

Minimizing human intervention in the development of basal ganglia-inspired robot control

doi:10.1080/11762320701492792

F. Montes-Gonzalez¹, T. J. Prescott²
and J. Negrete-Martinez¹

¹Department of Artificial Intelligence, Universidad Veracruzana, Sebastian Camacho 5, Xalapa, Veracruz, Mexico

²Department of Psychology, University of Sheffield, Western Bank, Sheffield S10 2TN, UK

Abstract: A biologically inspired mechanism for robot action selection, based on the vertebrate basal ganglia, has been previously presented (Prescott *et al.* 2006, Montes Gonzalez *et al.* 2000). In this model the task confronting the robot is decomposed into distinct behavioural modules that integrate information from multiple sensors and internal state to form ‘salience’ signals. These signals are provided as inputs to a computational model of the basal ganglia whose intrinsic processes cause the selection by disinhibition of a winning behaviour. This winner is then allowed access to the motor plant whilst losing behaviours are suppressed. In previous research we have focused on the development of this biomimetic selection architecture, and have therefore used behavioural modules that were hand-coded as algorithmic procedures. In the current article, we demonstrate the use of genetic algorithms and gradient-descent learning to automatically generate/tune some of the modules that generate the model behaviour.

Key words: Basal ganglia, action selection, reactive robotics, evolutionary robotics.

INTRODUCTION

The problem of action selection – deciding what to do next – is shared by both biological and artificial autonomous systems (Prescott *et al.* 1999). In order to learn more from biology some roboticists have started modelling aspects of animal behaviour in robots (see e.g., Dean 1998). To develop an animal-like robot (*an animat*), it is necessary to identify the behaviour that is required for the robotic task to be solved. A mechanism for generating this desired behaviour is then required. Some models attempt to do this in a direct and holistic fashion in which the complete behaviour is generated by a single system. A more common approach, which exploits the advantages of modular decomposition (Bryson 2000, Prescott *et al.* 1999, Prescott *in press*, Wagner and Altenberg 1996), is to break down the overall behaviour into sub-units – simpler action patterns

that can be assembled in sequence to obtain the target behaviour. In these circumstances an arbitration mechanism is needed to choose among the repertoire of behaviours available for selection at any given time. A large variety of *action selection* mechanisms have been proposed for use in arbitration (Bryson 2000, Maes 1995, Prescott *et al.* *in press*, Prescott *et al.* 1999). We have previously argued (Prescott *et al.* 1999, Prescott *in press*) that the design of the vertebrate brain makes use of a centralized selection system whose neural substrate consists of a group of fore brain and mid brain nuclei known as the basal ganglia. The present paper therefore focuses on the continued development and evaluation of our computational model of the vertebrate basal ganglia as a robot controller (Prescott *et al.* 2006, Montes Gonzalez *et al.* 2000).

With regard to the target behaviour of interest, it is common for roboticists to focus on foraging, social-cooperative, and social-competitive behaviours in building animats (Montes González *et al.* 2006). One of the main arguments for choosing these kinds of behaviour is to be able to take advantage of relevant studies in the biological literature. Next, a robotic platform for the implementation of the animat has to be both built and programmed. In our earlier work we focused primarily on the selection mechanism rather than on the component behaviours required

Corresponding author:

F. Montes-Gonzalez
Department of Artificial Intelligence
Universidad Veracruzana, Sebastian Camacho 5
Xalapa, Veracruz, Mexico-91000
Tel.: +52 228 8172957, +52228 8172855
Email: fmontes@uv.mx, jnegrete@uv.mx

to complete the task. In this paper, we use our existing model of the vertebrate basal ganglia as the arbitrator, and focus on the use of artificial evolution and learning to optimize the behavioural elements needed for a biologically inspired foraging task. The goal of the present research is to demonstrate that the model basal ganglia can work as effectively with evolved/trained components as with hand-coded ones, which should pave the way for the development of complete biological-inspired robot controllers that are constructed with minimal human intervention.

The remainder of the paper is organized as follows. A brief introduction to the genetic algorithms used in the reported experiments is provided in the section ‘evolution of robot behaviour’. Next, we introduce the design of each of the behaviours for the foraging task that is modelled on behaviour observed in laboratory rats (see Prescott *et al.* (2006) for more details). These components include searching a walled-arena for cylinders (*cylinder-seek*), collecting a cylinder in the robot gripper (*cylinder-pickup*), searching for, and following walls (*wall-seek*, *corner-seek*), and depositing a cylinder in a corner of the arena (*cylinder-deposit*). The use of our computational model of the basal ganglia for action selection is explained in later section. In the fifth section we present the results of integrating (evolved) neural networks-based behaviours with some hand-coded algorithmic behaviours. Next we offer a discussion about the present extent of our work. Finally, we provide some conclusions concerning the significance of the experiments described.

THE EVOLUTION OF ROBOT BEHAVIOUR

Evolutionary methods have been widely used in the development of controllers and elementary behaviours for robots performing survival tasks in semi-structured environments (Florenzano and Mondana 1994, Bajaj and Ang Jr. 2000, Nolfi 1997). The use of neural networks together with genetic algorithms (GAs) (Holland 1975) is also common practice for the development of robot control systems (Nolfi and Florenzano 2000). Typically, a population of robot controllers is represented using a fixed neural network topology with individuals differing according to the values of their weighted network connections. Next, the evolutionary operators – selection, crossover, and mutation – are applied in order to produce new offspring. The decoding of every genotype into individual phenotypes (neural controllers) is followed by the testing of individual robot controllers for a limited period of time (Balakrishnan and Honavar 1995), and the production of new offspring is halted when a satisfactory evaluation level is obtained. If we measure the fitness of all possible individuals, we will obtain a convoluted landscape – this is the search space over which the genetic operators are applied. Evaluating robot fitness is very time consuming; therefore, in the majority of cases a robot simulator is preferred to evaluate candidate controllers that are subsequently transferred to the real robot for final adjustment (Santos and Duro 2005).

