

# Building Aunt Hillary: Creating Artificial Minds with ‘Neural Nests’

Mike Reddy and Stuart Lewis

School of Computing, University of Glamorgan  
Treforest, Pontypridd., RCT. CF37 1DL Wales, UK  
{mreddy, sflewis}@glam.ac.uk  
<http://www.comp.glam.ac.uk/staff/mreddy/>

**Abstract.** This paper describes work in progress to enable a real robot to recreate trail following of ants engaged in pheromone-reinforced recruitment to food gathering. Specifically, it is proposed that development of a set of macro-behaviours for creating and following a trail can be achieved by use of micro-behaviours in a simulated environment to develop a novel neural architecture – the Neural Nest - for learning without explicit representation. A simulated ‘neural nest’ has been tested to determine the feasibility for ant colonies to encode higher-level behaviours for controlling a physical robot. In our experiments, the emergent behaviour from reinforcement of interactions between unsupervised simple agents, can allow a robot to sense and react to external stimuli in an appropriate way, under the control of a non-deterministic pheromone trail following program. Future work will be to implement the architecture entirely on the physical robot in real time.

## 1 Introduction

Swarm Intelligence (SI) research has been useful for applying models of adaptive behaviours of complex biological structures, consisting of simple entities (e.g. termites, ant nests, and beehives), to computational problems that require prohibitively expensive search algorithms [1]. One of the key strengths of SI is that elements have flexibility to interact in unpredictable ways, to facilitate optimisation. Specifically, Ant Colony Optimisation (ACO) [2] and Ant Colony Routing (ACR) [3] have been proposed, which use the analogy of pheromones in stylised environments to govern search and routing strategies. However, these techniques have, so far, been mostly applied to specific topological/geographical problems (e.g. the Travelling Salesman Problem [2], Communications Networks [3-5], etc). It should be noted that other approaches, such as novel neural networks [6] or evolutionary algorithms [7] may be more successful, where potential solutions are numerous, or uniformly distributed [1].

SI has also been applied to the problem of coordinating groups of autonomous robots to collaborate on a task [8]; the concept of simulated ants being mapped directly onto sets of robots situated in a simulated or real environment. The power of SI has, therefore, been in coordinating these situated agents at the macro-level without the

need for planning or explicit communications. However, most existing ant colony simulations do not attempt to address the complex emergent control of individual physical ants (robotic or biological). We believe that the strength of SI – its use of micro-components to optimise a complex macro-system – may also contribute to the control of individual micro-components, such as a single physical robot, rather than just simulating them as part of a larger macro environment. Many adaptive robot control architectures have been proposed, to control agents at the micro-level, that is, within the individual robot; these include traditional and recurrent neural networks and evolutionary programming, etc. (A good summary can be found in [9, 10]). Ethology, in the form of Tinbergen’s instinct centres [11] has inspired a learning classifier approach using Genetic Algorithms (GAs) [12]. However, there has been no attempt to make direct use of SI to control robots at the micro-level; namely to implement an autonomous learning robot using an Ant Colony System (ACS). Simulation of ant behaviour may yet prove useful for understanding the emergent properties of complex biological structures within the human brain. The neural nest, proposed in this paper, is an ant colony inspired implementation of a real-time neural network, without a pre-defined symbolic representation or the need for supervised learning. To allow emergent intelligence, it is, vital that a physical platform be available to reproduce the conditions for evolutionary development, and to enable evaluation of the ant-powered brain. Finally, we describe integrating the symbolic ant brain with the physical robot.

## 2 Previous Work

### 2.1 Aunt Hillary – A Reprise

In his seminal work “Gödel, Escher, Bach – an Eternal Golden Braid” [13], Douglas R. Hofstadter describes Aunt Hillary, a conscious ant colony, who consists of signal propagating and symbol manipulation activities using ants as base units interacting with the colony environment to shape the activities of ants into higher level concepts. This is not at odds with more symbolic representations. However, it is difficult to perceive how behaviours could evolve, when we only consider large grained views of intelligent systems. While Aunt Hillary serves principally as a clever analogy to explain how neurons can collaborate to produce symbol manipulation and, ultimately, human consciousness, the concept is an inspiring one, raising many interesting questions, especially when the brain can be considered a network of real-time processes (or neurons), with a flow of control from sensors to motors. Symbolic and sub-symbolic techniques make no attempt to describe how higher-level concepts could be implemented in the structures of the brain, and draw a sharp distinction with the physical nervous system. It must also be remembered that it is not just brains, but also bodies (sensors, actuators, etc.) that have co-evolved. Biological systems have not only learned to adapt, but also adapted to learn.

