
Book Section:
Prescott, T.J. orcid.org/0000-0003-4927-5390, Gurney, K.N. orcid.org/0000-0003-4771-728X, Montes-Gonzalez, F.M. et al. (2 more authors) (2002) The robot basal ganglia : action selection by an embedded model of the basal ganglia. In: Nicholson, L.F.B. and Faull, R.L.M., (eds.) The Basal Ganglia VII. Advances in Behavioral Biology (52). Springer , pp. 349-358. ISBN 9781461352075

https://doi.org/10.1007/978-1-4615-0715-4_35

To appear in Nicholson, L. F. B. & Faull, R. L. M. (eds.) *Basal Ganglia VII,* Plenum Press, 2002

THE ROBOT BASAL GANGLIA:

Action selection by an embedded model of the basal ganglia

Tony J. Prescott, Kevin Gurney, Fernando Montes-Gonzalez, Mark Humphries, and Peter Redgrave*

1. SUMMARY

Action selection is the task of resolving conflicts between multiple sensorimotor systems seeking access to the final common motor path. Recently,[1, 2] we proposed that the basal ganglia may act to provide a biological solution to the problem of selection. To test this notion we have implemented a high level computational model of intrinsic basal ganglia circuitry and its interactions with simulated thalamocortical connections.[3, 4] The computational model was then exposed to the rigors of 'real world' action selection by embedding it within the control architecture of a small mobile robot.[5] In a mock foraging task, the robot was required to select appropriate actions under changing sensory and motivational conditions, thereby generating sequences of integrated behavior. Our results demonstrate: (i) the computational model of basal ganglia switches effectively between competing channels depending on the dynamics of relative input 'salience'; (ii) its performance is enhanced by inclusion of anatomically inspired thalamocortical circuitry; (iii) in the robot, the model demonstrates appropriate and clean switching between different actions and is able to generate coherent sequences of behavior.

2. BACKGROUND

One of the problems in determining basal ganglia function is the apparent diversity of processes in which it takes part. Thus, the basal ganglia have been implicated in perception and cognition, working memory, and many aspects of motor function. However, one recurring theme in the literature,[1, 6] is that the basal ganglia are implicated in the problem of *action selection*. Stated informally, this is the problem of deciding

* Department of Psychology, University of Sheffield, Western Bank, Sheffield, S10 2TP, United Kingdom.

'what to do next', and concerns the resolution of conflicts between functional units that are physically separated within the brain but are in competition for behavioral expression. Our recent work[1, 2] has developed the idea of selection as a major unifying hypothesis of basal ganglia function, showing how it relates to known anatomy and physiology, and meets several high-level computational requirements. It is this premise that underlies all of the work reviewed in this chapter.

Cognitive psychologists have discovered that when an observable system has more than three interacting parts, it becomes very difficult for human minds to predict accurately how that system will change over time. Given the complexity of intrinsic basal ganglia interactions, with feedback loops both between and within the different nuclei, only a very limited understanding can be gleaned from informal, box-and-arrow-style, models. To achieve a deeper insight demands a computational approach involving quantitatively specified models, delimited by anatomical and physiological constraints. In our view, however, this standard *computational neuroscience* approach, in which specific brain systems are modelled in isolation of any wider context, still leaves many important questions unanswered. First, we are left wondering as to the best interpretation of the inputs and outputs of the model—we might choose to think of inputs as, say, 'sensory' signals, or of outputs as 'motor' signals, but such assignments are essentially ungrounded. Second, without locating a model within any wider context, we are unable to judge whether it can fulfil its hypothesised functional role within a more fully specified control architecture. Third, without any linkage to sensory and motor systems we may question whether a model could cope with noisy or ambiguous sensory data, or as part of a system challenged with co-ordinating the movements of real effector systems. Finally, without the context of multiple demands, such as the need to maintain physical integrity, avoid threats, discover and exploit resources, we will be unsure as to whether a model can meet some of the basic requirements for adaptive behavior. In order to begin to address these questions, our work has therefore taken the further step of studying computational models of the basal ganglia that are embedded within a robot control architecture, situated within a real environment, and faced with the task of co-ordinating a number of simple, but 'life-like', behaviors. In the remainder of this chapter we first review some of the findings and conclusions of our work on various computational neuroscience models of the basal ganglia, and then describe our recent efforts to develop and evaluate this embedded, robotic model of the basal ganglia.

