

Article

Not peer-reviewed version

Simulated Dopamine Modulation of a Neurorobotic Model of the Basal Ganglia

[Tony J. Prescott](#)^{*}, Fernando Montes Gonzalez, Kevin Gurney, Mark D. Humphries, Peter Redgrave

Posted Date: 8 January 2024

doi: 10.20944/preprints202401.0556.v1

Keywords: neurorobotics; basal ganglia; computational psychiatry; Parkinson's disease; dopamine; dopamine dysregulation; computational neuroscience



Preprints.org is a free multidiscipline platform providing preprint service that is dedicated to making early versions of research outputs permanently available and citable. Preprints posted at Preprints.org appear in Web of Science, Crossref, Google Scholar, Scilit, Europe PMC.

Copyright: This is an open access article distributed under the Creative Commons Attribution License which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

Article

Simulated Dopamine Modulation of a Neurorobotic Model of the Basal Ganglia

Tony J. Prescott^{1,*}, Fernando M. Montes Gonzalez², Kevin Gurney³, Mark D. Humphries⁴ and Peter Redgrave³

¹ Department of Computer Science, University of Sheffield, United Kingdom

² Departamento de Inteligencia Artificial, Universidad Veracruzana, Mexico

³ Department of Psychology, University of Sheffield, United Kingdom

⁴ School of Psychology, University of Nottingham, United Kingdom

* Correspondence: t.j.prescott@sheffield.ac.uk

Abstract: The vertebrate basal ganglia play an important role in action selection—the resolution of conflicts between alternative motor programs. Effective operation of basal ganglia circuitry is also known to rely on appropriate levels of tonic dopamine transmission. We investigated reducing or increasing the tonic level of simulated dopamine in model of the basal ganglia integrated into a robot control architecture engaged in a foraging task inspired by animal behavior. The main findings were that progressive reductions in the levels of simulated dopamine caused slowed behavior and, at low levels, an inability to initiate movement. These states were partially relieved by increased salience levels (stronger sensory/motivational input). Conversely, increased simulated dopamine caused distortion of the robot's motor acts through partially-expressed motor activity relating to losing actions; this could also lead to increased frequency of behavior switching. Levels of simulated dopamine that were either significantly lower or higher than baseline could cause a loss of behavioral integration, sometimes leaving the robot in a 'behavioral trap'. That some analogous traits are observed in animals and humans affected by dopamine dysregulation suggests that robotic models could prove useful in understanding the role of dopamine neurotransmission in basal ganglia function and dysfunction.

Keywords: basal ganglia; dopamine; robot; Parkinson's disease; dopamine dysregulation; neurorobotics; computational neuroscience; computational psychiatry

1. Introduction

The vertebrate basal ganglia are thought to play an important role in action selection—the resolution of conflicts between alternative motor programs [1-7]. The effective operation of basal ganglia circuitry, and its regulation of motor behavior, is also known to rely on appropriate levels of dopamine (DA) transmission [3, 8, 9]. For instance, DA antagonists, or dopamine-depleting brain lesions, have been found to impair a range of instrumental and spontaneous behaviors [10-16], effect the maintenance of behavior over time [10, 17], impair the initiation of movement [18, 19], reduce behavior switching [13, 20-23], and can induce bradykinesia (slowed movement) or akinesia (lack of movement) [24, 25]. Conversely, DA agonists have been shown to cause increases in behavior switching [21, 22, 26], or can lead to patterns of repetitive behavior (stereotypy) [27-30]. Human basal ganglia-related disorders such as Parkinson's Disease (PD), Schizophrenia, Attention Deficit Hyperactivity Disorder (ADHD), and Tourette's Syndrome (TS) are also known to involve abnormalities in the dopamine regulation of basal ganglia circuitry [31-35]. Nevertheless, in both humans and animals, there is still much to understand about how variation in tonic dopamine levels can have these different and variable effects on behavior.

In this article we show that when the tonic level of simulated dopamine in a robotic model of the basal ganglia [36] is significantly reduced or increased, relative to a baseline, a variety of behavioral outcomes are observed that provide interesting comparisons with the results of animal studies, and with some of the observed behavioral consequences of dopamine dysregulation in

disorders affecting the human basal ganglia. In this way, we hope that this article can contribute to the emerging field of computational psychiatry [37] and the investigation of models of psychopathology via robotics [38].

The structure of the article is organized as follows. Section 2 describes some principles that allow us to measure the effectiveness of action selection, provides a summary of the computational models developed by Gurney, Prescott and Redgrave [39, 40] and by Humphries and Gurney [41], that view the basal ganglia as a selection mechanism, and sets out some metrics that allow us to evaluate its performance against our principles. Section 3 describes the embedding of this model in the control architecture of a mobile robot as previously reported by Prescott, Montes Gonzalez et al. [36]. Section 4 describes some experiments with a non-embodied version of the model that provide fresh insight into the effects of tonic dopamine modulation on selection by this model. Section 5 then applies ethological methods to analyse the results of experiments with the robot embedding of the model in which we vary the simulated level of tonic dopamine. Finally, section 6 draws some comparisons with animal and human data, and discusses some of the implications of our study for the use of robotics modelling in neuroscience.

2. The basal ganglia viewed as an action selection mechanism

Requirements for effective selection

Given a set of competing and incompatible programs, the requirements for an effective action selection mechanism [3, 42] can be summarised as: (i) in selecting a winner, all else being equal, prefer the most strongly supported, or most *salient*, competitor as indicated by relevant external and internal cues; (ii) allow only one program to be expressed at a given time, this winner should be *cleanly selected* (i.e. allowed unrestricted access to the motor apparatus) and the losers should be prevented from interfering with its performance, termed lack of *distortion*; (iii) provide *clean switching*—a competitor with a slight edge over its rivals should see the competition resolved rapidly and decisively in its favor; and (iv) support *action maintenance*—a winning competitor may be required to remain active at lower salience levels than are initially required for it to overcome the competition. This latter characteristic, also termed *hysteresis* [43] or *behavioral persistence* [44], can prevent unnecessary switching, or ‘dithering’, between closely matched competitors.

Note that this view of action selection treats input salience as the ‘common currency’ according to which diverse behavioral options can be evaluated for purposes of selection—the selector doesn’t need to know what the option is, only how salient it is, with salience being determined by genetics and learning. We note that other ways of selecting between actions are possible that do not rely on salience computation, and that may well exist in the brain, these could operate in a complementary fashion to the centralised action selection mechanisms considered here (see [45] for further discussion).

A model of basal ganglia intrinsic circuitry

In a series of computational models, Gurney and co-workers [39-41, 46] showed that the intrinsic connectivity of the basal ganglia, shown Figure 1 (left) can meet many of these requirements for effective selection via a variety of mechanisms centred on: (i) a set of pathways from the *striatum*, the basal ganglia’s chief input nucleus, that can generate focused inhibition in basal ganglia output structures—the *substantia nigra pars reticulata* (SNR) and the *internal globus pallidus* (GPi) [2] (*entopeduncular nucleus* in rats); (ii) diffuse excitation of these output structures by the *subthalamic nucleus* (STN) [47], and (iii) regulation of the contrast between this focused striatal inhibition and diffuse STN excitation by the *globus pallidus* (GPe) [39, 48, 49]. The overall mechanism is one that selects by removing tonic inhibition of motor pathways provided by basal ganglia outputs, for selected actions only, whilst maintaining or increasing inhibition on non-selected actions, and is in accord with many theoretical models of basal ganglia viewed as an action selection mechanisms [e.g. 2, 3, 6]. The novelty of the Gurney et al. model included showing that intrinsic circuitry involving the

GP acts to regulate this selection effect, for instance, normalising the level of surround inhibition for different numbers of competitors [40].

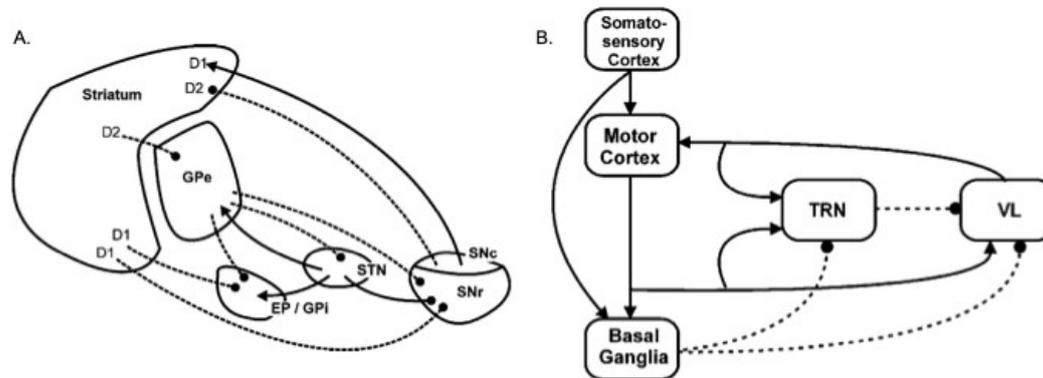


Figure 1. A. The connectivity, relative position, and relative size of the nuclei that comprise the vertebrate basal ganglia showing the separate projection targets of the D1 and D2 receptor striatal neurons as modelled by [39, 40]. B. The connection scheme of the extended basal ganglia model, as modelled by [41], incorporating feedback pathway to the cortex via the thalamus, the box labelled ‘basal ganglia’ contains the functional anatomy shown on the left. Solid lines depict excitatory pathway, dotted lines inhibitory pathways. Anatomical labels are for the primate brain. Abbreviations: GPe—Globus Pallidus external segment; GPi—Globus Pallidus internal segment (EP—entopeduncular nucleus in rat); STN—subthalamic nucleus; SNc—substantia nigra pars compacta; SNr—substantia nigra pars reticulata. TRN—thalamic reticular nucleus—VL—ventrolateral thalamus. Reprinted from Humphries and Gurney [41] figures 1 and 3 with permission.

The balance between the different intrinsic basal ganglia mechanisms is also thought to depend on the level of tonic dopamine expression which differentially impacts on striatal projection neurons with different receptor types [39, 50]. Specifically, striatal projection neurons can be separated into two broad classes. One population contains the neuropeptides substance P and dynorphin, preferentially expresses the D1 subtype of dopamine receptors, and projects primarily to the output nuclei (SNr and GPi). Activity in these ‘D1 neurons’ suppresses the tonic firing in basal ganglia output structures thus acting to select (disinhibit) target structures in the thalamus and brainstem [39, 51]. A second population of projection neurons contains enkephalin and preferentially expresses D2 subtype dopamine receptors. The inhibitory projection from these ‘D2 neurons’ constitutes the first leg of an indirect pathway to the output nuclei that has two inhibitory links (Striatum–GPe, GPe–STN), followed by an excitatory one (STN–GPi/SNr). The net effect of D2 activity is therefore to activate output nuclei increasing inhibitory control of the thalamus and brainstem [39, 52, 53]. Gurney et al. [40] demonstrated that simulation of increasing tonic dopamine in the model basal ganglia has the effect of increasing D1 neuron activity, reducing D2 activity, and consequently reducing activity in GPi/SNr. They concluded that raising tonic dopamine levels makes selection more ‘promiscuous’ increasing the likelihood that target motor pathways will be disinhibited, and potentially leading to ‘soft’ selection—the full or partial disinhibition of multiple channels.

A model of the extended basal ganglia

Humphries and Gurney [41] extended this intrinsic model, as shown in Figure 1 right, to include extrinsic feedback pathways via the *ventral thalamus* (VL) and the *thalamic reticular nucleus* (TRN). This new model provided improved selection, compared to the model of intrinsic circuitry alone, particularly with regard to generating clean selection with absence of distortion (partial expression of losing channels), and the ability to maintain selected actions through positive feedback provided by the *basal-ganglia-thalamo-cortical loop*.

Humphries and Gurney [41] provide a motivation for, and full description of, the extended basal ganglia model and a detailed description of its neurorobotic implementation is provided by Prescott et al. [36]. Here we note that all of these models are based on standard ‘leaky integrator’ units, where one unit is used to represent activity in a pool of neurons in each of the modelled nuclei illustrated in figure 1, and for each of the competing basal ganglia ‘channels’. The input to channel i of the model, denoted s_i indicates the instantaneous salience of that channel as computed by structures outside of the basal ganglia or by the striatal projection neurons themselves. The output for channel i , denoted y_i^{snr} indicates the instantaneous value of the inhibitory signal from the basal ganglia output nuclei to their targets elsewhere in the brain. Tonic dopamine modulation of the model is provided by a multiplicative factor in the equations specifying afferent input to the striatum, based on a variable parameter λ where $0.0 \leq \lambda \leq 0.5$. In striatal D1 units, where dopamine modulation *increases* synaptic efficacy, the effective weight is $(1+\lambda)$; in D2 units, where the effect is to *reduce* efficacy, the weight is $(1-\lambda)$. In the investigations described below we vary λ within the above range to simulate the effects of changing tonic dopamine modulation in the basal ganglia. Note that the intention is to model changes in dopamine that happen over longer time-scales and that we do not attempt, in this study, to model phasic short-latency dopamine responses that may also have an important effect on selection and are considered to be play a critical role in some forms of learning [8, 54].

Using basal ganglia outputs as selection signals

In order to consider the basal ganglia model as a model of selection we need to place some interpretation on the effects of basal ganglia outputs on targets in the brainstem and thalamus. As noted above, selection corresponds to basal ganglia removing inhibition from the winner(s) and increasing inhibition on the losers. We assume that for any given channel this effect varies between full disinhibition, partial inhibition, and full inhibition, and model this effect via a mechanism termed ‘shunting inhibition’ thought to capture some of the non-linear effects of the Gabaergic outputs from basal ganglia on their targets in vivo (see [36]). Specifically, for the i^{th} channel, we define the selection, or *gating*, signal e_i ($0 \leq e_i \leq 1$) as

$$e_i = \begin{cases} 0: \hat{e} < 0 \\ \hat{e}: 0 \leq \hat{e} \leq 1, \text{ where } \hat{e}_i = 1 - y_i^{snr} / y_c^{snr} \\ 1: \hat{e} > 1 \end{cases} \quad (1)$$

Here y_c^{snr} is a constant equal to the value of y_i^{snr} obtained when the basal ganglia model is run to convergence with zero salience input on all channels (in other words, the tonic output level when there are no active competitors).

Metrics for measuring effective selection

The gating signal, e_i , provides a useful normalized measure of selection efficiency that we can use to evaluate any given version of the model against our requirements for effective action selection. It is useful to define some qualitative labels for different values of e_i . Allowing a 5% margin from absolute limits, we define the selection state of the i^{th} competitor as *fully selected* if $0.95 \leq e_i \leq 1.0$, *partially selected* if $0.05 \leq e_i < 0.95$, and *unselected* if $e_i < 0.05$. It will also be useful to have specific metrics relating to the winning channel, hence, we define $e_w = \max_{\forall i} e_i$ as the *efficiency* of the current winner, $1 - e_w$ as its *inefficiency*, and

$$d_w = 2 \left(\sum_i e_i - e_w \right) / \sum_i e_i$$

as the level of *distortion* affecting the output of this winner. Note that d_w will equal zero when all other competitors have zero efficiency, will increase with the number of partially disinhibited losers, and will be 1.0 or greater if two or more channels are fully disinhibited (multiple winners). Inspired by ethological research [55], we will also describe an uninterrupted series of time-steps that share the same winner, and for which $e_w \geq 0.05$, as a single *bout* of behavior.

Finally, we note that the result of the basal ganglia selection competition, as a whole, can be summarised by the vector \mathbf{e} . Using the criteria just defined for single competitors we assign the following labels to the possible outcomes of the full competition as defined by the instantaneous value of \mathbf{e} :

Clean selection: One competitor fully selected, all others unselected.

No selection: All competitors unselected.

Partial selection: One or more competitors partially selected, no competitor fully selected.