## 2.2 AntNest – A Reprise


Real ants have limited senses and communication, but manage to evolve optimal solutions to search and routing problems, using pheromones trails, which result in a positive feedback effect. Several simulations of this activity exist, but most have a simple rule-based approach for controlling ant behaviour. Furthermore, these simulations are limited, especially where the strength, degradation and effectiveness of scents can be critical for success. Our previous work, AntNest, has been particularly successful due to using two scents: one for food-bearing ants (food scent), and one for foraging ants (nest scent) [14]. The revised rules in AntNest are as follows:

- 1) Ants leave the nest in a random direction, leaving a nest scent in the process.
- 2) Random search influenced by the food scent level, if any, leads the ant to food.
- 3) Some of the food is acquired and the ant begins to lay food scent towards home.
- 4) Random movements, influenced by the nest scent level, lead the ant homeward.
- 5) The ant then returns to step 1.

AntNest was originally conceived as a means for visualising Economics problems, by modeling a biological simulation of an ant colony; ants had several roles in addition to food gathering: nursing babies, clearing the dead from the nest, attacking enemies and sleeping. However, we were inspired by the similarity of the simulations to the architectures of neurons: axons and dendrites. Several features of current neural network architectures have been abstracted from the biological reality: real neurons react to rates of fire, and do so in a continuous time-frame; they are also excitatory or inhibitory in nature, whereas artificial neurones can generate positive and negative values, mediated by weightings (again positive or negative) at the receiving neuron. The research question was whether we could use competing ant colonies to implement a neural architecture that was more closely inspired by the qualities of real neurons: ant trails being analogous to neural links, while the rate of flow of ants from source to nest representing the rate of fire of a neuron.

## 3 Neural Nests – A Proposal

In this section, a connectionist reinforcement learning architecture is described, which may enable an autonomous robot to acquire motor-sensory navigation strategies without explicit representation. The use of an embedded ACS to implement a novel neural architecture has several advantages: Firstly, it allows us to experiment with combining the strengths of both SI and neural networks; Secondly, it allows the possibility of applying SI techniques in a direct manner, within a single entity, namely a physical robot ant in a real setting; Thirdly, it satisfies an aesthetic need to explore the ability for complex systems, such as the human brain, to evolve and manipulate meaningful symbols using simple components, without explicit prior representation as is the case for many other solutions to robot control.




**Fig. 1.** Neural Nest Internals: Food Sources and Motor Nests are passive entities inside the environment, which represent sensor inputs and motor outputs respectively. Ants collecting food and returning to the nest will cause the external robot to act accordingly depending on their behavioural properties

Furthermore, a tentative cognitive model of unsupervised neural learning is described – the Neural Nest – which uses a set of simple micro-agents to recreate the linkage and firing behaviours of biological neurons to demonstrate the feasibility of our approach to real autonomous robot control. It was recognised that a physical implementation in a real robot was needed as early as possible, rather than working purely with simulation, and then attempting to port to a robot afterwards. Simulation does not throw up the real-world errors and inconsistencies that are present with even the most simplest physical environments; this is especially important as we consider the body evolution to be key to brain evolution. However, the neural nest architecture is a slow one at present – certainly not capable of working in real-time – so a simulated robot was modelled in a closed environment. Initially, this robot had just two touch sensors and two motors.