3. COMPUTATIONAL NEUROSCIENCE MODELS OF THE BASAL GANGLIA

3.1. A 'System Level' Model of Intrinsic Basal Ganglia Processing

Just as any preliminary anatomical view of a neural system is best carried out at 'low magnification' in order to discover general patterns of connectivity, we argue that it is best, in the first instance, to adopt a 'system level' approach to brain modelling. In this coarse-grained view we are interested in the ensemble behavior of cells so that each relevant population is represented *in toto* by the smallest functional unit in the model. Armed with the hypothesis that the basal ganglia constitute a selection mechanism, we have therefore developed a system level model of the basal ganglia[3] constrained by the

known anatomy, and in which neural populations are represented by simple *leaky integrator units*[*].

In devising our model we have supposed that the brain is processing a large number of sensory and cognitive streams or *channels*, each one potentially carrying a request for action to be taken. For effective behavior to occur, the majority of these requests must be suppressed to allow the expression of only a limited number (perhaps just one). This channel-based scheme is consistent with evidence that basal ganglia input occurs via a series of afferent parallel processing streams that, at least where motor areas are concerned, display a somatotopic organisation. The action selection hypothesis of the basal ganglia suggests that the activity of cell populations in the striatum and subthalamic nucleus (STN) encodes the *salience*, or propensity for selection, of candidate actions. At the same time, the basal ganglia output structures, substantia nigra (SNr) and entopenduncular nucleus (EP) (in rats), are viewed as *gating* candidate actions via a reduction in their inhibitory output for 'winning' channels. When considered in isolation of the wider brain architecture, this action selection thesis is best restated in terms of the context-neutral problem of 'signal selection', in other words, the proposal is that large signal inputs at striatum and STN select for low signal outputs at EP and SNr.

From a signal selection perspective multiple mechanisms within the basal ganglia appear to be suitably configured to resolve conflicts between competing channels and provide the required clean and rapid switching between winners.[1, 6] First, there is the intrinsic property of striatal projection cells such that, at any given moment, a majority of cells are in an inactive 'down-state', and can only be triggered into an active 'up-state' (where they can fire action potentials) by a significant amount of coincident input. This bistable behavior could act as a first pass filter to exclude weakly supported 'requests'. Second, local inhibition within the striatum can selectively enhance the activity of the most active channels. Third, dopamine modulation may affect selectivity by differentially gating cortico-striatal inputs with D1 and D2 -type receptors (associated with excitation and inhibition respectively). All the above mechanisms act at the cellular or local circuit level. However, if the primary role of the basal ganglia is selection, we would also expect there to be mechanisms that act globally between nuclei. Computational theory suggests that a *feed-forward, off-centre, on-surround network* is one such candidate mechanism. In the basal ganglia, this type of selection circuit appears to be implemented by a combination of focused striatal inhibition of the output nuclei (the off-centre) and diffuse STN excitation of the same (the on-surround).[6] On closer examination, however, it appears that there are actually two such feed-forward networks in the basal ganglia intrinsic circuitry (see figure 1a,b), differentiated by the projection targets of two sub-populations of medium spiny cells. One instantiation makes use of EP/SNr as its 'output layer'; since this is clearly consistent with our signal selection hypothesis for the basal ganglia we designate this circuit the *selection pathway*. However, there is also a second implementation of the feed-forward architecture whose target is the globus pallidus (GP). The efferent connections of the GP are confined to other basal ganglia nuclei, thus it is not immediately clear in what sense this second implementation can contribute to the

[*] A simple model 'neuron', described by Arbib and others (see, e.g. *The Handbook of Brain Theory and Neural Networks*, p. 4-11, MIT Press, Cambridge, MA, 1995), that respects several important elements of neural computation such as the concept of a dynamic membrane potential.

overall selection task. This question can be resolved by supposing that this second subsystem forms a *control pathway* that functions to regulate the properties of the main selection mechanism. The control signals emanating from GP are evident when the two subsystems are combined to give the overall functional architecture shown in Figure 1c.