Distorted selection: One competitor only fully selected, at least one other partially selected.

Multiple selection: Two or more competitors fully selected.

3. A robot embedding of a model of action selection by the basal ganglia

Prescott et al. [36, 56] embedded the extended basal ganglia model [41] within the control architecture of a mobile robot in order to demonstrate that signal selection by the embedded model, as described above, could translate into effective action selection for an embodied agent expressing goal-directed behavior. This model was based on consideration of the typical behavior of a hungry rat placed in an open-topped arena with high sides (Figure 2a and Supplementary Video, part 1). In this situation, animals initially show fearful or thigmotaxic behavior—avoiding open areas in the centre of the arena, whilst exploring walls and corners. As animals become more accustomed to the novel environment, they show foraging behavior—collecting food pellets from a dish placed in the centre of the arena and typically consuming them in sheltered areas near the periphery. Salamone [10] showed that effective behavior switching in a similar environment is compromised by the dopamine antagonist Haloperidol and by dopamine-depleting lesions of the striatum, hence the task is an appropriate one for investigation of the effects of variation in simulated dopamine on robot action selection.

In the robot model of this task (Figure 2b and Supplementary Video, part 2), a table-top *Khepera* mobile robot with a gripper turret is placed in a rectangular arena with illuminated corners, to simulate safe places, and with small foil-covered cylinders to simulate food rewards. Fearful behavior is simulated as staying close to walls and corners; foraging involves searching for, locating, and picking up the cylinders; finally, consummatory behavior is modelled as carrying a cylinder to one of the two illuminated corners and depositing it there. To generate appropriate behavior, robot activity is decomposed into five sub-systems inspired by the ethological classification of behavior. Three of the five action sub-systems—*cylinder-seek*, *wall-seek*, and *wall-follow*—map patterns of input from the robot's sensors into movements that orient the robot towards or away from specific types of stimuli (e.g. object contours). These behaviors can be viewed as belonging to the ethological category of *orienting responses* or *taxes* [see e.g. 57]. The two remaining sub-systems—*cylinder-pickup* and *cylinder-deposit*—generate carefully timed movement sequences that achieve specific behavioral outcomes and are modeled on the ethological concept of a *fixed action pattern* (FAP) [58]. Each action sub-system generates its preferred action at a given moment in the form of a *motor vector* that specifies target values for the speeds of the two wheels, and for the positions of the gripper arm (raised/lowered) and gripper jaw (open/shut). In the case of the orienting responses, the preferred action is computed using the sensory information available to the robot at that moment, in the case of FAPs, the action specification can also depend on the current value of an internal clock.

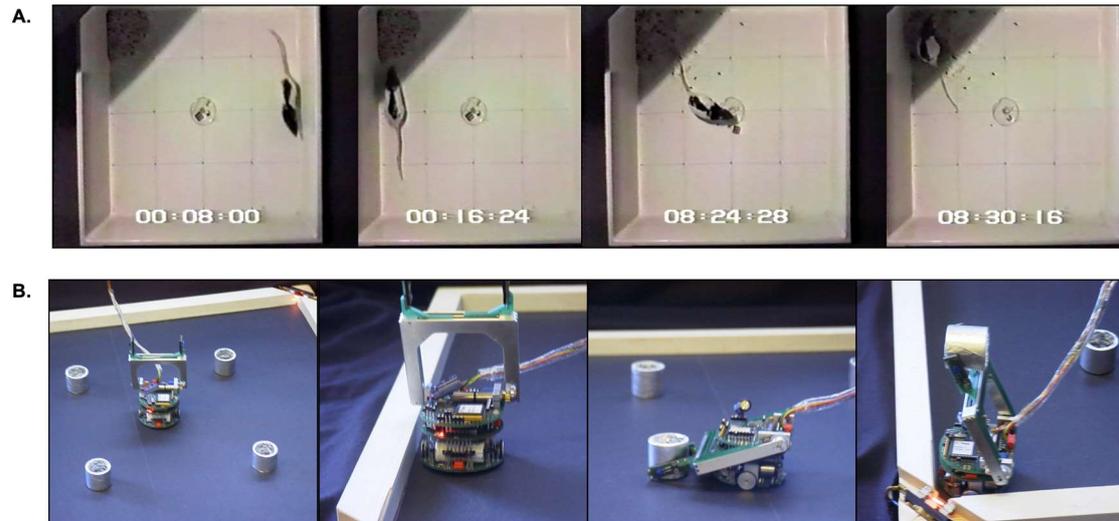


Figure 2. The model task. **A.** A hungry rat placed in an open arena will initially explore the periphery (frames 1 and 2) before eventually venturing into the centre (frame 3) to retrieve food pellets that are then consumed in a sheltered ‘nest’ corner (frame 4). **B.** In the robot these behaviors are simulated by seeking (frame 1) and following walls (frame 2) and by searching for and acquiring cylinders (frame 3) that are then deposited in the lit corner of the arena (frame 4) (see Supplementary Video, part 2).

In order to make appropriate action selection decisions, the robot needs information about relevant external and internal cues. Signals pertaining to external cues are computed by *perceptual sub-systems* from the raw sensory data available to the robot via an array of infra-red distance sensor signals, an ambient light sensor, and an optical sensor in the robot gripper. These sensory inputs are used to compute four bipolar signals indicating: the presence (+1) or absence (-1) of a nearby *wall*, *nest* area, *cylinder*, or of an object in the robot *gripper*. Internal state cues are provided in the form of two real-valued intrinsic drives, loosely analogous to hunger and fear, as calculated by two *motivational sub-systems*. In the model, ‘fear’ is calculated as a function of exposure to the environment and is reduced with time spent in the environment, whilst ‘hunger’ gradually increases with time and is reduced when cylinders are deposited in the nest corners of the arena.

Figure 3 shows how these different component sub-systems come together and interact with the embedded basal ganglia model. The model is composed of three parts: (i) the *robot* and its sensory and motor systems; (ii) the *embedding architecture*, that is, the set of *perceptual*, *motivational*, *action* sub-systems; and its interface to (iii) the *extended basal ganglia model*. Connections for the first of the five action sub-systems are shown, projections to and from the other action sub-systems are indicated by dotted lines.

As shown in figure 3, each action sub-system takes inputs from the perceptual and motivational sub-systems, and from an internally-generated *busy signal* (b) that is only non-zero if the action is currently selected, and that allows that sub-system to selectively boost its own salience. Based on these inputs the action sub-system generates a weighted sum (the weights are hand-tuned) that is an estimate of its own instantaneous *salience* (s) that is provided as an input to the embedded basal ganglia model. At the same time, the action-generating component of the sub-system calculates its preferred motor vector based on the robot’s sensor input and a *feedback signal* (f) from the component of basal ganglia model corresponding to the *ventrolateral thalamus* (VL). This feedback signal is used to update or reset the clock (C) for the action system (in the case of a FAP), and to trigger the *busy signal* that contributes to its salience calculation.

As noted above, for each action sub-system i , the output of the basal ganglia, y_i^{snr} , is converted into a gating signal, e_i , via equation 1 which is then used to scale the value of the motor vector for that action. An *integrator* module then sums all of the motor vectors and passes the aggregate vector through a limiter (L) that constrains all values to lie in the range 0–1, this vector is then converted

into the specific motor commands that control the robot. The full robot model operates on a series of discrete time-steps providing sensor updates and modifying its action output at a rate of approximately 7Hz. The embedded basal ganglia model is simulated using the Euler method and run to convergence for each time-step of the robot model.

Full details of the test environment, the robot sensor and motor systems, the embedding architecture components, and the simulation of the extended basal ganglia model, including their motivation in relation to the neuroscientific understanding of relevant brain sub-systems, are provided in [36], which also provides a broader discussion of the use of robotic models in neuroscience.

(i) Robot (ii) Embedding architecture (iii) Extended basal ganglia model

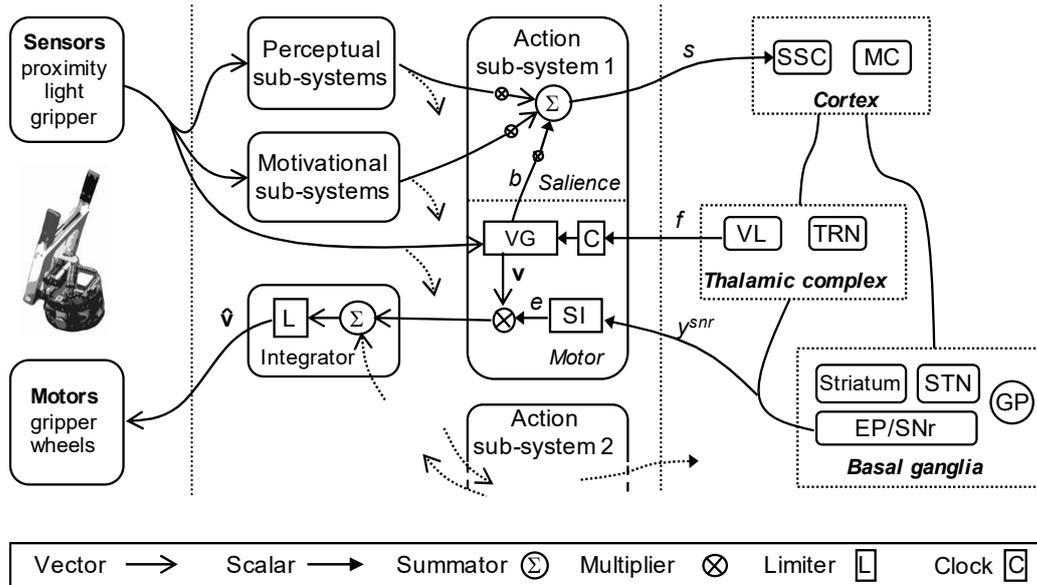


Figure 3. The embedded basal ganglia model. Abbreviations: VG—(motor) vector generator, SI—shunting inhibition (equation 1), e —gating signal, b —busy signal, s —saliency signal, f —feedback signal, y^{snr} —basal ganglia output, v —motor vector, \hat{v} —aggregate motor vector, SSC—somatosensory cortex, MC—motor cortex (other anatomical abbreviations as per figure 1). Reproduced with permission from [36].

To illustrate normal functioning of the model, figure 4 shows a single 240s run with the level of simulated tonic dopamine set at $\lambda = 0.20$. The top five lines of the plot show the value of the gating signal, e_i , for each of the five action sub-systems at each time step in the style of a behavioral ethogram. Comparing across the different actions, it is evident that the robot generates extended sequences of behavior with no more than one action sub-system fully selected at any given time. The efficiency of selected actions is at or near 100%, actions are performed over extended bouts (solid blocks of high efficiency) and the inefficiency of the winner (plotted as the sixth line of the plot) is generally near zero. In this run, the robot is initially fearful and seeks the wall (*wall-seek*), then switches into its *wall-follow* behavior. This can be viewed as forming a *higher-order sequence* of avoidance (A_v) behavior as labelled in the seventh line of the plot. The final line of the plot shows the activity of the model motivational systems. As the level of simulated fear gradually subsides, simulated hunger is increasing. As a result, at around 50s, the robot rapidly switches into its *cylinder-seek* behavior. When it subsequently locates a cylinder it switches to *cylinder-pickup*, then to *wall-seek* (this time carrying a cylinder), then *wall-follow*, and, when it finds a lit corner, *cylinder-deposit*. The higher-order action sequence beginning with *cylinder-seek* and ending with a successful deposit is labelled as *foraging* (Fo) in the plot. Releasing the cylinder has the effect of reducing simulated hunger such that the robot is again motivated principally by fear to perform its avoidance-related behaviors (*wall-seek* and *wall-follow*). However, the level of simulated hunger gradually rises which leads to two further higher-order foraging sequences interspersed by a period of *no behavior*. The absence of behavior

occurs when neither of the intrinsic motivations is sufficiently strong to trigger any action—the robot sits idle, just as the rat might wait quietly in the corner of the arena.

From the perspective of the observer, the robot's behavior appears to be integrated and purposeful, individual action bouts are assembled to larger sequences that successfully reduce its drives. In section 5, we will compare this example of effective action selection and integrated behavior with other runs in which the robot demonstrates various forms of behavioral disintegration as the result lowering or raising the level of simulated dopamine in the model basal ganglia.

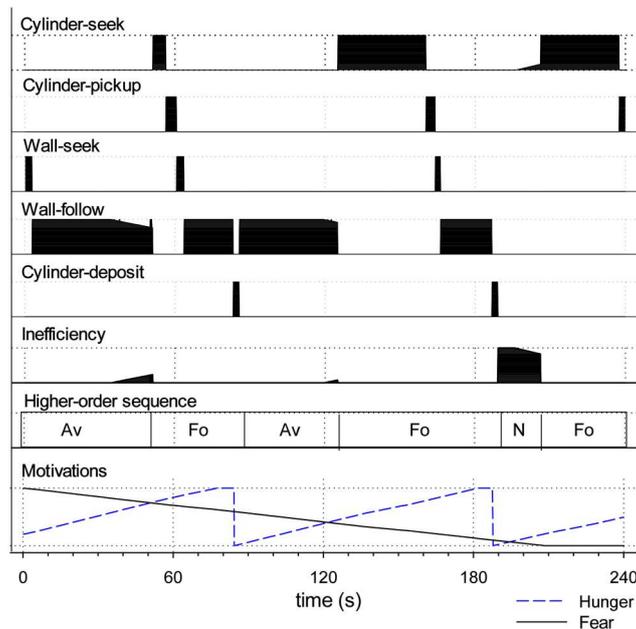


Figure 4. Bout/sequence structure of action selection in the robot model for a 240s trial ($\lambda = 0.20$), the first 100s is shown in the Supplementary Video, part 3. Each of the first five plots shows the efficiency (e) of selection for a given action sub-system plotted against time. The sixth plot shows the inefficiency of the current winner, the seventh the higher-order structure of the bout sequences, (av= avoidance, fo= foraging), and the final plot the levels of the two simulated motivations. All measures vary between 0 and 1 on the y-axis. The robot display appropriate bouts of behavior organized into integrated, goal-achieving sequences.

4. Tonic dopamine modulation in the extended basal ganglia model

Before presenting further results for the robot model, it is useful to investigate the response of a non-embodied version of the extended basal ganglia model to changes in tonic dopamine modulation as this will provide a useful yardstick for evaluating the embodied robotic version, and will help us to better understand the specific consequences that embodiment might entail. This investigation builds on prior studies of simulated tonic dopamine modulation [40, 41] by providing a fine-grained analysis across the spectrum of possible λ levels.

To better understand the effect of varying simulated dopamine on the selection properties of the extended basal ganglia model we simulated a five channel model, with two active channels, varying the salience s_1 in channel 1 systematically from 0 to 1 in steps of 0.01, then for each value of s_1 , varying the salience s_2 of channel 2 from 0 through 1, again in steps of 0.01. For each resulting salience vector $(s_1, s_2, 0, 0, 0)$ the model was run to convergence and the result classified according to the scheme set out in Section 2. Importantly, selection competitions were run in sequence from low values to high values. The activations levels of all leaky integrators in the model were initialized to zero for each new value of s_1 but thereafter, while that salience value was tested, were retained from one competition to the next. In other words, we simulated the situation where channel 1 was initially the only active channel, and *gradually increased channel 2 while holding channel 1 constant*, the goal being to simulate some aspects of the continuity of experience that we can expect in the robot model in

which the recent history of selection competitions may influence the current competition through hysteresis. Previous studies have established that the basal ganglia model, in both its original and extended forms, shows good selection properties, across a wide-range of salience pairings, with the simulated dopamine level set at around $\lambda = 0.20$; for this analysis we therefore looked at values of simulated dopamine ranging from 0 through to 0.5 in increments of 0.01.