The basic currency of the Neural Nest architecture is the collection of food from food sources linked to external sensors on a real or simulated robot. Nests are associated with external actuators (in this case, motors on the robot) and the collection of food and its return to the nest provides a signal processing capability that will control the robot's behaviour. The first experiment in the Neural Nest simulator was to see whether the robot would move more sensibly with two motor nests and two sensor food sources (See Figure 1). This first experiment showed the gradual grafting of the two touch sensors onto their respective motors. The early effect of this was to drive the robot into obstacles rather than away from them. However, over time a cross-link occurred that allowed the robot to be a little more effective at avoiding walls, when approaching them with one of the sensors. The second experiment involved adding light sensors to the robot, in order for it to engage in line following. We added food sources related to position of the sensor array on the front of the vehicle to the previous simulation arrangement, placing the new food sources in comparable positions to their placement on the physical robot described in Section 4. Again in simulation the architecture generated by the simulation did not show the bipolar symmetry that most animals display. This may be a fault in the model, such as scent evaporation, or faults in the simulated environment<sup>1</sup> but it may also be due to a very limited time available for the system to evolve a solution before testing. The model generated is shown in Figure 2. From observation of the simulated robot, what appeared to be happening was as follows:

- 1) The right motor stayed on continuously, so directions and steering came from the activities of the left motor, unless the right touch sensor was activated;
- 2) The left motor stayed on continuously, but reacted to both the left touch sensor and the left and centre light sensors.
- 3) When a line was detected by either of these light sensors the left motor would stall. If both detected a line the left motor would reverse, turning the robot left.

This gave a limited ability to sense and follow lines, while also maintaining the ability to react partially to the touch sensors. The effect of this architecture on the physical environment is detailed in Figure 4.




**Fig. 2.** Touch and Light Simulator: (a) The two motor nests search for food; (b) Neural architecture is complete with the left-hand nest also using the left light sensor

<sup>1</sup> This could be the random number generation skewing movement.

## 4 Robot AUNT<sup>2</sup> – The Macro-level

Robot AUNT is a tracked robot with two touch sensors on the front corners and three light sensors mounted facing down at the front of the canopy. The experiments ran on a white board placed on the ground and fenced off with walls, which the touch sensors would detect. A retractable white-board marker allowed the robot to lay a trail for the light sensors to follow. The initial problem to be faced was whether such a pen-based trail on a whiteboard could act as a suitable environment for testing pheromone following with a robot. Initial experiments were performed using a hand-coded program to test the feasibility of a real physical implementation of pheromone following. Note: This was done for various reasons, but the most significant was to test whether I could get the tracks to perform a rudimentary form of evaporation – the pen marks will not fade the way that a scent trail would, and there is only one robot, rather than thousands. In fact, the robot tracks provided an excellent mechanism for reproducing the effect of scent evaporation, even if not in the way it would in reality<sup>3</sup>.




**Fig. 3.** Robot AUNT: (a) Touch sensors mounted on the corner; (b) Three light sensors facing down; (c) A retractable pen mounted at the pivot point to mark trails

The key issue was then to determine whether the neural nest architecture could be used to control a physical robot. In order to do this effectively, an abstraction of the simulated network was implemented as a set of probabilistic choices in the control of robot motors, dependent upon sensor values. We simulated the arrival of ants as decisions to go forwards and backwards, with the delay between sensing and acting proportional to the average time for an ant to travel between the food source and the nest. Clearly, this is not a perfect reproduction of the actual nest architecture, but allowed a simple implementation on the robot. As was true for the simulation, the robot tended to veer to one side and occasionally trapped itself in corners. The next stage of the experiment was to implement the architecture described in Figure 2, using light and touch sensors. In order to do this, Robot AUNT would need to be able to lay trails autonomously. Modifications were made to the environment to allow the robot to detect whether it was at a food source (shiny surfaces) or the nest (black surface). The

<sup>2</sup> AUNT stands for Autonomous Unsupervised Neural Turtle – in honour of Aunt Hillary. The robot is a turtle rather than an ant because of its previous incarnation as a drawing robot.

<sup>3</sup> It is a little known fact that white-board marker is easier to wipe off once it is completely dry. The efficacy of the tracks to wipe the board clean was proved beyond doubt when we forgot to change the pen for one experiment and returned to find the board eerily clean.

program was modified so that the pen would be lowered once the robot had found food, and raised when it had returned successfully to the nest. Given that there was no help from a number of ants leaving trails, Robot AUNT did mark out trails to and from the nest (see Figure 4). It also performed slightly better at following lines than the simulation did<sup>4</sup>.