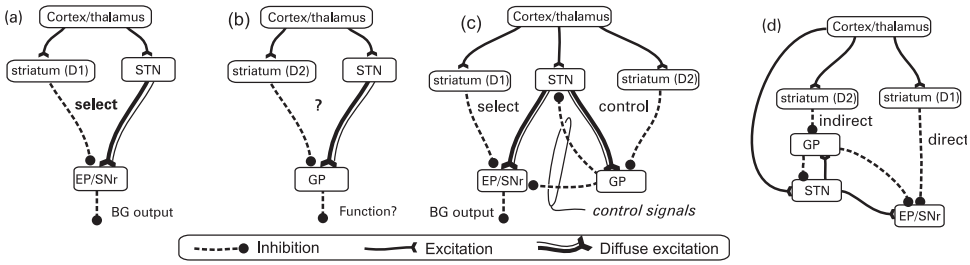

Figure 1. The new functional architecture (a–c) and the prevailing model (d).

Comparison of the new architecture with the prevailing direct/indirect model of the basal ganglia[7] (figure 1d) shows that the two schemes are structurally quite different. In the prevailing model, no mention is made of the diffuse STN projections underlying selection, while the new model emphasises their importance. Further, as originally proposed, the prevailing model does not include the GP to EP/SNr pathway and, as Smith et al.[8] have noted, incorporating this would introduce a multiplicity of 'indirect pathways' whose function is unclear.

In our computational neuroscience work we have operationalized the above circuit (figure 1c) as a multi-channel model whereby, for every basal ganglia nucleus, the neural population encoding each channel is simulated by a suitably configured leaky integrator unit. Analytical and simulation studies conducted with this model[3] demonstrate that it has the capacity to support effective switching between multiple competitors. In simulation, two or more channels of the model are provided with afferent input in the form of hand-crafted signals of different amplitude. Results show that the largest signal input always generates the smallest signal output (thus showing signal selection), and that the system rapidly switches from a currently selected channel to a competing channel that has a larger input. As shown in Figure 2, we have also been able to generate signal characteristics in the component circuits of our basal ganglia model (the insets in the top right of each graph) that follow similar temporal patterns to single-unit recordings of neural firing (larger graphs) in GP and SNr.

In addition to validating the model by demonstrating the desired computational properties (appropriate signal selection) and by showing similarities between model outputs and neurobiological data, this work also generated two novel functional hypotheses. First, our analysis of basal ganglia intrinsic circuitry concluded that the GP is a source of control signals for the rest of the basal ganglia (this view has also received support from recent data[8] indicating that GP innervates striatum thus ensuring it has access to all basal ganglia nuclei). In this context, our computational model suggests a specific role for GP control signals in relation to the STN whereby *negative feedback, via*

GP, functions to automatically scale the excitatory output of the STN with the number of active channels. A second hypothesis concerns the synergistic action of dopamine in the 'control' and 'selection' pathways of the basal ganglia. We have proposed that *a key function of dopamine in this mechanism is to regulate the ease of selection*. Increases in simulated dopamine seem to promote 'promiscuous' selection in which channels are more easily disinhibited, whilst reduced dopamine results in a 'stiffer' competition in which there are fewer winners and higher levels of general target inhibition.

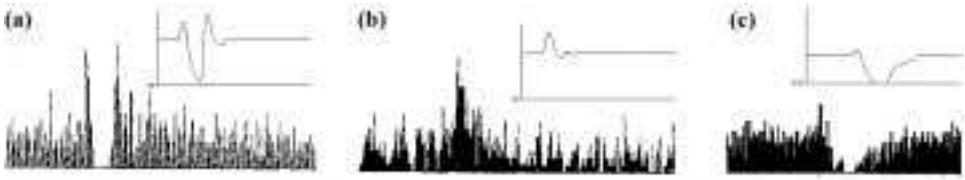

Figure 2. Model output compared with recordings in the GP (a, b) and SNr (c). Histograms from (a, b) Ryan, L.J. & Clark, K.B. *Experimental Brain Research* **86**, 641-651 (1991), by permission of Springer Verlag; (c) Shultz, W. *Journal of Neurophysiology* **55**, 660-677 (1986), permission the American Physiological Society.

3.2. Extensions of the Intrinsic Model

We have extended our model of basal ganglia intrinsic circuitry to include further anatomical constraints and, at each stage, the extended model has displayed improved selection properties, confirming the underlying hypothesis of selection. Thus, the inclusion of complete basal ganglia-thalamo-cortical loops[4] led to the hypothesis that the thalamic complex (ventro-lateral thalamus and thalamic-reticular nucleus) is a further selection circuit that promotes several desirable selection features including cleaner switching between channels of closely matched salience. In recent unpublished work we have shown that the intrinsic model can also accommodate new data on striato-pallidal projections, an extension that also appears to enhance the selectivity of the system.