Figure 5a shows the percentage of action selection competitions, across the 500,000 (50x100x100) runs, falling into each of the selection classes— *clean selection*, *no selection*, *partial selection*, *distortion*, and *multiple selection*. Values of λ below 0.01 result in no selection, while in the range 0.04–0.15 partial selection predominates, from 0.15 upwards the majority of competitions end in clean selection with a peak around 0.22, distorted selection begins to appear with values above 0.2, and multiple selection occurs with levels of 0.25 and greater.

Figure 5b shows the average values of efficiency and distortion across all runs at a given level of λ . These graphs indicate that average efficiency increases gradually reaching its maximal value (1.0) at $\lambda = 0.23$, distortion increases gradually from zero beginning at around $\lambda = 0.15$ and reaching 0.2 by $\lambda = 0.5$.

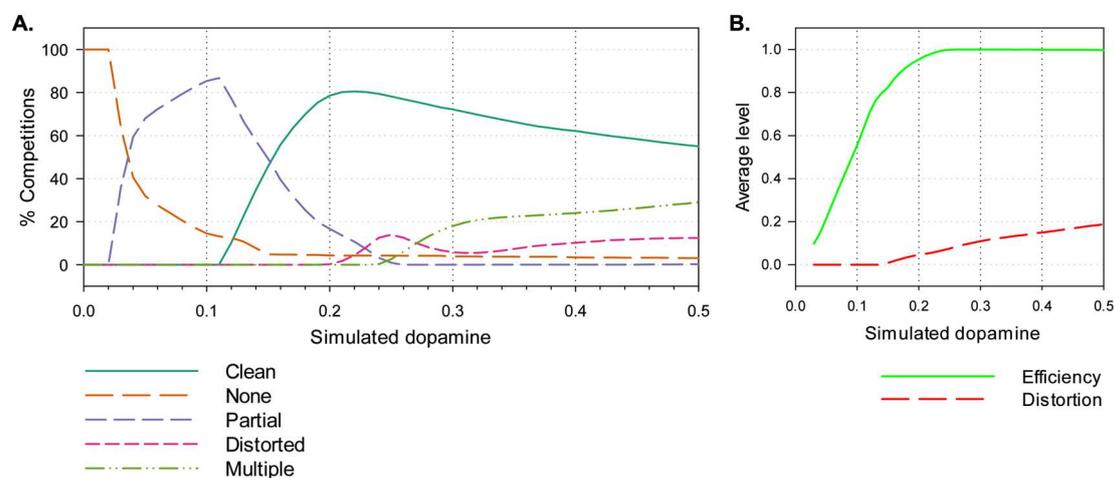


Figure 5. A. The percentage of selection competitions falling into different classes of selection outcome for values of simulated dopamine, λ , ranging from 0 through to 0.5 in increments of 0.01. Data was obtained through an exhaustive search of a two-dimensional salience space. Partial selection is predominant for low dopamine values, distortion and multiple-selection evident at high dopamine values. **B.** Average efficiency and distortion across all runs at each level of λ .

Figure 5 shows the average outcome at different levels of λ across all possible (s_1, s_2) dyads. In order to better understand the interplay between salience, simulated dopamine and selection, in figure 6 we show the outcome of the simulation for five specific values of simulated dopamine ($\lambda = 0.06, 0.12, 0.22, 0.31, 0.40$) but indicating the boundaries of different classes of selection outcomes on the (s_1, s_2) plane. For clean selection only, the plots also distinguish between selection of channel 1 (which is active first) and of channel 2 (which then competes for selection against channel 1).

Several properties of figure 6 are worth noting. First, at all levels of λ , there is little or no selection at very low salience levels. This is largely as a consequence of the threshold value of the model striatal input neurons which serves to weed-out weakly salient inputs. Second, with low λ (e.g. 0.12), clean selection (C1 or C2 in figure 6) occurs, if at all, only when there is a high salience input in just one channel, otherwise partial selection is the norm. Third, at all simulated dopamine levels there is no clean selection for strong, evenly matched, salience values (top-right corner of all plots). With low values of λ (0.06, 0.12) the outcome is no selection or partial selection of one or both channels, while with high values (0.31, 0.4) the result is distortion of the selected channel or multiple selection. The dotted line in the central plot ($\lambda = 0.22$) is shown to illustrate the extent of hysteresis in the model: channel 1 wins many selection competitions (encroaches across the diagonal) in which channel 2 salience is greater, purely because it was activated first.

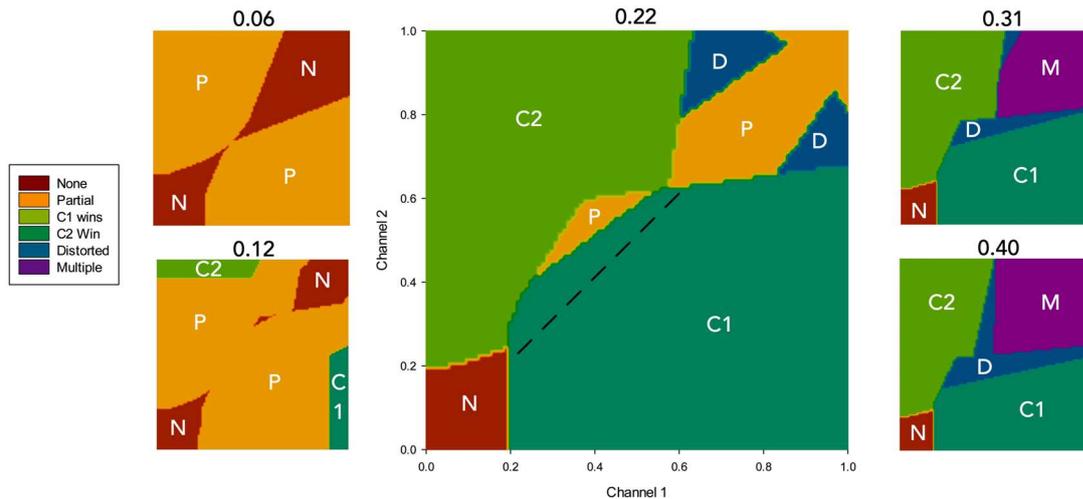


Figure 6. Selection boundaries in two-dimensional salience space for sample levels of simulated dopamine—very low ($\lambda=0.06$), low (0.12), intermediate (0.22), high (0.31), and very high (0.40). For each plot the salience of channel one is shown on the x-axis, and that of channel two on the y-axis ranging from 0.0–1.0 (shown only for the central plot). Labels indicate: N—no selection, P—partial selection, C1—clean selection of channel 1, C2—clean selection of channel 2, D—distortion, M—multiple selection.

To further our understanding of hysteresis in the model, the simulation results described above were reclassified to show the extent to which channel 1, which is always active first, is preferred to channel 2 irrespective of the selection outcome. Thus, the result of each competition was rescored as either a *channel 1 win* ($e_1 > e_2$), a *channel 2 win* ($e_2 > e_1$), a *tie* ($e_1 = e_2 \neq 0$), or *no selection* ($e_1 = e_2 = 0$). Figure 7a shows the results of this reclassification, and reveals that hysteresis is a property of the model for all but the lowest levels of simulated dopamine modulation ($\lambda \leq 0.06$) with channel 1 consistently winning up to 10% more competitions than channel 2.

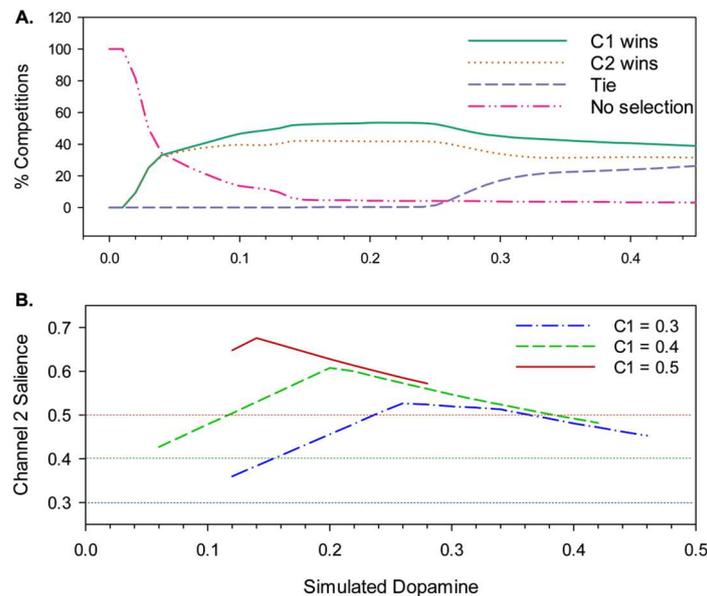


Figure 7. A. Selection outcomes in the disembodied model re-classified as a channel 1 win, a channel 2 win, a stand-off (no selection), or a tie. Channel 1 (c1) wins substantially more competitions than channel 2 (c2) for all but the lowest levels of simulated dopamine. **B.** The level of channel 2 saliency, s_2 , required for channel 2 to prevail (i.e. $e_2 > e_1$) against a channel 1 saliency, s_1 , of 0.3, 0.4, or 0.5, for different values of λ . Data is shown only where there is a clear switch from channel 1 to channel 2

with increasing s_2 (i.e. without an intervening interval of no-selection or multiple selection). The degree of hysteresis varies depending on λ and s_1 , with the value of λ that generates maximum hysteresis decreasing with increasing s_1 .

However, this is still not the full story. Figure 7b shows a further measure of hysteresis—the level of channel 2 salience required to overcome a given level of channel 1 salience—for three different initial, fixed levels of s_1 . The plot shows that hysteresis is governed by a complex interaction of λ with salience, specifically, for values of s_1 in the range 0.3-0.5 the degree of hysteresis first increases with increasing λ , peaks, and then decreases; at its maximum, channel 2 salience needs to reach 176% of the channel 1 salience in order to win the selection competition. The peak λ value for hysteresis also changes for different values of s_1 —as the salience of the selected channel increases, the value of λ at which hysteresis is maximal goes lower.

We conclude that the relatively flat level of hysteresis shown across a broad range of λ values in figure 7a masks a significance dependency on salience. This outcome can be explained by understanding that hysteresis in the model occurs as a consequence of activity in the basal-ganglia-thalamo-cortical feedback loop (via VL and TRN in Figure 1). Activity in this loop increases in proportion to reduced basal ganglia output, in other words, it increases with selection efficiency. With low values of λ , partial selection (low efficiency) predominates for low or intermediate salience values. This outcome results in less positive feedback via the thalamo-cortical pathway than with high salience competitions. Consequently, when λ is low, hysteresis will be maximal with high salience. In contrast, high λ levels result in high efficiency selection with comparatively low-levels of salience input thus generating substantial positive feedback and strong hysteresis. However, high-level salience competitions can result in the partial or full disinhibition of multiple channels (distorted or multiple selection), a consequence of this is increased TRN inhibition of the VL thalamus for the winning channel resulting in a significant reduction in thalamocortical feedback for that channel. This means that with higher levels of λ , the current winner can be more vulnerable to interrupt by its competitors.

In figure 5, clean selection for the disembodied model was above 75% in the range $0.2 \leq \lambda < 0.3$, fell steeply to zero in the lower range $0.0 \leq \lambda < 0.2$ and more gradually (to 55%) in the higher range $0.3 \leq \lambda \leq 0.5$. Defining these ranges as, respectively, intermediate, low, and high λ , and building on the analysis just described (and in earlier explorations in [36, 40, 41]), we can make the following hypotheses concerning the effects of varying simulated dopamine in the robotic model:

- *Hypothesis 1 (h1)*. At intermediate levels of λ ($0.2 \leq \lambda < 0.3$) we should expect to see a high proportion of clean selection with selected behaviors fully disinhibited and competing behaviors fully suppressed.
- *Hypothesis 2*. At low levels of λ ($0.0 \leq \lambda < 0.2$) we should expect a predominance of partial selection or no selection (very low λ) and consequently the slowing or absence of movement.
- *Hypothesis 3*. For high levels of λ ($0.3 \leq \lambda \leq 0.5$) we should expect to see reduced inhibition of losing channels, leading to distorted or multiple selection, and resulting in motor commands that mix the movement requests of more than one action sub-system.
- *Hypothesis 4*. At both low and high levels of λ , we should expect to see changes in the hysteresis of selected channels modulated according to the nature of the salience competition (e.g. whether the salience of competing channels is high or low, or evenly matched) as illustrated in figure 7b. Changes to hysteresis can be expected translate into consequences for action maintainence and for the timing of behavioral switching.

With respect to each of these hypotheses, the observed behavior of the robot may depend on a variety of factors related to its embodiment (discussed further below) and the requirement to generate sequences of integrated behavior. Moreover, whereas the above analysis was based on an exhaustive search of an essentially two-dimensional salience space, the robot model samples *behavior-dependant* trajectories through a five-dimensional salience space. The actual outcomes with respect to hypotheses 1-4 are therefore only partially predicatable from the disembodied model and to be further determined from observation.

5. Selection in the neurorobotic basal ganglia model

Based on our analysis of the disembodied model we decided to test the robot for 30 trials each at low, intermediate and high simulated dopamine levels, with five trials, each lasting 120s, at each of 18 different values of λ : low= 0.03, 0.06, 0.09, 0.12, 0.15, 0.18, intermediate= 0.20, 0.21, 0.22, 0.23, 0.25, 0.28, and high= 0.31, 0.34, 0.37, 0.40, 0.43, 0.46. The robot started each trial in the centre of the arena, facing one of the four walls, with four cylinders placed 18cm diagonally in from each corner (figure 2, right).

In each, trial, which typically consisted of around 800 robot time-steps, the outcome of the basal ganglia selection competition, at each time-step, was classified according to the selection criteria specified above. For each λ value, the percentage of time-steps resulting in each type of selection outcome was then averaged across all five trials regardless of the behavioral outcome on individual trials (which we consider next). The results of this analysis are shown in figure 8a-e together with a plot of average efficiency and distortion across the different λ levels (8f).

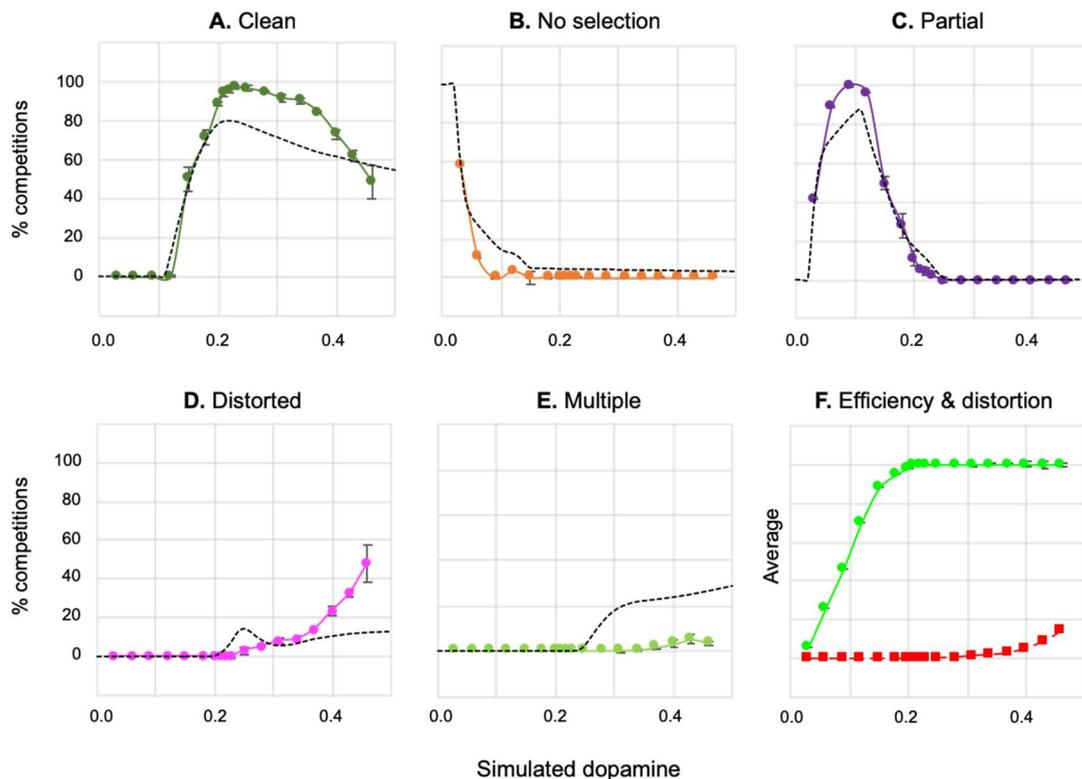


Figure 8. A–E. The percentage of selection competitions falling into different classes of selection outcome for values of simulated dopamine ranging from 0.03 through to 0.46. Data is obtained by averaging across five 120s trials of robot behavior, for each of the eighteen λ levels tested. Standard error bars are shown. Black dotted lines show comparable outcome from the non-embodied model (figure 5). Comparison of selection properties of the non-embodied and robot models shows more clean, partial, and distorted selection in the robotic model and less no selection or multiple selection. F. Average efficiency (green) and distortion (red) across all runs at each level of λ .