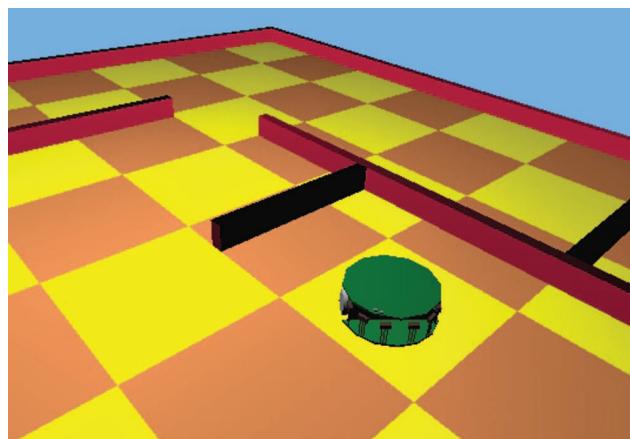


Figure 1 The *wall-seek* and *corner-seek* behaviours were evolved using a 3D robot simulator.

Usually the representation of a new robot controller is coded into a chromosome, where each locus (or position) takes a finite possible value (allele). This representation corresponds to the genotype that in most of the cases is directly encoded into the phenotype. For the initial population random controllers are spawned, their fitness evaluated, and the GAs’ operators applied. The *selection* operator chooses to breed the fittest individuals into the next offspring using operators such as the crossover and the mutation. As an example, *tournament selection* chooses to breed new individuals from the winners of local competitions. The fittest individuals at each stage of the competitions are often inserted into the new generation intact – a practice called *elitism*, which guarantees that the best solution found so far is not lost. *Crossover* is an operator that takes two individual encodings and swaps their contents around a couple of random points along the chromosome. *Mutation* occurs with a probabilistic flip of some of the chromosome bits of new individuals (in general, with a random change of the alleles of the genes). Although, mutation produces inferior individuals, better individuals may occasionally emerge.

THE DESIGN OF INDIVIDUAL BEHAVIOURS

In our experiments we have used the small (55 mm diameter), ‘table-top’ Khepera (Mondana *et al.* 1993) robot together with a commercial robot simulator known as Webots (Cyberbotics 2006) (Figure 1). The Khepera has two DC motors for controlling the movement of the wheels and eight infrared sensors distributed around the body of the robot (but predominating at the front). For the foraging task the Khepera is also equipped with a standard gripper turret.

Network configuration for component behaviours

The controller for each of three component behaviours (*wall-seek*, *corner-seek*, *cylinder-seek*) is a fully connected

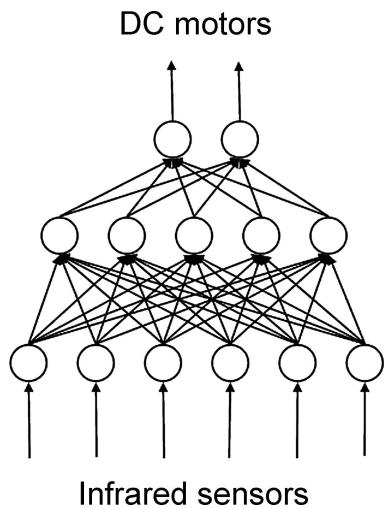


Figure 2 The behaviours *wall-seek*, *corner-seek* and *cylinder-seek* share the same neural network topology.

feedforward multilayer-perceptron neural network with no recurrent connections (Figure 2).

The topology of each neural network is six neurons in the input layer, five neurons in the hidden layer, and two in the output layer. The sigmoid transfer function is used at the hidden and the output neurons. Signals from the six infrared sensors towards the front of the robot are used as input to the network. Each signal, which is in the range 0 to 1023, with higher values indicating greater proximity, is first made binary by using a collision threshold $th_c = 750$. This step facilitates the transference of the neural controller to the robot, since this can be achieved by simply adapting the collision threshold to the real sensors. To generate robot movement, the two real-valued outputs of the neural network, which are in the range 0 to 1, are scaled to the ± 20 values required for driving the left and right DC motors driving the wheels.

The weights of each neural network are directly encoded into a vector w of 40 elements, the first 30 representing the 6×5 connections from the input units to the hidden layer and the last 10 the 5×2 connections from the hidden layer to the output units. A genetic algorithm is used to optimize these weights starting with random values $-1 < w_i < 1$ for all elements. A single vector representation is thus used to define each of the individuals in the population. The initial population, G_0 , consists of $n = 100$ neural controllers. *Elitism* is employed to replicate the two best individuals from one generation to the next. A *tournament selection* also allows four random parents to be chosen from $(n/2) - 1$ competitions. The chosen parents are bred in pairs with a random crossover point generated with a probability of 0.5. Each individual in the new population is affected with a mutation with a probability of 0.01. Fitness is tested by running each individual in the Webots simulator for about 20 seconds using the fast-speed mode. The initial location and orientation of the simulated robots are also randomized across trials.

