bibliography).

Week 4 2005:

The second meeting with supervisor, in which a two dimensional simulation was discussed. The implementation of this two dimensional simulation of a simple Braitenberg vehicle was then suggested. The week was spent programming the simulation of a simple Braitenberg vehicle capable of performing photo-taxis.

Week 5 2005:

Problems with two-dimensional simulation discussed. Experiment involving a robot catcher confined to one dimension of movement was then put forward. This was discussed briefly and then continuous time recurrent neural networks were explained for rest of meeting.

The week spent perfecting the two-dimensional simulation, fixing the original problems and thinking about the one-dimensional simulation.

Week 6 2005:

Photo-taxis two-dimensional simulation completed. One-dimensional simulation discussed in more detail.
Sensor mathematics were given, and briefly explained. A potential evaluation function was briefly discussed and continuous time recurrent neural networks were again brought up.
Tried unsuccessfully to get the one-dimensional simulation fully working, however many aspects such as bilaterally symmetrical neural network were implemented.

Week 7 2005:

Meeting involved a detail discussion of the sensor mathematics and the evaluation function.
Successfully programmed the simulation of a Braitenberg one-dimensional ball 'catching' agent.
However had no luck when trying to evolve a simple feed forward neural network on the same task.

Week 8 2005:

Evaluation function discussed in explicitly detail. All aspects of the genetic algorithm details also explained in depth (including liner mapping).

Whole week spent trying to evolve robot to catch a falling object yet still only partially successful.
CTRNN implementation delayed however the problems that were encountered had a knock on effect of allowing me to fully analyse the workings of the genetic algorithm and a simple feed forward network and forced me to trial various analytical methods ahead of time.

Week 9 2005:

No meeting with supervisor, Met with postgraduate Eduardo Izquierdo-Torres to discuss evolving problem and project ideas.

Week spent trying to solve the evolving problem. Identified that the problem was with the physical attributes of the robot rather than the evolutionary algorithm or neural network. Due to a heavy workload only a small amount of work was completed, the project will obviously overrun into the Christmas holidays.

Week 10 2005:

No meeting with supervisor. Finished interim report.

Second term:

Week 2 2006:

Week spent organizing work completed over Christmas. Christmas spent improving programming in terms of flexibility and complexity. In addition many trials were run using feed forward neural networks and an attempt was made understand CTRNN's fully and get an implementation working correctly. Much reading also accomplished within the holiday.

Week 3 - 6 2006:

Many problems implementing the continuous time recurrent neural network however after much discussion with supervisor and help from Eduardo Izquierdo Torres eventually the agent's control system worked correctly. In addition successfully accomplished evolution of agents capable of relational discrimination.

Week 7 – 8 2006:

Started formal experimentation and the analysis of neural behaviour, changed from using excel to Matlab for viewing and manipulating data. Discussed various evolution problems and experimentation methodologies with supervisor. Started formally writing up the report.

Week 9 2006:

Results almost too good to be true, evolution seemed too easy. On closer analysis the results were very strange consequently most of the week spent reviewing programming, and debugging. Eventually the problem was tracked back to a bug in the agents' sensory system. Consequently all previous results were invalid, in this case the problem was a result of a modification made previously when experimenting with the sensory systems and was never changed back. (Other course work in this week meant less time and concentration spent on project and as a result this modification was forgotten!). The project was set back considerably; however after making sure everything was working correctly the evolution of different agents was restarted.

Week 10 2006:

Continued evolution of agents and report write up. Draft report shown to supervisor and outline agreed for final report.

Easter Holiday:

Holiday started with a considerable amount of work to achieve! Evolution of absolute and relational agents completed, experimentation also completed and results written up. Computer failure late at night caused a set back at the end of the second week, much of the beginning of the report had to be rewritten from scratch. Eventually the results were written up, the last two weeks were spent drawing conclusions and finalising the report, in addition an attempt was made to evolve an agent capable of selecting the middle of three stimuli. Beginning of summer term spent preparing presentation.

Personal Conclusion:

Apart from the scientific goals, a personal aim was to gain some experience in evolving and simulating simple robots/agents. The project has thus introduced to me to various evolutionary robotic techniques many of which would not have been learnt if this project had not been undertaken. Although the programming requirement was modest compared with some software engineering projects it has still allowed me to gain some experience with planning, handling and debugging larger applications. In addition it has emphasised how difficult it can be to keep track of modifications and thus how easy it is to make mistakes. Most importantly however is the experience gained in managing time and in recording and analysing data. Large amounts of raw data were generated; not only taking considerable time to analyse but also requiring care when recording (On several occasions ambiguous or poorly labelled files and folders led to confusion). Problems were also encountered when deciding which data to present and write up, the limited word count restraining the data that could be presented and described. Looking back, much time was spent understanding and implementing novel concepts and in hindsight much more planning should have been accomplished in terms of thinking through the programming structure and experimental procedure. However, the project has allowed me to gain some experience in tackling a scientific investigation and overall a satisfying result has been obtained.

Appendix C:

(Source Code)

Part 1) Java Source Code for Visual Orientation Experiments. - All the classes are printed for this experiment, this is the base code on which other experiments were built on (Sim, Run and ReadGenes remain the same for each experiment).

Part 2) Java Source Code for Size Discrimination Experiments. - Many similar classes shared with above, only those modified are printed. It must be noted that small modifications were made to each class according to the specific experiment (i.e. changes made appropriately when evolving and testing the large size discriminator, the small size discriminator, the absolute size discriminator and the non- absolute size discriminator) these changes are described in the main body of the report and the areas of code that were frequently modified are commented suitably within the source code.

Part 3) Java Source Code for' Middle Size' Discriminator Experiment. - Again shares the same classes above, those modified are printed.

Part 4) Standard Matlab code used for the visualisation of the experimental data.

Part 1) Java Source Code for Visual Orientation Experiments

Part 2) Java Source Code for Size Discrimination Experiments

Part 3) Java Source Code for' Middle Size' Discriminator Experiment.

Part 4) Standard Matlab code used for the visualisation of the experimental data.