3.3. Lower Level Models

We have also made two initial studies that demonstrate the extra explanatory power of models that incorporate lower level data, and which allow for the patterning of signals in trains of action potential ('spikes'). First,[9] we modelled the oscillatory pacemaker discovered by Plenz and Kitai[10] in an organotypic co-culture of GP and STN. This work demonstrated the importance of modelling non-linear dendritic arithmetic and intrinsic membrane currents. Most recently (unpublished work) we have re-implemented the thalamo-cortical-basal ganglia model using populations of simple spiking neurons. This study showed that pairs of closely matched basal ganglia outputs (derived from closely matched saliences) are resolved into well-defined bouts of individual channel selection. This finding seems consistent with the rapid change in cognitive set seen in schizophrenia, and the patterns of increased behavior switching seen in children with attention deficit and hyperactivity disorder.

4. A ROBOT MODEL OF ACTION SELECTION BY THE BASAL GANGLIA

The modelling work considered above serves to demonstrate signal selection by the basal ganglia rather than action selection *per se*. To show that the basal ganglia model is able to act as an effective action selection device we have argued that it needs to be embedded in a real-time sensorimotor interaction with the physical world. An important goal has therefore been to construct an embedded basal ganglia model in which selection occurs between multiple, physically-realized behaviors in a mobile robot. To produce an embedded model of this type, simulations of specific brain structures and pathways must be interfaced with the additional control system components needed to create a working robot control architecture. We believe that it is not necessary for all components of this full architecture to be directly modelled on the nervous system. Instead, we have focused on building interfaces to our model neural systems that provide appropriate input signals and that can convert the outputs of our neurobiological models into signals that are useful for robot control.

Our current program of research has been inspired by observing the behavior-switching of adult rats placed in an unfamiliar rectangular arena containing a centrally-located dish of food pellets. The initial behavior of such animals is typically exploratory and defensive. Specifically, experimental animals tend to stay close to the arena walls (thigmotaxis), show a strong preference for the corners of the arena, and show little or no visible interest in food consumption. As the animal becomes more accustomed to the novel environment, hunger-related behaviors become more apparent, a common behavior on locating the food dish being to carry food-pellets back to one of the corners of the arena to be consumed. Our initial efforts to create an embedded basal ganglia model have focused on producing a similar, if much simplified, problem setting for a small mobile robot. Our wheeled robot, which possesses a gripper-arm and a ring of infra-red distance sensors, is placed in a square, walled arena in which a number of small cylindrical objects are also placed. The cylinders substitute for food pellets, so the collection and consumption of food is modelled by collecting cylinders and depositing them in the corners of the arena. The control architecture of the robot includes five sensorimotor sub-systems, or *behaviors*, which it can switch between at any time. These are: searching for cylinders (*cylinder-seek*), picking up a cylinder (*cylinder-pickup*), looking for a wall (*wall-seek*), looking for a corner (*corner-seek*), and depositing the cylinder in a corner (*cylinder-deposit*). Each behavior operates independently to generate a continuous stream of motor signals that are directed toward the motor systems. So, for instance, *cylinder-seek* will use infra-red signals to judge the direction toward the nearest cylinder and will generate a motor signal that specifies movement in that direction. The robot, and some of the elements of its behavioral repertoire are shown in Figure 3.

The task for the embedded basal ganglia model is to arbitrate at each moment in time between the five available behaviors and to generate a pattern of action selection over time that results in coherent sequences of activity. For these experiments we used the 'system level' simulation of basal ganglia intrinsic circuits (section 3.1), extended (as described in section 3.2) to include models of the ventral thalamus and thalamic reticular nucleus. The following provides a brief account of the mechanisms needed to embed this model basal ganglia within a wider robot architecture, for further details see Montes Gonzalez et al.[5]