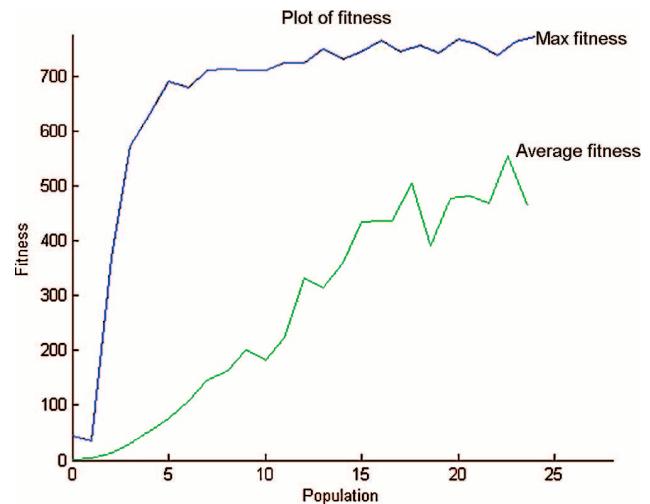


Figure 3 The predefined maximum fitness for the *wall-seek* behaviour is reached after 23 generations.

Fitness functions for component behaviours

Wall-seek: The behaviour for finding a wall can be seen as a form of obstacle avoidance in which the robot traverses the arena until a wall is reached whilst negotiating around obstacles such as cylinders. The fitness formula (Equation (3.1)) for this behaviour was

$$f_{c1} = \sum_{i=0}^{3000} \text{abs}(ls_i)^*(1 - \sqrt{ds_i})^*(1 - \max_ir_i), \quad (1)$$

where for iteration i : ls is the linear speed in both wheels (the absolute value of the sum of the left and right speeds), ds is the differential speed on both wheels (a measurement of the angular speed), and \max_ir is the highest infrared normalised value. The use of such formula rewards individuals who travel quickly in a straight line while avoiding obstacles. The use of walls in the centre of a simulated environment, such as the one in Figure 1, favours the evolution of effective wall-finding behaviour. The plot of the fitness, over 23 generations, and the average fitness for this behaviour is shown in Figure 3.

Corner-seek: Typically, the selection mechanism will stop the *wall-seek* behaviour when a wall has been reached and then switch into a behaviour that involves following the wall until a corner is located (*corner-seek*). This behaviour requires that the robot moves along a trajectory that keeps it within sensor range of the wall, but without actually touching it, until a corner is found. Note that, if the straight path to the corner is blocked, the robot must also manoeuvre around the encountered obstacle. The fitness formula (Equation (3.2)) employed for the behaviour *corner-seek* was as follows:

$$f_{c2} = f_{c1}^*(tgh)^2 \quad (2)$$

In this formula $tgh(tgh)$ – the tendency to remain close to walls – is calculated as the fraction of the test period for which an individual is close to a wall. This

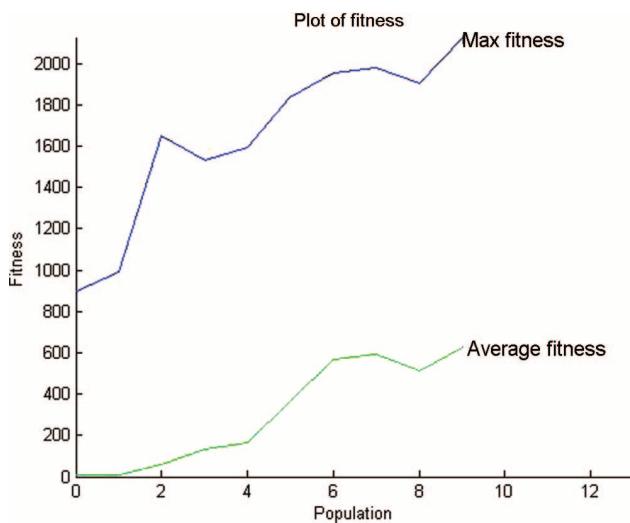


Figure 4 A coherent *corner-seek* behaviour is evolved after nine generations.

formula therefore favours individuals that travel close to walls whilst avoiding collisions. Figure 4 shows the plot of fitness and the average fitness for the corner-seek behaviour over nine generations. The selection mechanism generally operates to deselect the *corner-seek* behaviour once a corner has been found.

Cylinder-seek: This behaviour causes the Khepera to explore the open spaces in the square arena; when it locates a cylinder (identified by a distinctive pattern on the infrared sensor array) it should drive close to it and stop. The final component of this behaviour requires manoeuvring the robot body so that it is positioned with the cylinder exactly in the middle of the two most-frontal infrared sensors, as this allows unambiguous identification of the cylinder (which might otherwise be perceptually aliased with a wall). This final element of *cylinder-seek* proved to be difficult to evolve. To overcome this problem, we trained the previously described network topology using backpropagation and a set of fifteen training patterns (mapping sensor input into required motor output) corresponding to situations the robot was likely to encounter when positioning itself in front of a cylinder. In training good generalization was obtained after 1818 generations when the total error drops below 0.02 (Figure 5).