These results shows the expected similarity between the selection profiles for the robotic and non-embodied models, nevertheless there are some important differences. These include, in the robotic model, an increased proportion of partial selection at low λ levels ($0.03 \leq \lambda \leq 0.12$), of clean selection at intermediate and moderately-high levels ($0.2 \leq \lambda \leq 0.4$), and of distorted selection at high levels ($0.3 \leq \lambda \leq 0.46$). There is also an almost complete absence of multiple selection at high λ levels. Whilst average efficiency is similar across the robotic and disembodied models, the robot model overall has less distortion except at the highest λ levels. In the intermediate range of simulated

dopamine ($\lambda = 0.20$ – 0.29) clean selection for the robotic model is in the range 89–95% compared to 73–81% clean selection for the disembodied model.

These results largely reflect the fact the robot model is spending little time sampling the very high salience areas of the state-space, or the very low salience areas, compared to the exhaustive search conducted for the disembodied model. This was confirmed by an analysis of salience values across 15 runs (one at each level of λ) which found that 95% of selection competitions were in the range 0.3–0.75 for the winning channel and 0.2–0.7 for the strongest losing channel (see also [36] for a plot of how the salience space is sampled by the robot model). Note that there may also be up to five channels with non-zero salience at any time compared to just two in the disembodied model.

Effects of simulated dopamine modulation on behavioral outcome

Our previous robotic study of the basal ganglia [36] showed that an embedded basal ganglia model was able to generate integrate behavior in our biology-inspired foraging task and for a specific intermediate value of simulated dopamine ($\lambda = 0.20$). In the current study, we address the question of how varying simulated dopamine impacts behavioral integration, and seek to describe and understand a variety of distinctive patterns of behavioral disintegration that arise when simulated dopamine is reduced or increase relative to this baseline.

To begin this analysis a simple binary classification scheme was developed and applied to the 90 robot trials described above, evaluating each trial according to its success in achieving higher-order behavioral goals. Specifically, we define ‘integrated behavior’ for this task as constituting, at minimum, successful avoidance in the initial ‘high fear/low hunger’ phase, and a successful foraging sequence in the later ‘low fear/high hunger’ phase. Operationally, we define:

- (i) *successful avoidance* as activity resulting in the discovery of a wall (ignoring any cylinders encountered en route) followed by movement some distance along the wall’s length, and
- (ii) *successful foraging* as activity resulting in the deposition of a cylinder in a ‘nest’ area.

This classification scheme proved to be sufficiently simple to be applied during live observation of robot behavior, in addition, automatic logs were recorded detailing the robot’s sensory, motivational, and basal ganglia state at each time-step, and the bout structure of its behavioral selections, allowing us to reconstruct and analyse the robot’s behavior post hoc.

The outcome of our initial analysis was as follows. Seven levels of simulated dopamine (0.20–0.28 and 0.37) were scored as generating successful behavior on all five trials, five levels (0.03–0.12 and 0.46) were unsuccessful on all trials, and the remaining six levels (0.15, 0.18, 0.31, 0.34, 0.40, 0.43) generated a mixture of successful and unsuccessful trials.

In order to better understand what was happening at levels of λ that generated mixed results, a quota sampling strategy was implemented in which further trials were conducted until a total of five successful trials, at each of these levels, had been achieved. This required between 1 and 11 trials per level, resulting in an additional 26 trials. Figure 9 shows the total trials (9a) the overall success rate (9b) at different levels of λ , across all 116 trials, assessed against the criteria of success in both avoidance and foraging. Figure 9c shows a more detailed analysis of types of failures under the low and high λ regimes that we describe further below.

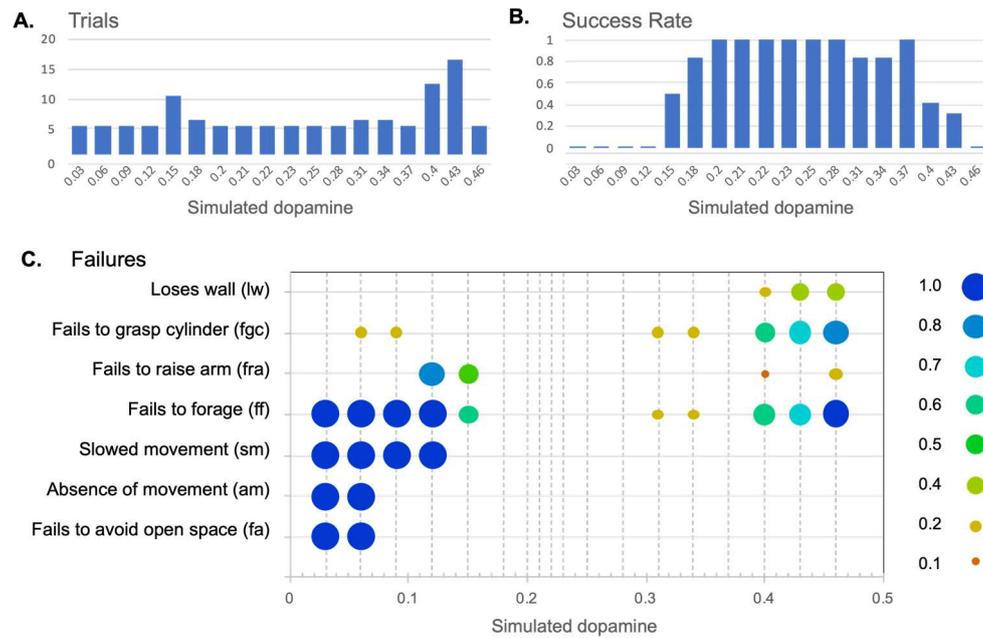


Figure 9. Total trials (A) and success rate (B, 0.0-1.0) in achieving avoidance/foraging different levels of simulated dopamine (λ). C. Evidence of disintegrated behavior at different levels of λ . The bubble plot shows the proportion of trials at each value of λ that resulted in the observed failure type. See text for further details.

Figure 9b confirms that the in the range of intermediate λ values (0.2–0.28), that generates a high proportions of clean selection in figure 8, the robot also reliably generates integrated sequences of behavior. The absence of any failures in the 30 trials in this range provides a 95% confidence level that the failure rate for this class of models is 10% or less.

In the remainder of this section we consider the nature of the failures in behavioral integration that occur with levels of λ below or above this intermediate range then explore the effects of simulated dopamine modulation on the timing and frequency of behavior switching. Figure 9c provides an analysis of the types of failure of behavioral integration observed at different level of λ , and as described in table 1. Figure 10a-e shows some example runs, recorded with low and high λ , that help to illustrate the robot behavior observed at different levels of simulated dopamine.

In the statistical analyses reported below we use an alpha value of 0.05 and report significance values as two-tailed. If Levene's test is significant then we report "equal variances not assumed" and provide adjusted degrees of freedom and p-values.

Table 1. Types of behavioral disintegration in the neurobotic basal ganglia model.

<i>Failure to meet success criterion</i>	
Fails to avoid open space (<i>fa</i>)	Failure with respect to criterion (i) above.
Fails to forage (<i>ff</i>)	Failure with respect to criterion (ii) above.
<i>Behaviors typically leading to <i>fa</i> or <i>ff</i></i>	
Absence of movement (<i>am</i>)	Failure to express movement despite being motivated. Typically leads to <i>fa</i> as the robot fails to leave open space.
Fails to raise arm (<i>fra</i>)	Fails to lift the arm after grasping a cylinder. Typically leads to <i>ff</i> as the lowered arm

	blocks the infrared sensors ability to detect the environment.
Fails to grasp cylinder (<i>fgc</i>)	Fails to lower the arm sufficiently to grasp a cylinder (therefore grasping at air). This can lead to <i>ff</i> as when the robot fails to grasp the cylinder it then immediately looks for a cylinder which can lead to repeated cycles of <i>cylinder-seek</i> followed by (unsuccessful) <i>cylinder-pickup</i> .
<i>Forms of behavioral disintegration typically not leading to fa or ff</i>	
Slowed movement (<i>sm</i>)	Scored when behavior, such as wheeled movement, was slowed to 75% or less of usual speed (as measured by the output motor signal).
Loses wall (<i>lw</i>)	Losing contact with the wall while expressing the <i>wall-follow</i> behavior. Scored as occurring if contact was lost a minimum of four times in sequence (since occasional losses can occur due to sensor noise).

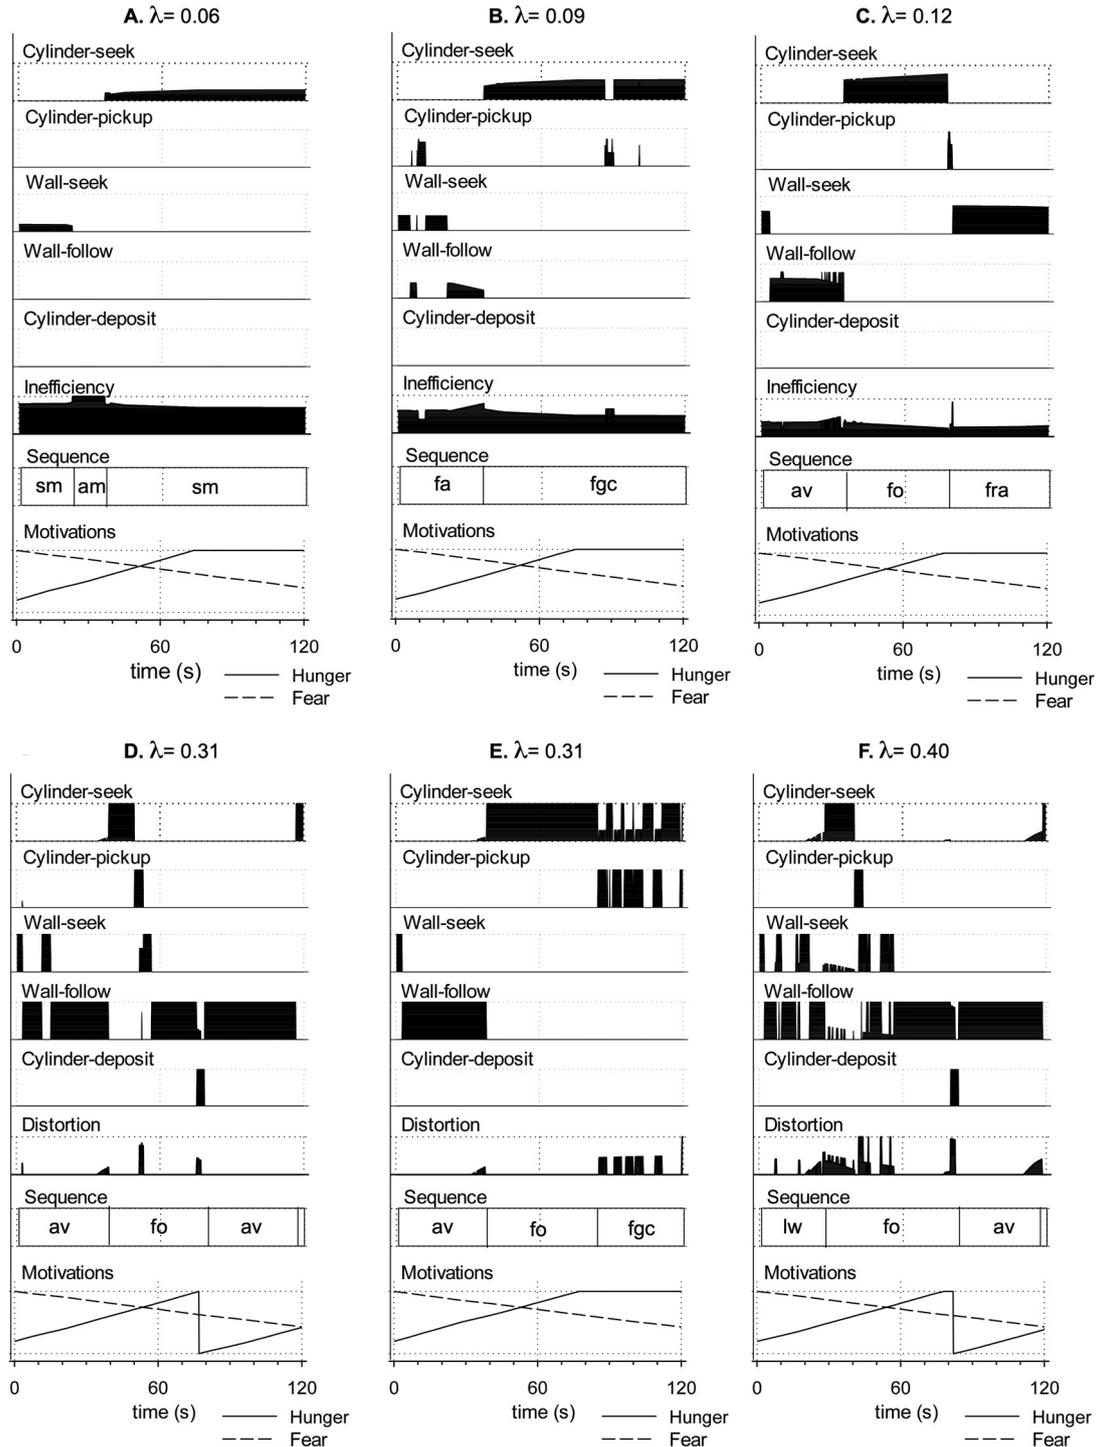


Figure 10. Bout/sequence structure of action selection in the robot model for three 120s trials with low simulated dopamine: A: $\lambda = 0.06$; B: $\lambda = 0.09$; C: $\lambda = 0.12$ and three with high simulated dopamine: D: $\lambda = 0.31$; E: $\lambda = 0.31$; F: $\lambda = 0.40$. Graph layout is as described for figure 4, except that distortion, d_w , of the winning action, replaces inefficiency for panels D–F (as inefficiency is always zero in these trials). Labels in the ‘sequence’ plot show successful avoidance (*av*), foraging (*fo*), or different forms of behavioral disintegration as per table 1. With low simulated dopamine the robot shows slowed movement (*sm*) and absence of movement (*am*). Inefficient selection can also causes premature deselection leading to the failures to grasp the cylinder (*fgc*) or raise the gripper-arm (*fra*) in plots B and C. With high values of λ distortion of the selected behavior, by the motor output of losing competitors becomes a significant issue. Distortion in the run shown in plot D has only benign effects, but in the run shown in plot E causes behavioral disintegration as the robot fails to grasp a cylinder

(*fgc*) despite multiple attempts. The run shown in plot F demonstrates that there is high frequency of behavior switching with high levels of simulated dopamine, in this case because distortion causes the robot to repeatedly lose track of the walls (*lw*). See text for further discussion.