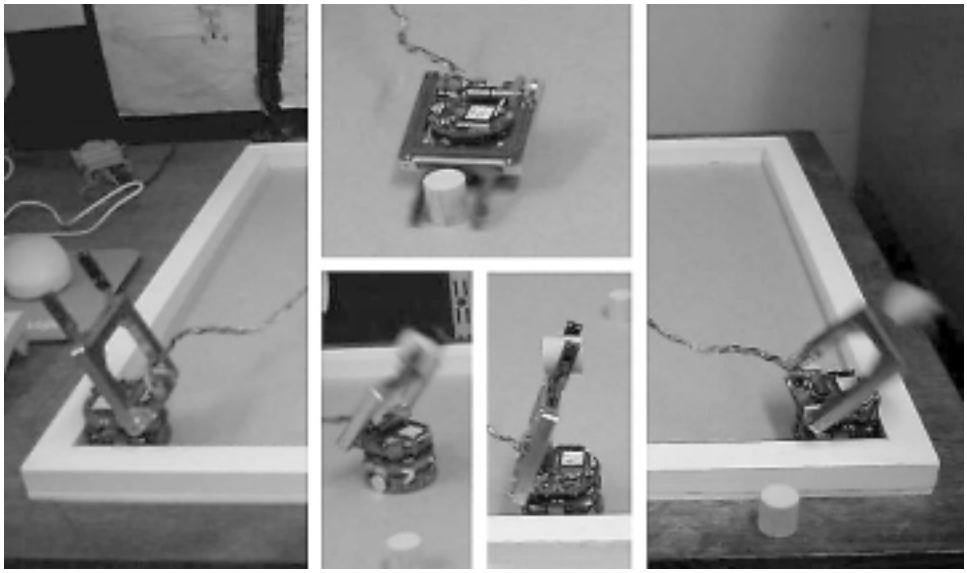

Figure 3. The mobile robot used to evaluate the embedded basal ganglia model. The photographs show various elements of the robot's behavior relating to finding, carrying, and depositing cylinders; finding walls; and following walls to find corners. See text for further explanation.

4.1. Determining Salience

Striatal medium spiny neurons have large dendritic trees and may receive input from many different regions of the brain. This can be viewed as a mechanism for assimilating relevant signals from a variety of sources that can then be integrated to generate a measure of salience for the activity associated with a specific basal ganglia channel.[1] In the embedded model, this mechanism is approximated by calculating the salience of each behavior as a weighted sum of relevant system variables including *perceptual*, *motivational*, *positive feedback*, and *efference copy* signals. Each of these sources of salience input is briefly described below.

Perceptual inputs to an action selection mechanism should indicate the presence (or absence) in the immediate environment of the relevant affordances for different behaviors. In the robot model, perceptual inputs are computed from the raw sensory data provided by the infra-red and gripper-arm sensors, and include signals indicating the presence/absence of a nearby wall, corner, or cylinder, or of an object in the robot gripper.

An action selection mechanism also requires information about intrinsic state, indicating, for example, the current level of energy reserves. In the current model two simple intrinsic drives, loosely analogous to 'hunger' and 'fear', are calculated. 'Fear' is a function of exposure to the environment and is reduced by time spent exploring the environment, whilst 'hunger' gradually increases with time and is reduced when cylinders are deposited in the corners of the arena.

Positive feedback has been widely discussed in the biological literature as an effective means for avoiding unnecessary and rapid switching between alternative behaviors (dithering). We have recently proposed[1] that basal ganglia-thalamo-cortical loops may act to provide a positive feedback pathway that can maintain an appropriate level of salience in a selected behavior. Specifically, ventral thalamic neurons are tonically active and have reciprocal excitatory connections with cortical neurons that provide input to the basal ganglia. Disinhibition of a thalamic channel by the basal ganglia should therefore increase the excitation directed at the corresponding population of basal ganglia input neurons. This positive feedback loop has been explicitly simulated in the robot model and provides an additional salience input to the basal ganglia.

Much of the input to the basal ganglia comprises collateral branches from fibres projecting to motor regions of the brainstem and spinal cord. Furthermore, electrophysiological studies have noted that activity changes in the BG often occur slightly after the beginning of EMG activity. One possibility, we propose, is that efference copy signals to the basal ganglia may be important for accurate timing of behavioral switching. In our embedded basal ganglia model, each behavior is capable of generating an intrinsic 'busy' signal that can contribute to its own salience. Experiments with the robot have shown that such signals may be important for controlling the maintenance and termination of a selected behavior.