Fixed action patterns

The *cylinder-pickup* and *cylinder-deposit* behaviours differ from the three described above in that they are timed sequences of action that are triggered by an initial sensory stimulus. These behaviours are modelled on the ethological concept of a fixed action pattern (see (McFarland 1993)). For instance, cylinder-pickup requires that the gripper is opened, the robot then reverses to create free space in front of the body, the gripper is closed (hopefully around the cylinder), and the arm is moved back into the upright position. Cylinder-deposit requires a sequence of lowering

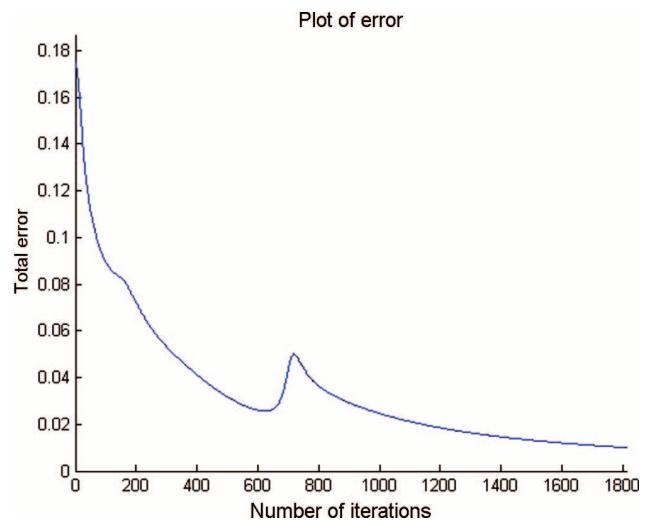


Figure 5 The neural network correctly classifies the 15 training patterns after 1818 iterations with a total-error below 0.02 for the *cylinder-seek* behaviour.

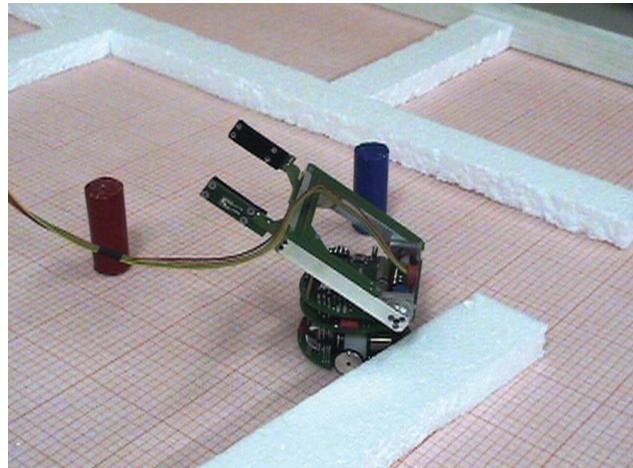


Figure 6 The Khepera robot is set in the middle of an arena with four wooden cylinders.

the arm, opening the gripper (releasing the cylinder), then raising the arm. It is our intention to attempt to evolve sequential controllers for both these behaviours, however, for the time being they have been programmed as algorithmic routines.

Transfer from the simulated Khepera to the physical robot

Once all of the mentioned behaviours were appropriately evolved/designed in simulation they were implemented and tested on the real Khepera (Figure 6). This transfer generated no significant problems indicating that the evolved/trained behaviours are robust to the differences between the simulated and real physics of the robot-environment interaction.

ACTION SELECTION IN THE BASAL GANGLIA MODEL

In earlier research we have demonstrated a computational model of the intrinsic circuitry of the vertebrate basal ganglia (Gurney *et al.* 2001a, 2001b, Humphries and Gurney 2002) embodied within a robot control system performing the task outlined in section three using hand-coded behaviours (Prescott *et al.* 2006, Montes Gonzalez *et al.* 2000). In addition to demonstrating the benefits of modularity provided by a central action selection model, these experiments have also proved sufficiency of the basal ganglia model as a selection mechanism (Montes Gonzalez *et al.* 2000). The importance of the basal ganglia in natural action selection becomes evident when we observe that these nuclei are an archaic feature common to all vertebrate animals (Prescott *et al.* 1999). The output structures of the basal ganglia generate, by default, a continuous stream of inhibition targeted at motor systems in the brainstem and cortex. Selection of a motor pattern occurs through the release of this inhibition. This occurs through activation of neurons in the basal ganglia's main input structure, the *striatum*, which in turn produces a focused inhibitory signal temporarily and selectively suspending some of the inhibitory signals generated by the output structures. At the same time that inhibition is released on the selected behaviour it is increased on non-selected behaviours. A further key feature of the model is the presence of positive feedback loops that allow a selected behaviour to enhance its own tendency to be selected. This feature endows the model with the capacity for behavioural persistence that helps it to avoid situations of rapid switching ('dithering'). At the implementation level, our computational model of basal ganglia replicates the functional anatomy of these nuclei (i.e., their inter-connectivity and connection type), and uses leaky integrators to simulate the input/output properties of target neural populations. In the vertebrate brain, the selection functions of the basal ganglia are modulated by tonic (slow-varying) changes in the level of the neuro-modulator *dopamine*. Pathological variations in dopamine occur in certain brain disorders such as Parkinson's disease and can significantly disrupt selection function in the basal ganglia. In the experiments described below, this parameter was therefore held at a fixed, intermediate level suitable to ensure normal selection functionality.

Figure 7 shows how the model basal ganglia is embedded in the robot control architecture. Our target foraging activity is modelled on observations of the behaviour of hungry rats placed in an arena containing a small, centrally located dish of food. Initially these animals are fearful, preferring walls and corners to the centre of the arena; later they will traverse the arena to collect food from the dish that is then consumed in a corner.