Behavioral consequences of low simulated tonic dopamine ($\lambda < 0.2$)

Slowed movement and periods of inaction. In section 4 we showed that the model basal ganglia generates partial (low efficiency) selection for low levels of simulated dopamine. Since our robotic model employs the basal ganglia output as a gate on targeted motor systems, the consequence of partial selection in behavioral terms should be that this gate is not fully opened for winning competitors; motor acts should be slowed or even extinguished altogether. This expectation, noted as hypothesis 1 above, was borne out in our study (see figure 9c) which saw the expected translation of partial/weak selection into slowed movement (*sm*) for all runs at λ level 0.12 or lower. At $\lambda=0.06$, 0.03 the robot is moving too slowly to complete a successful foraging sequence in the time allowed, thus failing altogether on the criterion for successful avoidance (*fa*). Periods during which the robot makes no movement (*am*), despite being otherwise sufficiently motivated, are seen at $\lambda=0.06$ (average of 14s per trial, compared to 2s for intermediate levels of λ) and for longer spells at $\lambda=0.03$ (average of 38s per trial). Note that it is possible to distinguish between the dysfunctional absence of movement due to low λ as seen in figure 10a, and its appropriate absence during periods of low motivation (as in the period of no selection for $\lambda=0.20$ in figure 4). The Supplementary Video (part 4) shows an example of slow movement and no movement for an example run with $\lambda=0.10$.

Premature deselection. In the range $\lambda=0.06$ – 0.15 behavior can break down as the result of the premature deselection of an ongoing behavior, this can be seen as a failure of persistence or action maintenance. At $\lambda=0.09$ or below this typically occurred during the initial wall-*seek* bout leading to an absence of movement and failure to reach the wall as noted above. A further point of vulnerability was seen in the range $\lambda=0.09$ – 0.15 and occurred when the robot attempted to execute the *cylinder pickup* FAP but either failed to grasp the cylinder (*fgc* in 9b) or failed to raise the gripper arm at the end of cylinder-pickup bout (*fra* in 9c). An example of the *fgc* failure is shown in the Supplementary Video (part 5). Failure to raise the gripper arm occurred in 80% of trials at $\lambda=0.12$ and 50% of trials at $\lambda=0.15$, and also resulted in a *behavioral trap*, as described in Appendix 1, where the robot detected its lowered arm as an obstacle and engaged in a slow circling behavior until the end of the trial.

Failures more likely at low salience levels. Our experiments show that, at low λ , weakly selected behaviors are typically not executed with sufficient vigour and can be vulnerable to interrupt, further investigation also shows support for hypothesis 4—that the effects of varying simulated dopamine can also depend on salience level. Specifically, comparison across the 10 trials at $\lambda=0.15$ shows that the variability in outcome (successful vs. unsuccessful) resulted from differences in the timing of the initial *cylinder-pickup* bout across trials—the robot encountered a cylinder, and initiated the *cylinder-pickup* FAP, significantly *later* in the successful runs ($M=66.7s$, $SD=6.88$) compared to the unsuccessful runs ($M=52.0s$, $SD=2.23$) (independent-samples t-test: $t(4.8)=4.557$, $p=0.007$, equal variances not assumed). Recall that the salience of *cylinder-pickup* increases with simulated ‘hunger’, which in turn increases gradually with longer search times. In other words, for those runs at $\lambda=0.15$ in which a cylinder is discovered quickly, and in which the robot is therefore more likely fail through premature deselection, the selection of the *cylinder-pickup* behavior is at a *lower* salience level than for the successful trials (longer search durations). This can be related to figure 7b which showed reduced hysteresis, and hence less behavioral persistence, for low values of λ (compared to intermediate values). More generally, in all low λ conditions, robot behaviors are executed more efficiently at higher salience levels, and therefore the symptoms of reduced simulated dopamine such as slowed movement are more pronounced when salience is low.

Behavioral consequences of high simulated tonic dopamine ($\lambda > 0.3$)

Distortion of winning channels by active losers. At high levels of λ the non-embodied model predicted reduced inhibition of the motor output from losing channels leading to distortion of the winning action (hypothesis 3). The behavioral consequences of distortion are visible in the robot model with levels of simulated dopamine $\lambda \geq 0.31$ and occasionally resulted in behavioral

disintegration for $\lambda=0.31, 0.34$ through failure to complete a foraging bout (*ff* in figure 9c). Likelihood of failure increased with very high levels of λ with more than 50% fails at $\lambda=0.4, 0.43$ and 100% fails at $\lambda=0.46$. At all of these λ levels, failure to forage was typically due to an inability to grasp a cylinder (*fgc*), however, other evidence of behavioral disintegration was also evident, particularly, difficulty in tracking walls (*lw*). Failure to grasp a cylinder often results in a second form of *behavioral trap* where the robot enters repeated cycles of *cylinder-seek* and (unsuccessful) *cylinder-pickup*, an example of this is shown figure 10e ($t=85-120s$), an example of this type of failure is shown in the Supplementary Video (part 6).

Failure more likely at high salience levels. That there was a mix of successful and unsuccessful runs, at some high λ levels, indicates that the impact of distortion on behavioral outcome can depend on circumstances. We illustrate this by comparing, in figures 10d and 10e, two trials with $\lambda=0.31$ showing that both successful foraging (10d) and disintegrated foraging (10e) are possible at this level. In 10d, the robot quickly locates a cylinder at $t=49s$, in 10e, the only unsuccessful run at this λ level, there is a much more protracted *cylinder-seek* search ending at $t=84s$ (see Appendix 1 for a detailed commentary and comparison). At higher λ levels (0.40 and 0.43), comparison of successful ($M=37.1s, SD=6.06$) vs. unsuccessful trials ($M=63.3s, SD=16.4$) shows that, on average, in successful runs the robots discovered a cylinder whilst foraging *earlier* than in unsuccessful trials (independent-samples t-test: $t(18)=-4.741, p<0.001$). This is the reverse of the situation with low λ —with high simulated dopamine it is the longer search bouts, giving rise to higher salience levels (from increasing ‘hunger’), that tend to result in greater behavioral disintegration. This again matches hypothesis 3—that the effect of varying simulated dopamine on behavior will depend upon salience levels—with contrasting effects seen at low and high λ levels.

From figure 7b, we can expect reduced hysteresis (behavioral persistence) for higher levels of λ , however, that figure also shows that increasing salience at high λ does *not* significantly impact on hysteresis. To understand why the robot performs better at lower levels of salience with high λ we therefore need to look beyond the basal ganglia model itself and to consider the influence of distortion on behavioral persistence via its effect on behavior. This is the topic of our final analysis.

Effects of distortion on behavioral persistence

A key property of the robotic model, that distinguishes it from the non-embodied simulation, is that selection outcomes have behavioral consequences that shape the robot’s subsequent sensory experiences. More specifically, the robot’s motor output, in part, determines its trajectory through the state-space of perceptual and motivational affordances for future selection competitions. Since varying the level of simulated dopamine can influence motor behavior by slowing movement or by merging partially selection actions with winning ones, it is interesting to establish whether this has any significant consequences for the selection behavior of the embodied model.

Here we explore this issue by examining the some effects of distorted selection on the timing and frequency of behavior switching. To assist this analysis an additional 90 robot trials were performed at all of the λ levels previously tested, but this time with a ‘winner-takes-all’ filter applied to the efficiency values of all sub-systems, such that the winning sub-system was always assigned an efficiency of 1.0, and all losers an efficiency of 0.0. In the following analyses the behavior of this winner-takes-all variant will be contrasted with the ‘soft’ selection generated by the standard model that allows multiple channels to influence motor output.

Timing of behavior switching. Our investigation of the non-embodied model showed significant hysteresis at almost all levels of simulated dopamine in the context of closely-matched salience competitions (figure 7), this should show up strongly, in the robot model in the initial transition from avoidance to foraging behavior. The key competitors at this point are wall-follow and cylinder-seek and the prime determinant of their relative salience, that eventually allows the latter to prevail, is a gradual, time-determined reduction in ‘fear’ alongside a steady increase in ‘hunger’. The length of the time leading up to this switch from avoidance to foraging therefore provides an measure of the operation of behavioral persistence in the model. Figure 11a plots this ‘time-to-switch’ measure against different levels of λ and shows the different outcomes observed with both the standard model

(from the original set of 90 trials) and the new winner-takes-all control. For each dopamine level we plot the average and standard error of the time-to-switch calculated over the five trials.

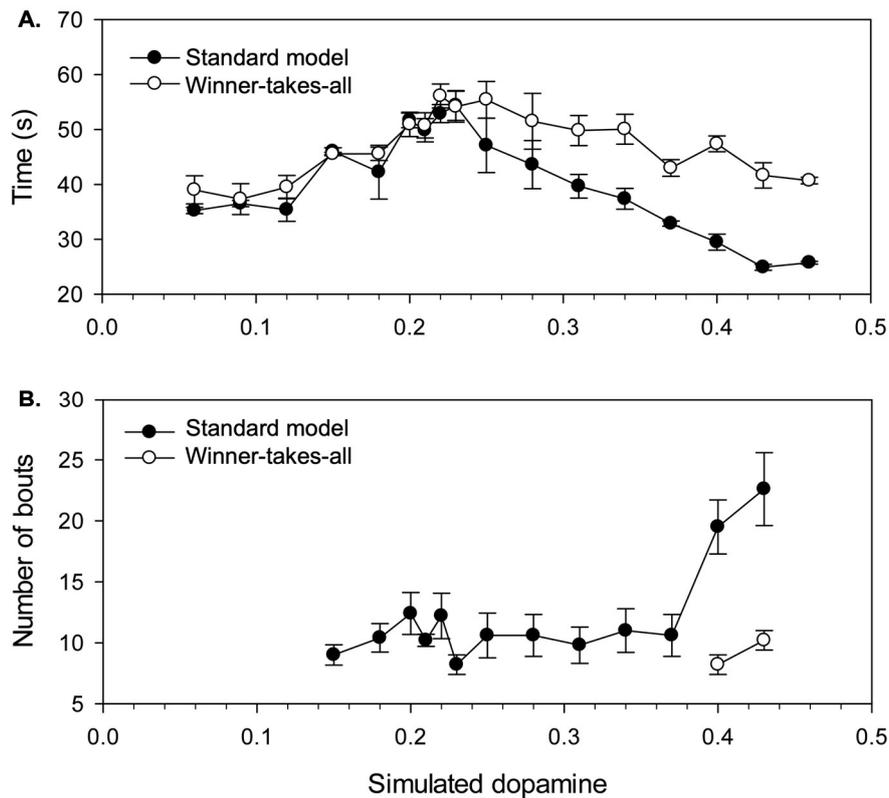


Figure 11. Comparing the standard ‘soft switching’ model of basal ganglia and a winner-takes-all variant on switch timing and frequency in the robotic model for different levels of simulated dopamine. **A.** ‘Time-to-switch’ from avoidance to foraging. The plot demonstrates that persistence (time-to-switch to foraging) varies with simulated dopamine and is affected by motor distortion at higher dopamine levels, in the case of the standard model only, leading to earlier switching (less persistence) compared to the winner-takes-all variant. **B.** Total number of bouts during the first avoidance and foraging sequences combined. Bout frequency is significantly increased at very high λ levels for the standard model only, indicating that distortion of motor behavior can cause more frequent switching. Each average is over five runs, bars show standard errors.

Comparison with figure 7b, shows that the graph for the winner-takes-all variant provides a good match to the degree of hysteresis found for a fixed salience (on the initial winning channel) of 0.4. Since the salience of wall-follow preceding the switch is typically in the range 0.3–0.4, this demonstrates that hysteresis in the embodied model basal ganglia generates a corresponding level of behavioral persistence under winner-takes-all conditions. However, the standard model generates an interesting difference from this result. Specifically, two-way ANOVA shows a significant interaction ($F(1,16)= 3.641, p<0.001$) between model type (standard vs. winner-takes-all) and λ . Posthoc comparisons for low, intermediate and high λ values, show a difference for high values only ($\lambda \geq 0.31$) where switching occurs significantly earlier in the standard model ($M= 31.7s, sd=6.26$) compared to the winner-takes-all variant ($M= 45.4s, sd= 5.66$) (independent-samples t-test: $t(58)= -8.92, p<0.001$). We conclude that, with higher λ , the distortion provided by losing channels can significantly reduce behavioral persistence in the robot over and above the reduction resulting from lower hysteresis in the embedded basal ganglia at higher levels of simulated dopamine.

Looking at figure 10 (panels d and e), which showed behavior for two trials with $\lambda=0.31$, we can observe, towards the end of the wall-follow bout (around $t=30$), a small, but gradually increasing, output on the cylinder-see channel. It is this ‘leakage’ of motor output from the cylinder-see subsystem that constitutes the difference between the standard and winner-takes-all versions of the

model. A key to understanding the effect of this distortion is to note that the wall-follow behavior is not especially robust, and is sometimes pushed off track by sensor noise or wheel-slip even when driven by a clean motor signal. The effect of the motor noise introduced by partial selection of cylinder-seek is therefore to increase the variability in the robot trajectory making it more difficult to maintain sensor contact with the nearby wall. In this situation, any loss of the wall percept due to distorted movement will lead to a rapid reduction in wall-follow salience and a switch to the cylinder-seek behavior.

Increased switching frequency with high simulated dopamine. If distortion makes some behaviors more vulnerable to interrupt, then we might also expect increased levels of behavior switching. To investigate this possibility, figure 11b illustrates one specific measure of switch frequency, the total number of bouts occurring during the first avoidance sequence and first foraging sequence of each trial. This measure was preferred to counting bouts (or switches) within a fixed time interval as it allows us to exploit a useful base-line—integrated behavior (by our earlier operational definitions) requires a minimum of seven bouts across these two sequences.

Since this measure can only be applied to trials containing a completed foraging sequence, this analysis only considered λ values in the range 0.15–0.43, and the graph plots the average and standard error of the number of bouts observed for the five successful trials at each simulated dopamine level. These data reveal that the performance of the robot is slightly above the base-line level of seven bouts across most of the range of simulated dopamine values, however, the number of bouts increases substantially for very high λ levels ($\lambda=0.40, 0.43$; $M=21.3$ bouts, $SD=4.73$). Moreover, as shown in figure 11b, comparison with winner-takes-all selection at these levels ($M=9.2$ bouts, $SD=1.99$) shows that the latter requires significantly fewer bouts (independent-samples t-test: $t(2.22)=4.33$, $p=0.041$, equal variances not assumed). We therefore conclude that the increased switching seen with the standard model is largely due to the distortion of motor output created by losing competitors. Figure 10e shows an example run with $\lambda=0.40$ that illustrates the increased frequency of bout switching (between wall-seek and wall-follow in $t=0-50s$) that can occur due to distortion with high simulated dopamine.

These analyses of the effects of increased λ on switch timing and frequency demonstrate that distortion in the robot model does not inevitably lead to a mixed motor output—trying to do two things at once—instead, its effect can be to make certain behavioral states more vulnerable to interrupt which can then lead to an increased frequency of behavior switching.

5. Discussion

Robotics can play an important role in neuroscience through its ability to create computational models of the nervous system that are *embodied*, that is, they control physical devices (robots) that exists in the world, and *situated*, that is, they must engage in real-time, and in closed sense-action loops, with the environments in which they are placed [59, 60]. Robotic models, like animals, can display *integrated behavior*, where they generate sequences of actions that are coherent with both their internal motivations and the unfolding dynamics of the world [45, 61]. Conversely, their behavior can become *disintegrated* when action sequences fall out-of-step with the affordances of the environment and they fail to achieve their goals [36]. The study of robotic models therefore offers opportunities for comparisons with animal and human behavior that differ from those that are available from the non-embodied models more typically studied in computational neuroscience. For instance, we can study them objectively, as behaving systems, without having to place an interpretation on their inputs and outputs [62, 63]. We can also examine the consequences for this observable behavior of specific interventions that simulate changes to the nervous system studied in relevant animals models, or that might arise in human neurological disorders.