**Fig. 4.** Robot AUNT II: (a) Searching for food by following lines from previous journeys; (b) Carrying food back to the nest, leaving a scent trail in white-board pen

## 5 Conclusions

A cognitive model of unsupervised neural learning has been described – the Neural Nest – which uses a set of simple micro-agents to control the macro-behaviours of a physical robot. We have seen a neural nest link touch and light sensors to motors in a way that is effective in both a simulated and real robot, with a degree of sensor fusion. Ant trails and food foraging have been used as the basis for a computational model of how neurons forge links and activate by rate of firing and excitatory and inhibitory behaviours. Results are encouraging, but tentative, with more quantifiable analysis needed to ensure that this is a valid technique for evolving robotic control systems. However, the current scent model needs some revision, as nests switch between different sensors, with only one good link at a time. Sensitivity analysis of the range of producible behaviours is also needed. Furthermore, the resulting robot search pattern is prone to the lack of positive reinforcement, where real ants would be surrounded by recruited colleagues in scent trail generation; our studies only used one macro-robot. Further research into the use of multiple robots may show more clearly the effectiveness of this macro and micro approach to robot control. While this may not produce the most efficient solution to the problem, it is our hope that such an evolutionary approach may also eventually lead to a better abstract representation of the link between low-level neural functioning and higher order behaviours of a robot.

<sup>4</sup> This is most probably due to the need to use eight fixed directions in the simulation, where a physical robot would be changing direction more uniformly using real motors.

## References

1. Bonabeau, E., Dorigo, M., Theraulaz, G.: *Swarm Intelligence: From Natural to Artificial Systems*. Oxford Univ. Press, New York, (1999).
2. Dorigo, M., Gambardella, L. M.: Ant colonies for the traveling salesman problem. *BioSystems* 43, 73-81 (1997).
3. Schoonderwoerd, R., Holland, O., Bruten, J., Rothkrantz, L.: Ant-based load balancing in telecommunications networks. *Adapt. Behav.* 5, 169-207 (1997).
4. Heusse, M., Guérin, S., Snyers, D., Kuntz, P.: Adaptive agent-driven routing and load balancing in communication networks. *Adv. Compl. Syst.* 1, 237-254 (1998).
5. Di Caro, G., Dorigo, M.: AntNet: Distributed stigmergetic control for communications networks. *J. Artif. Intell. Res.* 9, 317-365 (1998).
6. Chen, K.: A simple learning algorithm for the traveling salesman problem. *Phys. Rev. E* 55, 7809-7812 (1997).
7. Baluja, S., Caruana, R. In Frieditis, A., Russell, S.: *Proc. 12th Intl Conf. Machine Learning*. Morgan Kaufmann, Palo Alto. (1995). 38-46.
8. Kube, C. R., Zhang, H.: Collective robotics: from social insects to robots. *Adaptive Behavior* 2, (1994). 189-218.
9. Dorigo, M. (Guest ed): Special Issue on Learning Autonomous Robots. *IEEE Transactions on Systems, Man, and Cybernetics-Part B* Vol.26, No.3, (1996).
10. Dorigo, M., Colombetti, M.: *Robot Shaping: An Experiment in Behavior Engineering*. MIT Press. (1997).
11. Tinbergen, N.: *The Study of Instincts*. Oxford University Press. (1966).
12. Dorigo, M., Schnepf, U.: Genetics-based Machine Learning and Behaviour Based Robotics: A New Synthesis. *IEEE Transactions on Systems, Man, and Cybernetics*. Vol. 23, No. 1, (1993) 141-154.
13. Hofstadter, D.: Gödel, Escher, Bach: An Eternal Golden Braid.
14. Reddy, M., Kedward, P.: Ant Economics: What simple agents can teach us about supply, demand and eating your Dead. *Procs. of ANTS'2000 - From Ant Colonies to Artificial Ants: 2nd International Workshop on Ant Colony Optimization*. Brussels. Belgium. 8-9<sup>th</sup> September 2000.