The 'salience' (propensity for selection) of each of the five action sub-systems in the model is tuned to provide appropriate behavioural selections that simulate the avoidance-related and food-acquisition-related behaviour observed in these animals. More specifically, at any given

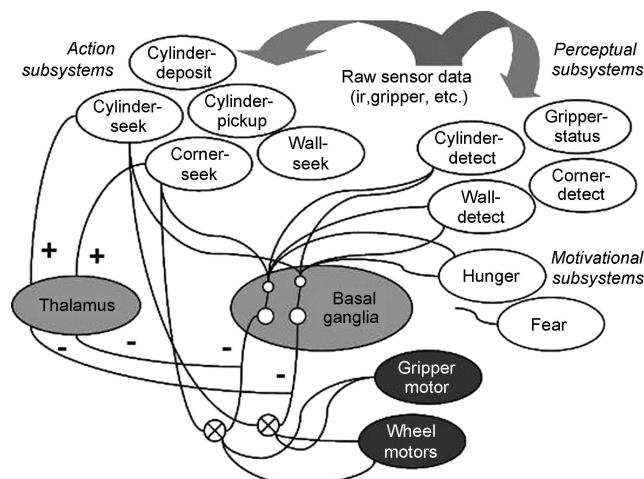


Figure 7 The basal ganglia compute the 'salience' for each action sub-system using inputs from perceptual, motivational, and action sub-systems. The selector resolves the competition and disinhibits relevant motor outputs (to gripper/wheel motors). Positive feedback is provided to the winning behaviour via a biomimetic loop modelled on connections from the basal ganglia to the cortex via the thalamus. Note that only a sub-set of actual connections are shown.

time, sub-system salience depends on the values of a number of *extrinsic* and *intrinsic* variables.

Extrinsic variables are calculated by *perceptual sub-systems* that use raw sensory signals from the robot to compute bi-polar values. These are indicated by the sub-systems *wall-detect*, *cylinder-detect*, *corner-detect* and *gripper-status* in the same figure. Thus, at each time-step the output of each perceptual subsystem is a bi-polar value indicating the presence (+1) or absence (-1) of the relevant target feature (e.g., a cylinder, a wall, a corner, or an object in the gripper) as determined by processing relevant sensory signals.

Intrinsic variables are computed by *motivational sub-systems* and are functions of recent experience and internal state. In our model these loosely approximate the drives of 'fear' (starts high and is reduced by time spent exploring the arena), and 'hunger' (gradually increases with time and is reduced when 'food' is 'consumed' – i.e., a cylinder is deposited in a corner of the arena). In the implementation of the model, the output of each motivational sub-system is a single scalar value in the range (0–1) which gradually decreases (fear) or increases (hunger) over time. Hunger is also reduced by a fixed amount when a cylinder is deposited in a corner of the arena.

Each action sub-system is also capable of generating an intrinsic variable, termed a 'busy signal' that can enhance its own salience during critical phases of activity. This signal is set to zero by default but can be given a value of +1 by the motor component of an active sub-system. A final contribution to the salience calculation is provided by feedback from within the basal ganglia circuit. Specifically,

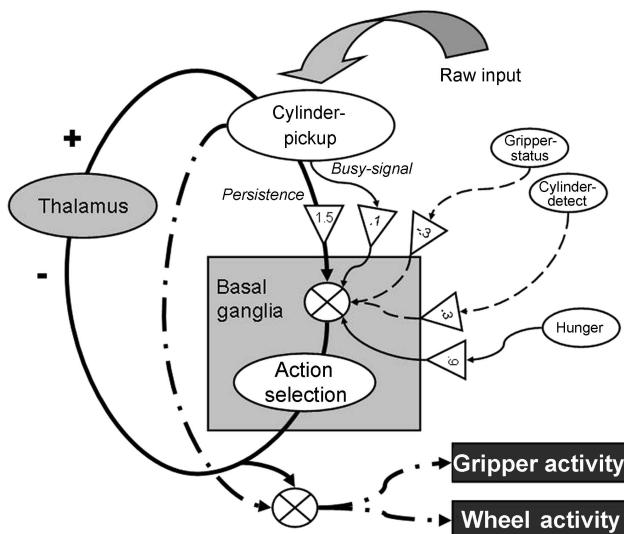


Figure 8 Connections in the path of only one channel for the cylinder-pickup sub-system with the basal ganglia are shown. Action selection may occur if the level of salience is enough to produce motor activity. The salience is input to the basal ganglia as the sum of the weighted persistence, busy signal, the perceptual sub-systems and the motivational sub-systems. A thick straight line in the figure represents the positive loop through the thalamus, the thick point-dashed line is the motor activity to the wheels and gripper. A thin straight line from the sub-systems represents a binary signal and a dashed one a bipolar signal.

in addition to gating the motor system, the output of the basal ganglia is known to complete a positive feedback loop, via the thalamus, to the cortical systems that bid for action selection (see Figure 8). In the current model, this *persistence* signal is differentially weighted for each behaviour. Note that the busy and persistence signals will be non-zero only for a currently active sub-system – this gives the currently selected behaviour a significant ‘edge’ over its rivals and so acts to reduce excessive switching (or ‘dithering’) between closely matched alternatives.

The salience or the urgency of each behaviour is found at each time-step by calculating the dot product of a vector containing the extrinsic and intrinsic variables, together with the busy and persistence signals, with a vector of hand-picked weights. For the experiments described below, the weights used for each behaviour were as follows: *cylinder-seek*: persistence +1.0, busy-signal 0.0, gripper-status -0.3, cylinder-detect -0.3, hunger +0.9; *cylinder-pickup*: persistence +1.5, busy-signal +0.1, gripper-status -0.3, cylinder-detect +0.3, hunger +0.9; *wall-seek*: persistence +1.0, busy-signal 0.0, wall-detect -0.3, gripper-status +0.3, fear +0.9; *corner-seek*: persistence +1.9, busy-signal 0.0, wall-detect +0.3, gripper-status +0.3, fear +0.8; and *cylinder-deposit*: persistence +1.5, busy-signal +0.1, gripper-status +0.75, corner-detect +0.75, hunger +0.5. For a further explanation of these parameters please refer to the work of Montes (Montes Gonzalez *et al.* 2000).