Effects of simulated dopamine modulation on robot behavior

In the current study we explored the capability of an embedded basal ganglia model to generated patterns of integrated behavior when operating across a range of simulated tonic dopamine levels (λ). The robot performed the intended avoidance and foraging behaviors successfully for a range of intermediate λ values (0.2-0.28), values below this range caused some

slowness of movement, in line with previous predictions from non-embodied models, with movement speeds falling below 75% of its intended vigor at around half of this range ($\lambda=0.12$), and with prolonged periods of no movement for very low λ values (0.06 or less). Some runs with low λ also resulted in the premature deselection of behavior. High values of λ (0.3 or greater) lead to some distortion of motor output as the result of partial (or full) selection of multiple competing action sub-systems.

We found that simulated dopamine modulation of action selection outside the intermediate range did not invariantly lead to behavioral disintegration, since its effects varied with the precise circumstances of the robot. Specifically, low λ systems functioned well (selecting cleanly) with high salience signals but poorly with weak salience inputs. Conversely, high λ systems generated cleaner selections at low salience levels. While expectations from non-embodied modelling (hypotheses 1-4 above) were borne out in the robot implementation, the performance of the robot, across the full range of λ values, was better than might have been predicted from prior analyses of the selection properties of the model basal ganglia. This result can be explained by the finding that the robot, through its behavior, “self-structures” its own input [64], sampling only a limited area of the state-space of salience competitions, and predominantly parts of the space that have better-than-average outcomes (in terms of effective selection).

Hysteresis in the non-embodied model translates into persistence in behavioral expression in the robot. Persistence varied in an interesting way with λ , in a manner only partially explained by the behavior of the embedded basal ganglia model. Persistence was maximal at an intermediate λ levels, with reduced persistence at both lower and higher levels that could be traced to the functioning of the basal ganglia-thalamocortical loop. For high λ , reduced persistence was also partly the result of motor distortion making the current behavior of the robot more vulnerable to interrupt. This is an outcome that was *not* predictable from the disembodied model. Very high levels of λ also produced an increase in behavior switching within extended sequences of goal-directed activity. Again, this result is not entirely predicted by the disembodied model which forecast a greater degree of distortion (mixed behavior) at high λ values as a result of partial or full selection of multiple competitors.

Dysfunction of basal ganglia dopaminergic function in animals and humans

Dysfunction of dopaminergic regulation of the basal ganglia is implicated in a range of neurological disorders [35]. In Parkinson’s Disease (PD), for instance, tonic dopamine depletion in the striatum is one of the primary drivers of symptoms, including those relating to impaired movement and difficulty in initiating movement [65]. In computational neuroscience models, the progressively debilitating effects of PD have been modelled as increased attenuation of tonic dopamine in the striatum [66-68]. ADHD, which is characterized by hyperactivity, impulsiveness, impaired attention, and executive dysfunction, has also been linked to dopamine dysregulation, and particularly, to increased levels of dopamine transporter that remove dopamine from the synapse [35]. This outcome has been modelled as resulting in a less pronounced (compared to PD) reduction in striatal dopamine [69]. In schizophrenia, by contrast, an up-regulation of dopamine is thought to underlie symptoms related to disorganization including expression of bizarre or inappropriate behavior [35, 70], this has been modelled as involving an increase in striatal dopamine [71]. Tourette’s syndrome which causes sufferers to make involuntary movements or sounds has also been characterized as a consequence of elevated striatal dopamine [71, 72]. Other motor dysfunctions such as chorea and dystonia have been hypothesized to involve a failure to inhibit unwanted movements in which dopamine dysregulation could be implicated [7]. Obsessive-Compulsive Disorder (OCD) is thought to involve hyperactivity in parts of the orbitofrontal cortex, and treatments involving dopamine antagonists have been found to augment the benefits of therapies involving serotonin reuptake inhibitors [73].

A large number of animal models have been developed to investigate the neurological bases for these disorders many of which have explored genetic, developmental, drug or lesion-induced alterations to the dopamine system [73-78]. Animal studies have also directly explored the role of dopamine in regulating action selection and motivated behavior [79-82]. In the remainder of this

discussion we briefly compare results with the robot model with animal studies and human neurological disorders thought to involve lowered or heightened levels of tonic striatal dopamine.

Dopamine-depleting interventions and neurological conditions associated with reduced striatal dopamine

Behavior execution. In animals, activational aspects of motivation, such as response rate, vigor, and persistence, are impaired at doses of DA antagonist that leave intact directional or goal-directed aspects of responding (for review see [9, 12, 16, 81]). In patients with PD, major symptoms include slowness in movement (bradykinesia), reduced size of movement (hypokinesia), and absence of movement (akinesia) [83]. Consistent with these findings, in the robot model, slowed movement was a visible consequence as λ was lowered below the intermediate range, often leading to more prolonged bouts of behavior as action sequences take longer to perform. As λ was further reduced, movements were only partially executed or even fully suppressed, despite high-levels of motivation.

Salience. In animal models, behavior evoked by events that have high biological salience are comparatively resistant to dysfunctional dopamine neurotransmission. Thus, complex learned responses to mild stimuli are more prone to disturbance than unlearned responses evoked by intense unconditioned stimuli [12]. Similarly, behavior directed by external sensory stimuli is less affected than internally motivated (interoceptive) behavior [15, 21]. Consummatory behaviors (e.g. eating, drinking) are less disrupted than preparatory behaviors (acts that lead to, or make possible, consummatory behaviors) [10, 16, 20, 84, 85]. For example, while lesions of the mesolimbic dopamine projection abolish food hoarding in rats, actual feeding and drinking remain relatively unaffected [85]. High levels of arousal evoked by painful or highly arousing stimuli (such as being plunged into a icy bath) can lead to the restoration of normal behavioral responses (such as swimming) in otherwise akinetic animals caused by lesions that effect the dopamine system [24, 86]. Patients with PD often show problems in initiating movement, however, salient visual stimuli such as stripes painted on the floor can facilitate initiation of walking and reduce the incidence of freezing of gait [87]. Patients with PD can also show “paradoxical kinesia” (close to normal movement) in times of acute stress, for example when escaping from fire [88]. Salience competitions appear to have a deleterious affect on patients with PD that is more marked than in controls, for instance, a stimulus such as a doorway can have an inhibitory effect on movement, causing some patients to freeze; irrelevant stimuli have also been found to increased reaction times in a manual response task [87]. More broadly, patients with PD can also have difficulty expressing two motor programs simultaneously [83, 89].

Our robot model casts interesting light on some of these findings. For instance, we found that, with low λ , behavioral selections made between highly salient competitors were less vulnerable to partial selection, or no selection, than those made on the basis of low salience competitions (figure 6). High levels of motivation also led to a general increase in salience for competing behaviors and consequently clean(er) selection. We also found that selection in the low- λ robot was impaired by increased salience of a competitor, in some situations this led to freezing where competitors were evenly matched (e.g. Supplementary Video, part 4). More generally, at low λ levels, selection of the winning channel was more impacted by the presence of activity in competing channels than in similar circumstances but with λ in the intermediate range.

Lack of persistence. Rats with reduced dopamine show difficulty in maintaining motivated behavior over time. For instance, Gaddy and Neill [17] showed that dopamine-deprived animals had impaired performance of behaviors requiring sustained effort, whilst Salamone [10] found increased frequency of unfinished feeding bouts (partially-eaten food pellets) and failure to carry food pellets to normal feeding loci. Patients with PD often make incomplete movements and can exhibit sudden freezing, they also show rapid fatigue and can have difficulty in maintaining a behavior over time. For example, in the case of hand writing, for many patients their letters become smaller and smaller (micrographia) before writing ceases altogether [90]. In the robot model we found that low λ makes the currently selected behavior more vulnerable to early deselection or interrupt, largely as the result of decreased thalamocortical feedback failing to maintain the selected behavior. A similar challenge could underlie the premature deselection of behaviors seen in PD [see 83] and the increased

distractibility, and lack of persistence, associated with ADHD. As illustrated in Figure 7b, hysteresis in the basal ganglia falls of quite quickly as λ is reduced, including for values in the intermediate range when salience is at a moderate level. This is consistent with the observation that individuals with ADHD show problems with behavioural persistence without the motor symptoms (bradykinesia etc) associated with more profound deficiencies in striatal dopamine.

Behavioral timing. Studies with animals provide inconsistent evidence regarding switching frequency and time to initiate behaviors, with outcomes varying with experimental set-up [10]. In the robot model we found that time-to-switch depends on the salience of the behavior and on that of its competitors. This may help explain inconsistent findings in humans and animals. For example, in PD there is evidence that while some visual saccades are slowed, others are made more rapidly (hyper-reflexively) than in controls. Through meta-analysis we previously demonstrated that latency to saccade was dependent on the size (eccentricity) of the saccade, with smaller saccades more likely to be hyper-reflexive [91]. We suggest that this outcome arises because the current fixation behavior is more vulnerable to early interrupt due to reduced hysteresis in the relevant basal ganglia loop.

Dopamine increasing interventions, and neurological conditions involving increased striatal dopamine

Response frequency and duration. Animals treated with dopamine agonists show increases response frequency alongside decreased response duration with increases in dose [92-94]. Seen in the context of our robot study this is consistent with our finding of reduced time to switch and increase in distractibility and number of bouts with high levels of λ (see figures 10e, 11).

Suppressing unwanted actions. A common feature of neurological disorders involving increased striatal dopamine is difficulty in suppressing unwanted actions and thoughts. These can include the more stereotyped forms of unwanted action or speech seen in Tourette's syndrome, as well as the short twitch-like movements seen in chorea and thought to resemble fragments of normal behaviors, and perhaps some of the intrusive thoughts and bizarre actions associated with schizophrenia. In the non-embodied basal ganglia model elevated λ levels resulted in simultaneous selection of multiple channels, an outcome that has some resemblance to dystonia. However, the robot model generated a somewhat different result including patterns of rapid switching between channels, indicating that interruption of ongoing behavior is made more likely by the motor interference generated by a partially selected competing channel. The more promiscuous forms of selection enable by higher dopamine levels mean that patterns of behavior whose salience activity is "bubbling below the surface" may find an opportunity for expression due to a momentary loss of attention or concentration.

Stereotypy and hyperactivity. At higher doses of DA agonist, animals typically express a narrower range of behaviors and can become fixated on certain action patterns, that have become known as stereotypies. These may be oral (e.g. licking, biting, and gnawing) but that can also include forms of repetitive movement, including running [95], that are matched to environmental affordances. For example, Kelley et al. [94], summarising results with a hole-board task, commented that "with the higher doses [of amphetamine], locomotor routes become shorter and animals focalize uniquely on the holes (but still maintaining some locomotion and shifting from hole to hole) [...] residual components of the original behavior remain, but their pattern is greatly altered" (p. 73). Dopamine transporter (DAT) knockout mice, which have levels of striatal dopamine that are elevated by 70%, show hyperactivity and reduced habituation when placed in a novel environment [96], while DAT knockout rats are less sensitive to reward than wildtype animals, and show rigidity of action choice, alongside, hyperactivity, choice pattern and compulsive stereotypies [97].

Dopamine agonist-induced stereotypy in animals has been seen as a model for schizophrenia—though schizophrenics typically do not exhibit motor stereotypies, their symptoms often do involve compulsive and repetitive patterns of behavior and thought [93]. Repetitive sequences of actions, including constrained exploration patterns within an open environment, have been observed in rats treated with the DA agonist quinpirole and have been compared to the rituals seen in people with obsessive compulsive disorder [95].

Qualitatively, the behavior of the robot model at the highest λ trials (e.g. figure 10f) bears some resemblance to patterns on behavior in hyper-dopaminergic animals—the actions of the robot sample a narrow range of the potential actions, resemble some elements of complete action patterns but are fragmentary, poorly organized and fail to achieve goals (see, e.g. Supplementary Video, part 6). The underlying cause of the behavioral disintegration is selection (full or partial) of multiple channels, leading to early interrupt of ongoing behavior or mixing and distortion of motor acts. In animals, removal of basal ganglia inhibition from the motor system, will lead to complex effects as selection of behavior is governing by multiple brain systems, including attentional mechanisms, which we might consider as ‘early’ selection, and brainstem and motor components, that may provide forms of ‘late’ selection [98].

Limitations and related work

The current model can be improved along a considerable number of lines. First, whilst the Gurney et al. model of basal ganglia employed here has been shown to have enduring appeal, there are multiple ways in which it has been improved and extended that could be integrated into a future robot embodiment (see [99]). For example, a richer model of D1/D2 receptor behavior (see [100]) could impact on the behavior of a robotic model, as has been investigated for a simulated robot by Bahuguna et al. [101]. There is also scope to develop the wider architecture. For instance, whilst the current model builds on our understanding of dorsal basal ganglia pathways, the ventral basal ganglia domain shows important similarities and differences, and significantly, plays a critical role in the regulation of dopamine neurons [102].

Our robotic modelling demonstrates the importance of understanding how selection circuitry interacts with wider sensorimotor systems in the brain. Elsewhere, we have explored this in the context of cortical and sub-cortical loops involved in the selection of eye movements in a robotic active vision model [103, 104], and in the control of whisker-guided behaviour in robots with moving vibrissae [105]. A simplified version of the current model has also been deployed in a commercial biomimetic animal-like robot controlled by brain-inspired layered architecture [106]. Other interesting work in this direction includes models of basal ganglia interactions with locomotor pattern generators systems such as those underlying lampbrush swimming [107]. For a more complete brain-inspired architecture that includes a basal ganglia model of action selection see [108].

The current model highlights the importance of understanding how drive systems in the brain interact with action selection mechanisms. In place of the proxy models of drives used here, future models could usefully explore drive models based on a more realistic model of energy management (e.g. [109]). Another interesting direction to explore is the interaction of the basal ganglia with other brain substrates involved in motivational and action selection. For example, in [110] we developed a layered model of the hypothalamus that models the interplay of hunger and satiety in a simulated foraging task, the model also operates to regulate the activity of simulated dopamine neurons in the ventral tegmental area. Variability in the tonic dopamine signal could be an interesting target for modelling as it is known to be impacted by task engagement, motivation and arousal systems, stress and reward [111-113] and has been shown here to have a significant interaction with salience in supporting effective selection. Finally, action selection is also impacted by other major neuromodulatory systems besides dopamine [114] as has been explored in a robotic model by Krichmar [115].

Conclusion

Neuroscience is faced with the challenge of interpreting the outcomes of animal studies in the context of limited evidence. For instance, in seeking to understand the role of the basal ganglia in action selection, in any given study, whilst we have some access to information about what behavior is being selected, we have very little insight into understanding what competing behaviors are being considered but are not being selected. Many of our measures of behavioral outcome are also entirely ambiguous with regard to mechanism, for example, perseveration of behavior could be as the result of increased salience for a behavior, increased positive feedback, or the failure of competing behaviors to interrupt. Whilst these alternatives could be disentangled through careful experimentation, the

transparency of the robot architecture, and the benefits of synthetic approach (see also [116]), allow us to precisely observe the operation of the underlying control systems and their role in generating observed behavior [59, 63]. Studying robot models, can therefore inspire us to think about target brain systems in a new light. For instance, the current robot model reminds us that the activity of non-selected competitors can have a critical influence on how selection competitions are resolved and how the resulting behaviors are expressed.

In our model system, as in animals including humans, we see an inverted U-shape relationship between successful performance of integrated behaviour and the level of tonic (simulated) dopamine. The robot with low simulated dopamine shows slowed movement or no movement reminiscent of the bradykinesia and akinesia seen in Parkinson's disease. With excessively high levels of simulated dopamine the robot displays hyperactivity and rapid switching between behaviours; symptoms that show some resemblance to hyper-dopaminergic outcomes in animals and humans. Inappropriate repetition of behaviour, or perseveration, is observed in psychiatric conditions and animal models associated with both reduced and elevated levels of striatal dopamine. Similarly, in our robot model we saw perseveration with low and high levels of tonic simulated dopamine, sometimes associated with a behavioural trap. In the latter case, this might involve the robot failing to complete an ongoing behavior leading to repeated cycles of behavioral initiation.