4.2. Gating Motor Output

The output of the model basal ganglia (EP/SNr) gates the motor signals produced by each behavior by reducing or increasing the inhibition on the corresponding motor pathway. In the embedded model, after gating by the basal ganglia, all motor signals are summed and the resulting aggregate signal used to control the wheel and gripper motors in the next time-step. Motor signals are expressed as a desired change in the current state, that is, as changes in wheel-speeds, gripper-arm elevation, and gripper-mouth position. This means that in the event of full basal ganglia inhibition of all channels, the aggregate command will have zero value and the current robot state will not change (i.e. it will freeze in its current position). In the event that one or more channels is partially (but not fully) disinhibited, the robot will act but its movements may be slowed by the resulting reduction in the size of the motor signals. Note that the motor signal generated by any losing behavior that is not fully inhibited by the basal ganglia will be combined with that of the winner. This mechanism allows for the possibility of distortion (the robot tries to do two things at once) in the event of ineffective suppression of competitors by the basal ganglia.

4.3. Results

Our initial results[5] have demonstrated appropriate and clean switching by the embodied basal ganglia model. In other words, given suitable salience weightings, the robot selects appropriate actions for the different circumstances it meets and generates sequences of integrated behavior over time. Some of the intrinsic activity of the robot model is illustrated in Figure 4 where the net salience input and corresponding basal ganglia output is shown for each of the five behaviors recorded over a short test period.

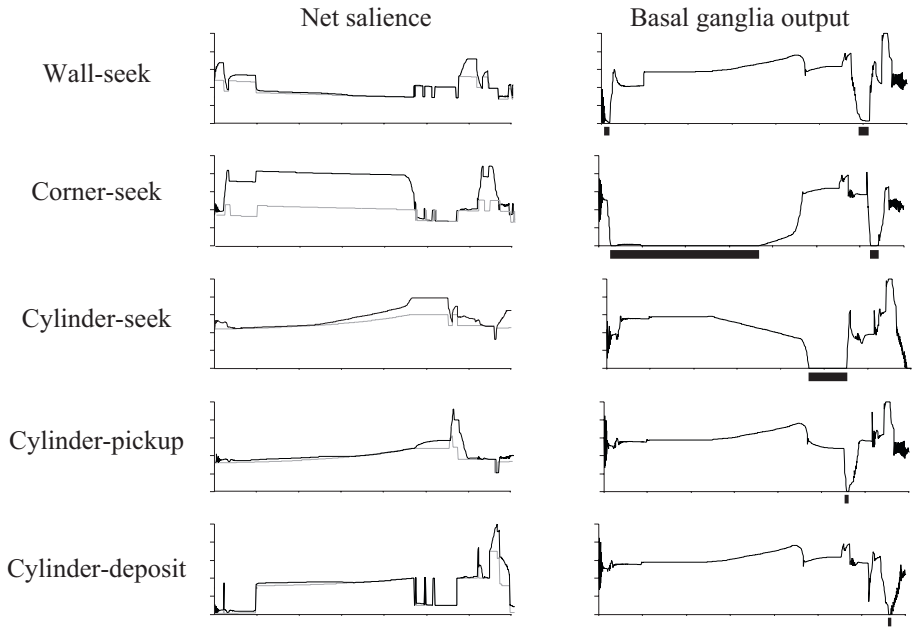

Figure 4. Basal ganglia activity relating to each of the five behaviors for a test period of approximately two minutes. Left: Net salience of each behavior. The lower *dotted line* shows the element of the salience due to perceptual and motivational inputs alone. The higher *filled line* shows the total salience including the effect of positive feedback signals. Right: Basal ganglia output. The bars beneath the graphs indicate the period during which each behavior is selected. The y-axis corresponds to nominal units of increasing signal strength, the x-axis increasing time. (For details see Montes Gonzalez et al.[7]).