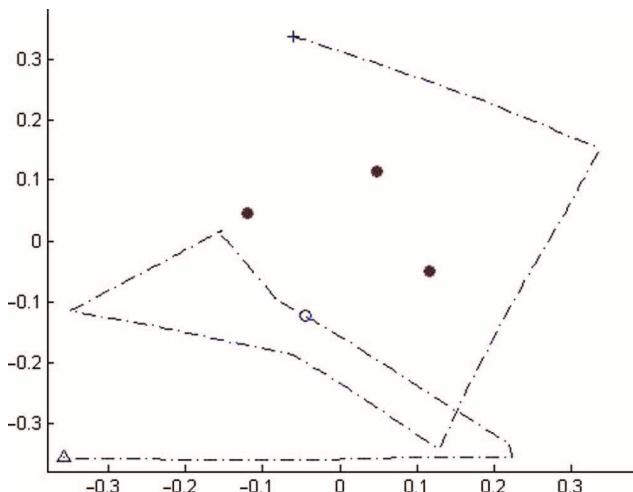


Figure 9 The Khepera robot is set in the middle of an arena with four wooden cylinders. A cross indicates the initiation of one foraging bout, and the triangle shows its termination. The labels in the figure show the selection of one of five possible behaviours.

The basal ganglia model runs within a main loop in which sensor readings are updated and motor commands are sent. At each time-step, the basal ganglia model computes salience of each behaviour as specified above then resolves the competition between behaviours disinhibiting the winning sub-system (usually the most salient).

EXPERIMENTS AND RESULTS

The foraging task was set in a four-walled square arena. A standard RS232 interface was used to connect the robot to the host computer. The simulated ‘food’ consisted of four wooden-cylinders that were initially placed in the middle of the arena. In each ‘foraging’ bout *cylinder-seek* is used to travel the arena searching for cylinders while avoiding obstacles, *cylinder-pickup* then grasps the cylinder, *wall-seek* is employed next to locate the nearest wall whilst avoiding obstacles, *corner-seek* then causes the robot to follow alongside a wall, finally, when a corner is reached, *cylinder-deposit* lowers and opens the gripper releasing the cylinder. A typical foraging sequence (one cylinder) is shown in Figure 9.

A behavioural ‘ethogram’ (Figure 10), and some related statistics (Table 1), show four sequential bouts of foraging preceded by a bout of avoidance behaviour (*wall-seek* succeeded by *corner-seek*) that is triggered by the initially high level of fear. Figure 10 also indicates the level of simulated ‘fear’ and ‘hunger’ to demonstrate the effect of these variables on selection. Salience levels and basal ganglia output are presented for each of the five action sub-systems in Figure 11. Whereas the salience of each sub-system (the weighted sum of intrinsic and extrinsic variables) can take negative or positive values, the basal ganglia output takes values between zero (selection) and one (not selected). The

Table 1 Elementary statistics for a representative run of the robot basal ganglia

Behavioral elements	Freq	Latency	TotDur	TotDur%	Mean	StdDev	StdErr	MinDur	MaxDur
1. Cylinder-seek	16.00	3.23	24.61	65.87	1.54	0.80	0.20	0.02	2.39
2. Cylinder-pickup	4.00	6.47	1.75	4.68	0.44	0.03	0.01	0.41	0.47
3. Wall-seek	103.00	0.06	3.25	8.69	0.03	0.15	0.01	0.00	1.28
4. Corner-seek	8.00	1.05	1.99	5.31	0.25	0.65	0.23	0.02	1.86
5. Cylinder-deposit	5.00	8.64	1.45	3.89	0.29	0.17	0.08	0.02	0.45
6. None	136.00	0.00	4.31	11.55	0.03	0.07	0.01	0.00	0.55
Total	272.00	0.00	37.36	100.00	0.14	0.43	0.03	0.00	2.39

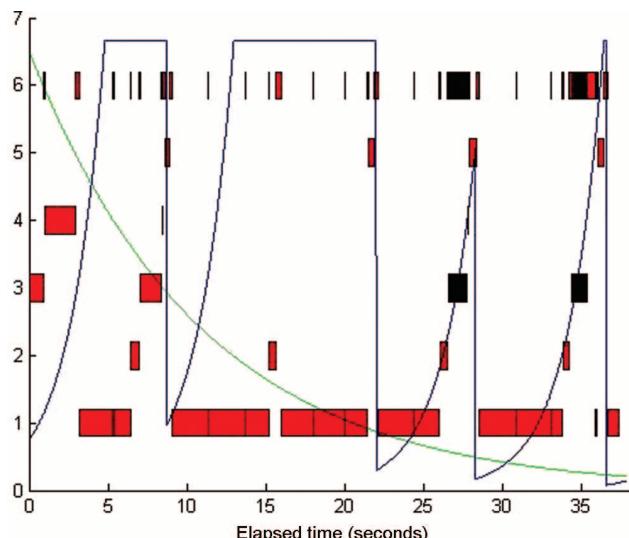


Figure 10 Ethogram for a typical run of the foraging task. The behaviours are coded as 1. cylinder-seek, 2. cylinder-pickup, 3. wall-seek, 4. corner-seek, 5. cylinder-deposit, 6. no action selected. Lines indicate the motivations ‘fear’ and ‘hunger’ (both varying between zero and one). Fear starts high when compared to hunger. Notice that because of the image resolution ‘no action selected’ seems to occur when selection actually does not occur. (figure available in colour on-line)

figure shows that winning behaviours are cleanly selected and losers fully inhibited throughout this trial.