Whilst there is much in this model that is oversimplified, we hope that it demonstrates the potential to apply robotics as a means to test models developed in computational psychiatry. Particularly, the differences between embodied and disembodied simulations investigated here, demonstrate that robotics can make observable some of the consequences of computational models that are not apparent when those models are tested in isolation.

Author Contributions: Conceptualization: Tony Prescott, Kevin Gurney, Mark Humphries and Peter Redgrave. Formal analysis: Kevin Gurney. Investigation: Tony Prescott and Fernando Montes Gonzalez. Methodology: Tony Prescott, Fernando Montes Gonzalez and Mark Humphries. Software: Tony Prescott and Fernando Montes Gonzalez. Validation: Tony Prescott. Writing – original draft: Tony Prescott. Writing – review & editing: Tony Prescott and Peter Redgrave.

Acknowledgements: This research was funded by the UK Engineering and Physical Sciences Research Council (EPSRC) under grant no. GR/R95722, by the European Union FET Flagship Human Brain Project (HBP-SGA3, grant no. 945539), and by Innovate UK Funding under the UK's funding guarantee scheme for the EIC Pathfinder project CAVAA (project no. 101071178).

Conflict of Interest Statement: TP is the co-founder and Director of two UK robotics companies, Consequential Robotics Ltd and Bettering our Worlds Ltd. Neither company stands to benefit from the publication of this research.

Appendix 1. Detailed commentary on robot behavior in figure 10

Low simulated dopamine, figure 10a-c

In Figure 10a, with $\lambda = 0.06$ (run 4), the robot starts towards the wall at an extremely slow pace but comes to a standstill a few centimeters short. Some time later the robot begins (again very slowly) to explore the arena for cylinders. This behavioural pattern can be understood as resulting from the altered competitive relationship between the avoidance and foraging systems in the embedded basal ganglia model. Initially, the salience for *wall-seek* is high, however, owing to the low level of simulated dopamine, it is only partially selected relative to its competitors. While the robot moves slowly towards the wall, the salience for *cylinder-seek* (driven by increasing 'hunger') begins to increase while that for *wall-seek* falls away (caused by the programmed reduction in 'fear' over time). At the point where the two saliences are close to parity the basal ganglia selection competition is resolved in favor of a stand-off (there is no selection of either action). Movement resumes later when the salience of *cylinder-seek* has increased further and is sufficient for it to be partially selected.

In Figure 10b, with $\lambda=0.09$ (run 5), the robot completes the avoidance sequence (*wall-seek* followed by *wall-follow*) albeit moving slowly, also notice that it is briefly distracted by detecting a cylinder en

route (showing that the robot is more distractable than normal). During the subsequent foraging sequence the robot detects the cylinder but the *cylinder-pickup* bout is affected by slowed movement and the arm is not lowered sufficiently to allow the cylinder to be grasped (*fgc*).

Figure 10c shows an example failure for a run with $\lambda=0.12$ (run 3). The breakdown in behavioral integration occurs at a point (around $t=80$), during the execution of the *cylinder-pickup* bout, where the cylinder has been grasped but the arm has not yet been raised to the vertical position (*fra* in 9c). Here, the detection of the cylinder in the gripper, combined with the reduced efficiency of the *cylinder-pickup* selection, brings about a reduction in salience, and loss of positive feedback, which causes that action to be prematurely deselected. After a momentary period of inactivity, and since the robot now holds a cylinder, *wall-seek* becomes salient and is selected. Unfortunately, the robot now detects its own, still lowered, gripper-arm as a nearby surface and engages in its normal response to this form of sensory input, during a *wall-seek* bout, which is to rotate anti-clockwise (turning out from the 'wall'). Behaviorally we observe that the robot engages in a slow anti-clockwise rotation and, since the gripper rotates with the robot and stays down, this leads to a continuous 'circling' behavior. This outcome can be considered to be a form of 'behavioral trap' resulting from circumstances where the robot's actions serve to maintain sensory inputs that drive a repetitive motor response.

High simulated dopamine, figure 10d-e

Figures 10d and 10e show two trials with $\lambda=0.31$ (runs 1 and 5) comparing successful (10d) and disintegrated (10e) outcomes.

In the successful run shown in 10d, there are a number of brief episodes of distortion of the selected action (note that the sixth plot from top shows distortion level), however, only one of these results in an outcome that is immediately apparent to an observer. This is the distortion of the latter part of a *cylinder-pickup* bout, by partial selection of the *wall-seek* sub-system (this occurs at approximately $t=52$ in 10d). The effect of this distortion is, in fact, relatively benign. As described for the low λ condition discussed above, a close salience competition arises once the cylinder is grasped by the gripper resulting in lowered salience for *cylinder-pickup* and increased salience for the next element of the foraging sequence *wall-seek*. However, whereas this situation resulted in reduced efficiency of *cylinder-pickup*, followed by premature deselection in some low λ trials, with increased λ , efficiency is not compromised. Instead *cylinder-pickup* remains fully selected until the action pattern has completed (raising the arm to the upright position), but *wall-seek* also begins to control the robot (or more specifically, the wheel motors) before the pickup move is finished. Once again (as in the low λ trial), the partially raised gripper-arm is detected as a nearby surface to which the *wall-seek* sub-system responds with an anti-clockwise turn. However, since the arm continues to be lifted out of the way by the still-active *cylinder-pickup* behavior there is no behavioral trap, instead a smooth transition is observed from the combined turning/lifting movement of the distorted behavior to the more usual straight-ahead movement generated by *wall-seek*. In other words, in this instance, significant distortion occurred but this did not jeopardise the integrated nature of the full behavioral sequence. Distortion also occurred during the *cylinder-deposit* behavior, at around $t=75$, this time through the failure to fully deselect the preceding behavior *wall-follow*, however, once again, the consequences of the distortion—wheel movements that serve to keep the robot close to the wall—do not interfere with successful completion of the *cylinder-deposit* bout, and the integrity of the foraging sequence is maintained.

A different outcome occurs in the trial in figure 10e. Here, after a relatively prolonged bout of search, *cylinder-pickup* is activated by the detection of a cylinder. However, the selection competition is not cleanly resolved, and the *cylinder-seek* sub-system is partially selected during repeated bouts of *cylinder-pickup*. The consequences of this distortion are not benign, instead the robot is driven forwards, towards the cylinder, at a point where it needs to move backwards to make room for the lowered gripper-arm. As a result the gripper-jaw is not correctly aligned to grasp the cylinder. The usual outcome in this situation is that the cylinder falls from the gripper-jaw or is grasped by a thin edge such that its presence is not registered by the optical sensor. In either case, the *cylinder-pickup* bout is not completed successfully and the robot re-engages the *cylinder-seek* routine. The

appearance of the robot through this episode is of frantic activity—it repeatedly tries to collect a cylinder, but excessive wheel movement means the manoeuvre is never successfully completed. Note that, this time, there *is* a form of behavioral trap—the failure to succeed in the initial *cylinder-pickup* bout leads to an repeating sequence of alternations between *cylinder-seek* and *cylinder-pickup*. Since the goal state of the foraging sequence is never achieved (depositing a cylinder in a ‘nest’ area), the motivation driving these behaviors saturates at a maximum, and the high levels of behavioral salience that initiated the distorted output are maintained. Whilst the benign form of distortion (produced by *wall-seek*) was observed in nearly all trials with dopamine levels of $0.31 \leq \lambda \leq 0.37$, the more damaging form (produced by *cylinder-seek*) was observed in the two fails trials at $\lambda = 0.31, 0.34$ and with increased frequency for trials with $\lambda \geq 0.40$.

Figure 10e shows an example run with $\lambda = 0.40$ (run 5) that illustrates the increased frequency of bout switching that can occur due to distortion with high simulated dopamine. In this run the robot has difficulty following the contour of a wall for any extended period, with both the avoidance sequence and the latter part of the foraging sequence including multiple alternating bouts of wall-*seek* and wall-*follow*.

References

1. Grillner, S., et al., *Mechanisms for selection of basic motor programs--roles for the striatum and pallidum*. Trends Neurosci, 2005. **28**(7): p. 364-70.
2. Mink, J.W., *The basal ganglia: Focused selection and inhibition of competing motor programs*. Progress In Neurobiology, 1996. **50**(4): p. 381-425.
3. Redgrave, P., T.J. Prescott, and K. Gurney, *The basal ganglia: A vertebrate solution to the selection problem?* Neuroscience, 1999. **89**: p. 1009-1023.
4. Prescott, T.J., P. Redgrave, and K. Gurney, *Layered control architectures in robots and vertebrates*. Adaptive Behavior, 1999. **7**(1): p. 99-127.
5. Balleine, B.W., M.R. Delgado, and O. Hikosaka, *The Role of the Dorsal Striatum in Reward and Decision-Making*. The Journal of Neuroscience, 2007. **27**(31): p. 8161.
6. Grillner, S. and B. Robertson, *The basal ganglia downstream control of brainstem motor centres—an evolutionarily conserved strategy*. Current Opinion in Neurobiology, 2015. **33**: p. 47-52.
7. Mink, J.W., *The basal ganglia and involuntary movements: Impaired inhibition of competing motor patterns*. Archives of Neurology, 2003. **60**(10): p. 1365-1368.
8. Schultz, W., *Multiple Dopamine Functions at Different Time Courses*. Annual Review of Neuroscience, 2007. **30**(1): p. 259-288.
9. Arber, S. and R.M. Costa, *Networking brainstem and basal ganglia circuits for movement*. Nature Reviews Neuroscience, 2022. **23**(6): p. 342-360.
10. Salamone, J.D., *Dopaminergic involvement in activational aspects of motivation - effects of haloperidol on schedule-induced activity, feeding, and foraging in rats*. Psychobiology, 1988. **16**(3): p. 196-206.
11. Salamone, J.D., M.J. Zigmond, and E.M. Stricker, *Characterization of the impaired feeding-behavior in rats given haloperidol or dopamine-depleting brain-lesions*. Neuroscience, 1990. **39**(1): p. 17-24.
12. Salamone, J.D., *Behavioral pharmacology of dopamine systems: a new synthesis*, in *The Mesolimbic Dopamine System: From Motivation to Action*, P. Willner and J. Scheel-Kruger, Editors. 1991, Wiley and Sons.
13. Bakshi, V.P. and A.E. Kelley, *Dopaminergic regulation of feeding-behavior .1. differential-effects of haloperidol microinfusion into 3 striatal subregions*. Psychobiology, 1991. **19**(3): p. 223-232.
14. Salamone, J.D., K. Mahan, and S. Rogers, *Ventrolateral striatal dopamine depletions impair feeding and food handling in rats*. Pharmacology Biochemistry and Behavior, 1993. **44**(3): p. 605-610.
15. Bury, D. and W.J. Schmidt, *Effects of systemically and intrastrially injected haloperidol and apomorphine on grooming, feeding and locomotion in the rat*. Behavioural Processes, 1987. **15**(2-3): p. 269-283.
16. Salamone, J.D. and M. Correa, *The Mysterious Motivational Functions of Mesolimbic Dopamine*. Neuron, 2012. **76**(3): p. 470-485.
17. Gaddy, J.R. and D.B. Neill, *Differential behavioral changes following intrastriatal application of 6-hydroxydopamine*. Brain Research, 1977. **119**: p. 439-446.
18. Cousins, M.S. and J.D. Salamone, *Involvement of ventrolateral striatal dopamine in movement initiation and execution - a microdialysis and behavioral investigation*. Neuroscience, 1996. **70**(4): p. 849-859.
19. Cousins, M.S. and J.D. Salamone, *Skilled motor deficits in rats induced by ventrolateral striatal dopamine depletions - behavioral and pharmacological characterization*. Brain Research, 1996. **732**(1-2): p. 186-194.
20. Koob, G.F., et al., *Effects of 6-Hydroxydopamine lesions of the nucleus accumbens septi and olfactory tubercle on feeding, locomotor activity, and amphetamine anorexia in the rat*. Journal of Comparative and Physiological Psychology, 1978. **92**(5): p. 917-927.

21. Cools, A.R., *Role of the neostriatal dopaminergic activity in sequencing and selecting behavioural strategies: facilitation of processes involved in selecting the best strategy in a stressful situation.* Behavioural Brain Research, 1980. **1**: p. 361-378.
22. Gelissen, M. and A. Cools, *Effect of intracaudate haloperidol and apomorphine on switching motor patterns upon current behavior of cats.* Behavioural Brain Research, 1988. **29**(1-2): p. 17-26.
23. Marin, C., et al., *Effect of long-term haloperidol treatment on striatal neuropeptides - relation to stereotyped behavior.* Brain Research, 1996. **731**(1-2): p. 57-62.
24. Marshall, J.F., D. Levitan, and E.M. Stricker, *Activation-induced restoration of sensorimotor functions in rats with dopamine-depleting brain lesions.* Journal of Comparative and Physiological Psychology, 1976. **90**(6): p. 536-546.
25. Teitelbaum, P., T. Schallert, and I.Q. Whishaw, *Sources of spontaneity in motivated behaviour,* in *Handbook of Behavioural Neurobiology*, P. Teitelbaum and E. Satinoff, Editors. 1983, Plenum Press: New York. p. 23-66.
26. Oades, R.D., *The role of noradrenaline in tuning and dopamine in switching between signals in the CNS.* Neuroscience and Biobehavioral Reviews, 1985. **9**(2): p. 261-282.
27. Iversen, S.D., *Striatal function and stereotyped behaviour,* in *Psychobiology of the Striatum*, A.R. Cools, Editor. 1977, Elsevier: Amsterdam. p. 99-117.
28. Kelley, A.E. and S.D. Iversen, *Substance P infusion into substantia nigra of the rat: Behavioural analysis and involvement of striatal dopamine.* European Journal of Pharmacology, 1979. **60**: p. 171-179.
29. Bakshi, V.P. and A.E. Kelley, *Dopaminergic regulation of feeding-behavior .2. differential-effects of amphetamine microinfusion into 3 striatal subregions.* Psychobiology, 1991. **19**(3): p. 233-242.
30. Langen, M., et al., *The neurobiology of repetitive behavior: Of mice....* Neuroscience & Biobehavioral Reviews, 2011. **35**(3): p. 345-355.
31. Jankovic, J.J. and E. Tolosa, eds. *Parkinson's Disease and Movement Disorders.* 4th ed. 2002, Lipincott, Williams, & Wilkins: Philadelphia.
32. Moore, H., A.R. West, and A.A. Grace, *The regulation of forebrain dopamine transmission: relevance to the pathophysiology and psychopathology of schizophrenia.* Biol Psychiatry, 1999. **46**(1): p. 40-55.
33. Brisch, R., et al., *The Role of Dopamine in Schizophrenia from a Neurobiological and Evolutionary Perspective: Old Fashioned, but Still in Vogue.* Frontiers in Psychiatry, 2014. **5**: p. 47.
34. Joel, D., *Current animal models of obsessive compulsive disorder: A critical review.* Progress in Neuro-Psychopharmacology and Biological Psychiatry, 2006. **30**(3): p. 374-388.
35. Klein, M.O., et al., *Dopamine: Functions, Signaling, and Association with Neurological Diseases.* Cell Mol Neurobiol, 2019. **39**(1): p. 31-59.
36. Prescott, T.J., et al., *A robot model of the basal ganglia: behaviour and intrinsic processing.* Neural Networks, 2006. **19**(1): p. 31-61.
37. Montague, P.R., et al., *Computational psychiatry.* Trends in Cognitive Sciences, 2012. **16**(1): p. 72-80.
38. Tolu, S., B. Strohmmer, and O. Zahra, *Perspective on investigation of neurodegenerative diseases with neurorobotics approaches.* Neuromorphic Computing and Engineering, 2023. **3**(1): p. 013001.
39. Gurney, K., T.J. Prescott, and P. Redgrave, *A computational model of action selection in the basal ganglia. I. A new functional anatomy.* Biological Cybernetics, 2001. **84**(6): p. 401-410.
40. Gurney, K., T.J. Prescott, and P. Redgrave, *A computational model of action selection in the basal ganglia. II. Analysis and simulation of behaviour.* Biological Cybernetics, 2001. **84**(6): p. 411-423.
41. Humphries, M.D. and K. Gurney, *The role of intra-thalamic and thalamocortical circuits in action selection.* Network: Computation in Neural Systems, 2002. **13**(1): p. 131-156.
42. Prescott, T.J., *Action Selection.* Scholarpedia, 2008. **3**(2): p. 2705.
43. Ludlow, A.R., *Applications of computer modelling to behavioural coordination.* 1983, London.
44. McFarland, D., *Problems of Animal Behaviour.* 1989, Harlow, UK: Longman.
45. Prescott, T.J., *Forced moves or good tricks in design space? Landmarks in the evolution of neural mechanisms for action selection.* Adaptive Behavior, 2007. **15**(1): p. 9-31.
46. Humphries, M.D., R.D. Stewart, and K. Gurney, *A physiologically plausible model of action selection and oscillatory activity in the basal ganglia.* J Neurosci, 2006. **26**(50): p. 12921-42.
47. Gillies, A. and D. Willshaw, *Models of the subthalamic nucleus: The importance of intranuclear connectivity.* Medical Engineering & Physics, 2004. **26**(9): p. 723-732.
48. Parent, A. and L.N. Hazrati, *Functional anatomy of the basal ganglia. II. The place of subthalamic nucleus and external pallidum in basal ganglia circuitry.* Brain Res Brain Res Rev, 1995. **20**(1): p. 128-54.
49. Suryanarayana, S.M., et al., *Roles for globus pallidus externa revealed in a computational model of action selection in the basal ganglia.* Neural Netw, 2019. **109**: p. 113-136.
50. Gerfen, C.R. and D.J. Surmeier, *Modulation of Striatal Projection Systems by Dopamine.* Annual Review of Neuroscience, 2011. **34**(1): p. 441-466.
51. Akkal, D., et al., *Responses of substantia nigra pars reticulata neurons to intrastriatal D1 and D2 dopaminergic agonist injections in the rat.* Neuroscience Letters, 1996. **213**(1): p. 66-70.