Note that for effective action selection a winning channel should have near zero output, and a losing one relatively high output. The behavior of the robot during this brief episode was as follows. The robot begins the experiment with a high level of simulated 'fear' that results in higher salience for *wall-seek* (thigmotaxis) than for other actions. After quickly finding a wall, *corner-seek,* whose salience is increased by the affordance of a wall, becomes more salient and is selected. The robot follows the wall until a corner is reached then waits while increasing hunger drives up the salience of *cylinder-seek*. Once the salience for *corner-seek* falls below that of *cylinder-seek* the robot switches to looking for cylinders. When it finds a cylinder, *cylinder-pickup* is selected, followed by *wall-seek* (this time carrying a cylinder), *corner-seek*, and finally *cylinder-deposit* (each behavior triggered by relevant perceptual affordances). The graphs of net salience show that in some circumstances, two or more behaviors can have similar levels of urgency thus generating a requirement for action selection by the basal ganglia. The gap between the dotted and filled lines on these graphs shows the difference that thalamic feedback makes to net salience—boosting a selected behavior so that it is less vulnerable to interrupts. The graphs of basal ganglia output show clean switching between action sub-systems with selected sub-systems more-or-less fully disinhibited.

4.4. Dopamine Regulation of Behavior Switching

There is an important strand of research suggesting that tonic levels of dopamine neurotransmission play an important role in behavior switching.[1, 11] Depending on the site and nature of the intervention, these effects include changes in the dominance relations between behaviors, reductions or increases in rates of switching, changes in the variability of behavior, and failure to complete behaviors. This data indicates that mild to moderate increases in dopaminergic neurotransmission tend to facilitate switching while comparable reductions in transmission may retard switching. This proposal is consistent with our findings with the intrinsic basal ganglia model (see 3.1) where simulated dopamine modulation was shown to regulate signal selection by differentially biasing the 'selection' and 'control' pathways. With the robot model we are able to observe the consequences of this modulation for actual behavioral sequences. For instance, lowering simulated dopamine below the normal tonic level causes the robot difficulty in initiating and in completing behaviors. On the other hand, an increase in simulated dopamine can cause the robot to select two behaviors simultaneously resulting in an inappropriate mixture of two activity patterns (such as repeated lifting and lowering of the gripper-arm whilst exploring the arena looking for cylinders). These behavioral outcomes show interesting similarities to some of the effects of dopaminergic treatments in animals, and to symptoms of dopamine-related human disorders such as Parkinson's disease and Tourette's syndrome. This suggests that, in addition to serving as a test-bed for the action selection hypothesis of basal ganglia function, the robot model could, in future, also provide a useful vehicle for investigating theories of basal ganglia dysfunction.

5. REFERENCES

1. P. Redgrave, T. J. Prescott, and K. Gurney, The basal ganglia: A vertebrate solution to the selection problem, *Neuroscience* **89**, 1009-1023 (1999).
2. T. J. Prescott, P. Redgrave, and K. Gurney, Layered control architectures in robots and vertebrates, *Adaptive Behavior* **7**(1), 99-127 (1999).
3. K. Gurney, T. J. Prescott, and P. Redgrave, A computational model of action selection in the basal ganglia: I & II. *Biological Cybernetics* **84**, 401-423 (2001).
4. M. D. Humphries and K. Gurney, A computational model of action selection in the basal ganglia: Thalamic and cortical interactions, *European Journal of Neuroscience* **12**, p. 508 (2000).
5. F. Montes Gonzalez, T. J. Prescott, K. Gurney, M. Humphries, and P. Redgrave, in: *From Animals to Animats 6: Proc. of the 6th Int. Conf. on the Simulation of Adaptive Behavior*, edited by J.A. Meyer (MIT Press, Cambridge, MA, 2000) , pp. 157-166.
6. J. W. Mink, The basal ganglia: Focused selection and inhibition of competing motor programs, *Progress In Neurobiology* **50**, 381-425 (1996).
7. R. L. Albin, A. B. Young, and J. B. Penney, The functional anatomy of basal ganglia disorders, *Trends in Neurosciences* **12**, 366-375 (1989).
8. Y. Smith, M. D. Bevan, E. Shink, and J. P. Bolam, Microcircuitry of the direct and indirect pathways of the basal ganglia, *Neuroscience* **86**, 353-387 (1998).
9. M. D. Humphries and K. Gurney, A pulsed neural network model of bursting in the basal ganglia, *Neural Networks* **14**, 845–863 (2001).
10. D. Plenz and S. T. Kitai, A basal ganglia pacemaker formed by the subthalamic nucleus and external globus pallidus, *Nature* **400**, 677-682 (1999).
11. P. Redgrave, T. J. Prescott, and K. Gurney, Is the short latency dopamine burst too short to signal reward error? *Trends in Neurosciences* **22**, 146-150 (1999).