DISCUSSION

In this work we have described and tested an action selection model that mimics the function of the basal ganglia in vertebrates. In order to do this, we have used both biomimetic and engineered components. In modelling brain systems we have naturally made many abstractions away from the biological details. For instance, we have vastly simplified the sensory information that the cerebral cortex receives via the thalamus in analogizing this to the information obtained from the sensors in the robot. Also, we have not attempted to model the neural control of movement in processing the motor commands that are sent to the robot. On the other hand, our model of the basal ganglia,

although high-level, does respect many neurophysiological constraints obtained from research on these structures. For instance, the intrinsic connectivity of the basal ganglia nuclei has been preserved as this has been shown to play an important role in the selection functionality of this brain circuit (Gurney *et al.* 2001a, 2001b).

The basal ganglia is one of two major brain systems involved in the ongoing control of movement, the other being the cerebellum (Houk *et al.* in press). Motor cortex sends information to both systems, and, in turn, both send back information to the motor cortex via the thalamus. The outputs of the basal ganglia and cerebellum have opposing effects in the sense that whereas the output of the cerebellum is excitatory, that of the basal ganglia is inhibitory. The role of the basal ganglia, as it has been assumed in this paper, is related to action selection. This includes both the initiation and maintenance of motor behaviour (Prescott *et al.* 2006). In contrast, the cerebellum is responsible for regulating the sequencing and duration of elementary movements. A simple gesture such as touching the tip of our nose requires not only that hand and arm movements are disinhibited by the basal ganglia, but also that the various segments of our hand and arm are smoothly and accurately deployed via regulation from the cerebellum. These two systems thus appear to play complementary roles and provide a balance in generating coordinated movement. Disturbances in either system frequently show up as movement disorders (Figure 12). At the moment we have not included yet a cerebellum component in our model but we expect its further development to proceed in that direction. One of the authors has also proposed that besides the action selection mechanism of the basal ganglia, an ‘action induction’ mechanism (Negrete-Martínez 2006) could also be present in the brain. Contrary to action selection, at a specific time, action induction provides a recommendation – from a fixed repertoire of which action could be executed.

In this work we have investigated the operation of the basal ganglia component using ‘normal’ levels of the simulated neuromodulator dopamine. However, previous computational models of the basal ganglia have shown that changes to simulated dopamine can have significant and interesting effects on selection (Gurney *et al.* 2001, Humphries and Gurney 2002). By embodying the basal ganglia in a robot we are able to directly observe the effects of dopamine modulation on behaviour, and thus, for

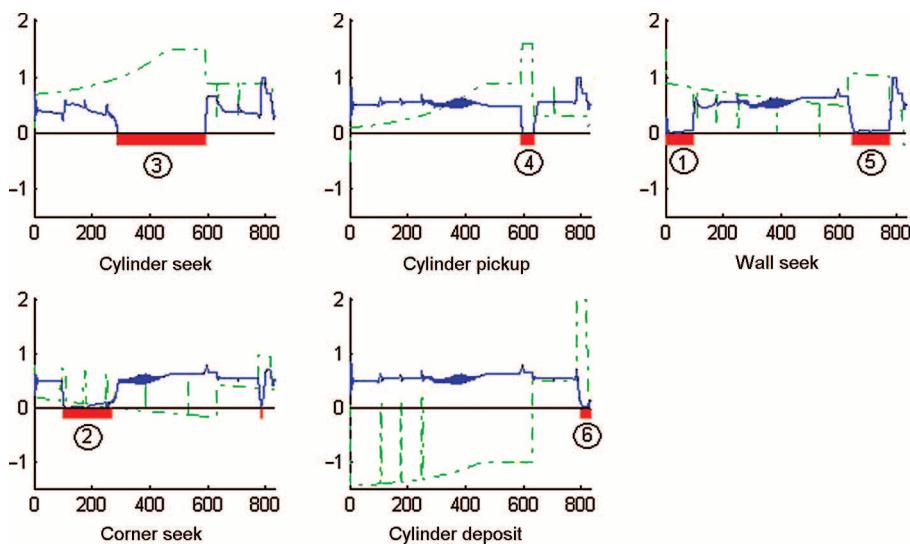


Figure 11 Basal ganglia activity is shown for a test period of 830 simulation steps, dashed (green) lines represent the salience for each of the five action sub-systems, whereas solid (blue) lines indicate the basal ganglia output that should be near zero activity for selection to occur. The bars below the x -axis indicate selection of that action sub-system. Numbers below the bars indicate the order in selection.

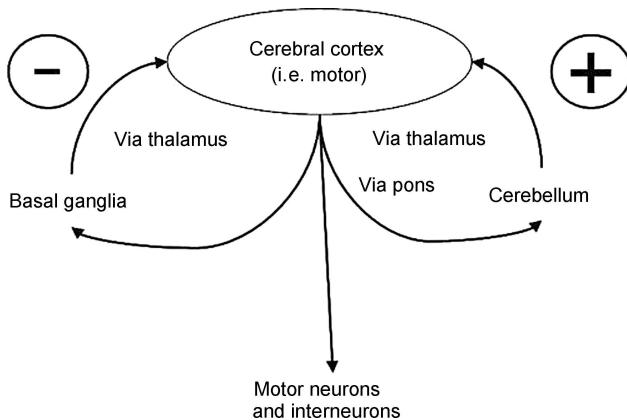


Figure 12 Excitatory and inhibitory nature of efferent connections from the basal ganglia, and the cerebellum, to the cortex via the thalamus.

instance, obtain a better understanding of movement related abnormalities seen in basal ganglia disorders involving dopamine (such as Parkinson's disease). We expect to investigate these ideas in future refinements to the model.