52. Cui, G., et al., *Concurrent activation of striatal direct and indirect pathways during action initiation*. *Nature*, 2013. **494**(7436): p. 238-242.
53. Tecuapetla, F., et al., *Complementary Contributions of Striatal Projection Pathways to Action Initiation and Execution*. *Cell*, 2016. **166**(3): p. 703-715.
54. Schultz, W., P. Dayan, and P.R. Montague, *A neural substrate for prediction and reward*. *Science*, 1997. **275**: p. 1593-1599.
55. Lehner, P.N., *Handbook of Ethological Methods*. 2nd ed. 1996, Cambridge, UK: Cambridge University Press.
56. Montes Gonzalez, F., et al., *An embodied model of action selection mechanisms in the vertebrate brain*, in *From Animals to Animats 6: Proceedings of the 6th International Conference on the Simulation of Adaptive Behavior*, J.A. Meyer, Editor. 2000, MIT Press: Cambridge, MA. p. 157-166.
57. Hinde, R.A., *Animal Behaviour: a Synthesis of Ethology and Comparative Psychology*. 1966, London: McGraw-Hill.
58. Lorenz, K., *Der Kumpan in der Umwelt des Vogels*. *Journal of Ornithology*, 1935. **83**: p. 137-213, 289-413.
59. Prescott, T.J. and S.P. Wilson, *Understanding brain functional architecture through robotics*. *Science Robotics*, 2023. **8**(78): p. eadg6014.
60. Hallam, J.C.T., et al., *Behaviour: perception, action and intelligence — the view from situated robotics*. *Philosophical Transactions of the Royal Society of London. Series A: Physical and Engineering Sciences*, 1994. **349**(1689): p. 29-42.
61. Brooks, R.A., *Coherent behaviour from many adaptive processes*, in *From Animals to Animats 3: Proceedings of the Third International Conference on the Simulation of Adaptive Behaviour*, D. Cliff, et al., Editors. 1994, MIT Press: Cambridge, MA. p. 22-29.
62. Mitchinson, B., et al., *Biomimetic robots as scientific models: A view from the whisker tip*, in *Neuromorphic and Brain-based Robots*, J. Krichmar and H. Wagatsuma, Editors. 2011, MIT Press: Boston, MA. p. 23-57.
63. Prescott, T.J., et al., *Embodied models and neurorobotics*, in *From Neuron to Cognition via Computational Neuroscience*, M.A. Arbib and J.J. Bonaiuto, Editors. 2016, MIT Press: Cambridge, MA. p. 483-512.
64. Verschure, P.F.M.J., T. Voegtlin, and R.J. Douglas, *Environmentally mediated synergy between perception and behaviour in mobile robots*. *Nature*, 2003. **425**(6958): p. 620-4.
65. Obeso, J.A., et al., *The basal ganglia in Parkinson's disease: Current concepts and unexplained observations*. *Annals of Neurology*, 2008. **64**(S2): p. S30-S46.
66. Frank, M.J., *Dynamic dopamine modulation in the basal ganglia: a neurocomputational account of cognitive deficits in medicated and nonmedicated Parkinsonism*. *J Cogn Neurosci*, 2005. **17**(1): p. 51-72.
67. Humphries, M.D., J.A. Obeso, and J.K. Dreyer, *Insights into Parkinson's disease from computational models of the basal ganglia*. *J Neurol Neurosurg Psychiatry*, 2018. **89**(11): p. 1181-1188.
68. Guthrie, M., C.E. Myers, and M.A. Gluck, *A neurocomputational model of tonic and phasic dopamine in action selection: A comparison with cognitive deficits in Parkinson's disease*. *Behavioural Brain Research*, 2009. **200**(1): p. 48-59.
69. Frank, M.J., et al., *Testing Computational Models of Dopamine and Noradrenaline Dysfunction in Attention Deficit/Hyperactivity Disorder*. *Neuropsychopharmacology*, 2007. **32**(7): p. 1583-1599.
70. Sonnenschein, S.F., F.V. Gomes, and A.A. Grace, *Dysregulation of Midbrain Dopamine System and the Pathophysiology of Schizophrenia*. *Front Psychiatry*, 2020. **11**: p. 613.
71. Maia, T.V. and V.A. Conceicao, *The Roles of Phasic and Tonic Dopamine in Tic Learning and Expression*. *Biol Psychiatry*, 2017. **82**(6): p. 401-412.
72. Singer, H.S., et al., *Elevated Intrasynaptic Dopamine Release in Tourette's Syndrome Measured by PET*. *American Journal of Psychiatry*, 2002. **159**(8): p. 1329-1336.
73. Xue, J., et al., *Midbrain dopamine neurons arbiter OCD-like behavior*. *Proceedings of the National Academy of Sciences*, 2022. **119**(46).
74. Jones, C.A., D.J.G. Watson, and K.C.F. Fone, *Animal models of schizophrenia*. *British Journal of Pharmacology*, 2011. **164**(4): p. 1162-1194.
75. Betarbet, R., T.B. Sherer, and J.T. Greenamyre, *Animal models of Parkinson's disease*. *BioEssays*, 2002. **24**(4): p. 308-318.
76. Blesa, J. and S. Przedborski, *Parkinson's disease: animal models and dopaminergic cell vulnerability*. *Frontiers in Neuroanatomy*, 2014. **8**: p. 155.
77. Dawson, T.M., H.S. Ko, and V.L. Dawson, *Genetic Animal Models of Parkinson's Disease*. *Neuron*, 2010. **66**(5): p. 646-661.
78. Schober, A., *Classic toxin-induced animal models of Parkinson's disease: 6-OHDA and MPTP*. *Cell and Tissue Research*, 2004. **318**(1): p. 215-224.
79. Wise, R.A., *Dopamine, learning and motivation*. *Nature Reviews Neuroscience*, 2004. **5**(6): p. 483-494.
80. Ikemoto, S. and J. Panksepp, *The role of nucleus accumbens dopamine in motivated behavior: a unifying interpretation with special reference to reward-seeking*. *Brain Research Reviews*, 1999. **31**(1): p. 6-41.
81. Salamone, J.D., et al., *Dopamine, Effort-Based Choice, and Behavioral Economics: Basic and Translational Research*. *Frontiers in Behavioral Neuroscience*, 2018. **12**: p. 52.

82. Berridge, K.C., *From prediction error to incentive salience: mesolimbic computation of reward motivation*. European Journal of Neuroscience, 2012. **35**(7): p. 1124-1143.
83. Berardelli, A., et al., *Pathophysiology of bradykinesia in Parkinson's disease*. Brain, 2001. **124**(11): p. 2131-2146.
84. Blackburn, J.R., A.G. Phillips, and H.C. Fibiger, *Dopamine and preparatory behavior .1. effects of pimozide*. Behavioral Neuroscience, 1987. **101**(3): p. 352-360.
85. Kelley, A.E. and L. Stinus, *Disappearance of hoarding behavior after 6-hydroxydopamine lesions of the mesolimbic dopamine neurons and its reinstatement with l-dopa*. Behavioral Neuroscience, 1985. **99**(3): p. 531-545.
86. Keefe, K.A., et al., *Paradoxical kinesia in parkinsonism is not caused by dopamine release. Studies in an animal model*. Arch Neurol, 1989. **46**(10): p. 1070-5.
87. McDowell, S.-A. and J. Harris, *Irrelevant peripheral visual stimuli impair manual reaction times in Parkinson's disease*. Vision Research, 1997. **37**(24): p. 3549-3558.
88. Schwab, R.S., *Akinesia paradoxa*. Electroencephalogr Clin Neurophysiol., 1972. **31**: p. 87-92.
89. Benecke, R., et al., *Performance of simultaneous movements in patients with parkinson's disease*. Brain, 1986. **109**(4): p. 739-757.
90. Wagle Shukla, A., et al., *Micrographia and related deficits in Parkinson's disease: a cross-sectional study*. BMJ Open, 2012. **2**(3): p. e000628.
91. Chambers, J.M. and T.J. Prescott, *Response times for visually guided saccades in persons with Parkinson's disease: A meta-analytic review*. Neuropsychologia, 2010. **48**(4): p. 887-99.
92. Rebec, G.V. and T.R. Bashore, *Critical issues in assessing the behavioral effects of amphetamine*. Neuroscience & Biobehavioral Reviews, 1984. **8**(1): p. 153-159.
93. Seiden, L.S., K.E. Sabol, and G.A. Ricaurte, *Amphetamine: Effects on Catecholamine Systems and Behavior*. Annual Review of Pharmacology and Toxicology, 1993. **33**(1): p. 639-676.
94. Kelley, A.E., M. Winnock, and L. Stinus, *Amphetamine, apomorphine and investigatory behavior in the rat: Analysis of the structure and pattern of responses*. Psychopharmacology, 1986. **88**(1): p. 66-74.
95. Eilam, D., *From an animal model to human patients: An example of a translational study on obsessive compulsive disorder (OCD)*. Neurosci Biobehav Rev, 2017. **76**(Pt A): p. 67-76.
96. Zhuang, X., et al., *Hyperactivity and impaired response habituation in hyperdopaminergic mice*. Proceedings of the National Academy of Sciences of the United States of America, 2001. **98**(4): p. 1982-1987.
97. Cinque, S., et al., *Behavioral Phenotyping of Dopamine Transporter Knockout Rats: Compulsive Traits, Motor Stereotypies, and Anhedonia*. Front Psychiatry, 2018. **9**: p. 43.
98. Allport, A., *Selection for action: Some behaviorial and neurophysiological considerations of attention and action, in Perspectives on Perception and Action*, H. Heuer and A.F. Sanders, Editors. 1987, Erlbaum: Hillsdale, NJ. p. 395-420.
99. Humphries, M.D. and K. Gurney, *Making decisions in the dark basement of the brain: A look back at the GPR model of action selection and the basal ganglia*. Biological Cybernetics, 2021. **115**(4): p. 323-329.
100. Humphries, M.D., et al., *Capturing dopaminergic modulation and bimodal membrane behaviour of striatal medium spiny neurons in accurate, reduced models*. Front Comput Neurosci, 2009. **3**: p. 26.
101. Bahuguna, J., P. Weidel, and A. Morrison, *Exploring the role of striatal D1 and D2 medium spiny neurons in action selection using a virtual robotic framework*. European Journal of Neuroscience, 2019. **49**(6): p. 737-753.
102. Humphries, M.D. and T.J. Prescott, *The ventral basal ganglia, a selection mechanism at the crossroads of space, strategy, and reward*. Progress in Neurobiology, 2010. **90**(4): p. 385-417.
103. Chambers, J., et al., *Mechanisms of choice in the primate brain: A quick look at positive feedback, in Modelling Natural Action Selection*, A. Seth, J.J. Bryson, and T.J. Prescott, Editors. 2011, CUP: Cambridge. p. 390-418.
104. Cope, A.J., et al., *Basal Ganglia Control of Reflexive Saccades: A Computational Model Integrating Physiology Anatomy and Behaviour*. bioRxiv, 2017.
105. Prescott, T.J., et al., *The robot vibrissal system: Understanding mammalian sensorimotor co-ordination through biomimetics, in Sensorimotor Integration in the Whisker System*, P. Krieger and A. Groh, Editors. 2015, Springer New York. p. 213-240.
106. Mitchinson, B. and T.J. Prescott, *Miro: A robot "mammal" with a biomimetic brain-based control system, in 5th International Conference of Biomimetic and Biohybrid Systems*, N.F. Lepora, et al., Editors. 2016, Springer International Publishing: Edinburgh, UK. p. 179-191.
107. Kamali Sarvestani, I., et al., *A computational model of visually guided locomotion in lamprey*. Biol Cybern, 2013. **107**(5): p. 497-512.
108. Verschure, P.F.M.J., C.M.A. Pennartz, and G. Pezzulo, *The why, what, where, when and how of goal-directed choice: neuronal and computational principles*. Philosophical Transactions of the Royal Society of London B: Biological Sciences, 2014. **369**(1655).
109. Girard, B., et al., *Where neuroscience and dynamic system theory meet autonomous robotics: A contracting basal ganglia model for action selection*. Neural Networks, 2008. **21**(4): p. 628-641.
110. Jimenez-Rodriguez, A. and T.J. Prescott, *Motivational Modulation of Consummatory Behaviour and Learning in a Robot Model of Spatial Navigation. in Biomimetic and Biohybrid Systems*. 2023. Cham: Springer Nature Switzerland.

111. Marinelli, M. and J.E. McCutcheon, *Heterogeneity of dopamine neuron activity across traits and states*. Neuroscience, 2014. **282**: p. 176-197.
112. Rice, M.E., J.C. Patel, and S.J. Cragg, *Dopamine release in the basal ganglia*. Neuroscience, 2011. **198**: p. 112-137.
113. Goto, Y., S. Otani, and A.A. Grace, *The Yin and Yang of dopamine release: a new perspective*. Neuropharmacology, 2007. **53**(5): p. 583-587.
114. Krichmar, J.L., *The Neuromodulatory System: A Framework for Survival and Adaptive Behavior in a Challenging World*. Adaptive Behavior, 2008. **16**(6): p. 385-399.
115. Krichmar, J., *A neurobotic platform to test the influence of neuromodulatory signaling on anxious and curious behavior*. Frontiers in Neurorobotics, 2013. **7**.
116. Hommel, B., et al., *No one knows what attention is*. Attention, Perception, & Psychophysics, 2019. **81**(7): p. 2288-2303.

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.