CONCLUSION

We have shown that the robot model of the basal ganglia provides effective selection when combined with non-homogenous behaviour patterns. The behaviours range from evolved and learned behaviours, through to algorithmic behaviours designed by hand. At present the parameters of the model that specify the salience calculations are hand-coded, these weights therefore constitute a further candidate for automated tuning which will be a focus of future research. Overall, this work paves the

way for the development of complete biological-inspired robot controllers that could be constructed with only minimal human intervention. We anticipate that the task of the designer could be limited to the specification of the decomposition of the task into component behaviours and their fitness functions, and of an overall metric for evaluating the performance of the complete controller.

ACKNOWLEDGMENTS

This work has been sponsored by CONACyT-MEXICO grant SEP-2004-C01-45726, and was also supported by a grant to Tony Prescott through the EU Framework 6 ICEA project.

REFERENCES

- Bajaj D, Ang MH Jr. 2000. An incremental approach in evolving robot behaviour. In *The Sixth International Conference on Control, Automation, Robotics and Vision (ICARCV'2000)*, Singapore.
- Balakrishnan K, Honavar V. 1995. Properties of genetic representation of neural architectures. In *The World Congress on Neural Networks (WCNN '95)*.
- Bryson J. 2000. Cross-paradigm analysis of autonomous agent architecture. *J Exp Theor Artificial Intell*, 12(2): 165–189.
- Cyberbotics (2006) Webots, Commercial Mobile Robot Simulation Software.
- Dean, J. 1998. Animats and what they can tell us. *Trends Cognitive Sci.*, 2(2): 60–67.
- Floreano D, Mondana F. 1994. Automatic creation of an autonomous agent: Genetic evolution of a neural-network driven robot. From Animals to Animats III. In *Proceedings of the Third International Conference on Simulation of Adaptive Behavior*, Cambridge, MA: MIT Press-Bradford Books.

- Gurney K, Prescott TJ, Redgrave P. 2001a. A computational model of action selection in the basal ganglia. I. A new functional anatomy. *Biol Cybernet*, 84(6): 401–410.
- Gurney K, Prescott TJ, Redgrave P. 2001b. A computational model of action selection in the basal ganglia. II. Analysis and simulation of behaviour. *Biol Cybernet*, 84(6): 411–423.
- Holland J. 1975. Adaptation in Natural and Artificial Systems, Ann Arbor: University of Michigan Press.
- Houk JC, Bastianen C, Fansler-Wald D, Fishbach A, Fraser D, Reber PJ, et al. (2007). Action selection and refinement in subcortical loops through the basal ganglia and cerebellum. *Philosophical Transactions of the Royal Society-B* 362: September issue.
- Humphries MD, Gurney KN. 2002. The role of intra-thalamic and thalamocortical circuits in action selection. *Network*, 13: 131–156.
- Maes P. 1995. Modelling adaptive autonomous agents. In Langton, CG (ed.), *Artificial Life: An Overview*. Cambridge, MA, MIT Press.
- Mcfarland D. 1993. Animal Behaviour, Harlow, Essex: Longman Scientific and Technical.
- Mondana F, Franzi E, Ienne P. 1993. Mobile robot miniaturisation: A tool for investigating in control algorithms. In *Proceedings of the 3rd International Symposium on Experimental Robotics*, Kyoto Japan, Springer, Verlag.
- Montes González F, Prescott TJ, Gurney K, Humphries M, Redgrave P. 2000. An embodied model of action selection mechanisms in the vertebrate brain. In Meyer, JA (ed.), *From Animals to Animats 6: Proceedings of the 6th International Conference on the Simulation of Adaptive Behavior*. Cambridge, MA: MIT Press.
- Montes González FM, Marín Hernández A, Ríos Figueroa H. 2006. An Effective Robotic Model of Action Selection. In Marín R et al. (eds.), CAEPIA 2005, LNAI 4177, pp. 123–132.
- Negrete-Martínez J. 2006. Self-inhibiting modules can self-organize as a brain of a robot: a conjecture. *Appl Bionics Biomech.*, 3(1): 23–27.
- Nolfi S. 1997. Evolving non-trivial behaviors on real robots: a garbage collection robot. *Robotics and Autonomous System*, 22: 187–198.
- Nolfi S, Floreano D. 2000. Evolutionary Robotics, Cambridge, MA: MIT Press.
- Prescott TJ, Montes González FM, Gurney K, Redgrave P, Humphries M. 2006. A robot model of the basal ganglia: behaviour and intrinsic processing, *Neural Networks*, 19: 31–61.
- Prescott TJ, Redgrave P, Gurney KN. 1999. Layered control architectures in robots and vertebrates. *Adaptive Behav.*, 7: 99–127.
- Prescott TJ. 2007. Forced moves or good tricks in design space? Landmarks in the evolution of neural mechanisms for action selection. *Adaptive Behavior*, 15.
- Prescott TJ, Bryson JJ, Seth A. 2007. Modelling natural action selection: Introduction to the theme Issue. *Philos Trans Royal Soc. B. Biol. Sci.*, Sep 29;362(1485): 1521–1529.
- Santos J, Duro R. 2005. Artificial Evolution and Autonomous Robotics (in Spanish), Ra-Ma Editorial Press.
- Wagner GP, Altenberg L. 1996. Perspective-complex adaptations and the evolution of evolvability. *Evolution*, 50(3): 967–976.

